# Speeding Code Execution with the Econ Cluster

For most tasks, it is often enough to execute code on your own computer. And for most tasks, using a single core of a processor is also enough. However, there are compute intensive tasks that can greatly benefit from a computer with several cores and access to a relevant amount of memory. For these tasks, your own computer might not suffice. Fortunately, the Economics Department at Duke has its own cluster of computers that Master's and PhD students can use.

# 1 Requesting Access to the Cluster

To access the cluster, you will need a username and password. The username for the Econ Cluster is the same as your University NetID username, but the password is different. If you do not have a password yet, or do not remember it (if you are a PhD student you were probably assigned a password during your 1st year at Duke), you can request a password by emailing help@econ.duke.edu.

# 2 Connecting to the Cluster

You can think of the Econ Cluster as a collection of several computers, which are managed by some centralized software. One of these computers is responsible for handling log in, and is known as the front-end node (or the login node). This node is also responsible for receiving and scheduling tasks on the other computers in the cluster.

To log in the Econ Cluster we will use the SSH protocol. This protocol was designed to allow two computers to securely communicate over an insecure network. If you are using a Mac or Linux-based operating system (like Ubuntu), then you already have what is required to use SSH. If you are on Windows, then you will need to install an SSH client, like PuTTY or any of the alternatives (they are all similar and straightforward).

## 2.1 Mac, Ubuntu or Linux

Open the Terminal program. You should see a window similar to the one depicted in Figure 1.

• • •	😭 guilhermesalome — -bash — 80×24	
Guilhermes-MacB	ook-Pro:~ guilhermesalome\$	

Figure 1: Terminal on a Mac.

You can type commands after the dollar sign, and the terminal will interpret the commands and execute them (REPL). We will discuss how to use the terminal in the next section.

To log in the Econ Cluster, we will the the ssh command. The ssh command uses the syntax:

### ssh username@hostname

1

Where username is the username you will use to log in, and hostname is the address of the computer host that you will connect to. If you are currently inside the Duke network, then the hostname is login.econ.duke.edu. In my case, I would execute:

### ssh gfs8@login.econ.duke.edu

Then, I am asked for my password to finish logging in (see Figure 2).



Figure 2: SSH Asking for Password.

After typing your password, you should be greeted with a welcome message and are now connected to the cluster (see Figure 3). You are connected to a bash terminal, which can take commands and will execute them. We will discuss these commands in the next section.

	(2) 6	guilhern	nesalom	e — ssh
builhermes-MacBook-Pro:~ guilherm gfs89login.econ.duke.edu's passwo Last login: Fri Jun 28 12:03:23 2	nesalome\$ s ord: 2019 from 1	sh gfs: 8.194.	801ogin 79.76	.econ.di
Quota report for your home dire	ectory			
ilesystem /econ/home/g/gfs8	Size 109	Used 1.9G	Avail 8.20	Use% 19%
Quota usage for your research o	directory			
Filesystem /econ/research/gfs8	Size 500	Used Ø	Avail 500	Use% 8%

Figure 3: SSH Welcome Message.

If you are outside the Duke network, then you first the need to **ssh** into the Duke network, and then **ssh** again into the Econ Cluster. To **ssh** into the Duke network, you should use login.oit.duke.edu as the hostname, but now the username and password are the same you use for logging into Duke websites (like Dukehub). After that, you can execute **ssh** again to log in the Econ Cluster (see Figure 4).

