



Powering Statistical Genetics with the Grid: Using GridWay to Automate R Workflows

John-Paul Robinson

Information Technology

Purushotham Bangalore

Department of Computer Science

Jelai Wang, Tapan Mehta

Department of Biostatistics

January 30, 2008

Grid Enabling Applications Workshop

Mardi Gras Conference 2008

Baton Rouge, LA



Outline

- Overview
- Our Plan
- Problems
- Solutions
- Project Status
- Future Directions



Statistical Genetics

- One area of study is the application of statistical methods to understand gene expression and how it relates to expressed traits
- Uses statistical analysis to determine the distribution of traits
- Attempts to build a mapping between general genetic traits and an individual's expression of those traits
- For example, which genes control the expression of diabetes?



UAB Statistical Genetics

- Section on Statistical Genetics (SSG) is part of the Department of Biostatistics in the School of Public Health
- Provide graduate programs in statistical genetics
- Support research use of statistical methods
- Conduct statistical methodology research
- Lead by Dr. David Allison




Statistical Genetics Workflow

- Applied Analysis
 - Real data from experimental results
 - Variables are regions of genome and may vary from 10K to 100K
 - Ex. identify region association with a disease
 - Occurs when experimental data is available
- Methodological Analysis
 - Random data generated for simulated analysis
 - Understand behavior of statistical methods used in real data analysis
 - 2k - 5k variables analyzed (maximum practical for existing, traditional 2 cluster resource pool)
 - Occurs repeatedly when methods are studied


Statistical Genetics Workflow Characteristics



- Data analysis code written in R
- Process level and MPI granularity
- Embarrassingly (obvious) parallel (workflow)
- Already benefits from cluster-scale computing (384 CPUs, ~3Tflops)



Single Process



MPI-Based Process



Statistical Genetics Workflow Issues

- Methodological analysis limited by available compute power
 - Accuracy improves when simulated data is same size as real data
 - Currently 2k-5k variables, ideally 10k-20k (two to ten times the current scale!)
- Manual, multi-cluster job management is problematic and inefficient
- Manual, multi-cluster R environment management does not scale
- Many inexperienced users excluded



Can “the Grid” Help?

- The workflow characteristics are ideal for grid migration – 1000's of isolated computations
- More compute power readily increases workflow throughput
- Management issues are not computation problems but scaling problems
- Signs point to Yes



R-Group Goals

- Assess feasibility of grid-based workflow
- Build larger computational pool to increase R program workflow throughput
- Expand accessibility of workflow leveraging web-based user interface to grid via GridSphere
- Migrate workflow as a test of UABgrid infrastructure



