



# Big Data Management with the Myria Cloud Service

Magdalena Balazinska

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING

UNIVERSITY OF WASHINGTON

http://www.cs.washington.edu/people/faculty/magda


UNIVERSITY of WASHINGTON
eScience Institute
ADVANCING DATA-INTENSIVE DISCOVERY IN ALL FIELDS

# Science is Facing a Data Deluge!

- Astronomy: High-resolution, high-frequency sky surveys (SDSS, LSST)
- Medicine: ubiquitous digital records, MRI, ultrasound
- Biology: lab automation, high-throughput sequencing
- Oceanography: high-resolution models, cheap sensors, satellites
- Etc.

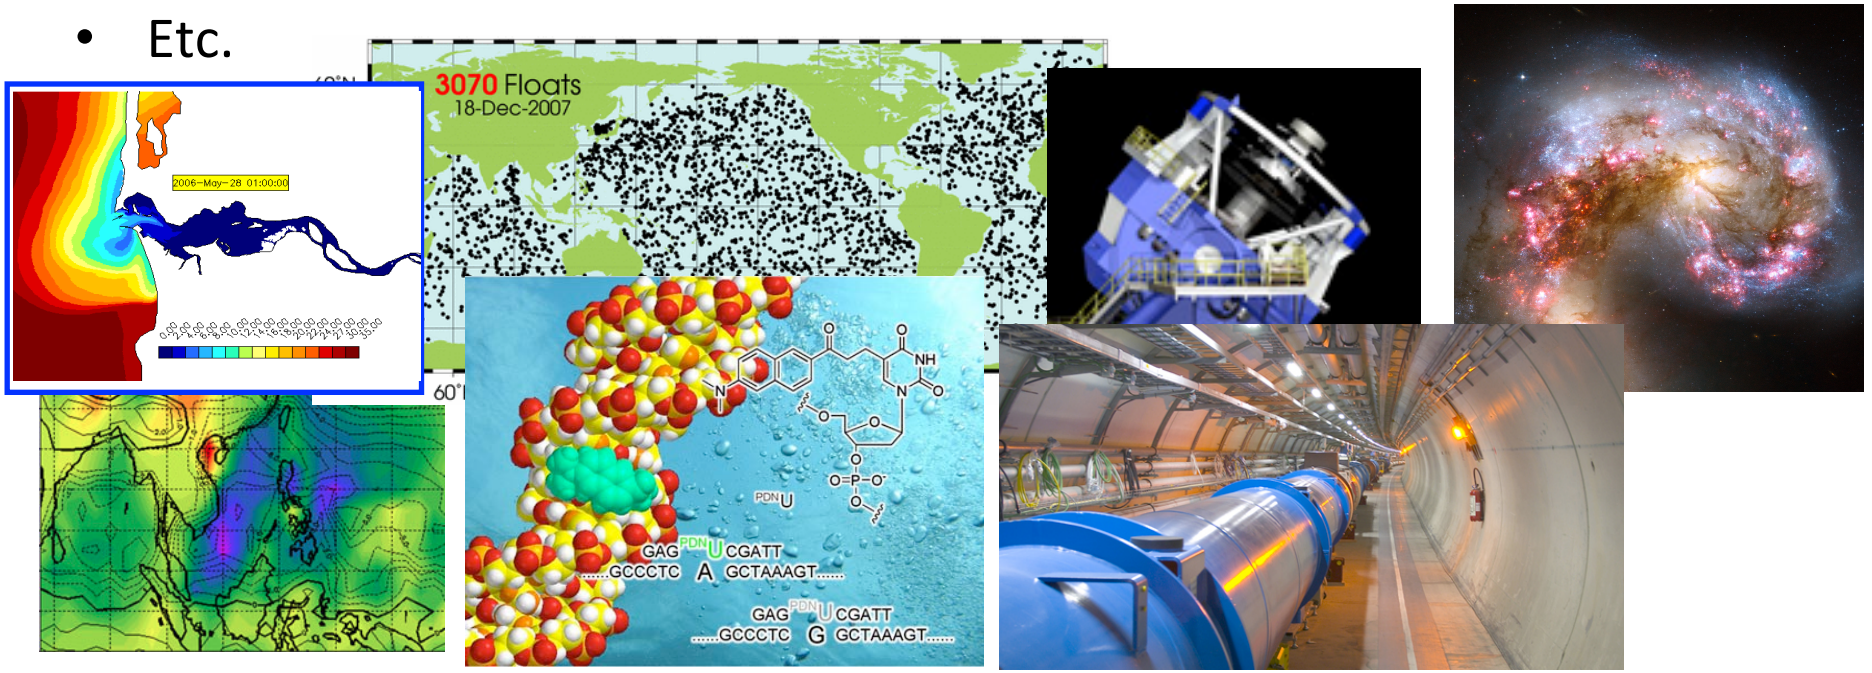

# Example: Astronomy

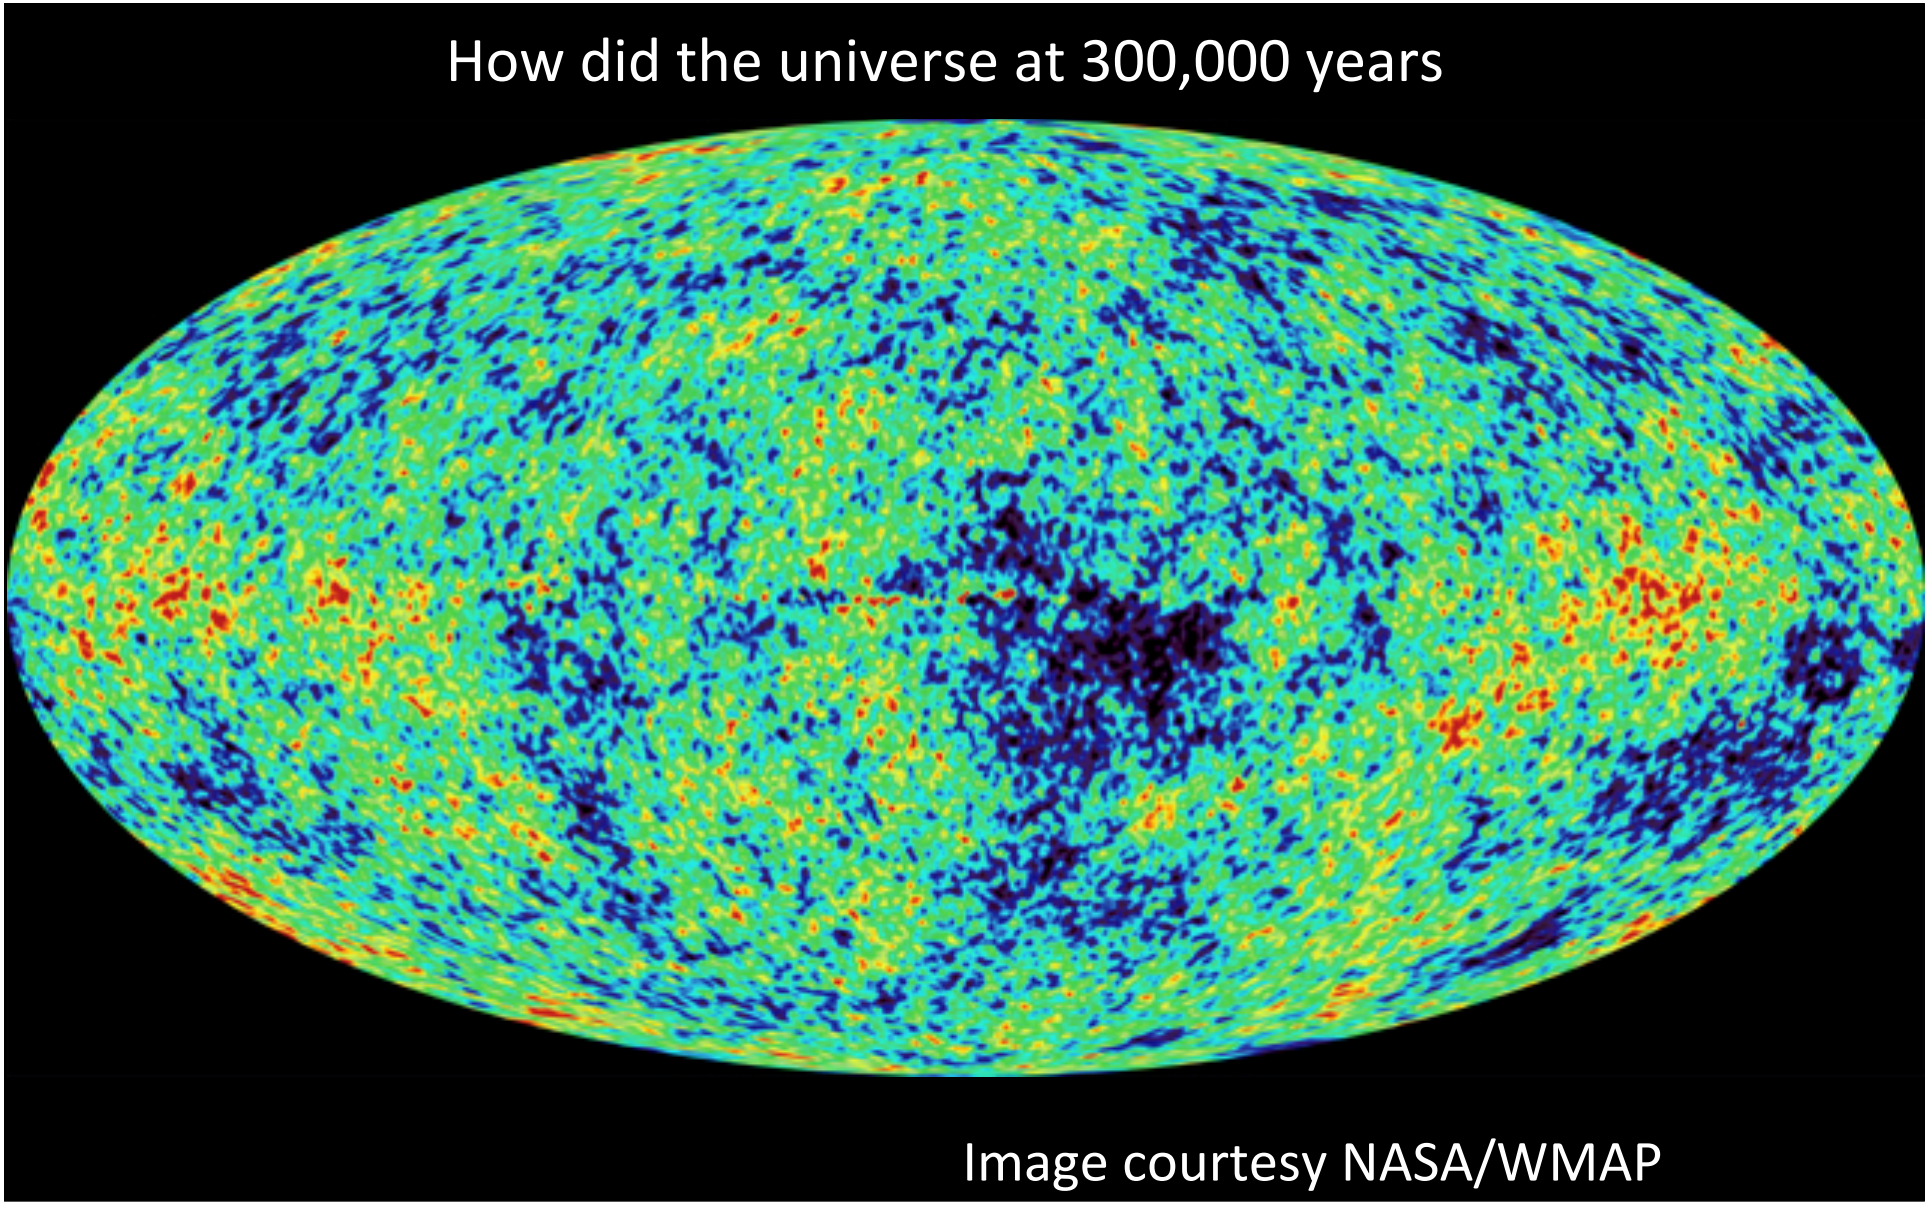


Slide from Andrew Connolly (Astronomy Dept.)


Slide from Andrew Connolly (Astronomy Dept.)

