

OSG Campus Infrastructures Community Workshop

Doing science with minimal
campus infrastructure

-
Experiences at UCSD

i.e. UCSDGrid

by Igor Sfiligoi (UCSD)

Background

- UCSD is a CMS T2 site
 - Implemented as an OSG site
- **We provide many cycles to opportunistic users**
 - i.e. non-CMS users
- **But we essentially were giving none to users from UCSD**
 - So we asked ourselves?
Why not?

How it started

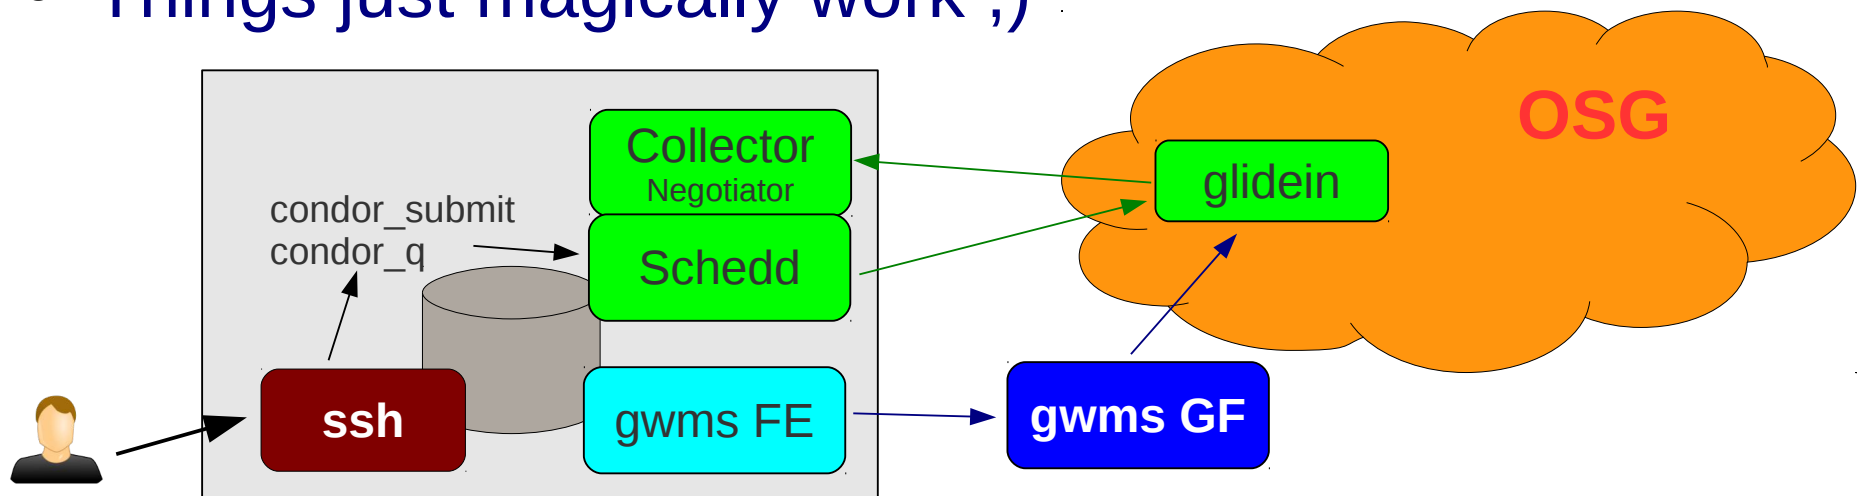
- So we started looking for local groups who would be interested in using OSG Late 2010
- We found the group from the Center for Theoretical Biological Physics (CTBP) who seemed interested
 - They were already running a small cluster
 - But they could definitely use more bursts of CPU
- Given our expertise, we went for a glidein solution
 - CTBP volunteered to provide the HW

Why glideins?

- **Grid infrastructure already in place**
 - And working fine for many VOs
 - So, not eager to re-invent the wheel
- Looking for a **cheap and easy solution**
 - Direct Grid submission not an option (of course)
 - **glideinWMS fit the bill perfectly**
- CMS uses it, too
 - So, plenty of experience

Our setup (at the time)

- Simple one-node glidein installation
 - Pilot proxy part of Engage VO
- Users log in using ssh
 - Told how to use Condor
 - Things just magically work ;)



A few more details

- Users told not to expect anything on the WN
 - **Have to bring with them everything they need**
- Users told to **use Condor file transfer**
 - No Grid data management tools
- **No Grid credentials** (i.e. x.509 proxies)
 - Cannot use glexec → all users share pilot UID
 - But user community small, so acceptable
 - No other use for proxies, so we went for simplicity

Glidein Configuration

- We have a pretty simple glideinWMS FE config
- Allow user choice in **two dimensions**:
 - UCSD only (default) vs “all of OSG”
 - UCSD only more uniform and faster networking (LAN)
 - 22h (default) vs 3h max job length
 - We noticed there were many 4h slots available to Engage
- Usual attribute convention
 - Requirements in the glideins only
 - User job requirements ~= True

<http://engage-fe.t2.ucsd.edu/vofrontend/policy.html>

Early user experience

- Got three major early adopters
- After a few false starts, they were happily churning their jobs
 - Up to 100k CPU hours in the first month!

