Cloudy Skies: Astronomy and Utility Computing

Ewa Deelman
Gurmeet Singh
University of Southern California
Information Sciences Institute

Work in collaboration with Miron Livny, UW Madison
Bruce Berriman, John Good, Montage Project, IPAC, Caltech

Funded by NSF under OCI-0722019

Parts of this work to appear in the Proceedings of SC 2008

Scientific Applications

- Complex
  - Involve many computational steps
  - Require many (possibly diverse resources)
- Composed of individual application components
  - Components written by different individuals
  - Components require and generate large amounts of data
  - Components written in different languages

Issues Critical to Scientists

- Reproducibility of scientific analyses and processes is at the core of the scientific method
- Scientists consider the “capture and generation of provenance information as a critical part of <…> generated data”
- “Sharing <methods> is an essential element of education, and acceleration of knowledge dissemination.”

NSF Workshop on the Challenges of Scientific Workflows, 2006, www.isi.edu/nsf-workflows06
Y. Gil, E. Deelman et al, Examining the Challenges of Scientific Workflows. IEEE Computer, 12/2007

Computational challenges faced by applications

- Be able to compose complex applications from smaller components
- Execute the computations reliably and efficiently
- Take advantage of any number/types of resources
- Cost is an issue
  - Cluster, CI, Cloud

Possible solution

somewhat subjective

- Structure an application as a workflow
  - Describe data and components in logical terms (resource independent)
- Can be mapped onto a number of execution environments
- Can be optimized and if faults occur the workflow management system can recover
- Use a workflow management system (Pegasus-WMS) to manage the application on a number of resources

Pegasus-Workflow Management System

- Leverages abstraction for workflow description to obtain ease of use, scalability, and portability
- Provides a compiler to map from high-level descriptions to executable workflows
  - Correct mapping
  - Performance enhanced mapping
- Provides a runtime engine to carry out the instructions (Condor DAGMan)
  - Scalable manner
  - Reliable manner
- Can execute on a number of resources: local machine, campus cluster, Grid, Cloud

Science-grade Mosaic of the Sky

Point on the sky, area

Image Courtesy of IPAC, Caltech

Generating mosaics of the sky (Bruce Berriman, Caltech)

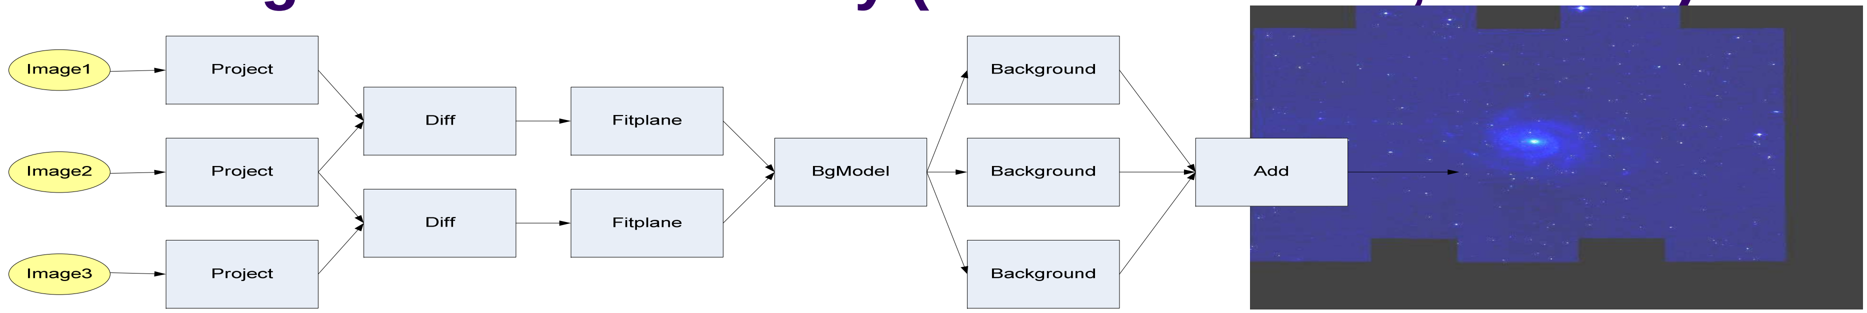

Size of the mosaic in degrees square*	Number of jobs	Number of input data files	Number of Intermediate files	Total data footprint	Approx. execution time (20 procs)
1	232	53	588	1.2GB	40 mins
2	1,444	212	3,906	5.5GB	49 mins
4	**4,856**	**747**	**13,061**	**20GB**	**1hr 46 mins**
6	8,586	1,444	22,850	38GB	2 hrs. 14 mins
10	20,652	3,722	54,434	97GB	6 hours

**The full moon is 0.5 deg. sq. when viewed form Earth, Full Sky is ~ 400,000 deg. sq.*

Montage/IPAC Situation

- Provides a service to the astronomy community
  - Delivers data to the community
  - Delivers a computational service to the community (mosaics)
- Has its own computing infrastructure
  - Invests ~ $75K for computing (over 3 years)
  - Appropriates ~ $50K in human resources every year to maintain hardware and do programming
- Expects to need additional resources to deliver services
- Wants fast responses to user requests

Cloudy Questions

- Applications are asking:
  - What are Clouds?
  - How do I run on them?
- How do I make good use of the cloud so that I use my funds wisely?
  - And how do I explain Cloud computing to the purchasing people?
- How many resources do I allocate for my computation or my service?
- How do I manage data transfer in my cloud applications?
- How do I manage data storage—where do I store the input and output data?