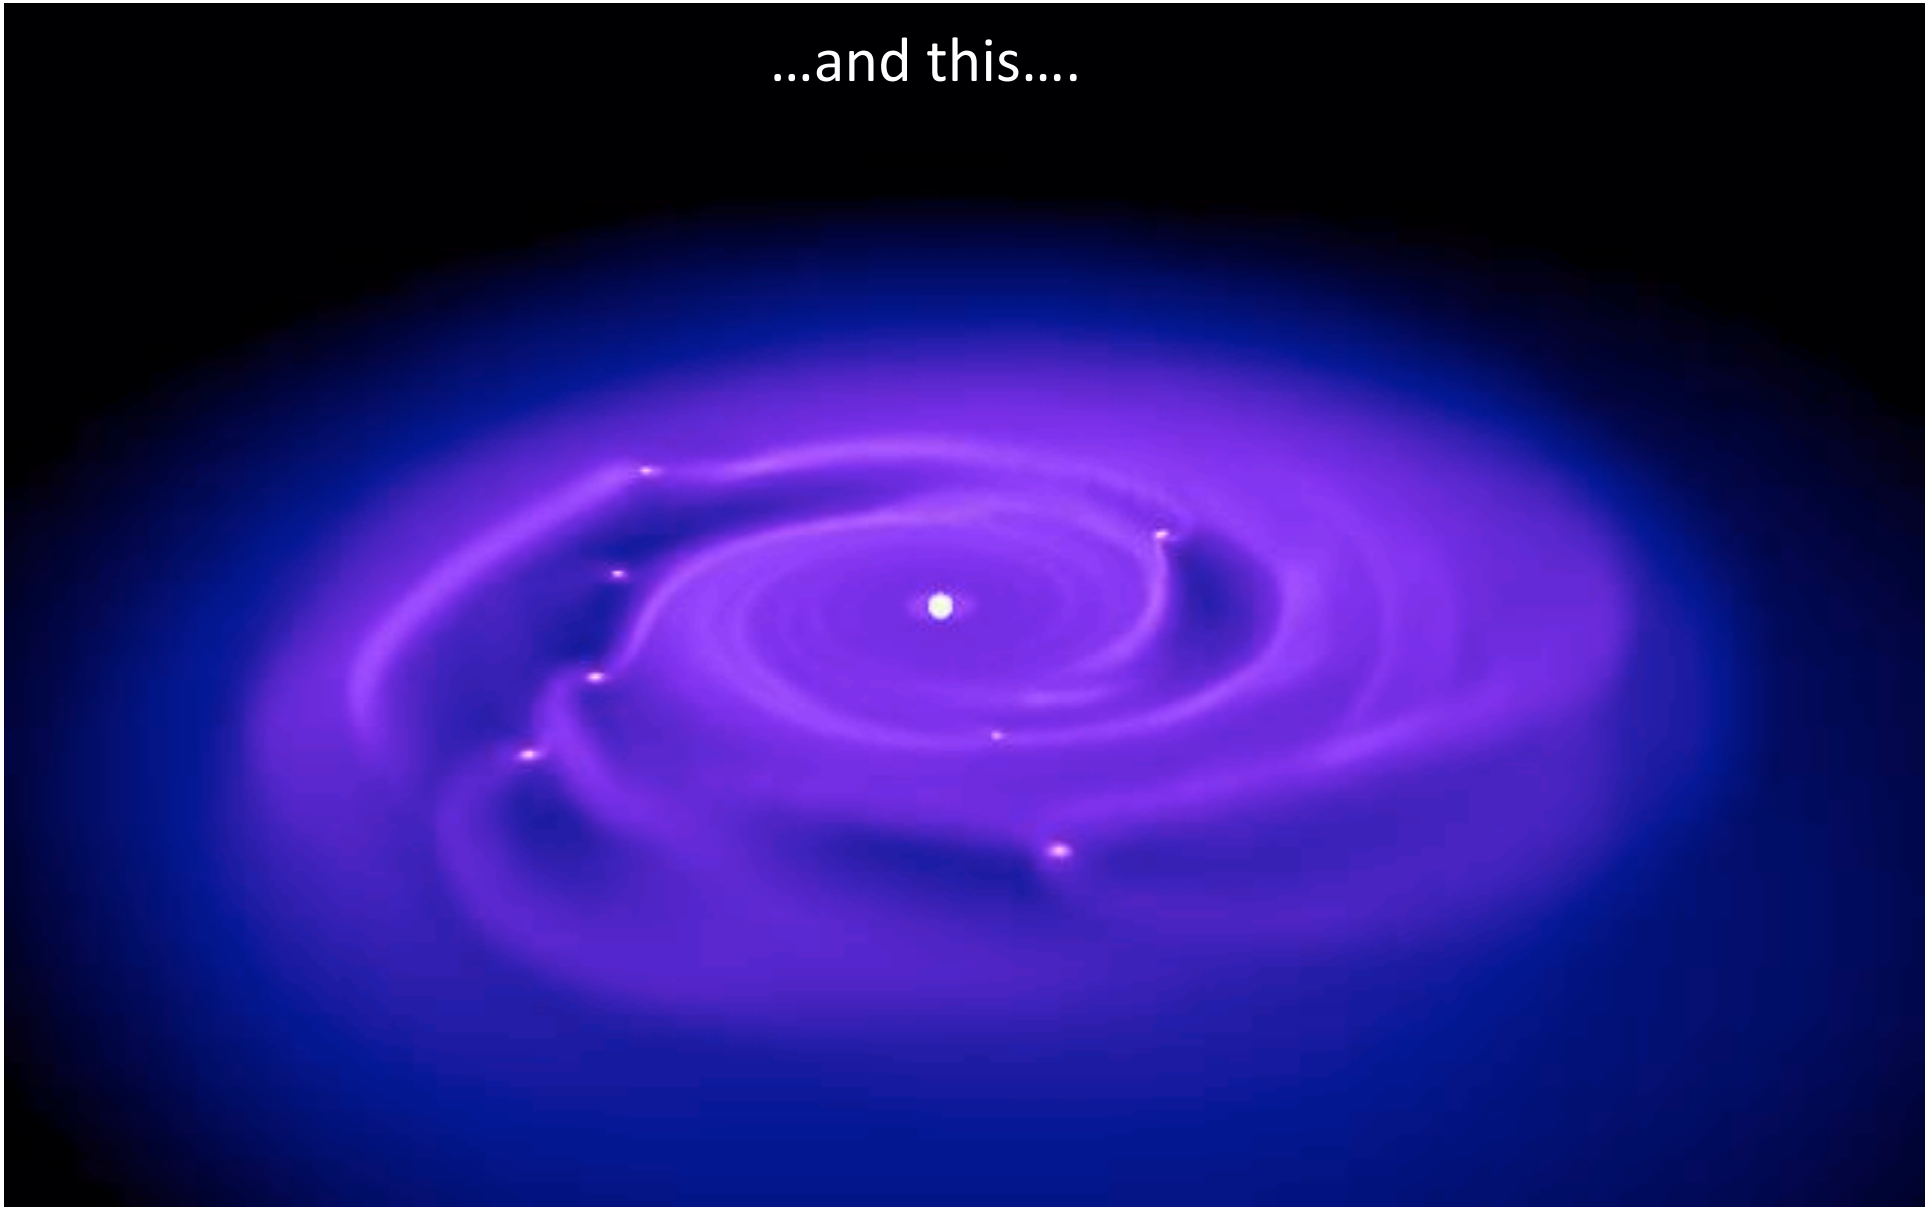


Slide from Andrew Connolly (Astronomy Dept.)


Slide from Andrew Connolly (Astronomy Dept.)

