



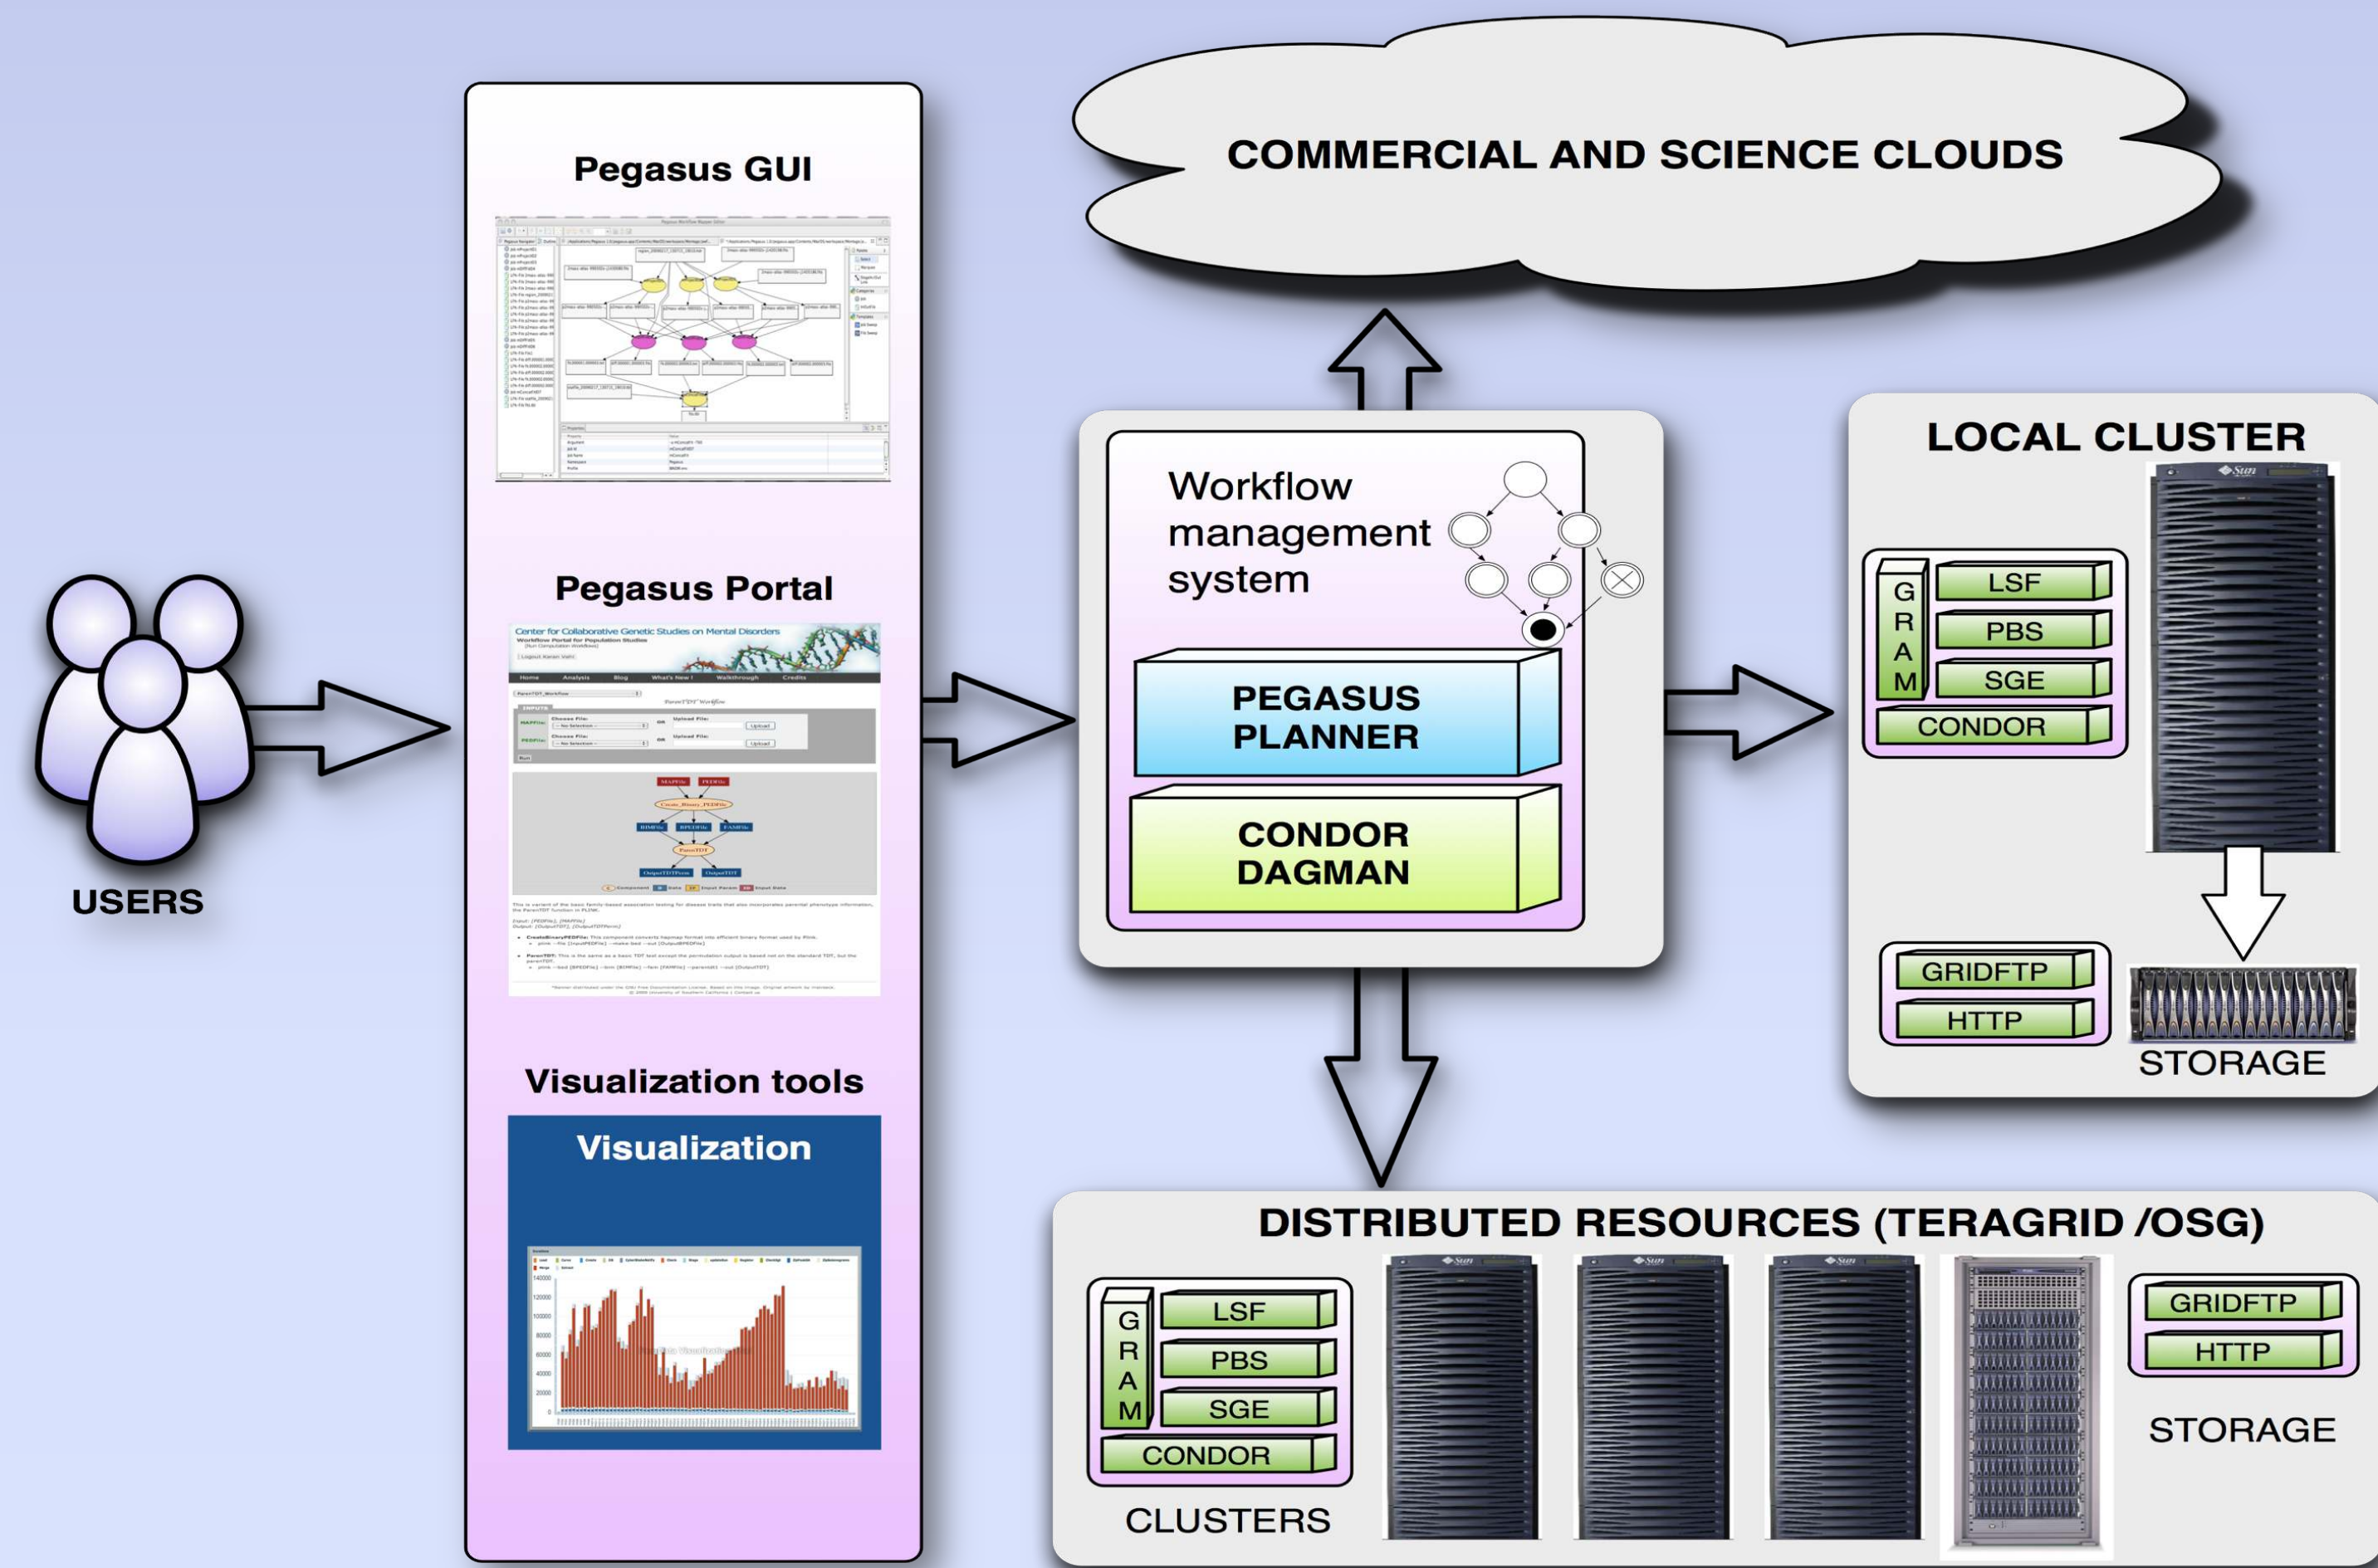
Pegasus WMS: Enabling Large Science On National CyberInfrastructure