Montage Infrastructure

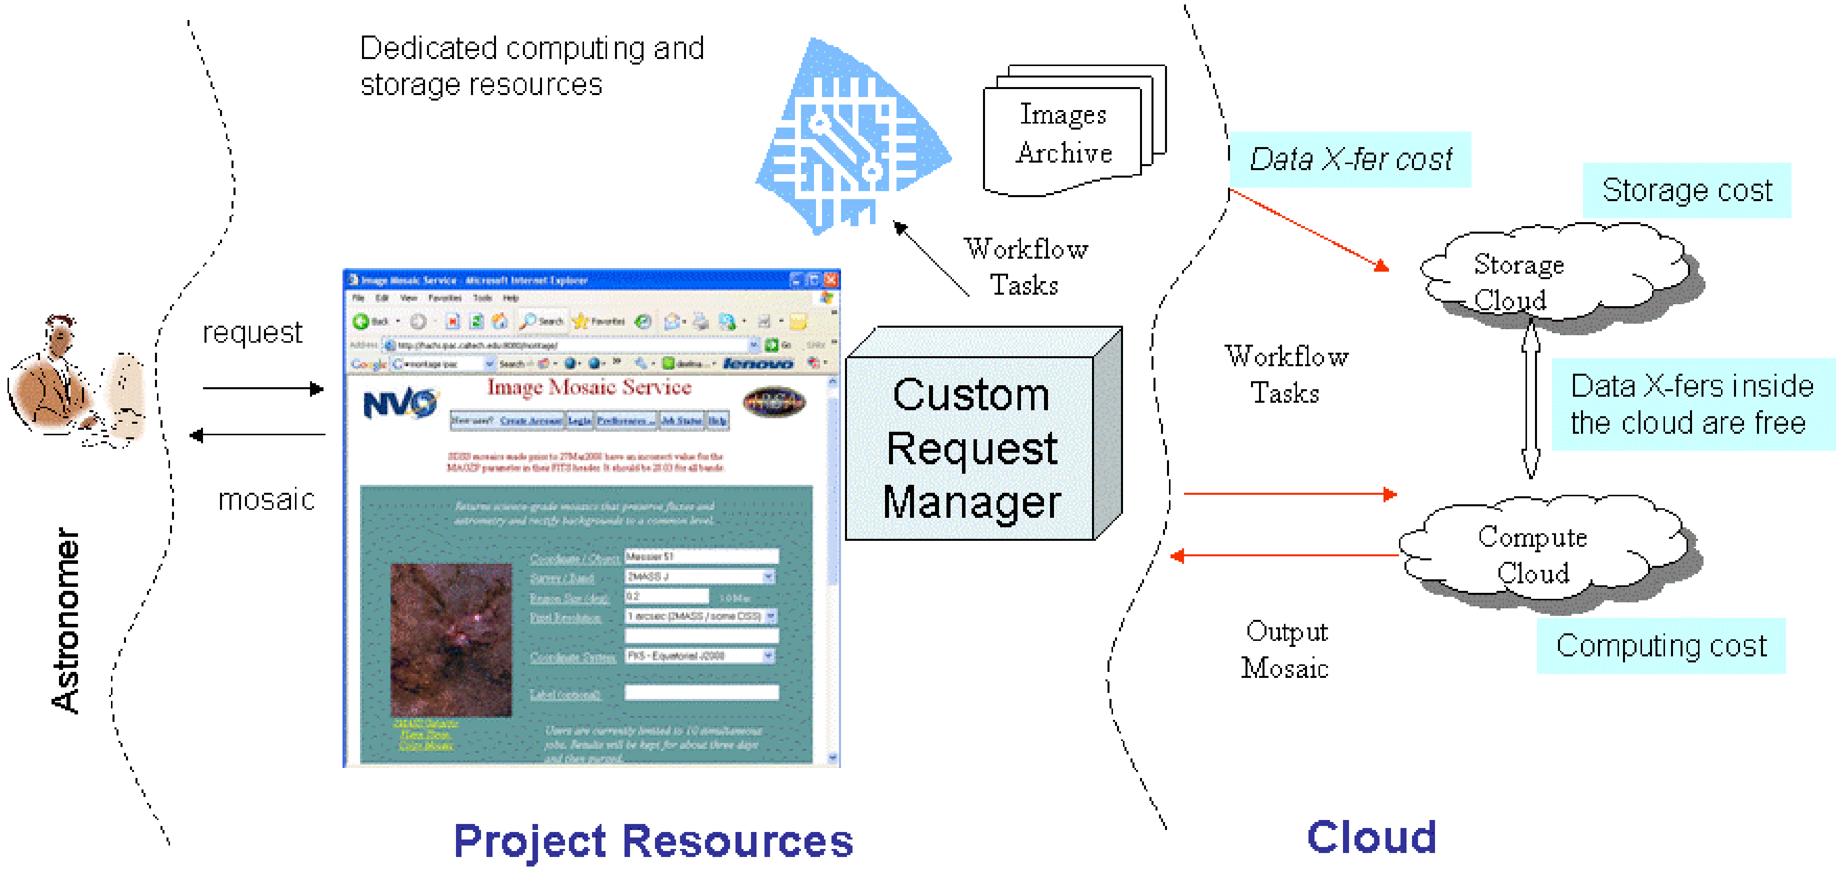

Computational Model

- Simulations done using a modified Gridsim simulator
- Based on Amazon's fee structure
  - $0.15 per GB-Month for storage resources
  - $0.10 per GB for transferring data into its storage system
  - $0.16 per GB for transferring data out of its storage system
  - $0.10 per CPU-hour for the use of its compute resources
- Normalized to cost per second
- Does not include the cost of building and deploying an image

How many resources to provision?

