Construction of a MODIS Scientific Data Reprojection and Reduction Pipeline in Windows Azure

Jie Li[1], Youngryel Ryu[2], Keith Jackson[3] , Catharine van Ingen[4], Deb Agarwal[3]

University of Virginia[1]
University of California, Berkeley[2]
Lawrence Berkeley Lab[3]
Microsoft Research[4]

Microsoft eScience Workshop 2009
October 17, 2009

eScience: A Data-intensive Computing Era

- Increasing data availability for scientific discoveries
  - Increasing throughput of large scientific equipments
  - Emerging large-scale inexpensive ground-based sensors
- Data-intensive computational models with increasing complexities

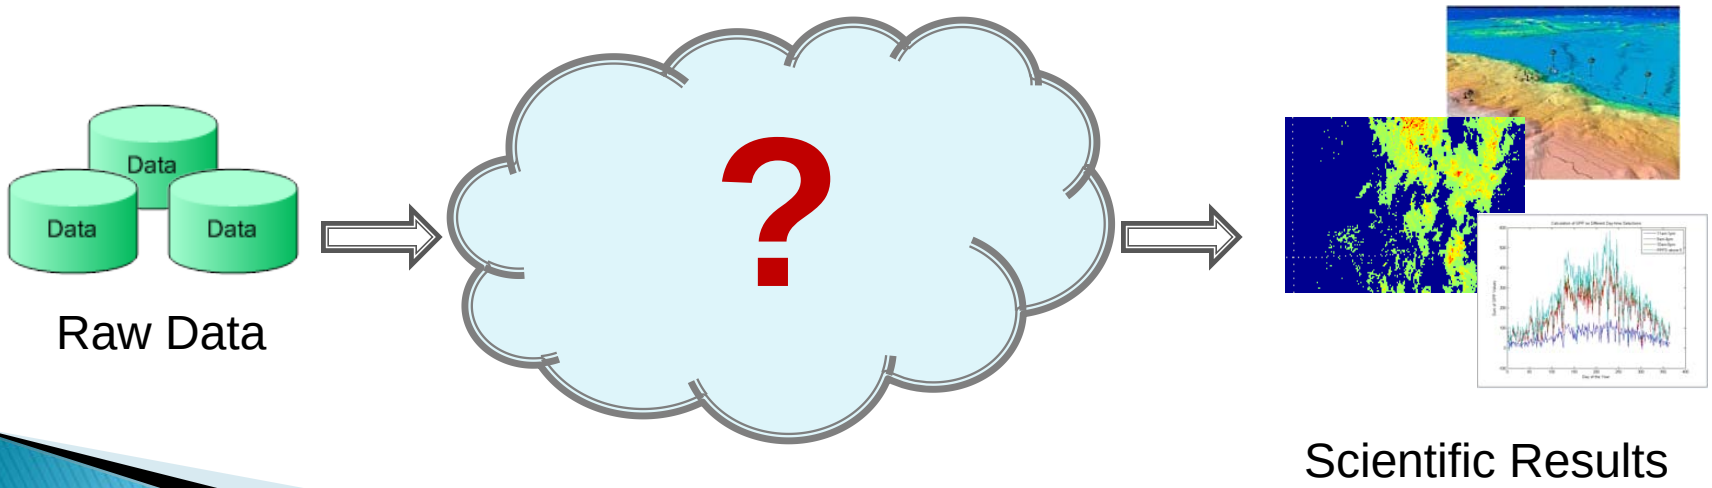

Barriers in Processing the Raw Data

- Computational Resource Entry Barriers
  - Limited local computational resources
  - Access to national-scale supercomputing centers may be limited and not flexible enough
- Data Entry Barriers
  - **Data Collection**: Multiple sources of data to be assembled
  - **Data Transformation**: Heterogeneous data formats to be unified
  - **Data Management**: Software to orchestrate the complex scientific workflows and manage the huge data sets

Do scientists need to be computer experts as well?

