

How I learned to Stop Worrying and Love Preemption

Condor Project
Computer Sciences Department
University of Wisconsin-Madison



Best thing about Condor?

The Community

Community as Metaphor

Two Questions

- > How much more room in OSG?
- > How long can jobs (effectively) run?


Open Science Grid Home page - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://www.opensciencegrid.org/

Most Visited Condor B&T Devel GCal WU Effort Certification Wiki #6 Bus

Open Science Grid Home page



Open Science Grid

A national, distributed computing grid for data-intensive research.

Search OSG using Google:

Search OSG at Work:

Home

About the OSG

Learn About Us

What We're Doing

Getting Started with OSG

Contact Us

OSG at work

Calendar of Events

Collaborative Twiki and Chat

Technical Documentation

Security

Software

Education

Information and help

Grid Operations Center

Documentation and Reference

Monitoring


Research Highlights

Case Study - Einstein@OSG
Einstein@Home, an application that uses spare cycles on volunteers' computers, is now running on the OSG.

PEGrid gets down to business
Grid technology enables students and researchers in the petroleum industry.

Linked grids uncover genetic mutations
Superlink-online helps geneologists perform compute-intensive analyses to discover disease-causing anomalies.

[View More Research Highlights and Vignettes...](#)


"Having adapted its SAMGrid data system to use the new GlideinWMS workload management system, D0 is now running its simulations successfully, both across OSG and the European EGEE sites. Since we now use the OSG tools to submit directly to EGEE sites, we can keep one set of code for all the grid infrastructures we use, making maintenance, monitoring, and jobs themselves much more efficient."
~Adam Lyon, D0 physicist and head of the Run II production support group

OSG News

Operational News & Announcements

Security News & Announcements

OSG Consortium All Hands Meeting 8-11 March 2010 at Fermilab

Campus grids secret to productive grid sites - ISGTW 31 Mar

Video of the week: Monitoring with the fishes - ISGTW 31 Mar

Q&A: Grid Colombia warms up - ISGTW 24 Mar

OSG Newsletter
[\[Archive\]](#)

Usage Stats

ATLAS Operations on the OSG July 2009 thru December 2009

Wall Clock Hours	Six Month Sum of Wall Clock Hours	Number of Jobs
10000	10000	10000
20000	20000	20000
30000	30000	30000
40000	40000	40000
50000	50000	50000
60000	60000	60000
70000	70000	70000
80000	80000	80000
90000	90000	90000
100000	100000	100000