Karan Vahi¹, Ewa Deelman¹, Gaurang Mehta¹, Mats Rynge¹, Fabio Silva¹, Prasanth Thomas¹, Jens Vöckler¹
Bruce Berriman²

¹ University of Southern California's Information Sciences Institute, ²Caltech IPAC

Pegasus WMS: <http://pegasus.isi.edu>

- Pegasus is a system for mapping and executing abstract application workflows over a range of execution environments.
- Workflows can run on a single system or across a heterogeneous set of resources.
- Pegasus can run workflows comprising between 1 to 10⁶ tasks.
- Pegasus WMS consists of three main components: the Pegasus mapper, Condor DAGMan, and the Condor schedd.
- The mapping of tasks to the execution resources is done by the mapper based on information derived for static and/or dynamic sources.
- The output is an executable workflow (also called the concrete workflow) that can be executed over a variety of resources.
- Pegasus adds and manages data transfer between the tasks as required.
- DAGMan takes this executable workflow and manages the dependencies between the tasks and releases them to the Condor schedd for execution.



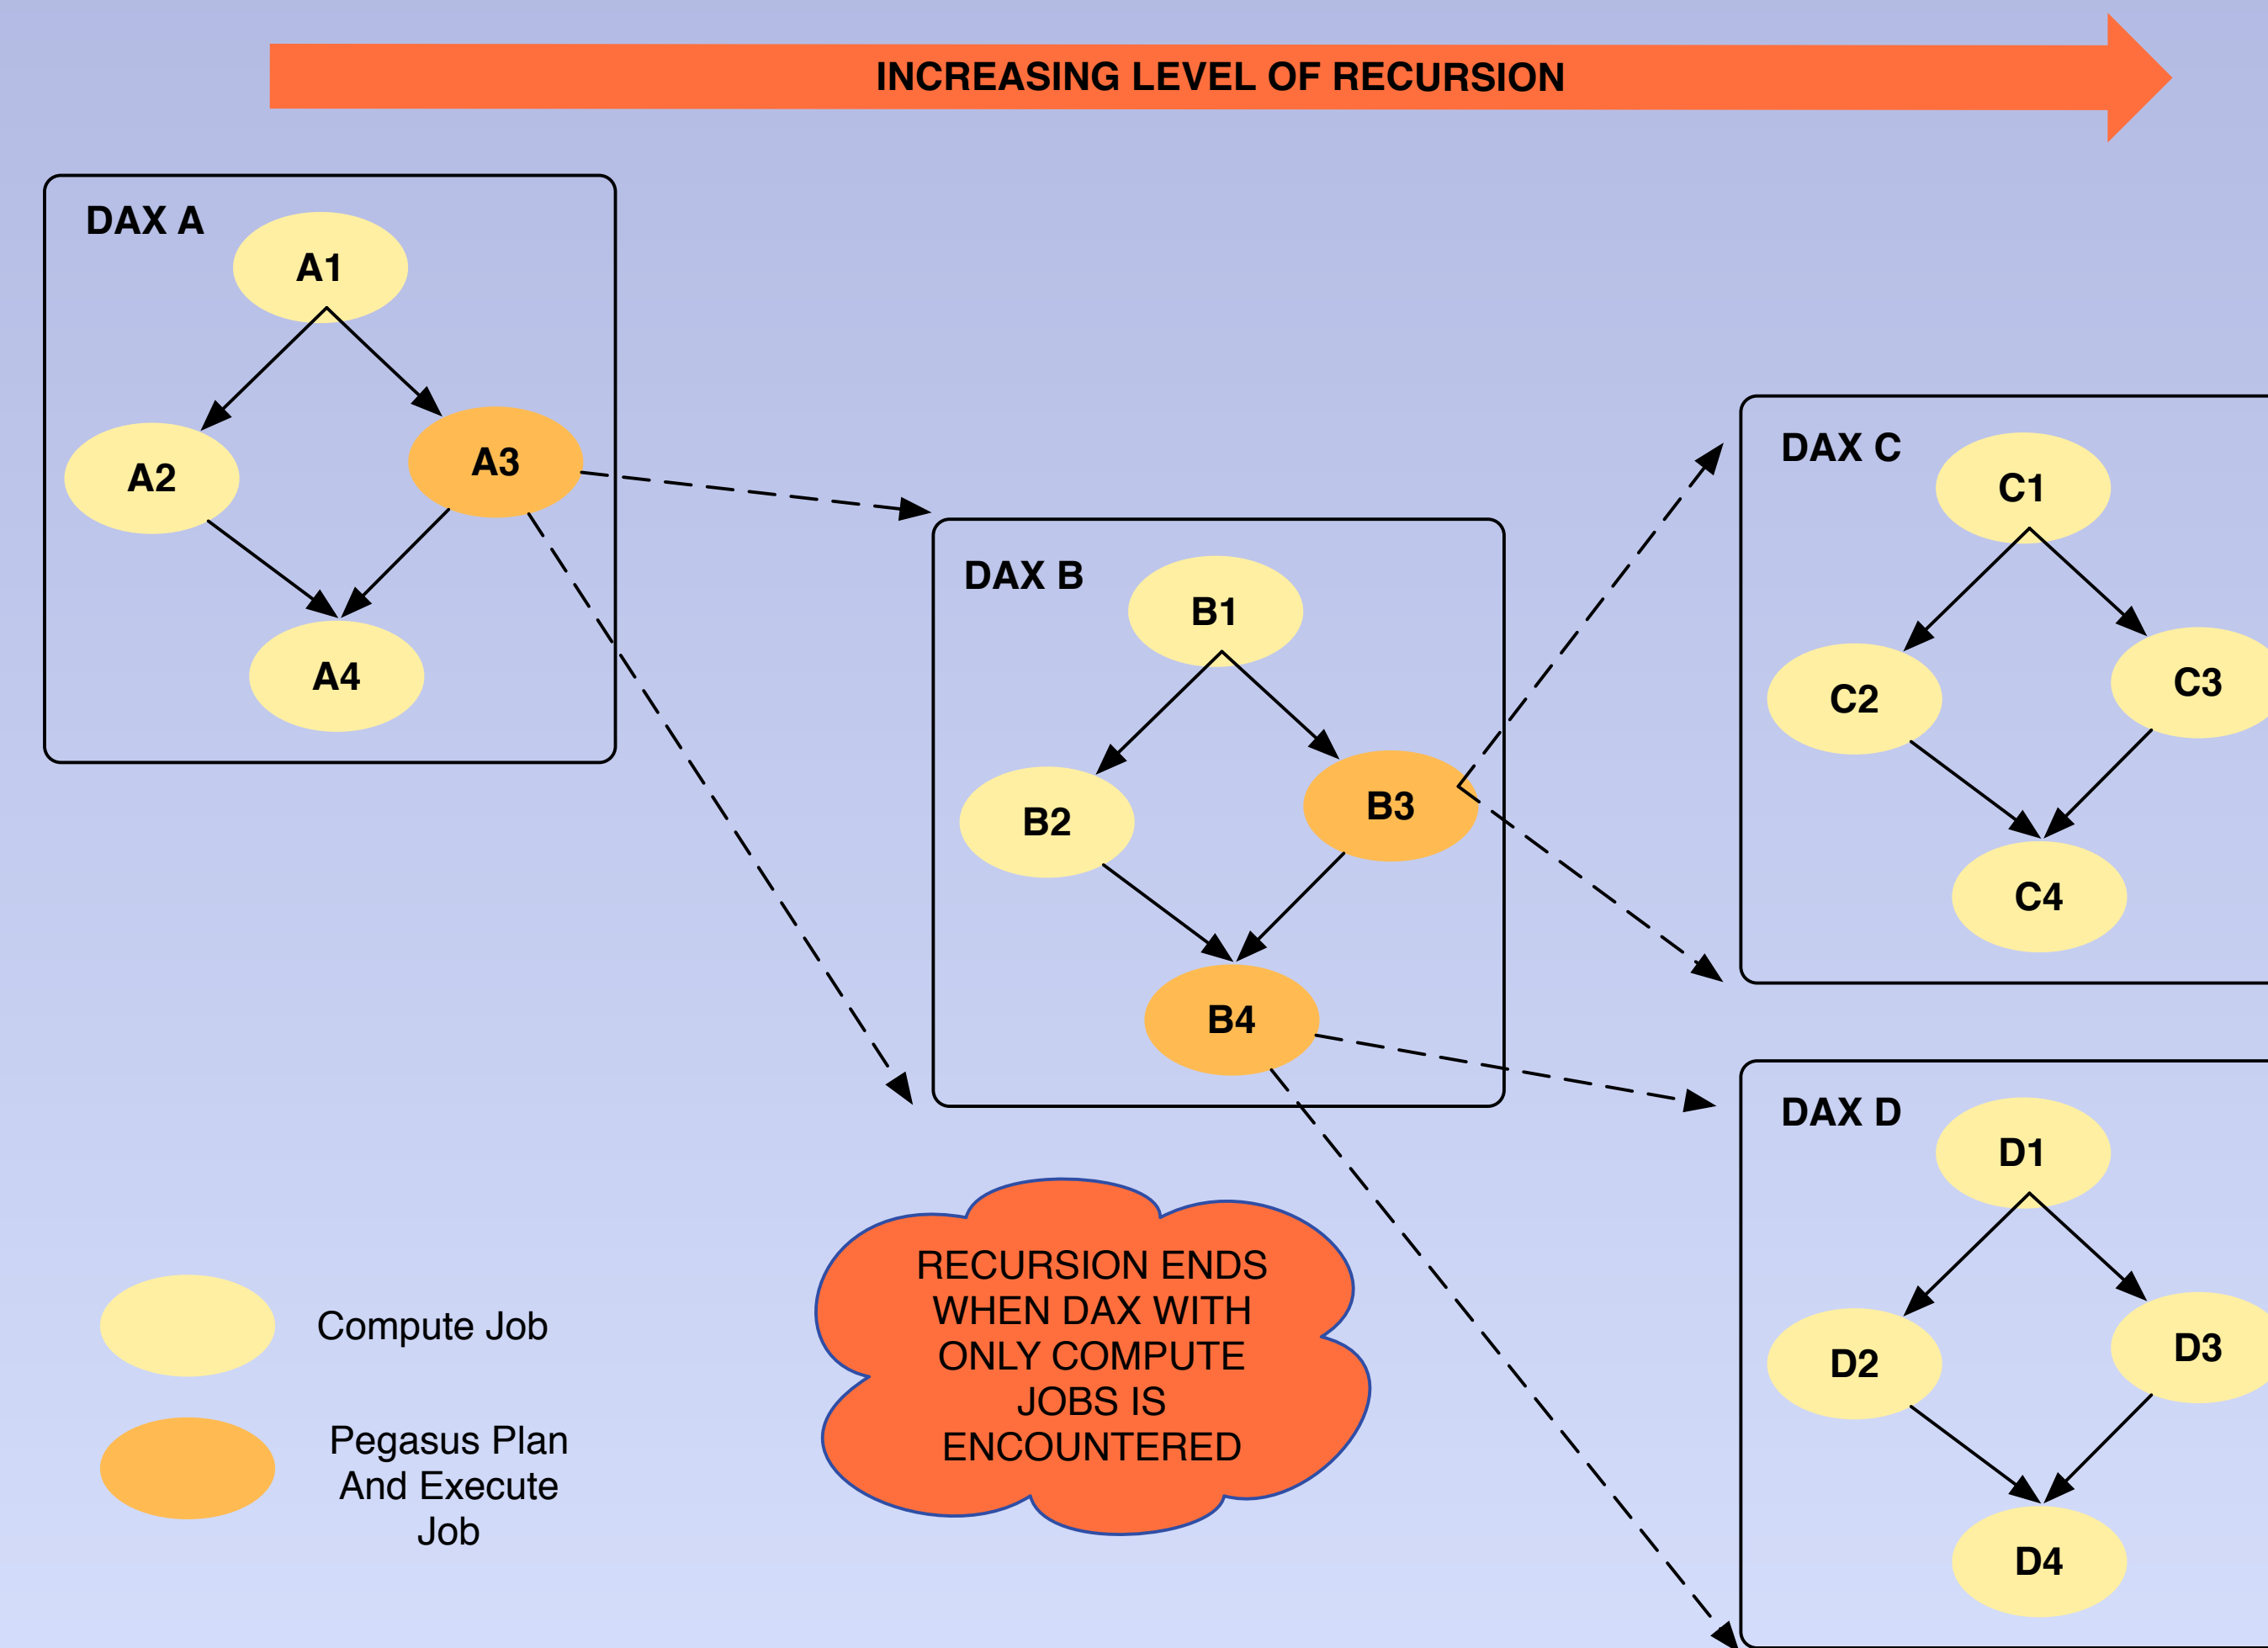
Pegasus Features

- Expresses workflow using simple JAVA, Perl or Python API
- Execute complex computations efficiently and reliably
- Retry computation in case of errors
- Express an analysis once, execute it on multiple resources simultaneously multiple times.
- Submit to different styles of resources (Clouds, TeraGrid, OSG, Campus Grids, Clusters, Workstation)
- Data Reuse Capabilities on basis of previously generated outputs.
- Stores provenance of data used, produced and which software was used with what parameters.
- Monitoring and Debugging tools to debug large scale workflows.
- Upcoming support for job and workflow notifications.
- Clustering of Tasks and optimized data transfers
- Supports variety of protocols (gridftp, scp , http, s3, irods) for data staging.
- Automatic data cleanup.
- Shell Code Generator to generate shell script instead of Condor DAGMan and submit files.

