

Workflow Technologies for Science Automation

Ewa Deelman


USC Information Sciences Institute

<http://www.isi.edu/~deelman>

Funding from the NSF and NIH

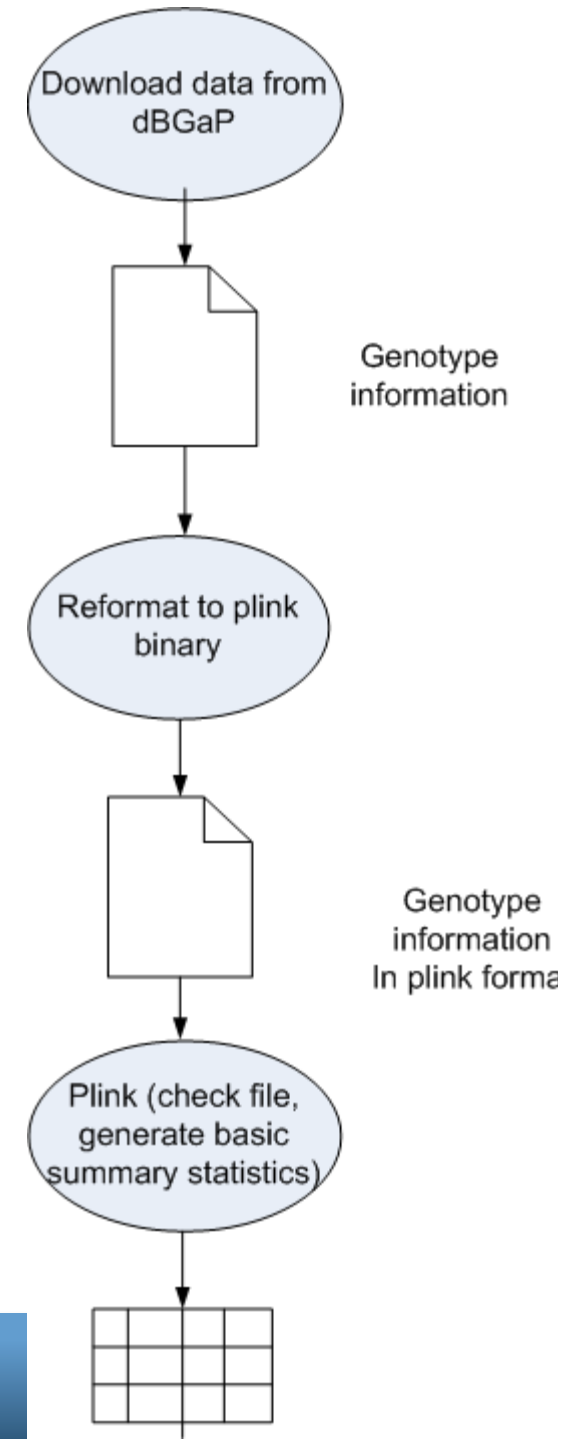
Outline

- **Scientific Workflows and Application Examples**
- **Workflow Management Systems**
- **Pegasus and its features**
- **Conclusions**




Scientific Workflows

- Structure an overall computation
- Define the computation steps and their parameters
- Define the input/output data, parameters
- Invoke the overall computation
- Reuse with other data/parameters/algorithms and share
- Workflows can be hidden behind a nice user interface (e.g. portal)



Benefits of Scientific Workflows (from the point of view of an application scientist)

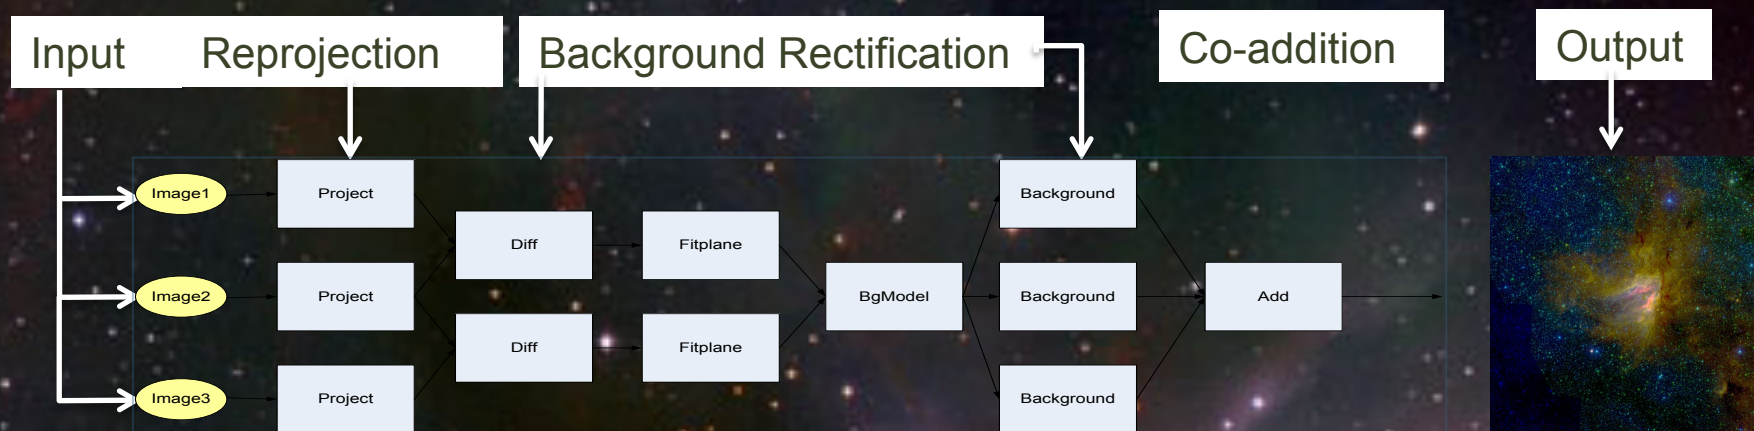
- Conducts a series of computational tasks.
 - Resources distributed across Internet.
- Chaining (outputs become inputs) replaces manual hand-offs.
 - Accelerated creation of products.
- Ease of use - gives non-developers access to sophisticated codes.
 - Avoids need to download-install-learn how to use someone else's code.
- Provides framework to host or assemble community set of applications.
 - Honors original codes. Allows for heterogeneous coding styles.
- Framework to define common formats or standards when useful.
 - Promotes exchange of data, products, codes. Community metadata.
- Multi-disciplinary workflows can promote even broader collaborations.
 - E.g., ground motions fed into simulation of building shaking.
- Certain rules or guidelines make it easier to add a code into a workflow.



Science-grade Mosaic of the Sky



Science-grade Mosaic of the Sky



Montage Workflow





Size of mosaic in degrees square	Number of input data files	Number of tasks	Number of intermediate files	Total data footprint	Cumulative wall time
1	84	387	850	1.9 GB	21 mins
2	300	1442	3176	6.8 GB	54 mins
4	685	3738	8258	18 GB	3 hours, 18 mins
6	1461	7462	16458	37 GB	7 hours, 7 mins
8	2565	12757	28113	64 GB	11 hours, 44 mins

