High Performance Computing Systems

Brandi Greer
Baylor University
Statistics Department

September 2009

BAYLOR
UNIVERSITY

What is Fozzie?

- Fozzie is one of the four high performance computing systems maintained by the Academic and Research Computer Services (ARCS) group at Baylor University.
- There are 32 computing nodes available on Fozzie used to run programs in batch mode.
- Processing capabilities on each node are comparable (a little faster) to computing times for the computers currently in the lab.

BAYLOR
UNIVERSITY

How will it help you?

- Continuous computing power without having to worry about computer reboots or class schedules
- Ability to queue as many jobs as you like that will be processed in a first come first serve order
- E-mail notification of job completion
- Access from anywhere
- Results compiled in one central location

BAYLOR UNIVERSITY

What's the catch?

- No WinBUGS capabilities
- Limited to using 3 nodes per individual (maybe)
- System is shared with entire university so there may be times of high traffic and limited availability

Getting Started - Mac

- Open Terminal

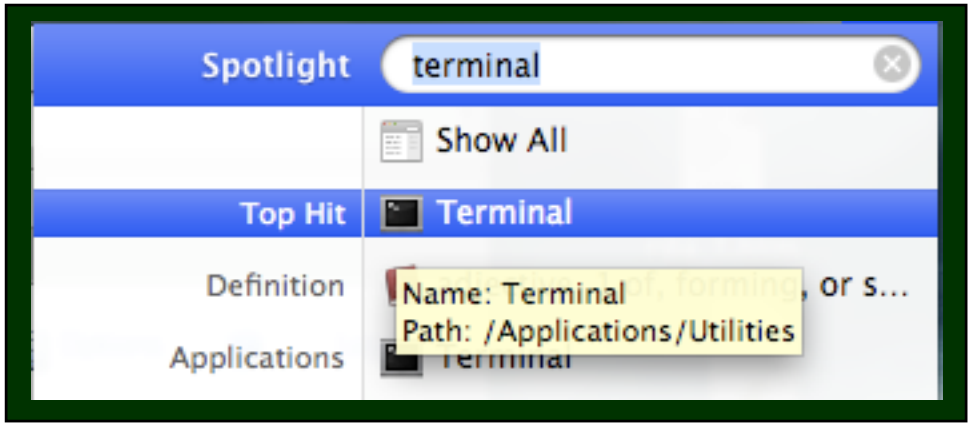

- In Terminal Type:

  ssh -l UserName fozzie.baylor.edu

  – where UserName is your Fozzie ID name and –l is a lower case L

BAYLOR UNIVERSITY

Getting Started - Mac

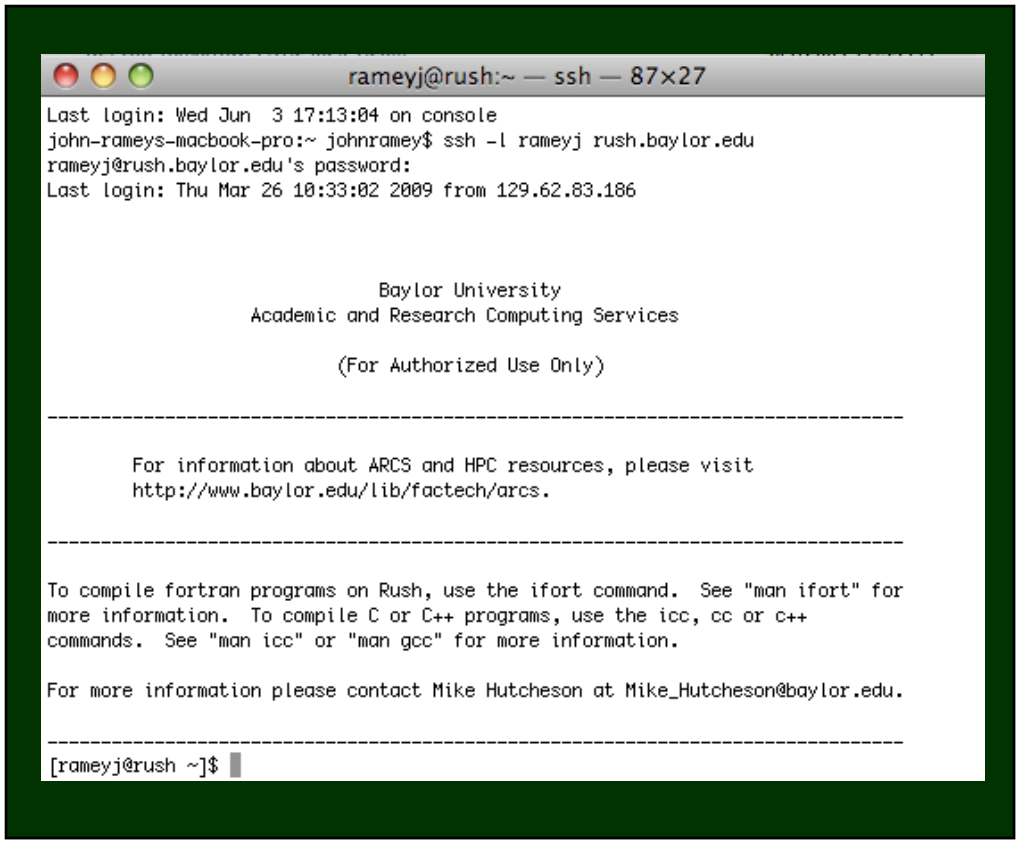

BAYLOR UNIVERSITY

Getting Started - Windows

- Download Putty and PSCP
  - http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html
  - www.freshmeat.com

Setting up Fozzie on Putty

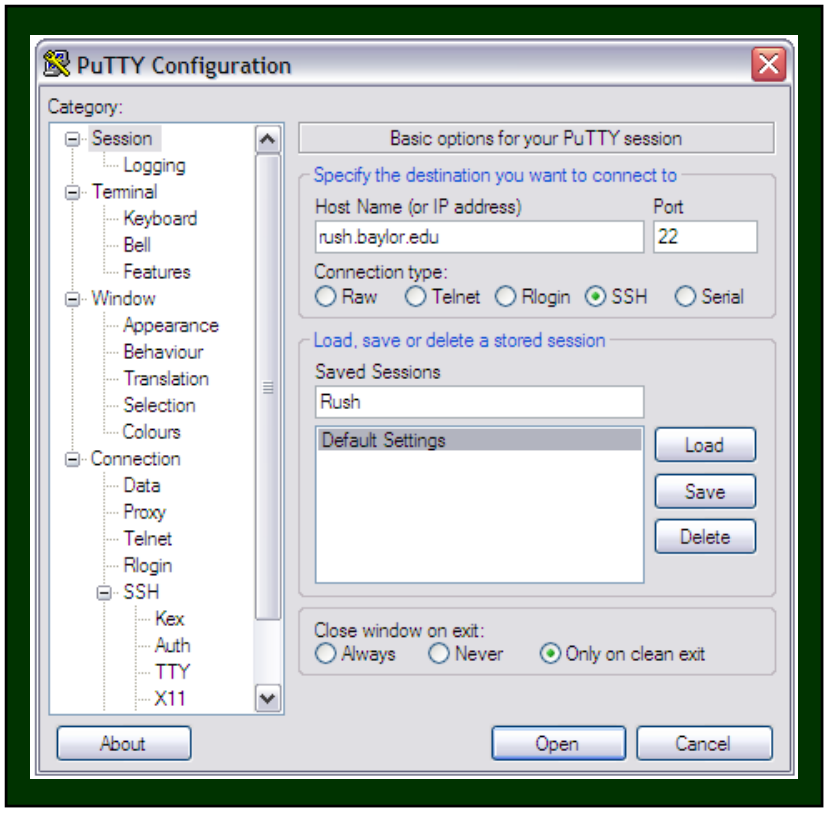

- You will need a username and password to log on.
- Open Putty
- Host Name: Fozzie.baylor.edu
- Connection type: SSH

BAYLOR UNIVERSITY

Setting up Fozzie on Putty

- Enable X11 forwarding
  - On left side of screen click on SSH, then X11, check “Enable X11 Forwarding”