# Method 1: Observe the Sky

## Large Synoptic Survey Telescope

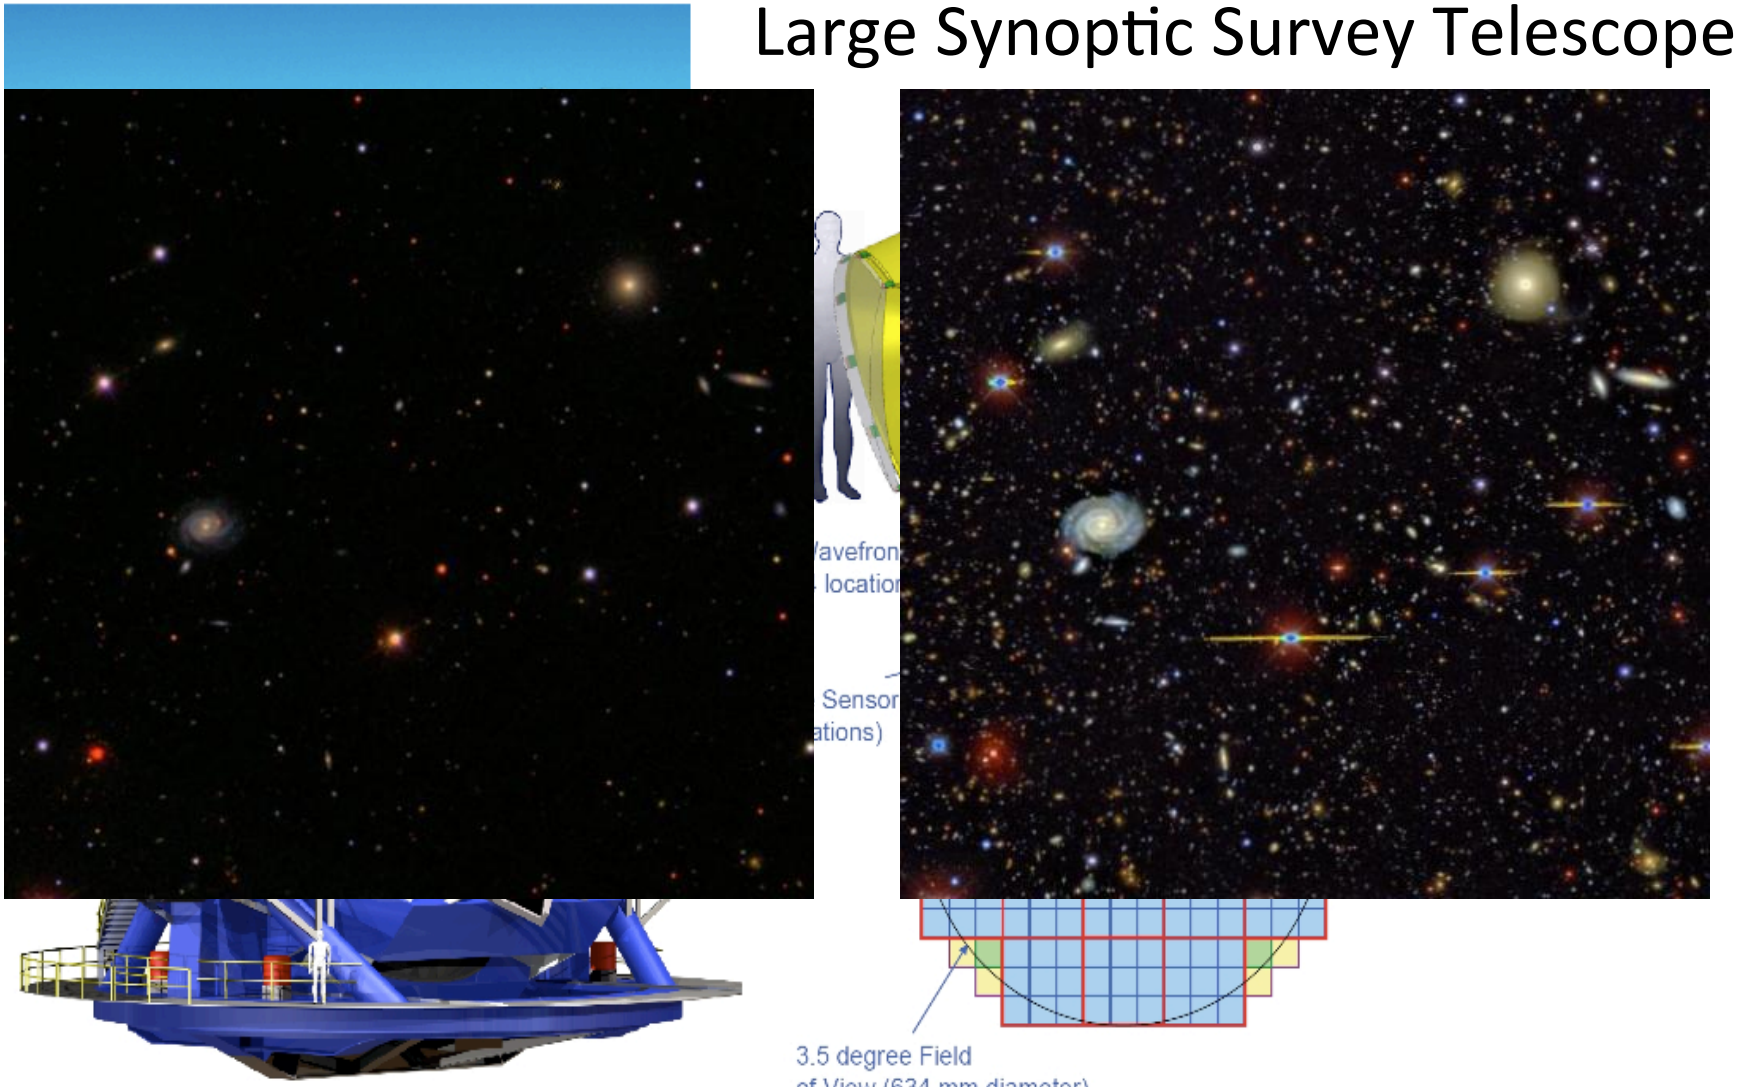


Slide from Andrew Connolly (Astronomy Dept.)

# Challenges from new astronomy

- Sloan Digital Sky Survey (SDSS)
  - 7 years of imaging
  - 8000 sq degrees of the sky (1/5$^{th}$)
  - 200 million stars and galaxies
  - **80 TB raw images**
- LSST data flow
  - 20,000 sq degrees every 3 nights
  - **40 TB of imaging per night**
  - $10^8$ sources a night ($10^3$ "events")
  - 1000 repeat observations over 10 years
  - 10 Petabytes of catalogs (10 years)
  - 100 PBs of images
  - 5 months to watch 1 year of data (HDTV)
  - Data public as soon as taken

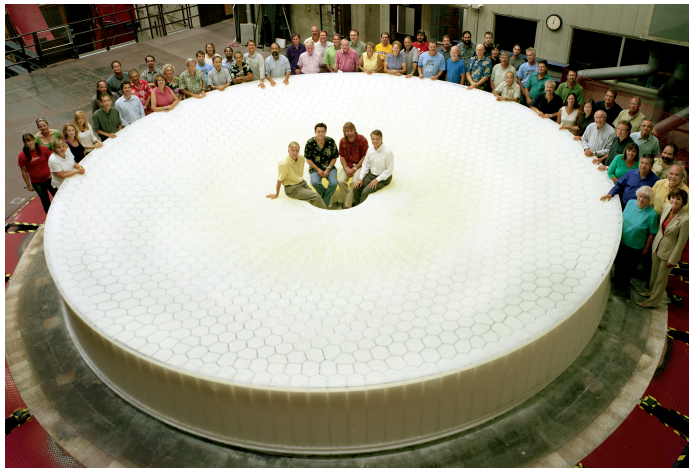

Slide from Andrew Connolly (Astronomy Dept.)

# Method 2: Simulate

- Evolution of large-scale structure in the universe
  - Universe is a set of particles (gas, dark matter, stars)
  - Output snapshot every few simulation timesteps

