


Diversity of Science on the OSG



Open Science Grid

Rob Gardner • University of Chicago
Ken Herner • Fermilab

OSG All Hands Meeting, University of California San Diego, March 6-9, 2017



Overview

- Review of opportunistic usage on OSG
- Highlights from electrical engineering, neuroscience, evolutionary biology, genomics, medical science & more!

● Summary of the Virtual Organizations

University of Pittsburgh, Indiana University, Tulane University, Syracuse University, University of California, Berkeley, Georgia Institute of Technology, University of Nebraska, Clemson University, University of Michigan, University of Wisconsin-Madison, University of Chicago, University of Illinois at Urbana-Champaign, Space Telescope Science Institute, University of New Mexico, University of Pennsylvania, University of Minnesota, Texas A&M University, University of Washington, SUNY at Stony Brook, UT Southwestern, University of Wisconsin, University of Arizona, Brookhaven National Laboratory, University of New England, Fermilab, University of Oklahoma, University of North Carolina at Chapel Hill, Colorado School of Mines, University of Illinois, Harvard Medical School, Georgia State University, Washington State University, Worcester Polytechnic Institute, Southern Illinois University, Michigan State University, University of Minnesota, Duluth, Western Washington University, University of Nebraska Lincoln, University of Nebraska - Lincoln, Rutgers, the State University of New Jersey, Los Alamos National Lab, College of Charleston, University of Wisconsin - Madison, Hobart and William Smith Colleges, Pennsylvania State University, University of California San Diego, University of Wyoming, Albert Einstein College of Medicine, Massachusetts Institute of Technology, Duquesne University, University of Nebraska-Lincoln, University of Missouri, University of Cincinnati, University of Oxford, University of Minnesota Duluth, Bryn Mawr College, Iolani School, University of Wisconsin Madison, University of Texas - Austin, Northwestern University, Rutgers University, Rochester Institute of Technology, University of California, San Francisco, Rutgers, Baylor College of Medicine, University of Notre Dame, International Center for Theoretical Physics, Siena College, Fermi National Accelerator Lab, UC Riverside, Georgia Tech, SLAC, Swarthmore College, OSG, University of California Santa Cruz, University of Tennessee, Knoxville, Open Science Grid, Columbia University in the City of New York, National Institute for Computational Sciences, Stanford University, University of Central Florida, New Jersey Institute of Technology, Oklahoma State University

Total Production (VO Summary)



VO Summary



Zoom Out

a year ago to Mar 6, 2017 23:59:59 UTC



Bin size

7d

VO

All

Job Type


Batch

Batch vs Payload Jobs

Total Wall Hours

1.375 Bil


Wall Hours by Usage Model by 7d



Total CPU Hours

956 Mil


Wall Hours by Site per 7d



Total Jobs

136.3 Mil

Jobs by Exit Code per 7d



This talk is about the opportunistic science filling the capacity

VO Summary w/o ATLAS and CMS



Total Wall Hours

390 Mil


Total CPU Hours

263 Mil


Total Jobs

31.3 Mil


Wall Hours by Usage Model by 7d



Wall Hours by Site per 7d



Jobs by Exit Code per 7d



About 28% of batch hours
in the past year were not
ATLAS or CMS



Opportunistic

Bin auto Filter +

Opportunistic Wall Hours per 7d




Total Opportunistic Wall Hours

176 Mil


Top 20 Sites Providing Opportunistic Usage

OIM_Site	OIM_Facility	Wall Hours
SU OrangeGrid	Syracuse University	24 Mil
FNAL USCMS Tier1	Fermi National Accelerator Laboratory	19 Mil
FermiGrid	Fermi National Accelerator Laboratory	19 Mil
Caltech CMS Tier2	California Institute of Technology	17 Mil
MWT2 ATLAS UC	University of Chicago	15 Mil
UCSD CMS Tier2	University of California San Diego	13 Mil
Nebraska-CMS	University of Nebraska	13 Mil
GLOW	University of Wisconsin	10 Mil
Nebraska-Omaha	University of Nebraska	9 Mil
Brookhaven ATLAS Tier1	Brookhaven National Laboratory	5 Mil
AGLT2	University of Michigan	4 Mil
Nebraska-Lincoln	University of Nebraska	4 Mil
UConn-OSG	University of Connecticut	3 Mil
OSG GOC	Indiana University	3 Mil
MIT CMS	Massachusetts Institute of Technology	3 Mil
Boston University ATLAS Tier2	Boston University	2 Mil
OU ATLAS	University of Oklahoma	1 Mil
Clemson IT	Clemson University	1 Mil
Purdue CMS	Purdue University	1 Mil
UC Davis CMS T3	University of California Davis	1 Mil

Opportunistic Wall Hours by Site (Top 20)





Opportunistic Wall Hours by VO (Top 20)





Opportunistic VO usage trends in 2016



Opportunistic Monthly Hours Spent on Jobs By VO in 2016



Opportunistic Wall Hours by VO in 2016



N.B. Mu2e folded into Fermilab in March

**Large ATLAS and CMS increases late in year:
Expected to continue**



osg connect



OSG Connect Summary




Zoom Out Last 1 year UTC

Bin size 7d Access Point osgconnect.net Project All User All

Total Wall Hours

52.1 Mil

Total Wall Hours for Top 10 Projects




	values
z2dqmc	9172329
FutureColliders	7823259
SourceCoding	7679865
seq2fun	5658408
BioGraph	4089017
AlGDock	2503503
VERITAS	2391222
PRTH	1907090
PainDrugs	1166063
PreBioEvo	1085337