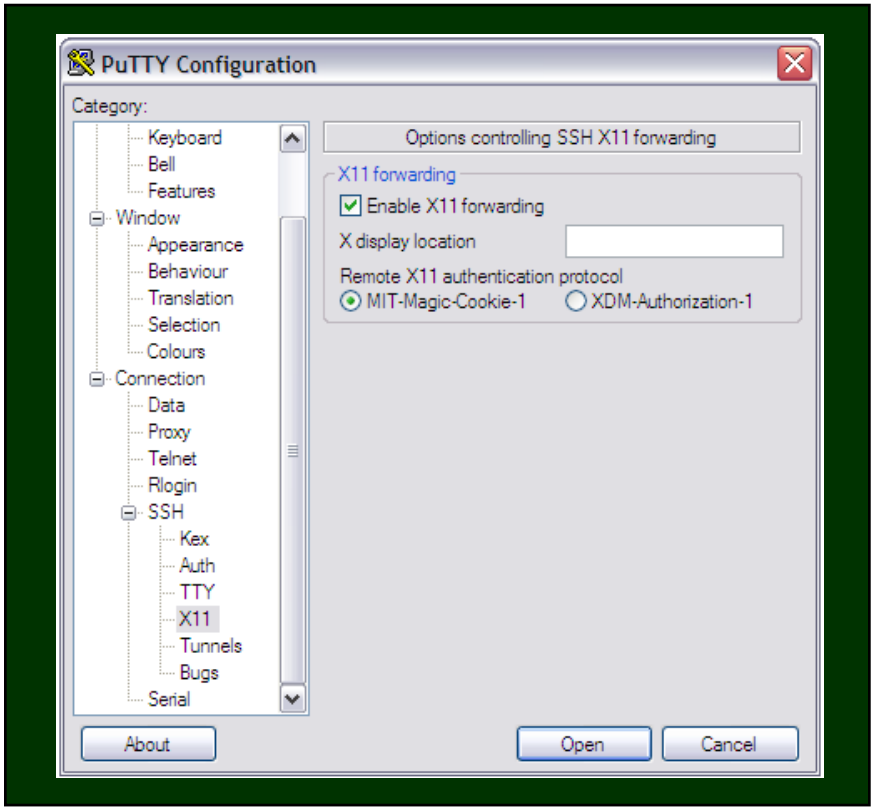

BAYLOR UNIVERSITY

R Programs

- Other than composing your program to run in batch mode, there is no special programming required for submitting jobs to R.
- Simply make sure you have things organized and set up as planned.
- Save files using either
  - write.table(width, file="width_test.txt", sep=" ",row.names=FALSE, col.names=FALSE)
    - This saves only the matrix “width” to “width_test.txt”
    - Use options to format
  - save(list = ls(all=TRUE), file = "widths.RData")
    - This saves all of the assigned data values in an Rdata file, which can be loaded as a workspace in R.

R Programs

- Shell files
  - Every file you submit will need a ".sh" shell file to run on Fozzie.
  - You can copy the text below into notepad and then save as a ".sh" file rather than a ".txt" file

```
#!/bin/bash
#PBS -S /bin/bash
#
# If this file is "run_r.sh", submit with "qsub run_r.sh"
#
echo "Starting R at `date`"
# Let example.r  be a file of R commands.
# Run R with this command file:
R --save <~/example/CI_for_mean_ex.txt> ~/example/int.out
echo "R run completed at `date`"
```

R Programs

- The shell code assumes you are in your home drive on Fozzie and want to execute the file “CI_for_mean_ex.txt” that is in the folder “example”.

- Then want the data to save to a file int.out in the same folder.
  - Note: In R, if you use the command “write.table” in your code, these files will be written to your home directory in Fozzie as they are completed. So the int.out file is irrelevant, but necessary.

Matlab Programs

- Be sure Matlab files end in “quit;”
- You can use the program to call other functions in separate files as long as these separate functions are in the same folder in your Fozzie account.

Matlab Programs

- Shell File

```
#!/bin/bash
#PBS -S /bin/bash
# If this file is "run_matlab.sh", submit with "qsub run_matlab.sh"
echo "Changing to directory from which PBS script was submitted."
cd $PBS_O_WORKDIR
echo "Current working directory is now: " `pwd`
echo "Starting MATLAB at `date`"
# Let myProgram.m be a file of Matlab commands.
# Be sure to include a "quit;" as the last line of the file.
# Run MATLAB with this command file:
/usr/local/bin/matlab -nodisplay -nodesktop -nojvm -nosplash < example.m >
    ex.out
echo "MATLAB run completed at `date`"
```

Matlab Programs

- The shell file assumes that you want to run the Matlab program “example.m” from your current Fozzie directory and want the results to be put in “ex.out” in the same directory.

