How scientific workflows help automate science and what we can do better?

Ewa Deelman

University of Southern California, School of Engineering
Information Sciences Institute
deelman@isi.edu

This work is funded by NSF awards #2018074 and #1664162.

We live in a heterogeneous world

HETEROGENEOUS APPLICATIONS

HETEROGENEOUS RESOURCES

HETEROGENEOUS USERS

Users

Level of Cyberinfrastructure Expertise

- Large collaborations
- Medium size groups
- Individual PIs and their students
- Educators and students
- The public

MeerKAT dish antenna, Karoo

SKA in South Africa (artist rendition)

https://www.skatelescope.org/africa/

James Webb Telescope Image

Heterogeneous Applications

- Non-monolithic
  - Mix of single core, multi-core/single node, multi-node parallel applications
- Written in a variety of languages
- Multi-task with data and control dependencies
- Iterative
- Distributed and parallel
- Accessing instruments, acquiring large amounts of data
- Simulating large phenomena
- Processing large amount of data/modeling

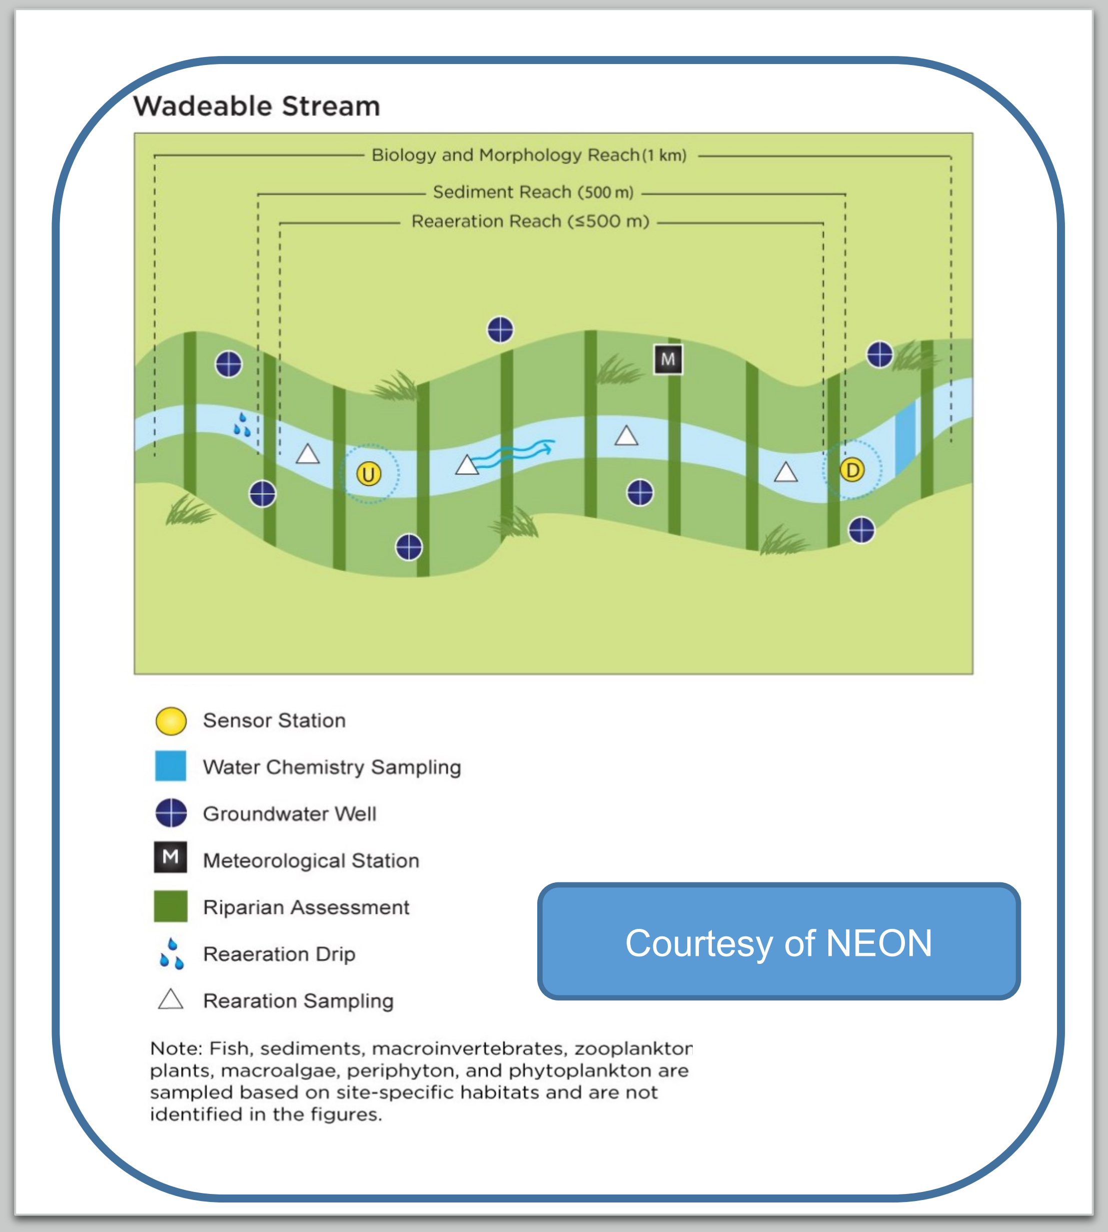

Heterogeneous resources

- Sensors
- Phone/tablets
- Instruments/drones/etc
- Edge devices
- Laptops
- Clusters
- Clouds
- HPC resources

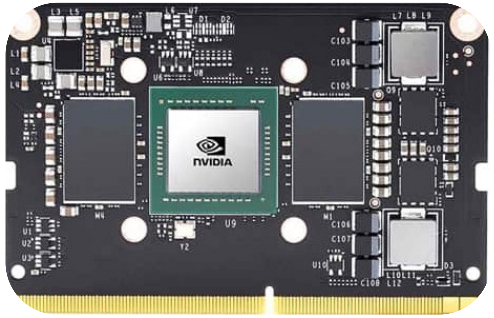

Jestson Nano

Frontier, ORNL

We started looking at supporting users and their applications in 2000

Virgo visit
September 2001

- https://www.ego-gw.it/public/about/welcome.aspx

LIGO circa 2001

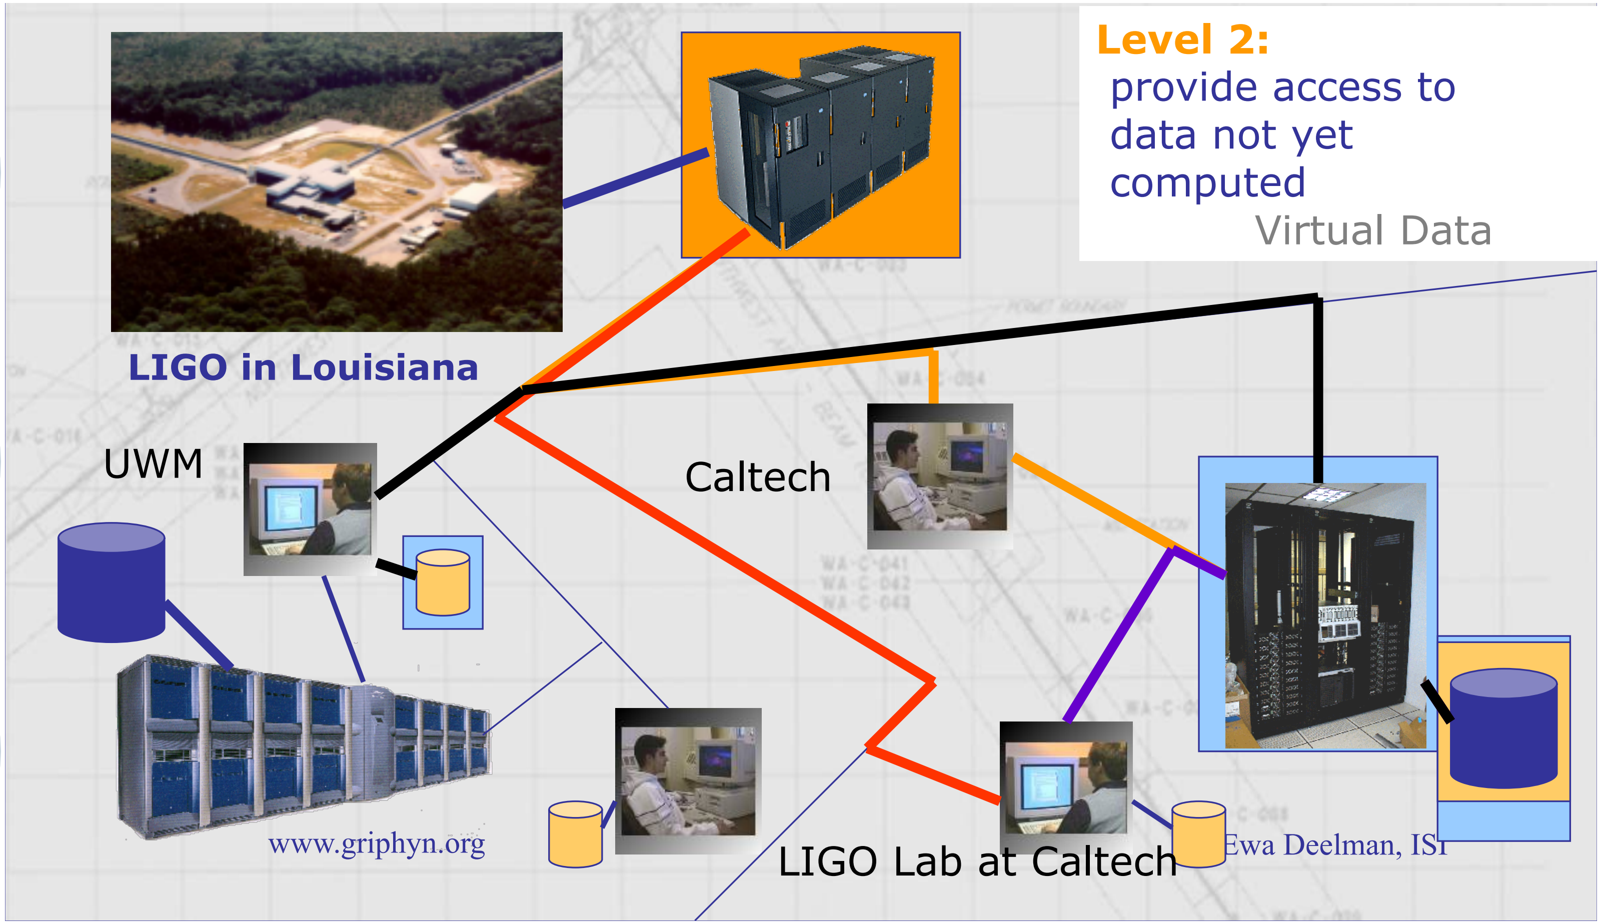

Presentation September 2001

Presentation September 2001

The Virtual Data Grid (VDG) Model

- Data suppliers publish data to the Grid
- Users request raw or derived data from Grid, without needing to know
  - Where data is located
  - Whether data is stored or computed
- User can easily determine
  - What it will cost to obtain data
  - Quality of derived data
- VDG serves requests efficiently, subject to global and local policy constraints

www.griphyn.org

Ewa Deelman, ISI

Virtual Data Scenario

- (LIGO) "Conduct a pulsar search on the data collected from Oct 16 2000 to Jan 1 2001"
- For each requested data value, need to
  - Understand the request
  - Determine if it is instantiated; if so, where; if not, how to compute it
  - Plan data movements and computations required to obtain all results
  - Execute this plan

GriPhyN Project, Ian Foster (PI), Paul Avery, Carl Kesselman, Miron Livny, (co-Pis)

Basic LIGO request

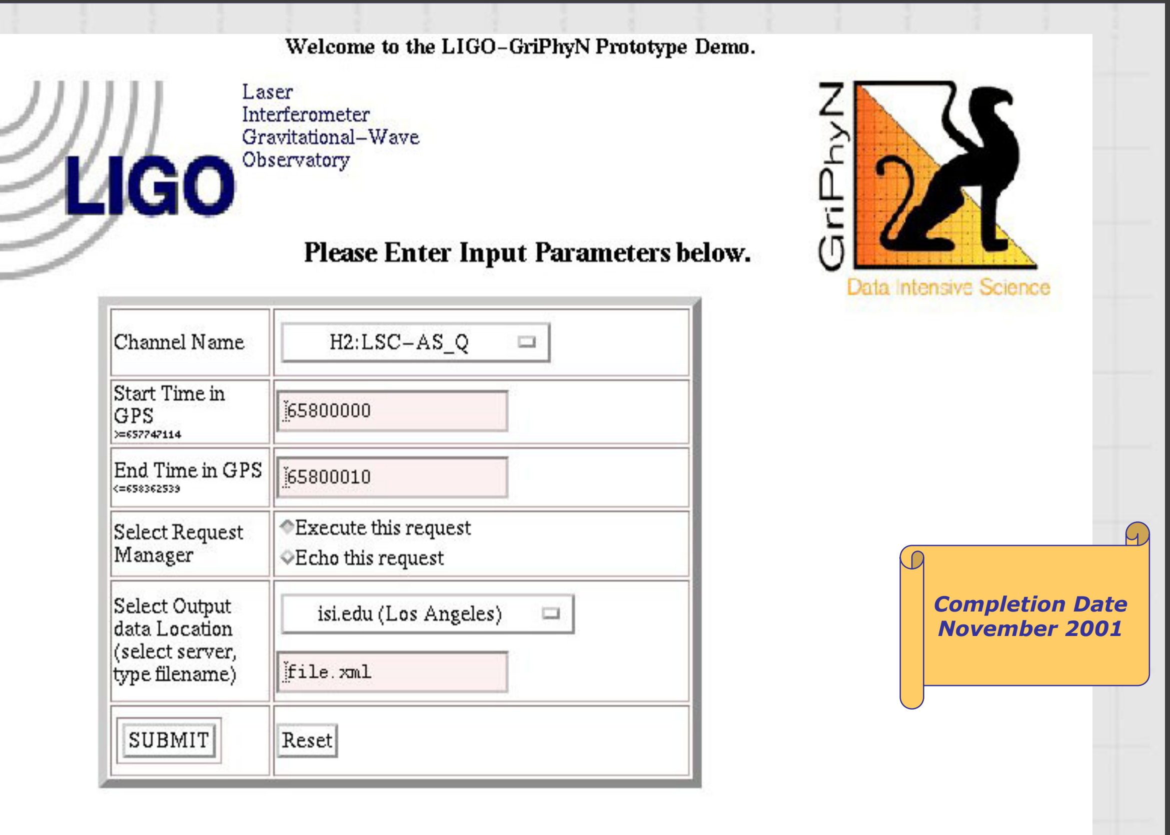

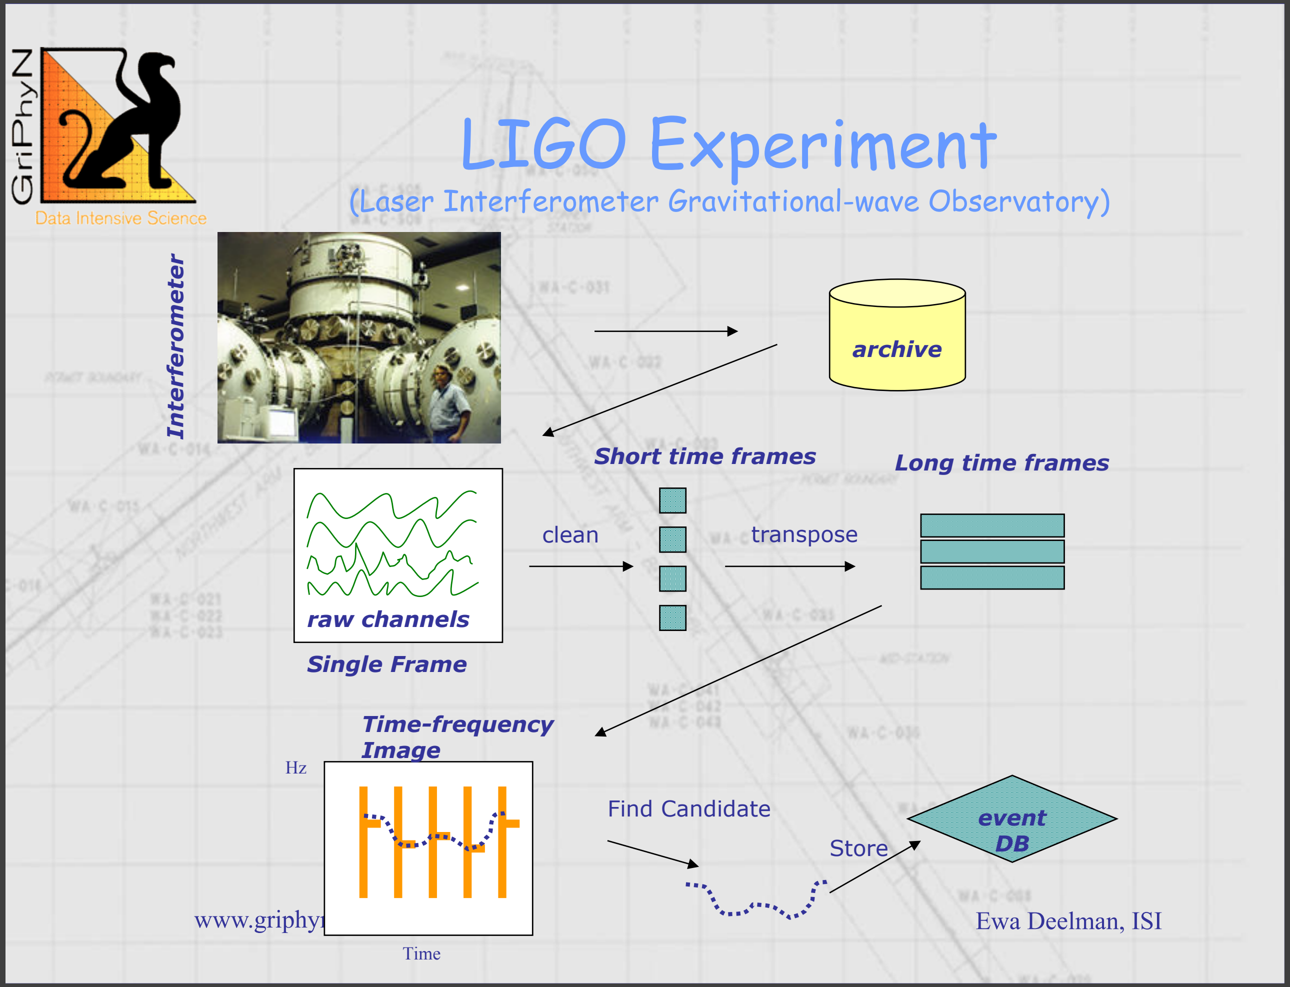

Used AI planning techniques to develop workflows


a Grid
Grid
Secur
LIGO and GriPhyN, Creating a Virtual Data Grid for Gravitational Wave Physics

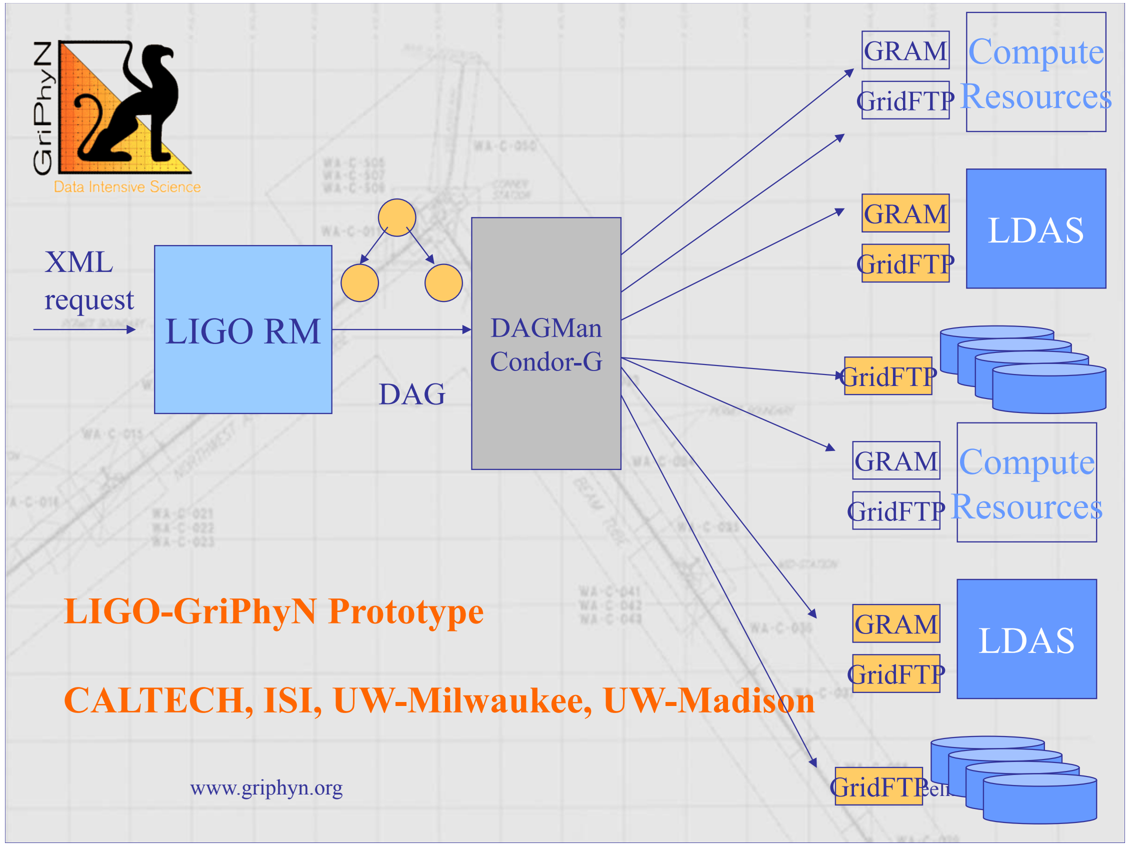
GriPhyN
Data Intensive Science
XML request
LIGO RM
DAG
DAGMan Condor-G
GRAM
GridFTP
Compute Resources
GRAM
GridFTP
LDAS
GridFTP
GRAM
GridFTP
Compute Resources
GRAM
GridFTP
LDAS
GridFTP
LIGO-GriPhyN Prototype
CALTECH, ISI, UW-Milwaukee, UW-Madison
www.griphyn.org

