18 September - Meeting

- Review high level objective, project plan
- Complete team roles
- Ensure everyone is set with team logistics:
    - Weekly Report
    - Drive, Todoist, Github
    - Who needs help learning Git? Short tutorial on git after the meeting.
    - Journaling
    - Related Work
- Overview of FPGAs by Dr. Ganesan
- Walkthrough of SCCs on Matlab as a first simple algorithm

Warp algorithm?? See Ganesan's papers
Ganesan suggests target applications

1. Database
2. Biology - Protein Groups

Jay Li - PhD student finishing his thesis and starting a job. Busy.
Han Yu - PhD student working on GPUs, computational biology. May help with project.

Two reasons to use new hardware (GPU or FPGA):

1. Data set so large we need more processing capability
2. Nature of processing itself is faster on new hardware

FPGAs: good in diverse processing - not good at floating point ops
GPUs: good for data parallelism - great at floating point ops

FPGAs - each logic block has configurable lookup tables
Schematic design given by hardware manufacturer
HDL = Hardware description language - VHDL, Verilog
Use *cores* (prebuilt VHDL components) and put them together with logic to write a VHDL program.
⇒ See **OpenCore** for examples of open source core components you can use in hardware design
Write in high level language, then use netlist and synthesizer to get the hardware code.

Link to accumulo user guide (also on Drive) (Accumulo examples)

Link to Xilinx download guide from Dr. Ackland for CpE 487. He uses Xilinx version 13 for some reason. I have 14.1 installed.

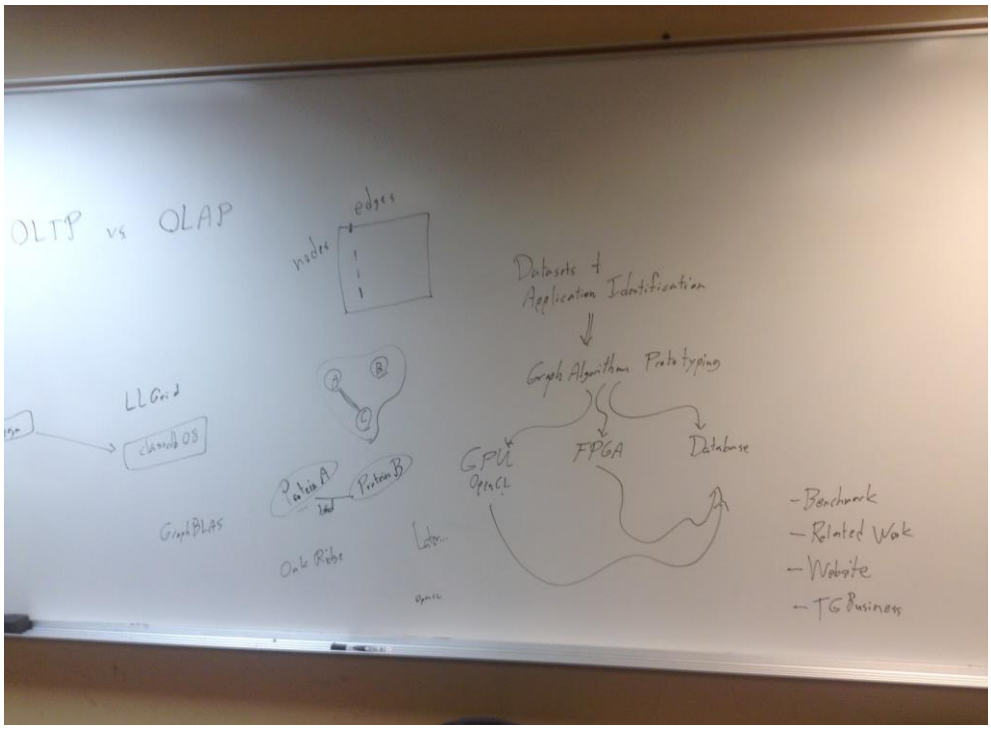

[All] Take a look at Dr. Ganesan's papers and references.

[All] Take a look at the Graph Computing webpage I showed today and use that as an example of datasets and analytic questions we can ask about datasets that we can answer in the form of a graph algorithm. **Next week, come with**

- **one dataset** you think is cool (you don't have to download an actual dataset, could be as vague as "social network data")
- **one question** you would like to ask about that dataset
- **one graph algorithm** that can answer that question
- **ideas for implementing it**, including whether you think a GPU, FPGA or database can help

You can write this in your journal

[All] Ensure you put your Github username inside your journal so that I can add you to the Stevens-GraphGroup organization.