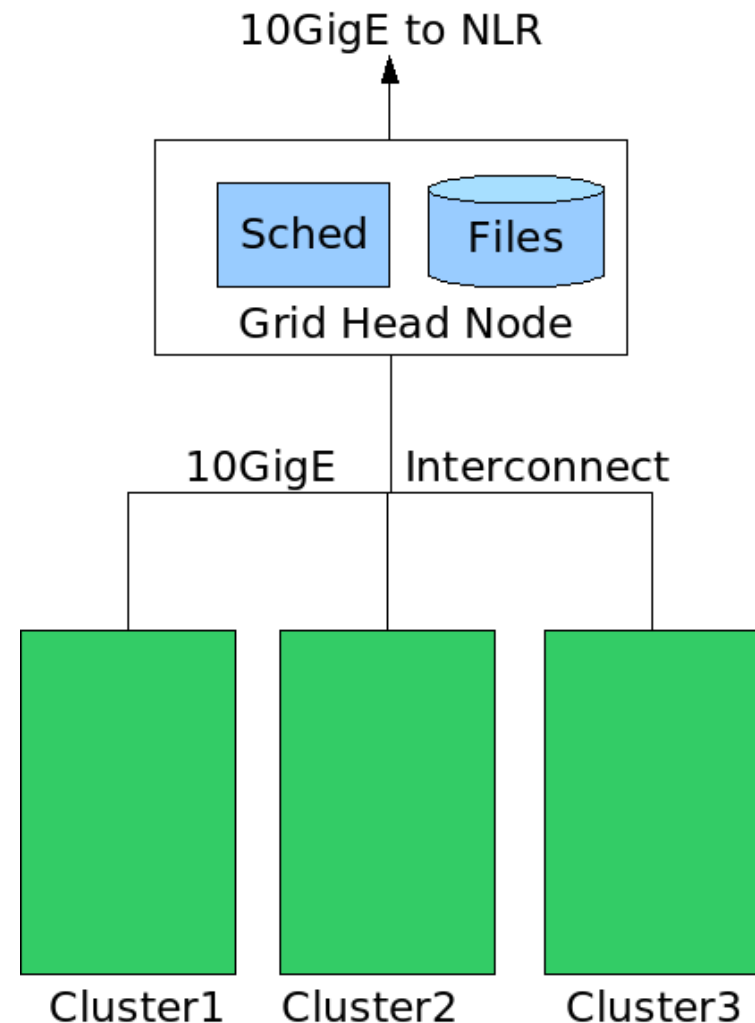
UABgrid Overview

- UABgrid Pilot launched at campus HPC Boot-Camp September 2007
- User-driven collaboration environment supports web and grid applications
- Leverages InCommon for user identification
 - SSO for web applications
 - Self-service certificate generation for Globus users
- Provides meta-cluster to harness on- and off-campus compute power using GridWay

UABgrid Meta-Cluster Design



- Grid head node for job and data staging
- On-campus ROCKS clusters supply compute power via Globus interface
- Additional compute power supplied by SURAgrid, OSG, or TeraGrid



UABgrid Meta-Cluster Specs

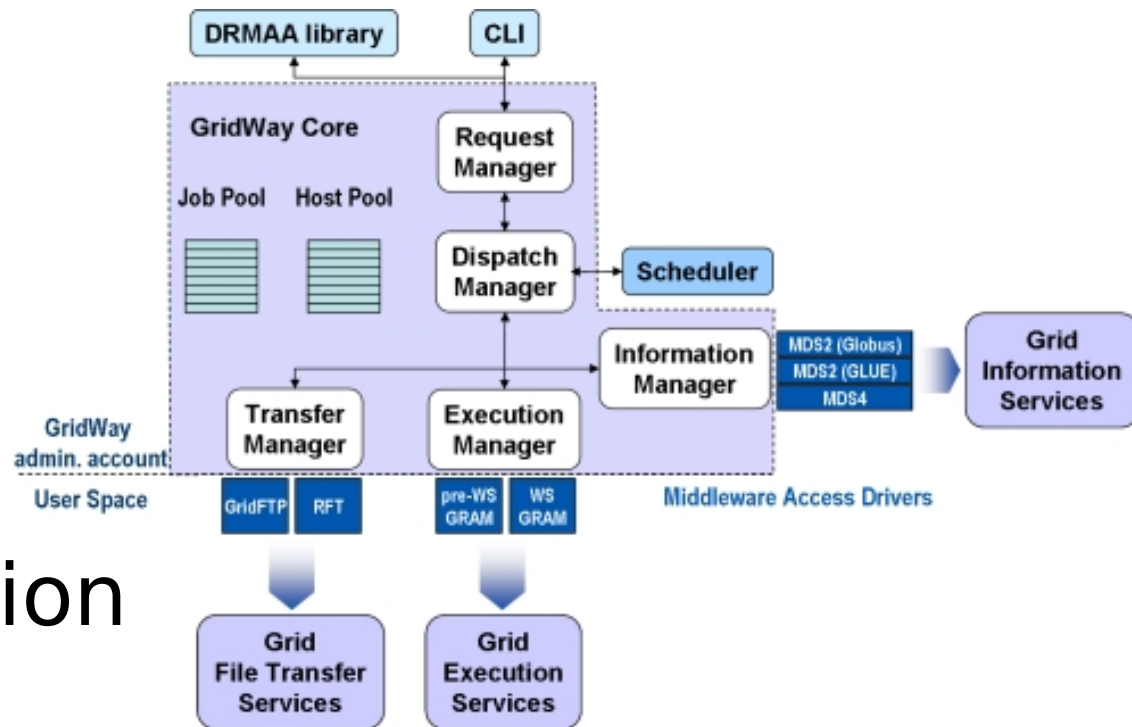


- Today
 - 3 campus clusters, 448 processors, 3+Tflops of power
- 2008 Q2 Targets
 - 5 campus clusters, 960 processors, ~8Tflops of power
 - Large-memory system at state super-computing center
- On Going
 - Engage SURAggrid, TeraGrid, and other suppliers for more compute power
 - Local expansion through campus HPC investments