SC'2002 Demo

Basic LIGO request

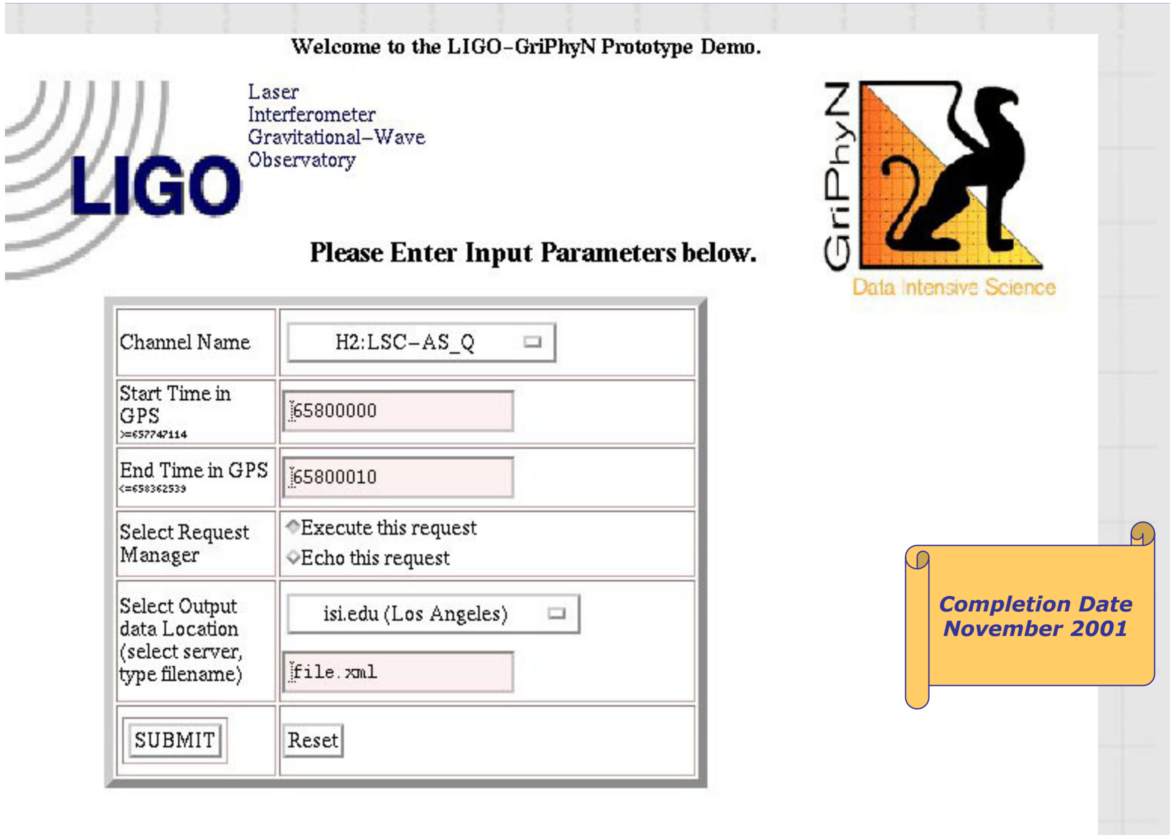

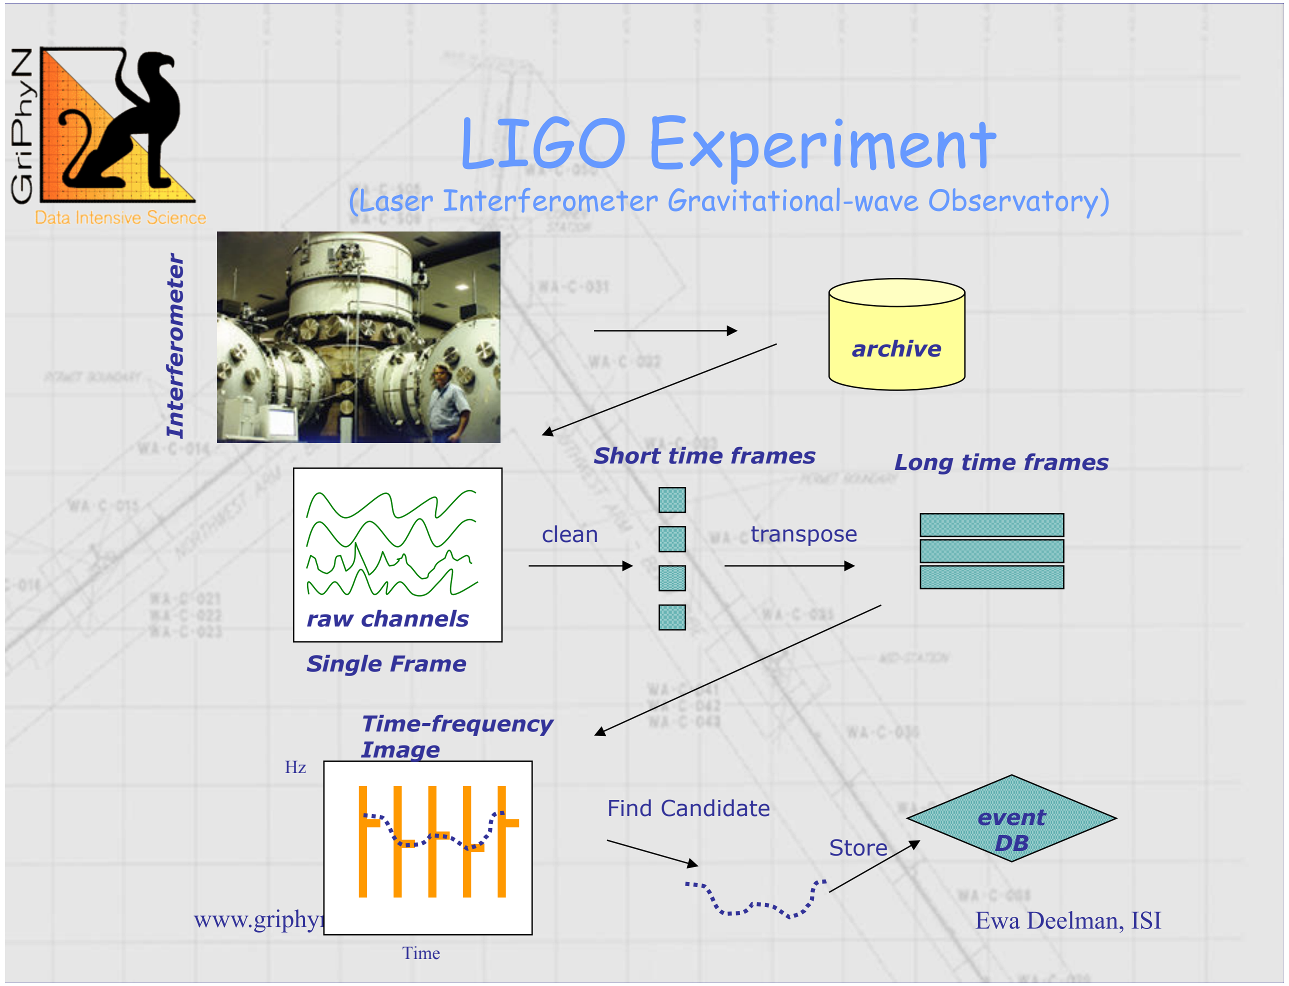

Lessons Learned

- Rewarding to work with real world problems
- Listen to the scientists needs – virtual data was a great concept, but too abstract
- Need to deal with distributed, heterogeneous data and resources
- Separation between workflow description and workflow execution

Focus

- Request planning and scheduling (performance)
- Task execution (fault tolerance)

Pegasus Workflow Management System

Workflow Challenges Across Domains

Describe complex workflows in a simple way

Access distributed, heterogeneous data and resources (heterogeneous interfaces)

Deals with resources/software that change over time

Ease of use. Ability to monitor and debug large workflows

Our Focus

- Separation between workflow description and workflow execution
- Workflow planning and scheduling (scalability, performance)
- Task execution (monitoring, fault tolerance, debugging, web dashboard)
- Workflow optimization, restructuring for performance and fault tolerance.

USC Viterbi School of Engineering Information Sciences Institute

Pegasus Workflow Management System

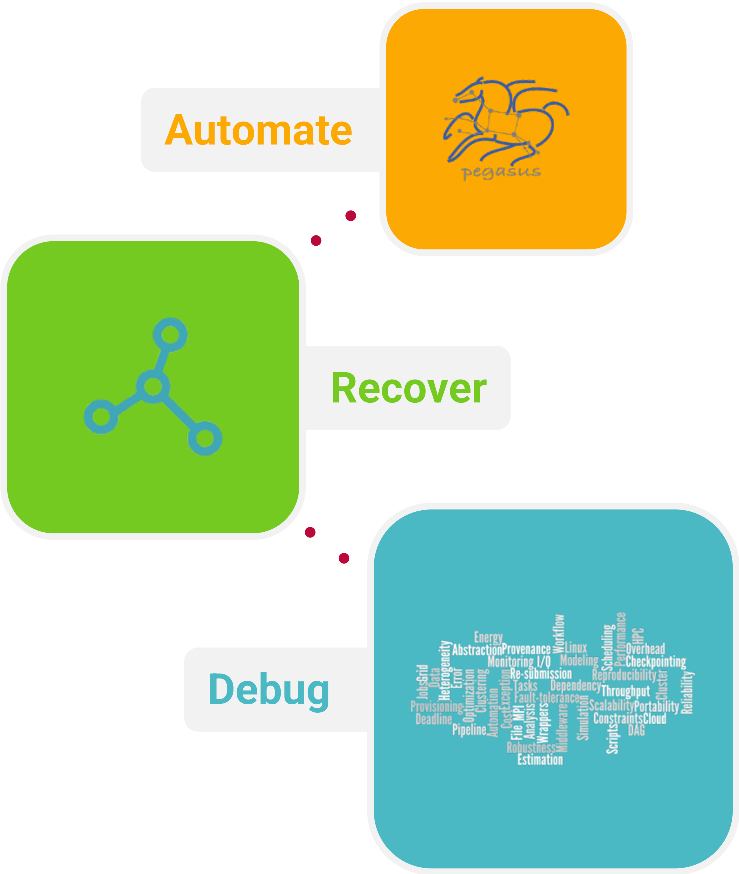

- **Automates Complex**, Multi-stage Processing Pipelines
- Enables Parallel, **Distributed Computations**
- **Automatically Executes** Data Transfers
- Reusable, Aids **Reproducibility**
- Records How Data was Produced **(Provenance)**
- Handles **Failures** to Provide Reliability
- Keeps Track of Data and **Files**

NSF funded project since 2001,
with close collaboration with HTCondor team

2006 Today…

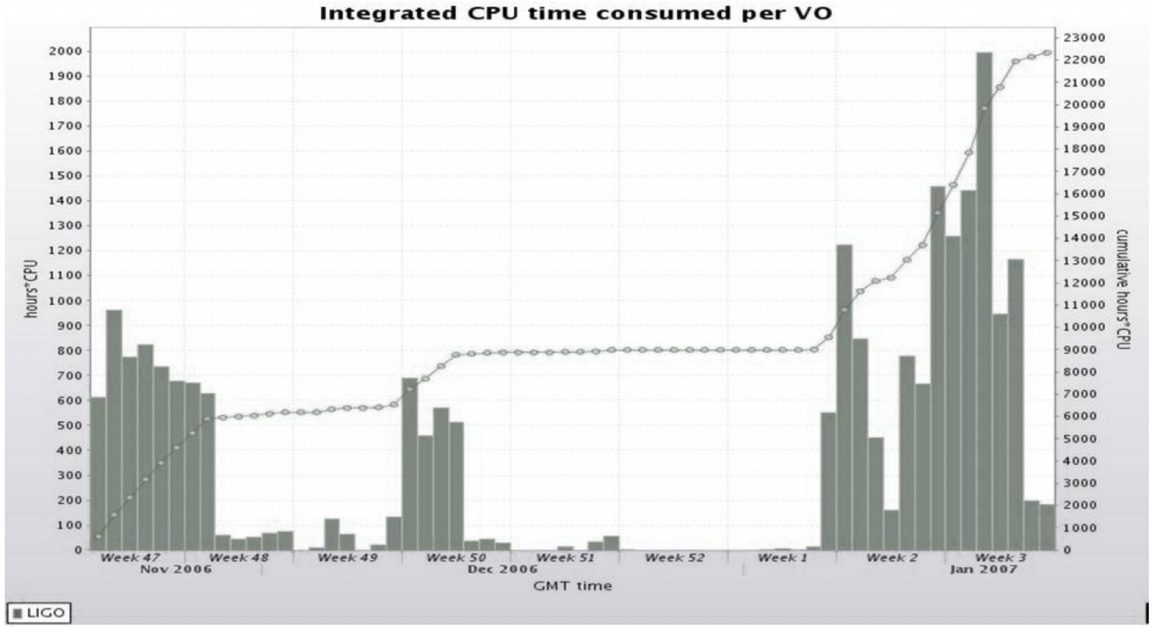

Figure 12: LIGO's CPU Hours Usage of OSG Resources.

2016: First detection of black hole collision

2017

Nobel Prize for LIGO

Multi-messenger neutron star merger observation

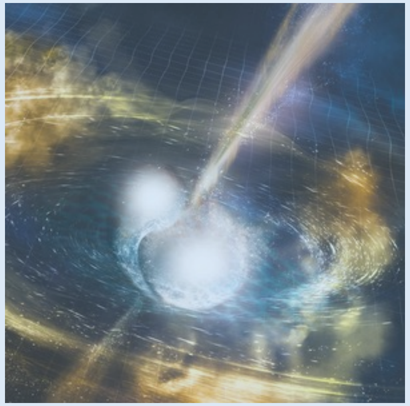

USC Viterbi
School of Engineering
Information Sciences Institute

2006 Today

ARTICLE

Powerful new res… TeraGrid *

Authors: D Scott McCaulay, … Craig A Stewart Authors Info

SC '06: Proceedings of the 2006 A…
es • https://doi.org/10.1145/1188…

ABSTRACT

Indiana University's 20.48 T… been made available to res… system finished 23rd on the… supercomputer owned and… available through the TeraG… of 512 IBM BladeCenter JS2… of ECC PC3200 SDRAM, 72G… a PCI-X Myrinet 2000 adapter for high-bandwidth, low-latency MPI applications.A significant portion

Home »Resources »News

NEWS

ORNL's Frontier First to Break the Exaflop Ceiling

May 30, 2022

FRANKFURT, Germany; BERKELEY, Calif.; and KNOXVILLE, Tenn.— The 59th edition of the TOP500 revealed the Frontier system to be the first true exascale machine with an HPL score of 1.102 Exaflop/s.

The No. 1 spot is now held by the Frontier system at Oak Ridge National Laboratory (ORNL) in the US. Based on the latest HPE Cray EX235a architecture and equipped with AMD EPYC 64C 2GHz processors, the system has 8,730,112 total cores, a power efficiency rating of 52.23 gigaflops/watt, and relies on gigabit ethernet for data transfer.

However, a recent development to the Frontier system has allowed the machine to surpass the 1 exaflop barrier. With an exact HPL score of 1.102 Exaflop/s, Frontier is not only the most powerful supercomputer to ever exist – it's also the first true exascale machine.

http://top500.org

* DOE's LLNL IBM BlueGene/L system, 280.6 TFlop/s was #1 on Top 500 in 2006

1. Resource-independent Specification

Input Workflow Specification YAML formatted

directed-acyclic graphs **Output Workflow**

Portable Description

Users do not worry about low level execution details

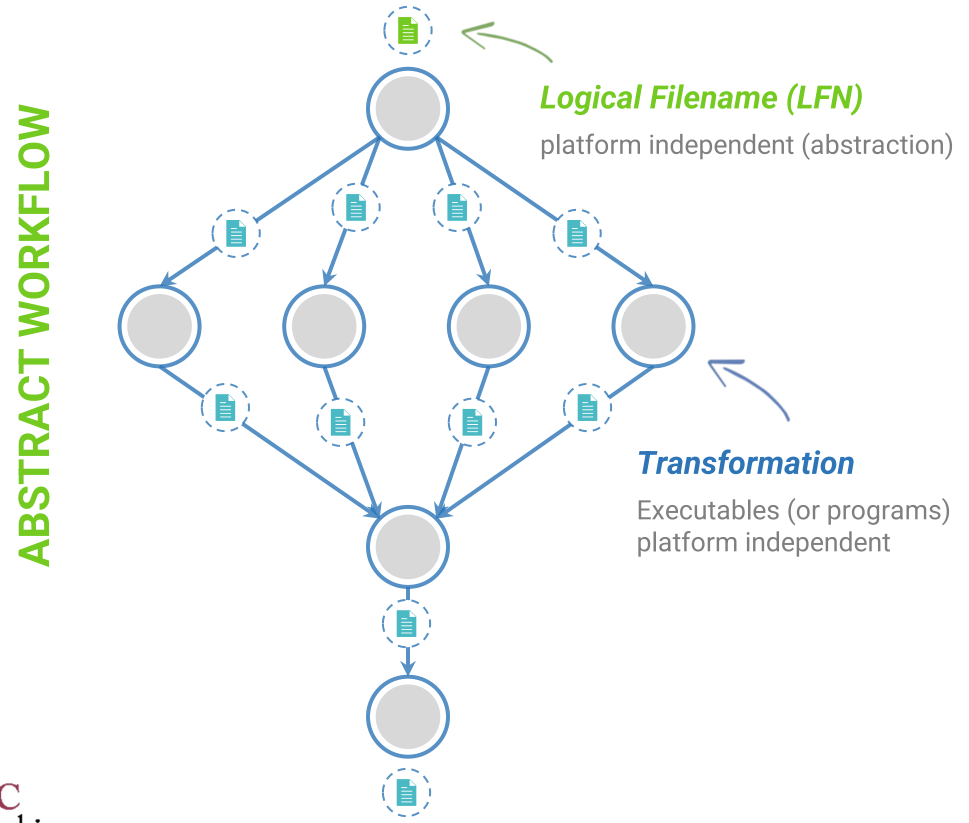

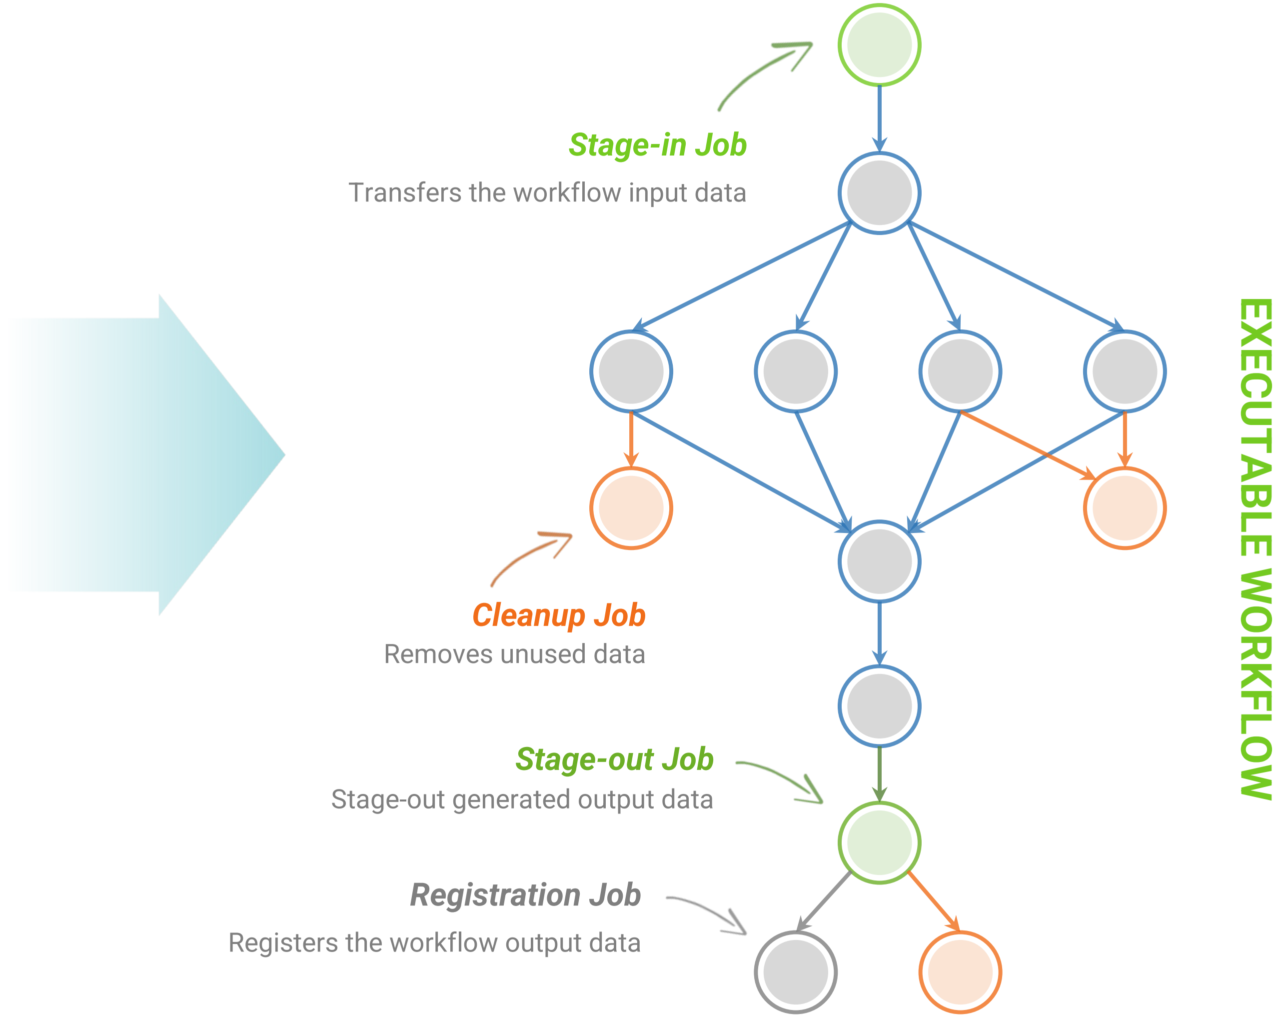

2006 Southern California Earthquake Center — 14 year later...

CyberShake:

What will the peak earthquake motion be over the next 50 years?

