

COMP/CS 605: Introduction to Parallel Computing

Homework 1: Using the Student Cluster/Unix Ops.

Mary Thomas

Department of Computer Science
Computational Science Research Center (CSRC)
San Diego State University (SDSU)

Due: 02/02/15
Posted: 01/24/15
Updated: 01/24/15

Table of Contents

- 1 **HW 1 Overview: Getting Started Using the Student Cluster**
 - HW 1.1: Course Tools & Accounts
 - HW 1.2: Unix operations
 - HW 1.3: Install Pacheco Test Code Library
 - HW 1.4: Running Serial Programs on Student Cluster
 - HW 1.5: Timing Serial Programs
- 2 **General Instructions**

HW 1 Overview: Getting Started Using the Student Cluster

- Due: 02/02/15
- This homework involves demonstrating that you can log onto the student cluster. Tasks:
 - access blackboard information for this class
 - join Google Group and mailing list
 - obtain a user account and logon on the student cluster
 - create the correct homework directory structure
 - perform simple unix operations to get you familiar with the system
 - install the Pacheco demo codes into your directories
 - compile, run, test, time a serial C program (histogram)

HW 1.1.a: Course Tools

Description: setting up and using required course tools.

- Access your Blackboard account
- Respond to the Google Group invitation.
- **What to "turn in":** image of BB access which will also appear in the "last accessed" column; image of Google Group welcome email (keep these small (1/2 page is fine)).

HW 1.1.b: Accessing the student cluster

Description: obtain a user account, and logon using ssh.

- Once you are confirmed in the class, we will create an account using your last name (or using first name initials + last name in the case of duplicate last names).
- This will usually be done after the second day of class.
- **What to turn in:** some image or evidence that you logged on: such as a listing of your directory on tuckoo

HW 1.1.c: Getting host information:

```
[mthomas@gidget:~/junkdir/tdir] nslookup edoras.sdsu.edu
Server:                2605:e000:1c03:7c:9284:dff:fee7:4908
Address:               2605:e000:1c03:7c:9284:dff:fee7:4908#53
```


Non-authoritative answer:

Name: edoras.sdsu.edu

Address: 130.191.17.50

- The cluster is on the internal SDSU campus network with no external login allowed. You can access the cluster from any on campus machine, including the Edoras Academic Computing system, edoras.sdsu.edu.
- Campus Access for SDSUid and VPN:
 - SDSUid: <http://sdsuid.sdsu.edu/>
 - SDSU Palo Alto VPN (GlobalProtect):
<http://edoras.sdsu.edu/vpn/>
 - SDSU Edoras Student Accounts: <http://library.sdsu.edu/computers-technology/hosting/instructor-class-accounts>

HW 1.1.d: Remote SSH login



- Launch *SSH* terminal on your computer
- *SSH* onto edoras:
`%ssh edorasUserName@edoras.sdsu.edu`
- *SSH* onto tuckoo:
`%ssh tuckooUserName@tuckoo.sdsu.edu`
- locate the class homework source code directory:
`/COMP605`

- HW 1.2.a.: Test these Unix operations:
 - `cat /etc/motd`
 - Note 1: try these from your home directory
 - `whoami, date, uname -a`
 - `cd ~, pwd`
 - `ls`, optional arguments [`-al`, `-R`]
 - `mkdir`
 - `chmod`, test arguments such as [`-R`]
- Create, compile and run a serial "Hello *user*" program in C (where *user* is your username).
- Compiler commands: use the specialized parallel library compiler commands
 - C code: `mpicc -o myprogram myprogram.c`
 - C code: `mpif90 -o myprogram myprogram.f90`
 - where is the command installed?
 - `%locate mpicc`
- **What to turn in:** evidence that you completed this: images, text file which contains the output captured, session output.

HW 1.2.b: Homework Directory Operations

Once you have account information, you will create a homework directory where I will look for all assignment material. You are free to create other directories for development and testing (e.g. dev), but these directories must *only* contain specific material for the assignment.

- log onto student cluster: tuckoo.sdsu.edu (130.191.127.136)
- create a homework directory call hw
- create a sub directory called hw1
- use unix command *chmod* to set the directory accesses so only you and the instructor can read/access the codes
- **What to turn in:** evidence that you completed this: an image, listing, etc.

HW 1.3.a: Install the Pacheco demo codes

Description:

- tar files are located in `/COMP605/pacheco_examples`
- you may want to read the Unix man pages or Web pages to learn about the commands for *tar* and *gzip*
- you will want to locate the histogram code in the IPP codebase (ch2)
- **What to turn in:** evidence that you completed this: directory listing.

HW 1.4.a: Compile & Run Serial Pacheco Histogram Code

- run the code for different variables:

usage: ./histogram

< *bin_count* >< *min_meas* >< *max_meas* >< *data_count* >

- use the following test cases (8 combinations):
 - *bin_count* = [5, 20]
 - *min_meas* = 1
 - *max_meas* = [500, 3000]
 - *data_count* = [50, 500]
- **What to turn in:** evidence that you completed this: screen image, or text copy of output.