Hunting Exoplanets with Kepler


<http://kepler.nasa.gov>

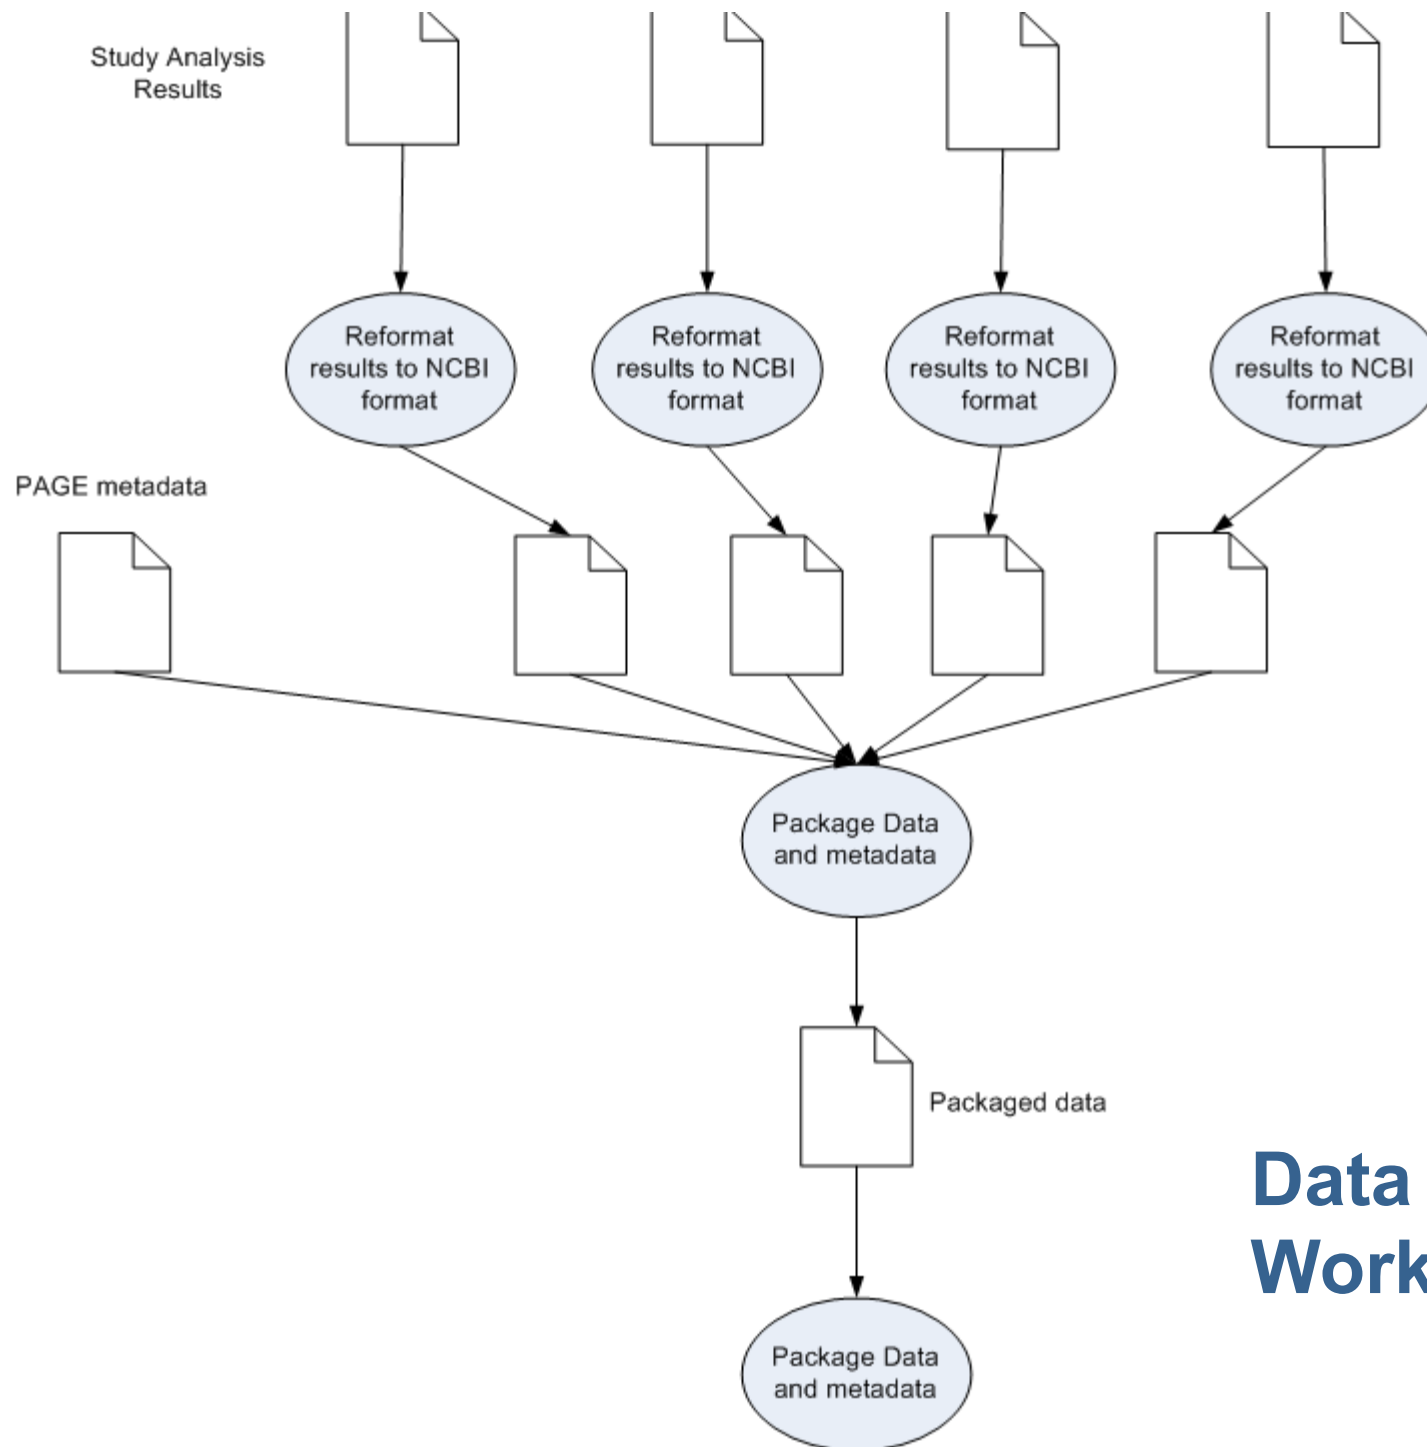
- Kepler continuously monitors the brightness of over 175,000 stars.
 - Search for periodic dips in signals as Earth-like planets transit in front of host star.
- Over 380,000 light curves have been released.
- Can take 1 hour to perform periodogram analysis of Kepler light curve
- Need to perform a bulk analysis of all the data to search for these periodic signals

- 210K input, 630K output files
- 210K tasks total




Kepler 6-b transit






Data Management Workflow



Outline

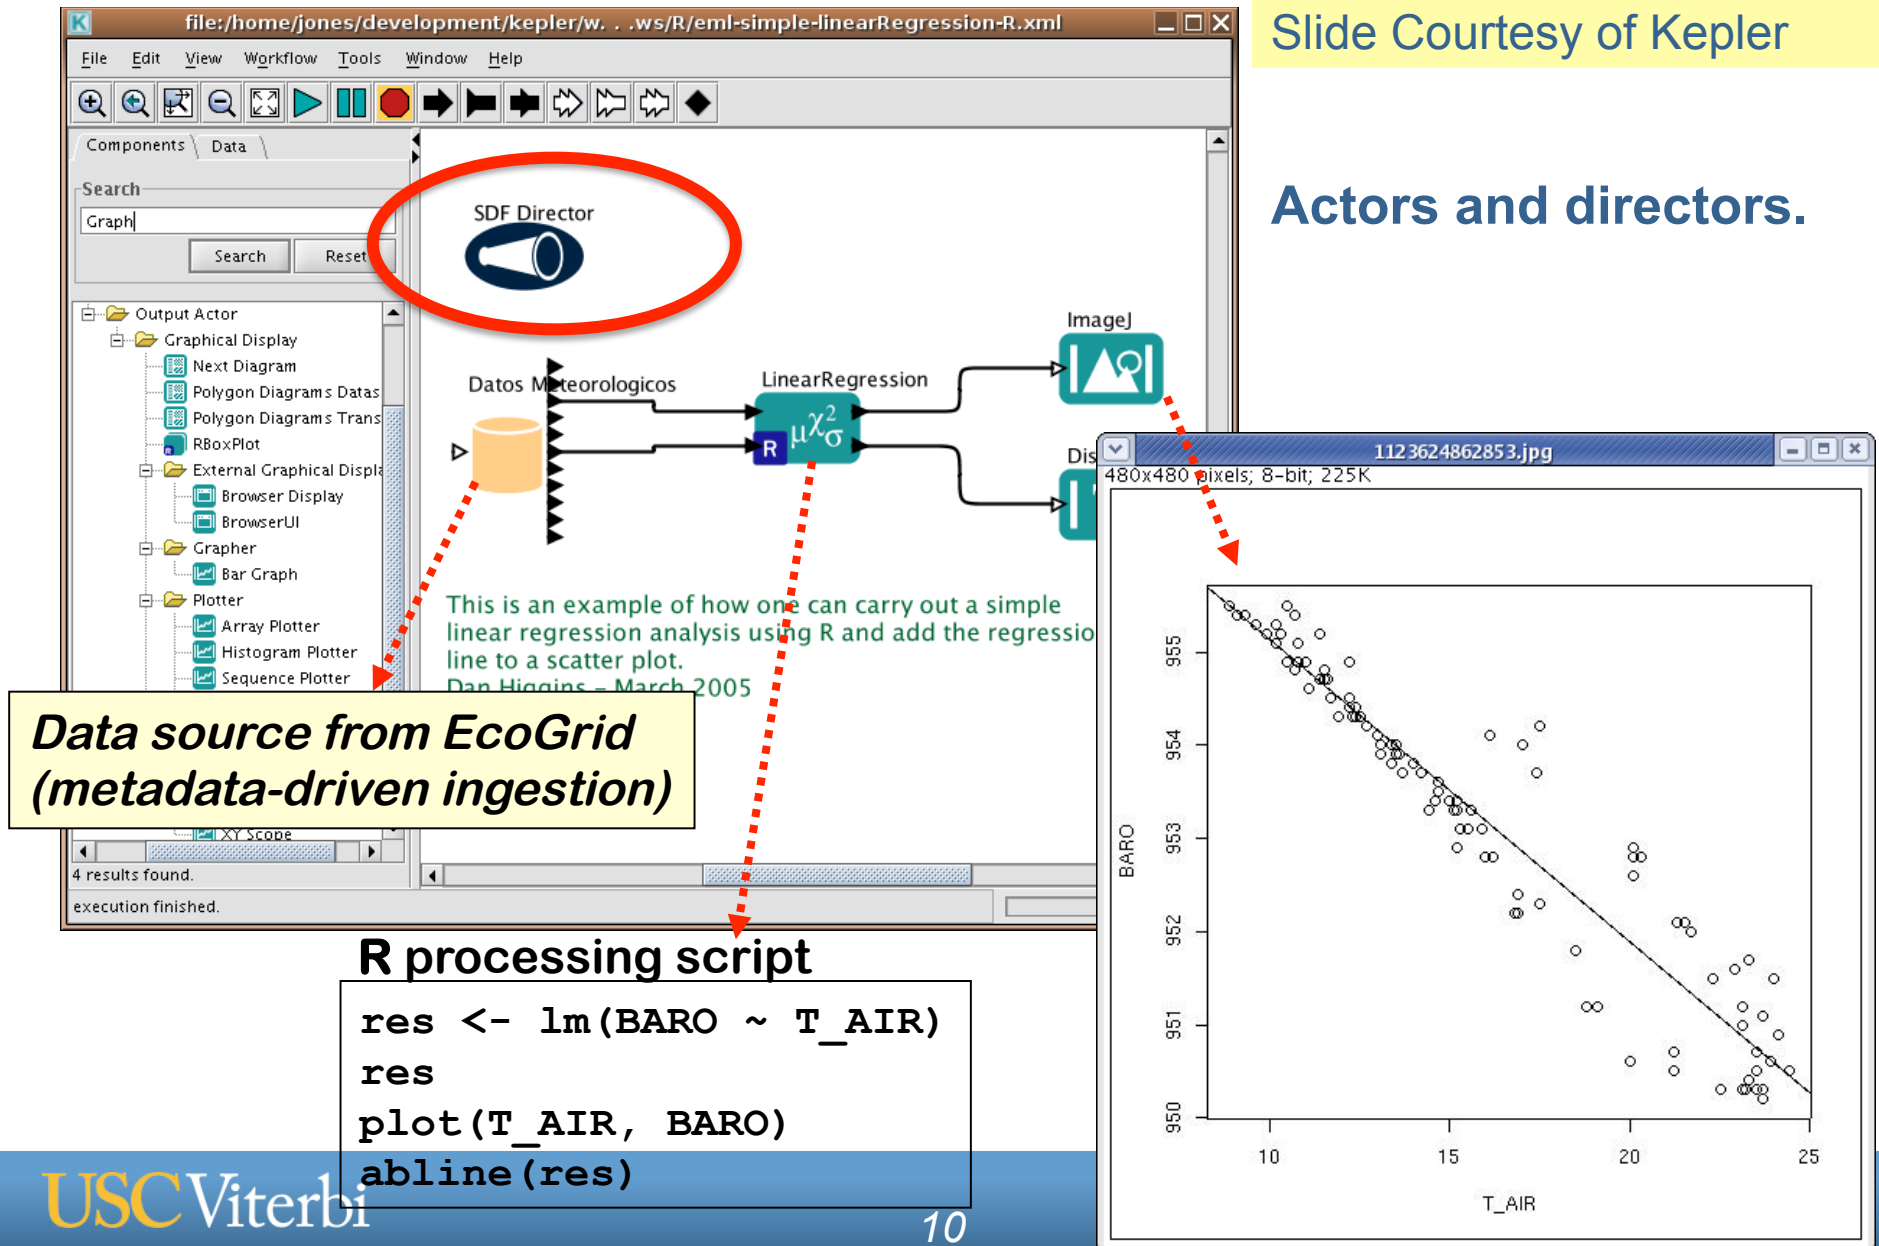
- Scientific Workflows and Application Examples
- **Workflow Management Systems**
- Pegasus and its features
- Conclusions