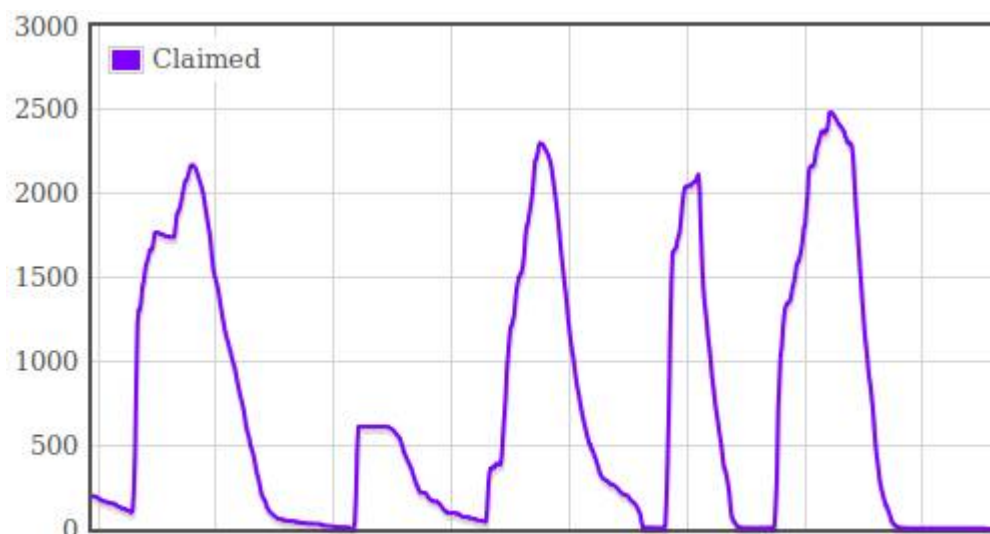
```
[frontend@osg ~]$ condor_userprio -all -allusers
```

User Name	Usage (hrs)	First	Last
-----	-----	-----	-----
jdurrant@osg.ctbp.ucsd.edu	28812.54	12/29/2010	1/10/2011
nschafer@osg.ctbp.ucsd.edu	2023.67	1/27/2011	2/16/2011
pocraig@osg.ctbp.ucsd.edu	101661.24	2/17/2011	3/10/2011
-----	-----	-----	-----
Number of users: ?	132502.88	12/16/2010	???

Spiky behavior

- While the integrated number of CPU hours was substantial, the demand was not uniform
 - The compute pattern was very spiky!

Perfect match for OSG!



- Iterative approach
 - User would get a batch of computation done
 - Then think over the results
 - Then submit a new bunch of jobs


CTBP contributes HW to T2

- CTBP was very pleased with the results
 - Science was being done way faster than before!
- So pleased, to contribute HW to the UCSD T2
 - Not far off what it was getting out of it on average
- And they were nice enough to not even ask for special privileges
 - Just trusting the system to deliver about the right amount on average

Since then

- Since then, we have grown the user base
 - Not dramatically, still in the O(10) range
 - But **expanded the sciences involved** (biology, astrophysics, climate, informatics, engineering)
- Usage patterns changed, too
 - Most users still have simple needs
 - But we have also demanding users
 - A high data user → gave him access to UCSD SRM
(1TB total, 100MB per job)
 - More complex workflows → taught a user to use dagman