of crashcritical uses only. You should have no spectral of private interval of the second state of the sec	nergen and a service. To rese on Off Virtual Computing Lab service. To rese re images will be added to this service in
f or arbitriid users and, You should have no expectation of privacy in your use of this interver. Use of this network constitutions constant to a storest for any purpose including crisinal protection. Sector 2014 and S	in the second se
<pre>f your use of this network. Use of this network constitutes consent to monitoring, retrieval, and disclosure of any information stored within the network for any purpose ficulating crisical prosection. Generative and the store of the</pre>	" resears on Tr Virtual Computing Lab service. To reas re images will be added to this service in
<pre># monitoring, retrieval, and disclosure of any information stored sithin the meters for any suppose including crisinal prosecution. ************************************</pre>	the # ######## com Of Virtual Computing Lab service. To rese re images will be added to this service in
<pre>returns for any purpose including crisinal prosection.</pre>	on OT Virtual Computing Lab service. To rese sadded to this service in
<pre>####################################</pre>	ssesses com OT Virtual Computing Lab service. To ress re images will be added to this service in
<pre>gfe8Diggin.eit.duke.edu's password: sail gin:Thu Aug 110:32:48 8026 from cpe-174-189-42-173.nc.res.rr.com Saif growisioned systems are new available for remote usage through the OIT ve your om vitual sachime places visit ven.duke.edu. Additional software i the cosing months. uvr/lin/rauth_ifile /tmp/Tauth_ofs8 dees not exist pfe8Diggin-ter:-13 [provid_ion] == sing sfe8Biggin.ecm.duke.edu sait login: Fri Jun 82 12:32:46 dops from login-teer-18.oit.duke.edu [Quata report for your heme directory</pre>	com OTT Virtual Computing Lab service. To ress re images will be added to this service in
Lest login: The Aug 11 80:32:48 3016 from ope-174-189-42-173, no. res.r.t.om Boff provisiond by Vetem are now available for renord usage through the OIT we your onn virtual section piese visit van.duke.edu. Additional software i har coming software i in / renorkutter for does not stati grabblogin.ecm.bl (renorkutter) = 5 sch grabblogin.ecm.duke.edu grabblogin.ecm.bl (renorkutter) = 5 sch grabblogin.ecm.duke.edu cast login: fri bar 25 Jibide 305 fram login-teer-18.oit.duke.edu cast renork your home directory	com OTT Virtual Computing Lab service. To rese re images will be added to this service in
Sof providend system are now available for reacts usage through the DTT or whar and other lanching pieces visit ven.duke.edu. Additional software i uver/bin/sufficiences.edu (the star star star star star star freeDigin/terr.ll (readwintin) - 5 sin gr#Blogin.econ.duke.edu stat login: Fri bm 28 12:13:164 000 from login-teer-10.oit.duke.edu stat login: Fri bm 28 12:13:164 000 from login-teer-10.oit.duke.edu	OTT Virtual Computing Lab service. To rese re images will be added to this service in
Filesystem Size Used Avail Use% /econ/home/g/gfs8 10G 1.9G 8.2G 19%	
Quota usage for your research directory	
Filesystem Size Used Avail Use%	
/econ/research/ofs8 50G 0 50G 0%	

Figure 4: SSH From Outside Duke Network.

### 2.2 Windows

To connect to the cluster when working in the Windows operating system, you will need to install an SSH client. I recommend installing PuTTY due to its simplicity and convenience (it is also free). After installing the software, you will use **ssh** but with a graphical user interface.

Open PuTTY and type login.econ.duke.edu on the Host Name window (see Figure 5). You can now click Open to connect to the cluster. You will be prompted for your username and your password.

alegoly.		
Session	Basic options for your PuTTY session	
Logging Terminal Keyboard Bell	Specify the destination you want to connect to Host Name (or IP address) Port 22	
- Window	Connection type: Raw Telnet Rlogin SSH Ser	ial
Appearance Behaviour Translation Selection	Load, save or delete a stored session Saved Sessions	
- Colours - Connection - Data - Proxy - Telnet - Riogin	Default Settings Load Save Delete	•
SSH Serial	Close window on exit: Always Never  Only on clean exit	

Figure 5: PuTTY SSH Client.

After typing your password, you should be greeted with a welcome message and are now connected to the cluster. You are connected to a bash terminal, which can take commands and will execute them. We will discuss these commands in the next section.

## 3 Using a Bash Terminal

Now that we are connected to the cluster, we can discuss how to use the terminal. The terminal window you see is a bash shell. A shell is just a user interface to the underlying operating system, and bash refers to the type of interface.

We can interact with the shell either by typing commands and executing them line by line, or by creating script files. We will cover some basic commands of the bash shell.

### 3.1 Working Directory

The command pwd prints the current working directory:

### pwd

You can change the working directory with the cd command:

```
# syntax: cd folder
1
2
  cd /econ
3
  pwd
  cd /econ/home
4
5
  pwd
  cd /econ/home/g/gfs8
6
7
  pwd
8
                                      # .. refers to the parent
  cd ..
     folder
9
  pwd
```

While . . represents the parent folder, there is a shortcut for your home folder as well:

```
1 cd ~
2 # ~ represents the home folder
3 # alternatively
4 cd $HOME
```

## 3.2 Creating Folders

You have permission to change things around only in your home folder. In my case, my home folder is the folder /econ/home/g/gfs8. We can create a new folder with the command mkdir:

```
1 # syntax: mkdir folder_name
2 mkdir Matlab
3 mkdir Test
4 # create multiple folders
5 mkdir A B C
```

## 3.3 Listing Files and Folders

We can list all files and folders inside a folder with the command ls.

```
1 # syntax: ls
2 ls
3 # display one file or folder per line
4 ls -1
```

You can get a full list of the options a command accepts by reading the manual page of the command. You can access the manual page of a command using man:

```
1 # syntax: man command_name
2 man ls
```