Transferring Files

- Download & install FileZilla
  - ***http://filezilla-project.org/***

- Log on to your Fozzie account
  - Host: fozzie.baylor.edu
  - Username & Password : use your assigned userID and your created password
  - Port: 22 (for SFTP)

Transferring Files

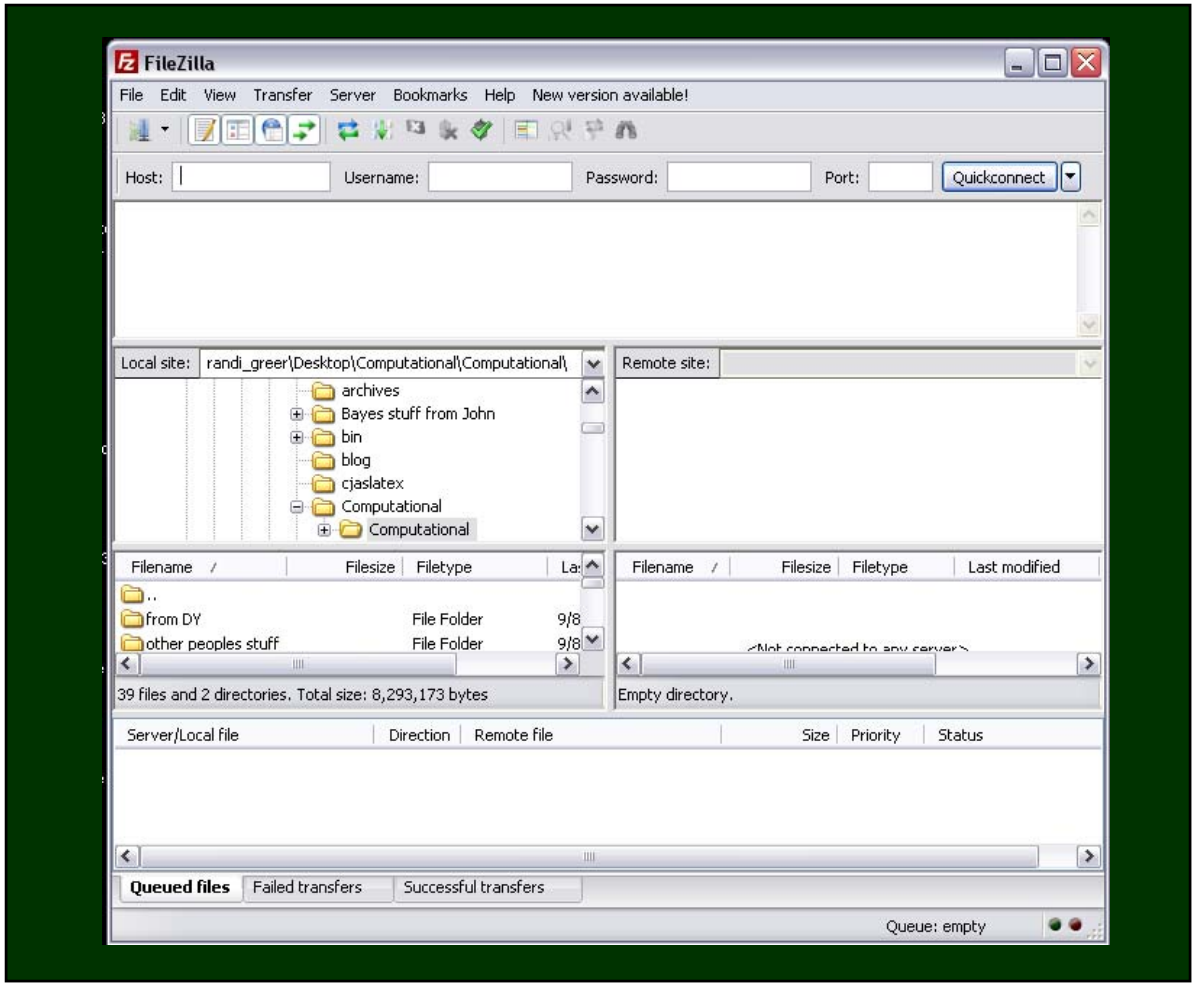

