High-Throughput High-Resolution Cryo-EM on the Cheap

EM close to the bone

Scott Stagg
Institute of Molecular Biophysics
Florida State University

Computational needs

- Data collection
- Data storage
- Preprocessing
- Refinement
- Data transfer and archiving

Computational environment - two models

- Do it all yourself
  - Advantages
    - Atomic control over environment
  - Disadvantages
    - Challenges getting necessary expertise
    - $$$
- Distributed
  - Advantages
    - Rely on existing expertise
    - Distributed costs
  - Disadvantages
    - No control over computational environment

FSU solution - partnership with FSU Research Computing Center (RCC)

- RCC hosts
  - Database
  - Leginon/Appion website
  - Primary storage
  - High performance computing (HPC)

FSU Pricing

- 1 normalized compute unit (NCU) - $221.43 / 5 yrs
  - With 26 GB storage
  - 32 cores – $7,085.76
  - 832 GB storage
- 1 TB high-performance storage - $1450 / 5 yrs
- 1 TB archival storage - $55-$110 / 5yrs

The trouble with not controlling the environment

- Annual breakage
  - FSU HPC upgrades the OS and libraries every year.
  - This essentially breaks everything until I have a chance to recompile/update/debug all the software packages of interest
- Appion example
  - FSU HPC – CentOS 7, Apache 2.4.6, PHP 5.6
  - Parts are incompatible with Appion
    - Since we host on FSU HPC, we have a hard time getting the environment set up right

Docker and Containers

- Docker
  - “wraps up a piece of software in a complete filesystem that contains everything it needs to run: code, runtime, system tools, system libraries – anything you can install on a server. This guarantees that it will always run the same, regardless of the environment it is running in.”
- Using containers allows us to separate OS release from required software packages
  - Now OS, PHP, Apache, and MySQL are all independent of each other

Potential for the future

- Using containers, will be able to ship data together with the complete user interface

