

# Abel tutorial for beginners

---

*Katerina Michalickova*

*November 5<sup>th</sup> , 2014*

*The Research Computing Service Group*

*<http://www.uio.no/hpc>*

## Contents

1. Resources .....	3
2. Log into Abel .....	3
2.1 From a Windows machine.....	3
2.2 From a Linux machine .....	4
2.3 From a Mac .....	5
3. Home area on Abel.....	6
3.1 Home directory.....	6
4. Unix command line environment.....	7
4.1 Tutorials .....	7
4.2 Shell scripting exercise.....	7
5. Prepare and run your first job.....	9
5.1 Software and modules .....	9
5.2 Job script .....	9
5.3 Submit your job and follow the progress .....	11
5.4 Job output .....	12
6. Prepare a job using scratch area .....	13
6.1 Job script .....	13
6.2 Upload you own file .....	14
7. Parallel jobs .....	17
7.1 Arrayrun .....	17
7.3. Arrayrun job scripts.....	18
7.3 Arrayrun results.....	20

## 1. Resources

This tutorial is intended as a start for people who would like to use Abel. Links below point to a complete user guide and information about the system.

Abel page: <http://www.uio.no/english/services/it/research/hpc/abel/>

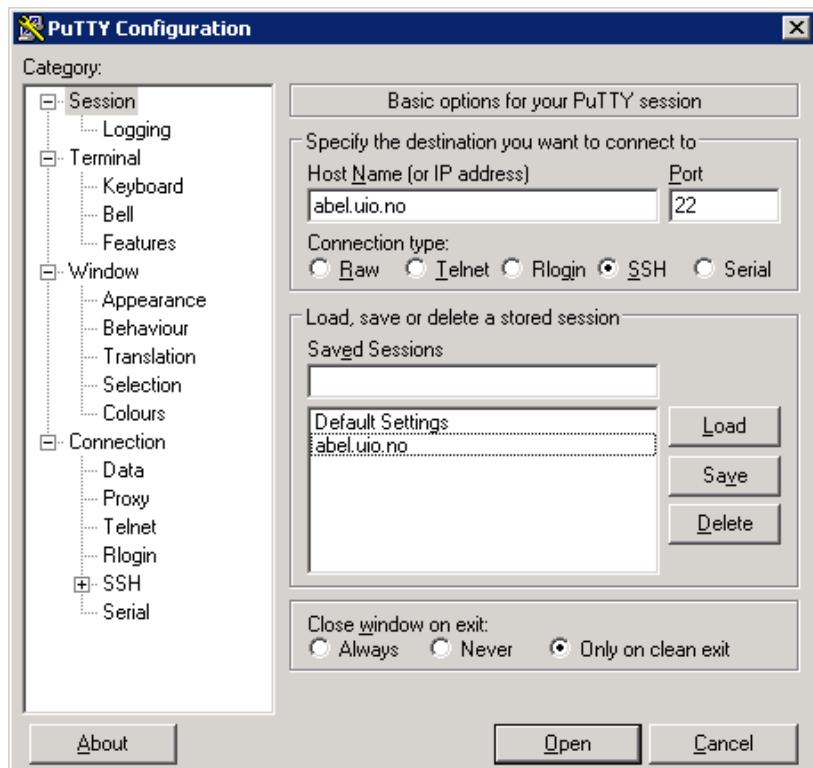
Abel user guide: <http://www.uio.no/english/services/it/research/hpc/abel/help/user-guide/>

## 2. Log into Abel

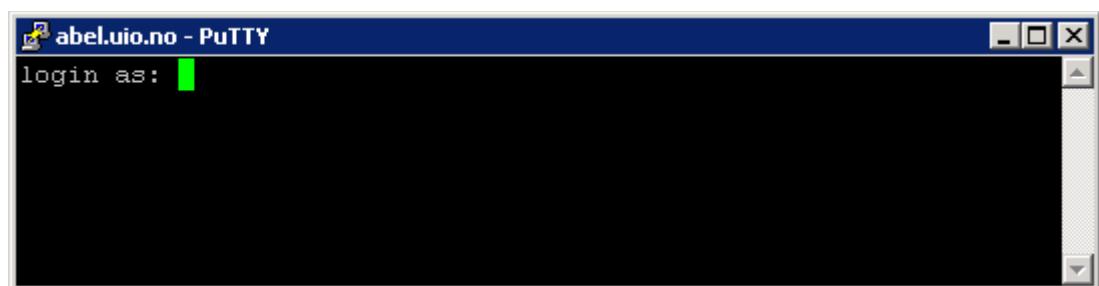
### 2.1 From a Windows machine

Download putty, a windows ssh client, from <http://www.putty.org/>

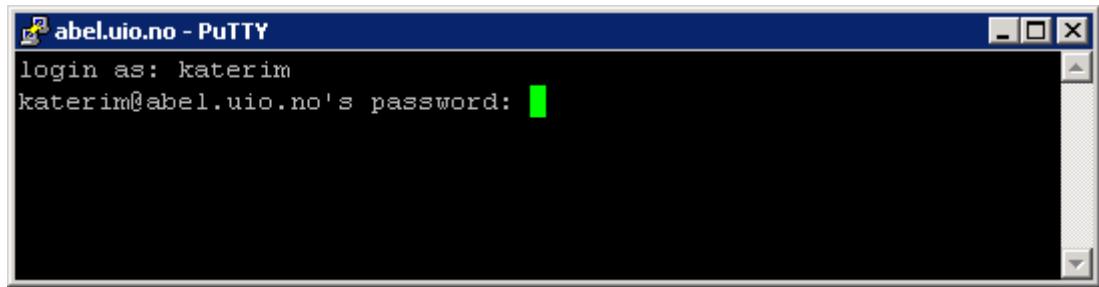
Start putty and type “**abel.uio.no**” into the host mane dialog box:



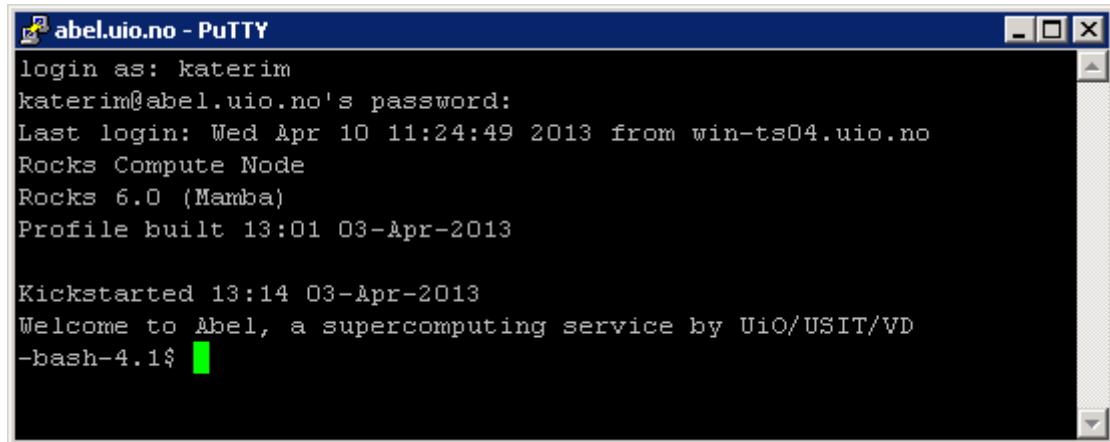
Press Open button:



Type in your UiO user name (and press Enter):