## 3.4 Deleting Files and Folders

We can delete an empty folder with the command **rmdir**.

```
1 # syntax: rmdir folder_name
2 rmdir Test
3 # remove multiple folders
4 rmdir A B C
```

To remove a non-empty folder we need to use the more versatile command rm with the option -r:

```
1 # create a non-empty folder
2 mkdir Test Test/A Test/B
3 # check it is non-empty
4 ls -1 Test
5 # equivalent to
```

```
cd Test
6
7
   ls -1
8
  cd ..
9
  # try to remove with rmdir
10
  rmdir Test
                                       # error
11
  # use rm −r
12
  rm -r Test
                                       # works
```

The command  ${\tt rm}$  can also be used to remove files.

### 3.5 Creating Files

We can create empty files with the touch command:

```
1
  # syntax: touch file_name
2
  touch test.txt
3
  # create multiple files
  touch a.csv b.jpg
4
  # remove files
5
6
  rm test.txt
7
  # remove multiple files
8
  rm a.csv b.jpg
```

We can also create and edit files. To edit a file we need an editor. Some of the editors available inside the bash terminal are: nano, vim and emacs. You can create or edit a file with these editors by typing the name of the editor followed by the name of the file. The editor nano is the most straightforward to use, however, the editors vim and emacs are extremely powerful and might be worth to learn if you often use your computer to type text into files.

```
# create a new file with nano
1
2
  nano data.csv
3
  # nano will now open with an empty file
4
  # type in:
5
  # 1,21,0
  # 2,28,25000
6
7
  # 3,35,70000
8
  # then use Ctrl-O to save the file, and then Ctrl-X to exit
     nano
9
  # create another file
10
  nano description.txt
11
  # type in:
12
  # id,age,income
```

### 3.6 Inspecting Files

You can quickly inspect the contents of a file with the cat command.

```
1 # check name of files
2 ls -1
```

```
3 # see contents of data.csv
4 cat data.csv
5 # see contents of description.txt and data.csv
6 cat description.txt data.csv
```

If the file is too big and you do not want to display its entirety on the screen, you can use the head and tail commands. The head command displays the first few lines of a file, while the tail command displays the last lines of a file.

```
# see first lines of the file
1
2
  head data.csv
3
  # see only first two lines of the file
  head -n 2 data.csv
4
  # see only the first line of the file
5
6
  head -n 1 data.csv
7
  # see last lines of the file
8
  tail data.csv
  # see only very last line of the file
9
  tail -n 1 data.csv
10
  # display first line of multiple files
11
  head -n 1 description.txt data.csv
12
```

### 3.7 Copying and Moving Files

To copy files use the cp command.

```
1 # syntax: cp source_file target_file
2 # create a copy of data.csv
3 cp data.csv data_copy.csv
4 # create a copy of the data inside a data folder
5 mkdir Data
6 cp data.csv Data/
7 ls Data
```

Files can be moved with the mv command.

```
1
  # remove the copy of data.csv from the Data folder
2
  rm Data/data.csv
3
  ls Data
  # move the original data file to the Data folder
4
  mv data.csv Data/
5
  # check the file changed folder
6
7
  ls
8
  ls Data
```

## 4 Bash scripts

A bash script is just like a Matlab script: it is a text file containing commands that should be executed line by line. We can create a bash script by creating a file with the extension .sh. Let's use the command echo to print strings to the terminal window and save them in a script file.

```
1 # syntax: echo string
2 echo "Hello there!"
3 echo "How are you?"
```

Save it in a file (nano greetings.sh):

```
1 # greetings.sh
2 echo "Hello there!"
3 echo "This is a bash terminal."
4 echo "Welcome."
5 head data.csv
```

We can execute this file with the command source. The source command takes a script file and executes it line by line in the current shell.

```
1 # syntax: source script_name
2 source greetings.sh
```

And you should see the three messages displayed in the shell.