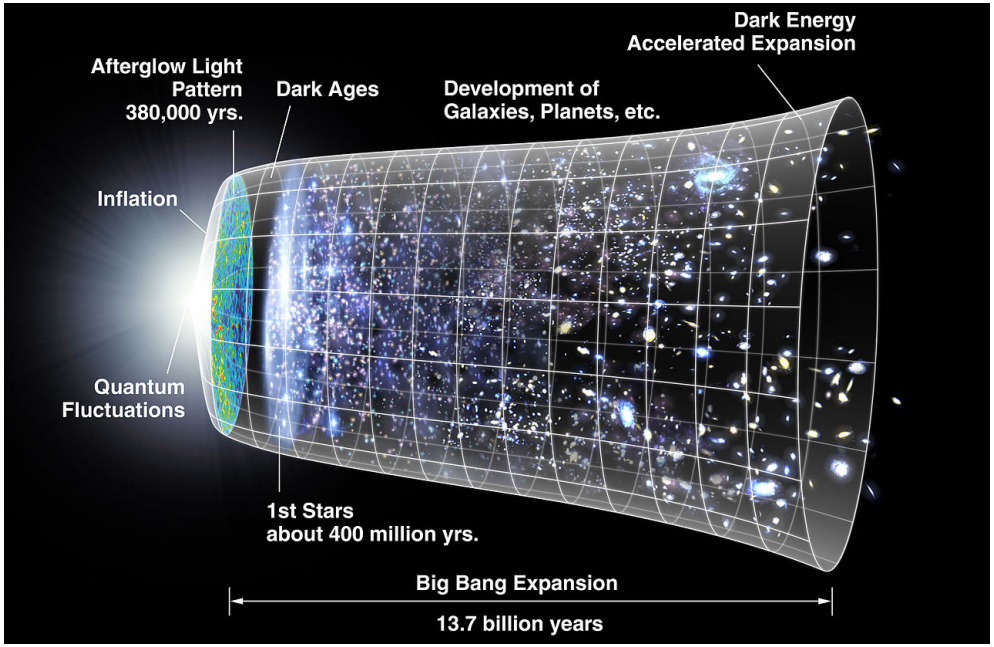


Typical simulations today

Spring 2014: 5TB

Fall 2014: 200TB

UW N-body group

# N-Body Simulations Data

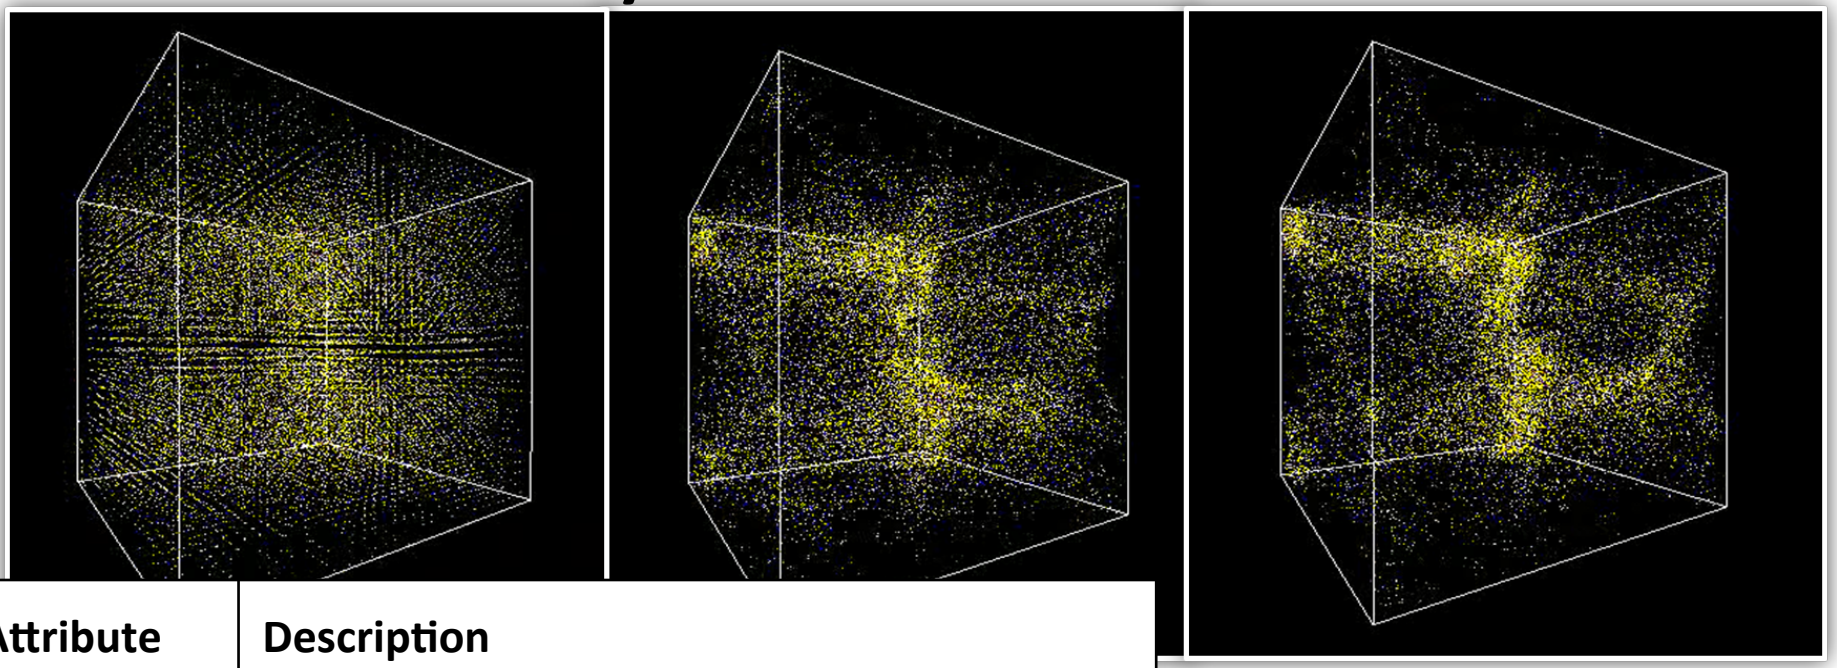

| Attribute | Description |
|---|---|
| iOrder | unique identifier |
| X, Y, Z | position in Cartesian coordinates |
| Type | type of particle: either dark, star or gas |
| Grp | halo group particle belongs in |
| Time | Timestep the particle belongs in |

Universe simulation
46 snapshots
5 TB

# Big Data Need in All Sciences!

- Astronomy: High-resolution, high-frequency sky surveys (SDSS, LSST)
- Medicine: ubiquitous digital records, MRI, ultrasound
- Biology: lab automation, high-throughput sequencing
- Oceanography: high-resolution models, cheap sensors, satellites
- Etc.

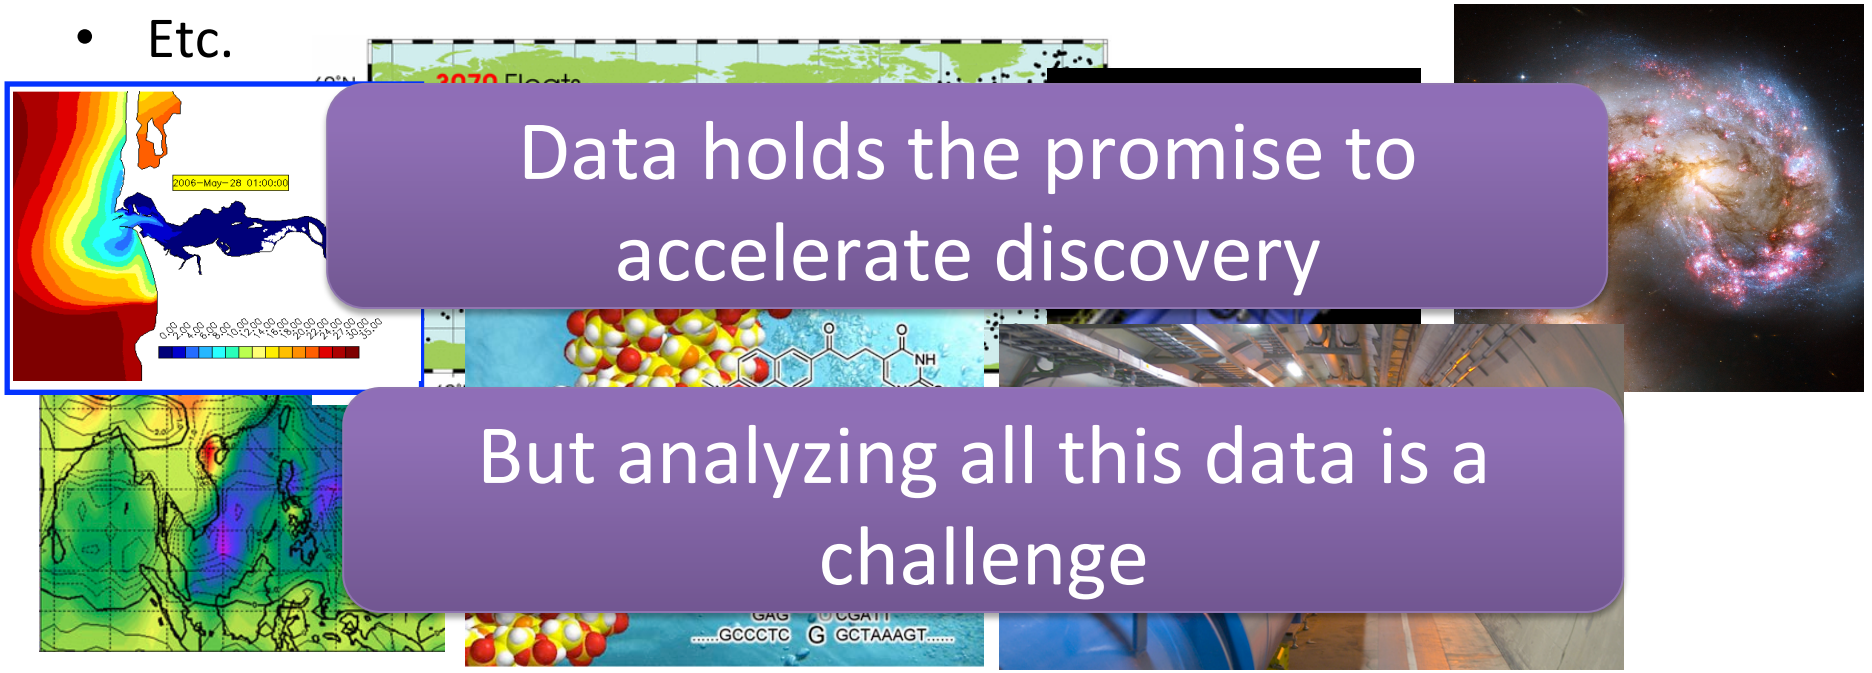


Data holds the promise to accelerate discovery

But analyzing all this data is a challenge

# Industry is Facing a Data Deluge!

- Clickstreams, search logs, network logs, social networking data, RFID data, etc.
- Examples: Facebook, Twitter, Google, Microsoft, Amazon, Walmart, etc.