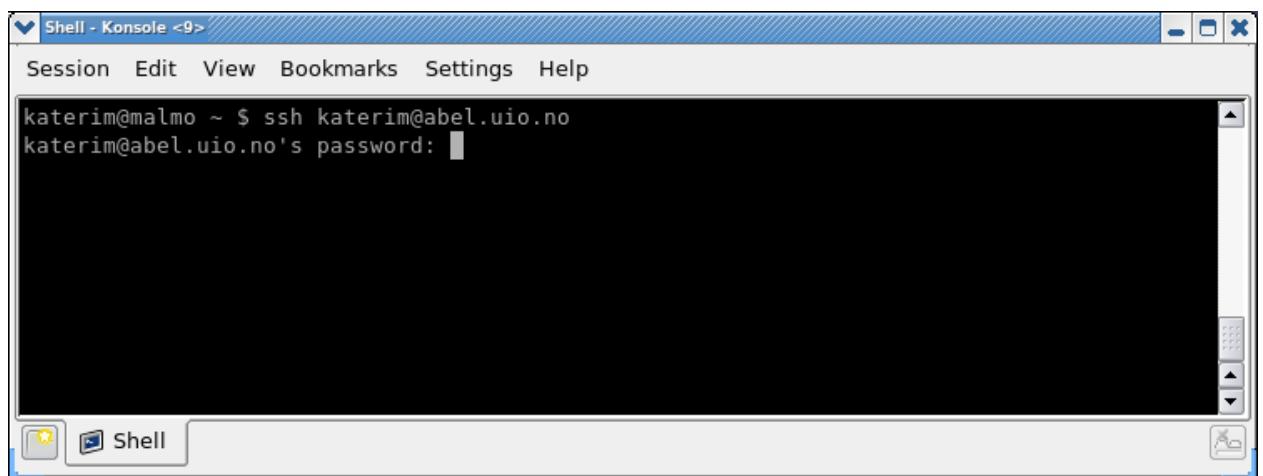


Type in your UiO password (and press Enter):

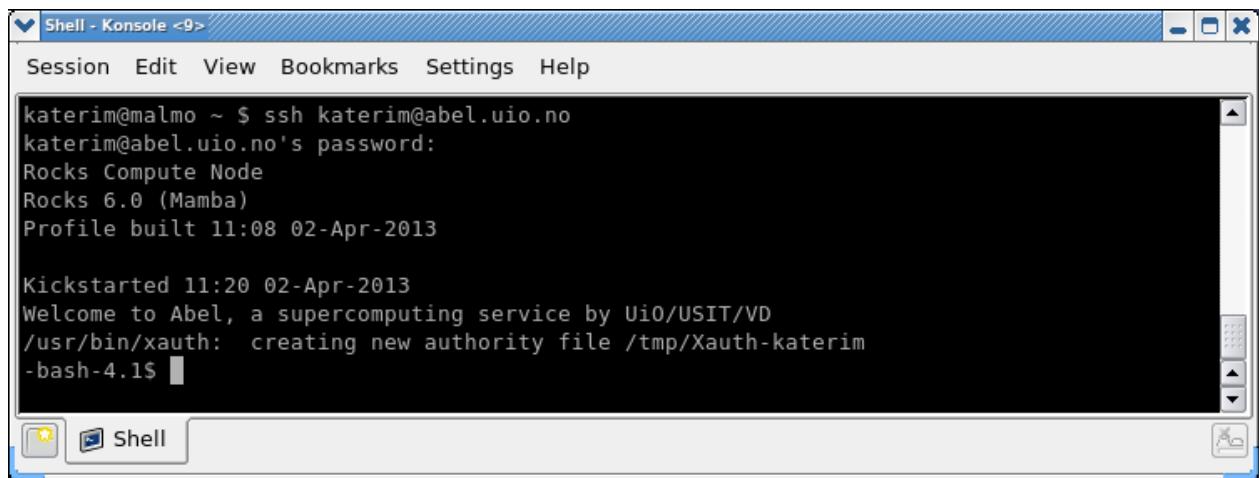


## **2.2 From a Linux machine**

Open a terminal and type “`ssh your_UiO_user_name@abel.uio.no`”:



Type in your UiO password:



The screenshot shows a terminal window titled "Shell - Konsole <9>". The window has a menu bar with "Session", "Edit", "View", "Bookmarks", "Settings", and "Help". The main pane displays the following text output from an SSH session:

```
katerim@malmo ~ $ ssh katerim@abel.uio.no
katerim@abel.uio.no's password:
Rocks Compute Node
Rocks 6.0 (Mamba)
Profile built 11:08 02-Apr-2013

Kickstarted 11:20 02-Apr-2013
Welcome to Abel, a supercomputing service by UiO/USIT/VD
/usr/bin/xauth:  creating new authority file /tmp/Xauth-katerim
-bash-4.1$
```

The bottom of the window shows a toolbar with icons for "Shell" and other functions.

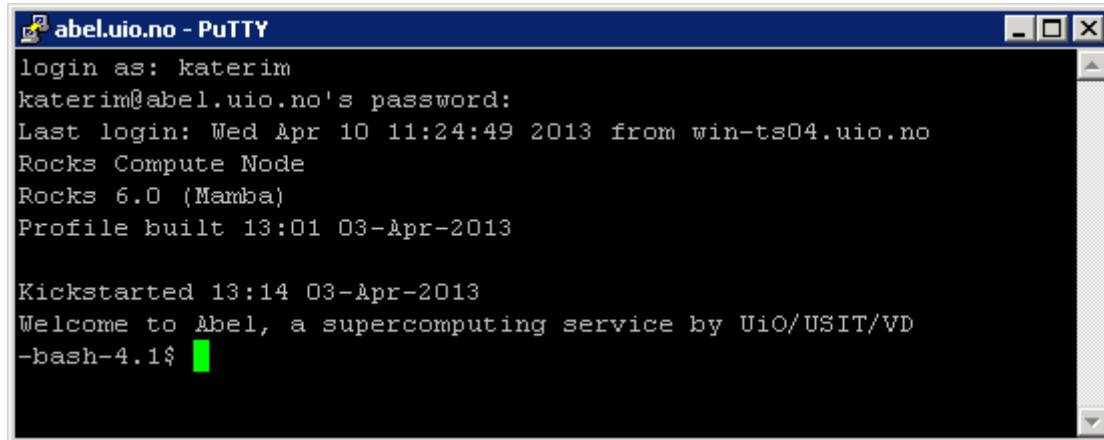
### **2.3 From a Mac**

Open terminal app and type type “`ssh your_UiO_user_name@abel.uio.no`”. Proceed in the same way as on the linux system.

### 3. Home area on Abel

#### 3.1 Home directory

You are logged in and it looks like this:

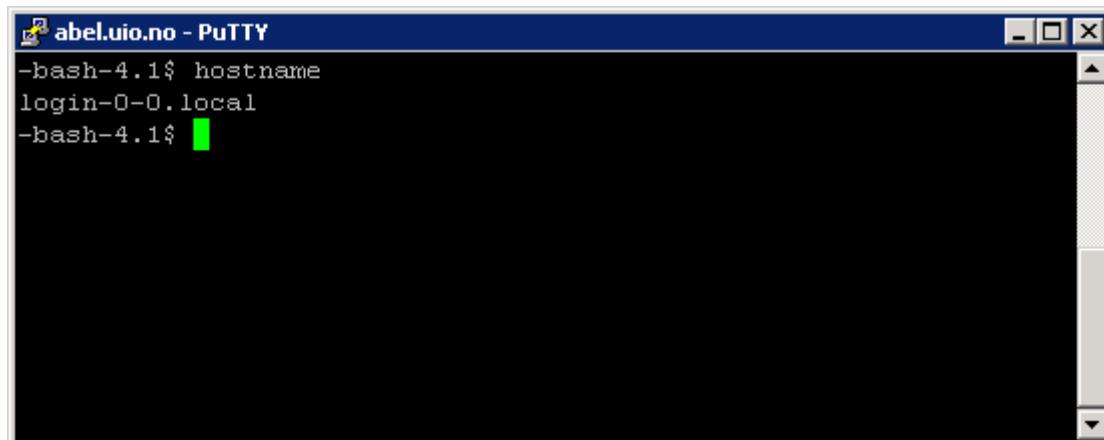


The screenshot shows a PuTTY terminal window titled "abel.uio.no - PuTTY". The session has been successfully established, displaying the following system information:

```
login as: katerim
katerim@abel.uio.no's password:
Last login: Wed Apr 10 11:24:49 2013 from win-ts04.uio.no
Rocks Compute Node
Rocks 6.0 (Mamba)
Profile built 13:01 03-Apr-2013

Kickstarted 13:14 03-Apr-2013
Welcome to Abel, a supercomputing service by UiO/USIT/VD
-bash-4.1$
```

