

The Anatomy of a Science Engagement:  
Analyzing the Human Connectome Project (HCP) Datasets using  
GPUs

John-Paul Robinson, MS  
Research Computing Day  
October 13, 2017

Everything's going so well!!!



11:52am

[VNC] shows error:

" No connection could be made because the target machine actively refused it". Is there any thing wrong at the server end? What is the problem? Thanks.

11:52am

[VNC] shows error:

" No connection could be made because the target machine actively refused it". Is there any thing wrong at the server end? What is the problem? Thanks.

12:03pm

I can access cheaha via an ssh session but I am trying to use a VNC session and its frozen. Can't even get a terminal to open.

12:13pm

I noticed that Cheaha is taking an unusually long time to respond to tab completion, ls, and it's often several seconds before what I type appears in the command line.

11:52am

[VNC] shows error:

" No connection could be made because the target machine actively refused it". Is there any thing wrong at the server end? What is the problem? Thanks.

12:03pm

I can access cheaha via an ssh session but I am trying to use a VNC session and its frozen. Can't even get a terminal to open.

12:13pm

I noticed that Cheaha is taking an unusually long time to respond to tab completion, ls, and it's often several seconds before what I type appears in the command line.

11:52am

[VNC] shows error:

" No connection could be made because the target machine actively refused it". Is there any thing wrong at the server end? What is the problem? Thanks.

12:48pm

The viewer is lagging and freezing for extended periods. The connection has also been terminated during a session on several occasions this morning


12:03pm

I can access cheaha via an ssh session but I am trying to use a VNC session and its frozen. Can't even get a terminal to open.

WTH!?!?

## Round Trip Time for 10.0.0.11:41794 → 164.111.161.5:22

rsc-ssh-tcp-errors-20170213.dump.gz



Click to select packet 270 (8.494s len 0 seq 9017 ack 8425 win 500) → 304 pkts, 16 kB ← 259 pkts, 73 kB