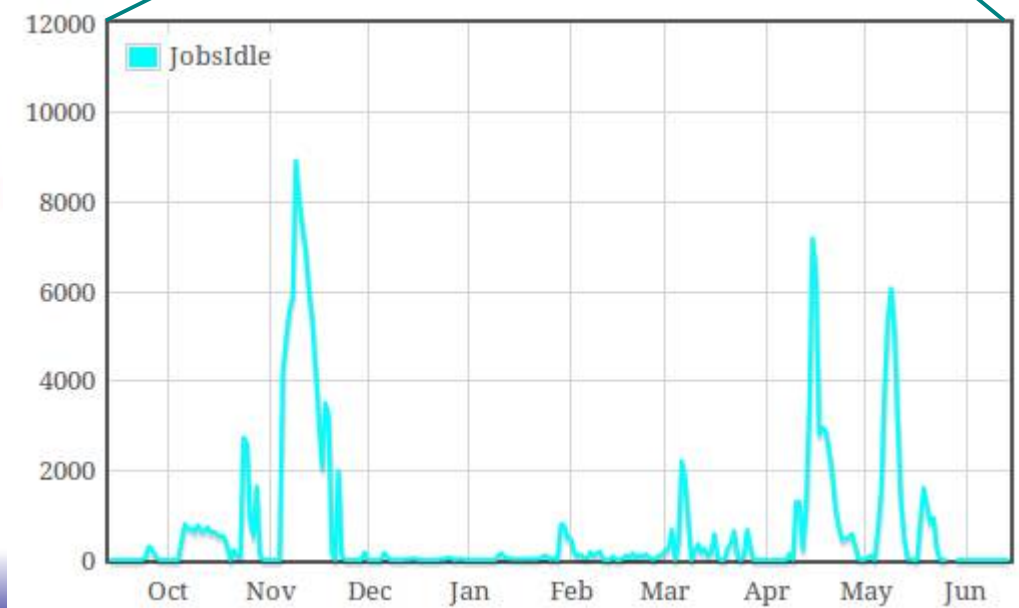
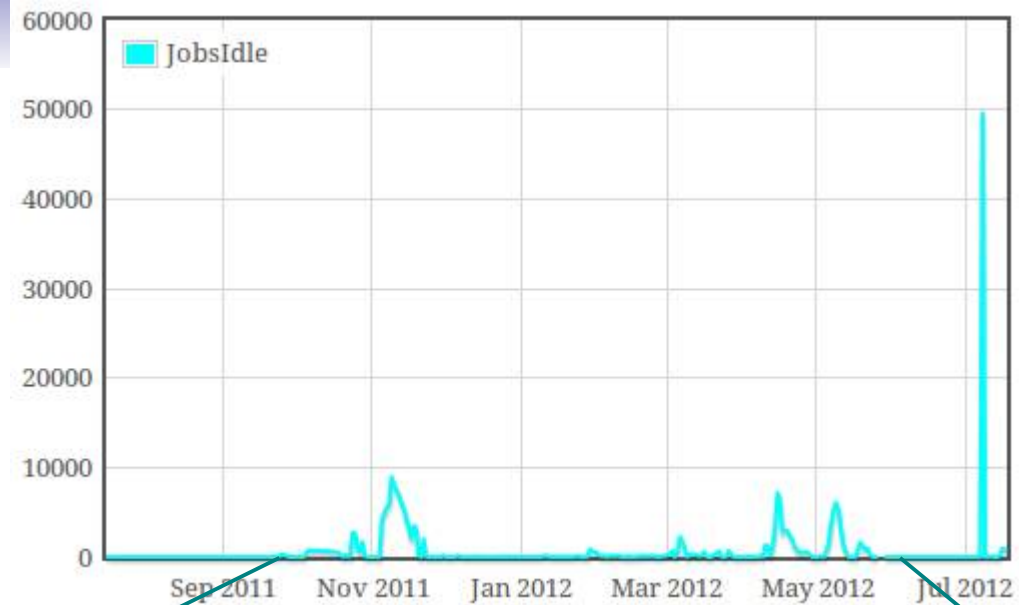
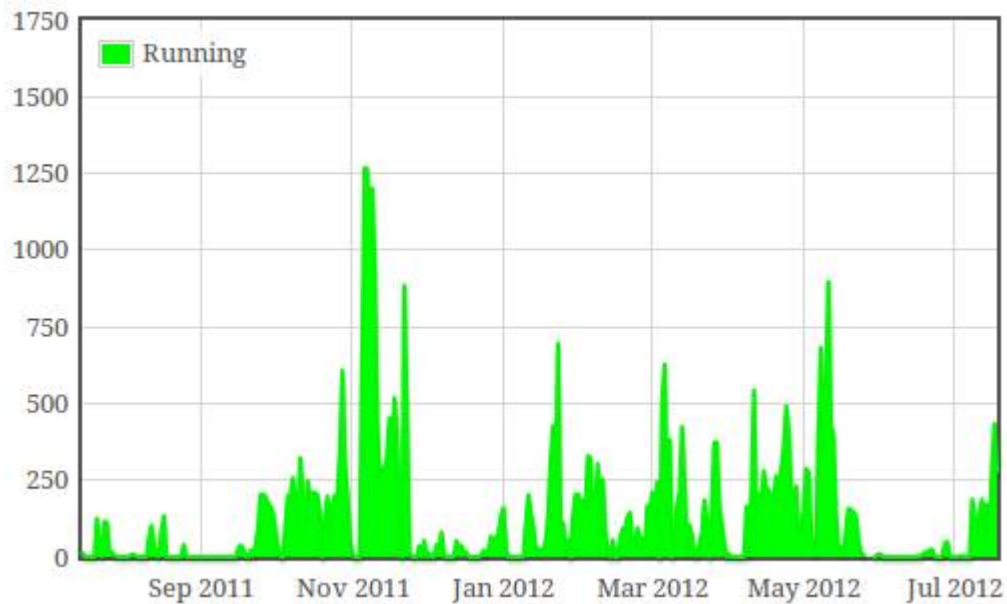
So he got a
Grid cert, too