Total Jobs

41.3 Mil


Wall Hours per 7d



Unique Projects per 7d




Total Jobs per 7d



Unique Users

283


Unique Users per 7d



Unique Projects


81

▼ By Field Of Science



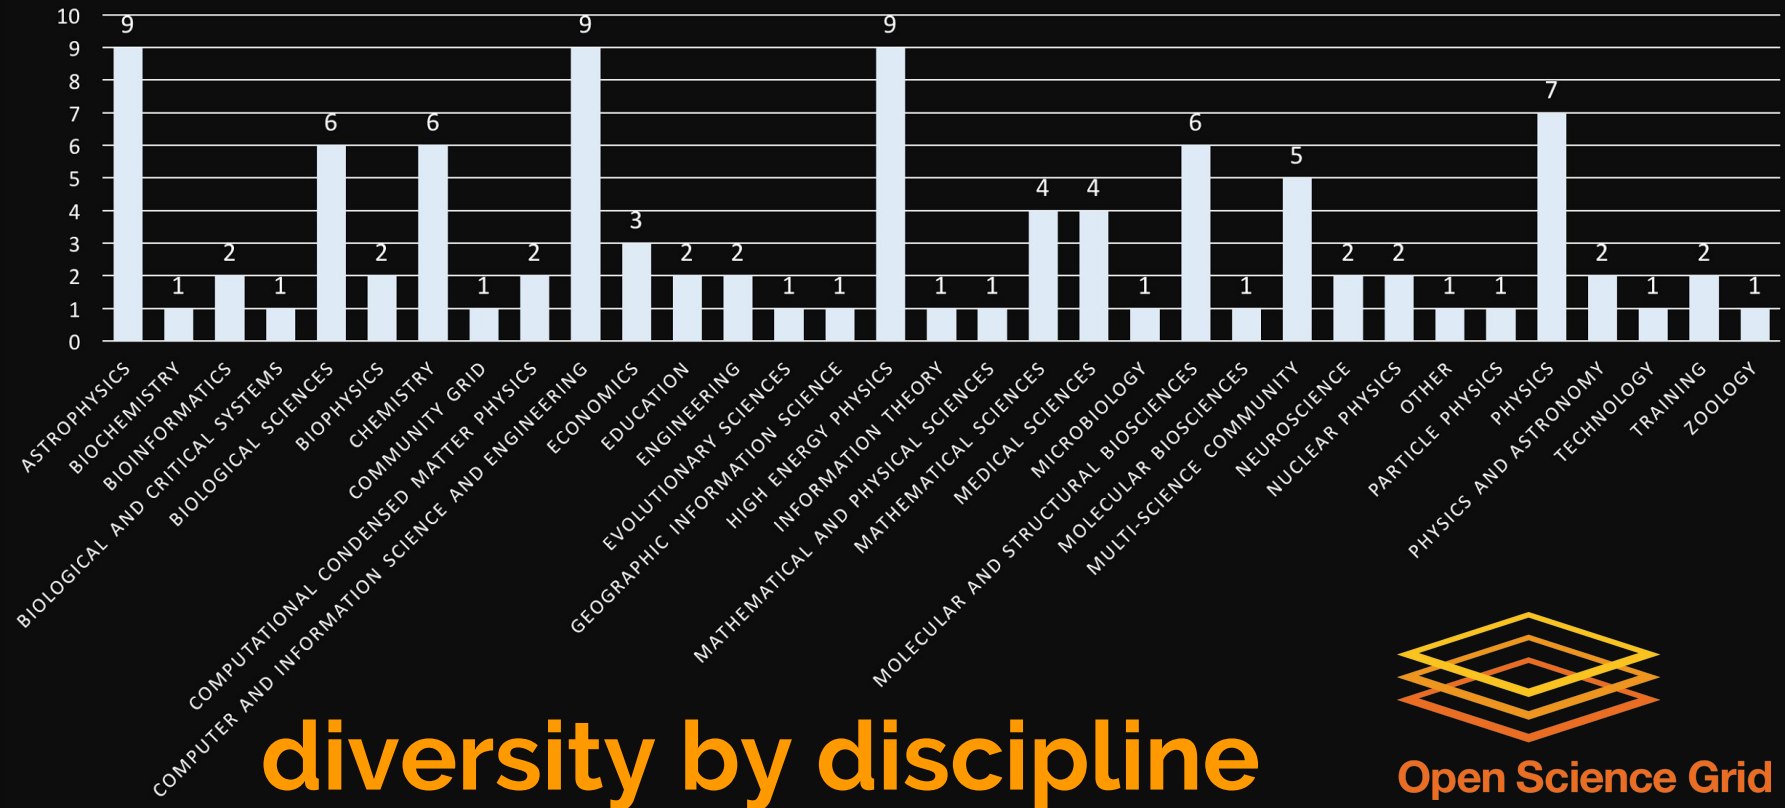
by wall
hour
usage

FIELDS OF SCIENCE 2016 **OSG CONNECT, XD, DIRECT**




PROJECTS

OSG CONNECT, XD, DIRECT 2016




Open Science Grid

INSTITUTIONS 2016
OSG CONNECT, XD, DIRECT

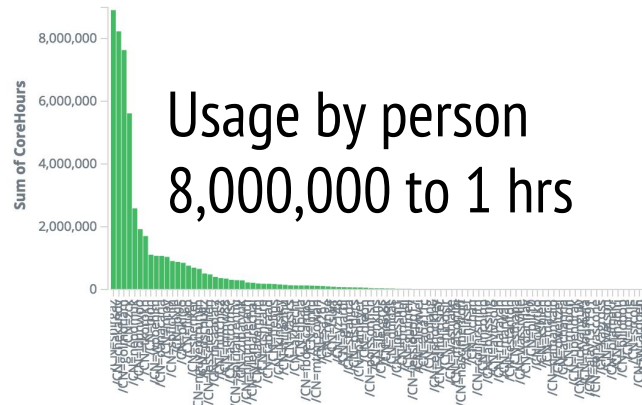


diversity by institution



- z2dqmc
- FutureColliders
- SourceCoding
- seq2fun
- BioGraph
- VERITAS
- AlGDock
- PRTH
- PainDrugs
- EvoSims
- PreBioEvo
- ConnectTrain
- numfpi
- DemandSC
- EvoProtDrug
- IceCube
- EIC
- CombinedPS
- nicesims
- molcryst
- Perchlorate
- NSNM
- SysBioEdu
- microphases
- atlas-org-uchicago
- fluidsim
- LiuLab
- uchicago
- RicePhenomics
- SNOplus
- duke-boolnet
- OSG-Staff
- freesurfer
- GeoTunnel
- holosim
- UserSchool2015
- SBGrid
- NSLS2ID
- ContinuousIntegration
- EvolvingAI
- TDASTats
- N/A
- MS-EinDRC
- SouthPoleTelescope
- mab
- duke-CMT
- xenon1t
- Unknown
- Paniceae-trans
- PathSpaceHMC

osg-connect-users-bar

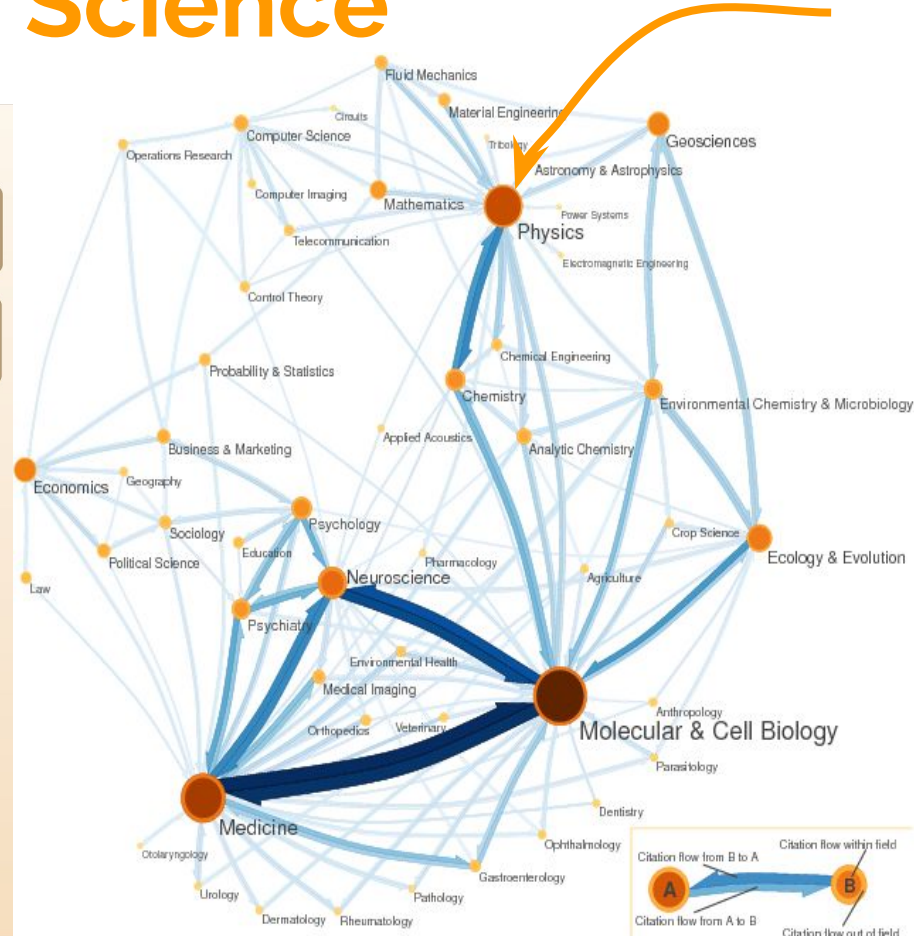



diversity by job scale

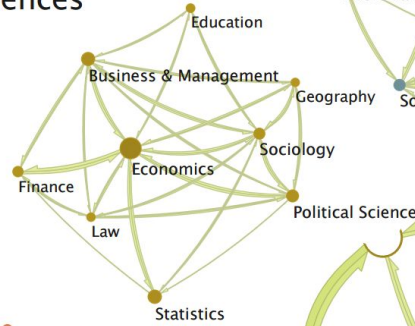
Science sampler

With apologies for the many projects we've not included. Selection here was almost random, weighted heavily by materials supplied by speakers here.