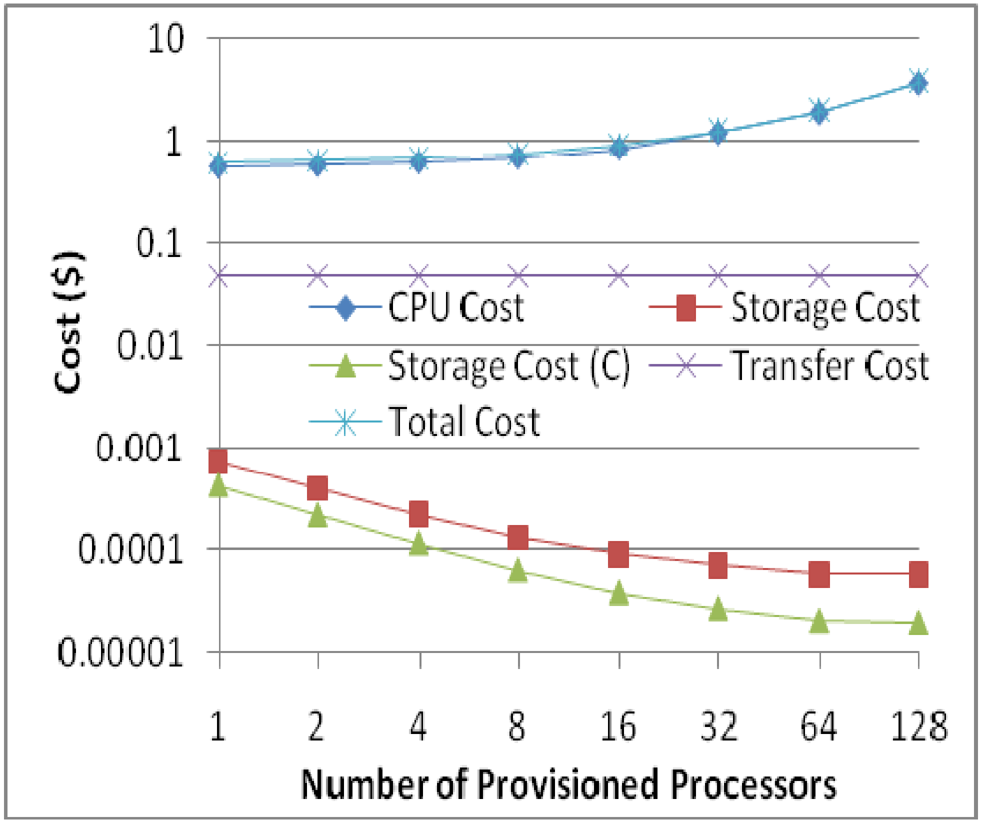

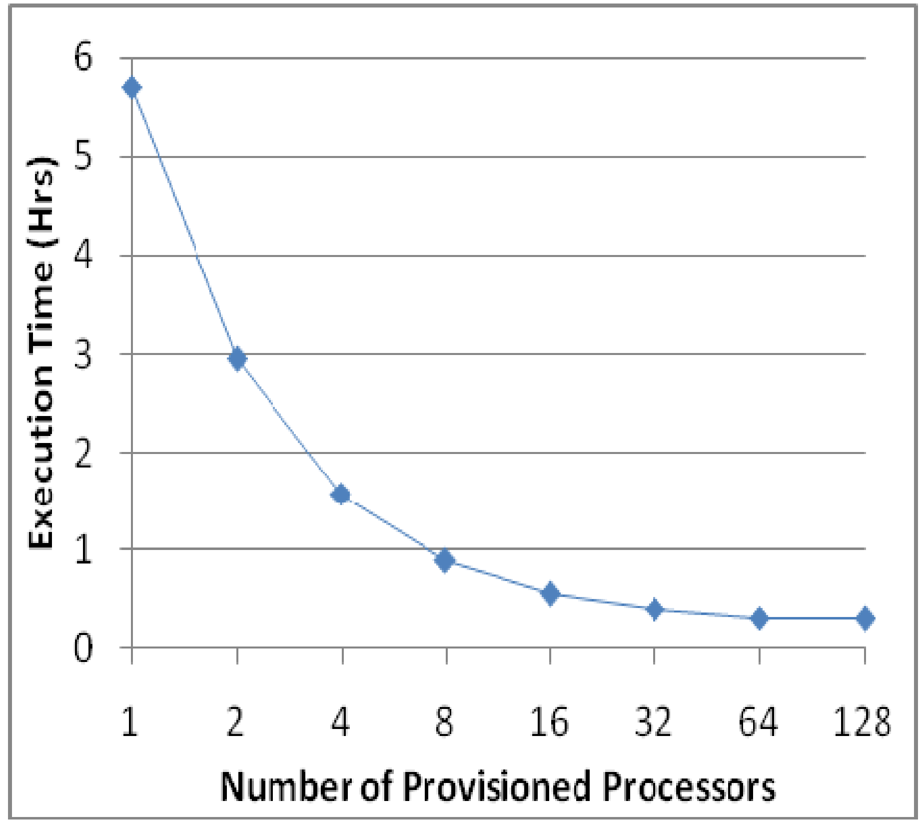

Montage 1 Degree Workflow 203 Tasks

60 cents for the 1 processor computation versus almost $4 with 128 processors, 5.5 hours versus 18 minutes