Type

Stream

Mouse ☒ drags ☐ zooms

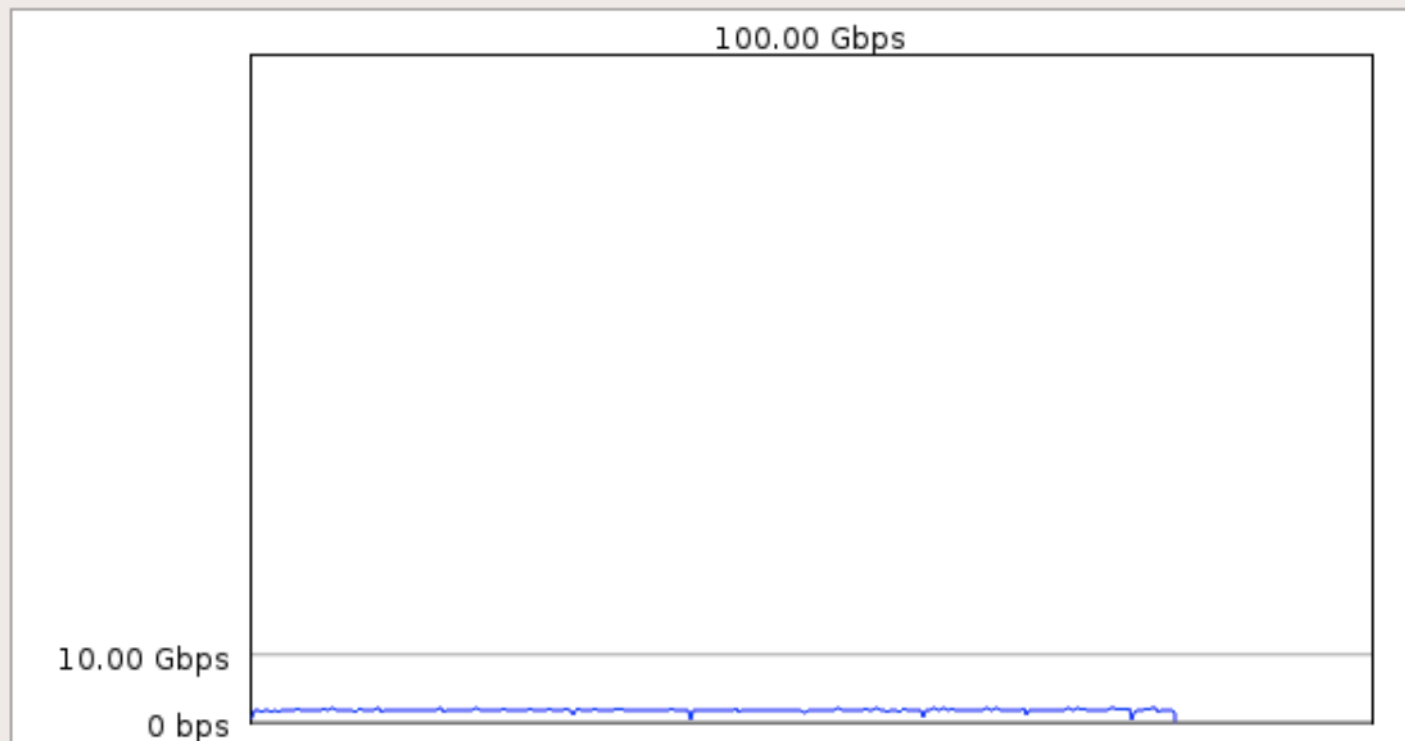


Transfer Monitor

100307\_3T\_Structural\_preproc.zip (+199)



Done



Fair

# Impact of 1Gbps Data Flow

- Background:
  - UAB Campus Network and Firewall upgraded in Fall 2016
  - Campus now has 100Gbps to Internet and individual flows from off campus can reach multi-gigabit speeds
  - Campus core now has 40Gbps capacity
  - Traditional scp transfers only reach ~200Mbps flows and old campus board limited flows to 100's Mbps
  - Aspera Connect used by Human Connectome Project able to push 1.4Gbps flows
- Root cause:
  - Cluster network link to core needed upgrade from 1Gbps to 10G building speed
- New resources:
  - Science DMZ can transfer ~5 Terabytes per hour (so far...)

Everything's going so well!!!



I am wanting to start an interactive job in cheaha, but cannot get past the “waiting for job to start” stage.

Members of our lab are experiencing difficulty getting an S interactive terminal to load today.

# *Brainhack Global*

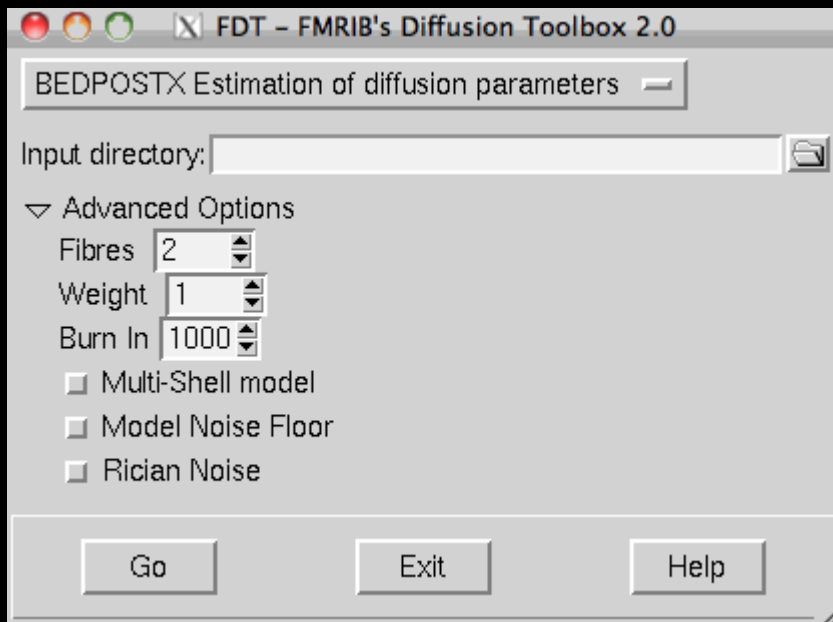
March 2-5, 2017

# Toolkits for Brain Image Processing



FSL is a comprehensive library of analysis tools for FMRI, MRI and DTI brain imaging data. It runs on Apple and PCs (both Linux, and Windows via a Virtual Machine), and is very easy to install. Most of the tools can be run both from the command line and as GUIs ("point-and-click" graphical user interfaces). – FSL docs