Transferring Files

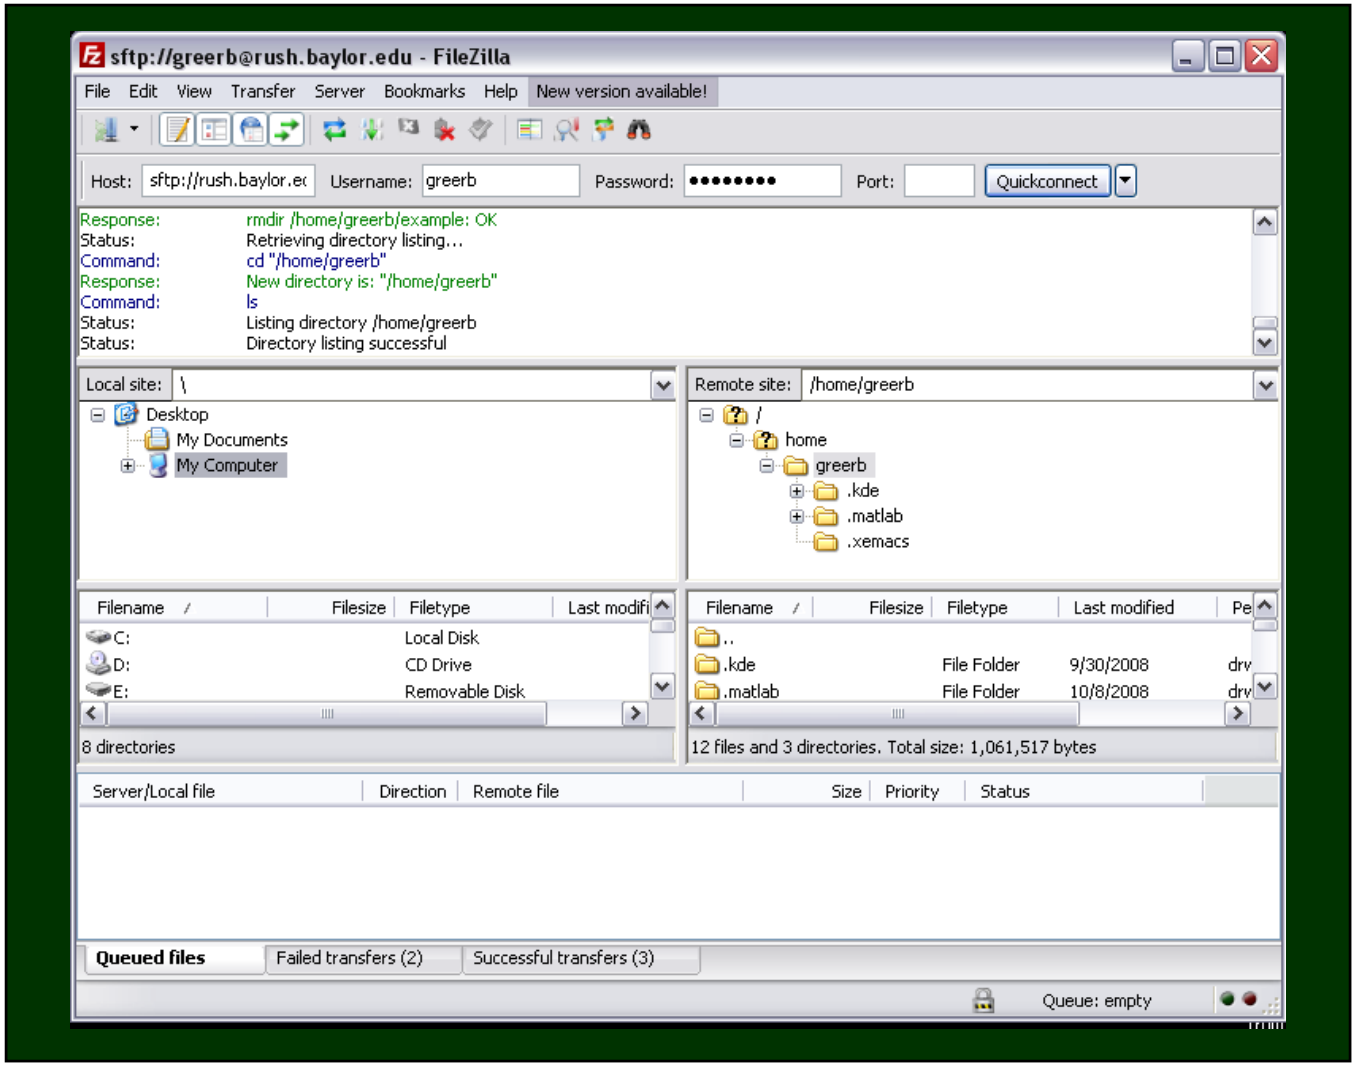

BAYLOR UNIVERSITY

Transferring Files

- Find the file or folder you wish to upload to Fozzie on the computer in the “Local Site” area of the FileZilla screen
- Drag and drop into the desired folder on your Fozzie account.