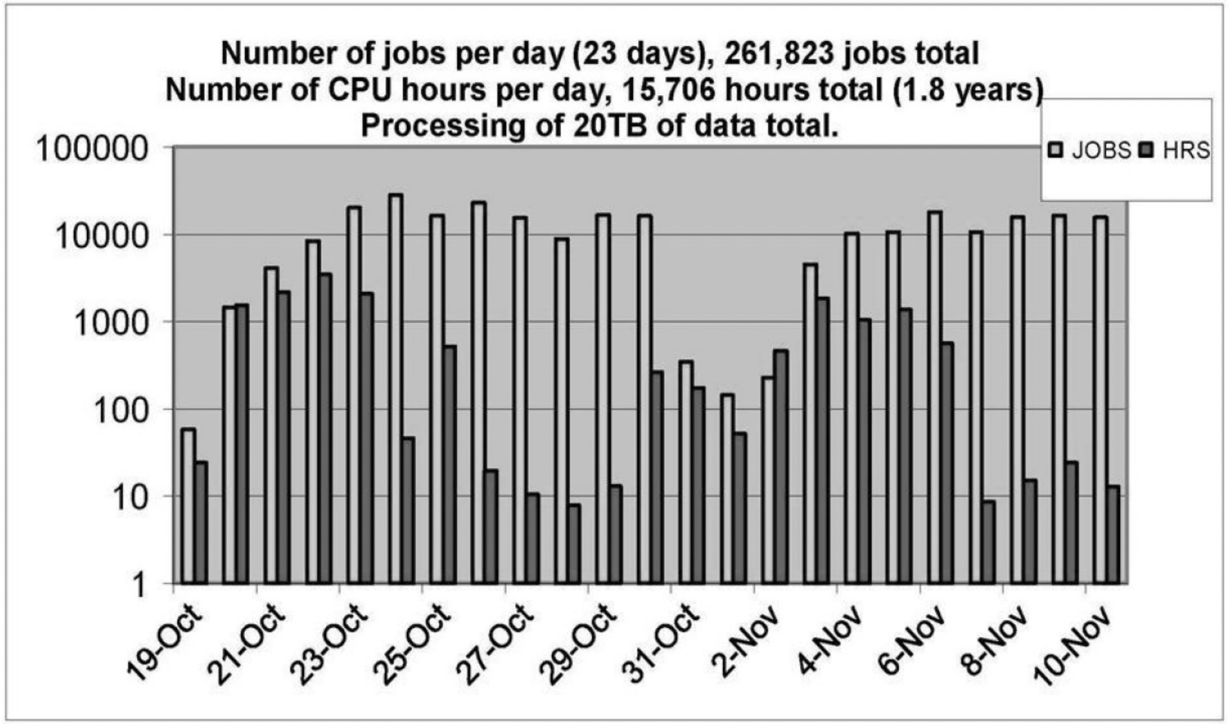

Figure 11: Distribution of Seismogram Tasks in the SCEC CyberShake Workflow.

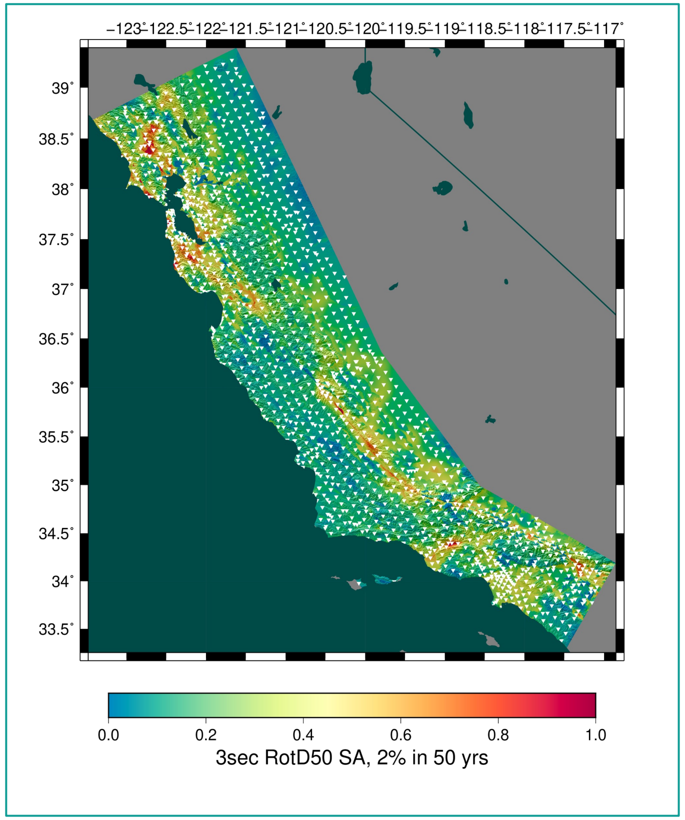

- **120 million** core-hours
- **39,285** jobs

- **1.2 PB** of data managed
- **157 TB** of data automatically transferred
- **14.4 TB** of output data archived

Mix of MPI and HTC postprocessing jobs

NCSA Blue Waters | OLCF Titan

Pegasus

2. Submit locally, run globally

- **Pegasus WMS ==** **Pegasus planner (mapper) + DAGMan workflow engine + HTCondor scheduler/broker**
  - Pegasus maps workflows to target infrastructure (1 or more resources)
  - DAGMan manages dependencies and reliability
  - HTCondor is used as a broker to interface with different schedulers
- **Planning converts an abstract workflow into a concrete, executable workflow**
  - Planner is like a compiler
  - Optimized performance
  - Provides fault tolerance
- **Can leverage distributed and heterogeneous CI**

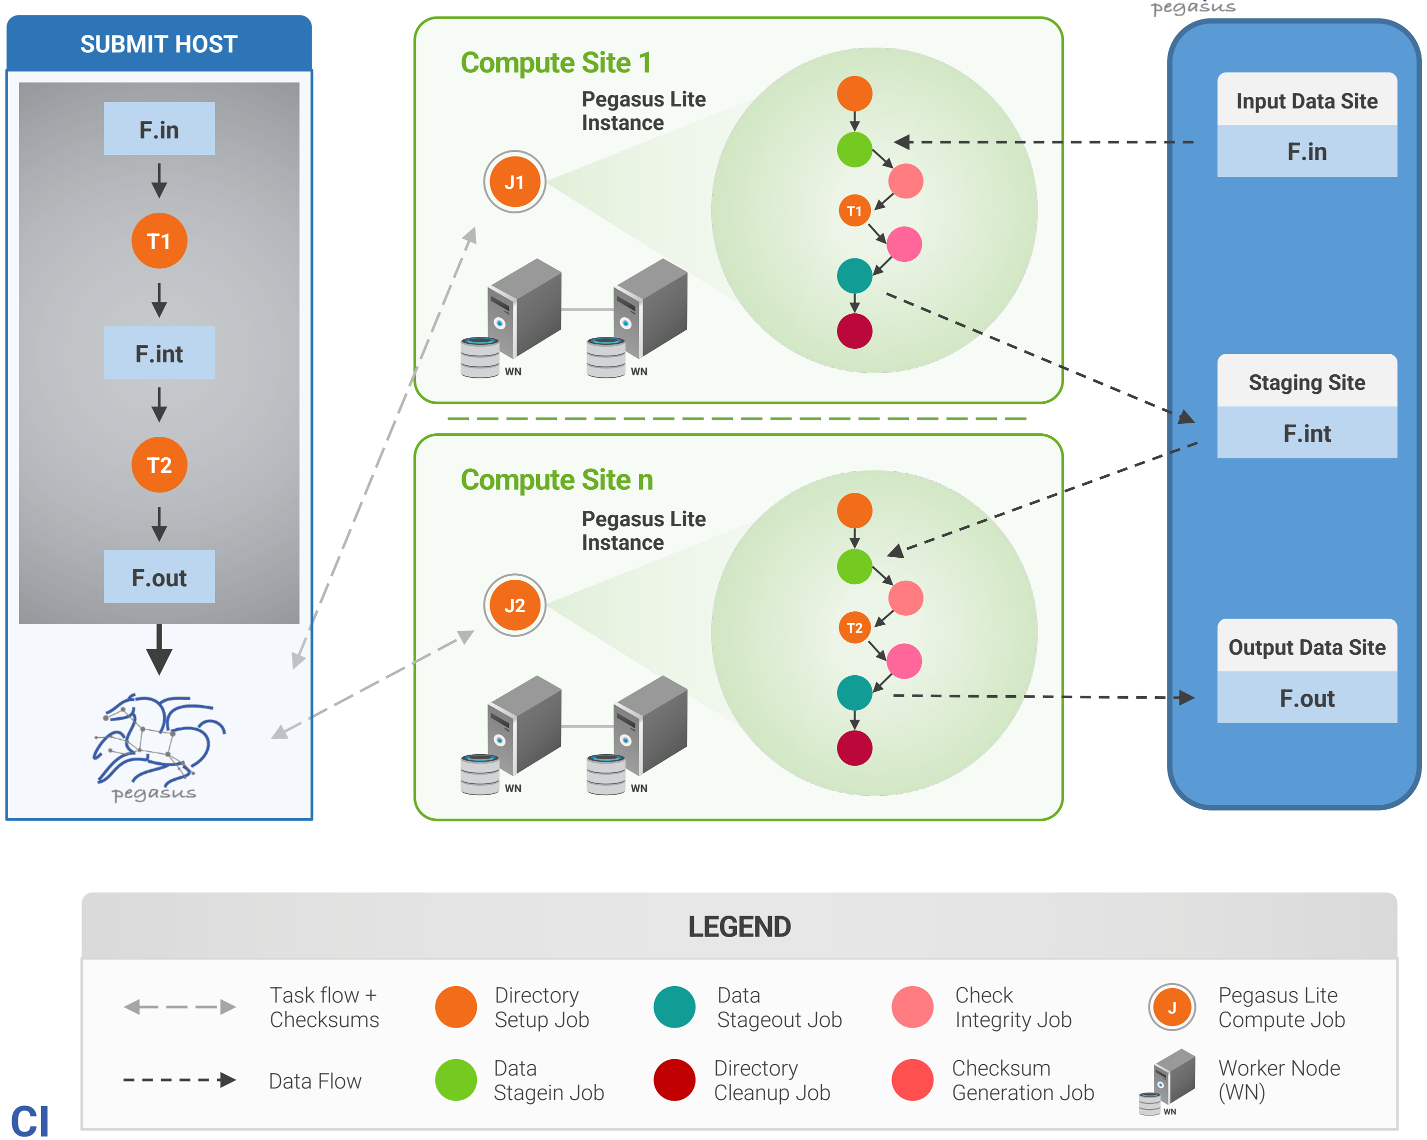

USC Viterbi School of Engineering Information Sciences Institute

Vermont EPSCoR Integrated Assessment Model

Integrated Assessment Models (IAMs)

IAMs project the impact of policy scenarios on socio-environmental systems (SESs) and thus, provide policy makers with data to make informed policy decisions.

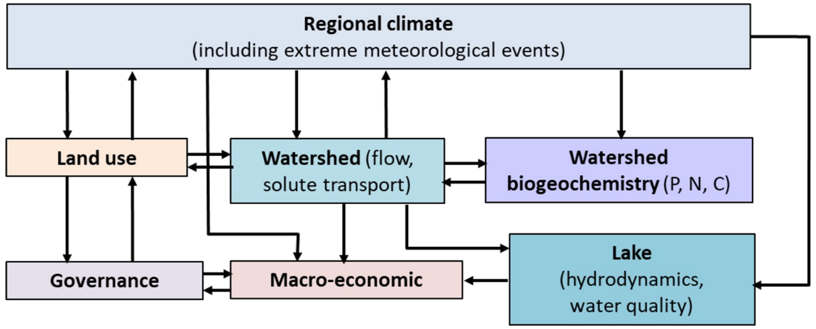

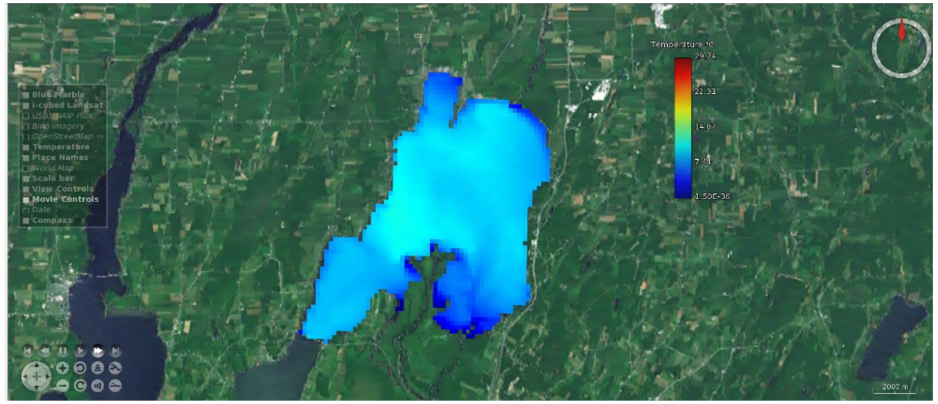

Temperature in Missisqoui Bay, Lake Champlain

Heterogeneous Execution Platforms – University of Vermont (UVM) and UCAR's Cheyenne

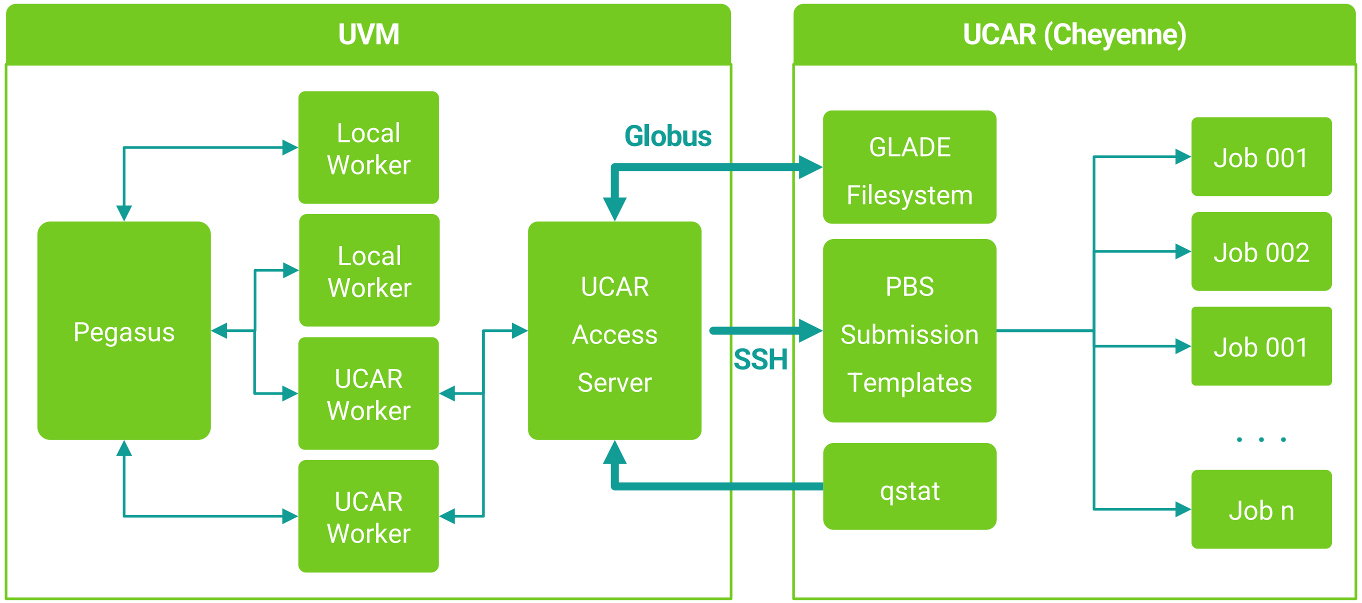

Parallelism achieved over different policy scenarios (max of 176 so far)

Each decade run sequentially in a cascade
to enable bidirectional asynchronous feedback between models

176 Scenario, 60 Year IAM Run:
~ 10 days wall time, ~ 25,000 jobs, ~ 3.5TB data generated

This material is based upon work supported
by the National Science Foundation under Grant No. OIA 1556770

USC Viterbi School of Engineering
Information Sciences Institute

2008: Feasibility of using Clouds for science

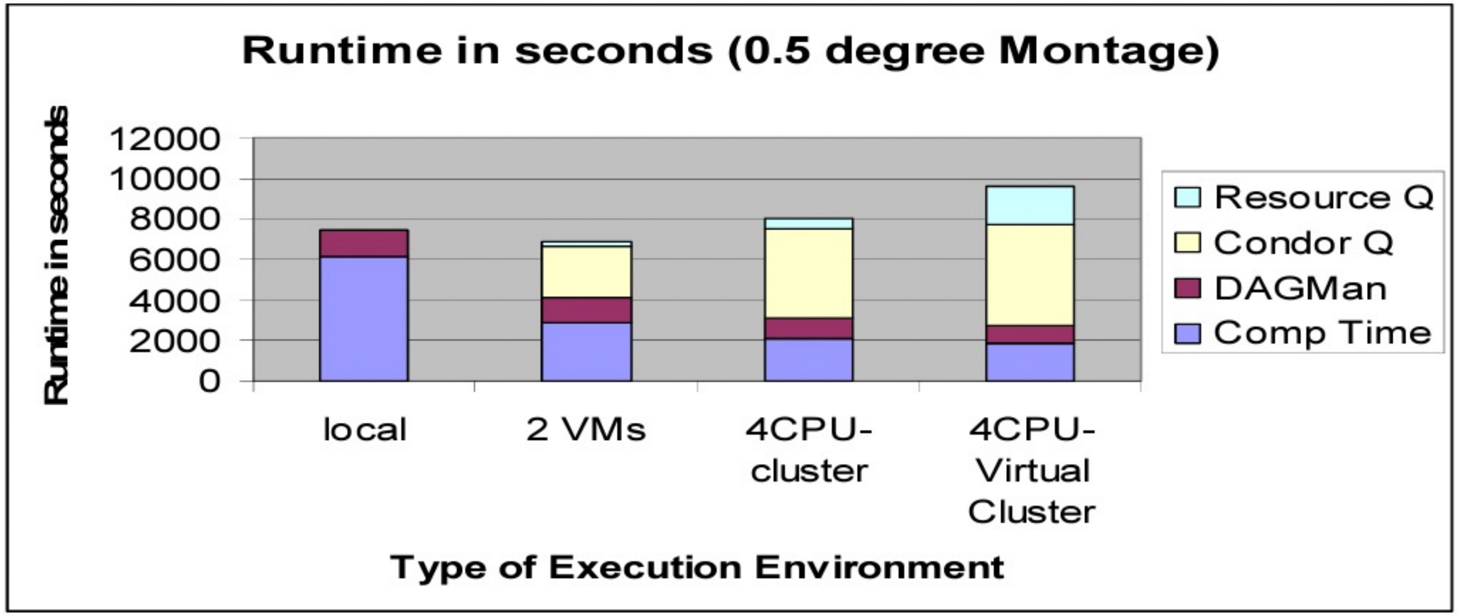

Figure 4: Overheads in the Workflow Execution in Various Environments.