We can use a shell script to execute several programs, including programs in other languages. To do so, the shell must be able to find other programs. For example, when we typed **nano** before, the shell searched for the program **nano** and then executed it with the parameters we passed it. We can check whether the terminal can find a program with the command **which**.

```
1 # syntax: which program_name
2 # if the program can be found, then the shell
3 # displays the path to the program binaries
4 which nano
5 which emacs
6 # if the program cannot be found, then nothing is displayed
7 which foo123
```

If the program cannot be found, then it could either not be available in the machine, or it could be outside of the path of the terminal. The path is a list of folders where the terminal searches for programs. If the program cannot be found in those folders, then the terminal does not return anything. However, if you know in which folder the program lives, then you can specify the full path to it. Alternatively, you can add its folder to the search path (this is left for another time, as it introduces some complications). For example, in my machine matlab is not in the search path, but it can be found by specifying the full path:

```
1 # on local machine
2 which matlab
3 # returns nothing
4 # but we can specify the full path
5 which /Applications/MATLAB_R2019a.app/bin/matlab
```

On the cluster, however, matlab should be in the search path:

```
1
  ssh gfs8@login.econ.duke.edu
2
  which matlab
3
  # returns /usr/local/bin/matlab
```

Let's create a Matlab script that will perform some computation and save the results.

```
% matlab computation.m
1
2
  % Do some computation
3
  x = rand(200, 1);
  y = rand(200, 1);
4
5
  res = x.*y.^2 + 1;
6
  % Save results
7
  save('results.mat', 'res');
```

We can now write a bash script that will interact with the operating system and then execute the Matlab script. To execute the Matlab script we would call matlab test.m. However, matlab also takes some options that can speed up its execution. We can pass a script to be executed with matlab by using the option -batch and specifying the script name, without the extension .m. The complete set of options matlab takes is described in this reference page.

```
1
2
```

4

```
# matlab from bash.sh
  echo "Executes the Matlab script: matlab_computation.m"
3
  # execute Matlab script
  matlab -batch "matlab computation"
```

We can now execute this script with source matlab from bash.sh, and the script will be executed. Notice that at the end of the script, Matlab saves the results in the file results.mat, which later we can import into our computer to analyze.

#### Slurm: Scheduling Tasks 5

Bash scripts are important because they are how we can interact with the Econ Cluster. Remember, the Cluster is simply a collection of computers being managed by some software. The software that manages the Econ Cluster is the Slurm Workload Manager.

In the previous section, we logged in the log-in node of the cluster. At that node, we can interact with the cluster via the terminal, and even execute some Matlab code via a bash script. However, Slurm will not allow us to execute a lot of code, or code that takes too long to run. Instead, Slurm allows us to submit tasks for it to run. That is, we can give Slurm a bash script, and it will allocate the script to some computer in the cluster, run it, and save the results in your home folder. In doing so, Slurm allows all users of the Cluster access to powerful computers, but there may be a queue.

#### 5.1Shebang

The script matlab from bash.sh is almost ready to be submitted for execution with Slurm. It is only missing a shebang line. The shebang is the very first line of a text file used as a script, which specifies the interpreter that should be used when executing the file. While we could execute our script with **source**, by adding a shebang to the file, we can execute it as an executable file. Create the following test script:

```
1 #!/bin/bash
2 # test_shebang.sh
3 echo "Hello!"
```

We can still execute it with source:

```
1 source test_shebang.sh
```

But now, we can execute it as an executable file:

```
1 # make the file executable
2 chmod +x test_shebang.sh
3 # execute it
4 ./test_shebang.sh
```

The first line of the script tells the terminal that it should use **bash** to interpret its contents. Slurm needs this shebang to correctly submit the script.

### 5.2 Submitting to the Cluster

Modify the matlab\_from\_bash.sh so that its first line is #!/bin/bash. Delete the results.mat file from the folder.

We can submit the matlab\_from\_bash.sh script to Slurm with the command sbatch.

```
1# submit script to Slurm2sbatch matlab_from_bash.sh
```

If there is no immediate issue with the script, Slurm will schedule the script for execution in one of the computers (nodes) of the cluster. The scheduled script is called a job. Slurm will print the job number on the screen.

### 5.3 The Queue

After submitting a script for execution, you should have its job number. You can use this number to check what is the state of the job. We can use the command squeue to see all jobs currently scheduled and being executed. See Figure 6 for the output of squeue

ofs89login=01 ~ls	soueue							
JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)	
33759	common	matlab f	afs8	PD	0:00	1	(None)	
33742	conmon	estima1	rd123	R	17:51:18	1	bafcnm-01	
33743	contron	estima2	rd123	R	17:49:26	1	comp-node-18	
33744	connon	estima3	rd123	R	17:47:37	1	comp-node-19	
33745	common	estima4	rd123	R	17:45:50	1	bafonm-01	
33746	common	estima5	rd123	R	17:44:02	1	bafcnm-02	
33747	contron	estima6	rd123	R	17:42:12	1	comp-node-18	
33748	contron	estima7	rd123	R	17:40:24	1	comp-node-19	
33749	contron	estima8	rd123	R	17:38:36	1	bafcnm-01	
33750	common	estima9	rd123	R	17:36:48	1	bafcnm-02	
33751	common	estima10	rd123	R	17:35:01	1	comp-node-18	
33753 gfs80login-01 ~]\$	common	matbatch	vv34	R	11:56:48	1	comp-node-19	

Figure 6: Slurm queue after submitting a job.