Which machine are you on? Type “**hostname**”:



The screenshot shows a PuTTY terminal window titled "abel.uio.no - PuTTY". The user has run the "hostname" command, which returns the following result:

```
-bash-4.1$ hostname
login-0-0.local
-bash-4.1$
```

You are on one of the two so-called login nodes of the Abel computer cluster. Abel consists of 600+ computers and only two are available for login. If you are interested, you can read about Abel technical specifications here:

<http://www.uio.no/english/services/it/research/hpc/abel/more/>

When you login into Abel, you always find yourself in your “home directory”. It is a place in the directory tree where you keep your files. Regardless what machine you are on, this directory is available to you. Type “pwd” (print working directory) to see the full path:

## 4. Unix command line environment

### 4.1 Tutorials

Users need some knowledge of the command line environment and shell scripting to efficiently work on Abel. Below is a link to the RCS command line tutorial for those who missed it.

[http://www.uio.no/english/services/it/research/hpc/courses/linux-command-line/unix\\_tutorial\\_nov2014.pdf](http://www.uio.no/english/services/it/research/hpc/courses/linux-command-line/unix_tutorial_nov2014.pdf)

For more examples, I recommend to review online tutorials, find one that suits you and read through:

<http://www.ee.surrey.ac.uk/Teaching/Unix/>

<http://people.ischool.berkeley.edu/~kevin/unix-tutorial/toc.html>

<http://www2.ocean.washington.edu/unix.tutorial.html>

<http://www.youtube.com/watch?v=v4YpWACs6Ts>

### 4.2 Shell scripting exercise

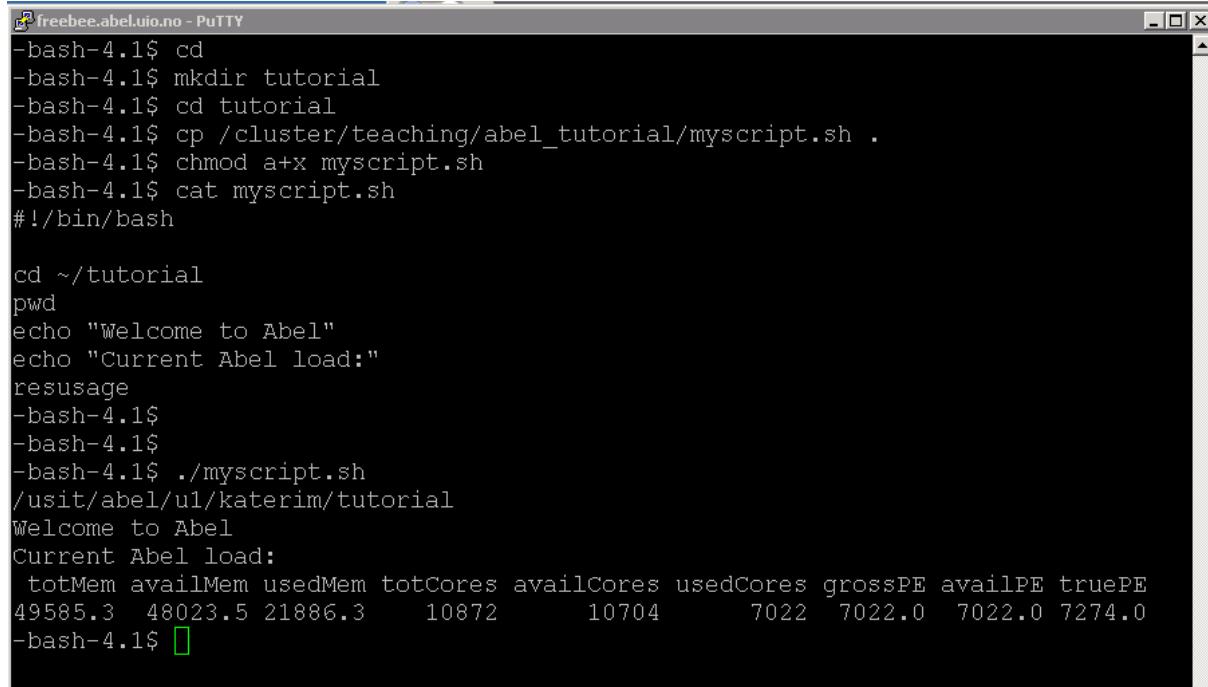
In order to run jobs on Abel, you will have to write job scripts. Job scripts are essentially shell scripts and this section contains a brief overview.

Shell script is a series of Unix commands written in a text file. Each command is on a separate line. Upon execution, all commands are executed sequentially. Shell scripts can be used for programming as they support flow control statements and variables.

Make sure that you are in your home directory. You can type “`cd ~`” (or “`cd`” for short). Make a directory named “tutorial” (“`mkdir tutorial`”) to keep all tutorial materials together. Change working directory to “tutorial” (“`cd tutorial`”).

Copy an example script called `myscript.sh` from the common area to your tutorial directory (“`cp /cluster/teaching/abelTutorial/myscript.sh .`”). Change permissions so the script can be executed; the command “`chmod a+x myscript.txt`” makes the file executable for everybody.

Examine the script. Note the first line of the script “#!/bin/bash” (#! is called a “hashbang”) tells the program loader to use an interpreter /bin/bash. Run the script by typing “**./myscript.sh**” and examine the output of the script. The script produced a current directory path (output of “pwd”), welcoming message and statistics of Abel load (output of “resusage”).



```
-bash-4.1$ cd
-bash-4.1$ mkdir tutorial
-bash-4.1$ cd tutorial
-bash-4.1$ cp /cluster/teaching/abel_tutorial/myscript.sh .
-bash-4.1$ chmod a+x myscript.sh
-bash-4.1$ cat myscript.sh
#!/bin/bash

cd ~/tutorial
pwd
echo "Welcome to Abel"
echo "Current Abel load:"
resusage
-bash-4.1$
-bash-4.1$
-bash-4.1$ ./myscript.sh
/usit/abel/ui/katerim/tutorial
Welcome to Abel
Current Abel load:
 totMem availMem usedMem totCores availCores usedCores grossPE availPE truePE
49585.3 48023.5 21886.3    10872      10704      7022  7022.0 7022.0 7274.0
-bash-4.1$
```