4 Degree Montage

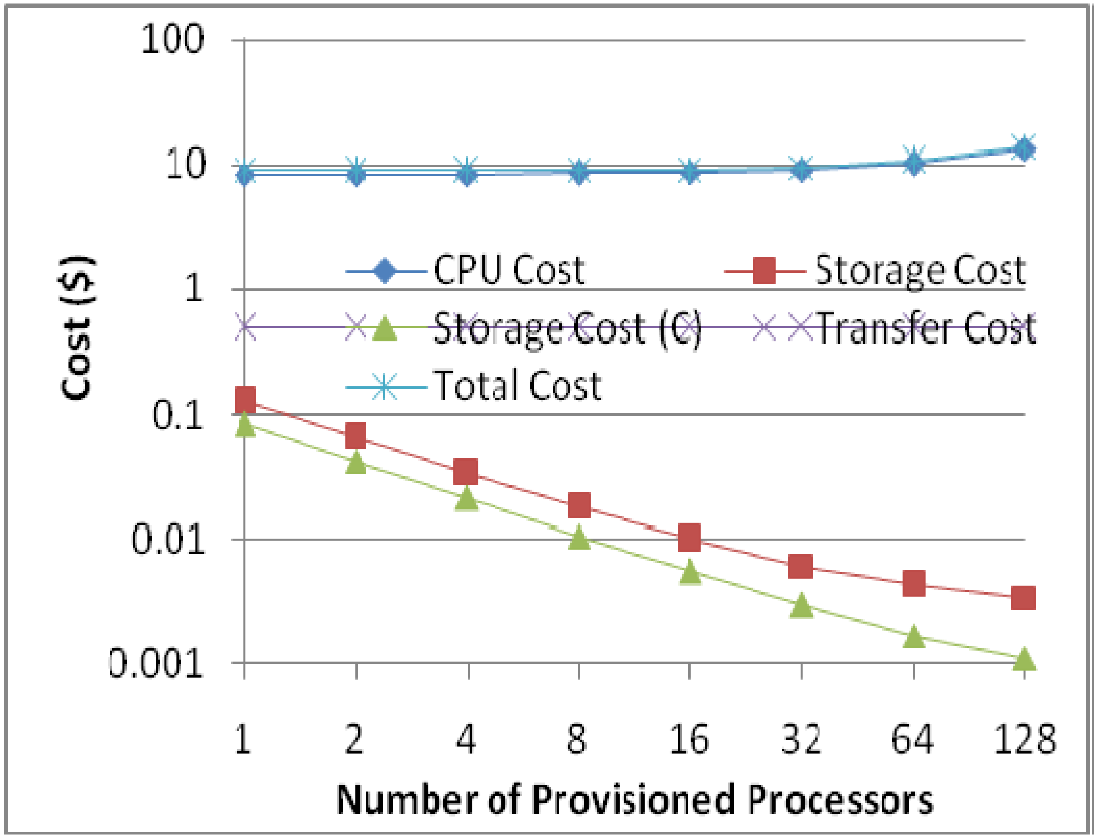

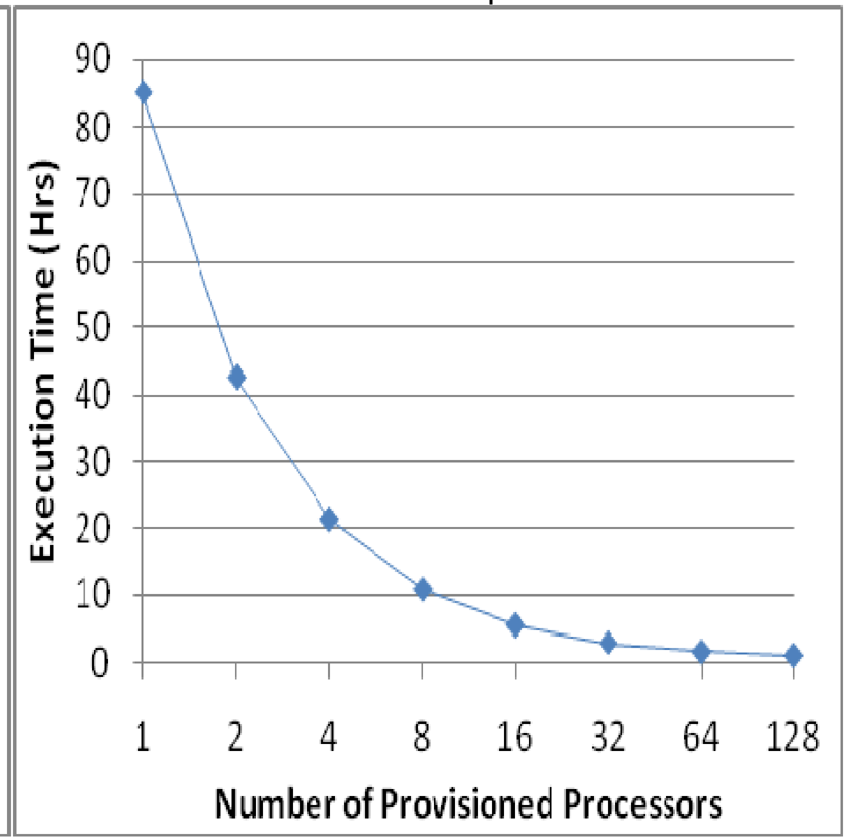

3,027 application tasks

1 processor $9, 85 hours; 128 processors, 1 hour with and $14.

How to manage data?