Usage Type at OSG Sites

Average Number of CPUs Released

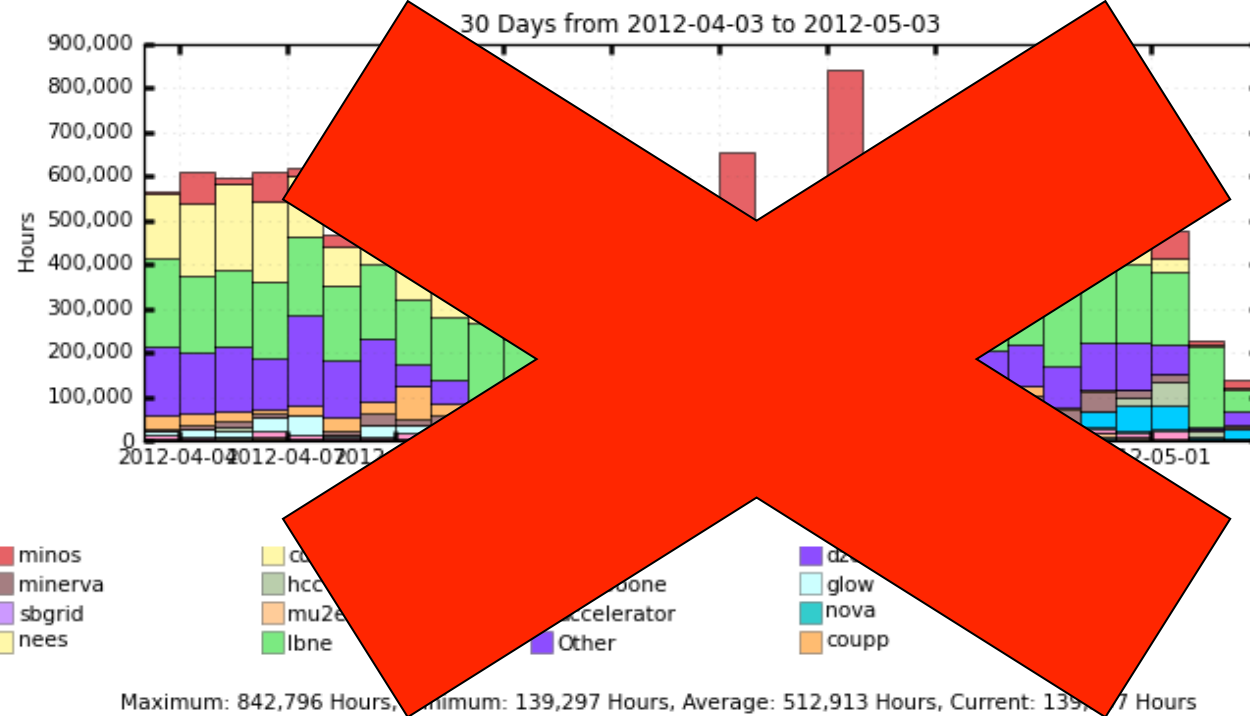
Problems Moved

View Live Grid Status

Done



Not an OSG talk



3 method for performance analysis

- > Modeling
 - Difficult in opportunistic pools
- > Simulation
 - Difficult in opportunistic pools
- > Measurement