The job we just submitted has the number 33759. You can see in the second line after the command squeue that the job has been scheduled. The NAME shows the name of the script submitted, USER shows the user who submitted the job (net id), ST shows the state of the job (PD stands for pending, while R stands for running). The TIME shows for how long the job has been running. In the case of our job it has not started yet, so TIME is 0:00. Notice that some other users have jobs running for more than 17 hours! The NODELIST describes the reason why the job is in the queue. In the case of the job we submitted, (None) means the job will be executed next. Usually, if there are not enough resources available to execute your script, you will see a (Resources), which indicates Slurm is waiting for other scripts to finish before executing yours. When the script is being executed, NODELIST will show the node (computer) where the script is running.

If we have multiple jobs running, then we can check the status of the jobs submitted only by us (not by everyone). We can do so with the option -u username, which displays squeue but only for the specified username.

squeue -u gfs8
----------------

1

See Figure 7 for an example.



Figure 7: View of Slurm queue for a specified user.

In this case I submitted several jobs. The queue shows they are all running, some are running on different nodes, and they have been running for a few seconds already. If you run squeue -u gfs8 again, then you might see fewer jobs. This is because when a job is completed, it does not show up in the queue anymore.

### 5.4 Details on a Job

We can use the **sacct** to get information on running and completed jobs we have submitted to Slurm. Running **sacct** will display details for all jobs submitted by you on a certain period of time. If you have the job number, then you can use pass it with the option -j to get details about that specific job (see Figure 8):

sacct -j 33759



Figure 8: View details of a submitted job.

If you want even more information about a specific job, then you can use the command sacct\_ec:

```
sacct_ec -j 33759
```

Figure 9 displays the output of sacct\_ec. Notice it shows the node where the script was executed, and even the start and end time of the script.

```
    Control = 10 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 00000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 + 00000 + 0000 + 0000 + 0000 + 0000 + 0000 + 0000 +
```

Figure 9: View even more details of a submitted job.

These commands are useful when you submit jobs that might take a long time to execute, or when debugging scripts that are failing to execute.

### 5.5 Canceling a Job

You can cancel a job using the scancel command. For example, if the job number is 33759, then scancel 33759 will cancel the job. You can only cancel jobs you started yourself. If your script is taking longer than expected to complete, then you might have some bug in your code that is causing it to hang. In this case, canceling the job might be required.

### 5.6 Outputs of the Job

When a job is submitted and starts running, a file with the name slurm-XXXXX.out is created. The XXXXX represents the number of the corresponding job. This file contains whatever your script prints to the screen. For example, when we executed the matlab\_from\_bash.sh script, it printed a couple of lines on the terminal. If we submit this script to Slurm, then the lines that would be printed on the terminal are saved on the slurm-XXXXX.out file.

The .out file can be used as a log of what is happening in your script. You can use it to debug your program, since debugging in the cluster is not as straightforward as debugging in your local machine (you cannot stop the execution of the code and inspect variables, for example).

The Matlab script we submitted also created a file containing the results. This file is also saved in our home folder. When we submit a job with Slurm, our home folder is directly accessible by the job, and behaves as a local drive.

## 6 Slurm: Partitions and Nodes

Slurm organizes the cluster in partitions, where each partition is a set of compute nodes (computers used to run code). We can get an overview of the partitions available in the Econ Cluster with the sinfo command. Figure 10 displays the output of sinfo.



Figure 10: Overview of cluster partitions.

Notice there are two partitions, common\* and common-lm. Different partitions might have different purposes. For example, there could be a partition for debugging code, and another for actually submitting tasks. In this case, the two partitions are for computing. The default partition is marked with an asterisk. That is, the partition common is where jobs are submitted to by default.

There are four compute nodes in the default partition. Some of them are in use, other are idle, leading to the mix state. In the common-lm partition there are two nodes, which are idle. The NODELIST column gives the names of the nodes. The nodes in the common partition are bafcnm-01, bafcnm-02, comp-node-18 and comp-node-19. The nodes in the common-lm partition are bafclnm-01 and bafcnm-02.

Notice that the column TIMELIMIT says infinite. This is the time limit to which jobs are subject. In this case, there is no time limit.

We can display information organized by node instead of partition with the option -N. We can also use the option -1 to display extra information.

sinfo -N -l

Figure 11 displays the output of sinfo -N -1.



Figure 11: Overview of cluster nodes.