Not all users interested

- Those were the success stories
- We also had a few users that came, looked around and then left
 - OSG environment perceived as too restrictive
- Major complaint was the lack of MPI
 - Single thread would take weeks at end
 - HTPC did not seem appealing
(and we had few slots at the time anyhow)

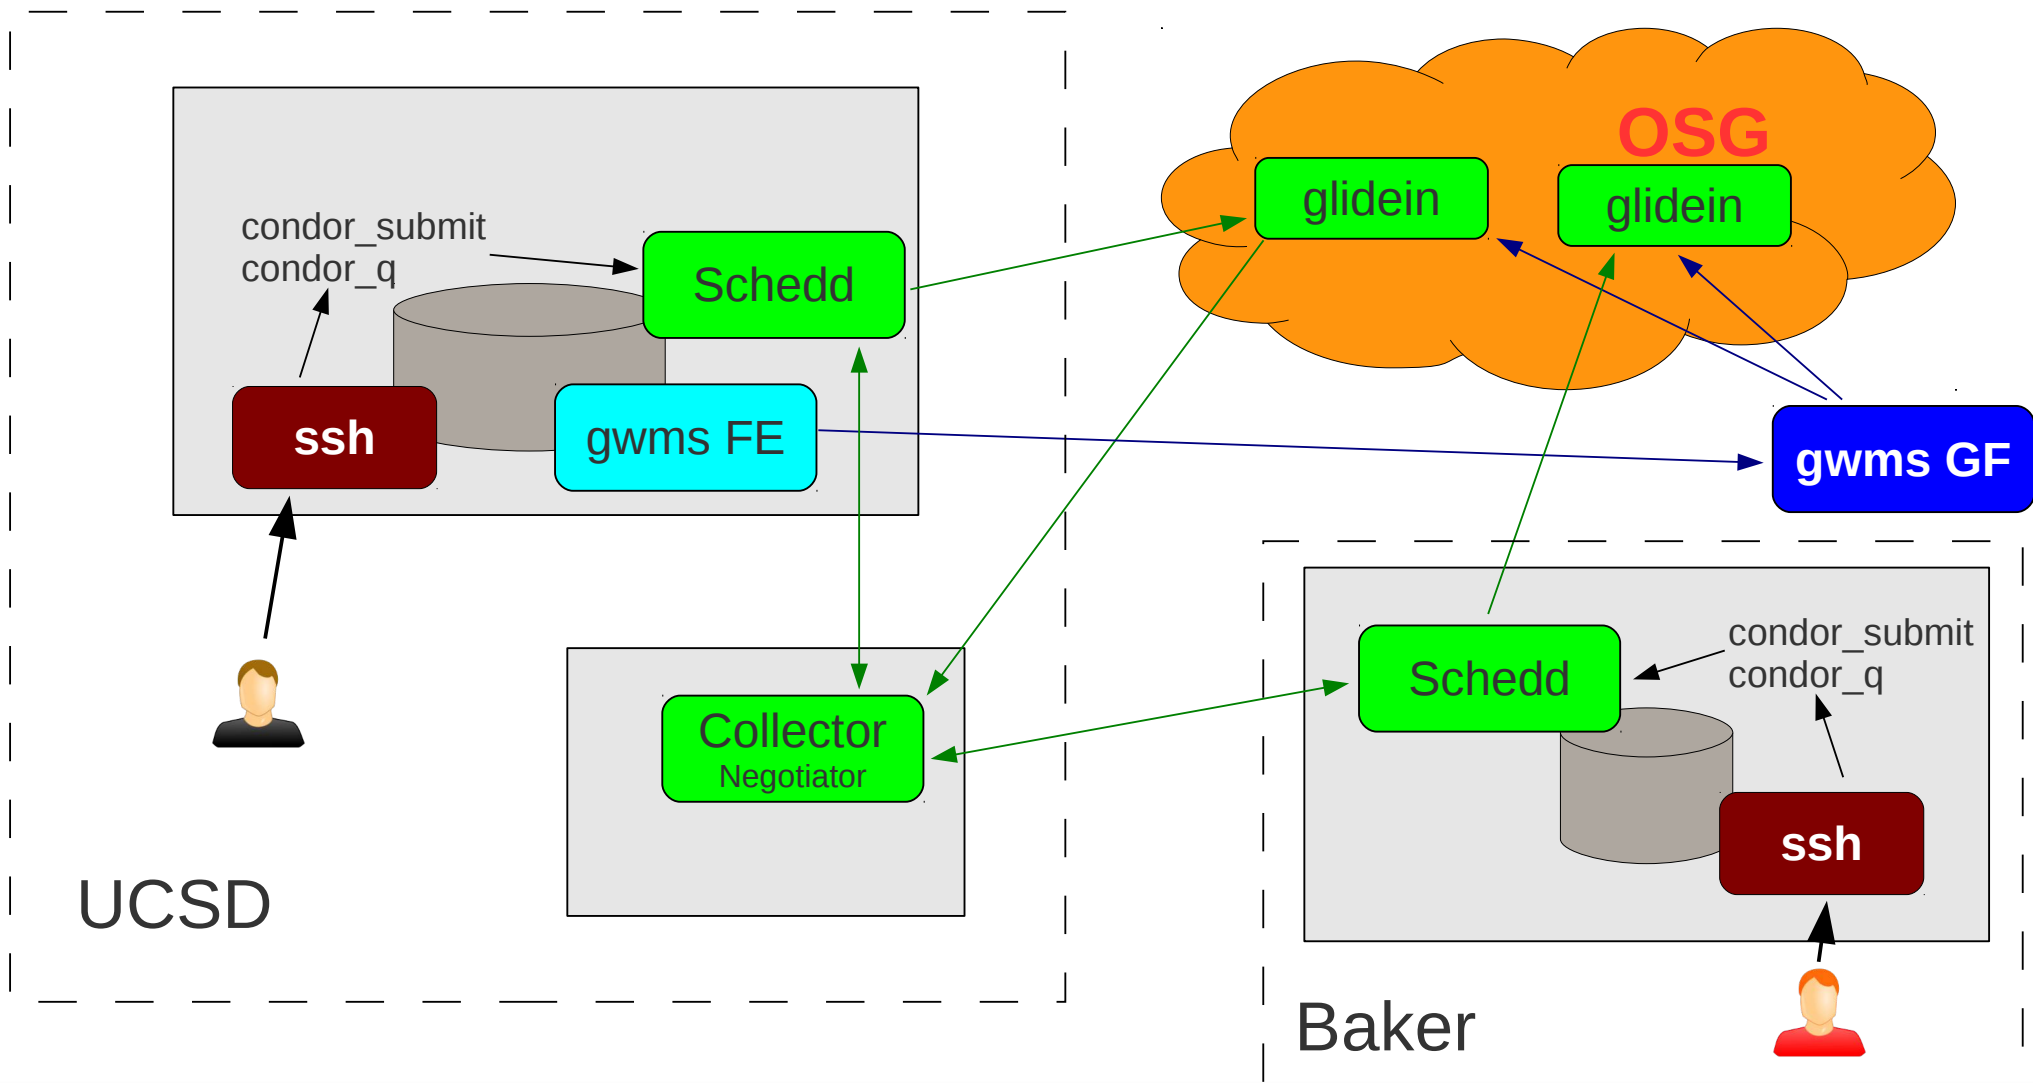
A few stats



Going beyond local users

- We recently became the glidein FE for Baker Lab, too
 - We provide Collector and FE
 - They provide they own schedd
 - We still have ours, of course
- We thus have two nodes at UCSD for this
 - The single node would not scale for combined UCSDGrid + BakerLab peak load

Current Setup



Not much changed

- At UCSD, not much changed
 - Apart from adding the new node
- FE config basically the same
 - We just added more time bins (with max being 44h)

Operational experience

- The **effort invested by our group very small**
- The system basically runs by itself
 - Occasional update of SW and the host proxy
- **Direct user support the most time consuming**
 - Simply because most of them never used more than a couple nodes at a time before
 - But it is just initial ramp up
 - After that, they are essentially self sufficient

Science showcase

- Follow a few slides from the major users
- Will not go into details
 - They are here just to give you an idea of the science being done

Biology - 1

The role of symmetry of tertiary structure in protein evolution

For each of about ~2000 protein domains exhibiting rotational symmetry in tertiary structure, we:

1. Determined the subunit comprising the symmetry
2. Identified other proteins, in different superfamilies, containing that subunit

What can the conservation of these subunits across distant proteins tell us about evolution?