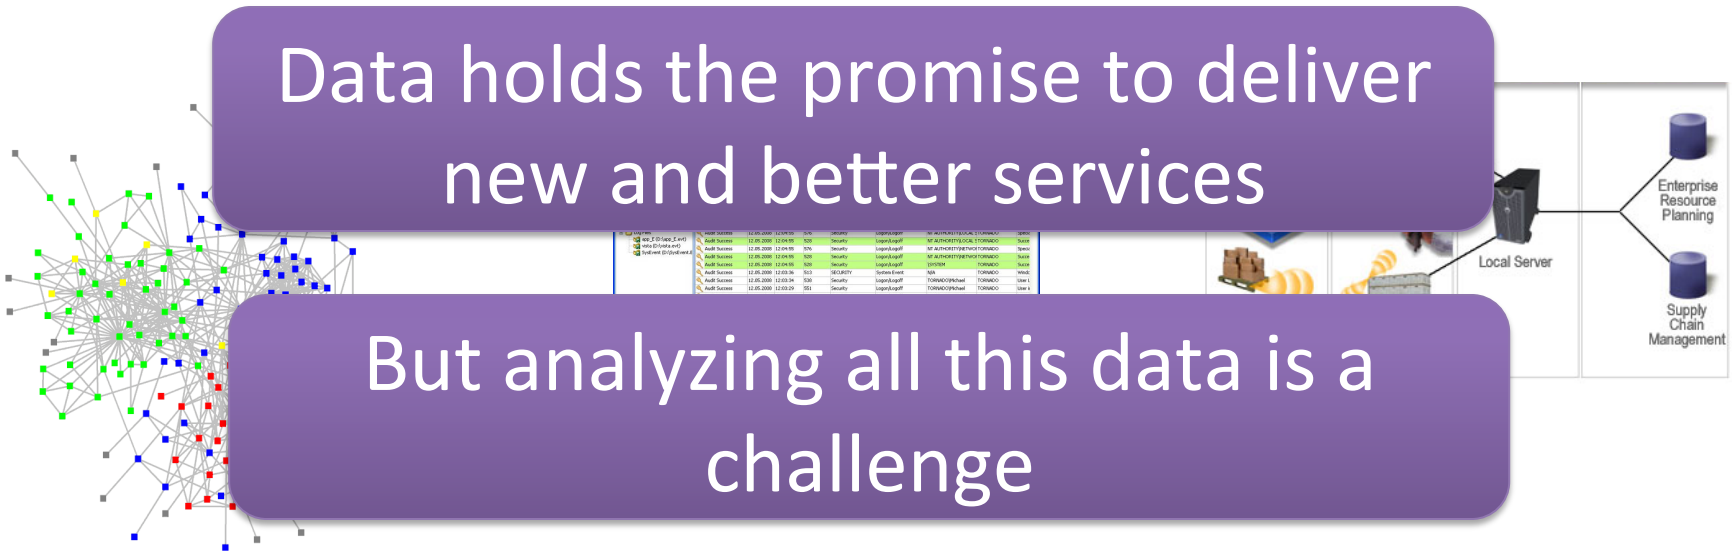

# Challenging Application Requirements

Exciting and challenging requirements

Often generalize beyond campus

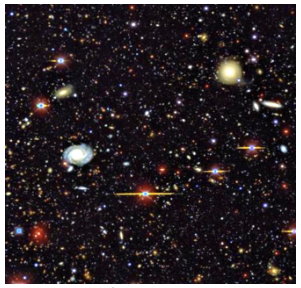

Telescope image:

1. Iterative data cleaning
2. Objects extraction
3. Classification

Picture from Deep Lens Survey (DLS: Tyson)

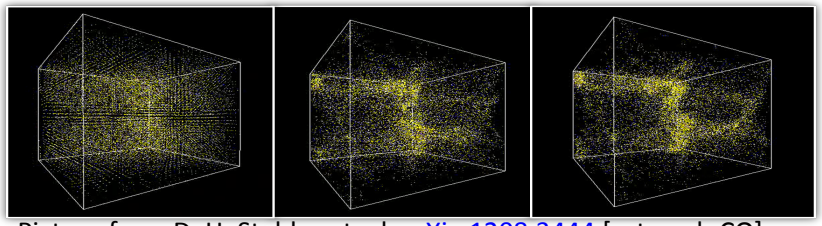

Picture from D. H. Stalder et. al. arXiv:1208.3444 [astro-ph.CO]

N-body simulation data:

1. Manage hundreds of TB of data
2. Data clustering to extract galaxies
3. Graph analytics to study galaxy evolution

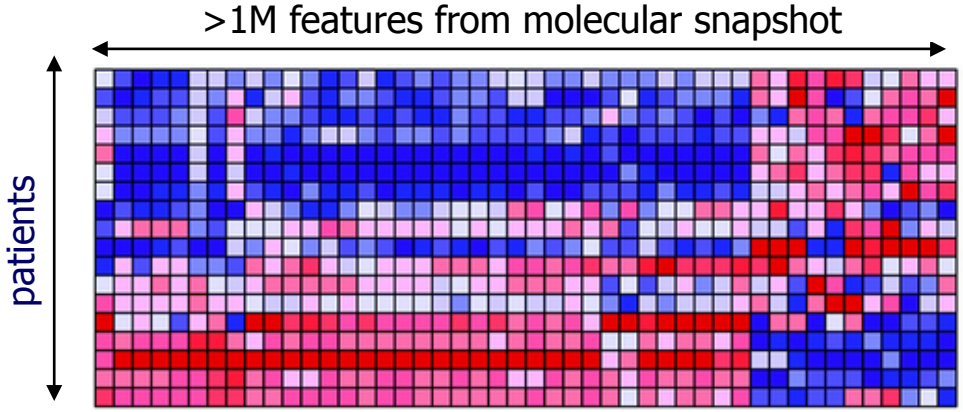


Picture from Su-In Lee

Genome data processing:

1. Linear algebra on large matrices
2. Novel machine learning algorithms

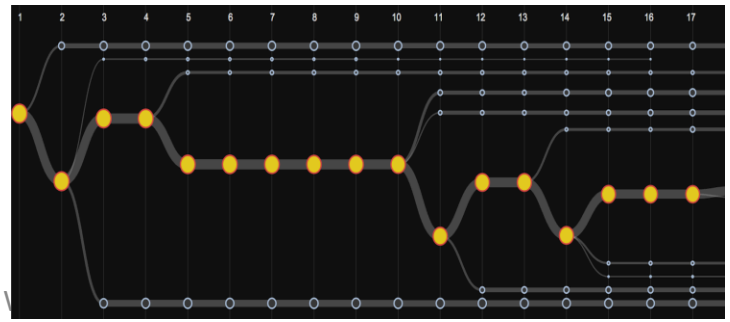

# How Can We Analyze All this Data

- Excel? Limit of 65,536 rows by 256 columns
- Write my own Java/Python/other program?
  - Not all data fits in memory
  - May want to run in parallel in a cluster
  - Do I really want to write a program for each question?
- Use an existing database management system
  - Today's data **V**olume, **V**elocity, and **V**ariety are breaking these systems… new systems are appearing
- Build a new database management system!

# Myria Big Data Management Service

Myria is a Cloud service: Just open browser and go!