MODIS Basics

- Moderate Resolution Imaging Spectroradiometer
  - A key instrument onboard 2 satellites: **Terra** (1999) and **Aqua** (2002)
  - Viewing the entire Earth's surface every 1 to 2 days
  - Acquiring data in 36 spectral bands (groups of wavelengths)
  - Separated into multiple data products (Atmosphere, Land, etc)
  - Playing a vital role in the development of validated, global, interactive Earth system models

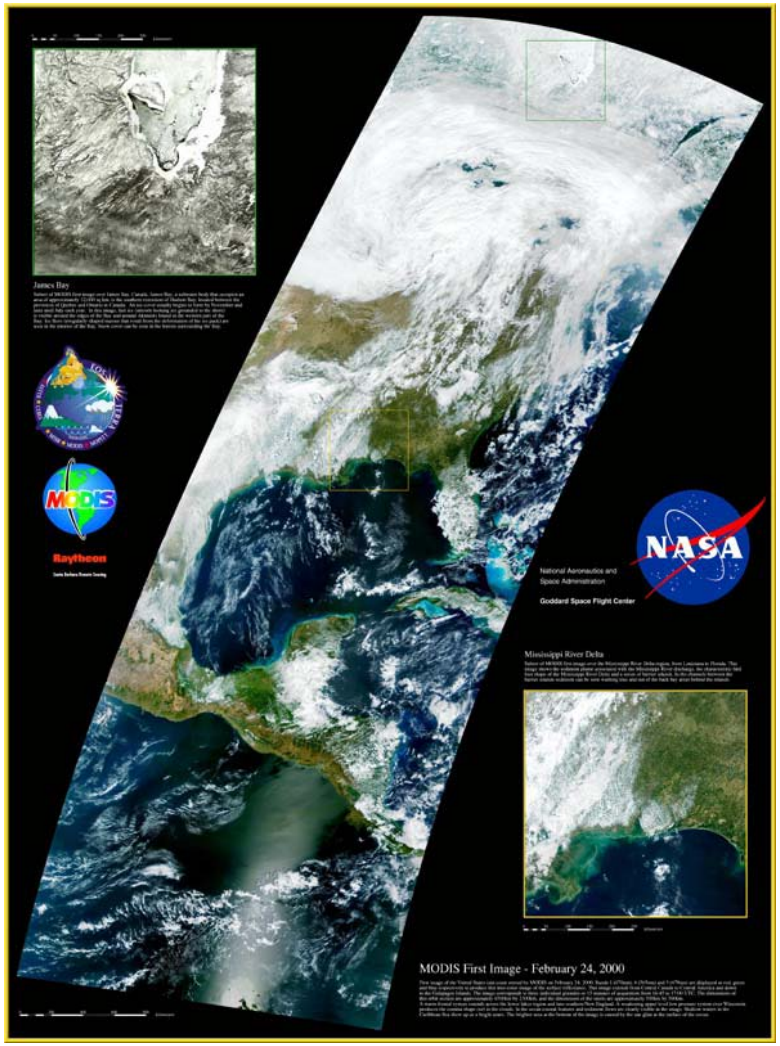

http://aqua.nasa.gov/doc/viz/media/aqua_orbit_sm.mpg

MODIS Application Examples

"I want surface temperature across the globe at a 100KM resolution as input for my climate model."

"I want monthly Leaf Area Index (LAI) for 2003 over the Metolius Ameriflux tower sites to compare with my ground sensor data."

"I want 4 atmosphere products, 5 land products across the US for all years at 1KM resolution and the finest time granularity available to investigate earth evaportranspiration model."

Our first working example

Data Entry Barriers for Using MODIS Data

- Data Collection
  - MODIS source data are published across multiple FTP sites with metadata separately maintained
- Data Transformation
  - Data products have different time granularities and space imaging resolutions
  - Different project types: "*Swath*" and "*Sinusoidal*"
- Data Management
  - For our working example, we want to process totally 10 years of interested MODIS data products for US area
  - Need to download and process ~585K source data files (embarrassingly parallel), requiring thousands of CPU hours and several TBs of data storage

Our Cloud-based Solution

- A MODIS Data Reprojection and Reduction Pipeline in Microsoft Windows Azure cloud computing platform
  - Lower resource entry barriers
  - Hide the complexities in data collection, reprojection and management from domain scientists
  - A generic **Reduction Service** for scientists to upload arbitrary executables to perform scientific analysis on reprojected data.