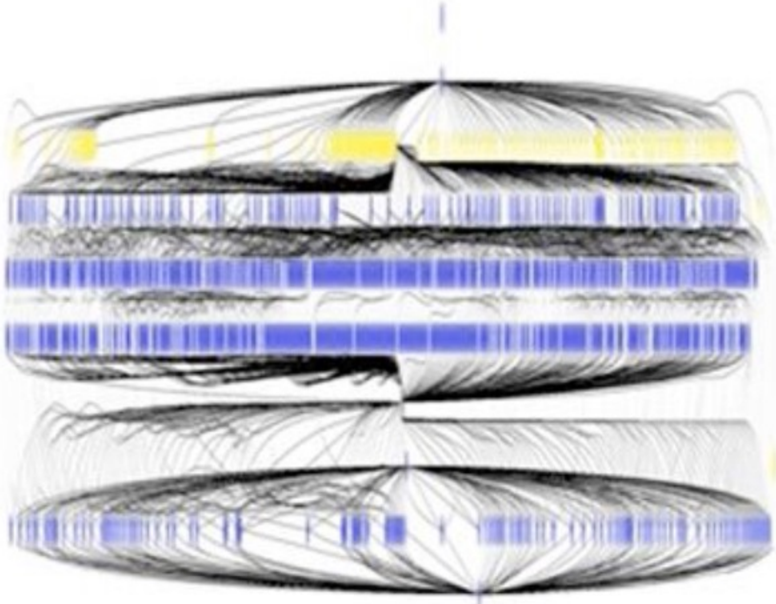

Bruce Berriman, CALTECH

3. Flexible Data Staging Configurations

HTCondor I/O (HTCondor pools, OSG, …)

Worker nodes do not share a file system
Data is pulled from / pushed to the submit host via HTCondor file transfers
Staging site is the submit host

Shared File System
(HPC sites, XSEDE, Campus clusters, …)

I/O is directly against the shared file system

Non-shared File System (clouds, OSG, …)

Worker nodes do not share a file system
Data is pulled / pushed from a staging site, possibly not co-located with the computation

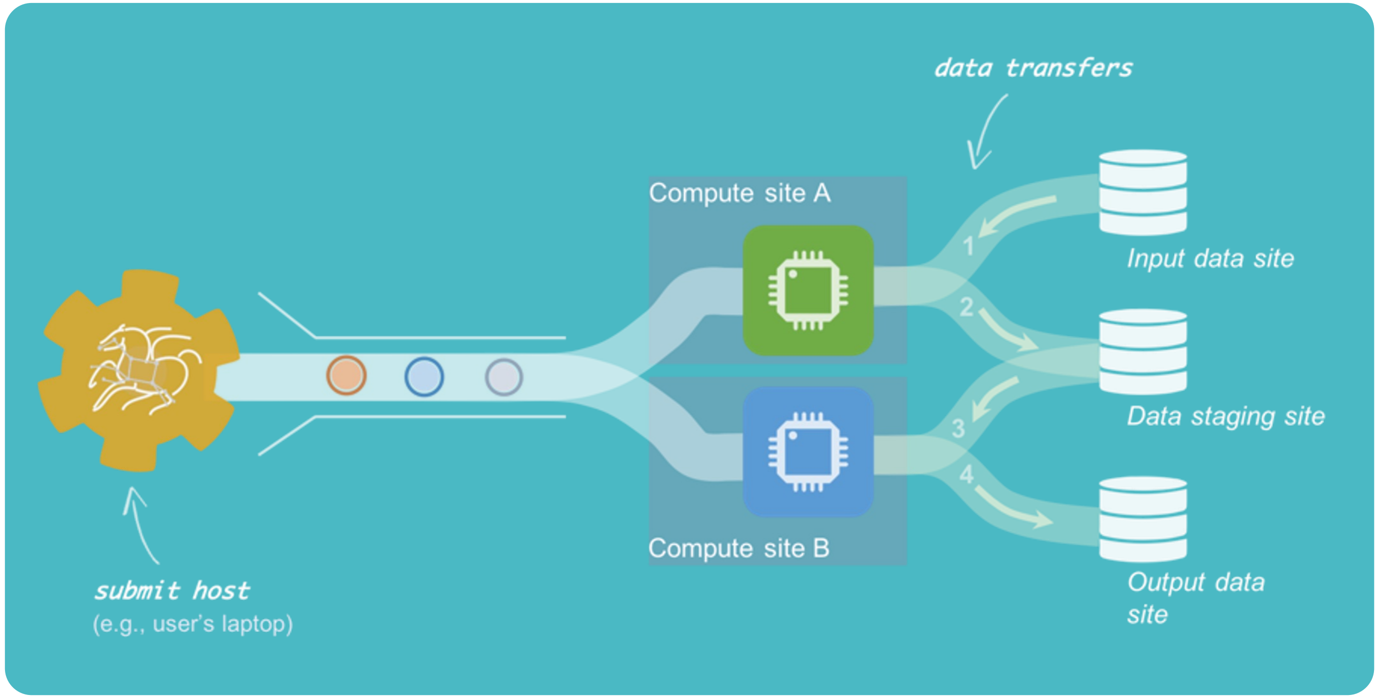

4. Flexible Data movement Pegasus-transfer

Pegasus' internal data transfer tool with support for a number of different protocols

- **Directory creation, file removal**
  - If protocol can support it, also used for cleanup
- **Two stage transfers**
  - e.g., GridFTP to S3 = GridFTP to local file, local file to S3
- **Parallel transfers**
- **Automatic retries**
- **Credential management**
  - Uses the appropriate credential for each site and each protocol (even 3rd party transfers)

HTTP
SCP
GridFTP
Globus Online
iRods
Amazon S3
Google Storage
SRM
FDT
Stashcp
Rucio
cp
ln -s

Edge-2-Cloud Applications

CASA: Collaborative and Adaptive Sensing of the Atmosphere

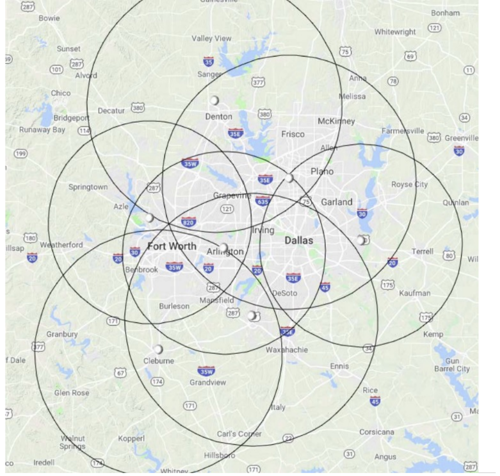

http://www.casa.umass.edu/

- Has deployed a network of short-range Doppler radars
- Compute and data repositories at the edge, close to the radars
- Use on demand cloud resources to scale up their computations

OOI: Ocean Observatories Initiative

- Has deployed a variety of sensors in the Atlantic and the Pacific oceans to study the oceans and the marine life
- Hydrophone sensors have been deployed in three locations in the state of Washington
- The Orcasound community initiative is using them to study Orca whales in the Pacific Northwest region.

https://www.orcasound.net/

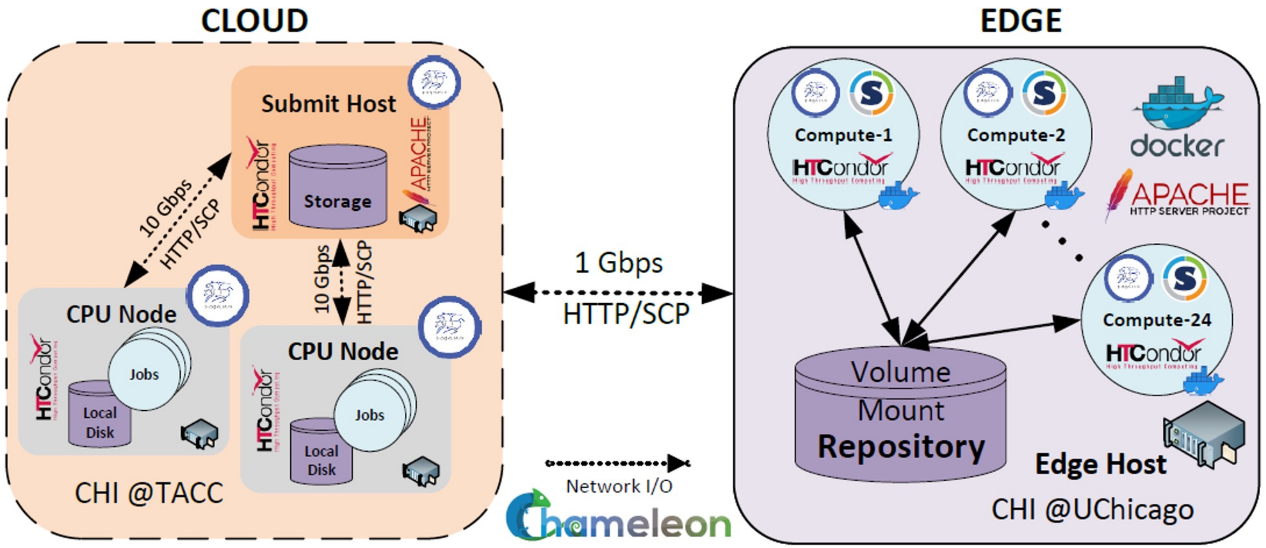

USC Viterbi School of Engineering
Information Sciences Institute

5. "Up and down" integrations with diverse Cyberinfrastructure, common languages, and Portal/GUI interfaces

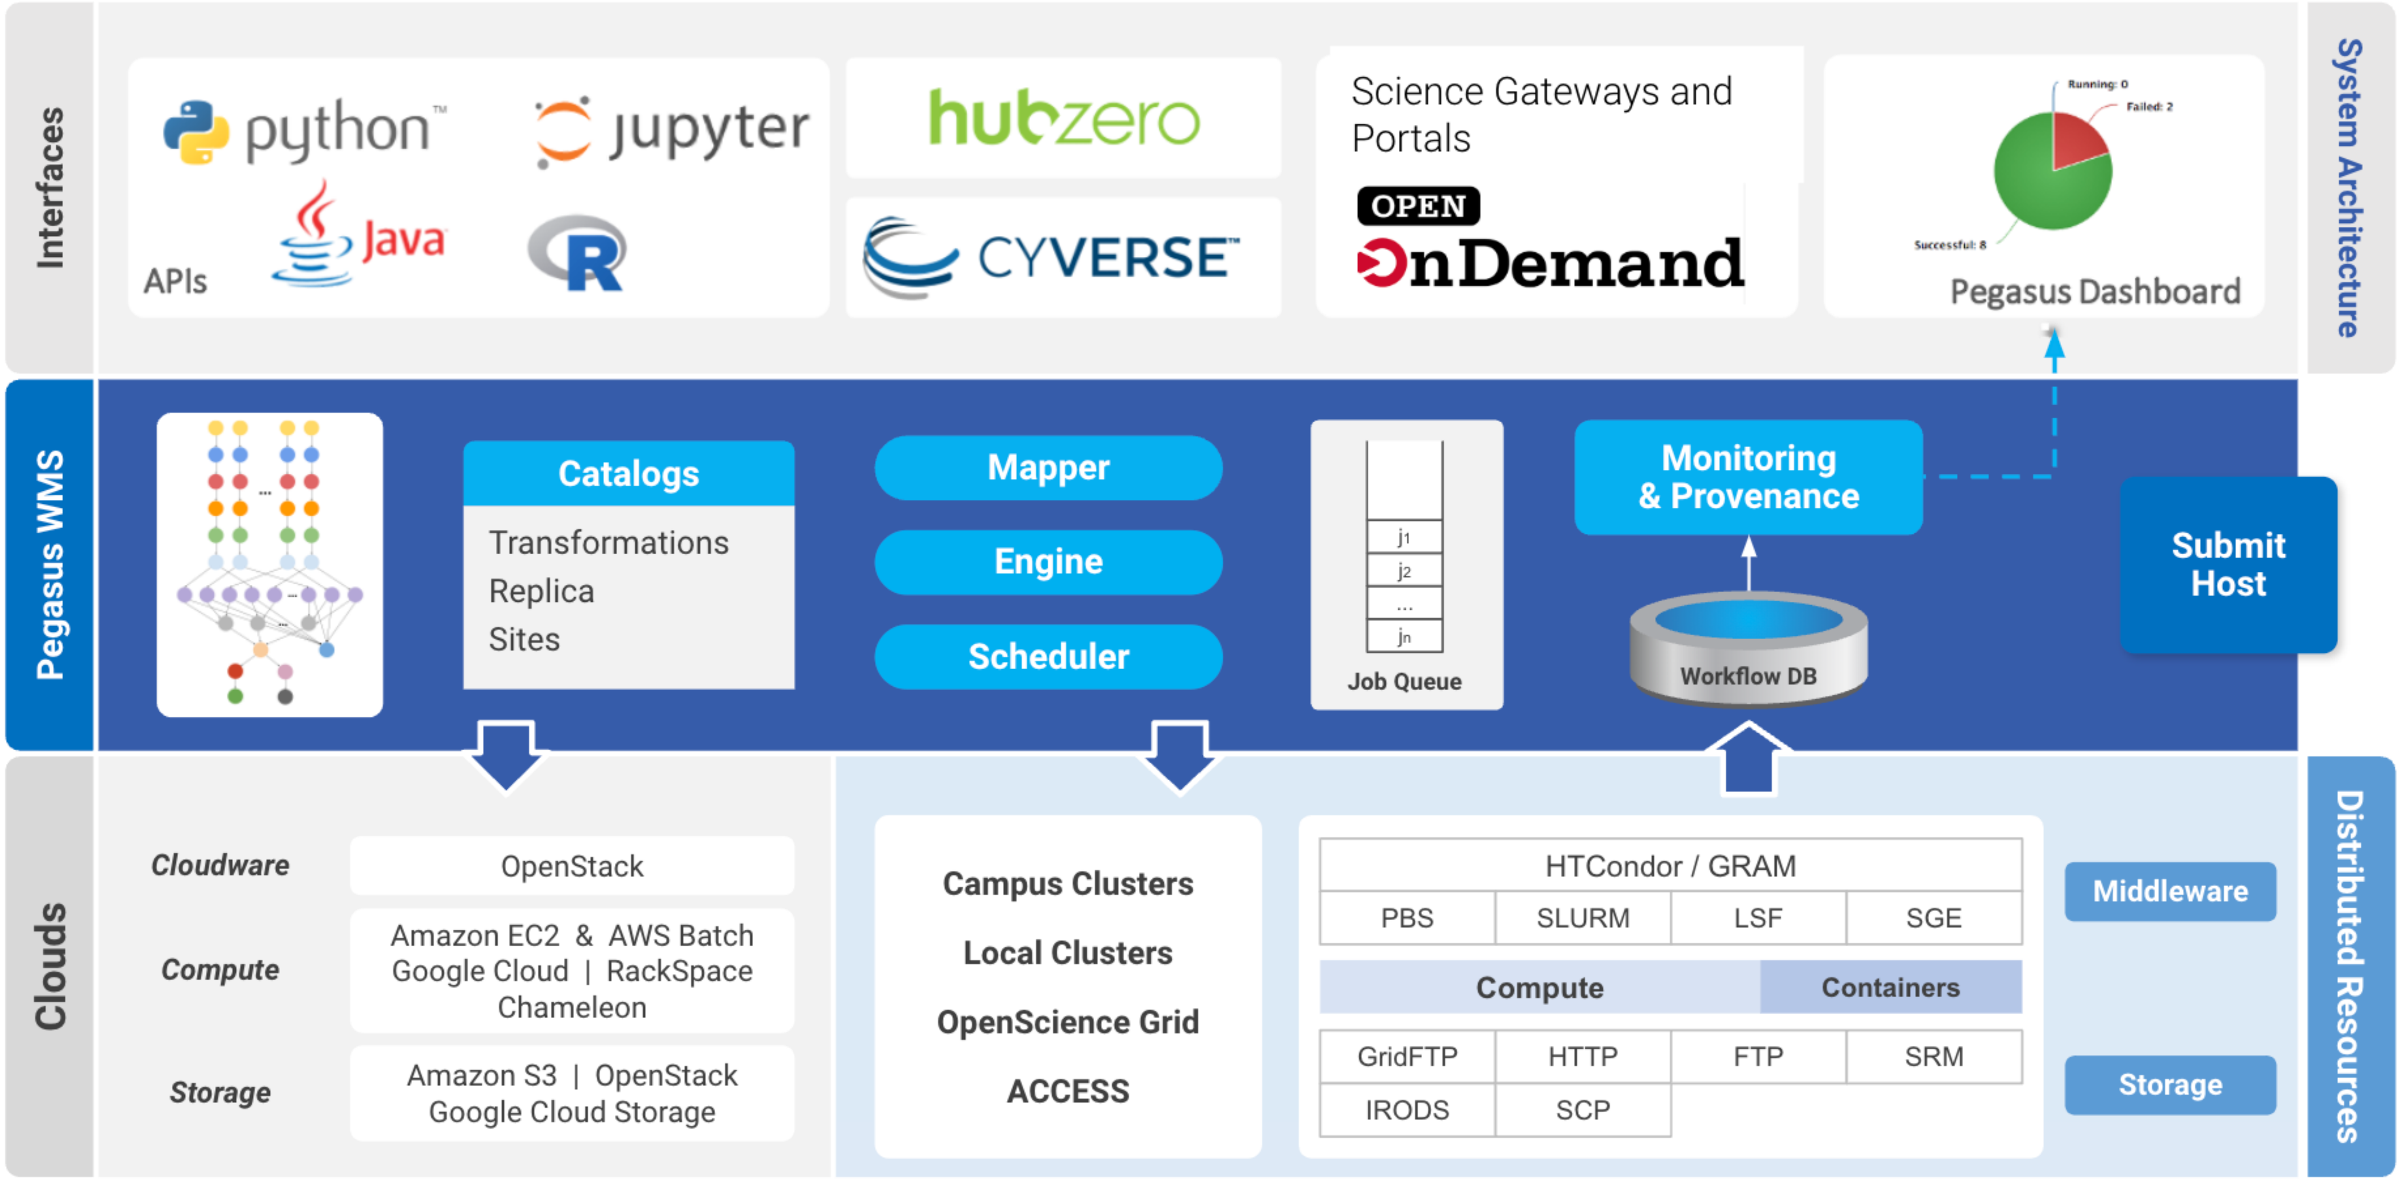

Event Horizon Telescope

Bringing Black Holes into Focus

8 telescopes: 5 PB of data

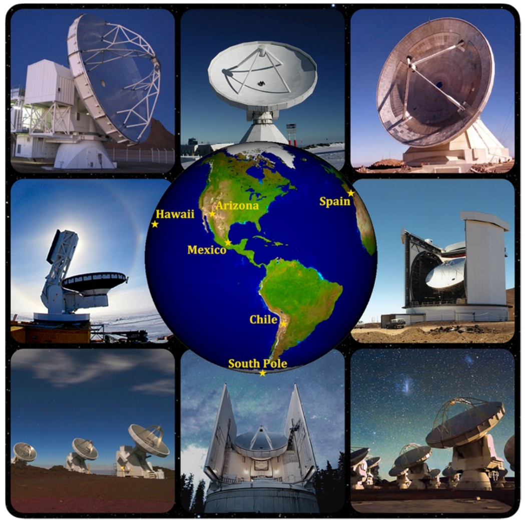

60 simulations: 35 TB data

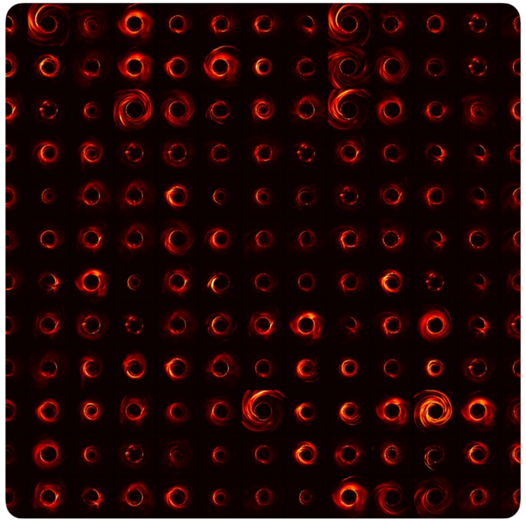

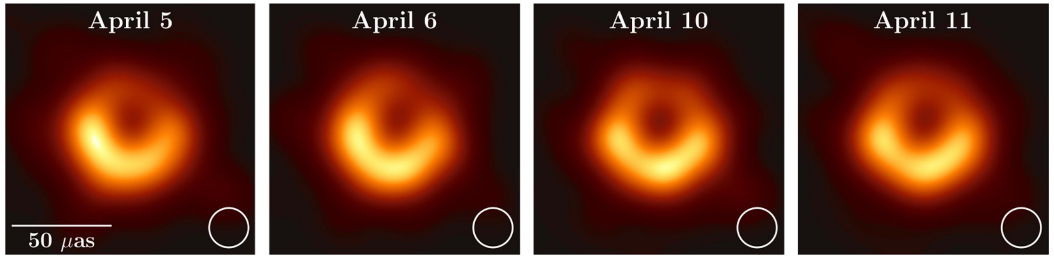

First images of black hole at the center of the M87 galaxy

480,000 jobs - 2,600,000 core hours

#15 in all OSG projects in last 6 months

#2 in all OSG astronomy projects in the last 6 months

Pegasus-SYMBA Pipeline

Physically accurate synthetic observation data from simulations are keys to develop calibration and imaging algorithms, as well as comparing the observation with theory and interpreting the results.

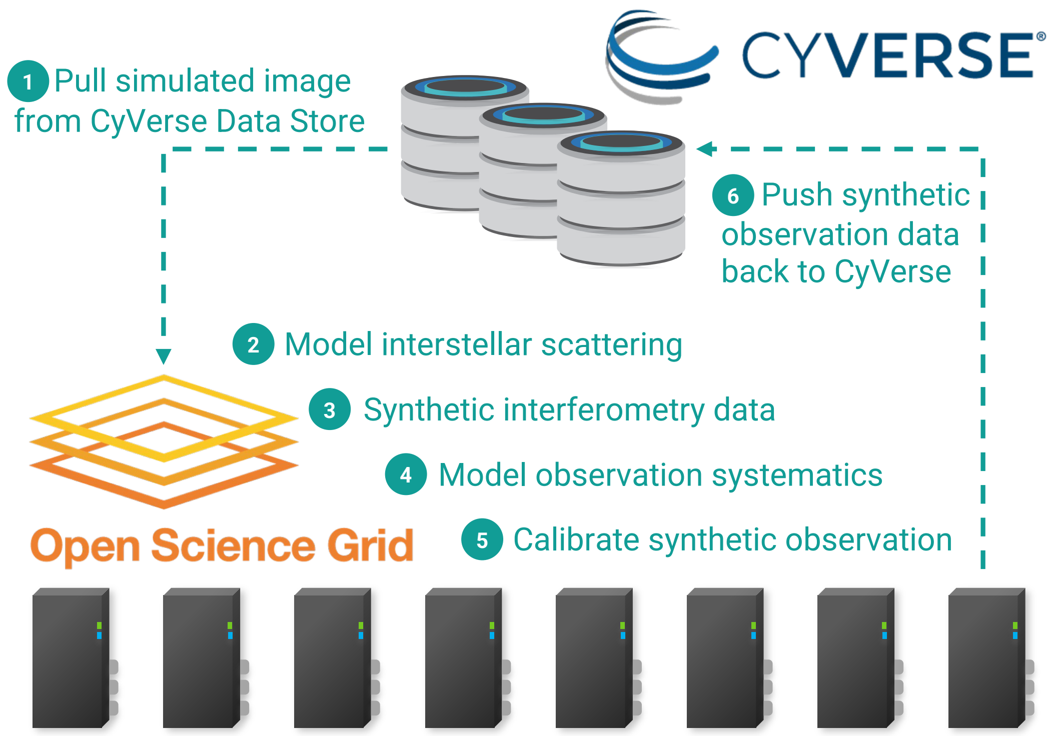

Improve constraints on Einstein's theory of general relativity by 500x

USC Viterbi School of Engineering Information Sciences Institute

https://pegasus.isi.edu

XENONnT - Dark Matter Search

Two Workflows

Monte Carlo simulations and the main processing pipeline.

Workflows execute across Open Science Grid (OSG) & **European Grid Infrastructure (EGI)**

Rucio for data management

MongoDB instance to track science runs and data products.