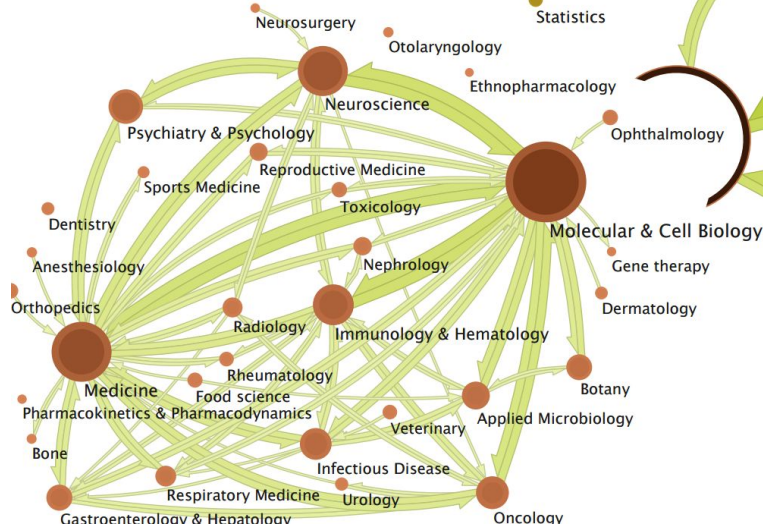
Science of Science



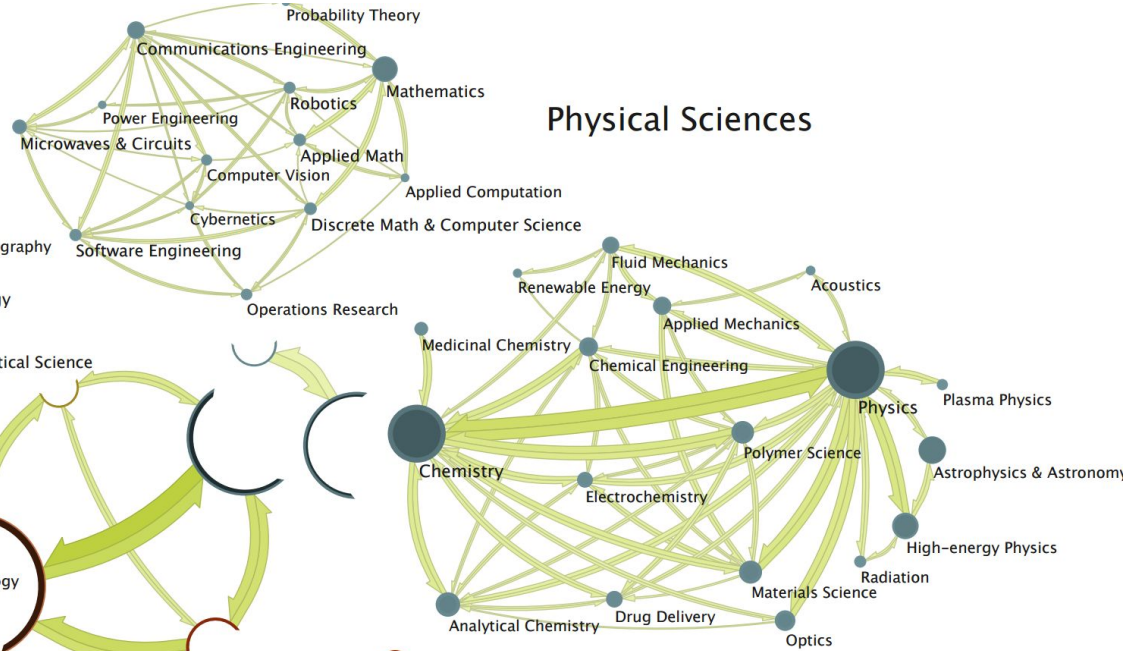
Social Sciences




Life Sciences



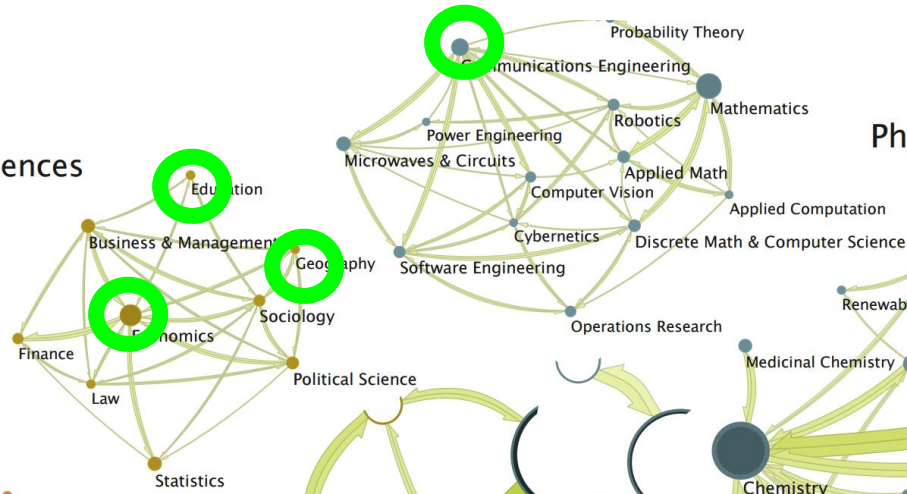
Physical Sciences



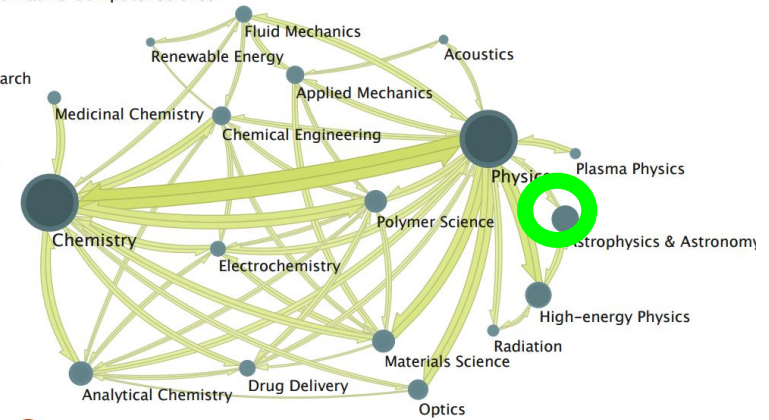
Ecology & Earth Sciences




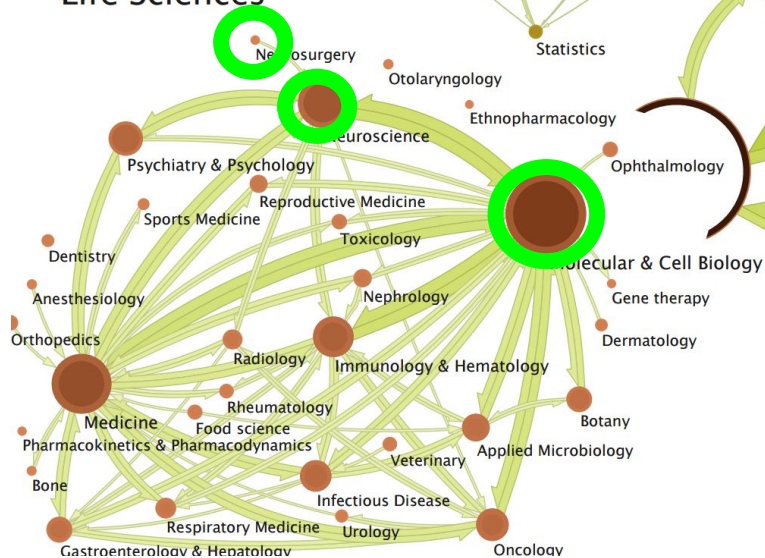
Social Sciences



Physical Sciences



Life Sciences



Ecology & Earth Sciences

- FASTQ files are mapped to a reference genome and converted to a BAM alignment file.
- BAM files can be mined for gene expression vectors that can be bundled into a gene expression matrix (GEM).
- GEMs are a stable data structure that can be mined for differentially expressed genes (DEGs) or used to construct Gene Co-expression Networks (GCNs)

Raw DNA Sequence → Gene Interaction Graph

```
@SRR001666.1 071112_SLXA-EAS1_s_7:5:1:817:345 length=36
GGGTGATGGCCGCTGCCGATGGCGTCAAATCCCACC
+SRR001666.1 071112_SLXA-EAS1_s_7:5:1:817:345 length=36
IIIIIIIIIIIIIIIIIIIIIIIIIIIIII9IG9IC
```



OSG-GEM

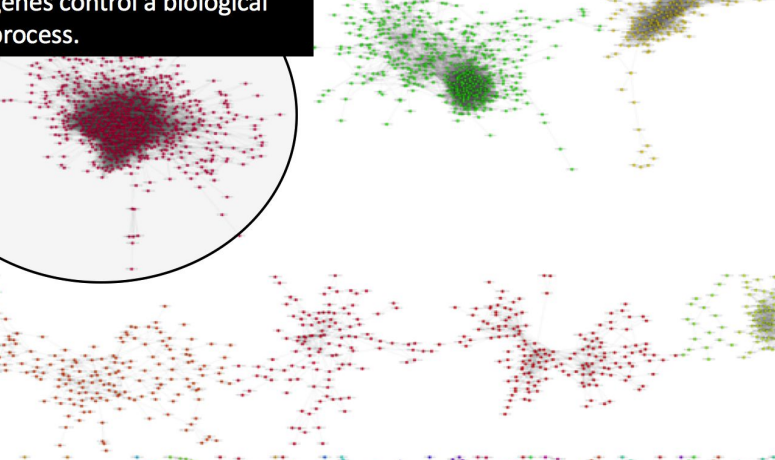
GeneID	ERR12353	ERR12355	ERR12362	ERR12366
GRMZM2G092856.T01.v6a	9.40	20.25	4.65	0.00
GRMZM2G0811275.T01.v6a	0.00	0.38	0.00	0.45
GRMZM2G0811273.T01.v6a	5.84	0.00	0.00	9.06
GRMZM2G099778.T01.v6a	0.45	1.65	0.00	0.88
GRMZM2G094757.T01.v6a	1.31	1.38	1.57	0.00
GRMZM2G093999.T01.v6a	9.00	1.08	1.80	0.65