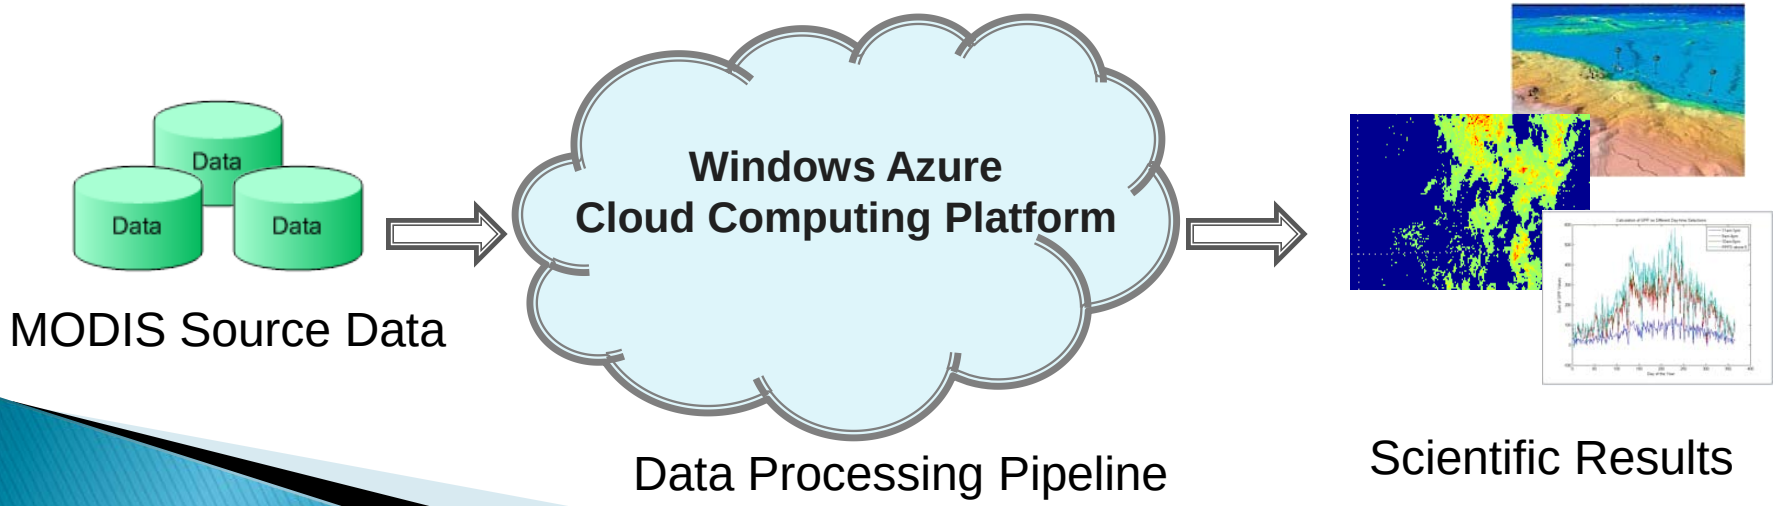

MODIS Data Service Demo

http://modisdataservice.cloudapp.net/

Windows Azure Platform Basics

- Hosted Services
  - **Web Role**: A Web application accessible via an HTTP and/or an HTTPS endpoint
  - **Worker Role**: A background processing application
  - Dynamic scalability of role instances
- Storage Services
  - **Blob service**: Storage for entities such as binary files and data files
  - **Queue Service**: A reliable, persistent queue model for messaging within and between hosted services
  - **Table Service**: Structured storage in the form of tables, with simple query support