Notice that some of the nodes have 512 GB of memory, and the nodes in the default partition have 64 GB of memory. The nodes in the common partition have more CPUS than the nodes in the common-1m partition. The S:C:T column displays the number of sockets in each motherboard<sup>1</sup>, the number of cores in each processor, and the number of threads in each core. If you multiple these numbers you get the number in the CPUS column. In terms of parallel computing, the number of CPUS is more or less equivalent to the number of parallel process that can be used with parfor.

It is important to know that when you submit a job to Slurm, it does not mean the job will use an entire compute node to execute. That is, a single compute node can execute multiple jobs at the same time. For example, a single job might use just one of the CPUS of a compute node, so that the other CPUS can be allocated to other jobs. It is also possible to have a job use multiple CPUS, including CPUS from more than one

<sup>&</sup>lt;sup>1</sup>The socket is the physical space in a computer's motherboard where the physical CPU is placed. Most common motherboards only have one socket. However, motherboards for cluster computers usually have more than one socket. This is the case for the computers in the Econ Cluster. Some of them have 4 sockets, while others have 8 sockets. That is, a single computer can have 4 or 8 different physical CPUs.

compute node. On the next section we will see how to request CPUS and memory when submitting jobs to Slurm.

# 7 Slurm: Requirements and Directives

When it comes to speeding up code execution, we saw that Matlab allows us to use more than one processor core to speed up loops with **parfor**. With the Econ Cluster, we can submit a job and request a certain number of CPUS to be available. Additionally, we can also request a certain amount of memory, so that we do not run into memory issues when using **parfor**. We will see how to specify these requirements using directives.

## 7.1 Directives

When we submit a script with sbatch, Slurm checks the script for directives. Directives are lines that start with #SBATCH. We can add directives that tell Slurm what resources are required to run the script. Slurm then processes these directives and waits until the resources are available to start executing your script. This allows us to specify, for example, how many CPUS and memory we need to run the script.

## 7.2 Requesting Memory and CPUs

Let's modify the matlab\_from\_bash.sh script to add some directives.

```
#!/bin/bash
1
2
  # matlab_from_bash.sh
3
  # require 12 GB of memory for this job
4
  #SBATCH --mem=12G
  # require 4 CPUS for this job
5
  #SBATCH --cpus-per-task=4
6
7
  echo "Executes the Matlab script: matlab_computation.m"
8
  # execute Matlab script
9
  matlab -batch "matlab_computation"
```

The first directive requires a large amount of memory (12 GB). Remember that the more workers you have, the more memory you need, since Matlab needs to distribute all information in your workspace to all of the workers. If you run into problems executing a parallel code in the cluster, it might be because you requested too little memory. If this is the case, either request more memory or reduce the number of workers you require. By using the --mem directive, we know that when our code is executed, at least 12 GB of memory will be available.

There is a caveat with the memory directive in Slurm. If at some point our code requires more than 12 GB of memory, Slurm will kill the execution of the job. This happens even if more memory is available at the node. That is, if the node has 64 GB of free memory, but your code uses more than the 12 GB that were requested, then Slurm will kill the job. This means that the **--mem** directive is binding. You must make sure that your program does not use more memory than what was requested.

The second directive requires a node that has at least four available CPUS. Contrary to the memory directive, the --cpus-per-task directive will allow your code to use more CPUS if available. That is, if Slurm starts executing your job and all CPUS of the node

are not in use, then your code may receive more CPUS than were requested. This allows us to use **parfor** to speed the execution of the code. If more resources are available, then we might even get more CPUS than requested.

Let's modify the matlab\_computation.m script to run a parfor loop.

```
1
   % matlab computation.m
2
   % Check how many cores were allocated to the job
3
   feature('numcores')
4
   % Start pool of workers
   ppool = parpool();
5
6
   disp(sprintf('Total Workers: %d', ppool.NumWorkers));
7
   % Run parfor with increasingly more workers
8
   for w = 1:ppool.NumWorkers
9
       disp(sprintf('Workers: %d', w));
       tic;
11
       parfor (i=1:10, w)
12
           pause(1)
13
       end
14
       toc;
15
   end
```

The first line tells Matlab to check how many cores are available. It prints out information on how many cores are available in total, and how many are actually being used by Matlab. The number of cores will determine how many workers can be used with parfor.

The second line initializes the pool of workers. By calling parpool() we create the pool of workers with the default number of workers. If you are worried about memory limitations (remember: more workers require more memory), then you can call parpool with the number of workers you require. For example, parpool(4) would initialize a pool of workers with only four workers, even if more are available.