Simple Kepler workflow using R

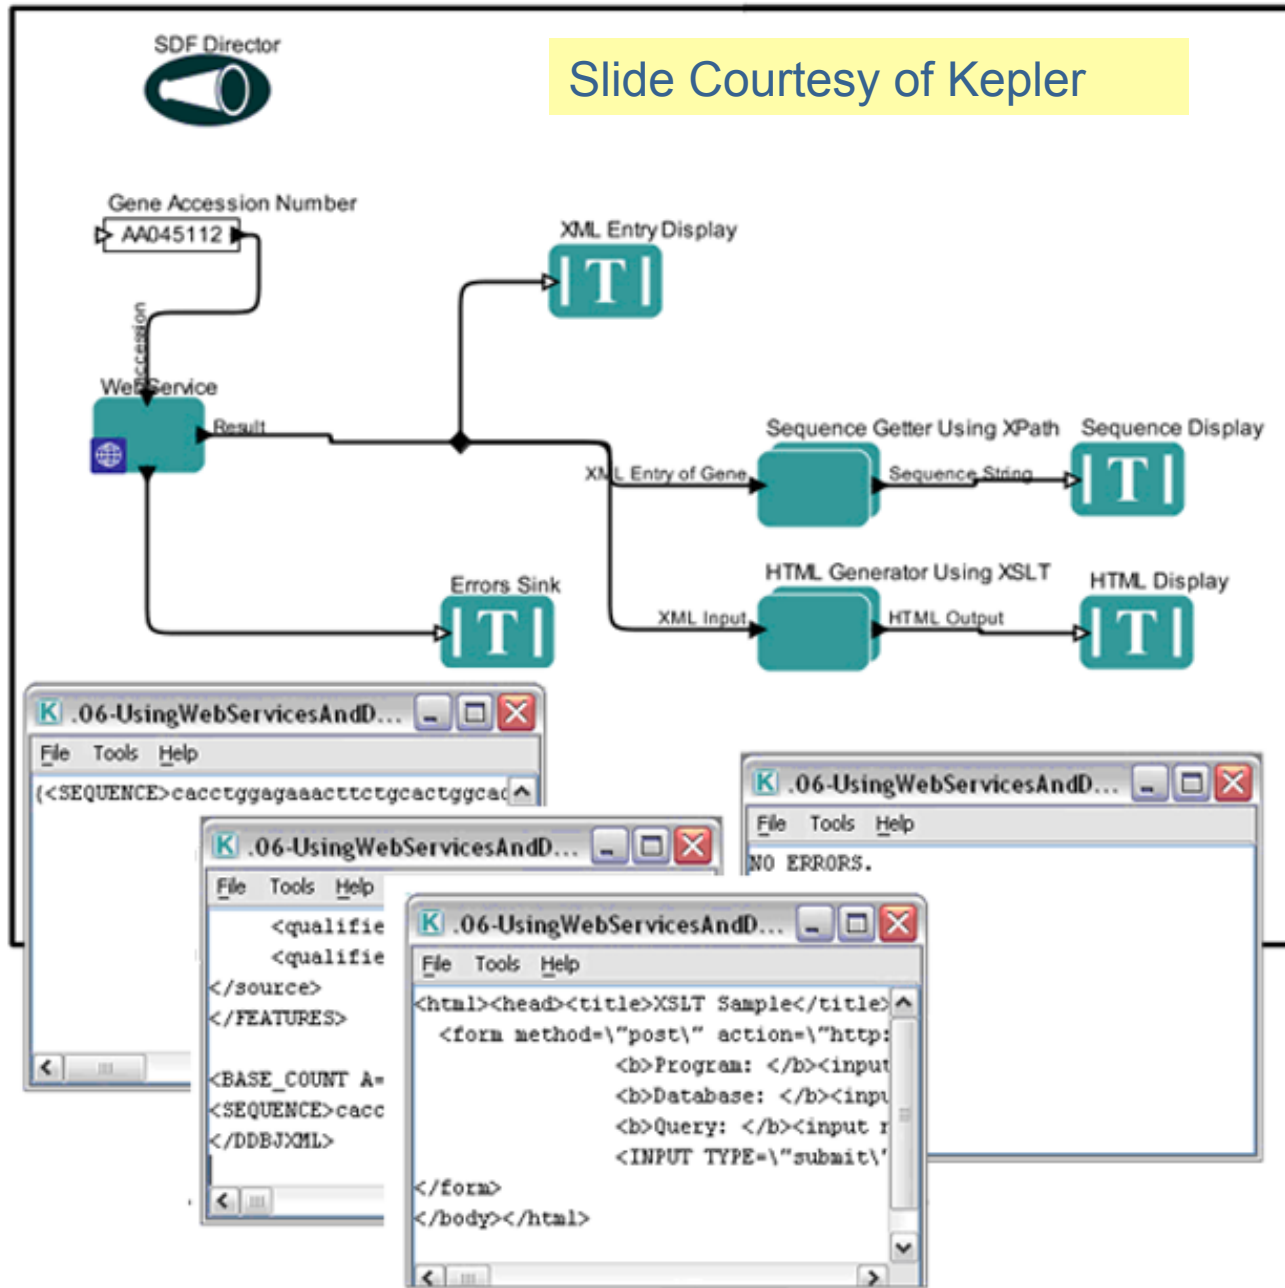
Slide Courtesy of Kepler

Actors and directors.



Web Services Workflow

The Web Services workflow uses Kepler's Web Service actor to invoke a genomics data web service, which accesses and queries a remote genomics database and returns a genetic sequence. The name of the sequence (i.e., the gene accession number) is passed to the Web Services actor by a String Constant actor.



Taverna Workflows

Workflow language specifies how bioinformatics processes fit together.

High level workflow diagram separated from any lower level coding –

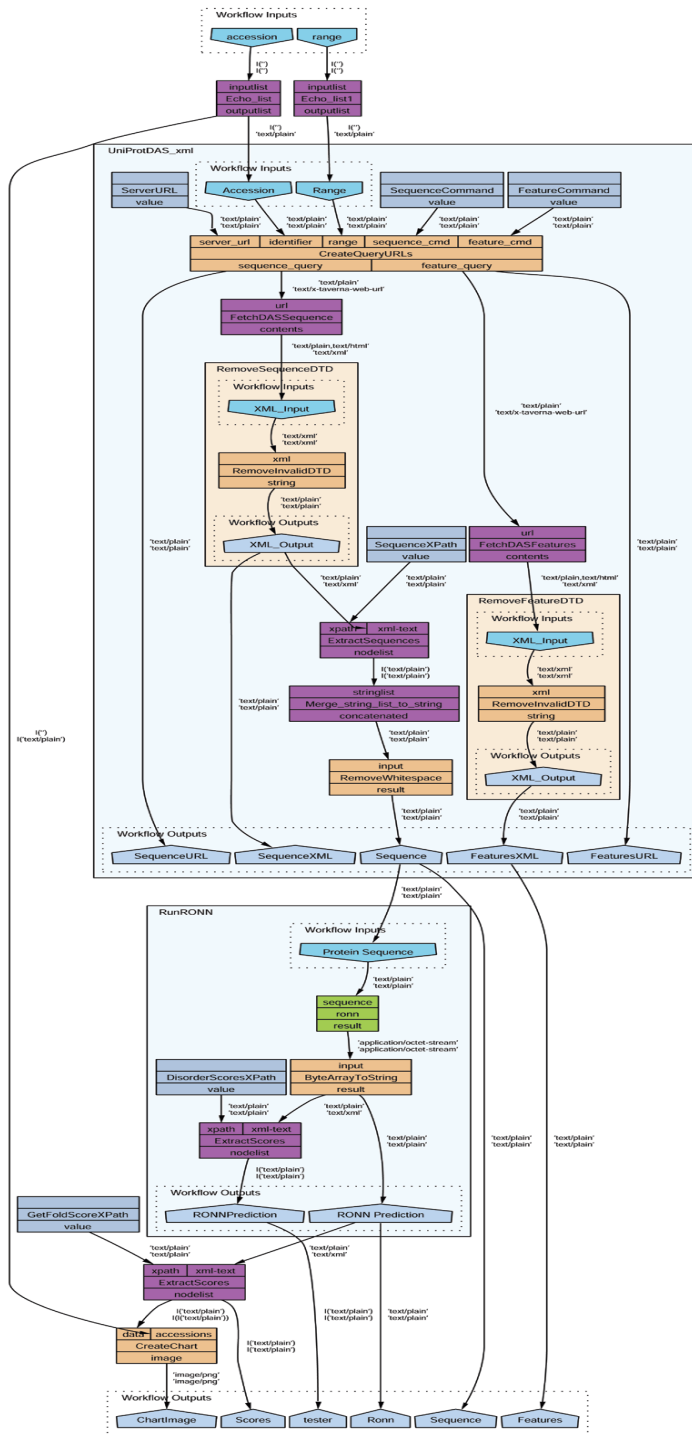
Easier to explain, share, relocate, reuse and repurpose.