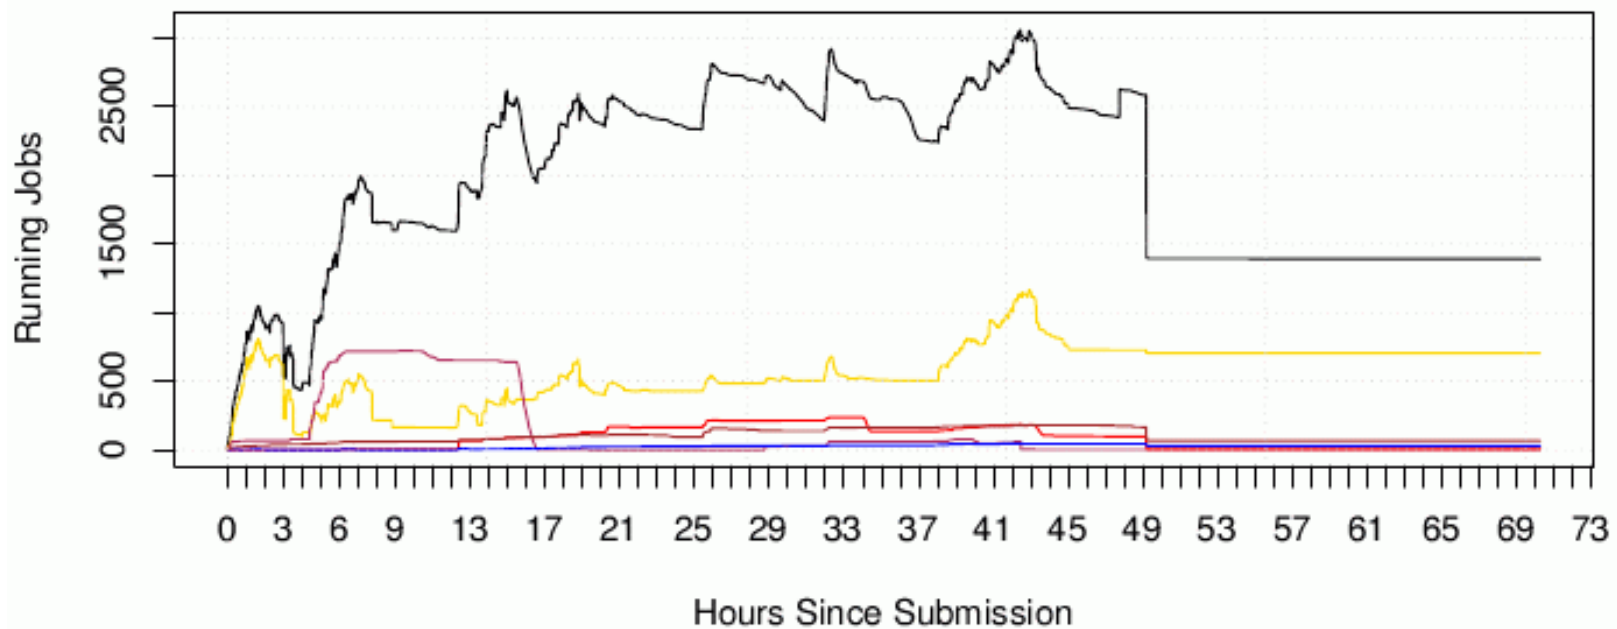
Measurement

- > Simply submit a ton of jobs
- > All measurement via user log

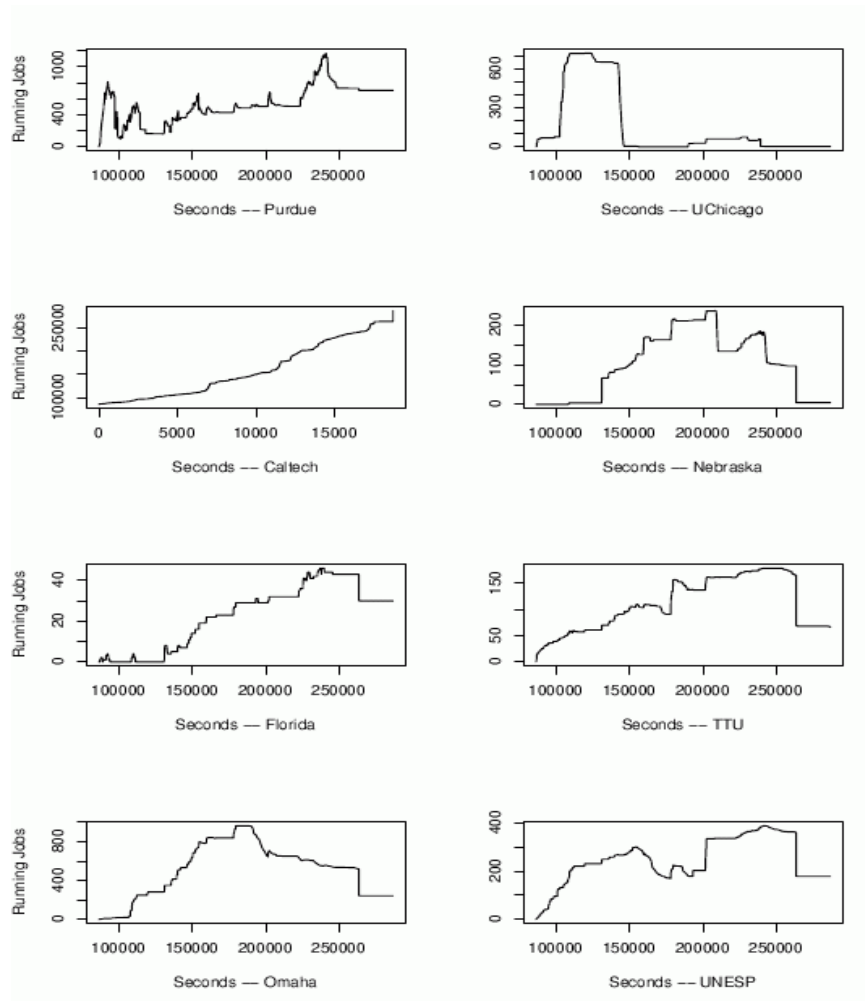
User log trick:

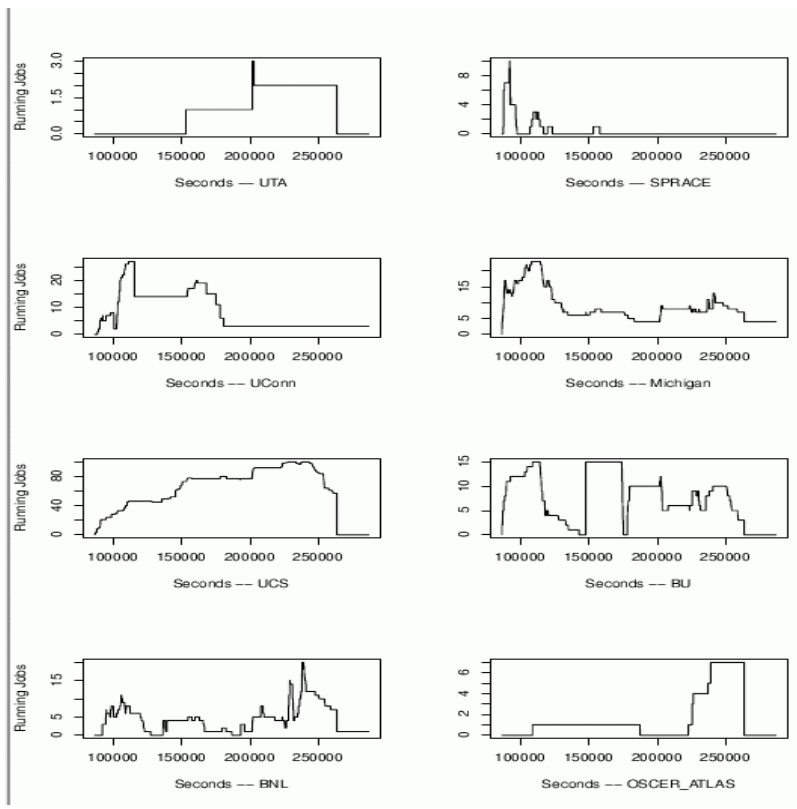
- > +TheSite="\$\$ (GLIDEIN_Site)"
- > job_ad_information_attrs = TheSite

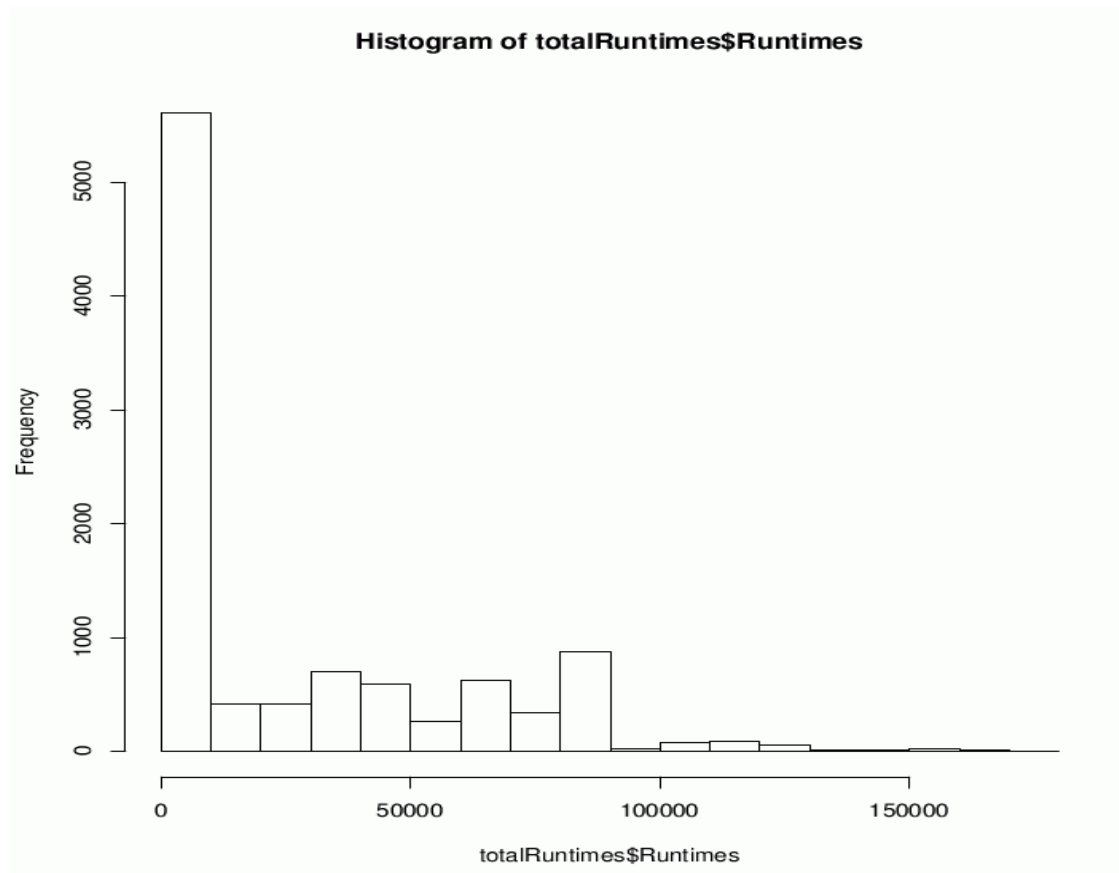
- > 028 (4.1.0) 02/17 18:38:03 Job ad information event triggered.
- > TriggerEventTypeName = 1
- > EventTypeNumber = 28
- > TriggerEventTypeName = "ULOG_EXECUTE"
- > Proc = 87
- > Subproc = 0
- > TheSite = "Purdue"
- > CurrentTime = time()
- > MyType = "ExecuteEvent"