MODIS Data Service Pipeline Overview

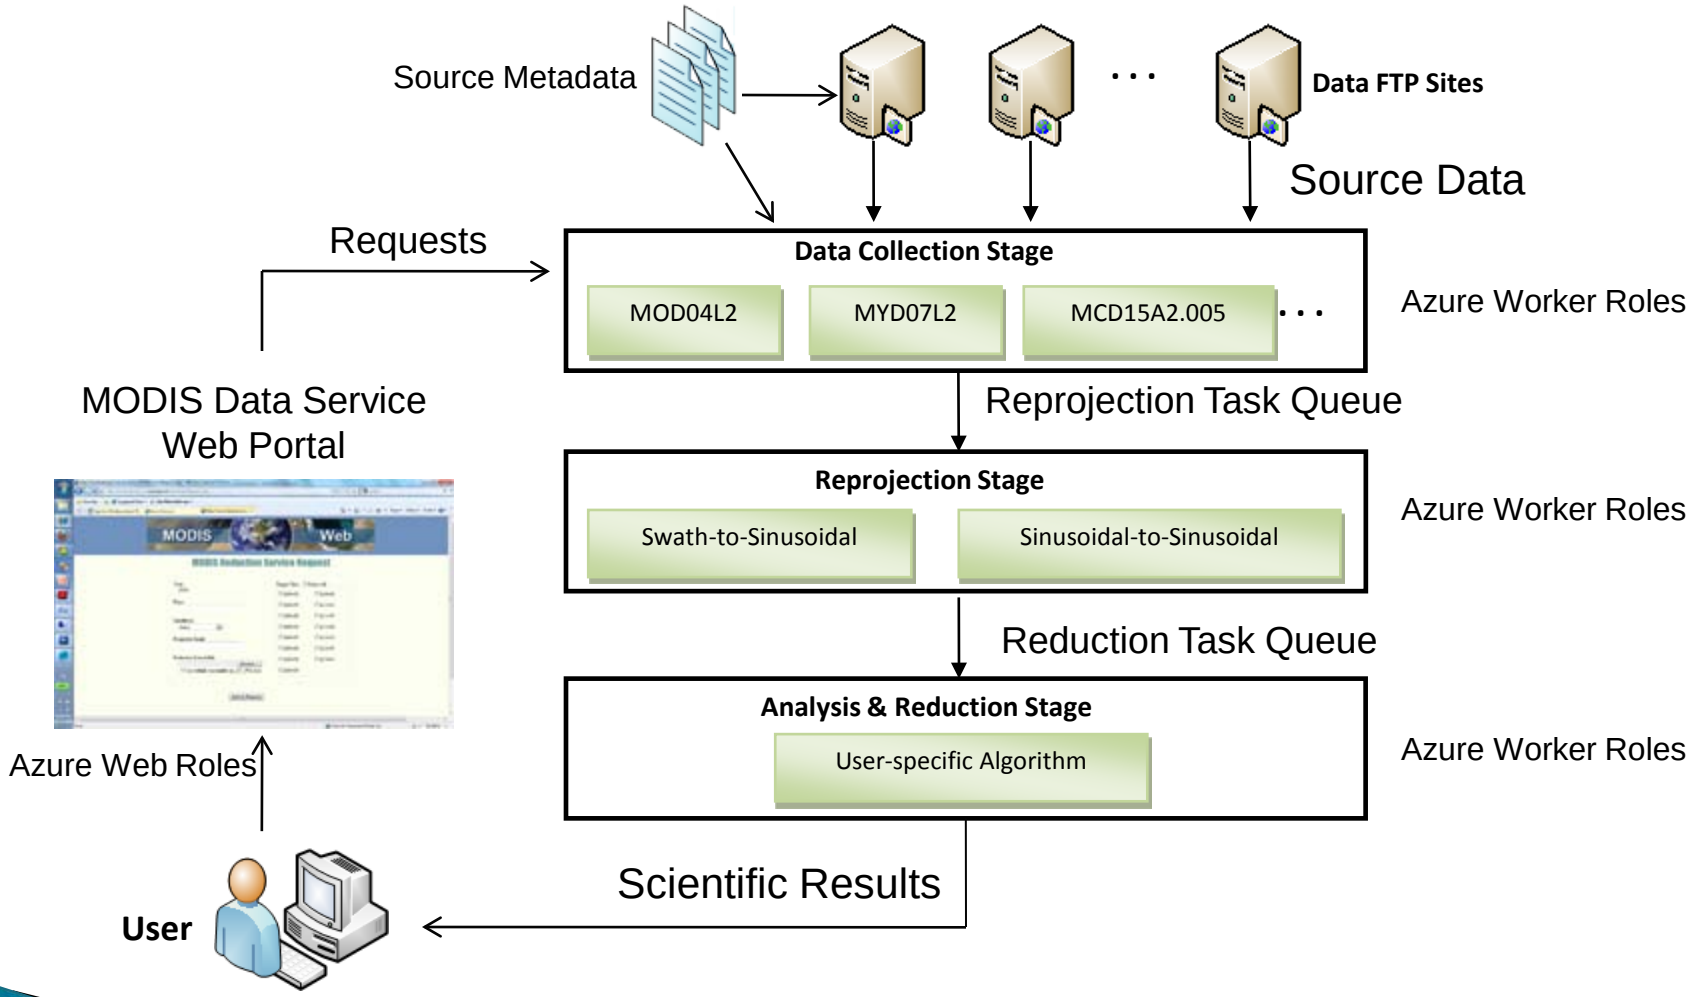

Data Collection

- Huge numbers of source data files, each covering a small fraction of geographic area.
- Metadata stored as Azure tables and support time- and area-based query