GRMZM2G080743.T01.v6a	5.21	1.66	21.96	1.19
GRMZM2G083153.T01.v6a	2.98	30.50	19.33	15.75
AC177838.2.FG0012.v6a	0.29	0.24	0.00	0.00

OSG-KINC



Gene Expression Matrix HPC Challenges


BIOLOGY: These connected genes control a biological process.

A complex network graph visualization. The nodes are small circles, and the edges are thin lines connecting them. The nodes are colored in several distinct groups: a large red cluster in the upper left, a green cluster in the upper right, a blue cluster in the lower left, a yellow cluster in the lower right, and a purple cluster in the bottom center. A circular inset in the top left corner provides a magnified view of the red cluster, showing a dense, interconnected web of nodes. The overall structure is highly interconnected, with many nodes having multiple connections.

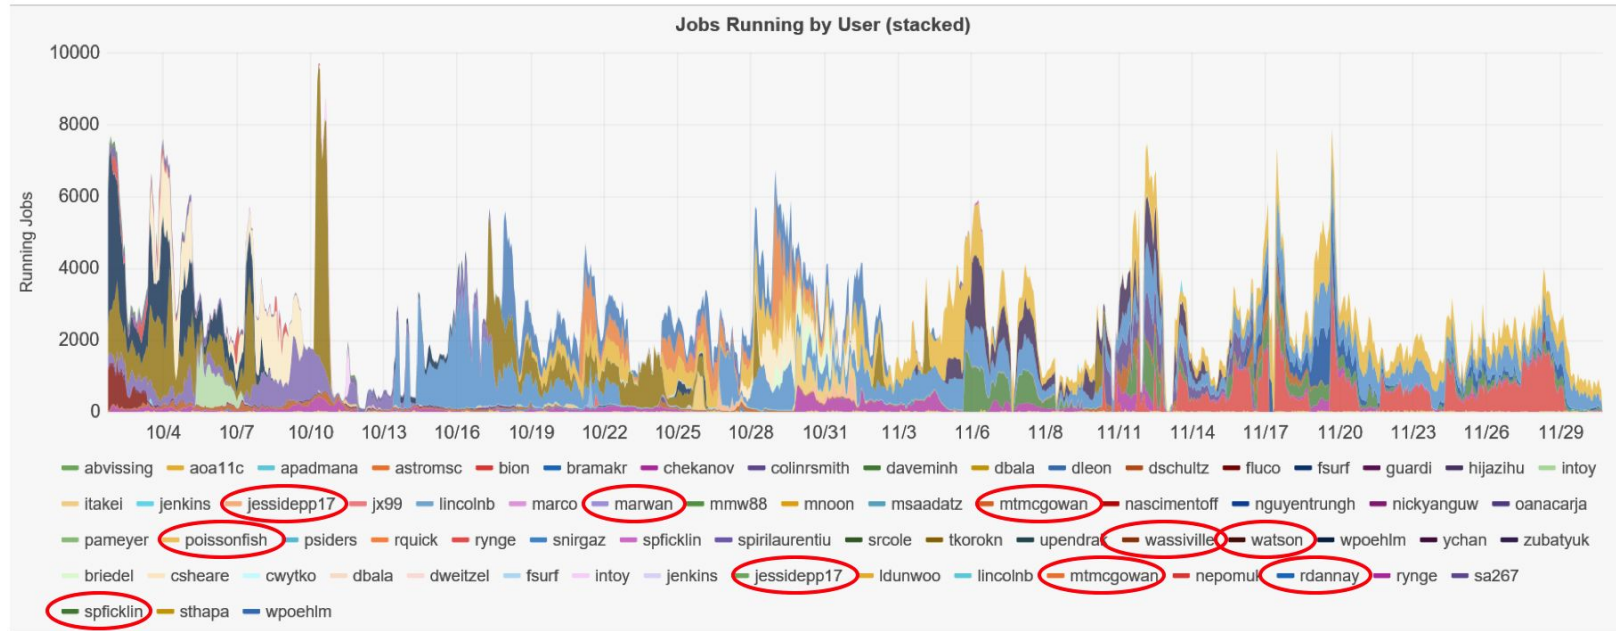
- **Memory Intensive**
 - Some Steps Require >16 GB RAM/dataset
- **Storage Intensive**
 - ~37 GB of intermediate files/dataset
- **Dataset Size Varies Widely**
 - 10M to 20M sequences/dataset
- **Data Transfer Issues (Tera-/Petabytes)**

```
@SRR001666.1 071112_SLXA-EAS1_s 7:5:1:817:345 length=36  
GGGTGATGGCCGCTGCCGATGGCGTCAAATCCCACC  
+SRR001666.1 071112_SLXA-EAS1_s 7:5:1:817:345 length=36  
IIIIIIIIIIIIIIIIIIIIIIIIIIIIII9IG9IC
```