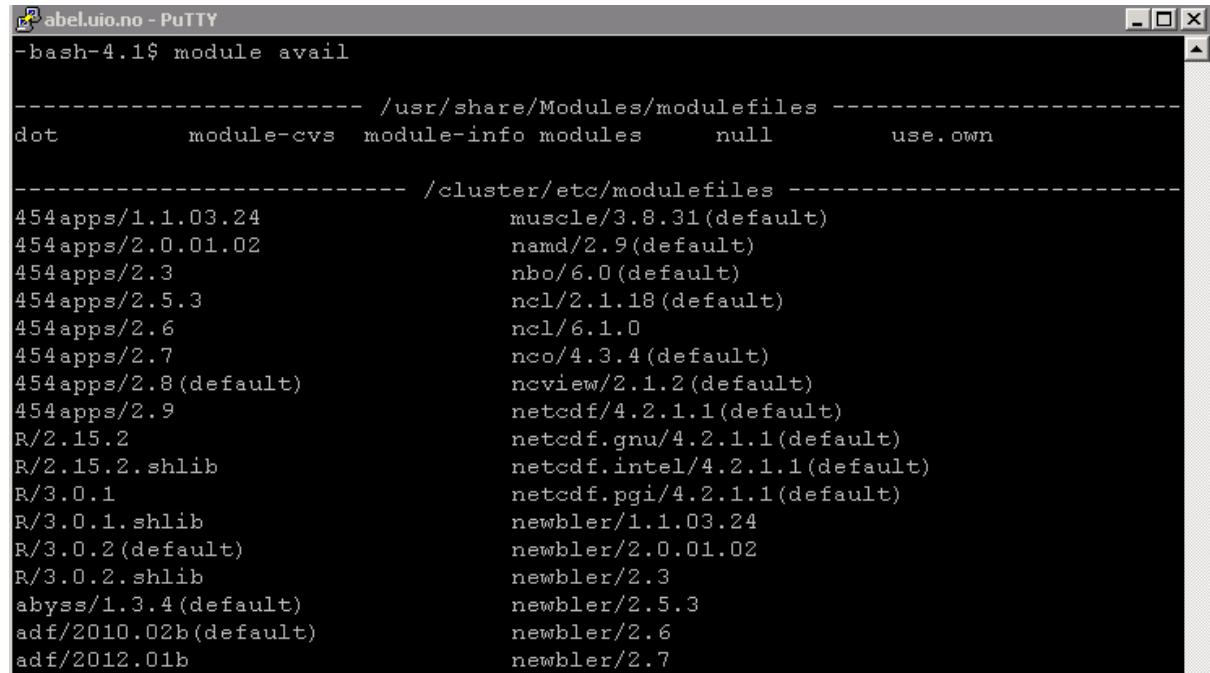
Read more about shell scripting -

[http://www.uio.no/english/services/it/research/hpc/courses/shell-scripting/20140325\\_intro\\_shell\\_scripting.pdf](http://www.uio.no/english/services/it/research/hpc/courses/shell-scripting/20140325_intro_shell_scripting.pdf)

## 5. Prepare and run your first job

### 5.1 Software and modules

Type “`module avail`” to see all software available on Abel. If you find a program that you want to use, type “`module load module_name`”. This sets up the \$PATH variable (and more) so the software is accessible to you. If you do not find the software that you are looking for, you can always place the executable in your home directory and use it from there.



```
abel.uio.no - PuTTY
-bash-4.1$ module avail

----- /usr/share/Modules/modulefiles -----
dot      module-cvs  module-info modules    null      use.own

----- /cluster/etc/modulefiles -----
454apps/1.1.03.24          muscle/3.8.31(default)
454apps/2.0.01.02          namd/2.9(default)
454apps/2.3                nbo/6.0(default)
454apps/2.5.3              ncl/2.1.18(default)
454apps/2.6                ncl/6.1.0
454apps/2.7                nco/4.3.4(default)
454apps/2.8(default)       ncview/2.1.2(default)
454apps/2.9                netcdf/4.2.1.1(default)
R/2.15.2                  netcdf.gnu/4.2.1.1(default)
R/2.15.2.shlib             netcdf.intel/4.2.1.1(default)
R/3.0.1                   netcdf.pgi/4.2.1.1(default)
R/3.0.1.shlib              newbler/1.1.03.24
R/3.0.2(default)           newbler/2.0.01.02
R/3.0.2.shlib              newbler/2.3
abyss/1.3.4(default)       newbler/2.5.3
adf/2010.02b(default)      newbler/2.6
adf/2012.01b               newbler/2.7
```

### 5.2 Job script

To submit a job (i.e. run a software), users have to communicate with the job manager on Abel. Job manager is a software that schedules jobs and oversees their execution on cluster compute nodes. On Abel, we use Simple Linux Utility for Resource Management – SLURM <https://computing.llnl.gov/linux/slurm/>.

Users communicate with the job manager using job scripts. Job scripts are shell scripts that contain the command that user wants to execute plus several job parameters. Job parameters inform the job manager on Abel about resources needed by the job. Without this information, the job cannot be scheduled.

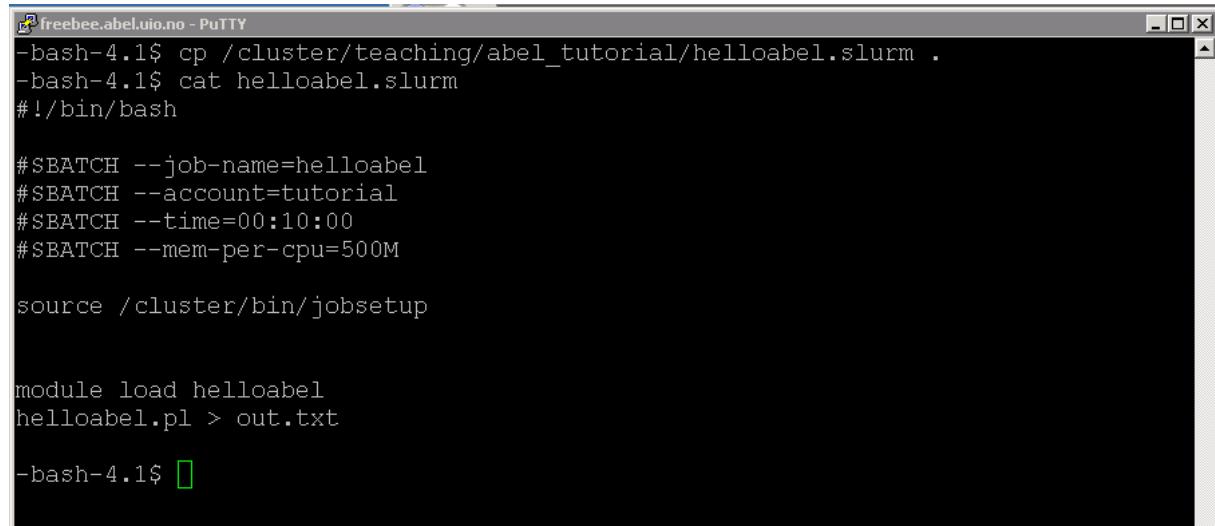
Abel documentation <http://www.uio.no/english/services/it/research/hpc/abel/help/user-guide/> contains a detailed description of job parameters and examples of job scripts.

Note of caution: Abel users are not allowed to use command line on login nodes to run jobs, this might render the login nodes irresponsible to other user login requests.

If you need interactive login, use qlogin command

(<http://www.uio.no/english/services/it/research/hpc/abel/help/user-guide/interactive-logins.html>). Qlogin reserves time on one of the cluster nodes, once the allocation is received, users can work directly on the command line of the node.