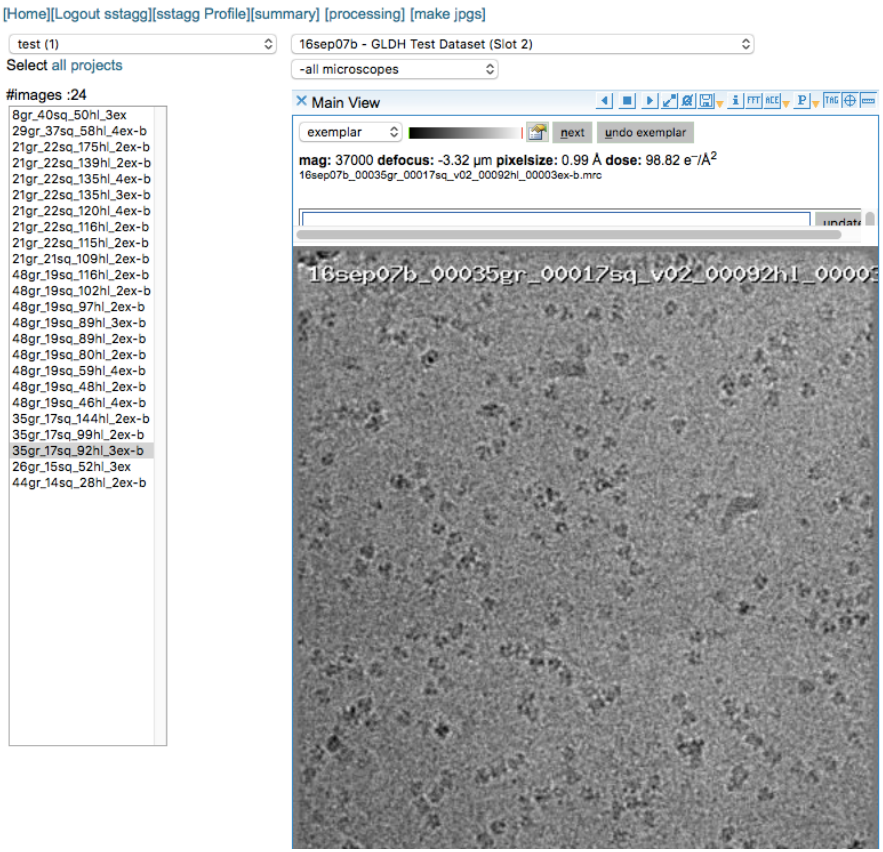

Modules for managing your environment

```
[sstagg@krios ~]$ module list
Currently Loaded Modulefiles:
  1) gnu-openmpi/1.10.2   8) relion1_4          15) direx
  2) eman2                9) simple             16) ctffind4
  3) de_process_frames   10) spider             17) localrec
  4) eman1               11) xmipp              18) scipion
  5) frealign            12) komodo             19) myamiss
  6) ihrsr               13) ffmpeg
  7) protomo             14) bsoft
[sstagg@krios ~]$
```

Modules for managing your environment

```
[sstagg@krios ~]$ module avail

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

---------------------------- /etc/modulefiles ----------------------------
g09test                       pgi-openmpi
gaussian09                    python3
gnu-mvapich2                  R/3.1.3
gnu-openmpi/1.10.2(default)   R/3.2.0
intel-mvapich2                R/3.2.5(default)
intel-openmpi                 stata/10(default)
matlab_dcs                    stata/13
orca                          stata/9
pgi-mvapich2

----------- /panfs/storage.local/imb/stagg/software/etc/modules -----------
bsoft              frealign           myami-3.1          protomo2_3_1
ctffind4           frealign9-08       myami-3.2          relion1_3
de_process_frames  frealign9-09       myami_container    relion1_4
direx              frealign9-10       myamidev           scipion
eman1              ihrsr              myamidevbeta       simple
eman2              komodo             myamiss            simple2
eman2_12           localrec           myamiweb           situs
eman2mpi           matlab2014tmp      openmpi-1-8-3      spider
ffmpeg             matlabtmp          protomo            xmipp
[sstagg@krios ~]$
```

Computational needs

- Data collection
  - Leginon
    - At minimum, requires a database and webserver
- Data storage
- Preprocessing
- Refinement
- Data transfer and archiving