# Bayesian Estimation of Diffusion Parameters Obtained using Sampling Techniques




bedpostx allows to model crossing fibres within each voxel of the brain. Crucially, bedpostx allows to automatically determine the number of crossing fibres per voxel. – FSL docs

# What's Going On?


- Pre-processing HCP900 data set for Brainhack conference
  - 900+ subjects with over 1 TB of data to process
  - BEDPOSTX preps data and submits work to xfibres binary
- Problems
  - BedpostX scripts don't detect Slurm cluster, submit non-parallel jobs
  - Each subject consumes ~520 compute hours (21 days / single core)
  - Running all subjects at once uses 900 cores (40% cluster) but still takes 21 days
  - Cluster partition limits only apply per-job, not across multiple jobs
  - Cluster can flood with work from one user




# BEDPOSTX is “Pleasantly Parallel”



**File size per Slice (Megabytes)**  
Tracks voxel count in slice



**RAM per Slice (Megabytes)**  
Tracks voxel count in slice

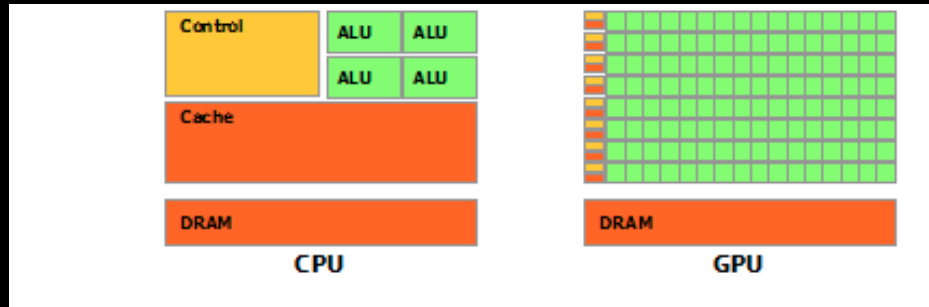


**Runtime per slice (Seconds)**  
on single core

# What Did We Do?

- Understand BEDPOSTX operation
- Port BEDPOSTX scripts to recognize slurm and submit parallel jobs
- BEDPOSTX splits image into slices, about 145 slices/image
- Each slice runs in one job on single core
- Reduce processing time to 7 hours per subject
- Problems
  - Each subject consumes 145 cores or 6 cluster nodes
  - Could consume entire cluster for two weeks (16 days)

# Enter the GPU



- CPUs use chip space for large, fast memory caches to speed execution of serial codes.
- GPUs use chip space to add 1000's of cores to speed execution of data-parallel codes
- Xfibres ported to GPU, 2013
- Cluster has 2x K80 GPUs with 2496 cores
- **GPUs process one subject in 2 hours**
- Still takes ~21 days to process whole data set because of limited GPU throughput (too few GPUs)
- CPUs are free for others

# Lessons

- Use the cluster as you need to accomplish your science goals
- If performance is not what you expect, contact us. (And we may contact you! :) )
- Fixes are not always immediate
- Sometimes they involve a lot of work and lead to extended collaborations
- Maintaining software requires person hours – all solutions shown here are still hand crafted and not in general code-base
- Need to grow software engineering practices for scientific computing

# Outcomes

- Demonstrated reality of “Big Data” and advantages of GPU computing for select problems
  - Even if you can’t use GPUs you should be happy your peers can!
  - GPUs help keep CPUs free for general purpose computing
- Contributed additional use-case to justify GPU expansion
- New GPUs can run BEDPOSTX on one HCP subject in 10 minutes
- Entire HCP900 data set could now be run in 8.5 hours!
- Strengthened capability of UAB research and support teams

Keep doing what you do.

We promise we'll keep doing what we do.

Special thanks to Kristina Visscher PhD, Sara Sims, and the Visual  
Brain Core