Submitting a Job to Fozzie

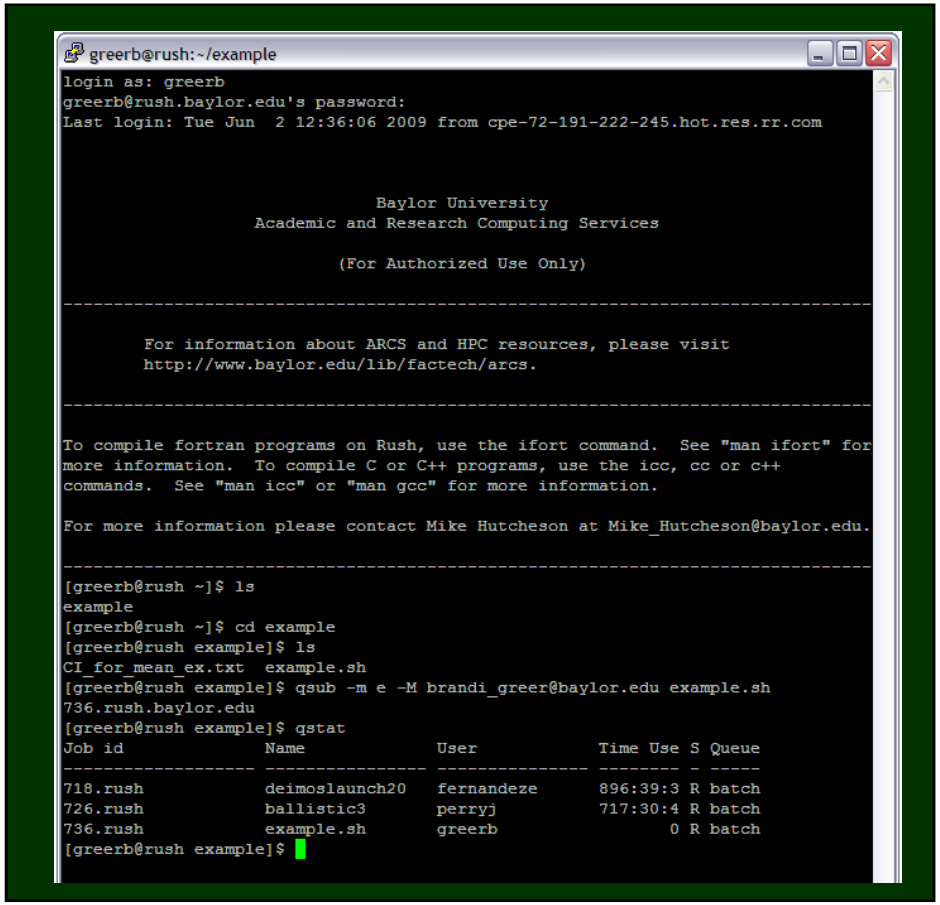

BAYLOR UNIVERSITY

Submitting a Job to Fozzie

- Use the “cd” command to get to the directory in Fozzie that contains both the program file and the shell file

- In Fozzie (Putty) type:

  qsub –m e –M your_email@baylor.edu example.sh

  - Qsub is the submit command
  - -m e –M your_email@baylor.edu tells Fozzie to email you when your job is complete
  - example.sh is the shell file that you want to execute

Retrieving Files from Fozzie

- Go back to Filezilla to drag and drop desired output files into your personal computer location.

Other Commands in Fozzie

- qstat – allows you to check the status of jobs on Fozzie
- qdel – used to delete a job in progress
  - Type “qdel 114” to delete job with job id #114
- cd – helps navigate through folders on your Fozzie account
  - “cd” alone will take you to your home directory
  - “cd folder1” will take you to folder1

Other Commands in Fozzie

- ls – lists all files in the current directory
- rm – remove a file from your directory
  - To remove an entire folder “rm –r folder1”
- For more file manipulation tips: http://linuxcommand.org/lts0050.php

Accessing Programs in Fozzie

- Aside from simply submitting simulations in batch mode on R and Matlab, you can also access these two programs via Fozzie.