Workflow \Leftrightarrow Model

Workflow is the integrator of knowledge

The METHODS section of a scientific publication

Slide Courtesy of Taverna



Galaxy– a bioinformatics workbench

Galaxy Interface: Analyze Data

Data analysis

The screenshot displays the Galaxy / WIBR web interface. The top navigation bar includes 'Galaxy / WIBR', 'Analyze Data', 'Workflow', 'Shared Data', 'Help', and 'User'. The left sidebar contains a 'Tools' menu with categories like 'Get Data', 'Text Manipulation', 'Filter and Sort', 'Join, Subtract and Group', 'Convert Formats', 'Extract Features', 'Fetch Sequences', 'Fetch Alignments', 'Get Genomic Scores', 'Operate on Genomic Intervals', 'FASTA manipulation', 'NGS: QC and manipulation', 'NGS: Mapping', 'NGS: RNA Analysis', 'NGS: SAM Tools', 'NGS: Peak Calling', and 'Workflows'. The central panel shows a 'Data display and tool's dialog window' with a text area containing genomic data. The right panel shows a 'History' window with a list of datasets, including '2: Filter sequences by length on data 1' and '1: AllIntrons.fasta'. A legend on the right side explains the status colors: Green for 'job is finished', Yellow for 'job is running', Gray for 'job is in queue', and Red for 'there is a problem'. A double-headed arrow at the bottom indicates the layout of the three main windows: Tools window, Data display and tool's dialog window, and History window.

Tools window

Data display and tool's dialog window

History window: datasets for each analysis are kept here

Processed data

Green: job is finished
Yellow: job is running
Gray: job is in queue
Red: there is a problem

4

Slide Courtesy of Galaxy

Creating Workflows

The screenshot shows the Galaxy web interface. The top navigation bar includes 'Analyze Data', 'Workflow', 'Shared Data', 'Help', and 'User'. The 'History' panel on the left displays a list of datasets. A context menu is open over the history, listing various actions. The 'Extract Workflow' option is circled in red.

History

format: sam, database: mm9
Info: Sequence file aligned.

1.QNAME	2.FLAG	3.RNAME	4.POS
@HD	VN:1.0	SO:unsorted	
@SQ	SN:chr1	LN:197195432	
@SQ	SN:chr10	LN:1	
@SQ	SN:chr11	LN:1	
@SQ	SN:chr12	LN:1	
@SQ	SN:chr13	LN:1	

History Lists

- Saved Histories
- Histories Shared with Me
- Current History
- Create New
- Clone
- Copy Datasets
- Share or Publish
- Extract Workflow**
- Dataset Security
- Show Deleted Datasets
- Show Hidden Datasets
- Purge Deleted Datasets
- Show Structure
- Export to File
- Delete
- Delete Permanently

6: FastQC.html

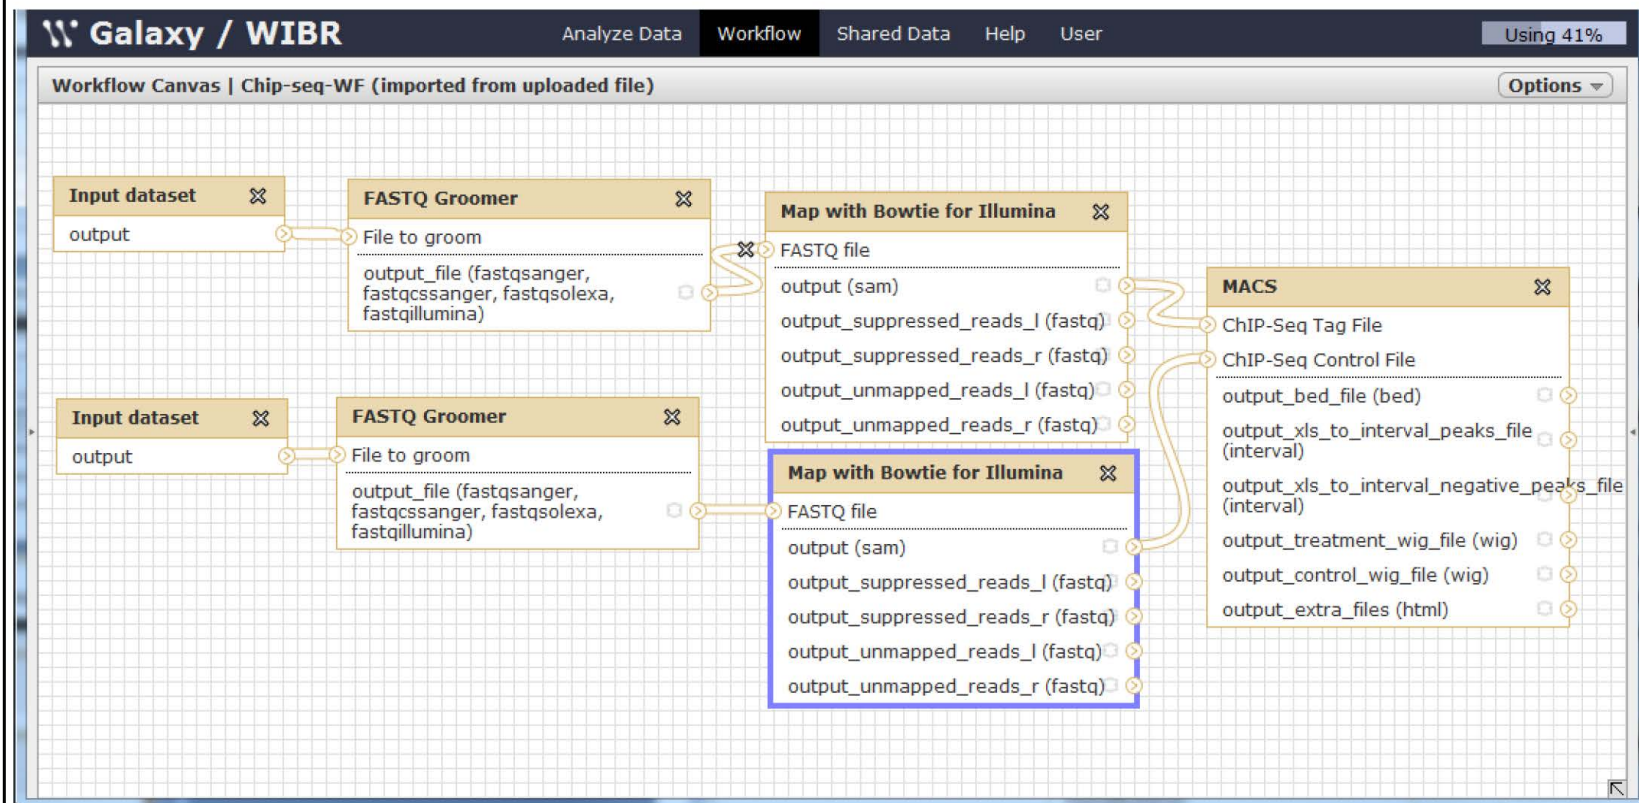
5: FastQC.html

4: FASTQ Groomer on dat

3: FASTQ Groomer on dat

Extract
workflow from
history

Workflow for ChIP-seq analysis



A workflow repository

Slide Courtesy of Myexperiment



[About](#) | [Mailing List](#) | [Publications](#)

[Log in](#) | [Register](#) | [Give us Feedback](#) | [Invite](#)

[Home](#) | [Users](#) | [Groups](#) | **[Workflows](#)** | [Files](#) | [Packs](#) | [Topics](#)

[Home](#) > [Workflows](#)

Workflows

Search filter terms

Filter by type

- ☐ Taverna 2 1091
- ☐ Taverna 1 562
- ☐ RapidMiner 237
- ☐ Kepler 43
- ☐ Bioclipse Scri... 34
- ☐ LONI Pipeline 26
- ☐ GWorkflowDL 24
- ☐ KNIME 22
- ☐ BioExtract Ser... 18
- ☐ SPARQL 12

Filter by tag

- ☐ example 223
- ☐ mygrid 103
- ☐ bioinformatics 102
- ☐ localworker 102
- ☐ cheminformatics 81
- ☐ benchmarks 77
- ☐ protein 75
- ☐ workflow 72

« Previous **1** [2](#) [3](#) [4](#) [5](#) ... [218](#) Next »

Sort by: Rank

Showing 2177 results. Use the filters on the left and the search box below to refine the results.

Search

Taverna 2



Pathways and Gene annotations for QTL region (7)

Original
Uploader



Paul
Fisher


Created: 19/11/09 @ 18:18:52 | Last updated: 07/09/12 @ 18:23:36

Credits: Paul Fisher

License: Creative Commons Attribution-Share Alike 3.0 Unported License

View

Download (v7)



This workflow searches for genes which reside in a QTL (Quantitative Trait Loci) region in the mouse, *Mus musculus*. The workflow requires an input of: a chromosome name or number; a QTL start base pair position; QTL end base pair position. Data is then extracted from BioMart to annotate each of the genes found in this region. The Entrez and UniProt identifiers are then sent to KEGG to obtain KEGG gene identifiers. The KEGG gene identifiers are then used to search for pathways in the KEGG path...