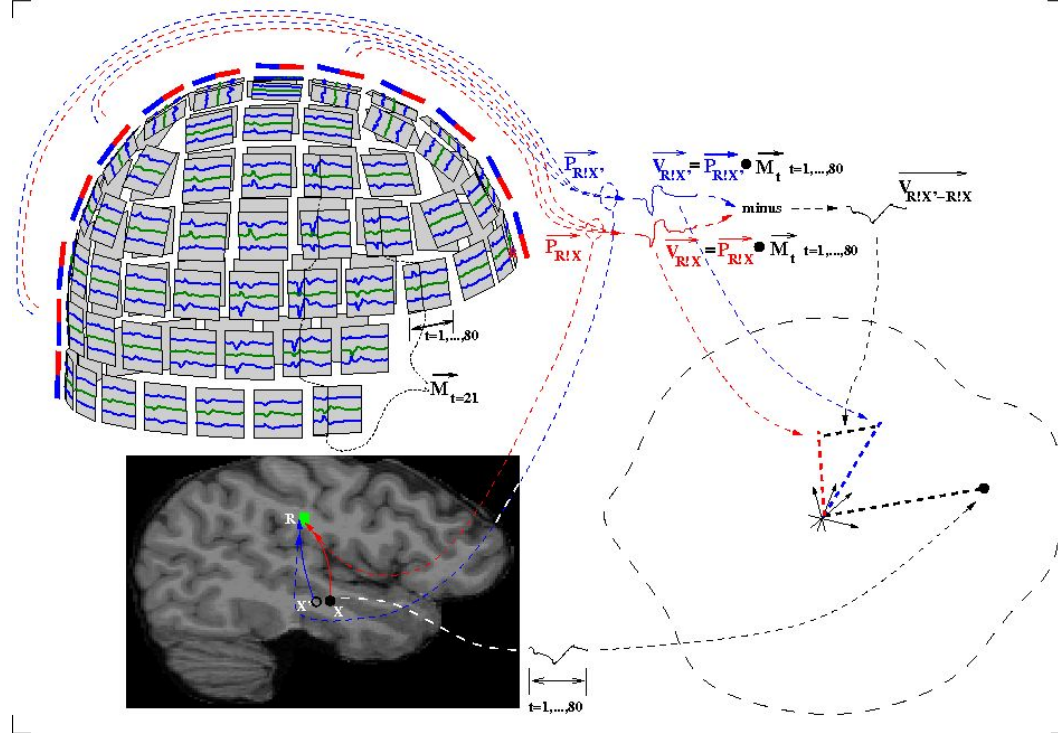

- Introduction to Systems Biology
- WSU HORT 503 Special Topics (3 credits)
- Offered Fall 2017
- 9 Students
 - Crop & Soil Sciences
 - Molecular Plant Sciences
 - Entomology
 - Biological Systems Engineering
- Why OSG?
 - Expose students to national computing resources
 - Anticipate large-scale genomics of the future.



- Success: SysBioEdu Usage from Oct-Dec 2016
 - Students launched jobs and they ran!



Functional Neuroimaging



Functional Neuroimaging..



- Don Krieger has been working with TEAM TBI at the University of Pittsburgh
 - **T**argeted **E**valuation, **A**ction and **M**onitoring of **T**ramatic **B**rain **I**njury
- TEAM TBI investigates the complexity of brain injury, and how targeted interventional strategies may improve outcome and function.



TeamTBI Functional Neuroimaging

David Okonkwo, Jim Becker, Sue Beers, Mickey Collins, Anthony Kontos, Malcolm McNeil, Walt Schneider



The Department of Neurological Surgery
at the University of Pittsburgh



TeamTBI Functional Neuroimaging

David Okonkwo, Jim Becker, Sue Beers, Mickey Collins, Anthony Kontos, Malcolm McNeil, Walt Schneider

Supported by: Department of Defense, Open Science Grid (NSF,DOE), Extreme Computing Environment for Science and Engineering (NSF).



TeamTBI Functional Neuroimaging

We record the magnetoencephalogram (MEG) from an array of 306 magnetometers sampled at 1 KHz typically for 30 minutes.



The Department of Neurological Surgery
at the University of Pittsburgh



TeamTBI Functional Neuroimaging

We record the magnetoencephalogram (MEG) from an array of 306 magnetometers sampled at 1 KHz typically for 30 minutes.

We utilize a unique deconvolution solver to identify and validate ($p < 10^{-12}$) each neuroelectric source within the brain with mm resolution.



TeamTBI Functional Neuroimaging

We record the magnetoencephalogram (MEG) from an array of 306 magnetometers sampled at 1 KHz typically for 30 minutes.

We utilize a unique deconvolution solver to identify and validate ($p < 10^{-12}$) each neuroelectric source within the brain with mm resolution.

The solver requires about 40 core hours per second of data. It typically yields about 500 simultaneously active sources for each 40 msec step through the data stream, i.e. 20,000,000+ sources from a single 30 minute recording session.



TeamTBI Functional Neuroimaging

We record the magnetoencephalogram (MEG) from an array of 306 magnetometers sampled at 1 KHz typically for 30 minutes.

We utilize a unique deconvolution solver to identify and validate ($p < 10^{-12}$) each neuroelectric source within the brain with mm resolution.

The quantity and resolution of these measurements is unprecedented.



TeamTBI Functional Neuroimaging

In December, we gained access to the CamCAN lifespan normative cohort in collaboration with the University of Cambridge (UK). These recordings were obtained from 628 volunteers ages 18-87.



TeamTBI Functional Neuroimaging

In December, we gained access to the CamCAN lifespan normative cohort in collaboration with the University of Cambridge (UK). These recordings were obtained from 628 volunteers ages 18-87.

Processing has been completed on 100 resting records using OSG including reserved glideins running on SDSC's Comet and PSC's Bridges.



TeamTBI Functional Neuroimaging

Consistent patterns of activation across 155 brain regions were identified from the CamCAN recordings with eigenvector analysis.

The Department of Neurological Surgery
at the University of Pittsburgh



TeamTBI Functional Neuroimaging

Consistent patterns of activation across 155 brain regions were identified from the CamCAN recordings with eigenvector analysis.

The brain patterns found in MEG recordings from volunteers with history of head injury are sufficiently different to classify them with 95% accuracy.


The Department of Neurological Surgery
at the University of Pittsburgh



TeamTBI Functional Neuroimaging

Consistent patterns of activation across 155 brain regions were identified from the CamCAN recordings with eigenvector analysis.

The brain patterns found in MEG recordings from volunteers with history of head injury are sufficiently different to classify them with 95% accuracy.




Department of Neurological Surgery
at the University of Pittsburgh



TeamTBI Functional Neuroimaging

Consistent patterns of activation across 155 brain regions were identified from the CamCAN recordings with eigenvector analysis.

The brain patterns found in MEG recordings from volunteers with history of head injury are sufficiently different to classify them with 95% accuracy.