Table 1. Data Product Properties for US Continent over 10 Years

	MO(Y)D04L2	**MO(Y)D05L2**	**MO(Y)D06L2**	**MO(Y)D07L2**	**MCD12Q1.005**
Data Group	Atmosphere	Atmosphere	Atmosphere	Atmosphere	Land
Project Type	Swath	Swath	Swath	Swath	Sinusoidal
#Source Data Files	65K	130K	130K	130K	150
Data Size	85GB	500GB	2TB	850GB	13GB
	MCD15A2.005	**MCD43B3.005**	**MO(Y)D11_L2.005**	**MO(Y)D13A2.005**	
Data Group	Land	Land	Land	Land	
Project Type	Sinusoidal	Sinusoidal	Sinusoidal	Swath	**Total**
#Source Data Files	7K	7K	113K	3K	585K
Data Size	13GB	110GB	337GB	40GB	3.9TB

Reprojection Process

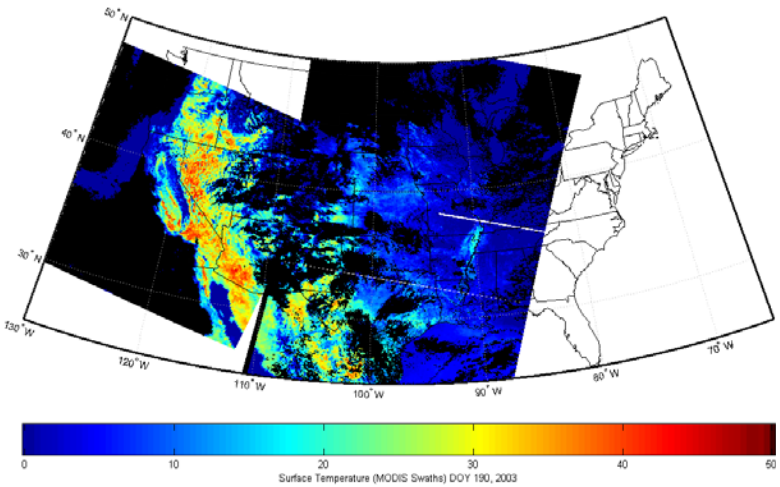

Source Data (Swath format)

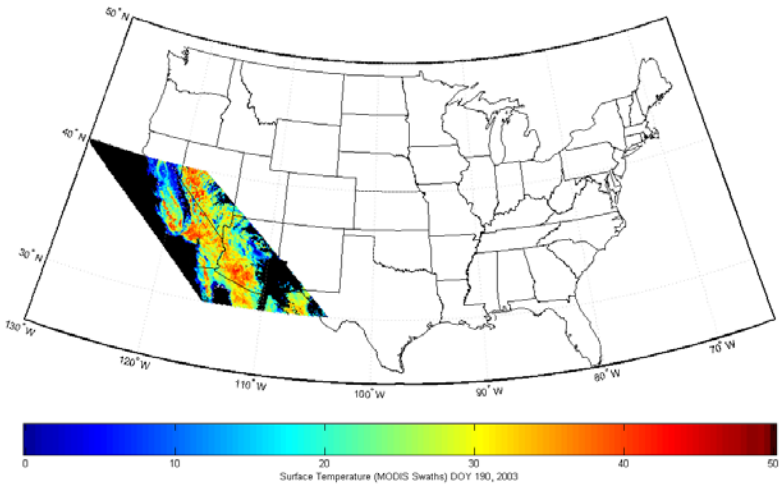

Reprojected Data (Sinusoidal format)

Reprojection Process(Cont.)