Accessing R

- Simply type “R” into the Fozzie prompt

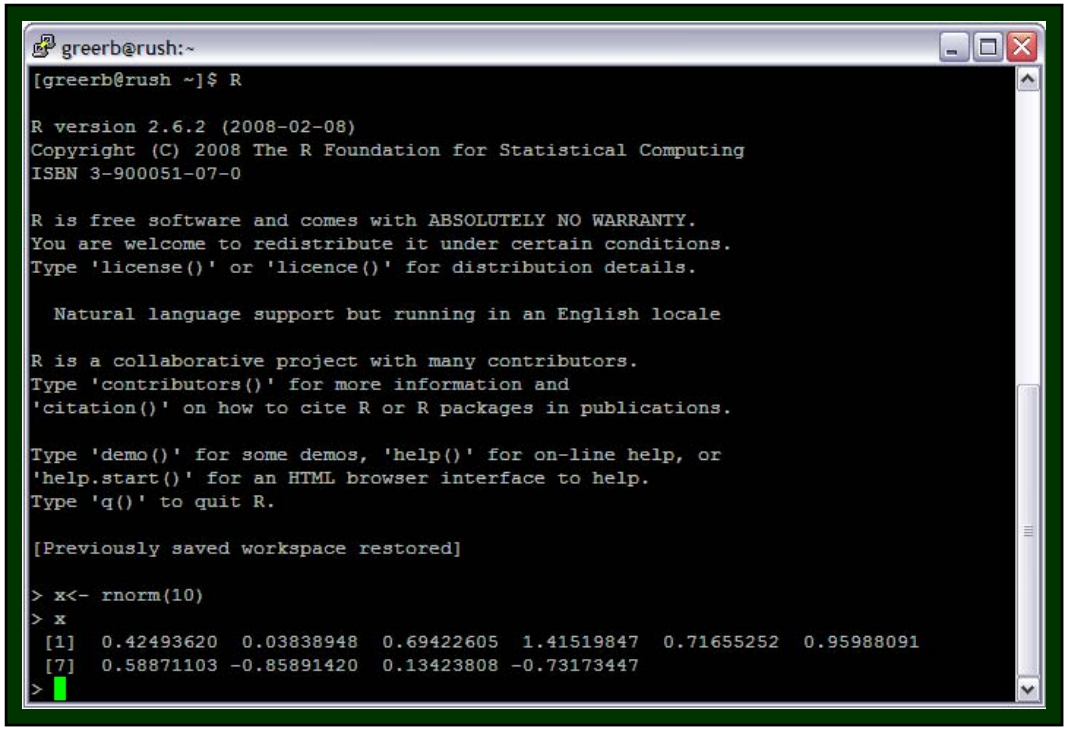

BAYLOR UNIVERSITY

Accessing Matlab

- First, install Xwin32
  - In the Run prompt while connected to the Baylor network, type:

\\bu-shares\BU-software

  - Select the “Starnet” folder
  - Then the Xwin32 folder
  - Copy the folder for the most recent version (9.4) to your desktop
  - Run the executable file and follow install instructions.

Accessing Matlab

- Open Xwin32
  - Two windows will appear
    - License

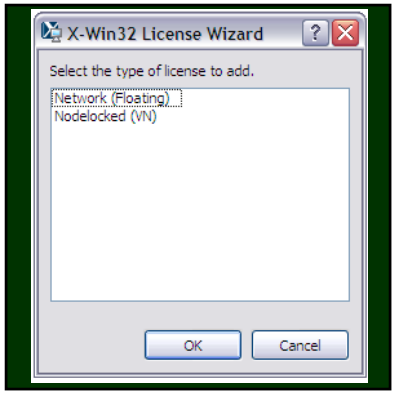

    - Welcome

Accessing Matlab

- License
  - Double click “Network(floating)”
  - Paste the key from the folder you downloaded
- Welcome
  - Create an SSH session
  - Name it whatever you like perhaps “Matlab”
  - The type is “ssh”

BAYLOR UNIVERSITY

Accessing Matlab

- Welcome(cont.)
  - Host: fozzie.baylor.edu
  - Use your Fozzie user name and password to login
  - Command: select Linux
  - “Launch” the session you just set-up
  - You are now in a GUI interface with Fozzie