Rating: 4.6 / 5 (10 ratings) | Versions: 7 | Reviews: 1 | Comments: 7 | Citations: 1

Viewed: 7002 times | Downloaded: 1228 times

Tags (22):

adasd | chromosome | data-driven | disease | ensembl | entrez | gene | genes | genotype | kegg | mouse | nbiconworkflows | pathway | pathway-driven | pathways | phenotype | qtl | shim |

New/Upload

Workflow GO

Log in / Register

Username or Email:

Password:

Remember me: ☐

OR

Use OpenID:

(eg: name.myopenid.com)

Log in

Need an account?
[Click here to register](#)

[Forgot Password?](#)

Popular Tags

25 tags

[\[All Tags\]](#)


[astronomy](#) | [benchmarks](#) |

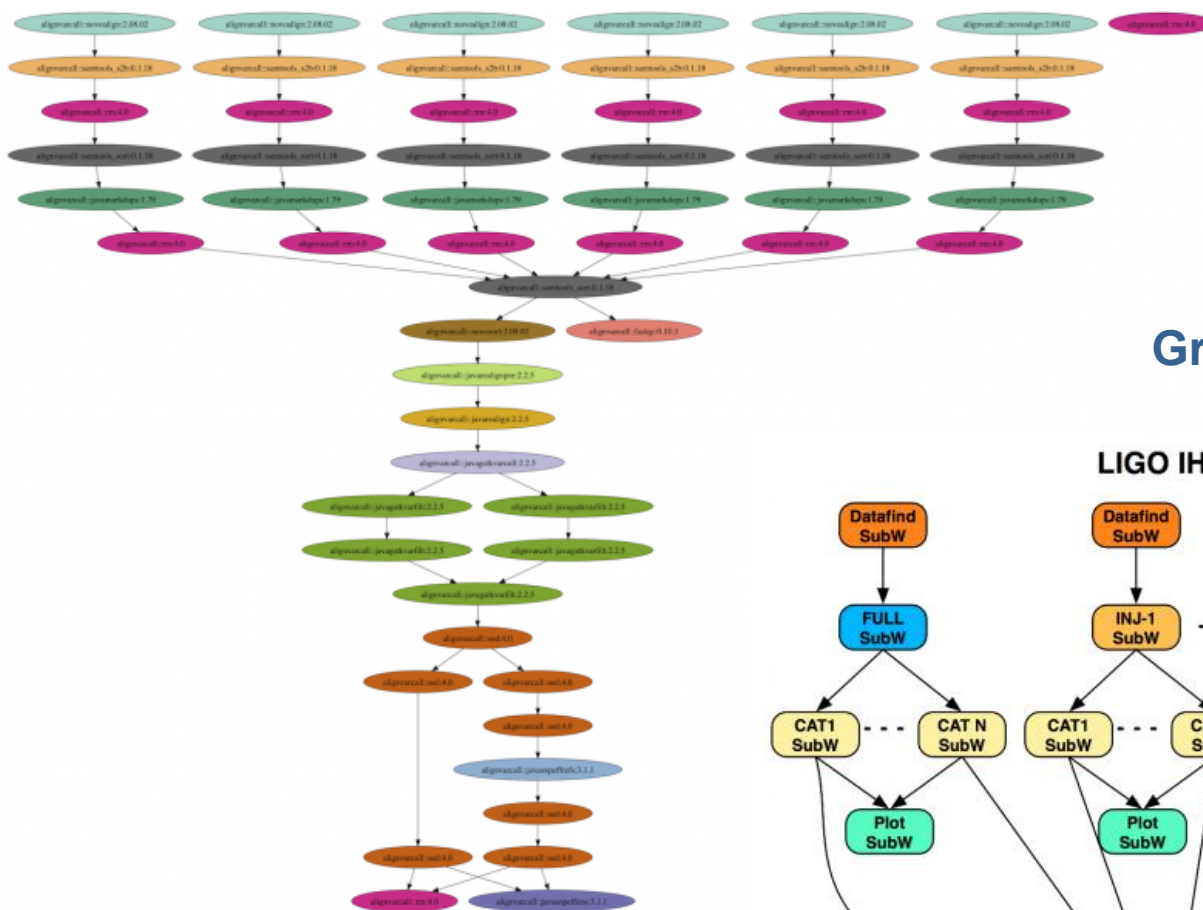
[bioinformatics](#) | [BLAST](#) |

[cheminformatics](#) | [data integration](#)

Outline


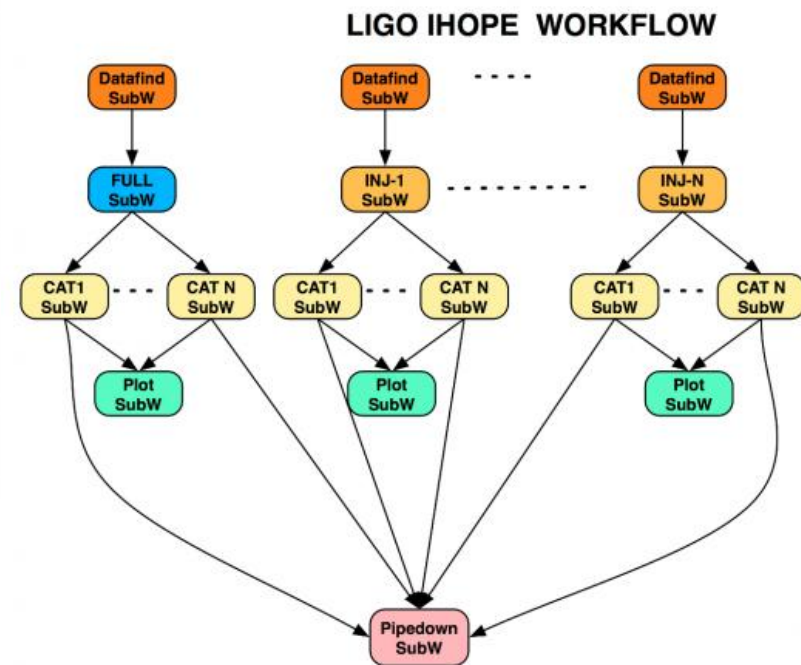
- Scientific Workflows and Application Examples
- Workflow Management Systems
- Pegasus and its features
- Conclusions





Some workflows are structurally complex

Gravitational-wave physics

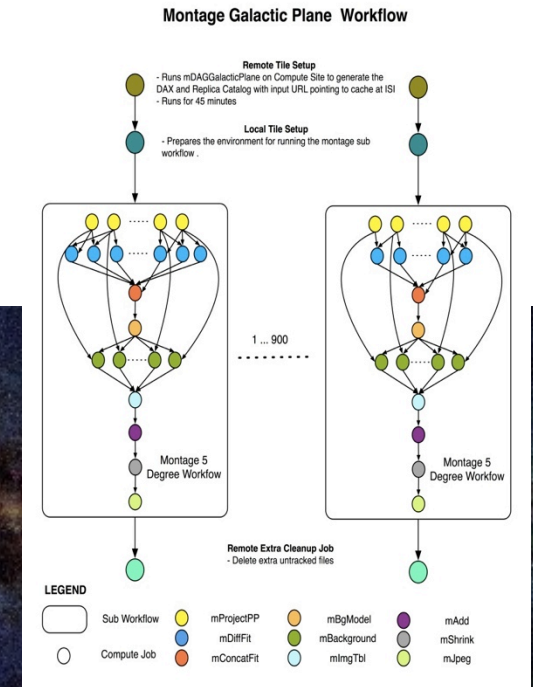


Genomics Workflow

LEGEND



Some workflows are large-scale and data-intensive




John Good (Caltech)

▪ Montage Galactic Plane Workflow

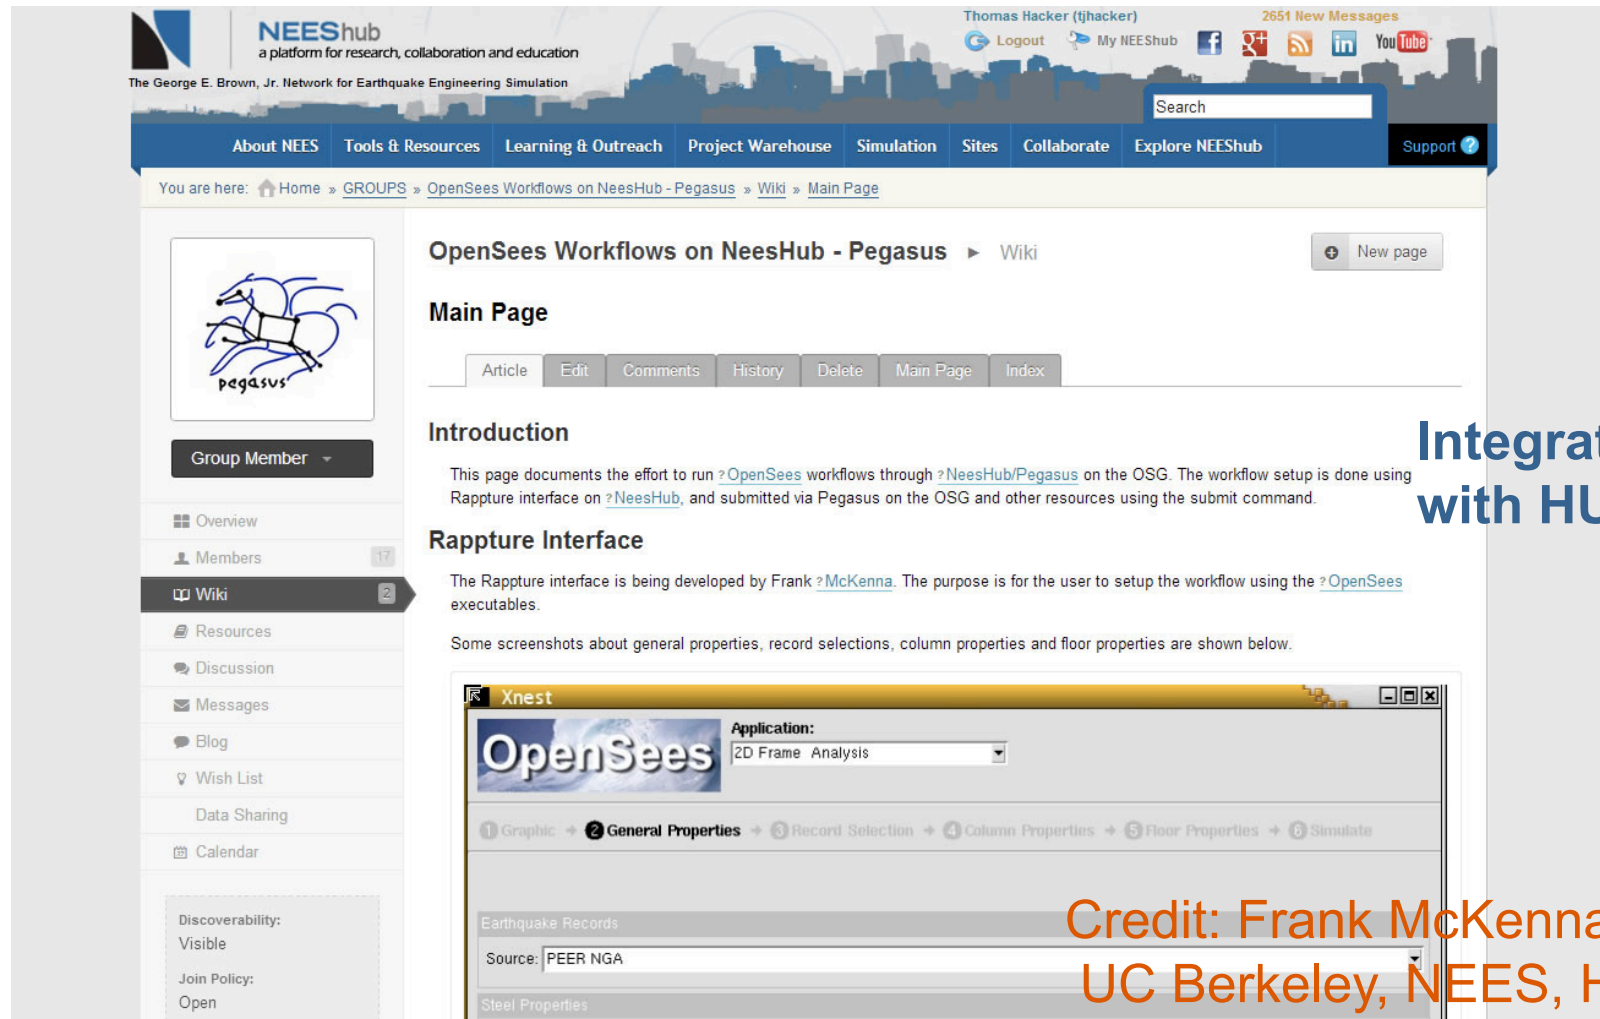
- 18 million input images (~2.5 TB)
- 900 output images (2.5 GB each, 2.4 TB total)
- 10.5 million tasks (34,000 CPU hours)