One of these two brain patterns used to classify the cohorts is specific to depression, the other to somatization. These are common symptoms of PTSD.




Large Scale Metagenomics.

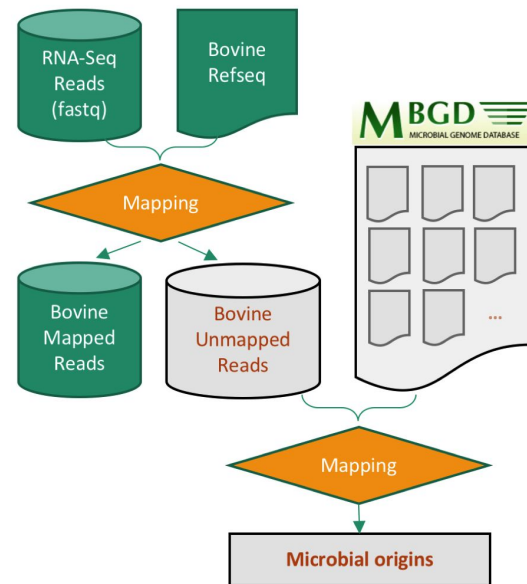


Motivation

- In a previous project, we have isolated exosomes from the commercial cow's milk and assessed the bovine molecules inside the exosomes.
- Moreover, we also found many unmapped reads are from microbial species.
- Thus, we designed a follow-up study to **understand the origin of microbial sequences in bovine milk exosomes.**
 - Metagenomics analysis



Pseudomonas cedrina



Large Scale Metagenomics..

Computational challenges

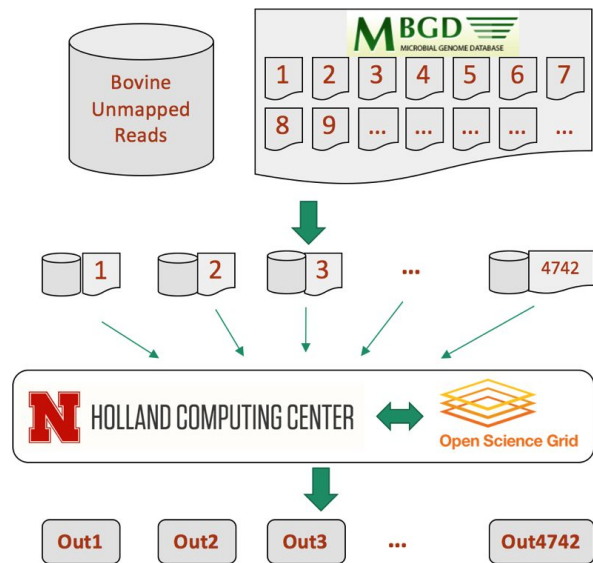
- A large number of target genomes
 - 4,742 genomes (size: 100KB ~ 58MB)
- Six samples contain over 100 million of bovine unmapped reads
 - Increasing the computing time
- In total, $6 \times 4,742 = 28,452$ mapping tasks

Question: Where to execute this many of jobs?

- Impossible for the lab-server (32 cores)
- Long pending time if submitted it to HCC clusters
 - Dynamic priority scheduling of users/groups
 - More jobs completed -> longer queue time

Perfect Fit of Open Science Grid (OSG)

- The tasks are independent to each other
- Limited file transfer
 - Total size of transferred files ~1GB
- Small memory consumptions
 - Memory < 2GB
- Short running time for each task
 - Maximum: 3 hours (HCC@UNL-Crane)
- Software is available on OSG
 - Pre-installed Bowtie and Tophat
 - No further configuration needed





Results

- Several microbial species were identified in bovine milk exosomes:
 - *Pseudomonas fluorescens*, *Pseudomonas chlororaphis*, *Pseudomonas poae*, *Enterobacter cloacae*, etc.
- Although some species have been reported in cow's milk before, this is the first time of identifying microbial sequences in milk exosomes
 - Potential transportability to other species through exosomes
- Based on the findings from this analysis, we have designed two experiments to further our understanding in this subject

Counterfactual Analysis.



- Economic analysis & public policy
- Considering "what if" scenarios in microeconomics
- Simulate firm/consumer behaviors

Using OSG to Evaluate Policy

We need to solve

$$\frac{\partial \Pi_{jt}}{\partial p_{jt}} + \beta \left[\frac{\partial \mathbf{s}_t}{\partial p_{jt}} \right]' \mathbb{E}_t \left[\frac{\partial V_j(\mathbf{s}_t, \mathbf{X}_t)}{\partial \mathbf{s}_t} \right] = 0,$$

for every firm and combination of \mathbf{s}_t and \mathbf{X}_t .

Need to compute $\mathbb{E}_t \left[\frac{\partial V_j(\mathbf{s}_t, \mathbf{X}_t)}{\partial \mathbf{s}_t} \right]$. How? Forward simulation.

OSG is what makes this possible for large state spaces.

Counterfactual Analysis..



Using OSG to Evaluate Policy

Outcome: making people more active decreases prices and it does not affect returns.

Case	Mean and 95% CI
Base simulation	6.195% [6.181%,6.210%]
No enrollment cost	3.666% [3.660%,3.671%]
No decision cost	3.837% [3.833%,3.842%]
No switching costs	2.607%

Note: The table reports the mean expected fees and 95% confidence intervals for the different scenarios under study using 10,000 random initial states.

Conclusions

1. Policy/transaction evaluation is critical in modern microeconomics.
2. Evaluations often require simulating consumer/firm behavior over a large state space.
3. In particular in the case of dynamic games, OSG could become a crucial tool, making the difference between being able to do something and not.

Simulating Source Coding.



- Data deluge - much of it mobile traffic
- Optical data compression
- Important for digital space and satellite communication & wireless data transmission

Monte Carlo Simulation for Next Generation Source and Channel Coding on OSG Connect

Ahmad Golmohammadi*, David G. M. Mitchell*,
Joerg Kliewer[†], and Daniel J. Costello, Jr.[‡]

* Klipsch School of ECE, New Mexico State University

[†] Dept. of Electrical and Computer Engineering, New Jersey Institute of Technology

[‡] Dept. of Electrical Engineering, University of Notre Dame




March 7, 2017

Simulating Source Coding..

- Whole system simulations: transmitter, decoder, receiver & stochastic noise
- Data compression & reconstruction algorithms

Data Compression

- In order to handle the vast amounts of data society will produce, we need **efficient and low-complexity** algorithms to reliably compress and reconstruct data.



- We want to **reconstruct** the sequence with as little **distortion** as possible with a practical scheme (lossy source coding)
- The ultimate limit of compression is known (Shannon 1958)


Simulating Source Coding...



- Sparse graphs can approach fundamental limits
- To verify the results, large Monte Carlo samples needed - "not possible without the OSG"

LDGM Codes for Lossy Source Coding

- We construct **codes** for lossy source coding using a small structured graph (protograph)
- A large graph can be obtained from a **protograph** by **graph lifting** with **lifting factor M**
 - The graph is copied M times and the edges randomly permuted following the graph structure
- Low complexity algorithms based on **belief propagation** can be defined for the sparse graph
 - Here, **messages** are passed iteratively forward and backward in the graph until we converge to a codeword \underline{z}



(3, 6)- regular



- Understanding evolution at molecular scale in DNA with combination of mathematical modeling and simulation
- How quickly does a genome fix a mutation?
- Role of randomness versus natural selection?



Environmental variation is commonplace yet unpredictable across biological systems from the adaptive immune system, the microenvironment in cancerous neoplasms, to populations of pathogens under drug pressure.