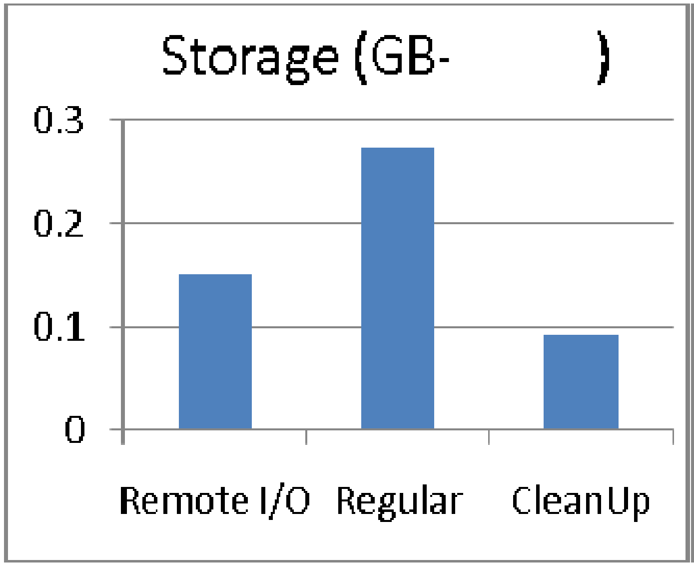

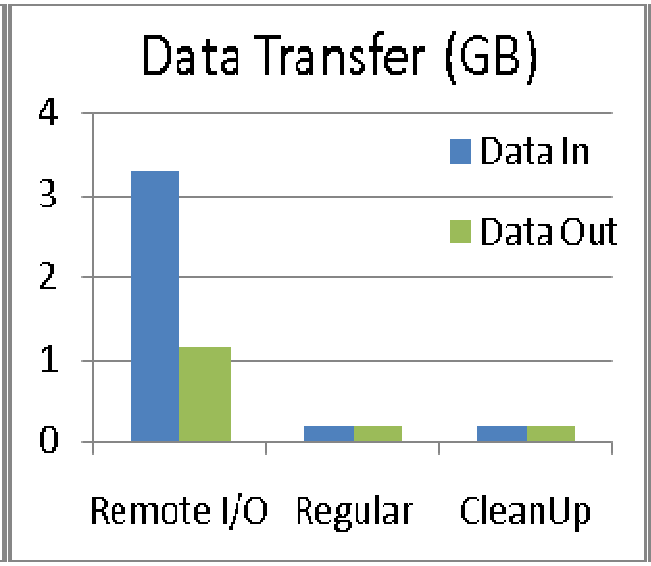

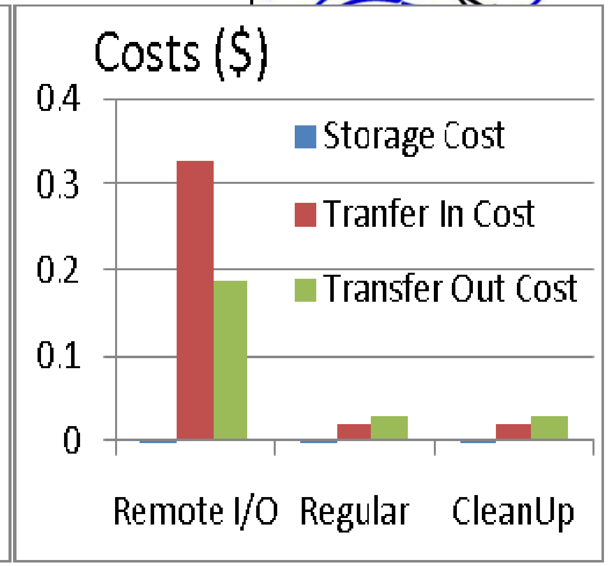

1 Degree Montage

4 Degree Montage