Basic setup - Leginon

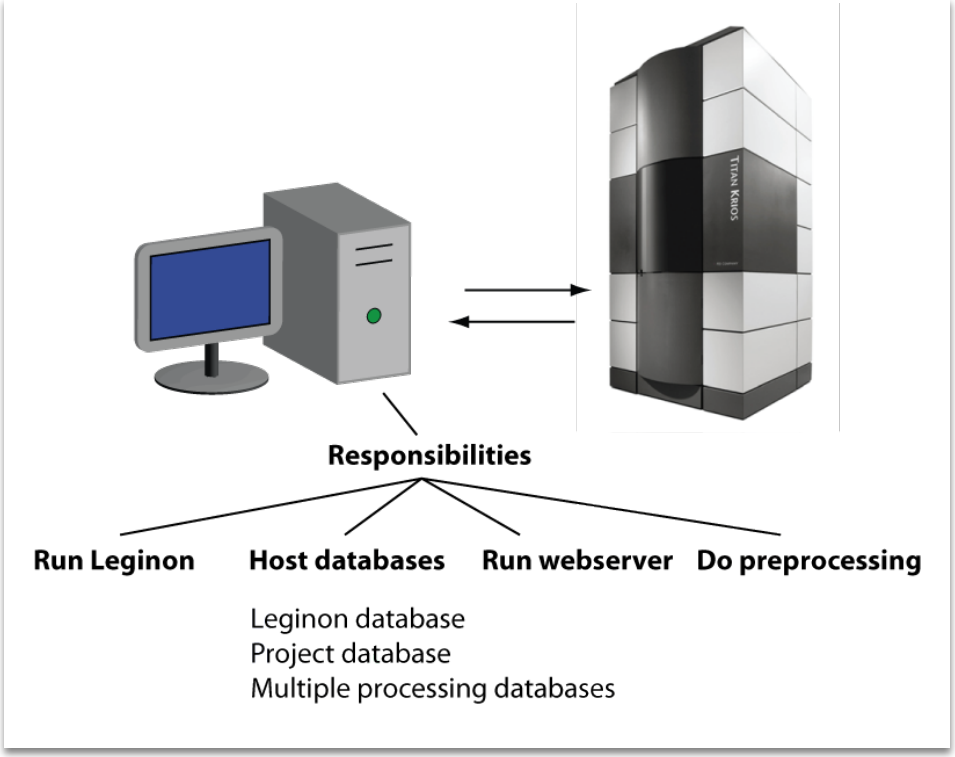

High-throughput creates two bottlenecks

- Disk space
- Processing and reconstruction
- Both problems are solved by collaborating with the high performance computing center at FSU