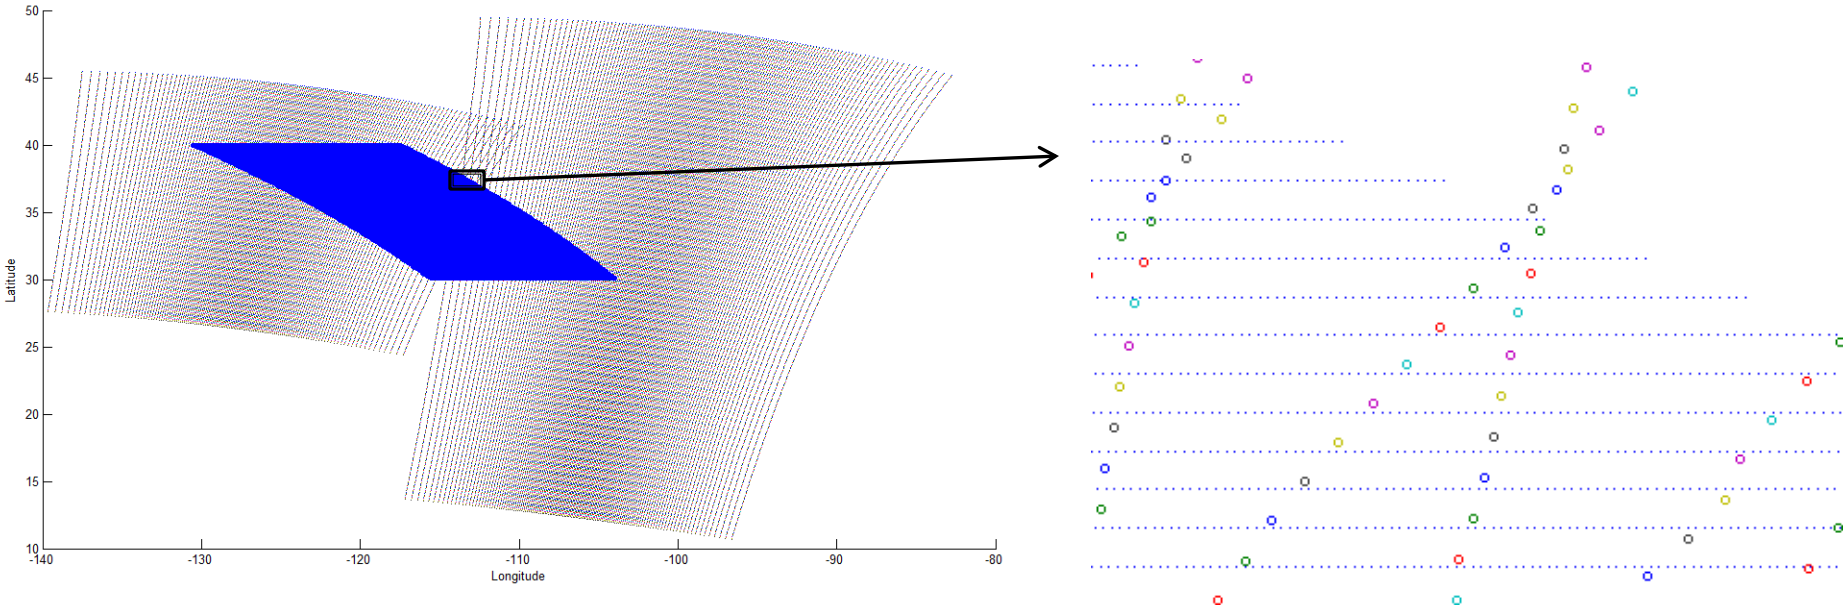

Early Experiences with Reduction Service

- Benefits of reduction algorithm upload support
  - Scientists are free to use up-to-date scientific processing tools
  - Reduce scientific process development time across multiple iterations
- Running MATLAB tool in Azure
  - Can't install MATLAB in Azure because Azure doesn't provide administrator privilege for its hosted service roles (and the potential license issues?)
  - But we can run stand-alone executables generated by the MATLAB Compiler
- Java-based scientific tools are also supported

Task Scheduling in Azure (Reprojection)