Accessing Matlab

- Type “matlab” into the command box

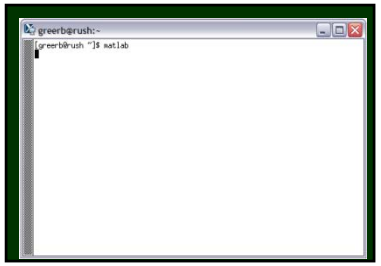

- You now have a GUI interface with Matlab

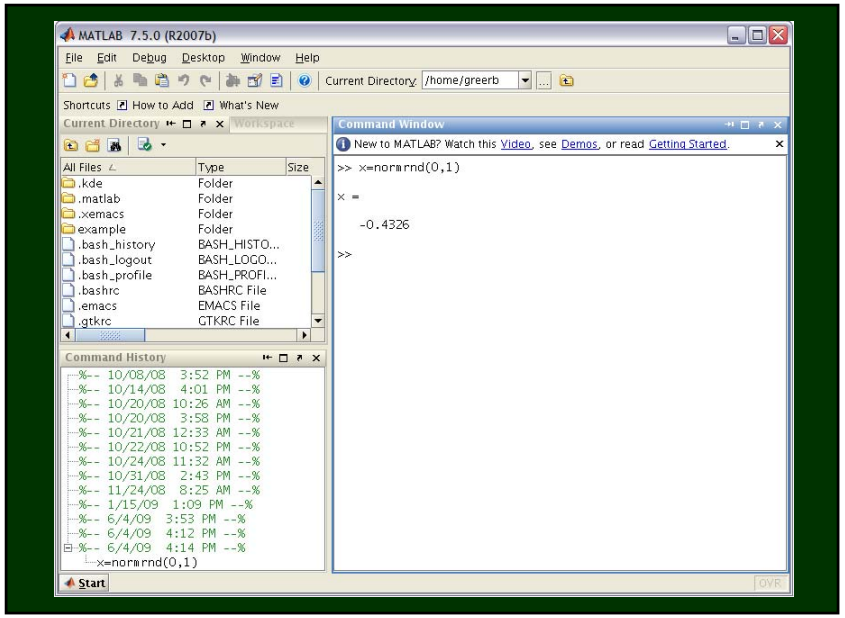

BAYLOR UNIVERSITY

Accessing R

- If you need plotting capabilities in R, you can access R through the Xwin GUI interface

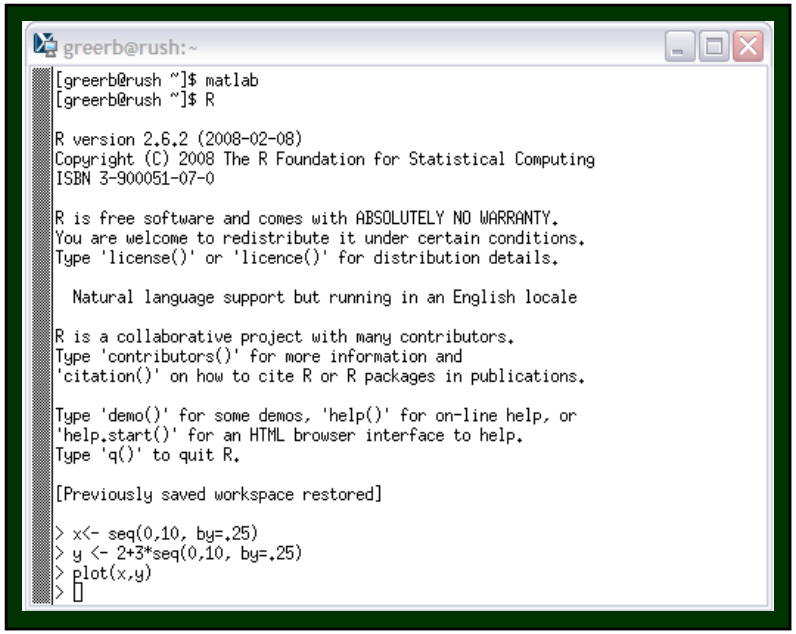

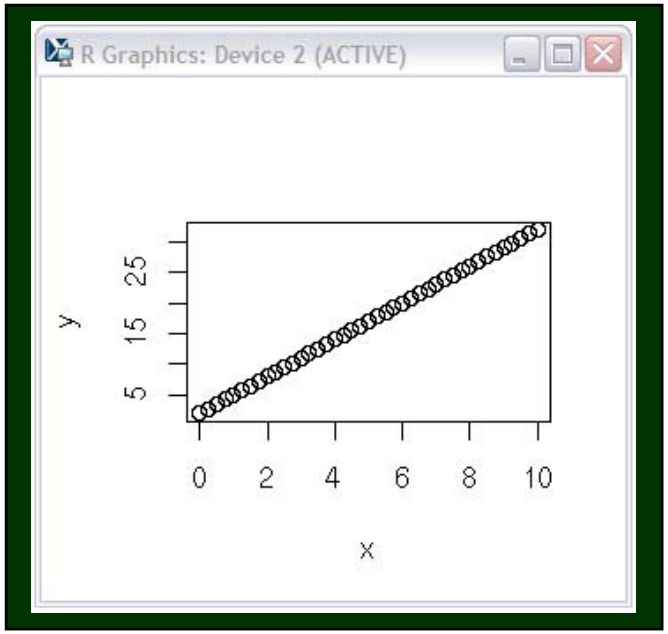

BAYLOR UNIVERSITY