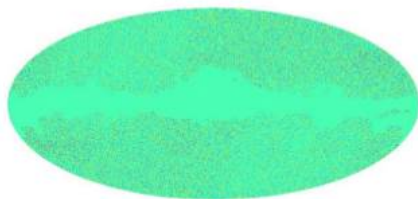


2gg6 is a transferase that has rotational symmetry of order 2

Astrophysics

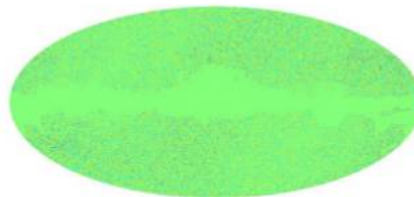
Primordial non-Gaussianity of Cosmic Microwave Background (CMB)

Filtered primordial CMB



WMAP mask and noise

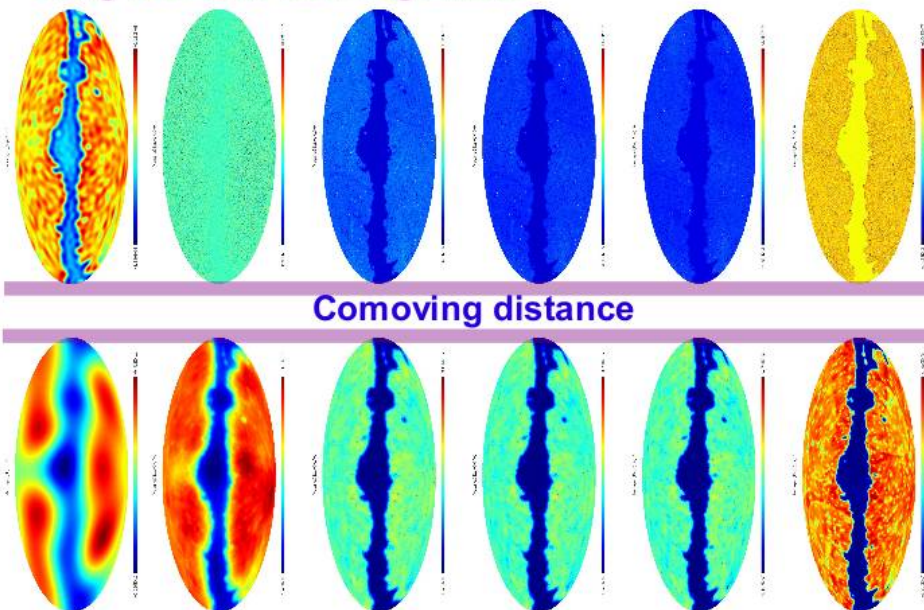
Filtered CMB with non-Gaussian contribution



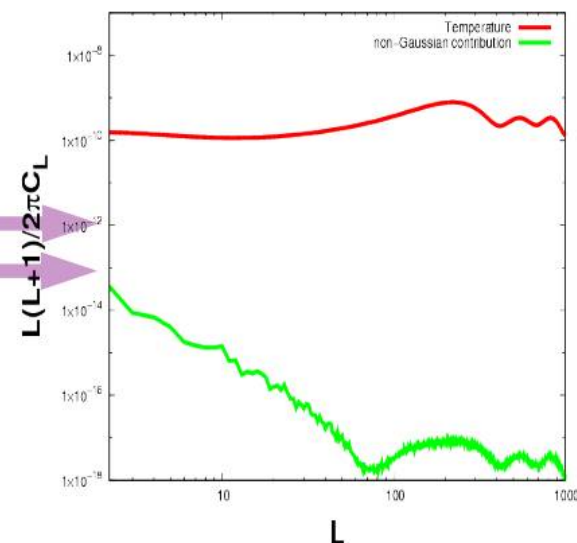
$$\Phi(\mathbf{x}) = \Phi^G(\mathbf{x}) + f_{\text{NL}} \left((\Phi^G(\mathbf{x}))^2 - \langle (\Phi^G(\mathbf{x}))^2 \rangle \right)$$