} × 17

▪ Need to support hierarchical workflows and scale




Sometimes you want to “hide” the workflow





The screenshot shows the NEEShub website interface. At the top, there's a header with the NEEShub logo and navigation links. Below the header, a search bar and a navigation menu are visible. The main content area displays the 'OpenSees Workflows on NeesHub - Pegasus' Wiki page. The page has a sidebar with links like Overview, Members, Wiki, Resources, Discussion, Messages, Blog, Wish List, Data Sharing, and Calendar. The main content area includes an introduction, a 'Rappture Interface' section, and a screenshot of the Rappture application window. The Rappture window shows the 'OpenSees' application with a dropdown menu set to '2D Frame Analysis' and a progress bar indicating the current step in the workflow.

Integration with HUBzero


Credit: Frank McKenna
UC Berkeley, NEES, HUBzero



Sometimes the environment is complex



Work definition



Local Resource

Campus Cluster

XSEDE


NERSC

ALCF


Open Science Grid

FutureGrid

Amazon Cloud



Sometimes the environment is complex



data

Work definition

Workflow
Management
System

Local Resource

work

Campus Cluster

XSEDE


NERSC

ALCF


Open Science Grid

FutureGrid

Amazon Cloud



Sometimes the environment is complex




data

Campus Cluster

XSEDE

Montage Galactic Plane



You don't want to recode your workflow

Work definition

Workflow
Management
System

Local Resource


work

ALCF

Open Science Grid

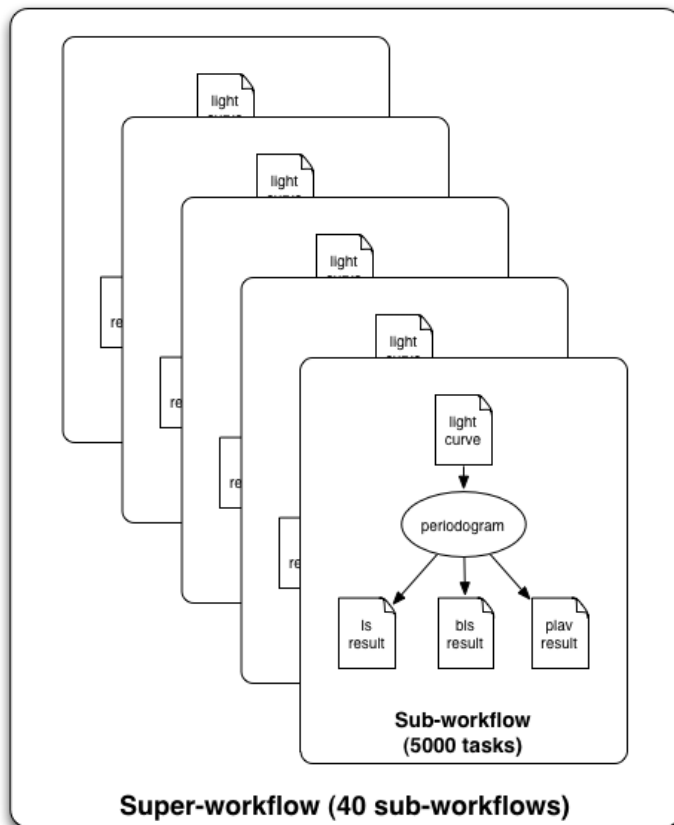
FutureGrid

Amazon Cloud




Sometimes the environment is just not exactly right

Single core workload




Cray XT System Environment /
ALPS / aprun

- Designed for MPI codes



Workflow Management


- You may want to use different resources within a workflow or over time
 - Need a high-level workflow specification
 - Need a planning capability to map from high-level to executable workflow
 - Need to manage the task dependencies
 - Need to manage the execution of tasks on the remote resources
- Need to provide scalability, performance, reliability



Pegasus

Workflow Management System (est. 2001)

- A collaboration between USC and the Condor Team at UW Madison (includes DAGMan)
- Maps a resource-independent “abstract” workflow onto resources and executes the “concrete” workflow
- Used by a number of applications in a variety of domains
- Provides reliability—can retry computations from the point of failure
- Provides scalability—can handle large data and many computations (kbytes-TB of data, $1-10^6$ tasks)
- **Infers data transfers, restructures workflows for performance**
- Automatically captures provenance information
- Can run on resources distributed among institutions, laptop, campus cluster, Grid, Cloud




Generating executable workflows


APIs for
workflow
specification
(DAX---
DAG in XML)

Java, Perl, Python

(DAX)



Abstract Workflow




Executable Workflow

LEGEND

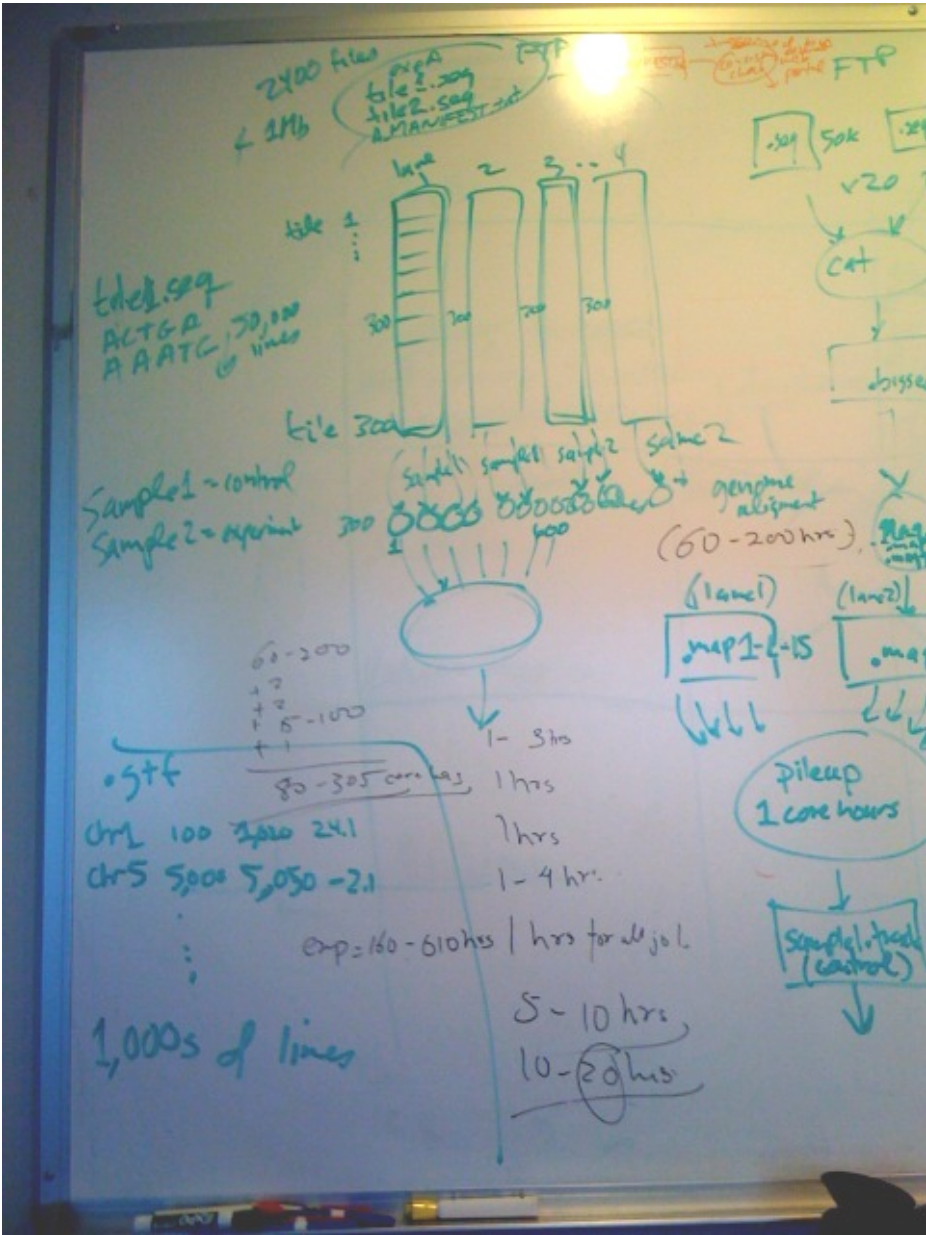
- Unmapped Job
- Compute Job mapped to a site
- Stage-in Job
- Stage-Out Job
- Registration Job
- Make Dir Job
- Cleanup Job

Planning Process

- Assume data may be distributed in the environment
- Assume you may want to use local and/or remote resources
- Pegasus needs information about the environment
 - data, executables, execution and data storage sites
- Pegasus generates an executable workflow
- Data transfer protocols
 - Gridftp, Condor I/O, HTTP, scp, S3, iRods, SRM, FDT (partial)
- Scheduling to interfaces
 - Local, Gram, Condor, Condor-C (for remote Condor pools), via Condor Glideins – PBS, LSF, SGE




Time to solution ~ 3 months





Execution on USC resources

Number of jobs





Advanced features

- Performs data reuse
- Registers data in data catalogs—supports checkpointing
- Manages storage—deletes data no longer needed
- Can cluster tasks together for performance
- Can manage complex data architectures (shared and non-shared filesystem, distributed data sources)
- Different execution modes which leverage different computing architectures (Condor pools, HPC resources, etc..)