Copy a simple job script to your tutorial directory (“**cp /cluster/teaching/abel\_tutorial/helloabel.slurm .**”) and look at the contents. This is an example of a minimal job script, i.e. script that contains minimum information for the SLURM controller to schedule the job.



```
reebee.abel.uio.no - PuTTY
-bash-4.1$ cp /cluster/teaching/abel_tutorial/helloabel.slurm .
-bash-4.1$ cat helloabel.slurm
#!/bin/bash

#SBATCH --job-name=helloabel
#SBATCH --account=tutorial
#SBATCH --time=00:10:00
#SBATCH --mem-per-cpu=500M

source /cluster/bin/jobsetup

module load helloabel
helloabel.pl > out.txt

-bash-4.1$
```

- The script contains **#SBATCH** keywords that communicate to the job manager various settings for the job. The minimal set of these involve an account, time and memory.
  - The parameter “**--account**” is an administrative group of Abel users that a user has a right to use. Each account has a number of resources (cpus and cpu hours) assigned to it. Type “projects” on the command line to see which account(s) you have access to.
  - The parameter “**--time**” is the real time you expect your program to run.
  - The parameter “**--mem-per-cpu**” is the RAM requirement for your job. The job manager has to know this so it can schedule jobs in the most efficient manner. If you exceed time or memory specified, your job is cancelled.
  - The parameter “**--job-name**” is not compulsory but it is useful if you want to follow your job easily. We recommend to use it.
- The “**source /cluster/bin/jobsetup**” command sets up environment variables for your job. This line must always be present.
- The “**module load helloabel**” command gives you access to the helloabel executable.
- The “**./helloabel.pl > out.txt**” command is the core of the script, it is the program that is executed on Abel. In this case, the output is captured into the file called out.txt.

### **5.3 Submit your job and follow the progress**

To execute the helloabel.pl script on Abel, submit the job using “`sbatch helloabel.slurm`”. The command submits the job into the job queue and returns a job identifier (jobid). Jobs can be followed as they are waiting in the queue and as they are executed. Type “`squeue`” and the whole job queue is displayed. To limit the output to your jobs only, type “`squeue -u user_name`”. Typing “`scontrol show job jobid`” produces detailed output including job parameters and the job script path. If in need to cancel a job, type “`scancel jobid`”.

```
abel.uio.no - PuTTY
-bash-4.1$ sbatch helloabel.slurm
Submitted batch job 1760453
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER      ST          TIME  NODES NODELIST(REASON)
1760453    normal  helloabe  katerim    PD          0:00      1  (Priority)
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER      ST          TIME  NODES NODELIST(REASON)
1760453    normal  helloabe  katerim     R          0:12      1 c11-28
-bash-4.1$ scontrol show job 1760453
JobId=1760453 Name=helloabel
  UserId=katerim(108773) GroupId=users(100)
  Priority=30009 Account=tutorial QOS=noturhigh
  JobState=RUNNING Reason=None Dependency=(null)
  Requeue=1 Restarts=0 BatchFlag=1 ExitCode=0:0
  RunTime=00:01:47 TimeLimit=00:10:00 TimeMin=N/A
  SubmitTime=2013-04-12T15:18:17 EligibleTime=2013-04-12T15:18:17
  StartTime=2013-04-12T15:18:47 EndTime=2013-04-12T15:28:47
  PreemptTime=None SuspendTime=None SecsPreSuspend=0
  Partition=normal AllocNode:Sid=login-0-2:26061
  ReqNodeList=(null) ExcNodeList=(null)
  NodeList=c11-28
  BatchHost=c11-28
  NumNodes=1 NumCPUs=1 CPUs/Task=1 ReqS:C:T=*:*:*
  MinCPUsNode=1 MinMemoryCPU=500M MinTmpDiskNode=0
  Features=(null) Gres=(null) Reservation=(null)
  Shared=OK Contiguous=0 Licenses=(null) Network=(null)
  Command=/cluster/home/katerim/tutorial/helloabel.slurm
  WorkDir=/cluster/home/katerim/tutorial

-bash-4.1$
```

- In the example above, the job id is 1760453.
- There are two instances of the `squeue` command output. In the first case, the status (ST) of the job is PD or pending as the job is waiting for available resource. In the second case, the job has been running on the node c11-28 for the last 12 sec.
- The example above also shows the output of the `scontrol` command that contains detailed information about the job. Note the last two rows that show the path to the working directory and to the job script. The `RunTime` and `TimeLimit` items (in the 5th row) tell you how much total time you have and how long the job has been running. If you are running out, let us know, we can extend the time.

## **5.4 Job output**

List the contents of your tutorial directory. The file **out.txt** contains the output of our program. The file **slurm-jobid.out** contains messages from the queuing system. It also contains time, place and other parameters of cluster execution. If there is a problem, error messages are logged in this file.

**Exercise:** Learn how to run a script that is in your tutorial directory.

- Copy a perl script called hellocal.pl from /cluster/teaching/abel\_tutorial to your tutorial directory.
- Prepare a job script to run hellocal.pl. First, copy ~/tutorial/helloabel.slurm into ~/tutorial/hellocal.slurm. Second, make changes to hellocal.slurm so it executes “hellocal.pl” from your directory (hint use ./ to tell the script that the program is in the current directory).
- Finally, submit the job and check the output files.

**Note:** Executing jobs in you home area does not require using module load command, instead you have to provide path to the executable.

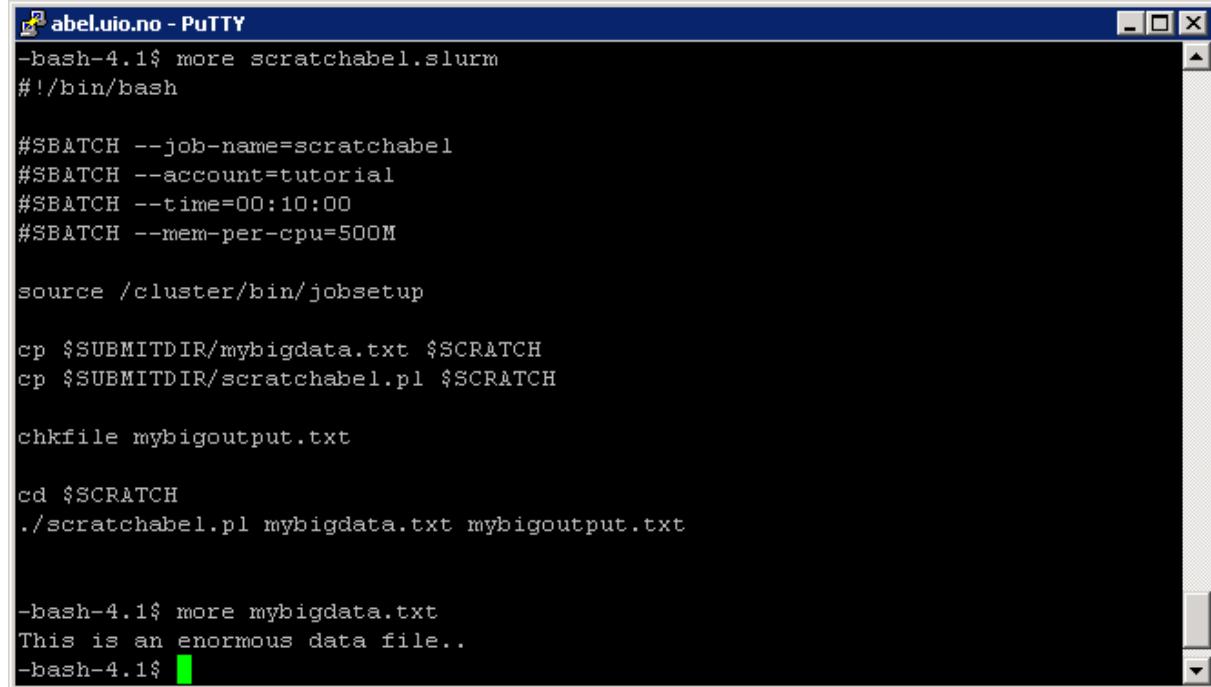
## 6. Prepare a job using scratch area

If your job reads and writes lots of data (and especially if this happens often and in small chunks), we recommend to stage the job execution onto a faster file system that is available on Abel. This partition is more suitable for I/O demanding jobs than your home (and slower) directory. On Abel, this partition is called /work and every job receives a temporal (scratch) space on this partition (specifically in /work/jobs/jobid.d subdirectory). This directory is removed when the job is finished. It cannot be used for data storage.