Test Inflationary models

Using OSG Grid to average files



Angular power spectra comparison

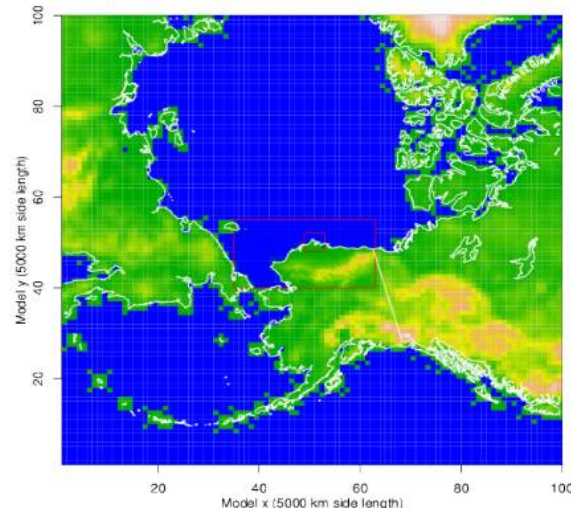


Climate science

Mülmenstädt/Climate science

Arctic regional modeling

- Arctic climate system is very sensitive to global warming because of strong positive feedback between temperature and surface reflectivity
- Correct parameterization of clouds in the Arctic is still elusive, and many phenomena (such as the coexistence of liquid water and ice in clouds as cold as -40 deg C over a span of several days) are the result of an interplay between several processes
- We are simulating 10 years' worth of regional weather systems over Alaska, eastern Siberia and the Arctic Ocean and comparing the simulated cloud properties with radar, lidar and radiometer observation – an unprecedented time span for model—data intercomparison
- The regional model runs in four-day chunks, with initialization and boundary conditions from a global weather analysis
- Each four-day chunk requires 12 h CPU time and is independent of all others → trivially parallelizable
- Boundary conditions, initialization and results are stored in Hadoop



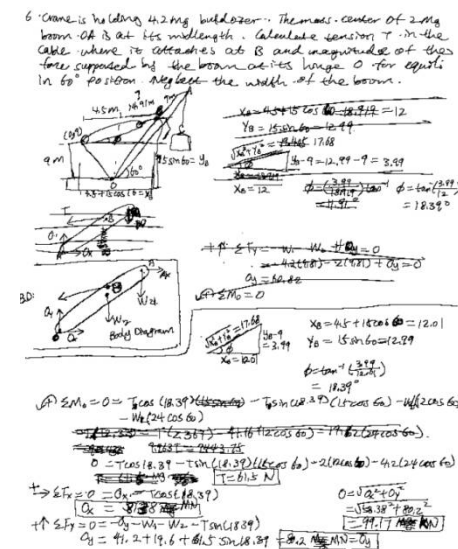
Properties of the stratocumulus cloud fields over the global oceans

- Marine stratocumulus exerts a strong cooling influence on global climate by reflecting incoming solar radiation and copiously emitting infrared radiation to space
- The occurrence of stratocumulus and its radiative properties are sensitive to perturbation by aerosols (natural and man-made)
- How sensitive they are to aerosols and how they will behave in a warmer climate are currently the largest sources of uncertainty in global atmospheric models
- Current-generation global models have limited vertical resolution and produce only a single layer of cloud
- This project uses satellite-borne imagers and laser profilers to measure the occurrence frequency of multi-layered stratocumulus
- Multi-layered stratocumulus may explain discrepancies between model and satellite estimates of the aerosol effects on clouds
- The satellite constellation orbits the Earth 17 times a day. Each orbit requires a few minutes of CPU time and is independent of all others → trivially parallelizable
- Satellite data and analysis results are stored in Hadoop