[All] Check that you can see the shared Todoist project. If you can't, make sure your email (that is, the email you login to Todoist with) appears correctly on this list:

Share options

Email or Name

Invite

BX Brian Cesar Tondreau Xhero91 PENDING
xhero91@gmail.com

HY Hefei Yang Youflyaway PENDING
youflyaway@gmail.com

JJ Jaroor Modi Jaroorm96 PENDING
jaroorm96@gmail.com

MM Mojie Yao Myao1 PENDING
myao1@stevens.edu

NN Narayan Ganesan Nganesan PENDING
nganesan@stevens.edu

XX Xin Li Xli26 PENDING
xli26@stevens.edu

XX Xuelian Liu Xliu43 PENDING
xliu43@stevens.edu


Me
dyhutchison@gmail.com

DR Di Ren
rendi6668@gmail.com

EC eric cherin
eric.cherin@gmail.com

PG Peiran Guan
pguan@stevens.edu

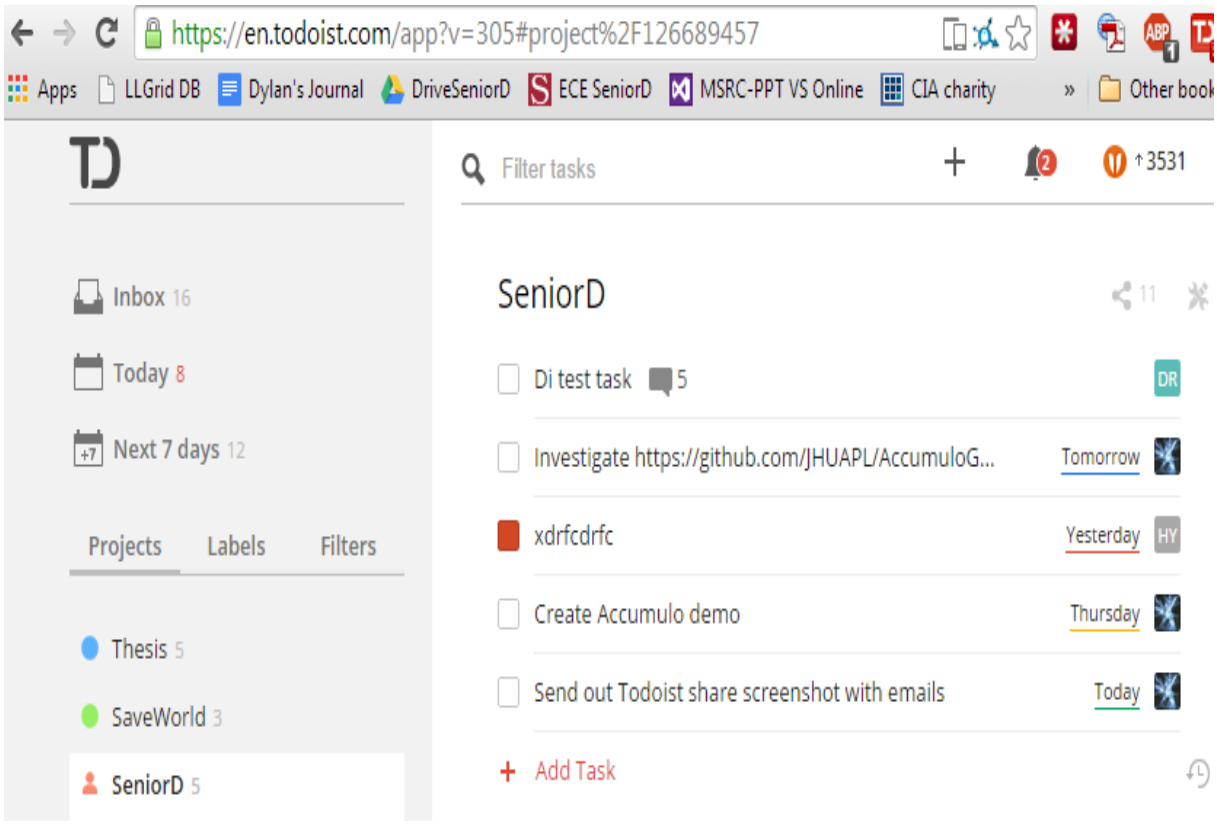

[Xin] Start website planning and development. Decide what framework you will use for the website; if not sure list a couple options with pros and cons and ask us next week. See if you can get an index file up on Github.
Create a repository named Stevens-GraphGroup.github.io under the organization and follow the steps here.

[Dylan] Prepare a demo of Accumulo with Twitter data for next week. Accumulo architecture review.