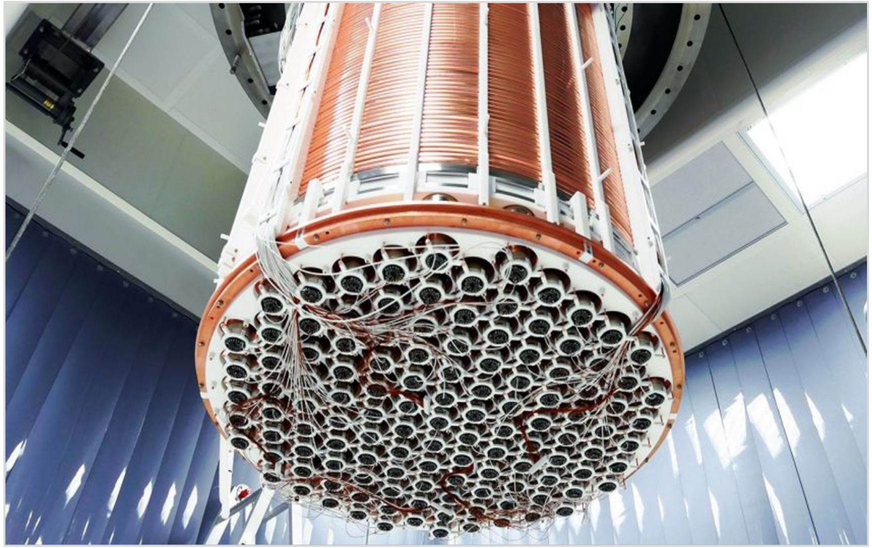

```
------------------------------------------------------------------------------
Type           Succeeded Failed  Incomplete  Total     Retries   Total+Retries
Tasks          4000      0       0           4000      267       4267
Jobs           4484      0       0           4484      267       4751
Sub-Workflows  0         0       0           0         0         0
------------------------------------------------------------------------------

Workflow wall time                                       : 5 hrs, 2 mins
Cumulative job wall time                                 : 136 days, 9 hrs
Cumulative job wall time as seen from submit side        : 141 days, 16 hrs
Cumulative job badput wall time                          : 1 day, 2 hrs
Cumulative job badput wall time as seen from submit side : 4 days, 20 hrs
```

The challenge is to manage the workflows accessing different data management systems.

Processing instrument data in real time

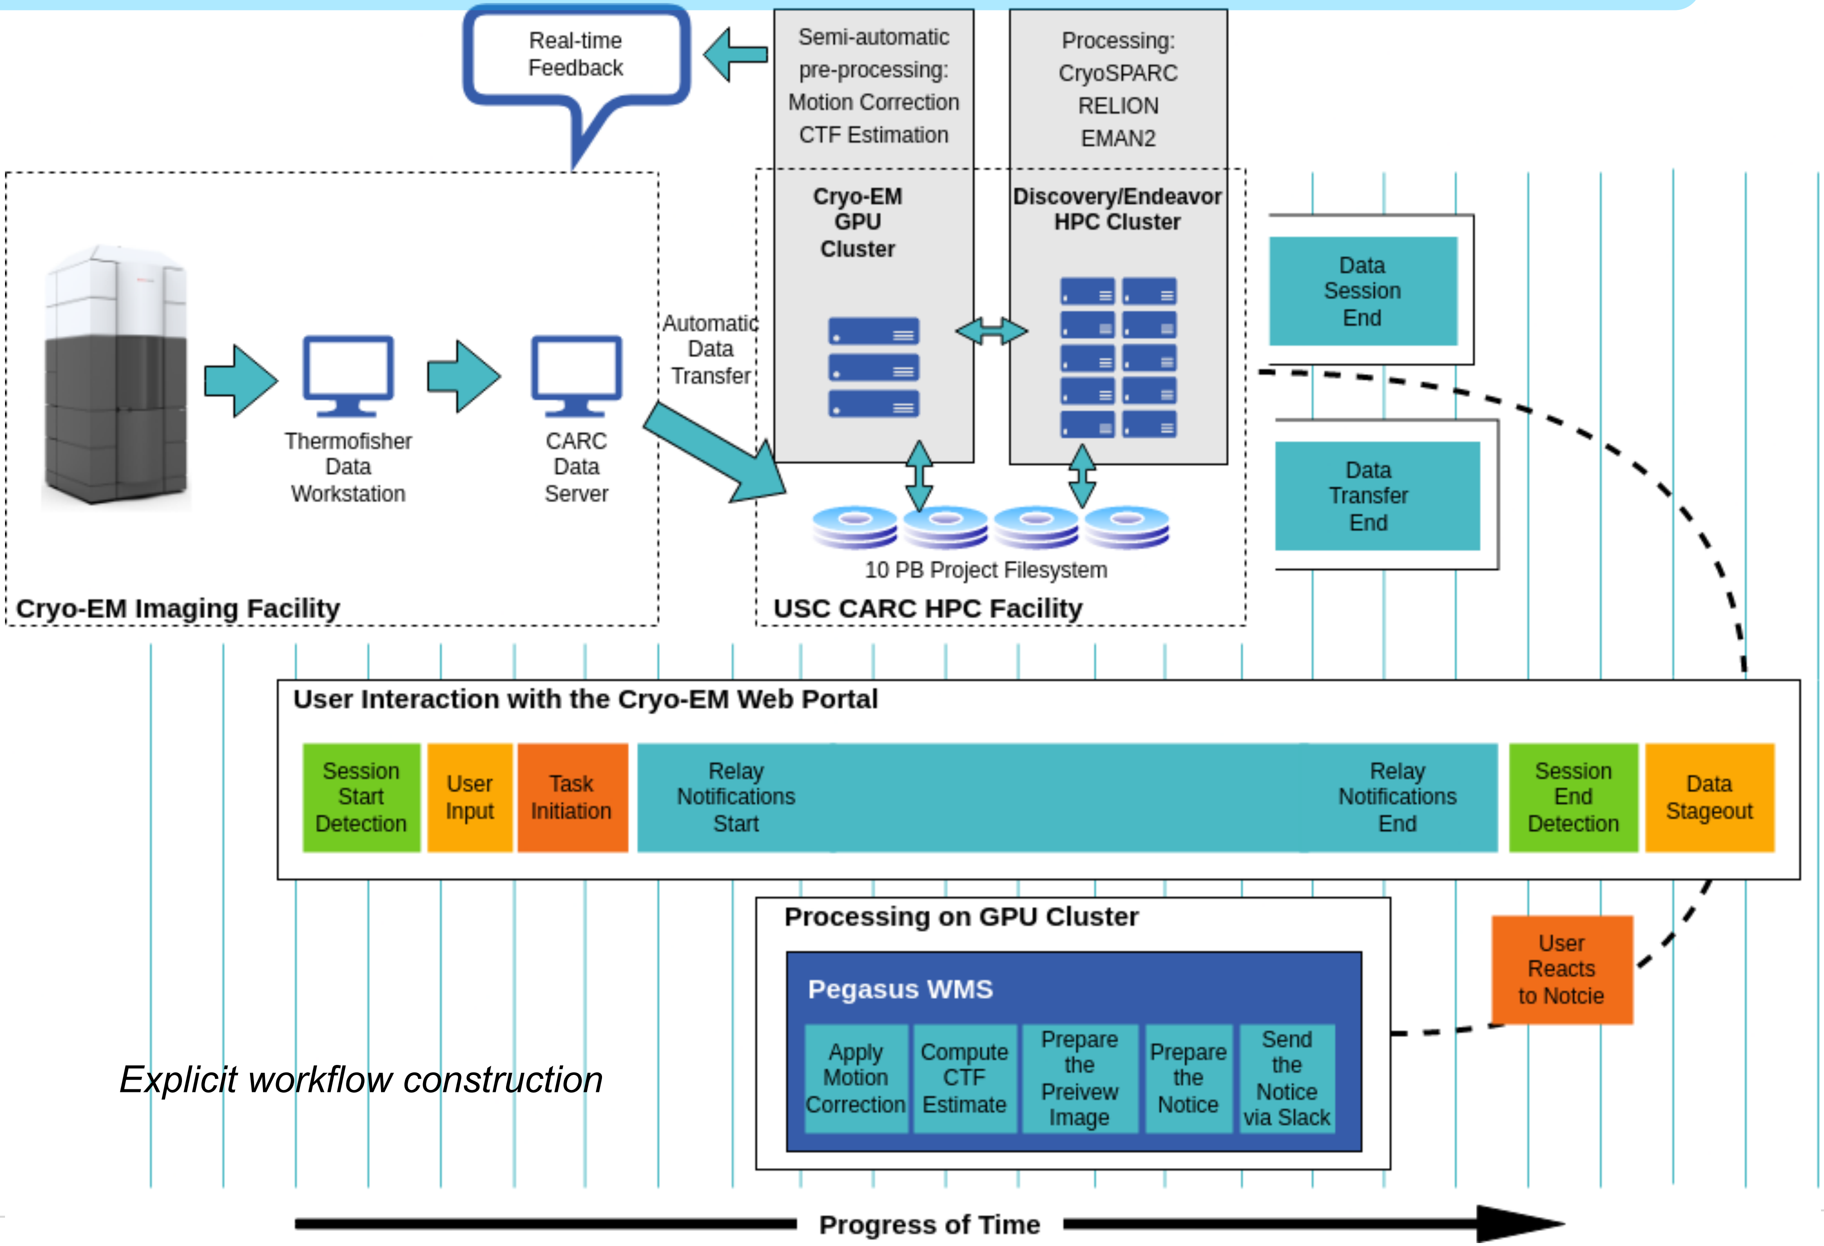

Hidden workflows

Explicit workflow construction

Community Archives: Galactic Plane Atlas

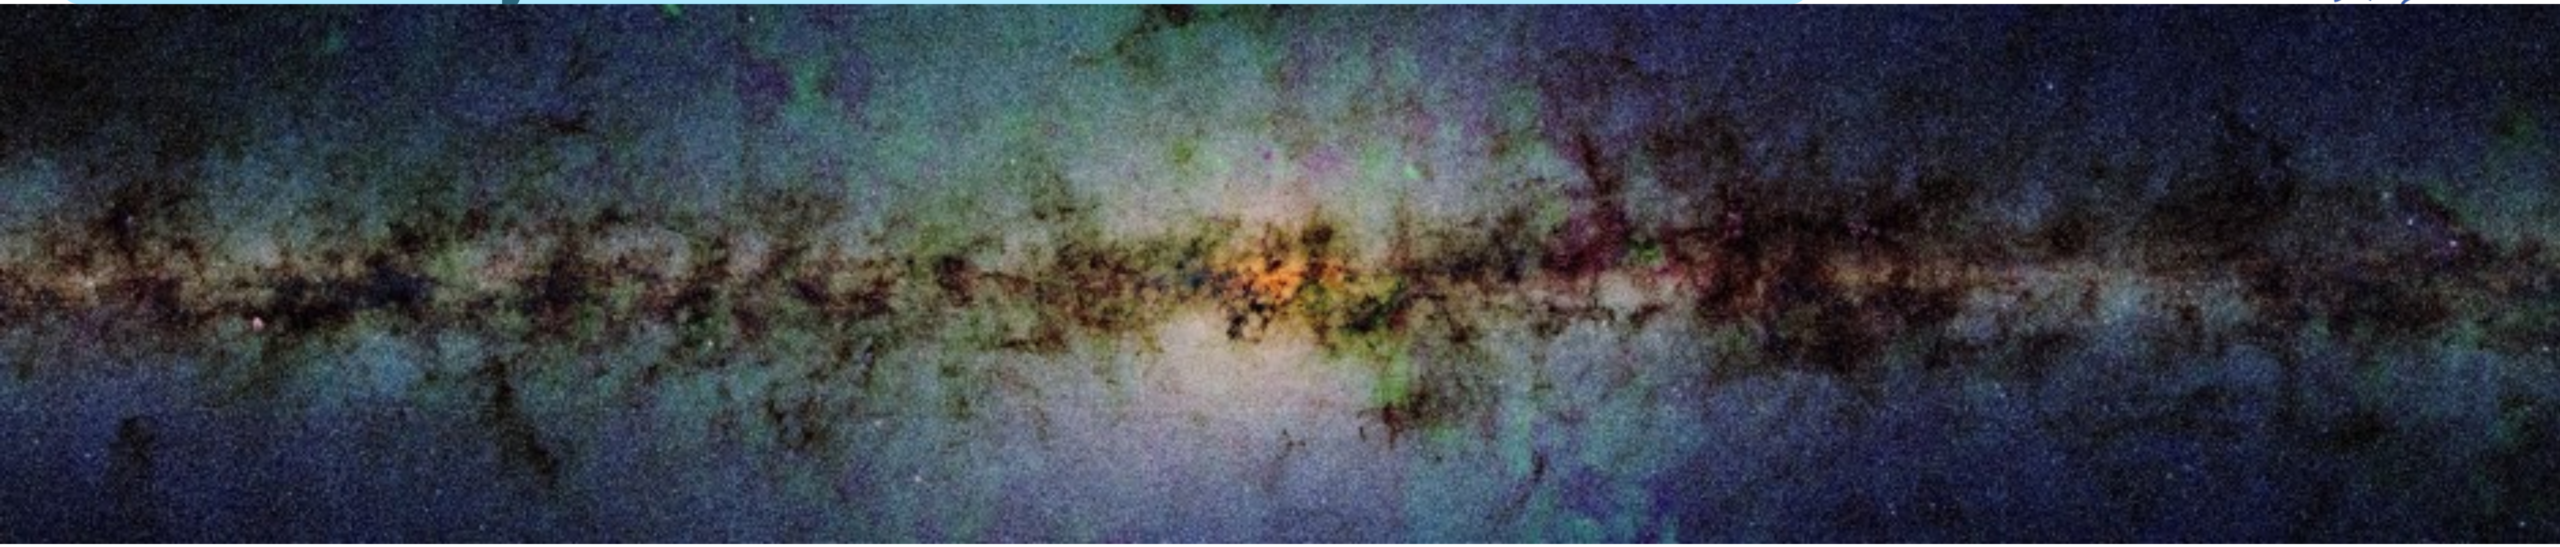

- 18 million input images (~2.5 TB)
- 900 output images (2.5 GB each, 2.4 TB total)
- Measuring the global star formation rate in the galaxy
- Studying the energetics of the interaction of molecular clouds with the interstellar medium
- Determining whether coagulation or fragmentation governs the formation of massive stars
- Assessing the supernova rate in the Galaxy

} × 17

Data Reuse prune jobs if output data already exists

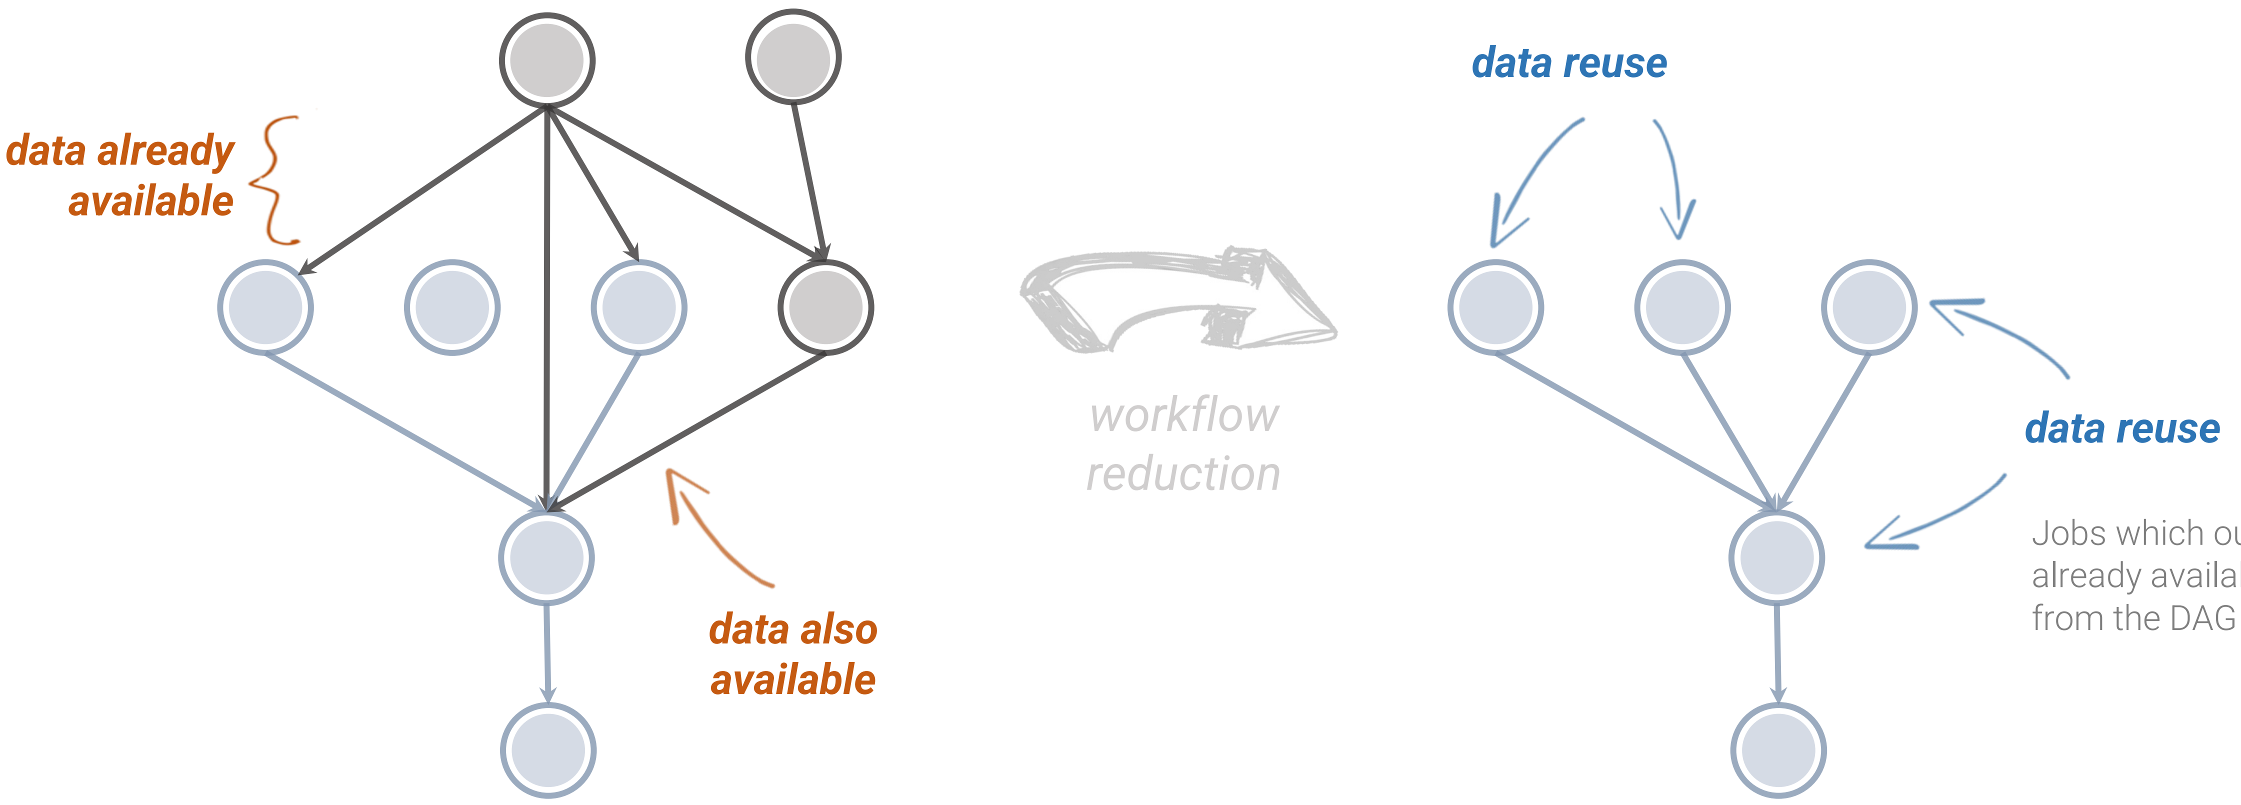

Jobs which output data is already available are pruned from the DAG

Performance optimization, Fault recovery strategy

Ensemble Manager

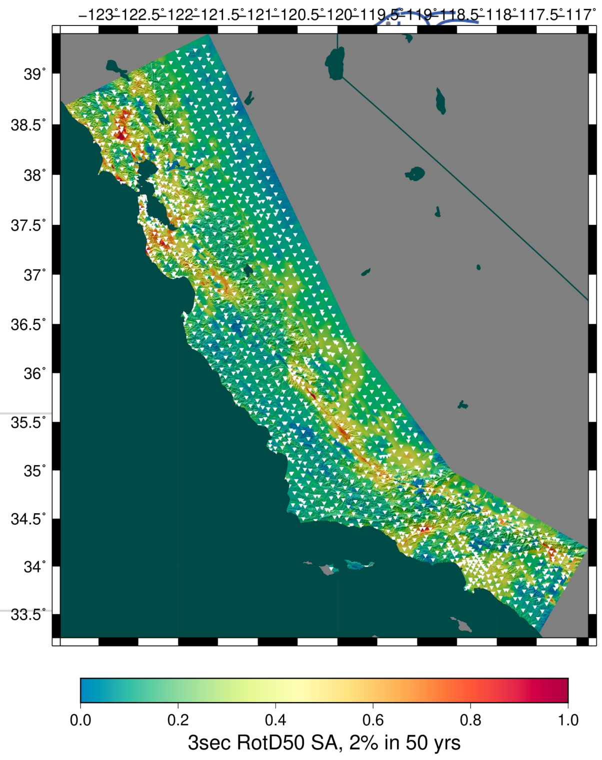

Allow users to submit a collection of workflows (ensembles)

- Automatically **spawn** and **manage** collections of workflows

Trigger submission of workflows

Cron workflow trigger

File pattern workflow trigger

Properties

- Workflows within an ensemble may have **different priorities**
  > *Priorities can also be changed at runtime*
- Ensembles may limit the number of **concurrent** planned and running workflows

Additional Actions

- Ensembles can be **paused, resumed, removed, re-planned**, and **re-executed**
- A **debugging** mechanism is also provided to investigate failures in workflow runs
- Actions can be performed both to ensembles and single workflows within ensembles

Types of Workflow Applications: **Supporting community-based analysis**

SCEC's CyberShake:
What will the peak earthquake motion be over the next 50 years?