GridWay Overview

- Grid meta-scheduler
- Manages job and data staging to clusters using Globus interfaces
- Familiar job template abstraction
- Pluggable schedulers and providers





Initial Solution

- It's an ideal grid use case
- R application already works on target cluster
- UABgrid meta-cluster already has working GridWay job submission to target cluster
- Should be straight forward migration of cluster (SGE) job scripts to GridWay job scripts
- What could possibly go wrong?

What Could Possibly Go Right?



- GridWay Issues
- Globus Issues
- MPI Issues
- R Issues
- Configuration Management Issues



GridWay Issues

- GridWay job templates stage the data and the executable
- Executable staging controlled with path naming
 - Relative path names cause staging
 - Fully qualified path names avoid it
- Fully qualified executable names are not portable - different cluster, different path
- Easily solved with shell script wrappers - delay executable name resolution until runtime



Globus Issues

- SGE is not a native job manager in Globus
- Requires error-prone post-install configuration on ROCKS clusters
- Job Type definitions (single, multiple, MPI) are not well defined and vary across sites and schedulers
 - SGE treats single jobs with $\text{CPU} > 1$ as array job
 - PBS treats it as a single job that reserved x CPUs



MPI Issues

- MPI job type in Globus has been the most challenging
- Globus SGE JobManager assume MPICH (MPIv1)
- R MPI codes assume MPIv2 support, prefer LAM/MPI or OpenMPI
- Exploring two solutions
 - Globus LAM/OpenMPI support
 - Rmpi MPICH support
- Simply does not work out of the box



R Issues

- Development is active with regular release cycles
- Feature rich environment for developing and using statistical methods
- It does for statistics what Perl does for text processing
- Large public library of tools and methods CRAN (akin to Perl's CPAN)
- Many jobs have compatibility matrix
- Not bad, just complex



Configuration Management Issues

- Compatibility matrix for R scripts hard to maintain
- Many clusters with identical configurations not available (not even with 2 clusters)
- Need solution that keeps configuration in the hands of those most interested
- Look to options that empower application specialists with configuration management framework...



Containerization

- Types of Containers
 - User Accounts
 - Virtual Machines
 - Java Boxes
- Account Container
 - Initial target because most common and addresses R application configuration
 - Allows for library dependency and site dependency configuration
 - Full continuum of deployment options from fully staged for each job to statically cached on resources



Reflections on Progress

- Would have liked to focus on workflow migration issues, instead we had to focus on component compatibility issues
- Initial goal was intentionally simplistic but provide surprisingly challenging
- Don't be surprised when infrastructure is not ready
- Problems can be solved individually, but all will need to be solved before it will work – *this is a production application*



R-Group Update

- Lifecycle Model
 - Feasibility Study
 - Requirements Analysis
 - Design Application Framework
 - Implementation by mid-2008
 - Maintenance
- Status
 - Feasibility study reasonably complete: definitely possible even with technology glitches to work out
 - Requirements for existing resources are known
 - Will continue to follow open source model: release early and release often



Future Plans

- ASA Resources
 - Add large memory resource via grid to address some R job requirements for large data structures
- SURAgrid and OtherGrid Resources
 - What resources are available for reliable or predictable compute power generation?
 - Explore NIH Biowulf “swarm” model
- Explore custom resources like BlueGene
 - Are there statistical computations that could benefit from porting to a massively parallel processor?
 - Could web-service provide R-level abstractions?



Addition Information

- UABgrid R-Group Project
 - <http://projects.uabgrid.uab.edu/r-group>
- UAB Section on Statistical Genetics (SSG)
 - <http://www.ssg.uab.edu>
- UABgrid
 - <http://docs.uabgrid.uab.edu>
- UAB
 - <http://www.uab.edu>



Contacts

- John-Paul Robinson <jpr@uab.edu>
- Purushotham Bangalore <puri@cis.uab.edu>
- Jelai Wang <JWang@ms.soph.uab.edu>
- Tapan Mehta <TMehta@ms.soph.uab.edu>