HW 1 Overview: Getting Started Using the Student Cluster

HW 1.4: Running Serial Programs on Student Cluster

HW 1.4.b: Running the histogram program

```
[mthomas@tuckoo pacheco]$ cd intro-par-pgming-pacheco/
[mthomas@tuckoo intro-par-pgming-pacheco]$ ls
total 876
drwx----- 3 mthomas mthomas 4096 Mar  1 16:19 .
drwx----- 4 mthomas mthomas 4096 Mar  4 11:22 ..
drwx----- 8 mthomas mthomas 4096 Feb 14 2014 ipp-source
-rwx----- 1 mthomas mthomas 880640 Oct 16 2012 ipp-source.tar
[mthomas@tuckoo intro-par-pgming-pacheco]$ cd ipp-source
[mthomas@tuckoo ipp-source]$ ls
total 52
drwx----- 8 mthomas mthomas 4096 Feb 14 2014 .
drwx----- 3 mthomas mthomas 4096 Mar  1 16:19 ..
drwx----- 2 mthomas mthomas 4096 Jan 19 2011 ch2
drwx----- 2 mthomas mthomas 4096 Apr 27 13:57 ch3
drwx----- 2 mthomas mthomas 4096 Apr 23 13:48 ch4
drwx----- 3 mthomas mthomas 4096 Nov  1 2012 ch5
drwx----- 2 mthomas mthomas 4096 Jan 15 2012 ch6
-rwx----- 1 mthomas mthomas 14515 May 26 2011 INDEX
drwx----- 7 mthomas mthomas 4096 Feb 14 2014 ipp-source
-rwx----- 1 mthomas mthomas 1694 Jan  7 2011 README
[mthomas@tuckoo ipp-source]$ cd ch2
[mthomas@tuckoo ch2]$ ls
total 20
drwx----- 2 mthomas mthomas 4096 Jan 19 2011 .
drwx----- 8 mthomas mthomas 4096 Feb 14 2014 ..
-rwx----- 1 mthomas mthomas 8638 Jan 19 2011 histogram.c
[mthomas@tuckoo ch2]$ mpicc -o histogram histogram.c
[mthomas@tuckoo ch2]$ ./histogram 10 1 1500 100
1.000-150.900: XXXXXXXX
150.900-300.800: XXXXXXXX
300.800-450.700: XXXXXXXXXX
450.700-600.600: XXXXXXXXXX
600.600-750.500: XXXXXXXX
750.500-900.400: XXXXXXXXXX
900.400-1050.300: XXXXXXXXXXXXX
1050.300-1200.200: XXXXXXXXXX
1200.200-1350.100: XXXXXXXXXXXXX
1350.100-1500.000: XXXXXXXXXXXXXXX
[mthomas@tuckoo ch2]$
```

HW 1.5: Timing the Histogram code

- Time how long the code takes to run as a function of Problem Size:
 $ProbSize = data_count$
Wallclock Time: T_{wall}
- use the following test cases:
 - $bin_count = 10$
 - $min_meas = 1$
 - $max_meas = 50000$
 - Vary $data_count = 10^n$, where $n = 0, 1, 2, N_{max}$
 - What is N_{max} ? Explain the limit.
- Modify how the code prints out the results: you don't need to "plot" the histogram
- You only need to printout the bins, ranges, and the count (not all the X's).
- Plot your test results using excel, Matlab, by hand: T_{wall} vs $ProbSize$

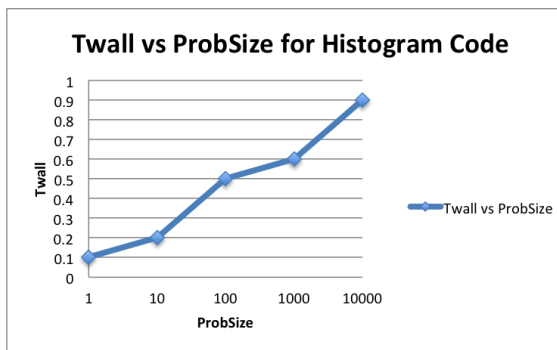
HW 1.5: Timing the Histogram code (cont.)

- Suggestions for timing:
 - Use C/Fortran internal timer
 - You can find example timer code in `/COMP605/tools/code_timer.c`
 - Avoid using the unix `date` function, it is not very sensitive

HW 1.5: Data Analysis and Presentation

A table of your test data and a plot of the results. Include labels.

ProbSize	T_{wall}
1	0.1
10	0.2
100	0.5
1000	0.6
10000	0.9



Homework

- Put homework into a directory:
HOME/ < your_username > /hw/hw1
- include the source code(s), compiled binaries
- see sections above for what to include
- Write a simple report (this can be TEXT, Word, PDF Doc).
- See each section for what to turn in.
- Turn in hard copy (condensed/minimal number of pages) at start of class.

Once the submission timeline has closed
DO NOT CHANGE THE FILE TIMESTAMPS!

HW directory listing example

```
[mthomas@tuckoo_hw_dir_ex]$ ls -R
.:
total 12
drwx----- 3 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 3 mthomas mthomas 4096 Aug 28 12:22 ..
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 username

./username:
total 20
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 3 mthomas mthomas 4096 Aug 28 12:23 ..
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 dev
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 hw
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 misc

./username/dev:
total 8
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..

./username/hw:
total 20
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..
drwx----- 2 mthomas mthomas 4096 Aug 28 12:24 hw1
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 hw2
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 hw3

./username/hw/hw1:
total 8
drwx----- 2 mthomas mthomas 4096 Aug 28 12:24 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..
-rw-rw-r-- 1 mthomas mthomas  0 Aug 28 12:24 data1.dat
-rw-rw-r-- 1 mthomas mthomas  0 Aug 28 12:24 file1.c
-rw-rw-r-- 1 mthomas mthomas  0 Aug 28 12:24 file2.c

./username/hw/hw2:
total 8
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..

./username/hw/hw3:
total 8
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..

./username/misc:
total 8
drwx----- 2 mthomas mthomas 4096 Aug 28 12:23 .
drwx----- 5 mthomas mthomas 4096 Aug 28 12:23 ..
[mthomas@tuckoo_hw_dir_ex]$
```