Southern California Earthquake Center

- Codes are collaboratively developed
- Codes are "strung" together to model complex systems
- Ability to correctly connect components
- Automating the flow of data (instead of emails)
- Automatic fault recovery and support for scalability

Useful information for:

- Building engineers
- Disaster planners
- Insurance agencies

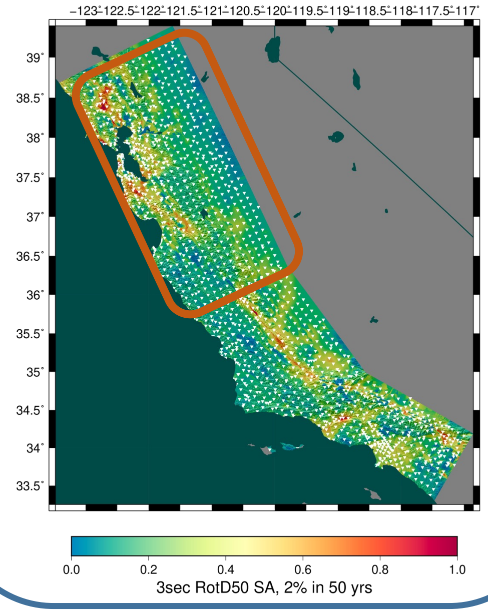

NCSA *Blue Waters*

OLCF *Titan*

- 120 million core-hours
- 39,285 jobs
- 1.2 PB of data managed
- 157 TB of data automatically transferred
- 14.4 TB of output data archived

Slide credit: USC's Southern California Earthquake Center

Increasing computational granularity:

Task clustering

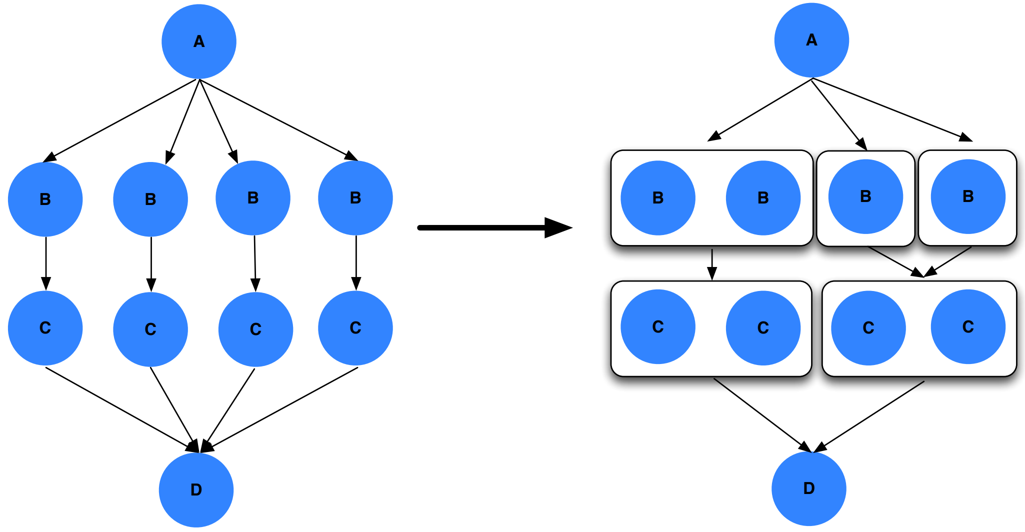

Partition the workflow into sub-workflows and send them for execution to the target system

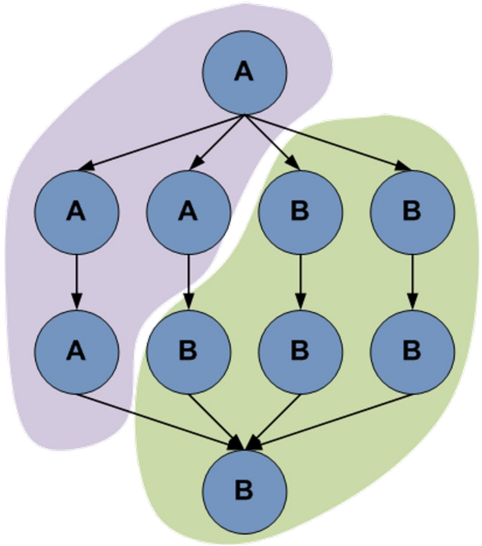

Use "pilot" jobs to dynamically provision a number of resources at a time

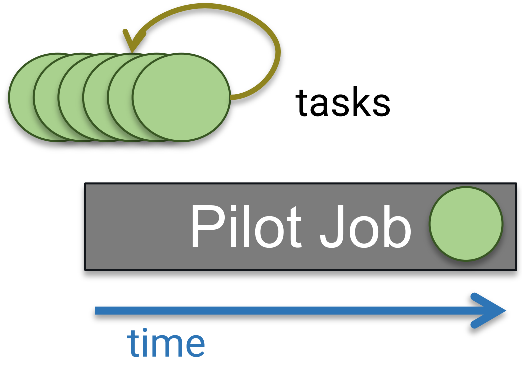

Handling heterogeneous workloads: Running HTC jobs on HPC systems…

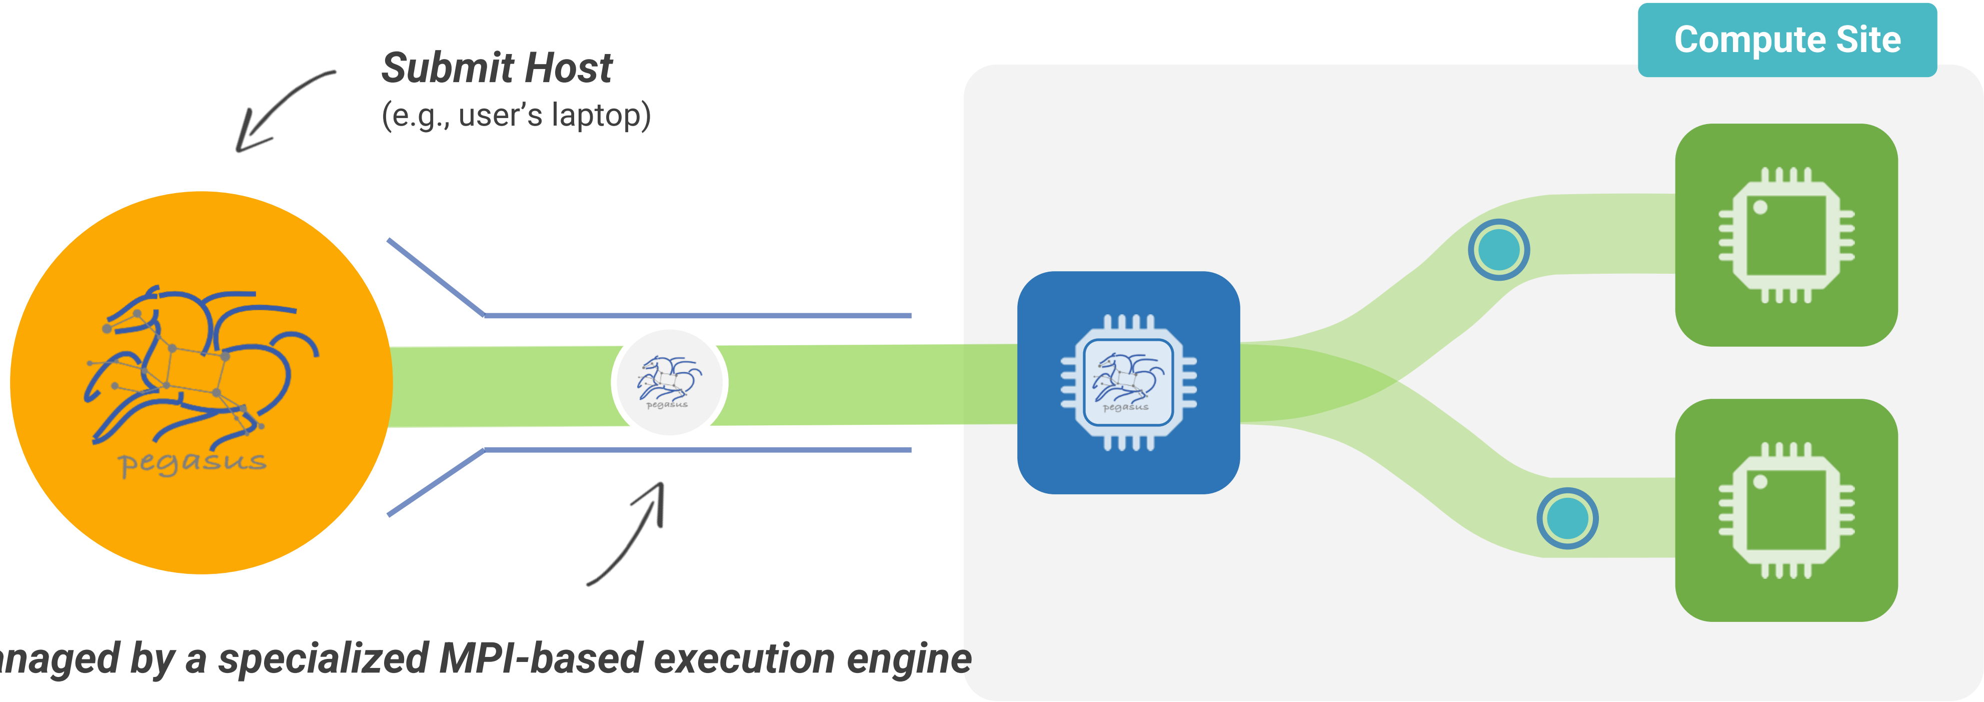

Workflow managed by a specialized MPI-based execution engine

Allows sub-graphs of a Pegasus workflow to be submitted as monolithic jobs to remote resources

Handling of large-scale workflows

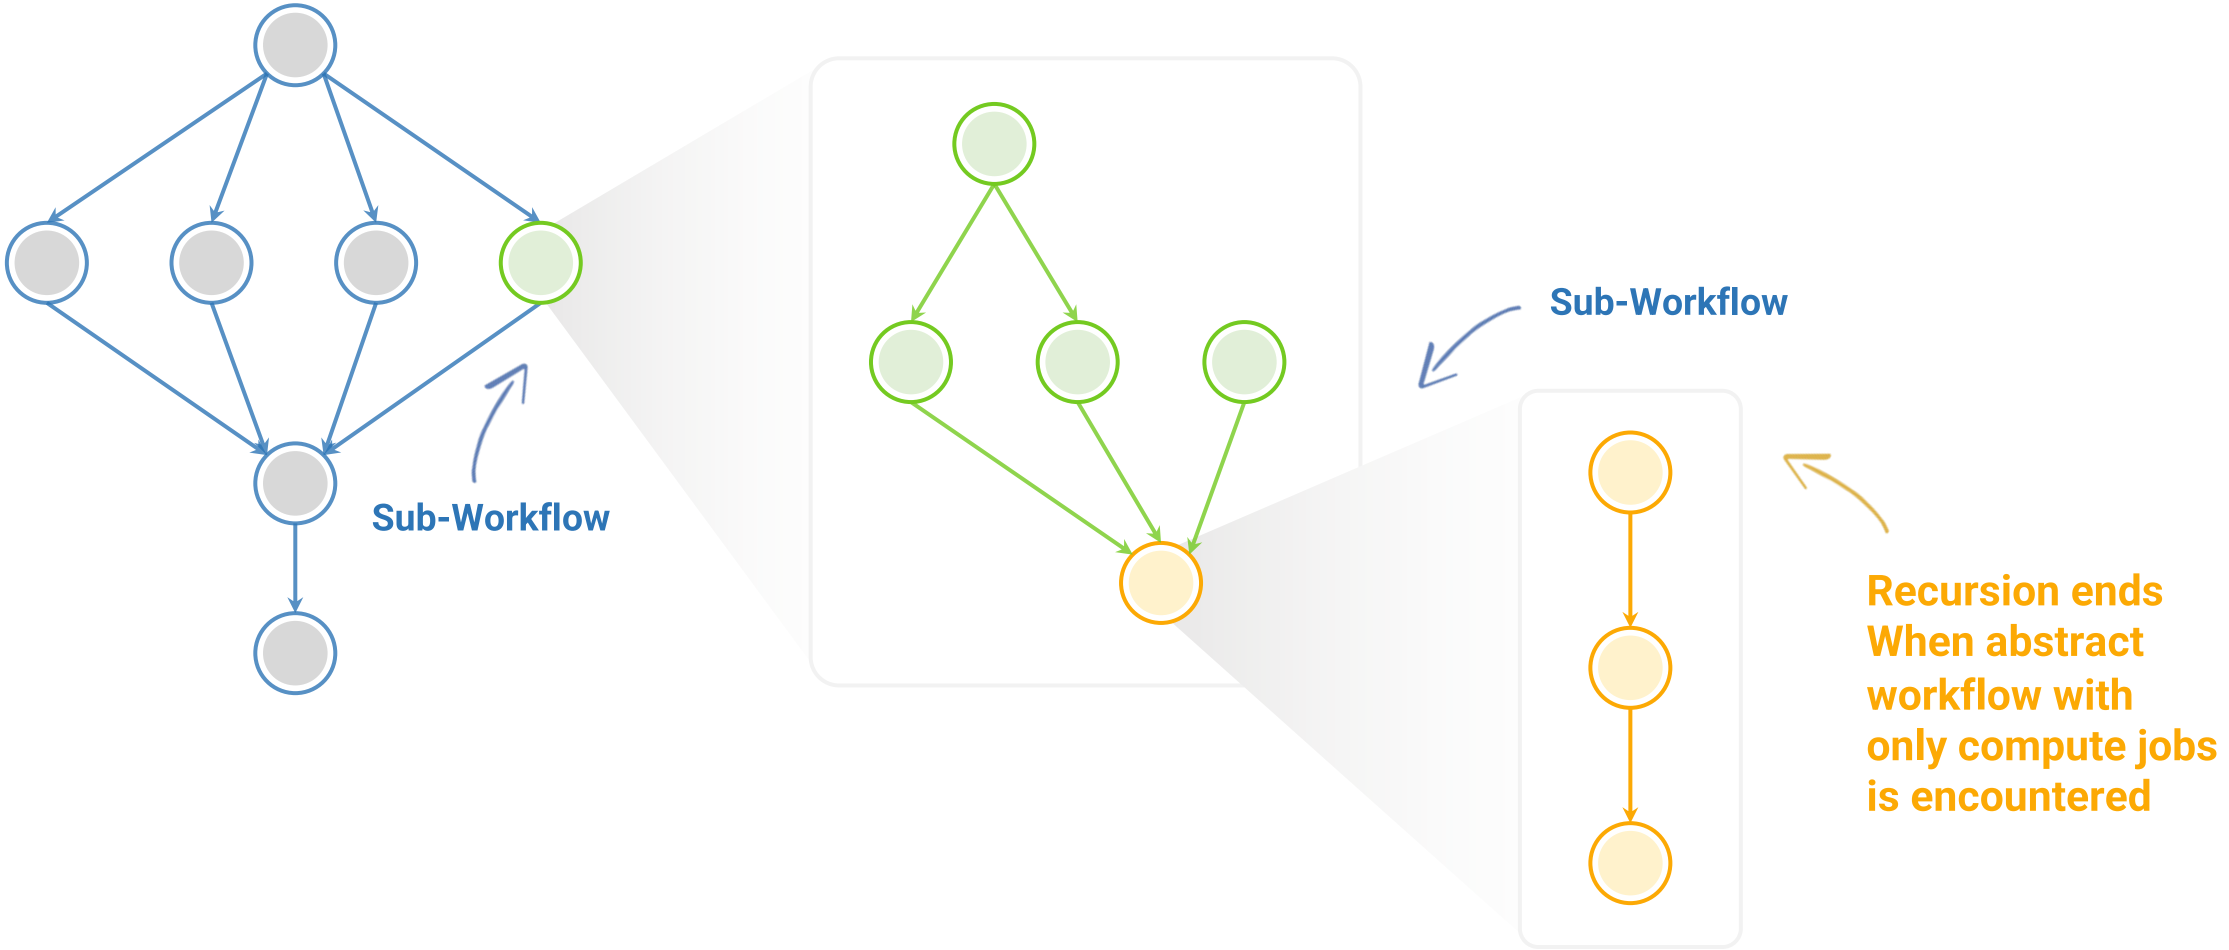

Montage Astronomy Workflow

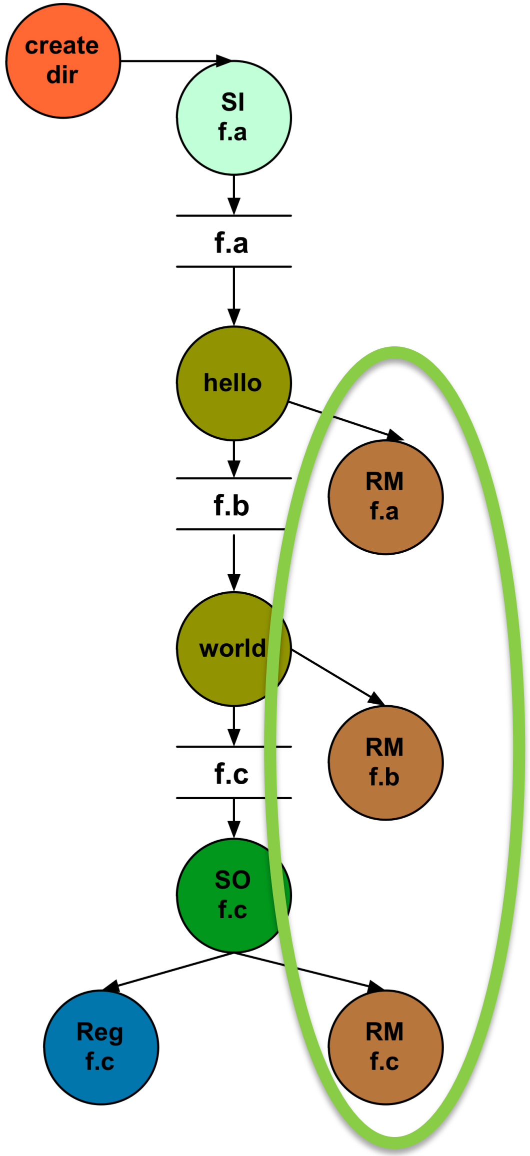

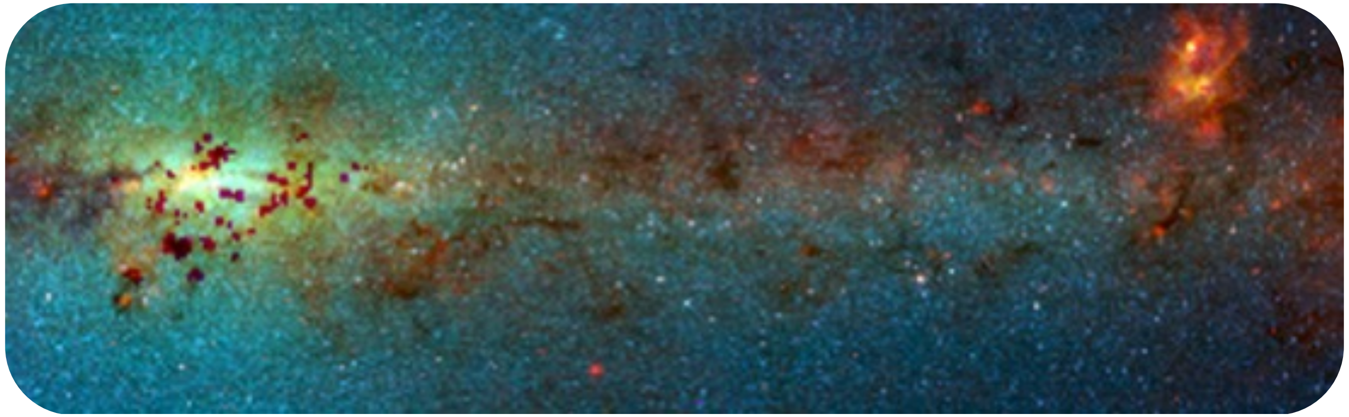

2Mass Mosaic, IPAC, Caltech

Automatically add tasks to “clean up” data no longer needed

1.25GB versus 700 MB

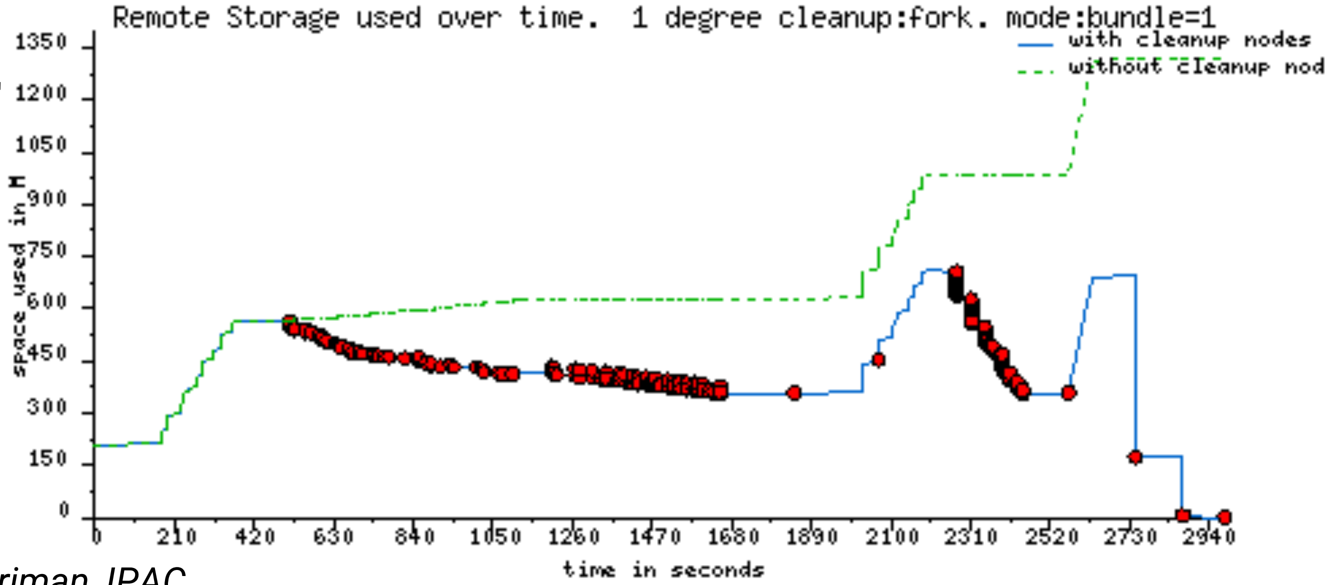

Collaboration with Rizos Sakellariou, U Manchester and Bruce Berriman, IPAC

Full workflow:
185,000 nodes 466,000 edges
10 TB of input data
1 TB of output data.