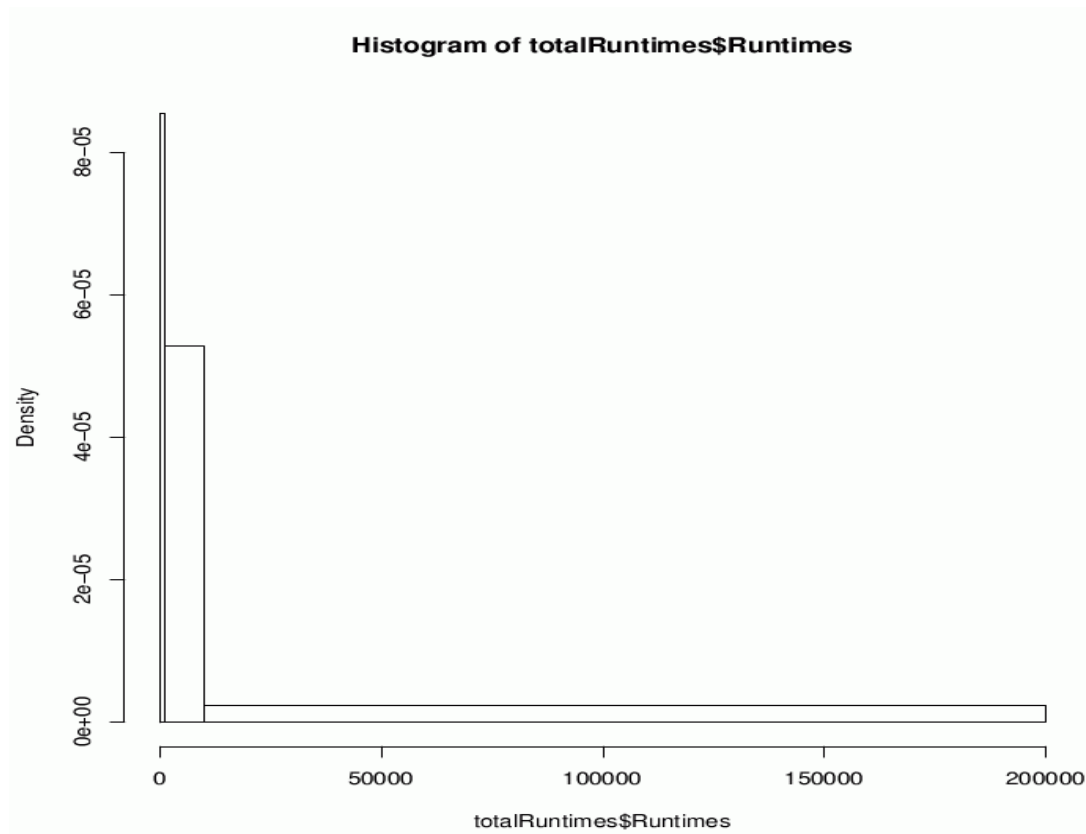


76,161 total cpu-hours












Preemption Probability





Wall clock slowdown



How to fix?

- > Mix and match
- > Opportunistic for throughput
- > Local resources for tail chopping

Summary

- > Opportunistic cycles are still very useful
- > Fairness and Throughput can be duals
- > Opportunistic + dedicated very powerful