The path to the scratch directory is contained in the SCRATCH environment variable that is instantiated when you submit your job. Every job creates several environment variables, e.g. SUBMITDIR that contain path to the directory containing the job script.

### 6.1 Job script

Copy the script “scratchabel.slurm” from the common area. In addition copy the files scratchabel.pl and mybigdata.txt. s



A screenshot of a PuTTY terminal window titled "abel.uio.no - PuTTY". The window displays a command-line session. The user has run the command "more scratchabel.slurm" which shows a shell script named "scratchabel.slurm". The script contains several "#SBATCH" directives specifying job parameters like name, account, time, and memory, followed by commands to source a setup file, copy input files to the scratch directory, and run the executable. After executing the script, the user runs "more mybigdata.txt" which outputs the message "This is an enormous data file..".

```
-bash-4.1$ more scratchabel.slurm
#!/bin/bash

#SBATCH --job-name=scratchabel
#SBATCH --account=tutorial
#SBATCH --time=00:10:00
#SBATCH --mem-per-cpu=500M

source /cluster/bin/jobsetup

cp $SUBMITDIR/mybigdata.txt $SCRATCH
cp $SUBMITDIR/scratchabel.pl $SCRATCH

chkfile mybigoutput.txt

cd $SCRATCH
./scratchabel.pl mybigdata.txt mybigoutput.txt

-bash-4.1$ more mybigdata.txt
This is an enormous data file..
-bash-4.1$
```

- The script takes advantage of environment variables SUBMITDIR (directory that you are submitting from) and SCRATCH (temporal directory assigned to your job). They are accessible to your job script and your executable during job execution.
- To use SCRATCH, you have to copy your executable and your input data files (mybigdata.txt and scratchabel.pl) there. This is accomplished using “cp \$SUBMITDIR/mybigdata.txt \$SCRATCH” and “cp \$SUBMITDIR/scratchabel.pl \$SCRATCH” lines.
- Register output file(s) that you want copied back from the scratch area using the chkfile command. This is the recommended way of doing so since this command is

always executed at the end of your job regardless if the job finished successfully or not. Alternatively, you could simply copy the result file at the end of your script (e.g. “cp \$SCRATCH/mybigoutput.txt \$SUBMITDIR”). This is fine if the job runs successfully. In case of a failed run, the last copy statement is not executed as opposed to chkfile that always runs.

- Before the job is executed, you have to change working directory to scratch area using “cd \$SCRATCH”.

Submit the script using “**sbatch scratchabel.slurm**”, note job id and check the /work/jobid.d subdir where you can see your script and data files. When done, check the output:

```
abel.uio.no - PuTTY
-bash-4.1$ sbatch scratchabel.slurm
Submitted batch job 1777530
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER     ST      TIME  NODES NODELIST(REASON)
1777530        normal scratcha  katerim    R      0:16      1 c3-4
-bash-4.1$ ls -l /work/1777530.d/
total 2
-rw-r--r-- 1 katerim users  32 Apr 14 17:19 mybigdata.txt
-rw-r--r-- 1 katerim users   73 Apr 14 17:19 mybigoutput.txt
-rwxr-xr-x 1 katerim users 436 Apr 14 17:19 scratchabel.pl
-bash-4.1$ more mybigoutput.txt
This is the copy of your big data file.

This is an enormous data file..

Executed on Sun Apr 14 17:19:04 2013 on compute-3-4.local in /work/1777530.d.
-bash-4.1$
```

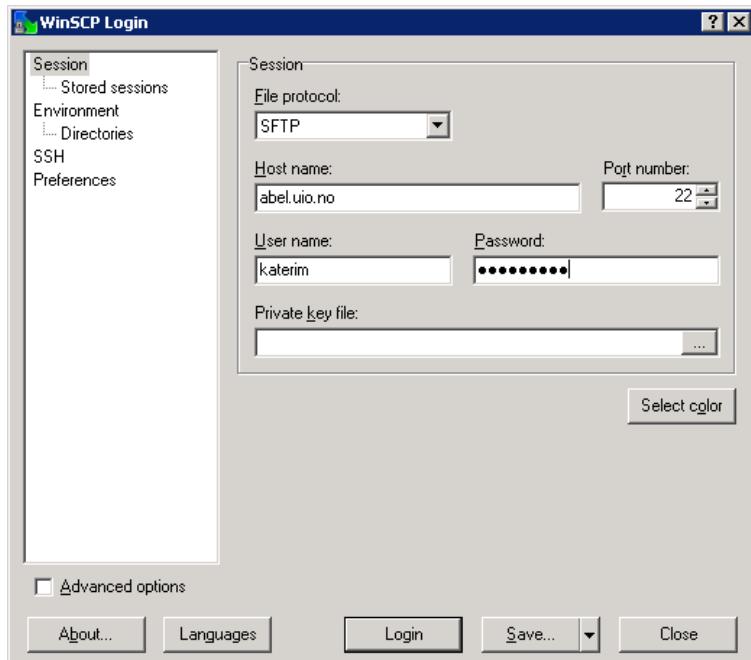
The script scratchabel.pl is only a demonstration of a principle, it simply echoes back your input data and adds information about time, node and current directory.

## **6.2 Upload you own file**

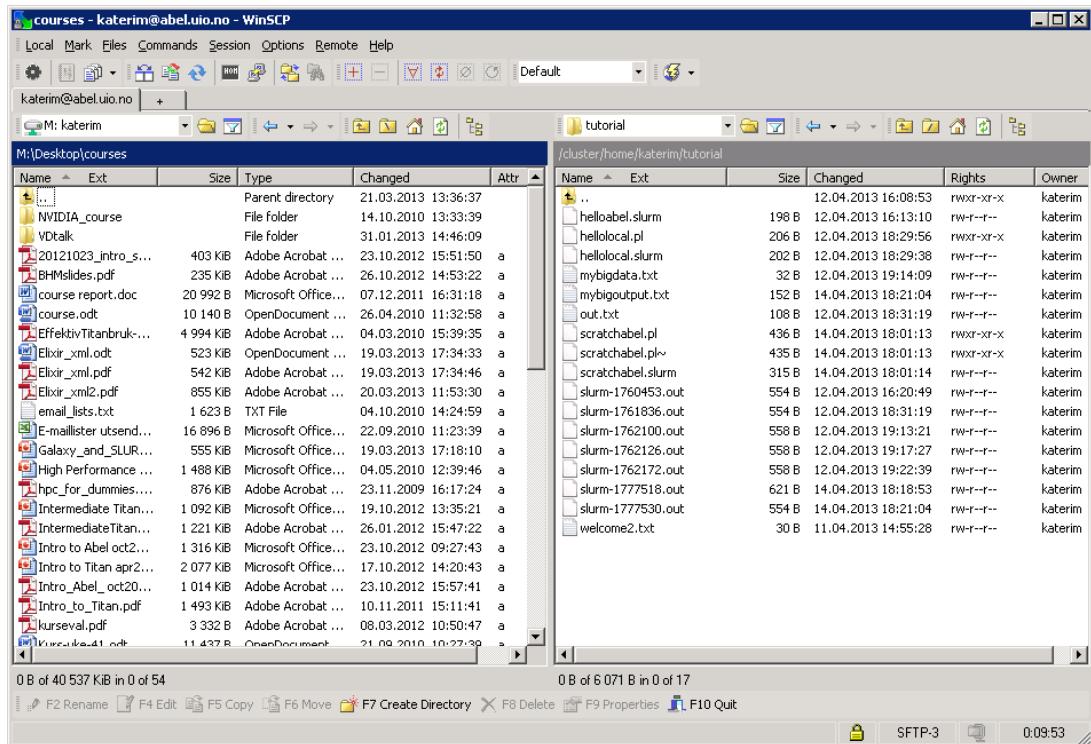
This section explains how to upload data from your local computer to Abel.