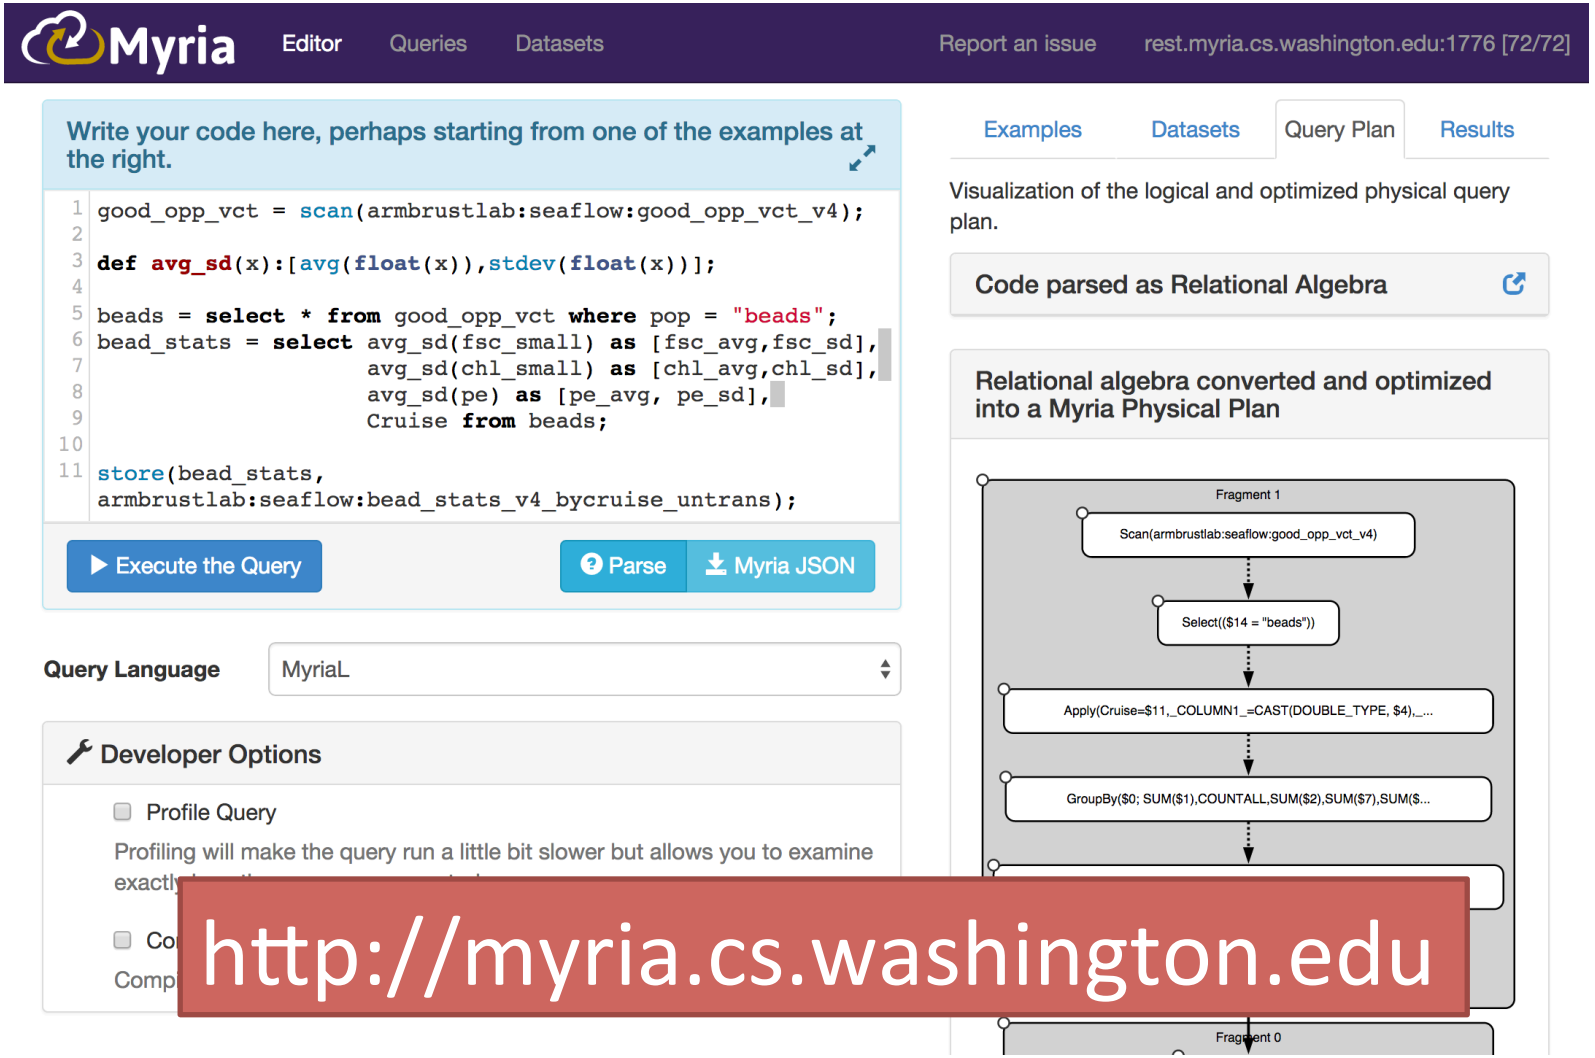


http://myria.cs.washington.edu

# Analysis through Declarative Queries

## Grab my data from Amazon S3

```
smallTable = load("https://s3-us-
west-2.amazonaws.com/uwdb/sampleData/smallTable",
csv(schema(column0:int, column1:int), skip=0));
…
store(smallTable, smallTable, [$0]);
```

## Perform some data manipulation

```
t = scan(public:adhoc:smallTable);
smallTableAggregated = select count(*) from t;
store(smallTableAggregated, smallTableAggregated);
```