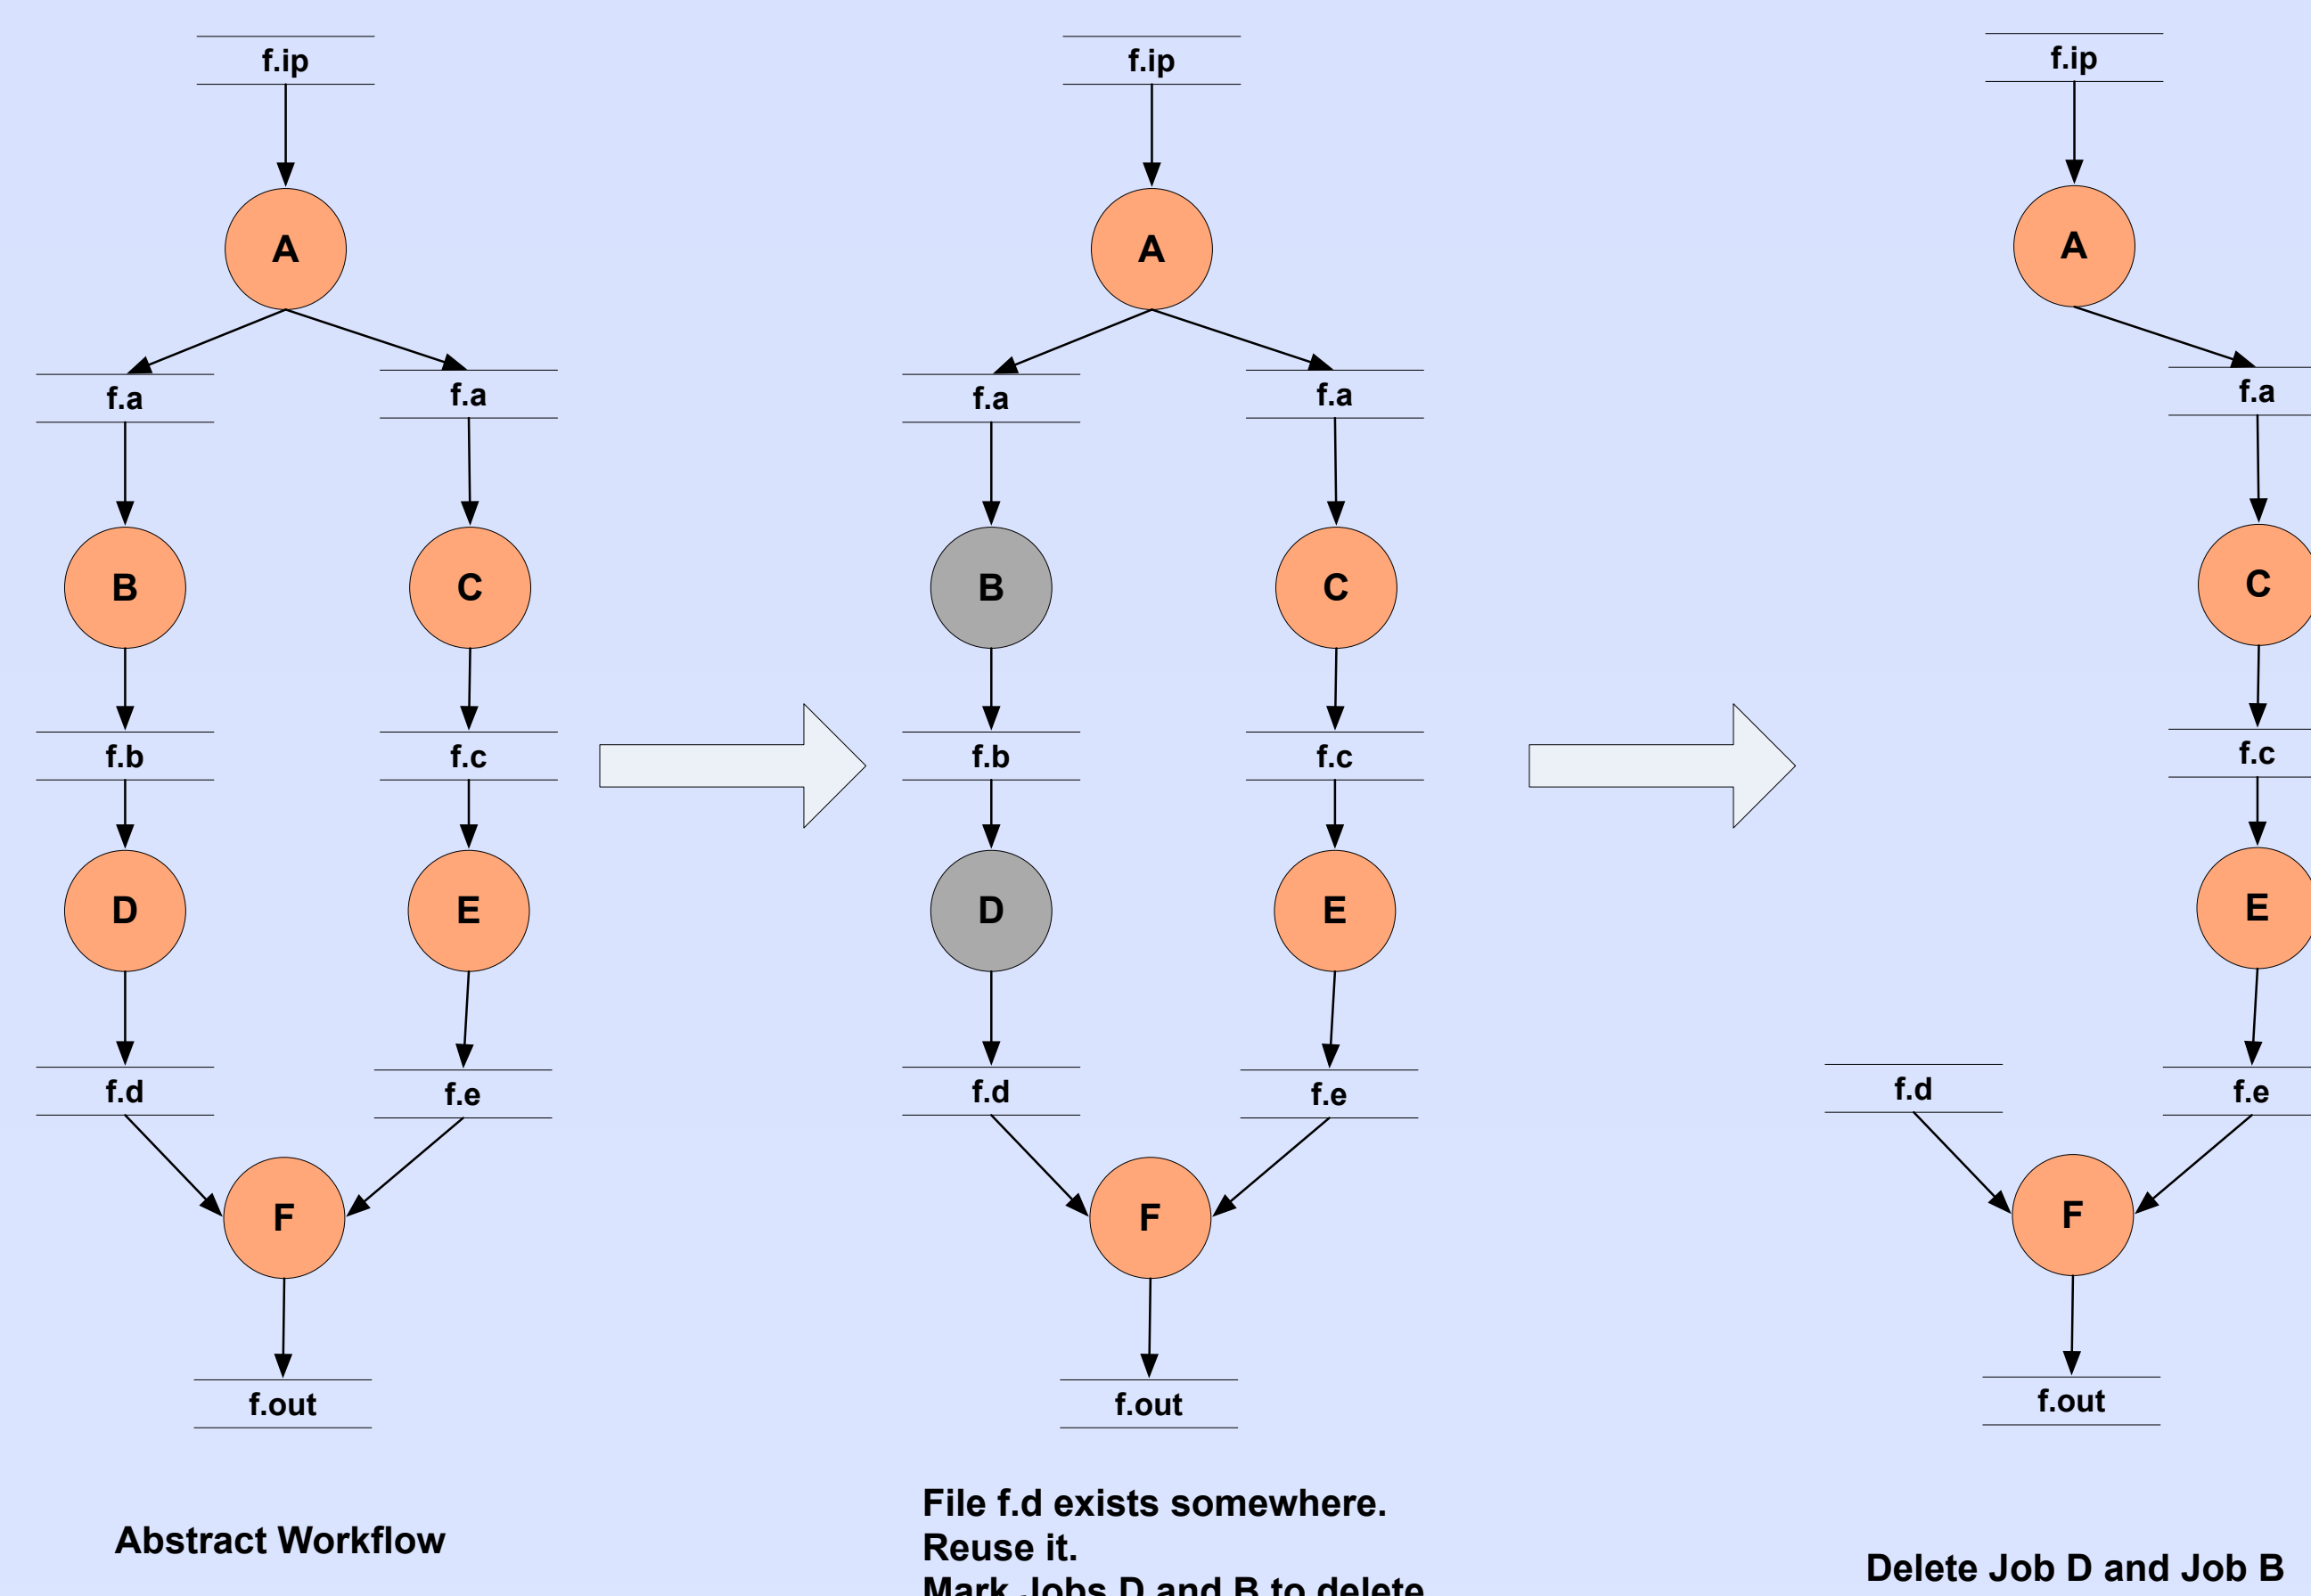
Hierarchal Workflows

- Multiple Level of Nested Workflows.
 - Workflow can consists of Tasks, Abstract Workflows (DAX) or DAGs.
- Allows users to scale workflows to order of 10⁶ tasks.
- Better management and scalability at Condor level.
- Each sub workflow is planned only when it is ready for execution.

HIERARCHAL WORKFLOWS



Data Reuse Capabilities



- Workflow output files are tracked in a Replica Catalog
- If a subsequent workflow refers to output files already generated, Pegasus will prune the workflow.
- Cascades the deletion of jobs in the workflow upwards wherever appropriate.
- Pegasus will stage-in the previously generated files instead of re-generating them.
 - Can be turned off i.e. all jobs in the workflow will be executed.
- Can reduce the computing time drastically for overlapping analysis.