How do populations survive environmental stochasticity? How do they manage to persist and keep one's footing on an ever-changing landscape?

Can organisms prepare for this environmental stochasticity?

Can evolution prepare populations for this environmental stochasticity?

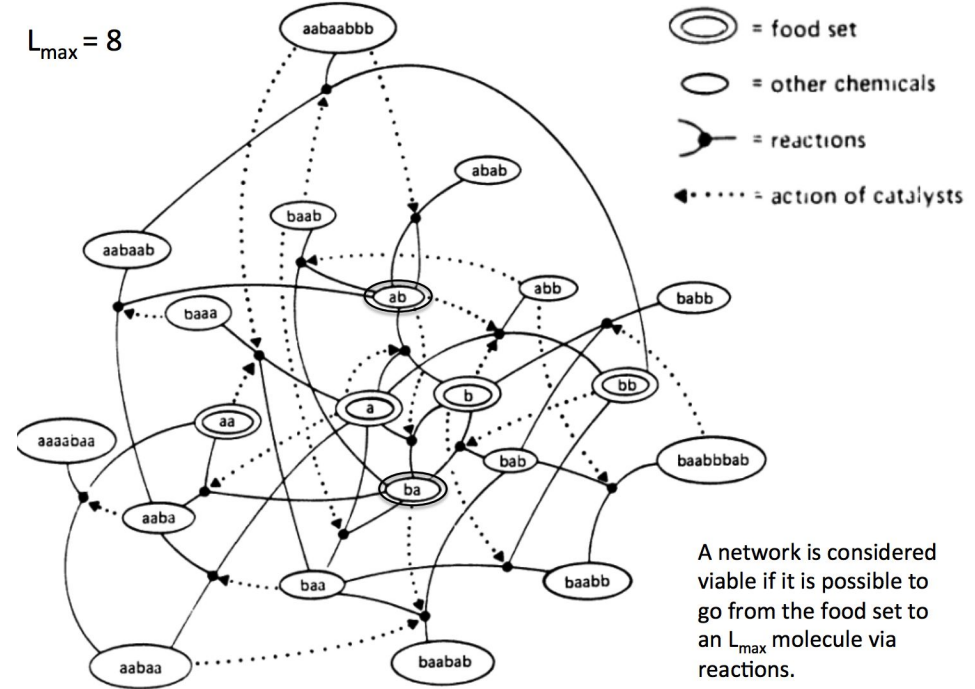
Evolving Strategies for Life..



Models of Prebiotic Evolution



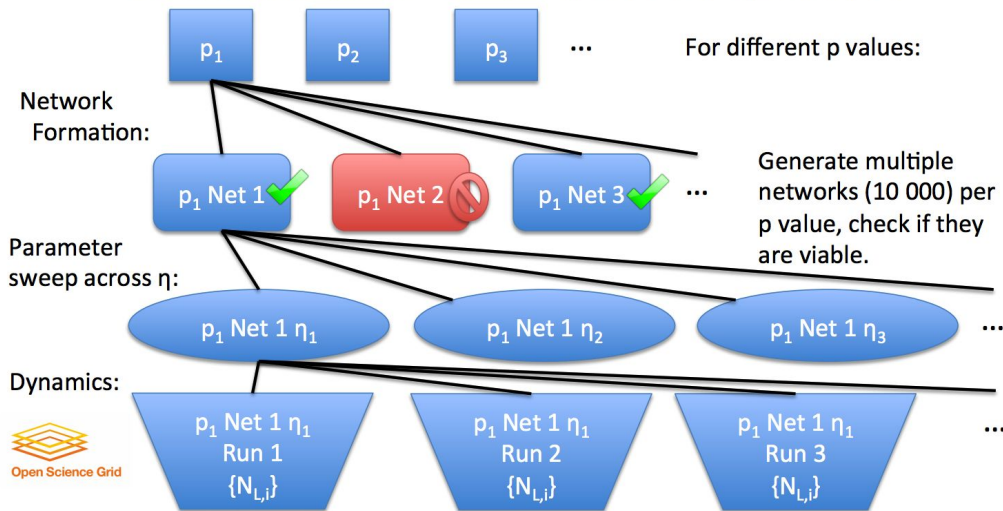
- Protein first origin of life model
- Network of interacting molecules assumed to be polymers
- Perhaps solve Eigen's paradox (low probability of randomly constructing "starter gene")



"The Origins of Order" – Stuart A. Kauffman

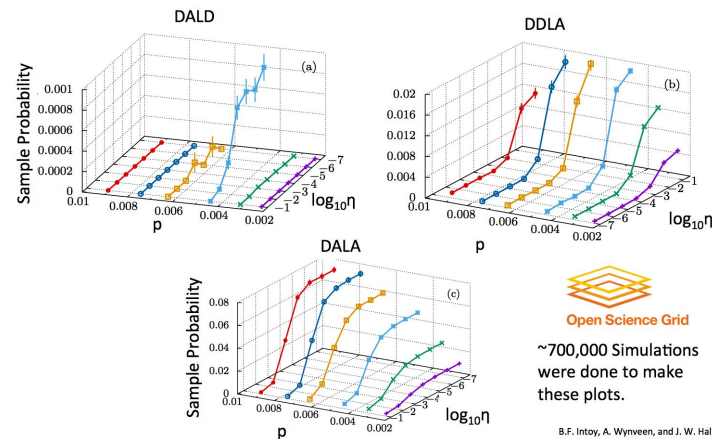


Simulation General Structure



- Do multiple dynamic simulations with random initial conditions using a given viable network generated by parameter p combined with reaction rates and diffusive value η .
- A steady state is then reached with polymer length and spatial distribution $\{N_{L,i}\}$.
- Analyze the $\{N_{L,i}\}$'s to determine whether the run was lifelike or not.

Probabilities of DALD, DDLA, DALA states as a function of p and η





B.F. Intoy, A. Wynnveen, and J. W. Halley
Physical Review E 94, 042424 (2016)

Protein Evolution






Understand the fundamental physical bottlenecks and dynamical behavior of protein evolution. Important questions include the extent of dominant pathways (convergent evolution) and phase transitions in evolutionary rates (punctuated equilibrium). These principals and their structural underpinnings can also be used to inform rational design of antibiotics that exploit bottlenecks in pathogen mutational response.



Analysis of Brain Rhythms

Large-scale brain recordings



- Healthy humans
 - Patients
 - Animals
-
- Sampling: 500-30000 Hz
 - Duration: 1 hour - 1 week
 - Channels: 1 - 250+
 - Several GB per subject



Analysis of Brain Rhythms..



Fourier Transform-based analysis



Free supercomputing for research: A tutorial on using Python on the Open Science Grid

Jan 3, 2017

Supercomputing resources typically cost money, but the [Open Science Grid \(OSG\)](#) provides high-throughput computing to any researcher in the US **for free**. Briefly, OSG users can run jobs on servers owned by dozens of academic institutions, whenever those servers are not actively running a job for its owners.

Besides cost, a second major barrier to entry for those who are new to supercomputing (or, specifically, Condor) is the necessary troubleshooting before we can actually run our jobs. The purpose of this tutorial is to provide a complete example for running Python jobs on the OSG. This example is nontrivial, in that it includes multiple data sets, public libraries (e.g. `scipy`), private libraries, and analyzing output. In complement to this tutorial, the OSG has [tutorials](#), a [structured class](#), and extremely helpful online support when you get stuck.


This tutorial goes through the steps of manually connecting to and running commands on the remote server, but see the **Fabfile** section at the bottom for how this can be automated on your local machine.