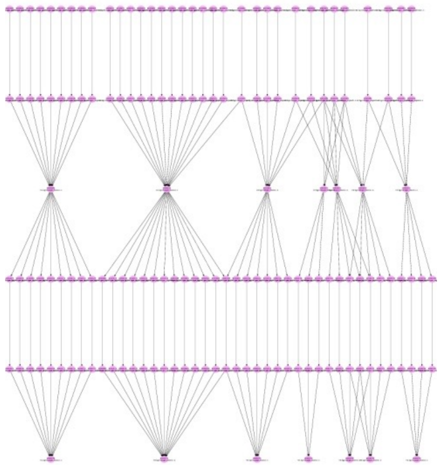

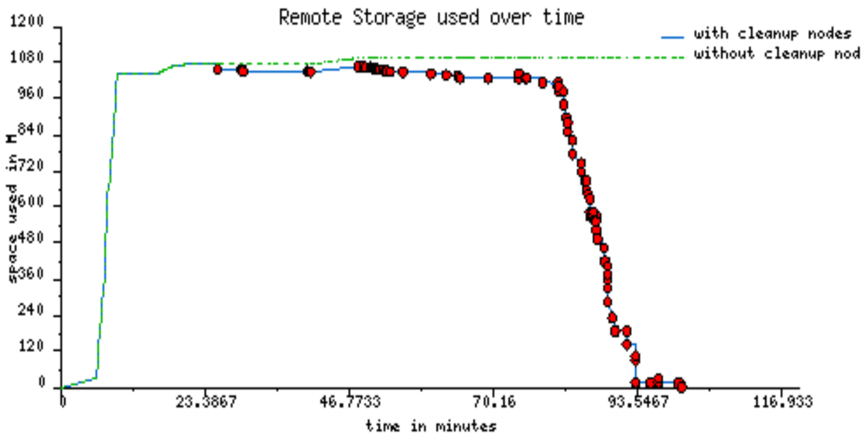

166 nodes

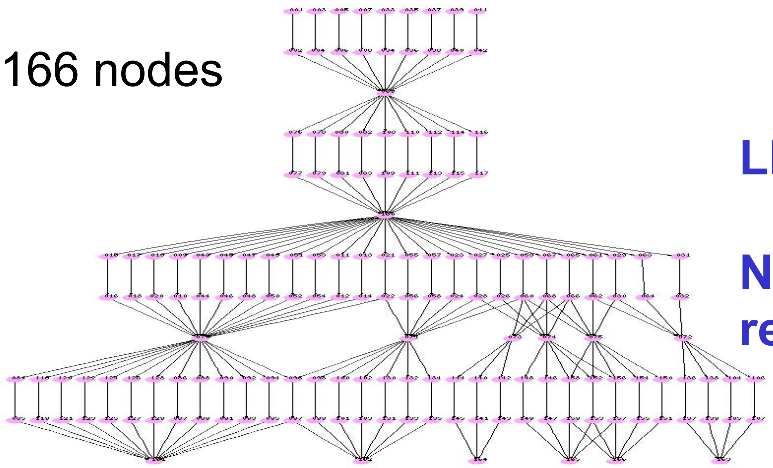

LIGO Workflows

Need additional restructuring

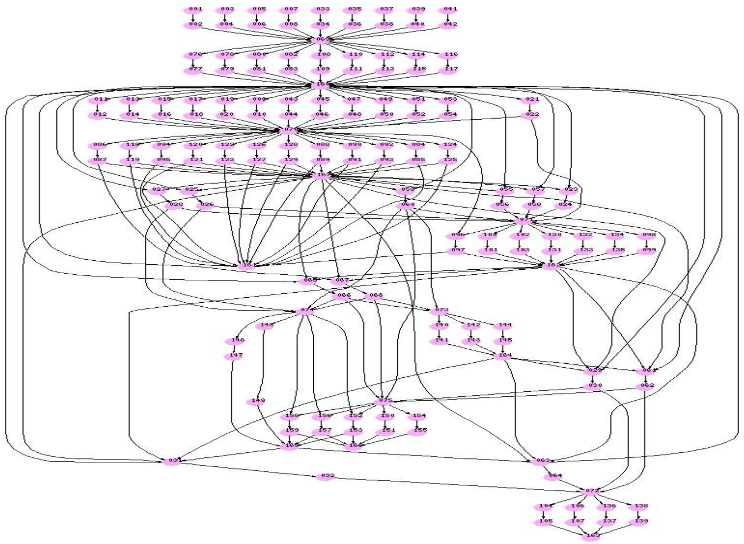

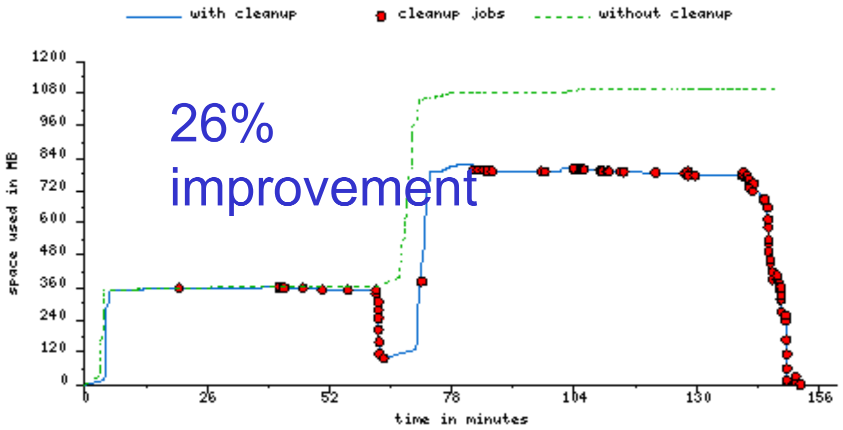

26% improvement

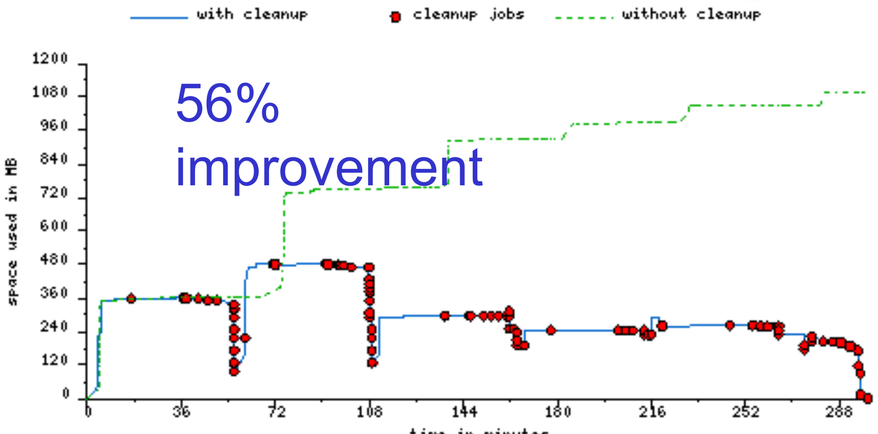

56% improvement

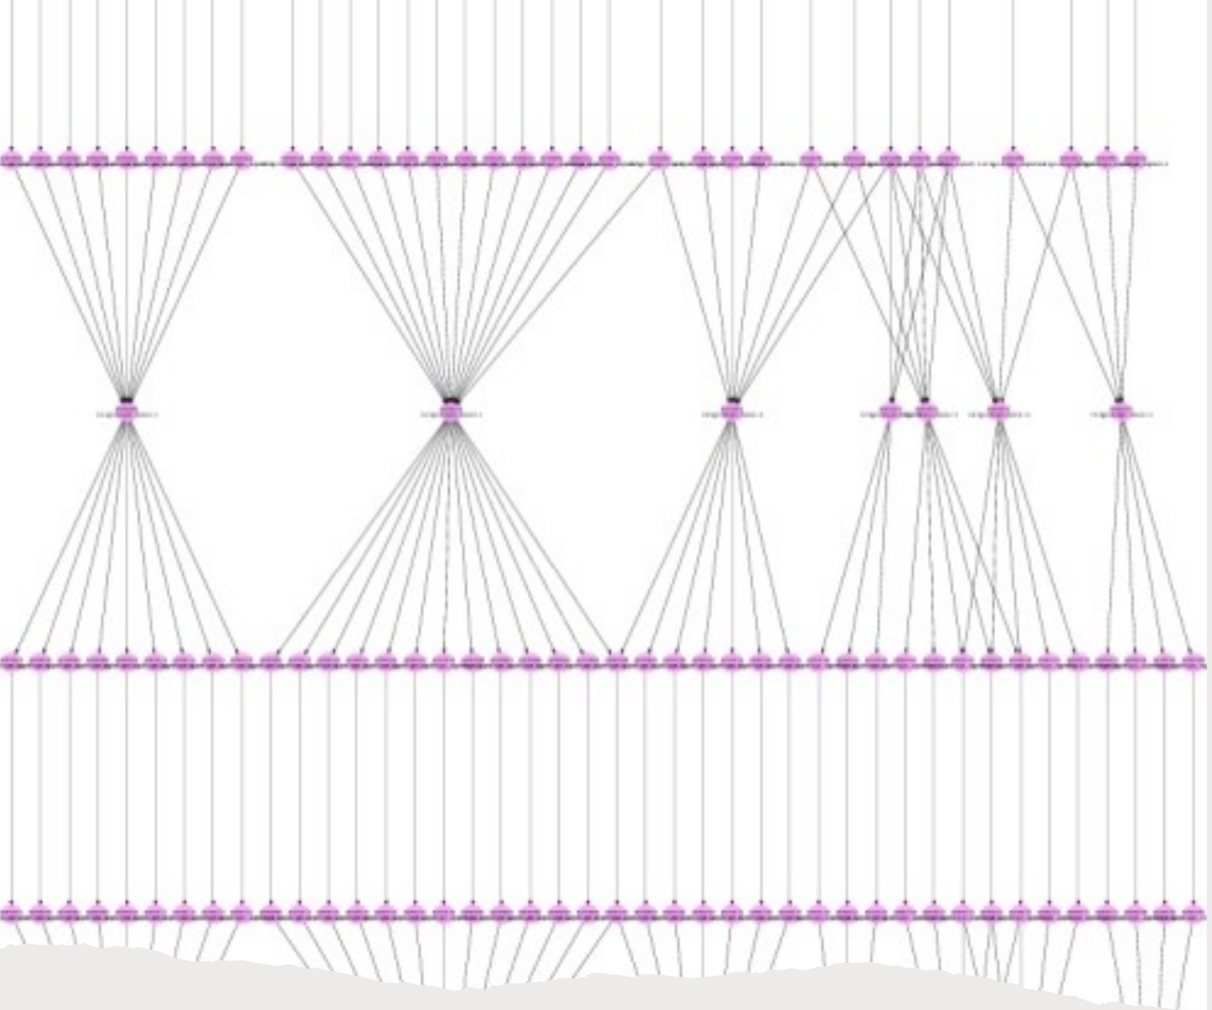

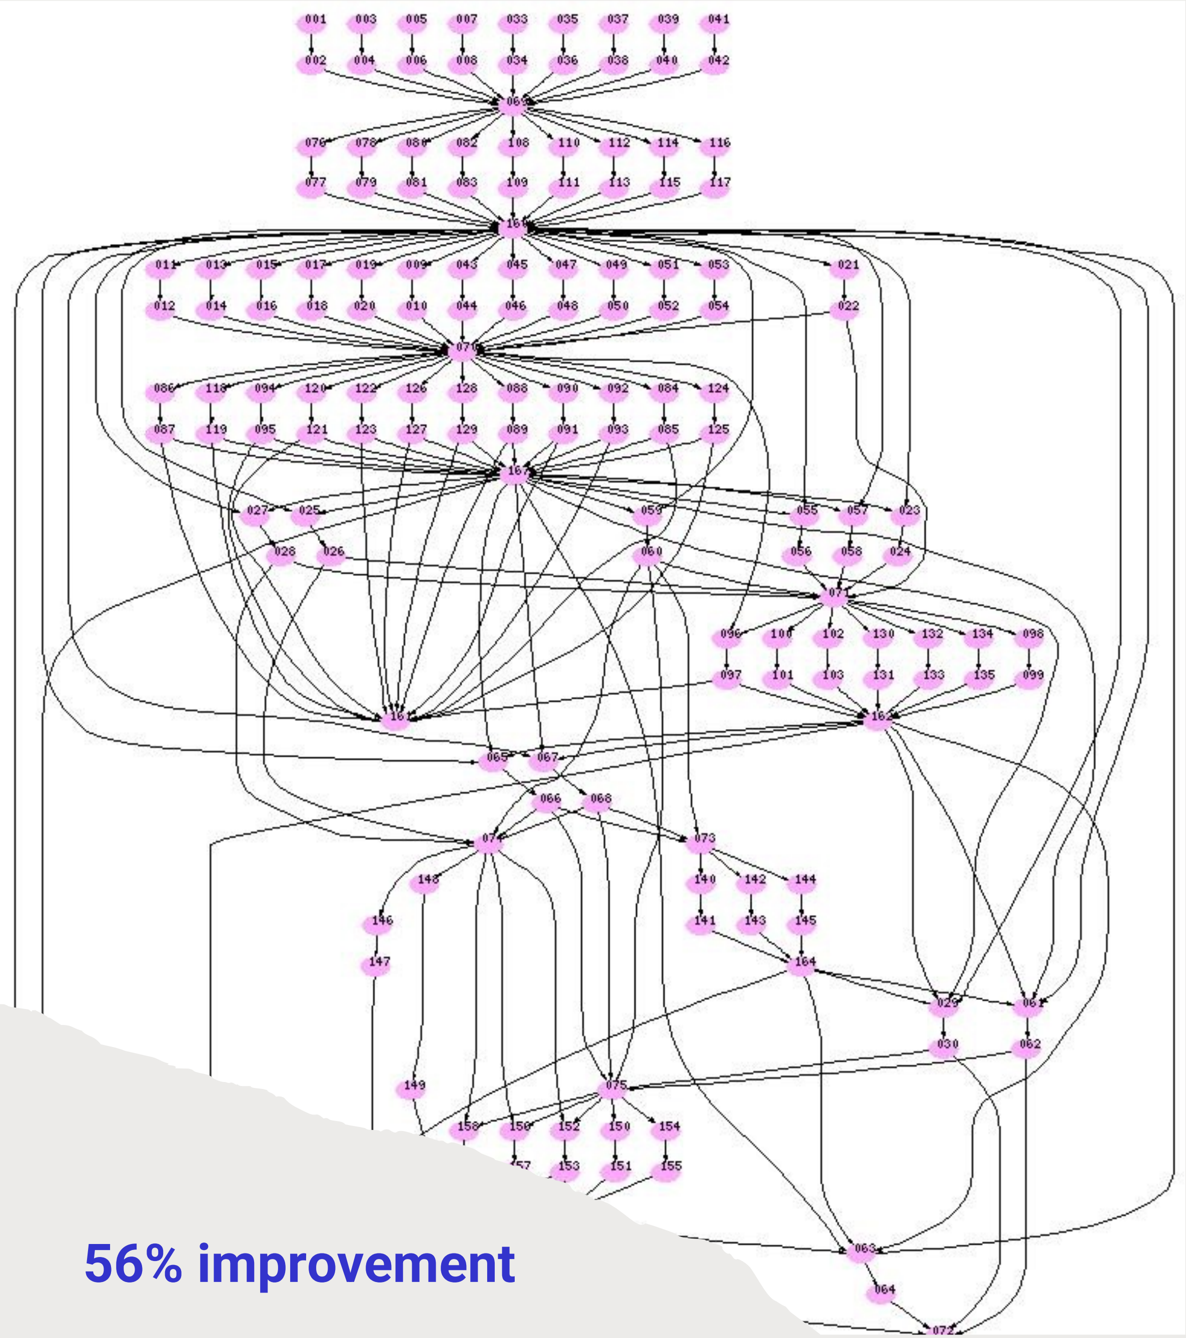

Is this a good thing?