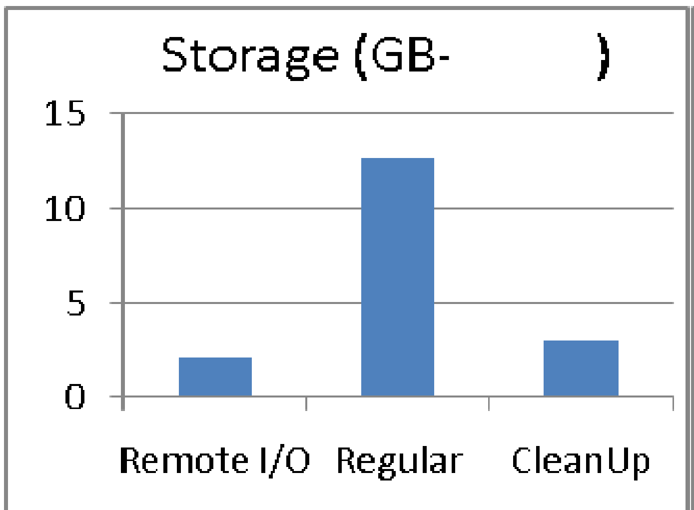

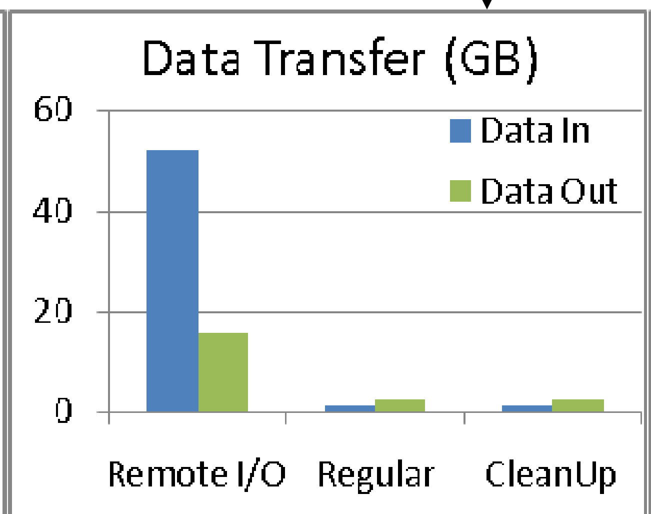

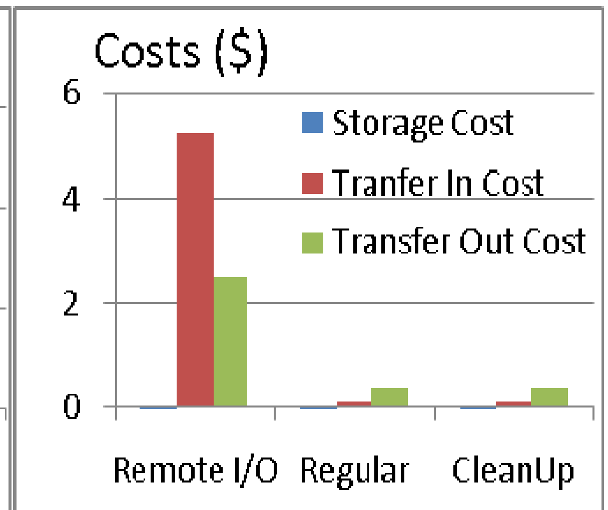

How do data cost affect total cost?