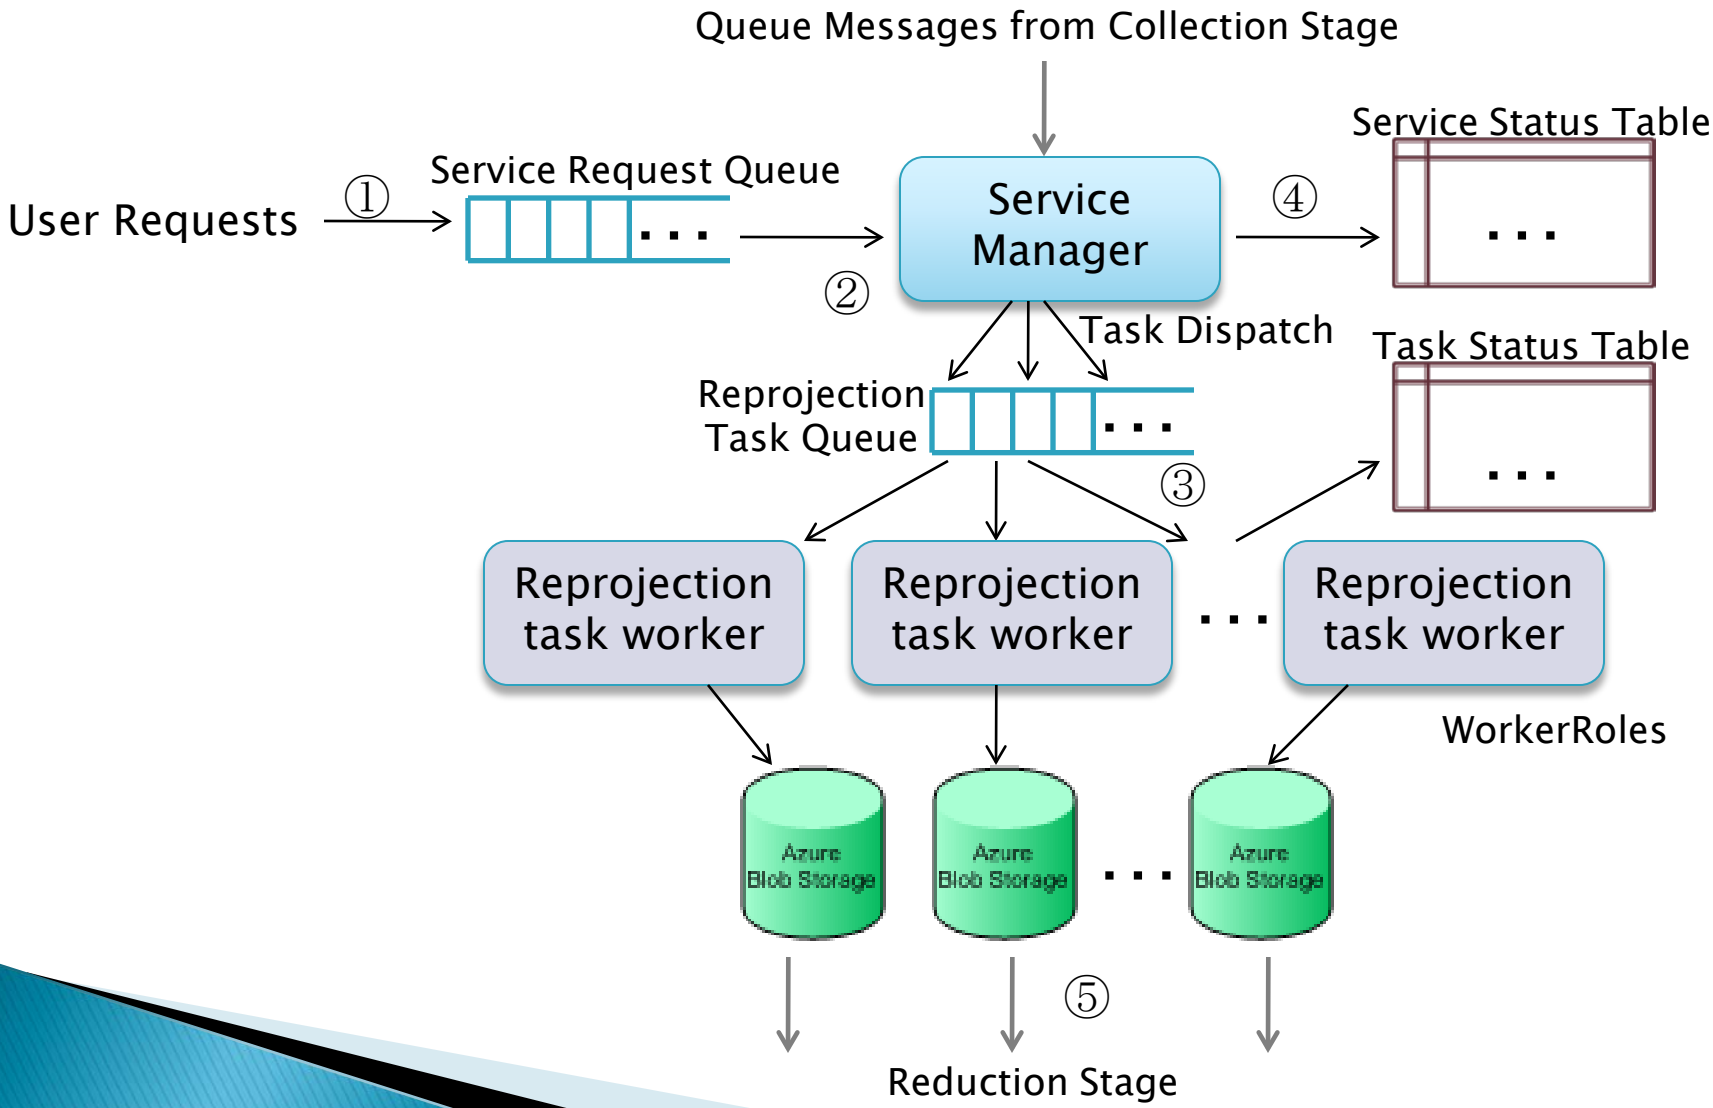

Fault Tolerance

- Tasks can fail for many reasons
  - Broken or missing source data files
  - Reduction tool may crash on execution
  - Unpredictable system errors may happen to worker roles
- The service monitor keeps track of the status of each active task
  - Restart a task on a different worker role if it doesn't finish in a certain amount of time
  - Abort a task if it still fails after 3 retries
  - Service status persistent in Azure tables, thus if service monitor crashes, it can resume status when it gets running again.

Performance

Table 2. Capacity of desktop machine and a single Azure instance

	Desktop	Azure Instance
Capacity	CPU: Intel Core2Duo E6850 @ 3.0GHZ Memory: 4GB Hard Disk: 1TB SATA Network: 1Gbps Ethernet OS: Windows 7 RC Build7100 (32-bit)	CPU: 1.5-1.7GHZ X64 equivalent processor Memory: 2GB Local Storage: 250GB Network: 100Mbps OS: Windows 2008 Server x64 (64-bit)