The third line of code displays how many workers are available in the pool. The foorloop iterates over the number of available workers. It starts with a single worker, and times the execution of a parfor-loop. Observe that if there is only one worker, then the parfor-loop is a regular loop. It will iterate over the variable i, pausing for one second at every iteration. When w is one, this parfor-loop should take about 10 seconds to complete.

On the next iteration in the outermost loop, the variable **w** becomes two. Now, we run the parfor-loop with two workers. This means that the command **pause(1)** will be assigned twice, once to each worker, so the loop will execute faster. Indeed, it should take about 5 seconds to execute, since we two workers paused at each time.

Since the number of workers in the cluster can be high, the time to execute the parforloop code above will get progressively smaller. After submitting the code with sbatch, we can inspect the output file. Figure 12 displays the output.



Figure 12: Output of Matlab script using parfor on the cluster.

In this case, even though we only requested four CPUS, many more were available. Indeed, we could create a pool of workers with twelve workers! Notice that there were actually fourteen cores available to Matlab, but the default number of workers from the parpool command is twelve. We could have created a pool with the fourteen workers by calling parpool(14). Alternatively, we could have called parpool(feature('numcores')) to use all available cores.

## 7.3 Multiple Jobs and Transferring Files

We have discussed how to submit a single job that uses many cores to speed up computation. An alternative paradigm is submitting several different jobs that use a single core, and then collecting all of their results.

For example, consider you have a statistic that you are interested in studying via bootstrap, but that takes a long time to be computed (potentially because of a large sample, or because the statistic is complicated). In this case, we could launch several different jobs, where each computes a bootstrap sample and the statistic just once. We can then accumulate the results to study the distribution of the statistic.

Let's create some fake data to bootstrap.

```
1
2
3
4
```

5 6

```
% generate_fake_data.m
mu = 20.5;
sigma = 5.8;
sample_size = 1000000;
sample = normrnd(mu, sigma, sample_size, 1);
save('fake_data.mat', 'sample');
```

We will use this data to compute the confidence interval for the mean via bootstrap. We need to get the data from our computer into the cluster. To do so, we will use the command scp, which is similar to the ssh command, but is used to copy files securely. If you are transferring a lot of files, or want a graphical user interface, then you could use the program Cyberduck, for example.

The syntax for **scp** is the following:

### 1 scp file\_to\_copy username@host:~

The command will take the file named file\_to\_copy and copy it to the folder ~ in the host. Remember that ~ represents your home folder.

At the time of writing, I am outside of the Duke network, so there are two steps to copy the data from my laptop to the Econ Cluster. First, I need to transfer the data from my laptop to the Duke login node. Then, I **ssh** into the Duke login node and transfer the data from there to the Econ Cluster.

```
# Locate the file on my laptop
1
2
  ls -1
3
  # fake data.mat is in the current working directory
  scp fake_data.mat gfs8@login.oit.duke.edu:~
4
  # will ask for your password and then start transferring the
5
     file
6
  # Now, ssh into the login node of Duke
  ssh gfs8@login.oit.duke.edu
7
8
  # check that fake data.mat is in the current working
     directory
9
  ls -1
  # transfer to the Econ Cluster
10
11
  scp fake data.mat gfs8@login.econ.duke.edu:~
12
  # ssh into the Econ Cluster
13
  ssh gfs8@login.econ.duke.edu
14
  # verify that the file was transferred
  ls -1
15
```

Now, let's create the bootstrapping script.

```
% bootstrap_mean_fake_data.m
1
2
  % Load data
  load('fake data.mat', 'sample');
3
  % Bootstrap a new sample
4
  ssize = length(sample);
5
  new sample = sample(unidrnd(ssize, ssize, 1));
6
7
  % Compute statistic
  statistic = mean(new_sample);
8
9
  % Save result to csv
10
  % Question: but what should be the name of the file?
```

The script loads the fake data, generates a bootstrap sample, computes the statistic and saves it in a .csv file. However, we have a problem. What should be the name of the .csv file? Since we are launching multiple jobs at the same time, how do we set the name of the file programmatically? Before answering these questions, let's create the bash script that will launch several jobs, where each job will execute the Matlab script above once.

```
1 #!/bin/bash
2 # multiple_jobs.sh
3 #SBATCH --cpus-per-task=1
4 #SBATCH --mem=1GB
5 #SBATCH --array=1-10
6 matlab -batch "bootstrap_mean_fake_data"
```

We can use the --array directive to submit multiple jobs with Slurm. The directive --array=1-10 will launch ten different jobs, numbered from 1 to 10. Each job will execute our Matlab script. Since each job has a different number, each job will create a separate .out output file. However, this output file is just what is displayed in the terminal, and is mainly used for debugging. What we want to recover is the file containing the statistic we computed.