Educational Informatics



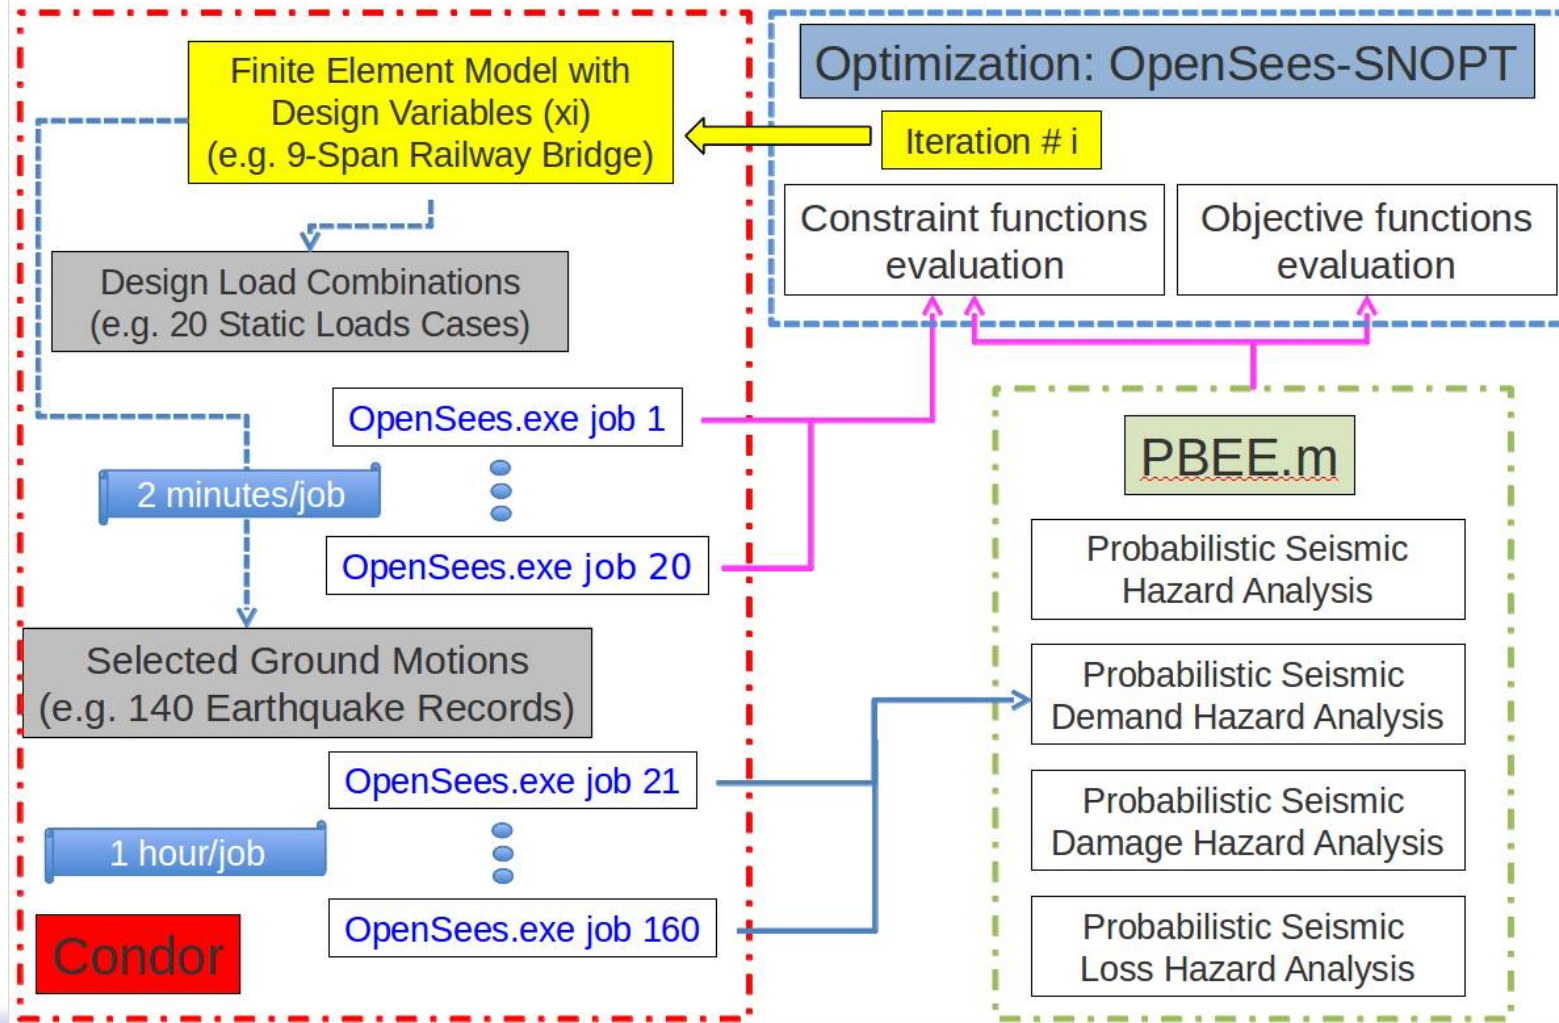
For the past three years, the Smart Tools Lab has provided undergraduate Mechanical Engineering students with Livescribe Digital Pens with which they completed their coursework. This has generated an unprecedented digital corpus of student work, comprising millions of strokes.



Using the OSG framework we have performed machine learning and data mining on this data, with exciting results. For example, we have trained statistical models capable of accurately predicting a student's grade by considering just the spatial and temporal organization of the solution. Additionally, we have identified significant patterns in the solution habits of successful students, leading to important discoveries about how students learn.

Engineering

Probabilistic Performance-based Earthquake Engineering Optimization Using Condor based in OSG



Summary

- UCSD invested a small amount of human resources to create a one-node glidein system for local users
 - But not from our department
 - CMS users well taken care of already
- Several scientists decided to try it out and liked it
 - Drastically increasing the amount of science they were able to do

Acknowledgments

- This document was sponsored by grants from the US NSF and US DOE, and by the UC system