Workflow Reduction


Data Reuse Workflow-level checkpointing



Abstract Workflow




File f.d exists somewhere.
Reuse it.
Mark Jobs D and B to delete




Delete Job D and Job B

LIGO Workflows:
185,000 nodes,
466,000 edges
10 TB of input
data,
1 TB of output
data.




Workflow Restructuring to improve application performance

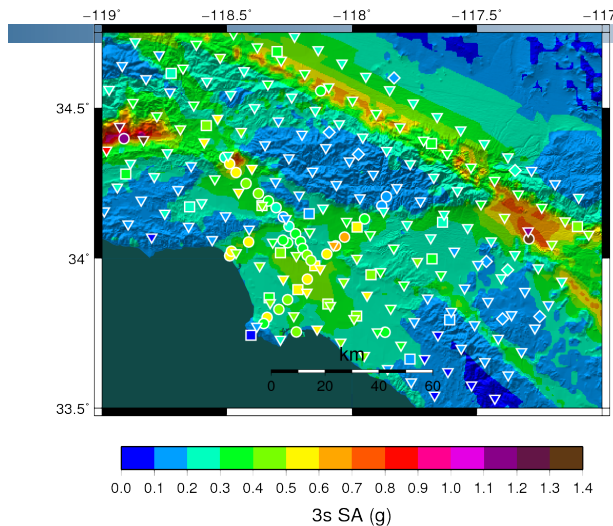
- **Cluster small running jobs together to achieve better performance**
- **Why?**
 - Each job has scheduling overhead – need to make this overhead worthwhile
 - Ideally users should run a job on the grid/cloud that takes at least 10/30/60/? minutes to execute
 - Clustered tasks can reuse common input data – less data transfers



Level-based clustering



Southern California Earthquake Center




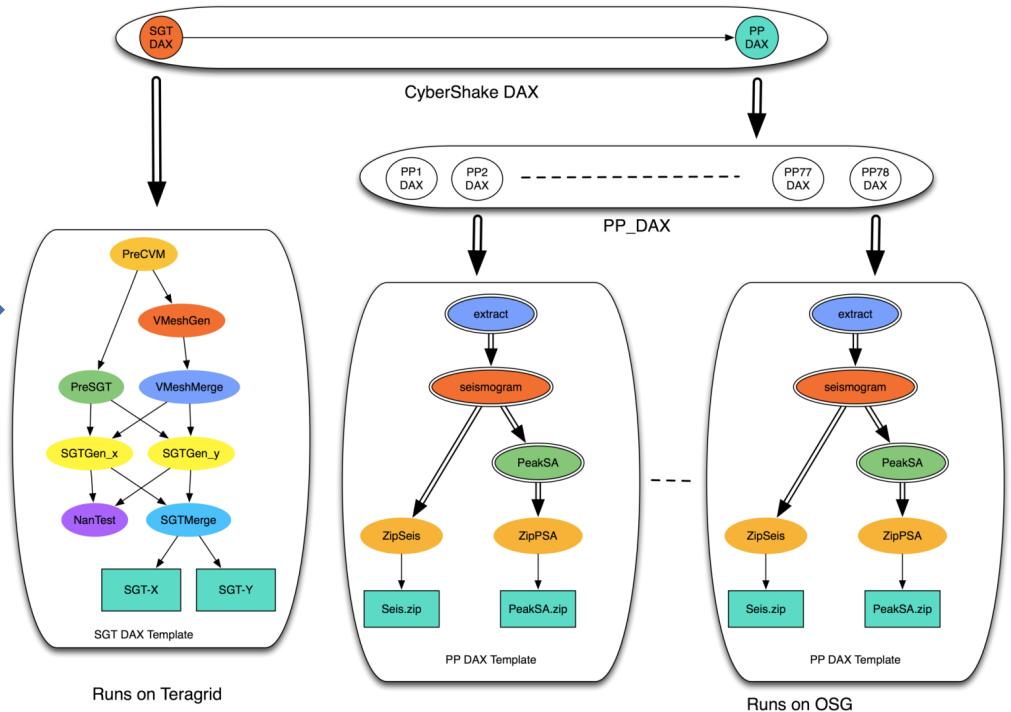
CyberShake PSHA Workflow

❖ Description



- ❖ Builders ask seismologists: “What will the peak ground motion be at my new building in the next 50 years?”
- ❖ Seismologists answer this question using Probabilistic Seismic Hazard Analysis (PSHA)

239 Workflows

- Each site in the input map corresponds to one workflow
- Each workflow has:
 - ❖ 820,000 tasks




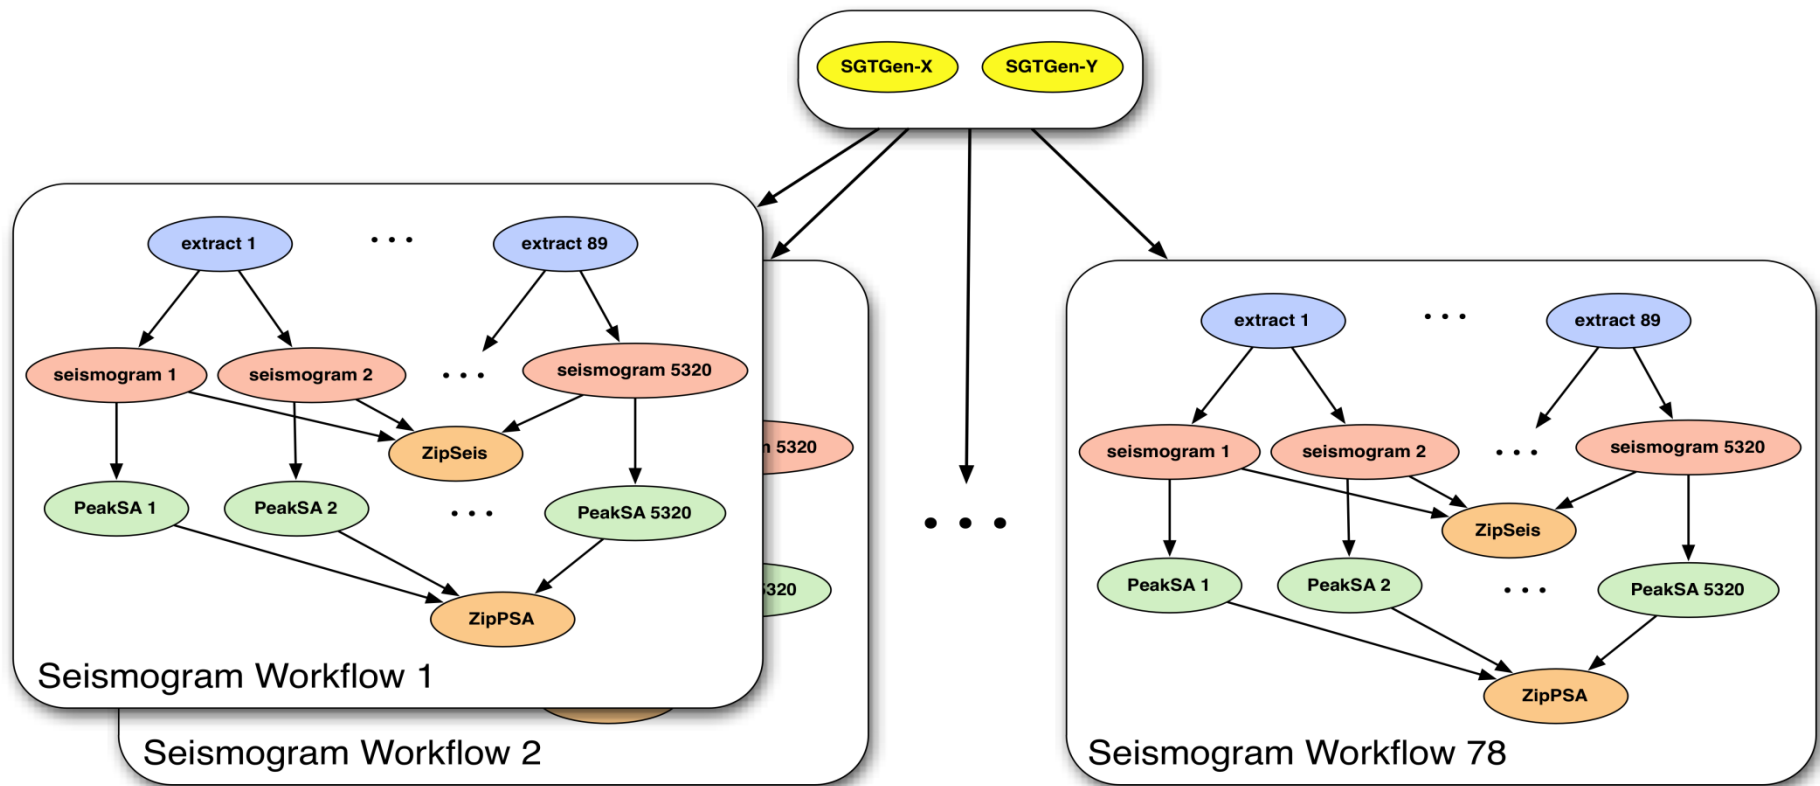
Workflows have different computational needs