Table 3. Processing time for 1500 reprojection tasks (Unit: hours)

	MOD04_L2	MOD06_L2	MYD11_L2.005
150 instances	0.30	0.85	0.44
100 instances	0.40	1.20	0.61
50 instances	0.76	2.25	1.12
Desktop	16.29	72.62	33.45

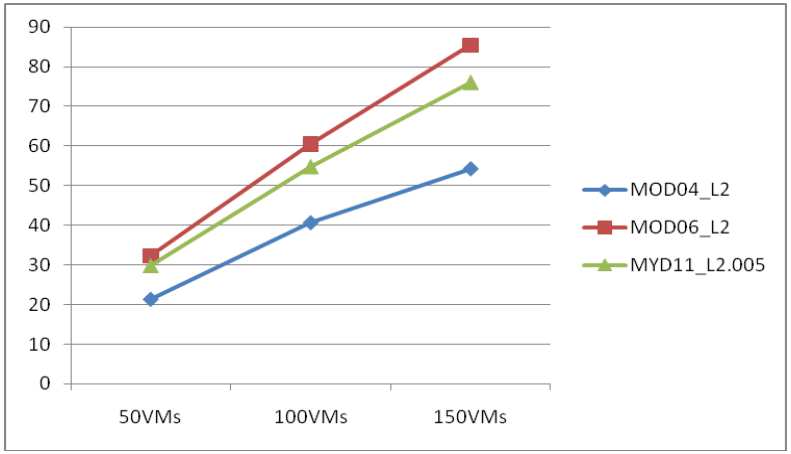

Fig. 1 Performance speedups over a single desktop

Conclusions

- Lower the resource and data entry barriers for scientists to conduct research on MODIS data
- Leveraging Window Azure cloud computing platform, construct a data processing pipeline with high flexibility and dynamic scalability
- Our architecture can be easily modified to incorporate other eScience applications with similar computation patterns

Future Work

- Improve the web portal interface to incorporate more general data processing specifications
- Validation of uploaded reduction algorithms and tools
- Apply MapReduce-enabled frameworks such as Dryad and DryadLINQ in our architecture

MODIS Data Service Check Out

Thank you!

&

Questions?