56% improvement

We cannot forget about the users

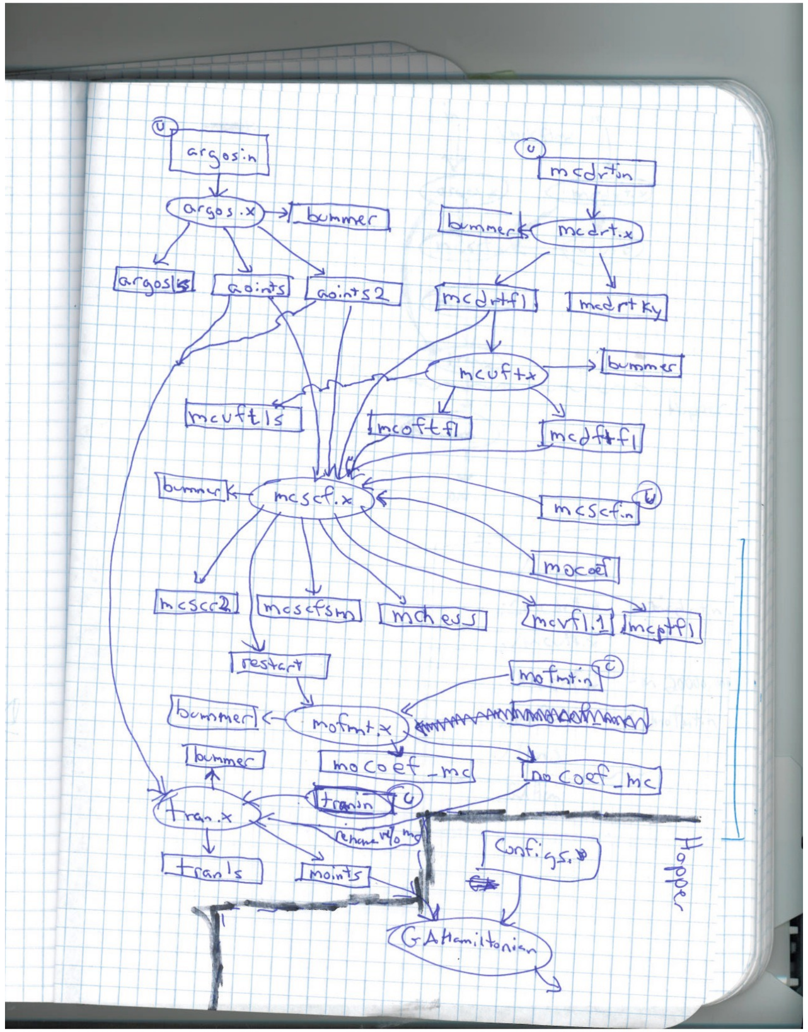

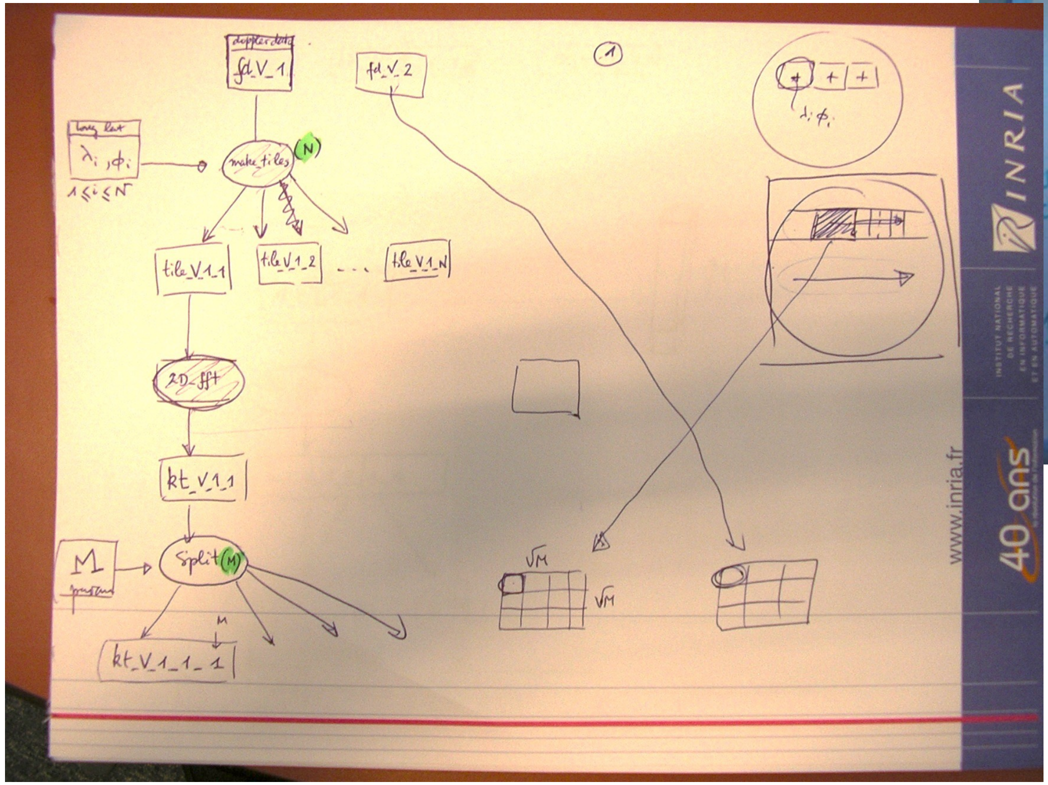

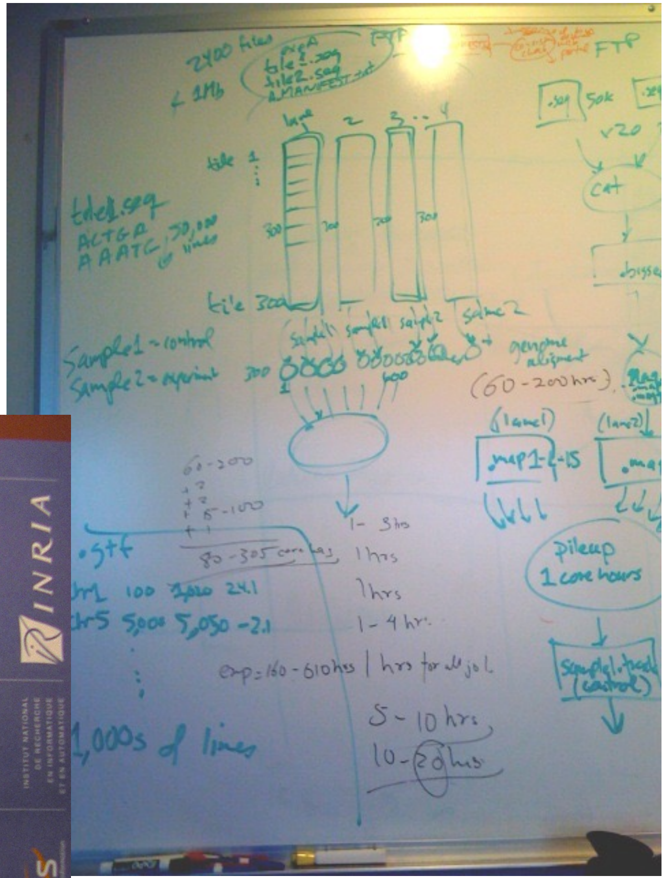

Running **Pegasus** workflows with **Jupyter**

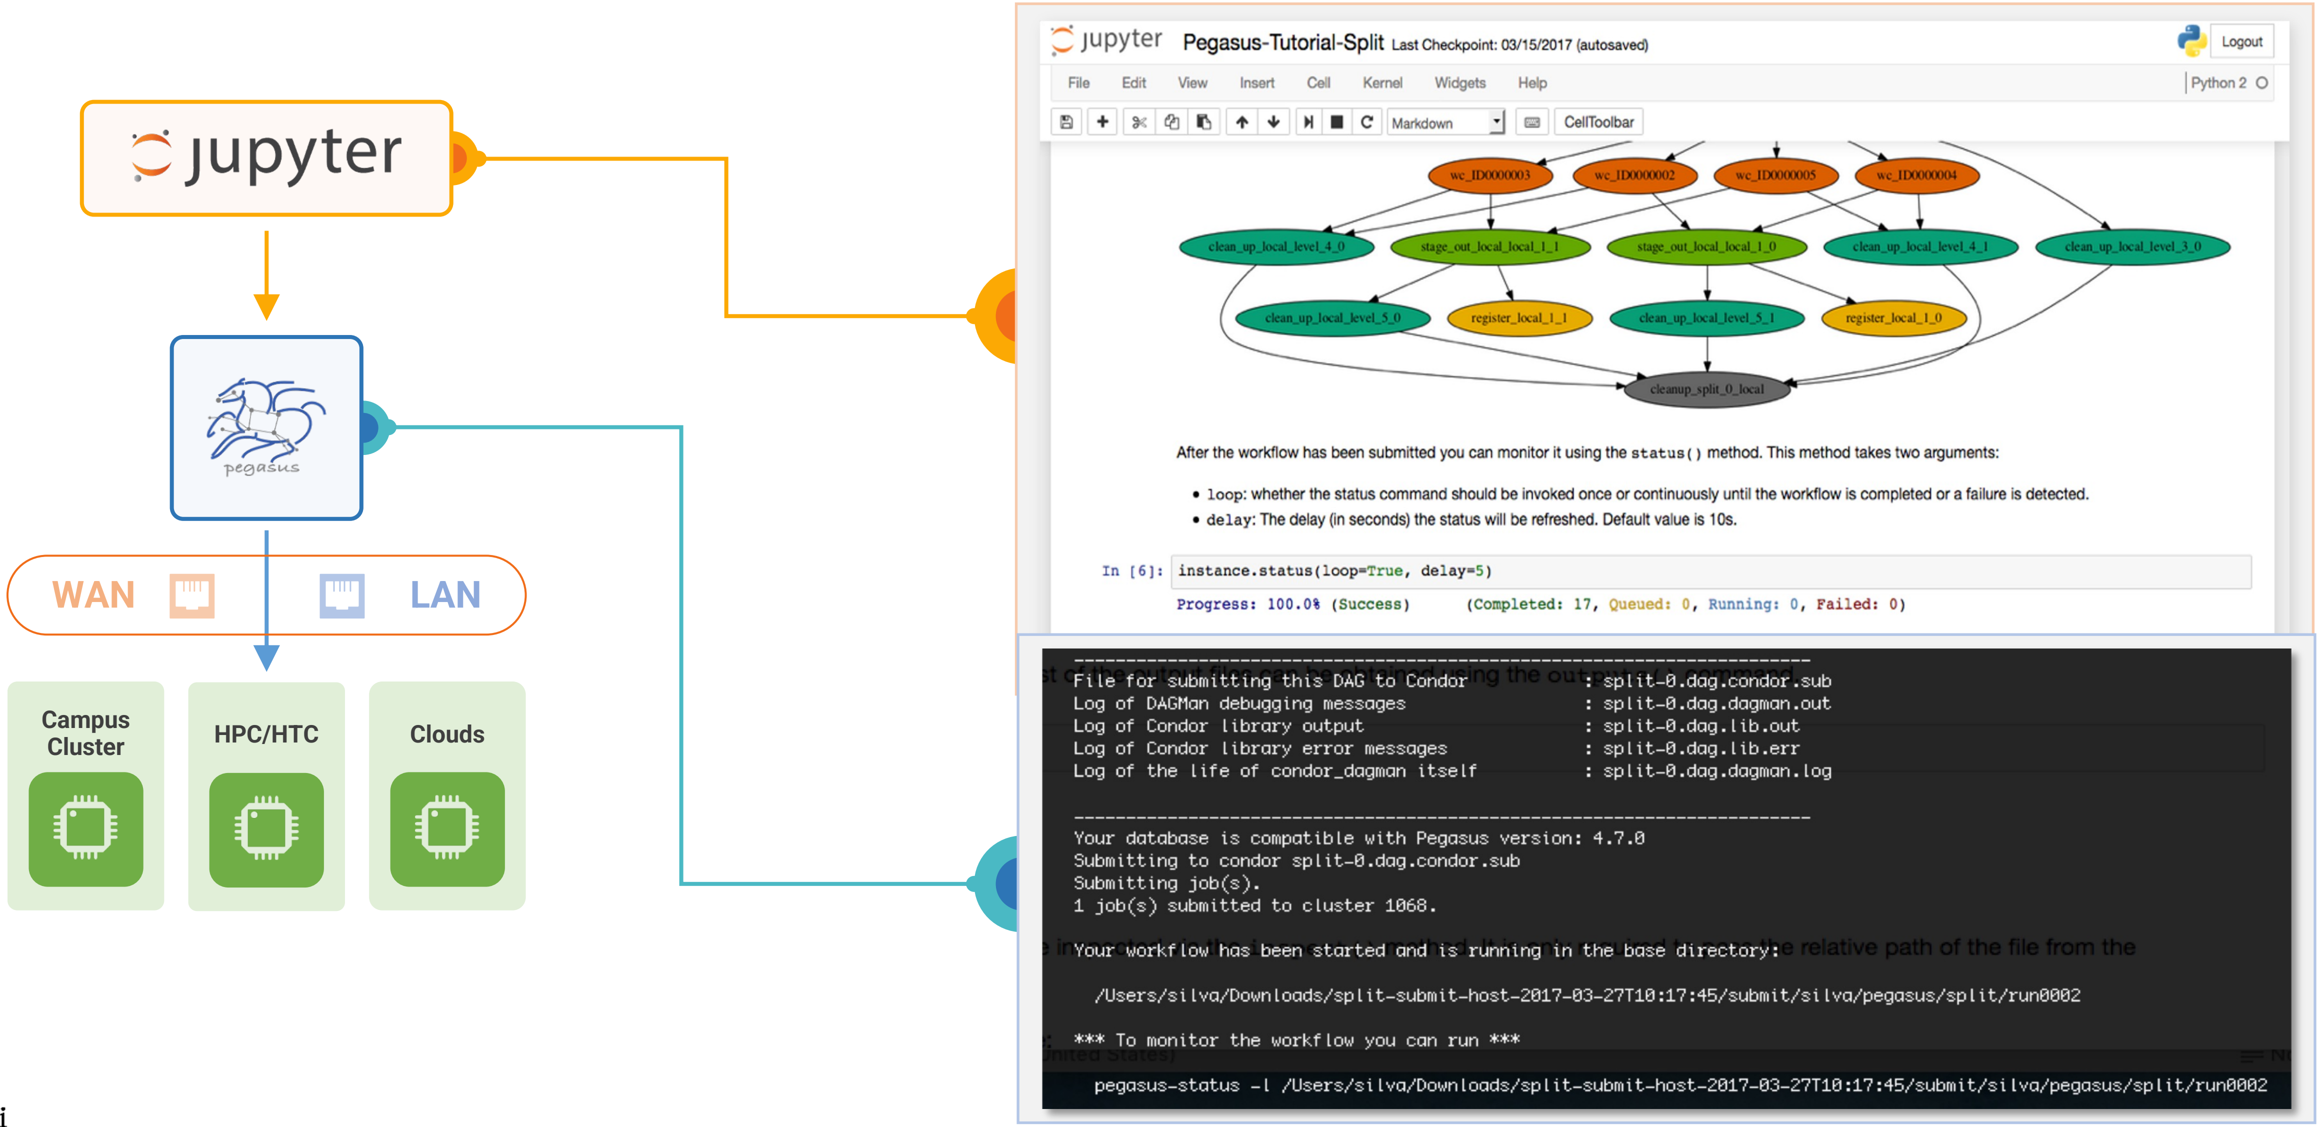

Better fault tolerance

Postscript

detects non-zero exit code output parsing for success or failure message exceeded timeout do not produced expected output files

Job Retry

helps with transient failures set number of retries per job and run

Checkpoint Files

job generates checkpoint files staging of checkpoint files is automatic on restarts

Rescue DAGs

workflow can be restarted from checkpoint file recover from failures with minimal loss

Anomaly Detection and Classification using Graph Neural Networks (GNN)

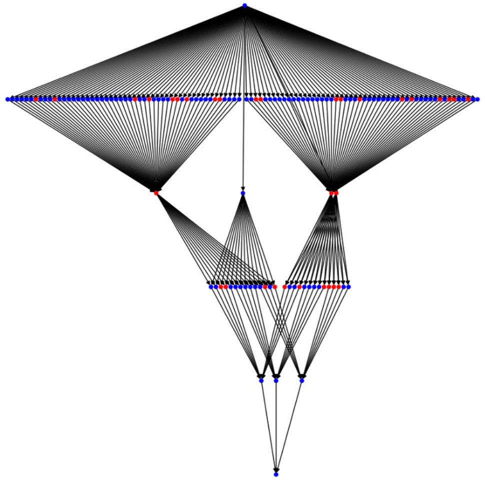

Figure: input data

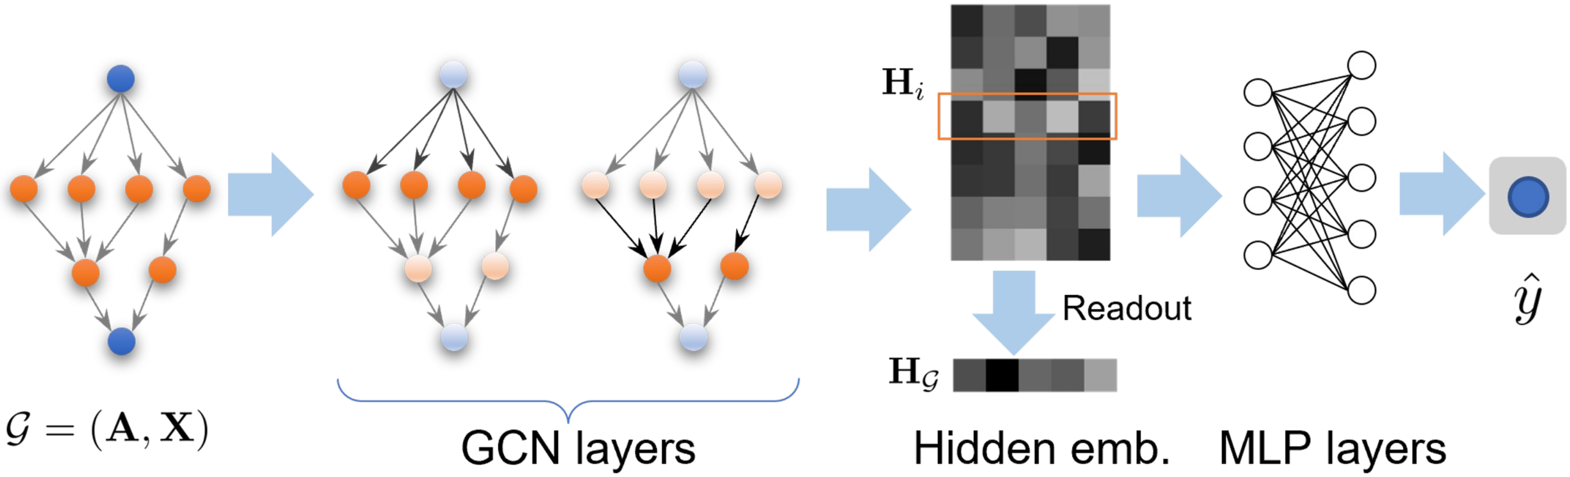

Figure: Graph neural networks architecture

Input: directed acyclic graphs (DAGs) represent simulated normal and anomaly workflows

Output: the normal/anomaly labels for workflow-level (entire graph) and job-level (single node)

Graph Neural Networks - performance

Available workflows

Single model for multi-workflows

Workflow	Binary				Multi-label
	Accuracy	F1	Recall	Precision	Accuracy
1000 Genome	0.917 ± .014	0.915 ± .019	0.921 ± .009	0.938 ± .010	0.882 ± .006
Nowcast w/ clustering 8	0.768 ± .009	0.715 ± .017	0.778 ± .023	0.768 ± .15	0.792 ± .009
Nowcast w/ clustering 16	0.837 ± .012	0.675 ± .020	0.815 ± .012	0.837 ± .011	0.830 ± .007
Wind w/ clustering casa	0.776 ± .002	0.652 ± .032	0.769 ± .021	0.776 ± .017	0.764 ± .19
Wind w/o clustering casa	0.781 ± .02	0.853 ± .013	0.800 ± .012	0.781 ± .008	0.886 ± .007
1000 Genome (partial anomaly)	1.000 ± .0	1.000 ± .0	1.000 ± .0	1.000 ± .0	1.000 ± .0
ALL	0.836 ± .006	0.878 ± .013	0.886 ± .011	0.856 ± .009	0.877 ± .008

Figure: Graph-level classification

Model	Acc.	Recall	Prec.	F1
SVM	0.622	0.622	0.667	0.550
MLP	0.874	0.874	0.875	0.874
RF	0.898	0.898	0.908	0.887
AlexNet	0.910	0.914	0.910	0.910
VGG-16	0.900	0.900	0.900	0.900
ResNet-18	0.910	0.916	0.910	0.910
Our GNN	**0.917**	**0.921**	**0.939**	**0.915**

Figure: Model comparison

SVM: Support vector machines (SVMs)
MLP: Multilayer perceptron with hidden layers (128, 128, 128)
RF: Random forest with maximum depth set to 3.
(AlexNet,…) Previous work: computer vision inspired DNN by generating Gantt charts from node features.

What can we do better? ML can make our systems "smarter"

- Anomaly detection
- Anomaly/error classification and attribution
- Predictive models of performance
- Better workflow adaptation based on failures and anomalies
- Workflow reuse or composition based on recommendations
- Many others... Can you interact with ChatGPT to create workflows?
- Challenges:
  - Collect enough (quality data, richness, balanced class representation)
  - Enough labeled data, need to augment data
  - Structure (normalize, scale, transform) the data in a way that is amenable to the application of current techniques (or develop new ones)
  - Select the appropriate ML algorithms or architectures
    - DL hyperparameter optimization (learning rate, #epochs, hidden layers, activations functions..)

Impact of AI on the Environment

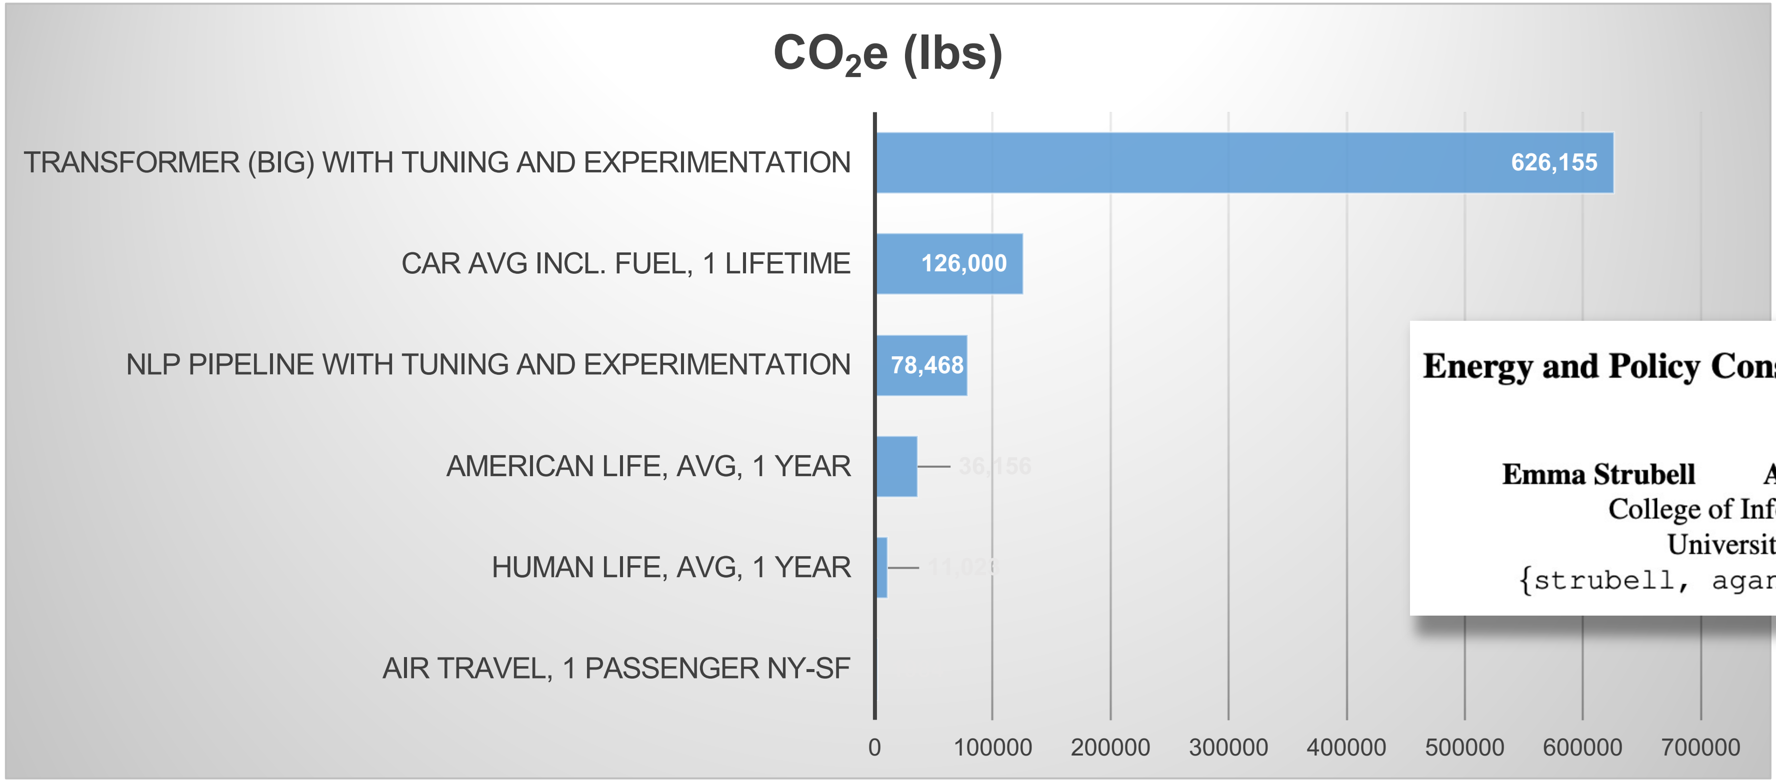

Energy and Policy Considerations for Deep Learning in NLP

Emma Strubell **Ananya Ganesh** **Andrew McCallum**
College of Information and Computer Sciences
University of Massachusetts Amherst
{strubell, aganesh, mccallum}@cs.umass.edu

https://arxiv.org/pdf/1906.02243.pdf

** Training 1 model on 1 GPU*

Growing Heterogeneity and Need for Automation

High Performance Computing Systems

- Complex
- Heterogeneous & specialized data storage
- Increasingly faulty

Distributed Systems

- Software Defined capabilities
- Programmable networks
- Specialized data storage

Clouds

- New platform for science
- Very heterogeneous
- Can be costly

EDGE DEVICES

Neuromorphic/Quantum architectures

Need to Manage The Resource Complexity

- Under constraints: resource capabilities, types, time, budget, **carbon footprint**
- Need to collect resource and application behavior/profiles and learn from them
- Faulty environment: detection and attribution
- Adaptive workload/resource management

Conclusions

- Pegasus' 5 principles for heterogeneity:
  - Resource-independent representation
  - Submit locally, run globally
  - Flexible data staging configurations
  - Flexible data movement
  - Up and down integration with other systems
- There is more to do! Workflow and resource management systems should continue to increase the level of automation and ease of use
- We need to be mindful of the environmental impact of our work

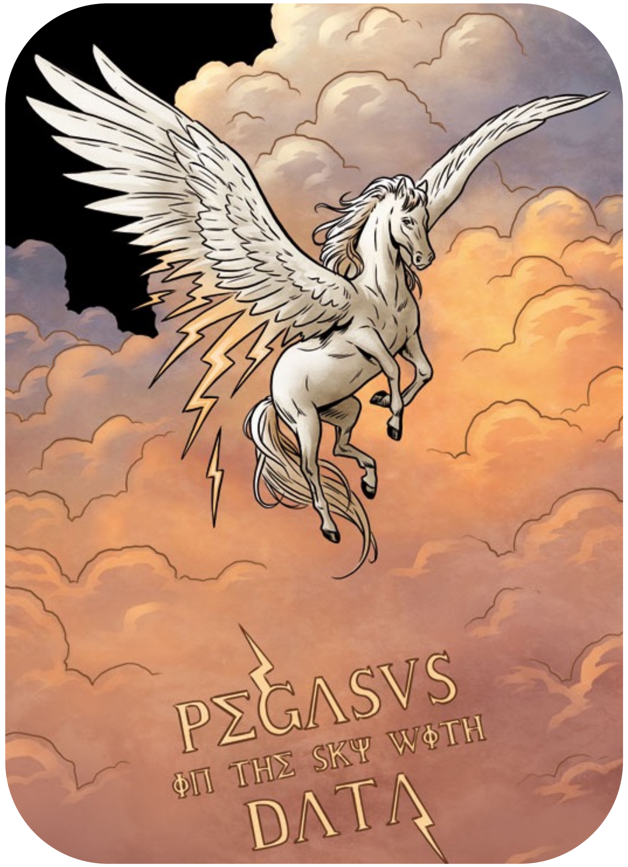

http://pegasus.isi.edu