### ***6.2.1 On Windows***

Download WinSCP from <http://winscp.net/eng/download.php>. Install and run. Type abel.uio.no into the “Host name” dialog box. Below, specify user name and password:



Press “Login” and manipulate your files using “drag and drop”:



### ***6.2.2 On Linux***

Open a terminal and use scp (secure copy) command. Type “scp myfile.txt your\_user\_name@abel.uio.no:~/tutorial” to copy a file from a current directory to abel tutorial directory. For large file the rsync command is preffered since, in case of interruption, it is able to resume where it left off; “rsync –z myfile.txt your\_user\_name@abel.uio.no:~/tutorial” uploads your file to the tutorial directory on Abel.

### ***6.2.3 On the Mac***

Open a terminal and use scp the same way as on Linux.

## 7. Parallel jobs

The strength of cluster computing lies in parallel processing i.e. executing many instances of the same executable at the same time. This is often needed when you have many input datasets or when you run simulations with different input parameters. It might often be possible to split a large input file into chunks and parallelize your job.

Abel offers a utility called “arrayrun” that can start a number of parallel jobs using the same job script.

Since the parallel execution is driven by the same job script from the same directory, the main consideration during setup is to manage input and output files in such a way that the correct files are read and the output does not get overwritten. The recommended way is to use the TASK\_ID environment variable in file names. This variable has unique value for each run; if you run your program 10 times using arrayrun command, the TASK\_ID ranges from 1 to 10.

### **7.1 Arrayrun**

In your tutorial directory, make a subdirectory called arrayruntest and cd into it (“`cd ~/tutorial; mkdir arrayruntest; cd arrayruntest`”). Copy the file arrayruntest.tar from /cluster/teaching/abel\_tutorial to your arrayruntest directory. Unpack all files by “`tar xvf arrayrun.tar`”.

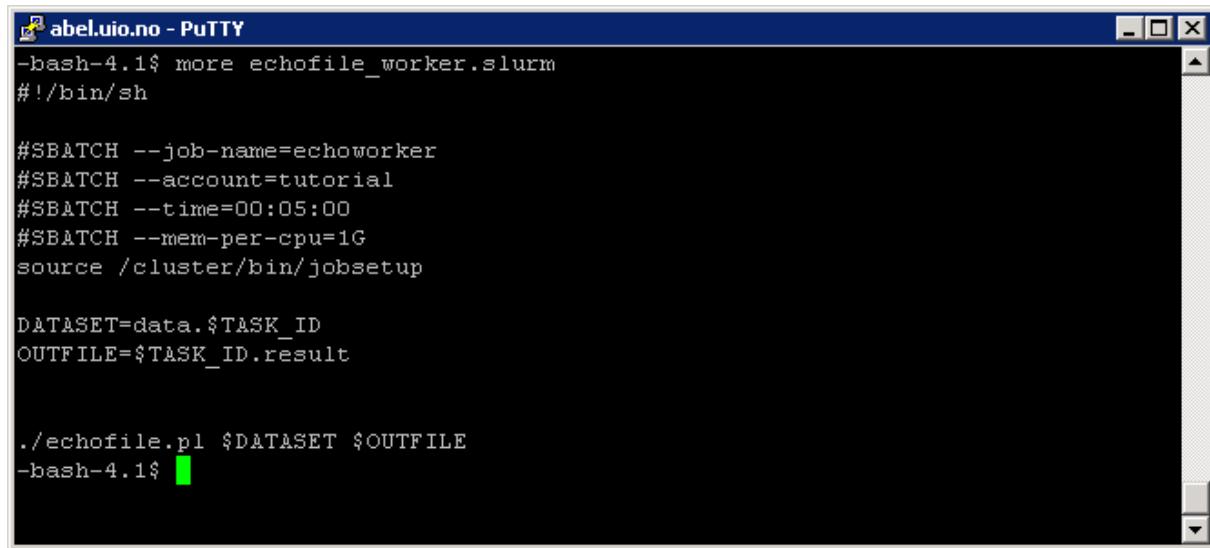
The archive (arrayruntest.tar) unpacks into datafiles, two job scripts and one executable. The perl executable, echofile.pl, is a very simple placeholder for your own program, it copies an input file to output and adds some info about time and place. There are ten data files called data.\* and two job scripts – echofile\_submit.slurm and echofile\_worker.slurm.

```
abel.uio.no - PuTTY
-bash-4.1$ tar xvf arrayruntest.tar
data.1
data.10
data.2
data.3
data.4
data.5
data.6
data.7
data.8
data.9
echofile.pl
echofile_submit.slurm
echofile_worker.slurm
-bash-4.1$ ls -l
total 1366
-rw-r--r-- 1 katerim users 706560 Apr 15 12:47 arrayruntest.tar
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.1
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.10
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.2
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.3
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.4
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.5
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.6
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.7
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.8
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.9
-rwxr-xr-x 1 katerim users 401 Apr 15 12:26 echofile.pl
-rw-r--r-- 1 katerim users 185 Apr 14 19:26 echofile_submit.slurm
-rw-r--r-- 1 katerim users 227 Apr 15 11:56 echofile_worker.slurm
-bash-4.1$
```

### **7.3. Arrayrun job scripts**

Every arrayrun needs one “worker” script and one “submit” script.

The worker script calls your program and it is similar to the scripts we have already used in this tutorial. The modifications consist of the file management using the TASK\_ID variable. Two extra variables DATASET for input and OUTFILE for output were added. Their names include the value of TASK\_ID. In this example in run 1, the input is called “data.1” and output “result.1” and in the last, 10th run “data.10” and “result.10” respectively.



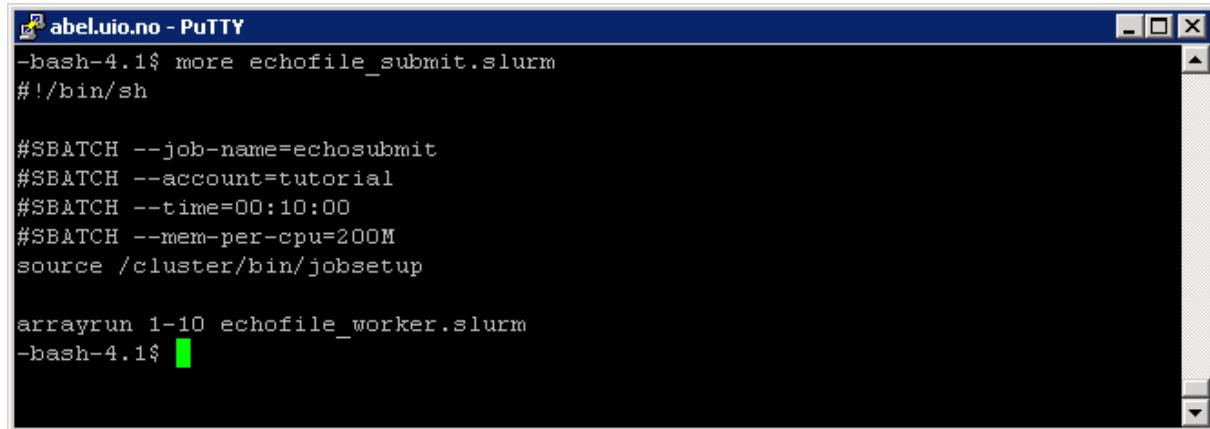
```
abel.uio.no - PuTTY
-bash-4.1$ more echofile_worker.slurm
#!/bin/sh

#SBATCH --job-name=echoworker
#SBATCH --account=tutorial
#SBATCH --time=00:05:00
#SBATCH --mem-per-cpu=1G
source /cluster/bin/jobsetup

DATASET=data.$TASK_ID
OUTFILE=$TASK_ID.result

./echofile.pl $DATASET $OUTFILE
-bash-4.1$
```

The submit script executes the arrayrun command. The arrayrun will submit all the jobs and then check on their status. It will be finished when all jobs are done. The submit script contains the same SBATCH settings as an ordinary job script. Note that the memory consumption for this job is low but the time setting is high. The time for the submit job should always be much longer than the time for individual jobs because the jobs are not started all at once but are released in batches. The command arrayrun needs parameters specifying the number of runs and the name of the worker script:



```
abel.uio.no - PuTTY
-bash-4.1$ more echofile_submit.slurm
#!/bin/sh

#SBATCH --job-name=echosubmit
#SBATCH --account=tutorial
#SBATCH --time=00:10:00
#SBATCH --mem-per-cpu=200M
source /cluster/bin/jobsetup

arrayrun 1-10 echofile_worker.slurm
-bash-4.1$
```

Submit the arrayrun by “`sbatch echofile_submit.slurm`” and then type “`watch squeue -u your_user_name`”. You can see the individual jobs being released by the submit script.