A FreeSurfer Workflow Service



- Widely used software suite for analysis of human brain MRI scans.
- Neurophysiology of depression, examining possible anatomical differences involved in ADHD, and studying autism



Computational model of the cortical and subcortical brain structures that form the basis of the BrainPrint, a system for representing the whole brain based on the shape, rather than the size of structures. (Martin Reuter, PhD, and Christian Wachinger, PhD, Martinos Center for Biomedical Imaging)

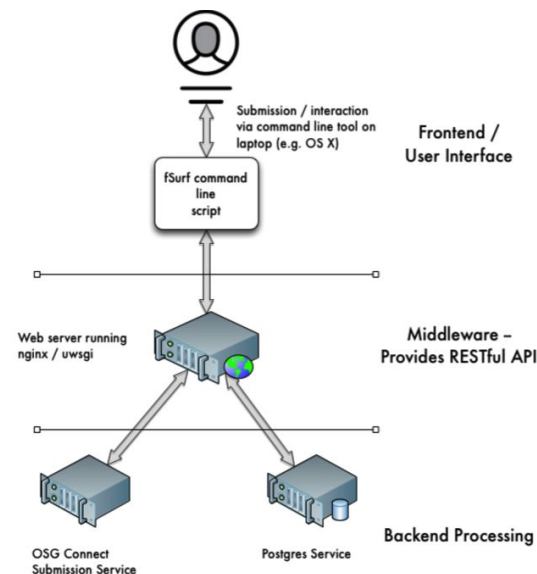
A FreeSurfer Workflow Service

Neuroscience



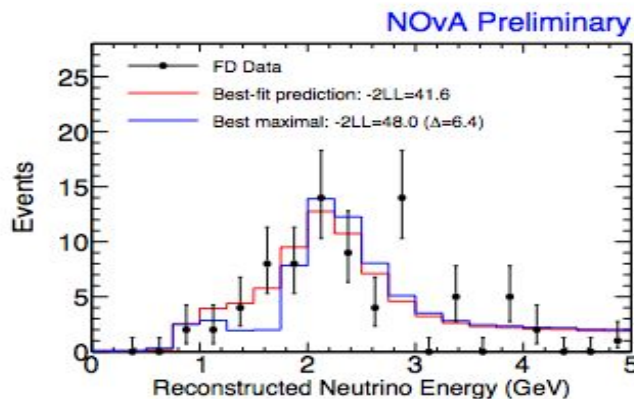
- Working with Don Krieger (Pittsburgh) to develop an OSG-based execution service
- Uses Pegasus
- Handles "standard" transforms and user options
- To be released this week!

- Front-end user interface
 - Command line script
 - Communicates to the execution service via RESTful-API commands
 - Allows user to submit, view, and remove workflows as well as download results
 - Middleware
 - NGinx / uwsgi server provides RESTful API
- Backend services (running on a single VM)
 - Postgres database
 - Tracks workflow status, locations of inputs and results
 - Stores user information
 - Pegasus submission service (using same infrastructure as OSG Connect) to



V0s





VO Highlights: From the smallest scales...



NOvA: Fermilab-based neutrino experiment



**MINOS+:
limits on
LEDs**





Mu2e: Lepton-flavor violation experiment

Nearly 60M opportunistic hours on OSG and counting

>500,000 in one day!

VO Highlights: From the smallest scales...




Au+Au event

GlueX: probing exotic mesons predicted by LQCD


STAR: Heavy Ion Physics

...to the largest...



Dark Energy Survey:


Discovery of dwarf planet-
Second-most distant known object in solar system
Techniques applied to ongoing Planet 9 search



Ice Cube:


Neutrino Observatory, also sensitive to
extremely high energy cosmic rays

...the completed to the still in planning...



Infrared Processing and Analysis Center:

NASA's archive for a host of IR/sub-mm astronomy missions, galaxy catalogs, Keck Observatory, and more!




[Tarun Souradeep (2016), Beyond Gravitational wave discoveries with LIGO-India]

LIGO India:

Additional detector will greatly
Improve localization of gravitational
wave sources

...and working in all corners of the globe.



VERITAS:

4 12m Cerenkov telescopes for
gamma ray astronomy: Arizona, USA



- 1st ton-scale experiment
- 3.2t of LXe, 2t in TPC
- All systems commissioned since Fall 2016
- Calibration and science data taking now ongoing

XENON1T:

Dark matter detector at Gran
Sasso National Laboratory, Italy



South Pole Telescope:

Microwave-millimeter
telescope

And in space!



Alpha Magnetic Spectrometer (AMS) mounted at the ISS

Photo credit: NASA

OSG User Support

Team




Bala



Emelie



Mats



Suchandra



Rob



Ken

<http://support.opensciencegrid.org/>



Welcome
Login Sign up

Home

Solutions

How can we help you today?

Enter your search term here...

SEARCH

+ New support ticket

📄 Check ticket status


User Support

Training

Intro to HTC on OSG
Connect in
Synthesis Center
tomorrow



Thank you!



opensciencegrid



user-support@opensciencegrid.org



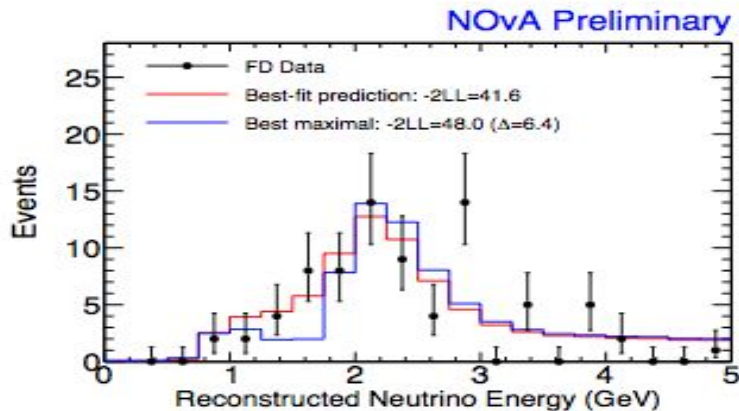
support.opensciencegrid.org




www.opensciencegrid.org/links

Extra slides

Selected results using the FIFE Tools

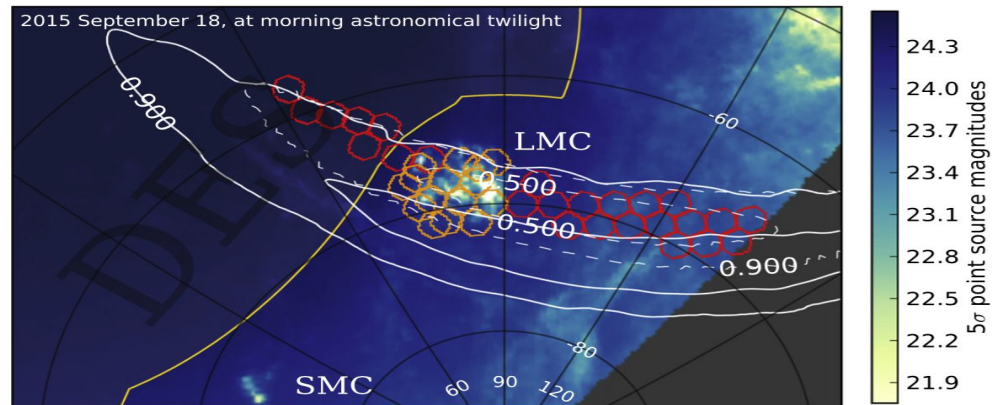


NOvA: excludes maximal mixing at 2.5 s.d.



Microboone:
first results

MINOS+:
limits on
LEDs



Dark Energy Survey: Optical follow-up
of gravitational wave triggers