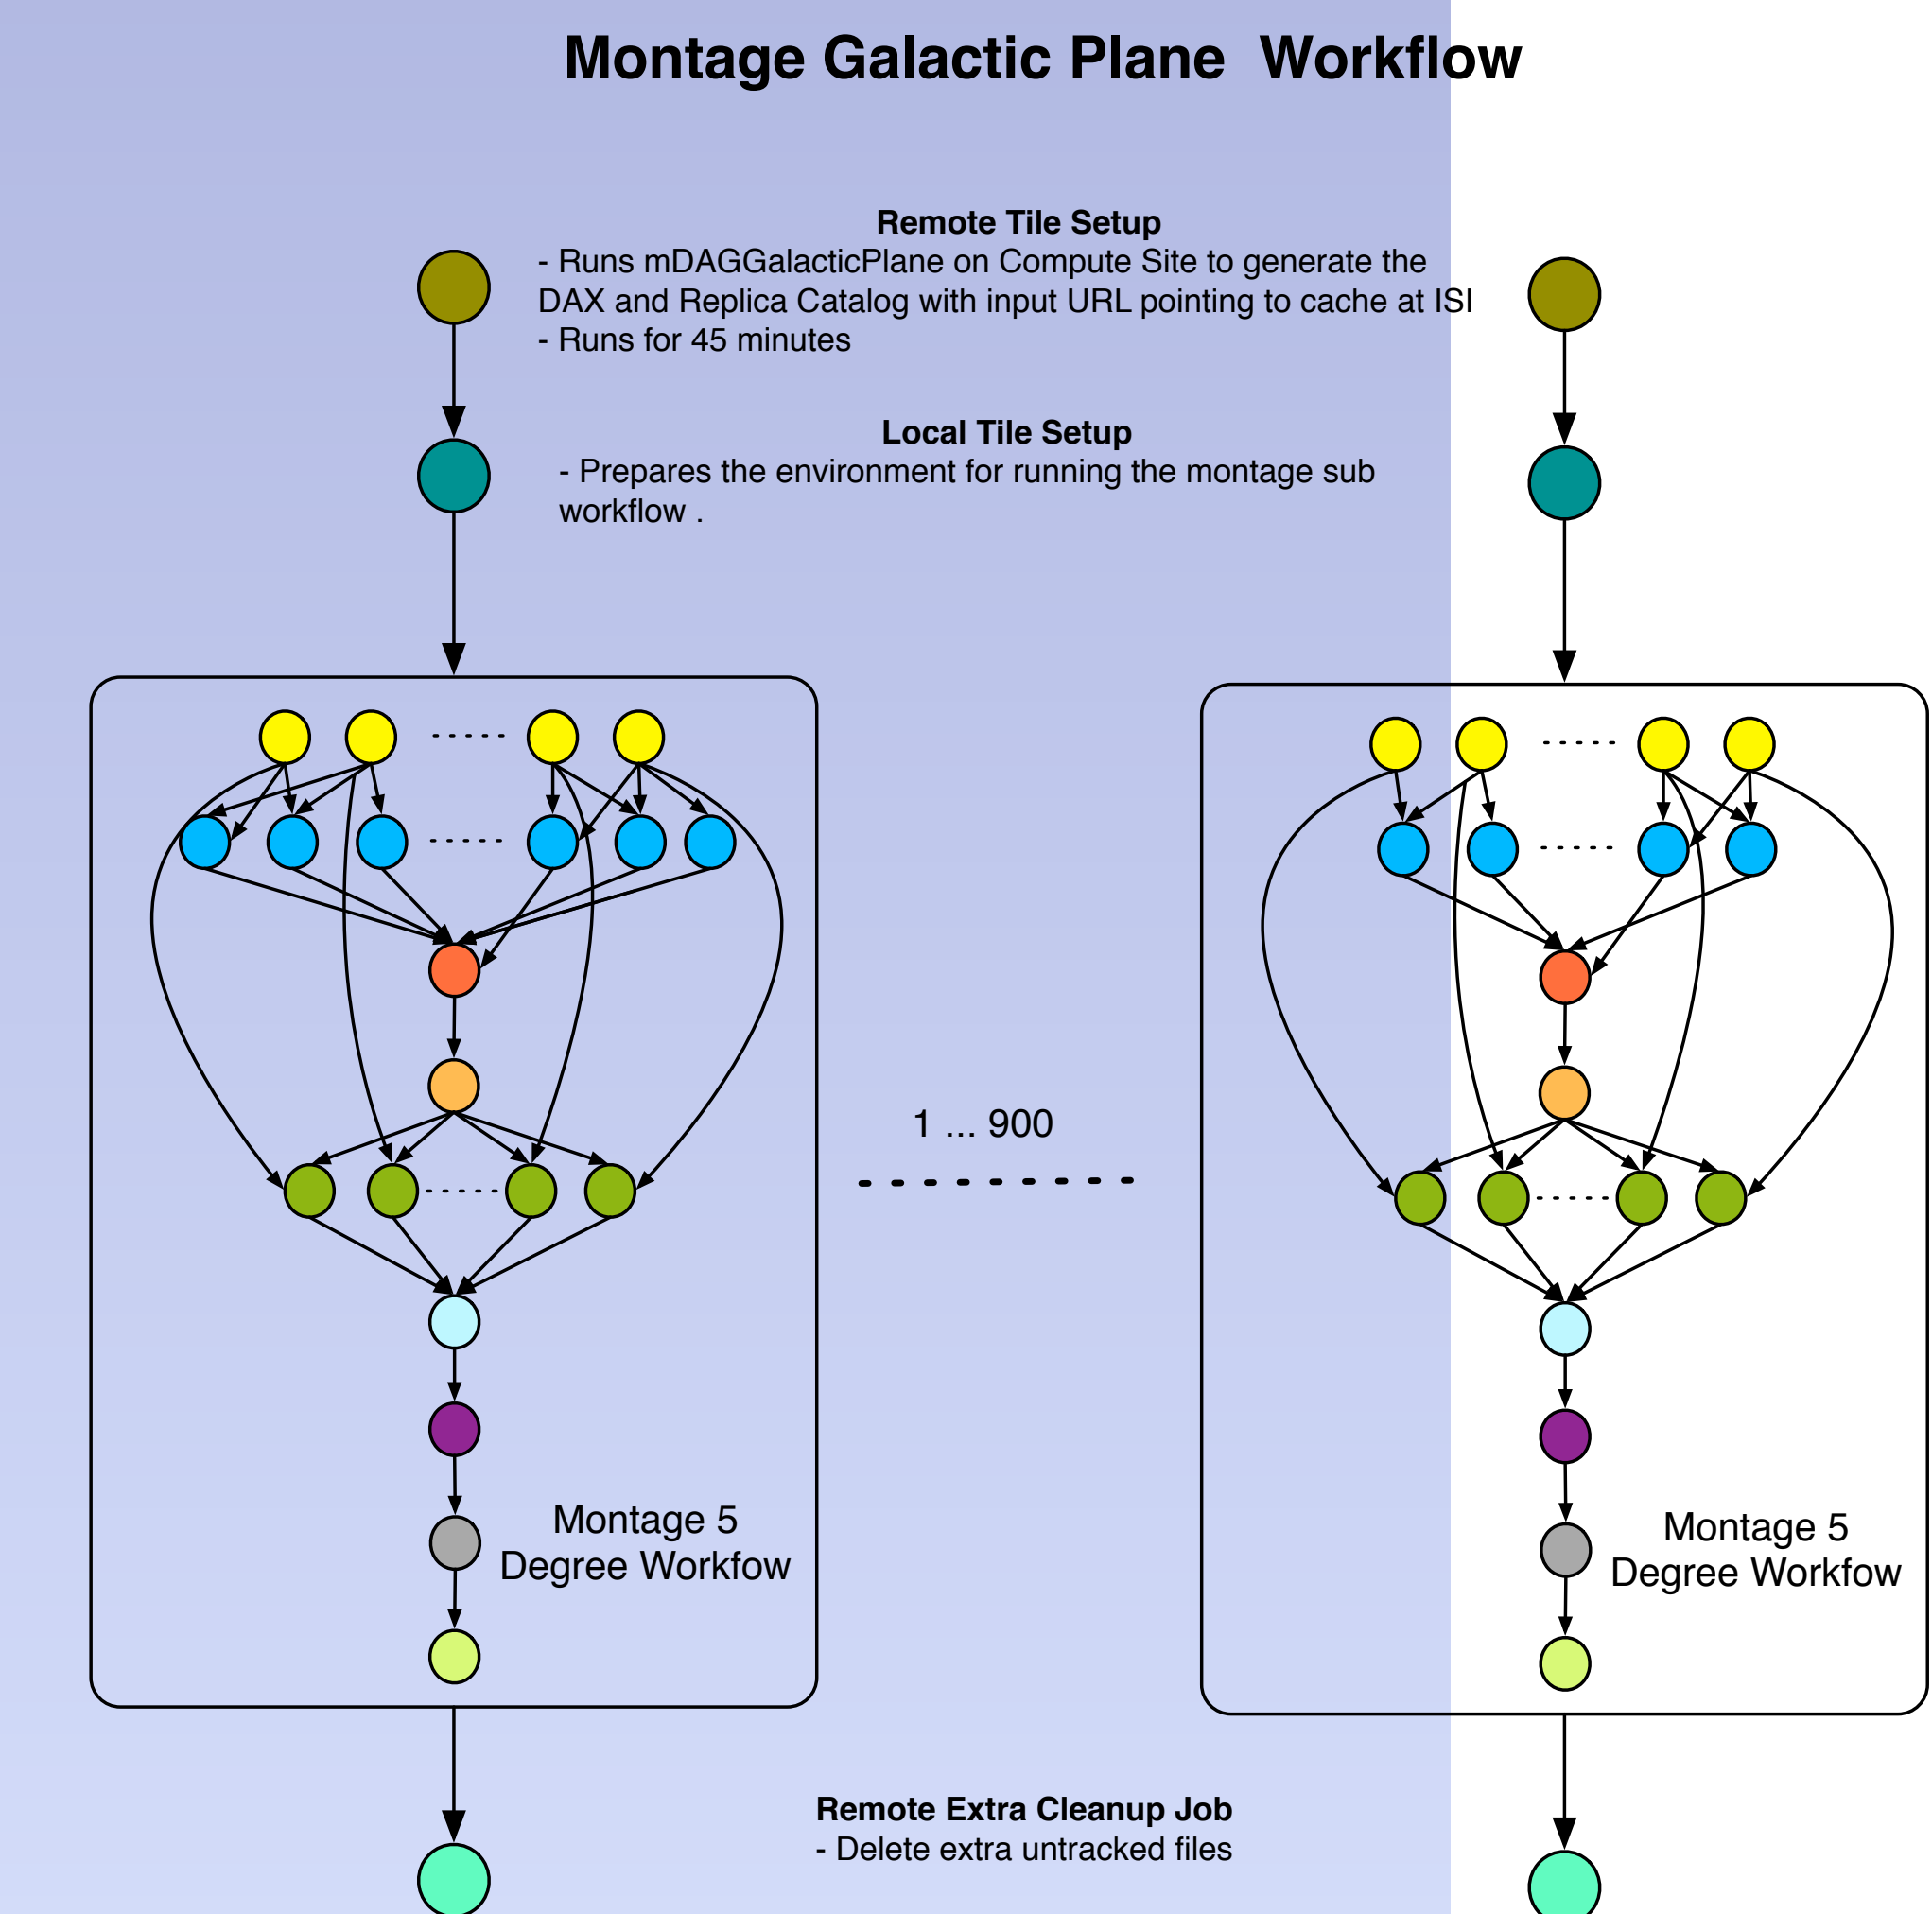
Acknowledgements

- Condor:** Miron Livny and Kent Wenger

Large Scale Astronomy Workflows on TeraGrid

- Compute and Data Intensive Workflows
- Utilize Pegasus Cleanup capabilities to clean up data from TeraGrid sites as workflows execute.
- Hierarchical Workflows with total number of jobs in 100's of thousands.

Montage and Galactic Plane



- Galactic Plane workflow is a workflow of Montage workflows
- Used to generate tiles 360°x40° around the galactic equator.
- A tile is 5°x5° with 1° overlap with its neighbors.
- Output datasets to be used in NASA sky and Google Sky.
- One workflow run for each of 17 bands (wavelengths).
- Each such workflows uses 3.5TB of input imagery(1.6mil files)
- Each workflow consumes 30K CPU hours and produces 900 tiles in FITS format - 4.7 TB of outputs
- Reference Run consisting of 230 tiles done on TACC Ranger

Proposed Runs

- Run Workflow corresponding to each of the 17 bands (wavelengths).
- Total Number of Data Files – 18 million.
- Potential Size of Data Output – 86 TB



NASA Star & Exoplanet DB Periodogram

- Used to find extra-solar planets using two methods:
 - Periodic dips in star brightness (luminosity) caused by planetary transits
 - Radial velocity “wobbles” caused by gravitational pull of planet that causes star’s spectrum to change
- Compute periodograms for 210K light curves released by Kepler project
- Use 3 algorithms (Lomb-Scargle, Box-fitting Least Squares, Binless phase-dispersion minimization) on each curve
- Ran entire dataset 3 times with 3 different sets of parameters