Since arrayrun can release many jobs, it is necessary to know how to cancel them if needed. To cancel arrayrun and all worker processes, use “`scancel jobid`”, jobid being the id of the submit script.

In the following figure, you can see output of the queue command. In the first instance of the submit script is waiting to run. In the second instance, the worker scripts have been released and are waiting to run. Finally, in the third instance, most of the worker scripts started to run:

```
abel.uio.no - PuTTY
-bash-4.1$ sbatch echofile_submit.slurm
Submitted batch job 1788940
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER      ST          TIME  NODES NODELIST(REASON)
1788940      normal  echosubm  katerim    PD      0:00      1  (None)
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER      ST          TIME  NODES NODELIST(REASON)
1788942      normal 1788940.  katerim    PD      0:00      1  (None)
1788943      normal 1788940.  katerim    PD      0:00      1  (None)
1788944      normal 1788940.  katerim    PD      0:00      1  (None)
1788945      normal 1788940.  katerim    PD      0:00      1  (None)
1788946      normal 1788940.  katerim    PD      0:00      1  (None)
1788947      normal 1788940.  katerim    PD      0:00      1  (None)
1788948      normal 1788940.  katerim    PD      0:00      1  (None)
1788949      normal 1788940.  katerim    PD      0:00      1  (None)
1788951      normal 1788940.  katerim    PD      0:00      1  (None)
1788952      normal 1788940.  katerim    PD      0:00      1  (None)
1788940      normal echosubm  katerim     R      0:22      1 c12-15
-bash-4.1$ squeue -u katerim
  JOBID PARTITION      NAME      USER      ST          TIME  NODES NODELIST(REASON)
1788949      normal 1788940.  katerim    PD      0:00      1  (Priority)
1788951      normal 1788940.  katerim    PD      0:00      1  (Priority)
1788952      normal 1788940.  katerim    PD      0:00      1  (Priority)
1788942      normal 1788940.  katerim     R      0:08      1 c12-15
1788943      normal 1788940.  katerim     R      0:08      1 c12-15
1788944      normal 1788940.  katerim     R      0:08      1 c12-15
1788945      normal 1788940.  katerim     R      0:08      1 c12-15
1788946      normal 1788940.  katerim     R      0:08      1 c12-15
1788947      normal 1788940.  katerim     R      0:08      1 c12-15
1788948      normal 1788940.  katerim     R      0:08      1 c12-15
1788940      normal echosubm  katerim     R      0:38      1 c12-15
-bash-4.1$
```

### 7.3 Arrayrun results

When you see no more jobs in the queue you can examine the results. Your arrayruntest directory should look like this:

```
achel.uio.no - PuTTY
-bash-4.1$ ls -l
total 2057
-rw-r--r-- 1 katerim users 68515 Apr 15 15:14 10.result
-rw-r--r-- 1 katerim users 69233 Apr 15 15:14 1.result
-rw-r--r-- 1 katerim users 68877 Apr 15 15:14 2.result
-rw-r--r-- 1 katerim users 68513 Apr 15 15:14 3.result
-rw-r--r-- 1 katerim users 69233 Apr 15 15:14 4.result
-rw-r--r-- 1 katerim users 68876 Apr 15 15:14 5.result
-rw-r--r-- 1 katerim users 68513 Apr 15 15:14 6.result
-rw-r--r-- 1 katerim users 69233 Apr 15 15:14 7.result
-rw-r--r-- 1 katerim users 68876 Apr 15 15:14 8.result
-rw-r--r-- 1 katerim users 68513 Apr 15 15:15 9.result
-rw-r--r-- 1 katerim users 706560 Apr 15 12:47 arrayruntest.tar
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.1
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.10
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.2
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.3
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.4
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.5
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.6
-rw-r--r-- 1 katerim users 69147 Apr 15 11:53 data.7
-rw-r--r-- 1 katerim users 68790 Apr 15 11:53 data.8
-rw-r--r-- 1 katerim users 68427 Apr 15 11:53 data.9
-rwxr-xr-x 1 katerim users 401 Apr 15 12:26 echofile.pl
-rw-r--r-- 1 katerim users 185 Apr 14 19:26 echofile_submit.slurm
-rw-r--r-- 1 katerim users 227 Apr 15 15:11 echofile_worker.slurm
-rw-r--r-- 1 katerim users 5975 Apr 15 15:16 slurm-1789315.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789334.out
-rw-r--r-- 1 katerim users 554 Apr 15 15:14 slurm-1789336.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789338.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789339.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789340.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789341.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789342.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:14 slurm-1789343.out
-rw-r--r-- 1 katerim users 552 Apr 15 15:15 slurm-1789344.out
-rw-r--r-- 1 katerim users 554 Apr 15 15:14 slurm-1789345.out
-bash-4.1$
```

You can see the result files (1-10.result) and also slurm log files. There is one log file for each run as well as for the submit script (11 all together). The biggest slurm log is always the one produced by the submit script.

Finally, we examine the content of the result files. They contain log of time and place of execution. Type “**grep Job \*.result**” to see this information from each file at once.

```
achel.uio.no - PuTTY
-bash-4.1$ grep Job *.result
10.result:Job 10: Mon Apr 15 15:13:51 2013 on compute-11-3.local
1.result:Job 1: Mon Apr 15 15:13:19 2013 on compute-9-30.local
2.result:Job 2: Mon Apr 15 15:13:19 2013 on compute-11-23.local
3.result:Job 3: Mon Apr 15 15:13:43 2013 on compute-11-3.local
4.result:Job 4: Mon Apr 15 15:13:20 2013 on compute-11-3.local
5.result:Job 5: Mon Apr 15 15:13:20 2013 on compute-11-3.local
6.result:Job 6: Mon Apr 15 15:13:19 2013 on compute-11-3.local
7.result:Job 7: Mon Apr 15 15:13:20 2013 on compute-11-3.local
8.result:Job 8: Mon Apr 15 15:13:51 2013 on compute-11-3.local
9.result:Job 9: Mon Apr 15 15:14:01 2013 on compute-11-3.local
-bash-4.1$
```

**Exercise: Try to run larger arrayrun:**

- edit the submit script to add more runs (pls do not go over 50, our queue is limited).
- get more input data files, use a tar file called bigdata.tar in /cluster/teaching/abel\_tutorial that contains more data
- change the input file name in your worker script
- increase the time limit in your submit script to one hour (--time=01:00:00)

Watch the queue, you see the jobs being released in batches. When your all is finished you can examine the time information in the output files. Can you see the batches of jobs being released by arrayrun?

Tip: To examine the results, use “grep Job \*.result | sort -n “. This will give you an overview about the arrayrun execution.