When we submit a job with --array, a variable is created with the number of the number of the job. We can pass this variable to the Matlab script and use it to save the statistic value. The variable name that holds the job number is SLURM\_ARRAY\_TASK\_ID. We can access its value in the bash script using the command \$SLURM\_ARRAY\_TASK\_ID. We can then pass this value to the matlab script:

Now, inside the boostrap\_mean\_fake\_data.m script we can access the variable job\_number:

```
1
   % bootstrap mean fake data.m
2
   % Set seed
   rng(job_number);
3
4
   % Load data
   load('fake data.mat', 'sample');
5
  % Bootstrap a new sample
6
7
   ssize = length(sample);
   new sample = sample(unidrnd(ssize, ssize, 1));
8
   % Compute statistic
9
   statistic = mean(new sample);
10
   % Save result to csv
11
  filename = strcat(num2str(job_number), '.csv');
12
  csvwrite(filename, statistic);
13
```

The variable job\_number is initialized by the bash script. We use it to create the .csv file, so that each job, which has its own unique number, will create a unique .csv file. This also has the advantage that if some job fails (a bug, or takes too much resource, or bad weather), then we can quickly find it. We also used the job\_number to set the seed for random number generation. This way when we generate the bootstrap samples they

are not all going to be the same, and the output will be repeatable if executed with the same job\_number.

We can submit the script to Slurm and verify that indeed 10 jobs are created:

```
# Submit script for execution
1
2
  sbatch multiple_jobs.sh
  # A job id is assigned.
3
  # Look at all jobs that were created:
4
5
  squeue -u gfs8
6
  # Notice that JOBID is xxxxxx_y, where xxxxxx is the job id,
     and
7
  # y is the job number specified in the --array directive
  # Observe that many slurm-xxxxxx_y.out files were created:
8
  ls -1 slurm-*
9
  # The * above expands slurm- to all existing filenames that
10
11
  # begin with slurm-
12
  # Jobs should be done now:
13
  squeue -u gfs8
14
  # Notice that the .csv files were created
15
  ls -1 *.csv
16
  # We can sort the names numerically with the option -v
17
  ls -1 -v *.csv
18
  # We can print them on the screen with cat
19
  cat *.csv
```

The directive --array takes a range of numbers, like 1-10, and launches jobs using those numbers to create the job ids. However, the range of numbers can be discontinuous. For example, the directive --array=1-3,5,7-10 would launch jobs with ids 1, 2, 3, 5, 7, 8, 9 and 10, skipping the ids 4 and 6. It is also possible to use --array=7 to launch a single job with that id number. The single value or discontinuous --array directives are useful for resubmitting a specific job that failed.

# 8 Assignment

**Problem 1** Modify the scripts multiple\_jobs.sh and bootstrap\_mean\_fake\_data.m so that a new folder is created to store the resulting .csv files. If the folder already exists, the files inside it should be deleted before the script runs.

**Problem 2** (Optional) Create a bash script that takes all of the .csv files outputted from multiple\_jobs.sh and creates a single .csv file with all of the results.

**Problem 3** Create a Matlab script that takes all of the .csv files outputted from multiple\_jobs.sh and creates a single .csv file with all of the results.

**Problem 4** Compare the advantages and disadvantages of the following takes on parallel computing with the cluster:

- Launching a single job that uses multiple cores;
- Launching several jobs that use a single core.

**Problem 6** Consider a deterministic growth model, where an agent decides between consumption  $(c_t)$  and investment in capital  $(k_t)$ , while maximizing his utility. We can write this problem as:

$$\max \sum_{t=0}^{\infty} \beta^{t} U(c_{t})$$
  
subject to 
$$\begin{cases} k_{t+1} = k_{t}^{\alpha} - c_{t} + (1-\delta)k_{t}, \forall t >= 0\\ k_{0} > 0 \end{cases}$$

Write the problem as a Bellman equation. Let  $U(c; \sigma) = \frac{c^{1-\sigma}-1}{1-\sigma}$ . Obtain the Euler equation for this problem in terms of the consumption c. Solve the problem by Value Function Iteration. Consider  $\sigma = 2$ ,  $\beta = 0.95$ ,  $\delta = 0.1$  and  $\alpha = 0.33$ . Use the steady state value of k to create a grid for the possible values of k, say 100 points between  $0.25k^*$  and  $1.75k^*$ . Start with a guess for V over the grid, for example V(k) = 0 for all k in the grid. Use the Matlab minimization function to solve for k. You may want to add the constraint that c should always be positive.

Use **parfor** and the Econ Cluster to speed up the solution of this problem. Compare the speed gain with solving the problem using your local computer.