- Data stored outside the cloud
- Computations run at full parallelism
- Paying only for what you use
  - Assume you have enough requests to make use of all provisioned resources

Cost in $

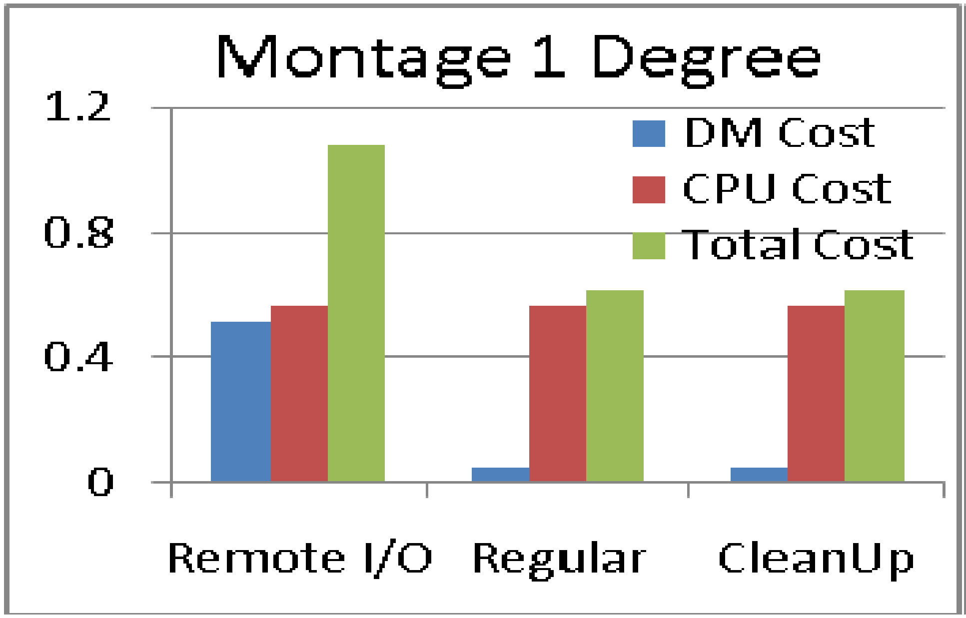

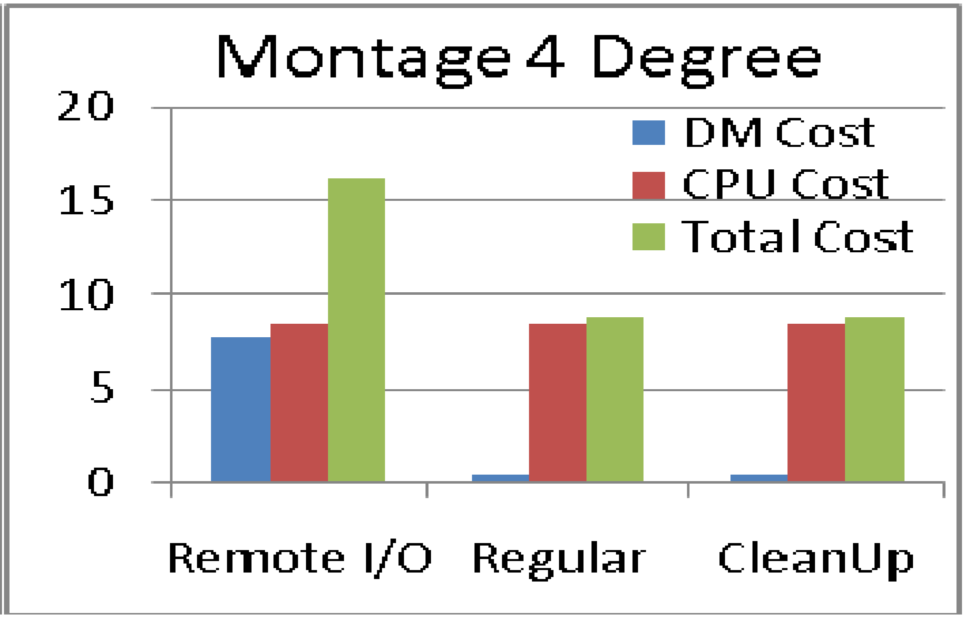

Where to keep the data?

- Storing all of 2 Mass data
  - 12 TB of data → $1,800 per month on the Cloud
- Calculating a 1 degree mosaic and delivering it to the user $2.22 (with data outside the cloud)
- Same mosaic but data inside the cloud: $2.12
- To overcome the storage costs, users would need to request at least $1,800/($2.22-$2.12) = 18,000 mosaics per month
- Does not include the initial cost of transferring the data to the cloud, which would be an additional $1,200
- Is $1,800 per month reasonable?
  - ~$65K over 3 years (does not include data access costs from outside the cloud)
  - Cost of 12TB to be hosted at Caltech $15K over 3 years for hardware

The cost of doing science

- Computing a mosaic of the entire sky (3,900 4-degree-square mosaics)
  - 3,900 x $8.88 = $34,632
- How long it makes sense to store a mosaic?
  - Storage vs computation costs

	Cost of generation	Mosaic size	Length of time to save
1 degree^2	$0.56	173MB	21.52 months
2 degree^2	$2.03	558MB	24.25 months
4 degree^2	$8.40	2.3GB	25.12 months

Remember virtual data from GriPhyN? Now we can quantify things a bit better.