Pegasus-MPI-Cluster

- A master/worker task scheduler for running fine-grained workflows on batch systems
- Runs as an MPI job
 - Uses MPI to implement master/worker protocol
- Works on most HPC systems
 - Requires: MPI, a shared file system, and fork()
- Allows sub-graphs of a Pegasus workflow to be submitted as monolithic grid jobs to remote resources





Enables earthquake scientists (SCEC) to run post-processing (single core) computations on new architectures (Kraken)




Workflow Monitoring - Stampede

- **Leverage Stampede Monitoring framework with DB backend**
 - Populates data at runtime. A background daemon monitors the logs files and populates information about the workflow to a database
 - Stores workflow structure, and runtime stats for each task.
- **Tools for querying the monitoring framework**
 - **pegasus-status**
 - Status of the workflow
 - **pegasus-statistics**
 - Detailed statistics about your finished workflow

Type	Succeeded	Failed	Incomplete	Total	Retries	Total+Retries
Tasks	135002	0	0	135002	0	135002
Jobs	4529	0	0	4529	0	4529
Sub-workflows	2	0	0	2	0	2

Workflow wall time : 13 hrs, 2 mins, (46973 secs)
Workflow cumulative job wall time : 384 days, 5 hrs, (33195705 secs)
Cumulative job walltime as seen from submit side : 384 days, 18 hrs, (33243709 secs)


Collaboration with Dan Gunter and Taghrid Samak, LBNL






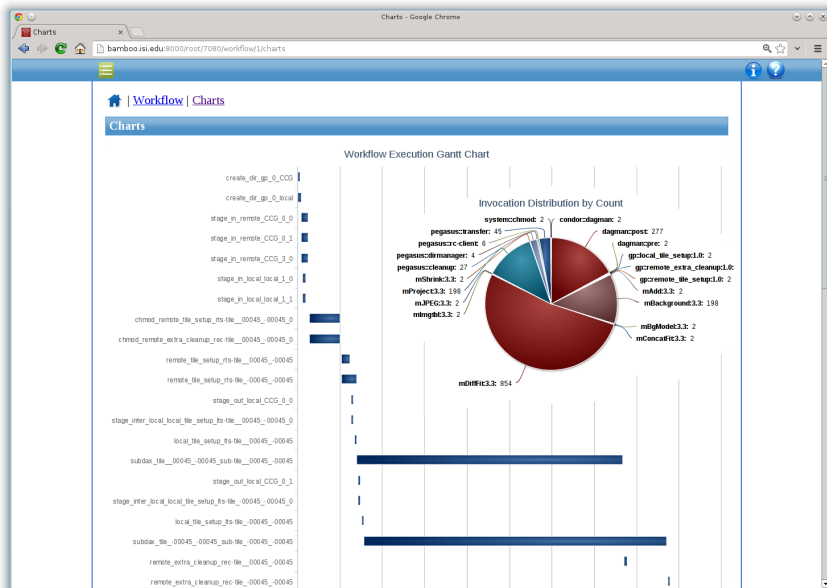

Workflow Monitoring Dashboard: pegasus-dashboard

- **A python based online workflow dashboard**
 - Queries the STAMPEDE database
- **Lists all the user workflows on the home page and are color coded.**
 - Green indicates a successful workflow,
 - Red indicates a failed workflow
 - Blue indicates a running workflow
- **Explore Workflow and Troubleshoot (Workflow Page)**
 - Has identifying metadata about the workflow
 - Tabbed interface to
 - List of sub workflows
 - Failed jobs
 - Running jobs
 - Successful jobs.

Collaboration with Dan Gunter and Taghrid Samak, LBNL




Workflow Monitoring Dashboard – pegasus-dashboard




Benefits of Pegasus

- **Provides Support for Complex Computations**
 - Can be hidden behind a portal
- **Portability / Reuse**
 - User created workflows can easily be run in different environments without alteration (XSEDE, OSG, FutureGrid, Amazon)
- **Performance**
 - The Pegasus mapper can reorder, group, and prioritize tasks in order to increase the overall workflow performance
- **Scalability**
 - Pegasus can easily scale both the size of the workflow, and the resources that the workflow is distributed over.




Benefits of Pegasus

- **Provenance**
 - Performance and provenance data is collected in a database, and the data can be summaries with tools such as pegasus-statistics, pegasus-plots, or directly with SQL queries.
- **Reliability**
 - Jobs and data transfers are automatically retried in case of failures. Debugging tools such as pegasus-analyzer helps the user to debug the workflow in case of non-recoverable failures.
- **Analysis tools**



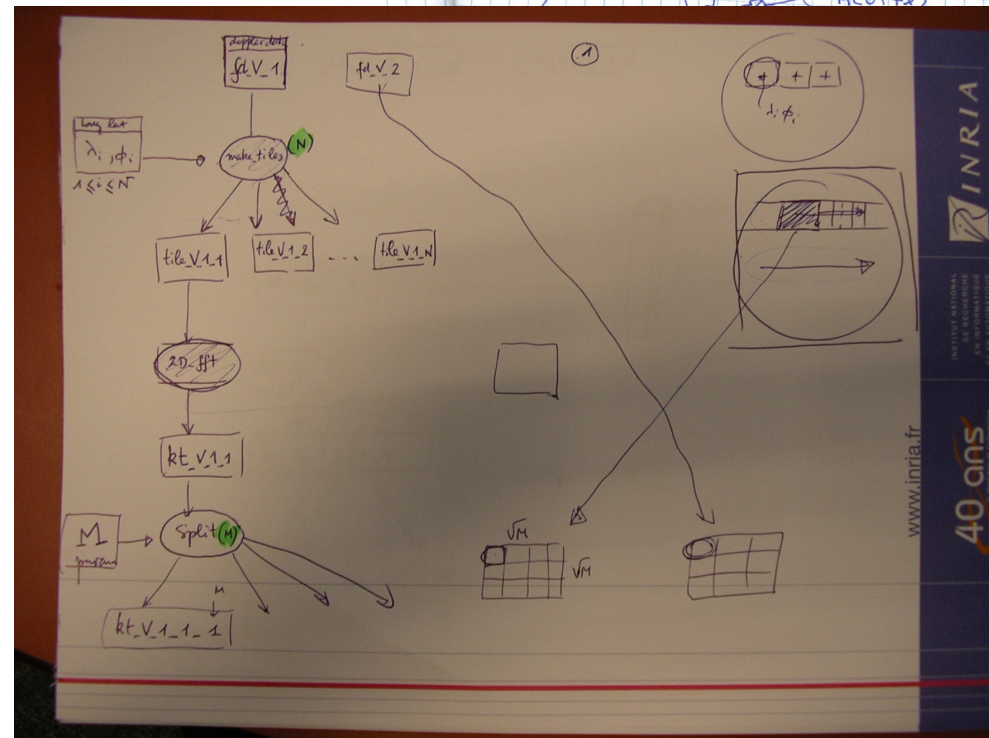
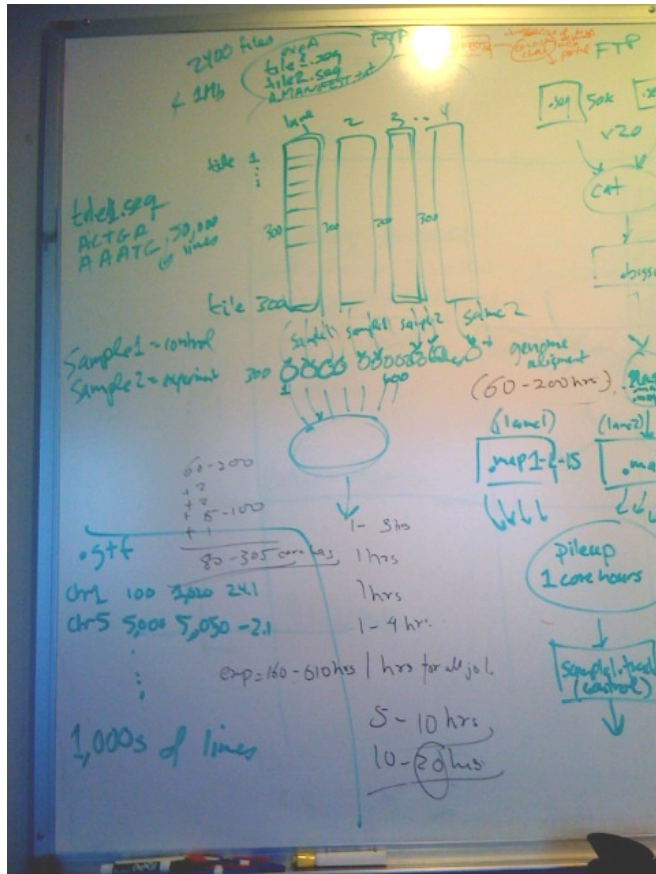
How to get started?

- Pegasus: <http://pegasus.isi.edu>
- Tutorial and documentation:
<http://pegasus.isi.edu/wms/docs/latest/>
- Virtual Machine with all software and examples
<http://pegasus.isi.edu/downloads>
- Take look at some Pegasus applications:
<http://pegasus.isi.edu/applications>
- Support: pegasus-users@isi.edu
pegasus-support@isi.edu



If you get stuck...

And you can draw....




We can help you!

Workflows Can Address Issues Critical to Scientists

- **Reproducibility** of scientific analyses and processes is at the core of the scientific method
 - Workflows give us the opportunity to provide reproducibility
- Scientists consider the “capture and generation of **provenance** information as a critical part of the workflow-generated data”
- “**Sharing** workflows is an essential element of education, and acceleration of knowledge dissemination.”[1]

[1] Y. Gil, E. Deelman et al, **Examining the Challenges of Scientific Workflows**.
IEEE Computer, 12/2007



Acknowledgments

- Pegasus Team
- Condor Team
- Colleagues at LBNL
- Funding agencies: NSF, NIH
- and everybody who uses Pegasus