Distributed setup

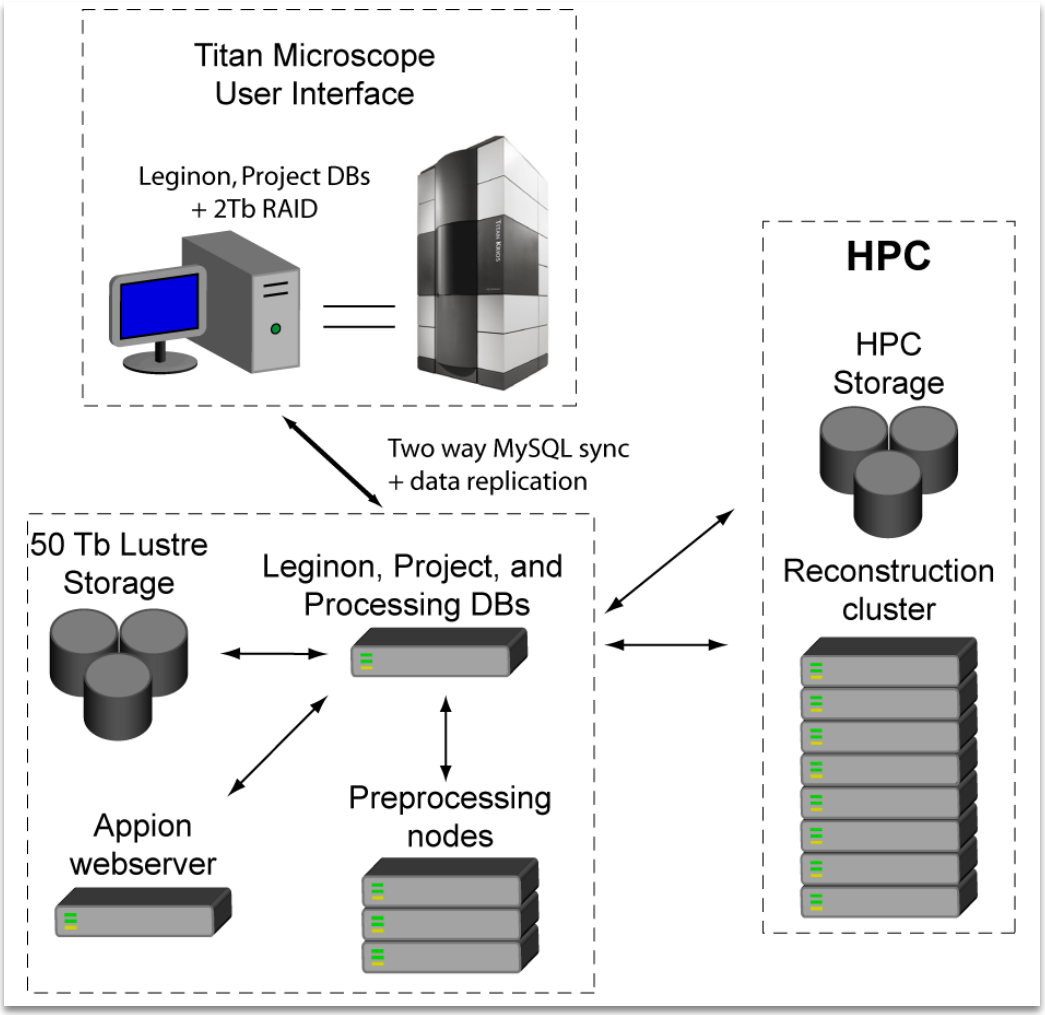

Schrum *et al.* JSB, 2012

Handling synchronization

- iwatch/inotify - take care of data replication
  - run in command line mode as well as in daemon mode
  - using an easy xml configuration file
  - can watch directory recursively and watch new created directory
  - can have a list of exceptions
  - can use regex to compare the file/directory name
  - can execute command if an event occurs
  - send email
  - syslog
  - print time stamp
- MySQL - capable of two-way database replication with proper setup

Schrum *et al.* JSB, 2012

Computational needs

- Data collection
- Data storage
  - Two types are needed – live and archival
- Preprocessing
- Refinement
- Data transfer and archiving

Multi-user facilities require high performance storage file systems for data collection and processing

- RAID
  - Pros – redundant (lower failure risk), lower cost, modest computational scalability
  - Cons – fixed volume size, limited simultaneous reads/writes
- Scalable file systems
  - Pros - central expandable volume, high parallel performance
  - Cons – cost
  - Examples
    - Lustre
    - Panasas
    - GlusterFS
    - GPFS
    - others

Archival storage may be cheap, but not high-performance storage

- 8 TB disk – $250
- 8 TB of high performance Lustre space for 5 yrs on HPC - $11,600

Computational needs

- Data collection
- Data storage
- Preprocessing
  - Frame alignment, CTF estimation, particle picking
- Refinement
- Data transfer and archiving

At FSU preprocessing is handled through Appion

- For SECM[4], we provide frame alignment as part of the service provided with data collection
  - Use DE frame alignment software – integrated in Appion
    - Tends to work better for integrated frames than motioncorr
  - Parallelized on HPC
  - 10 Gb “FASTLANE” fiber to HPC combined with parallel frame processing keeps up with data collection
- Also provide CTF estimation because it adds value to data collection
- Sometimes do particle picking to get statistics for user

Computational needs

- Data collection
- Data storage
- Preprocessing
- Refinement
  - Up to the user
- Data transfer and archiving

Computational needs

- Data collection
- Data storage
- Preprocessing
- Refinement
- Data transfer and archiving

Data transfer and archiving

- Our workflow
  - 10Gb from camera to HPC storage
    - Infiniband on HPC
    - Frame alignment on HPC
  - 10Gb from HPC storage to external hard drive
    - Ultimately bandwidth limited by disk speed without having RAID or parallel FS
  - Make copy of hard drive
  - Ship one copy to user
    - Keep other copy until user verifies that they have made local copy of data
    - Then ship user the other copy and wipe from local storage

In near future

- Florida HPCs recently completed an “invitation to negotiate” (ITN) for high volume storage solutions
  - Winning bid will provide storage for $22 per TB per year
- Still slightly more expensive than buying disks but will be redundant, secure, high availability, fast write speeds, and can facilitate automated transfer

What about cloud solutions?

- Preliminary investigations suggest that transfer speed is insufficient for frame storage
  - From Donny Shrum of FSU's HPC – Amazon S3 write speeds vary between 300K and 5 megabytes / s
  - We collect ~2 TB of frames per day – so 5 days to upload a single day of data, download speeds are slower
- Instead of frames, could store aligned, compensated, summed images on the cloud
- Data security?