Summary

- We started asking the question of how can applications/scientific workflows best make use of clouds
- Assumed a simple cost model based on the Amazon fee structure
- Conducted simulations
  - Need to find balance between cost and performance
  - Computational cost outweighs storage costs in the case of Montage
  - Storage can be expensive in the long run
- Did not explore issues of data security and privacy, reliability, availability, ease of use, etc

Will scientific applications move into clouds? (a myopic view)

- There is interest in the technology from applications
- They often don't understand what are the implications of using the new technology
- They need support from CS folks to navigate the dark clouds
- Need tools to manage the cloud
  - Build and deploy images
  - Request the right number of resources
  - Manage costs of individual computations
  - Manage project *costs—does your purchasing department know about Cloud computing?*
- Projects need to perform cost/benefit analysis

Issues Critical to Scientists

- Reproducibility – yes—maybe--through virtual images, if we package the entire environment, the application and the VMs behave
- Provenance – still need tools to capture what happened
- Sharing – can be easier to share entire images and data
  - Data could be part of the image

Getting tired of Clouds?

Black Hole Computing
a New Computing Paradigm!

Data and
Computing

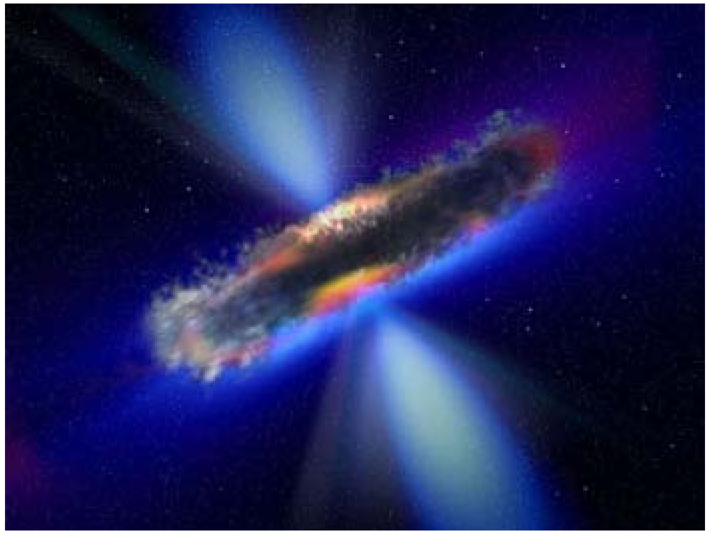

Results ?

http://www.nasa.gov/vision/universe/starsgalaxies/integral_blackholes.html

Black Hole Provisioning
New International Collaboration

LHC Experiment at CERN
http://www.cern.ch

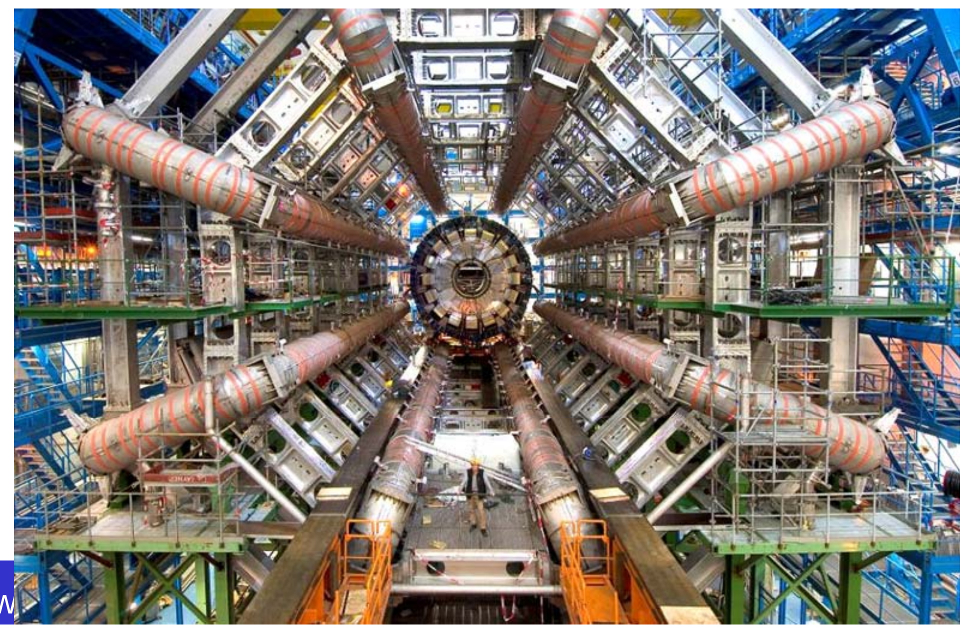

Relevant Links

- Pegasus-WMS: pegasus.isi.edu
- DAGMan: www.cs.wisc.edu/condor/dagman

- Gil, Y., E. Deelman, et al. *Examining the Challenges of Scientific Workflows.* IEEE Computer, 2007.
- Workflows for e-Science, Taylor, I.J.; Deelman, E.; Gannon, D.B.; Shields, M. (Eds.), Dec. 2006

- Montage: montage.ipac.caltech.edu/
- Condor: www.cs.wisc.edu/condor/