# Example: MyMergerTree

Cloud Service for analyzing galactic merger trees

Built on top of Myria

Used to analyze a 5TB dataset

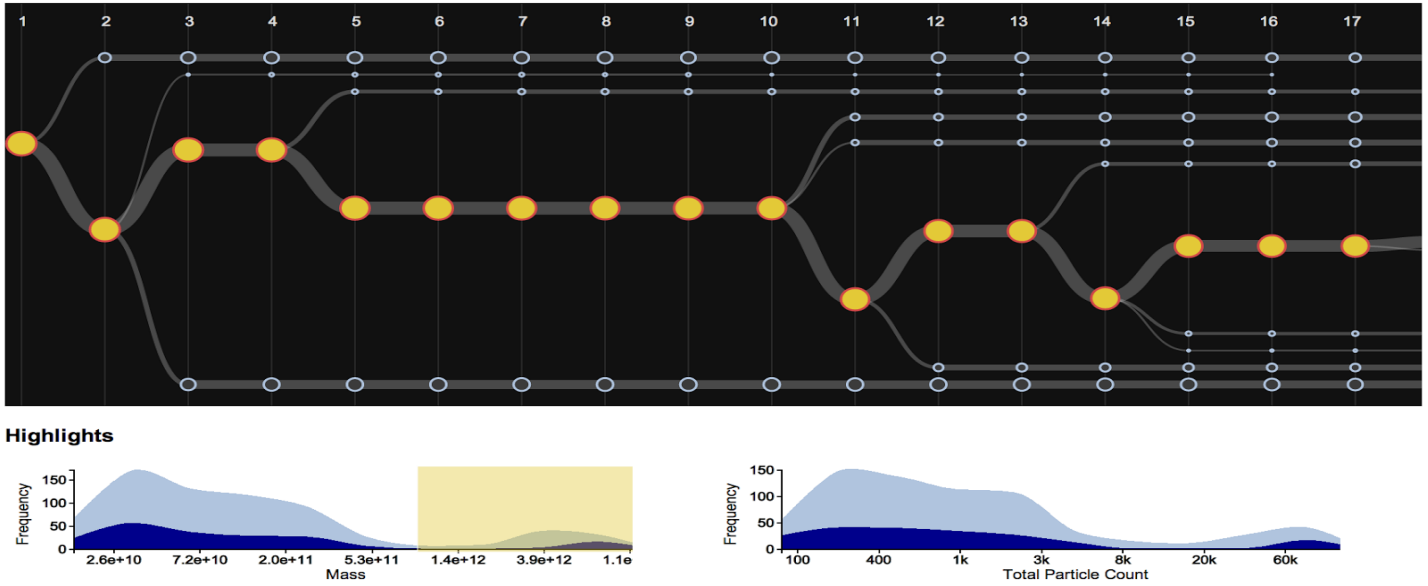

# Example Myria Applications

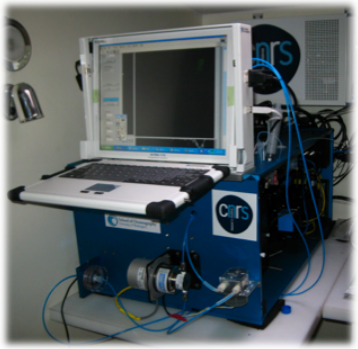

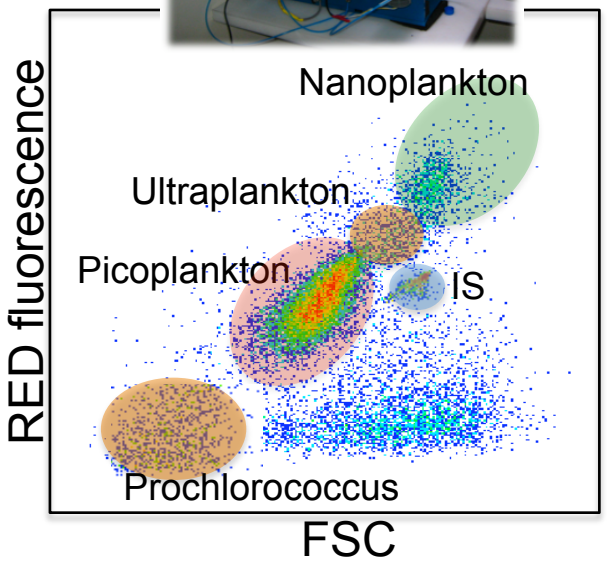


Environmental Flow Cytometry

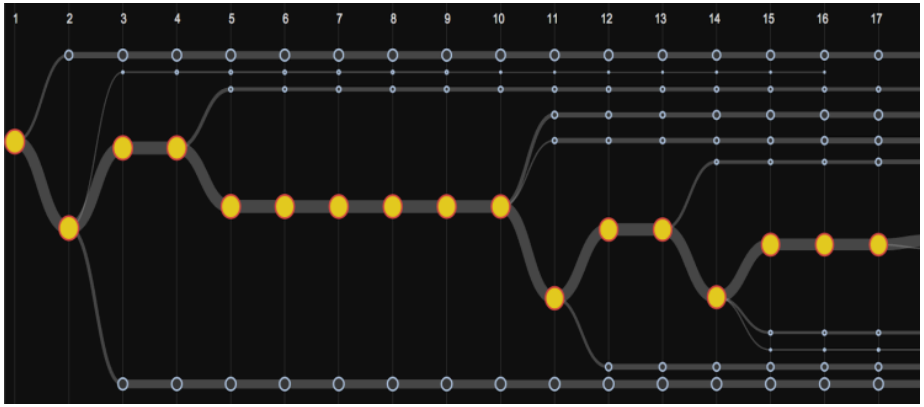


Galaxy Simulations

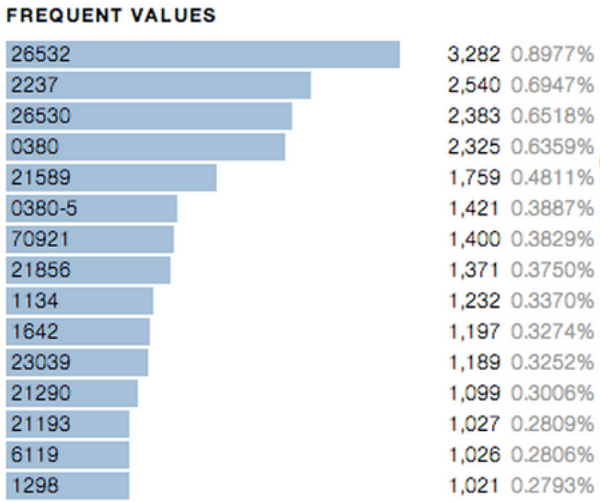


Retail Analytics

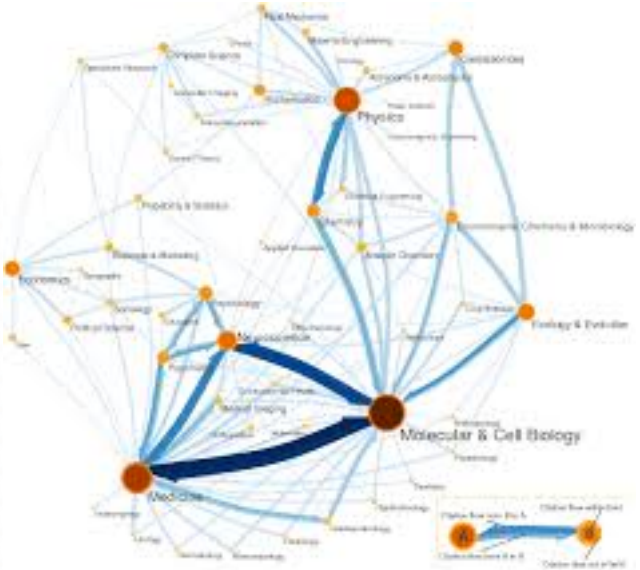


Bibliometrics

# Myria Is a Cloud Service

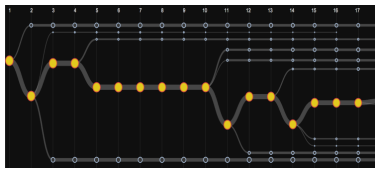

MyMergerTree

Browser | Python | Specialized Services

RACO

Query Translation, Optimization, and Orchestration

MyriaX

Parallel, Iterative, and Elastic Query Execution

MPI | SciDB | Graphs | NoSQL

# Myria is Expressive

## Myria is expressive

- SQL, Datalog, MyriaL

## Quick example (connected components)

```
edgesRaw = scan(astro:cosmo50:cosmo50CompleteEdges);
edgesAll = [from edgesRaw as E where E.NowGroup = 52
            emit E.CurrentTime, E.CurrentHalo, E.NextHalo, E.SharedParticlesCount];
edgesNew = [from edgesAll as E where E.CurrentTime == 1
             emit E.CurrentTime, E.CurrentHalo, E.NextHalo, E.SharedParticlesCount];
edgesResult = edgesNew;
I = [1 as i];
do
   delta = [from edgesNew as E1, edgesAll as E2, I
   where E1.NextHalo == E2.CurrentHalo and E1.CurrentTime == I.i and E2.CurrentTime == I.i+1
   emit E2.CurrentTime as CurrentTime, E2.CurrentHalo as CurrentHalo, E2.NextHalo as NextHalo,
   E2.SharedParticlesCount as SharedParticlesCount];
   edgesResult = unionall(delta, edgesResult);
   edgesNew = delta;
   I = [from I emit i+1 as i];
while [from I emit min(i) <= 10];
store(edgesResult, finalAnswer);
```

# Myria is a Parallel Data Management System

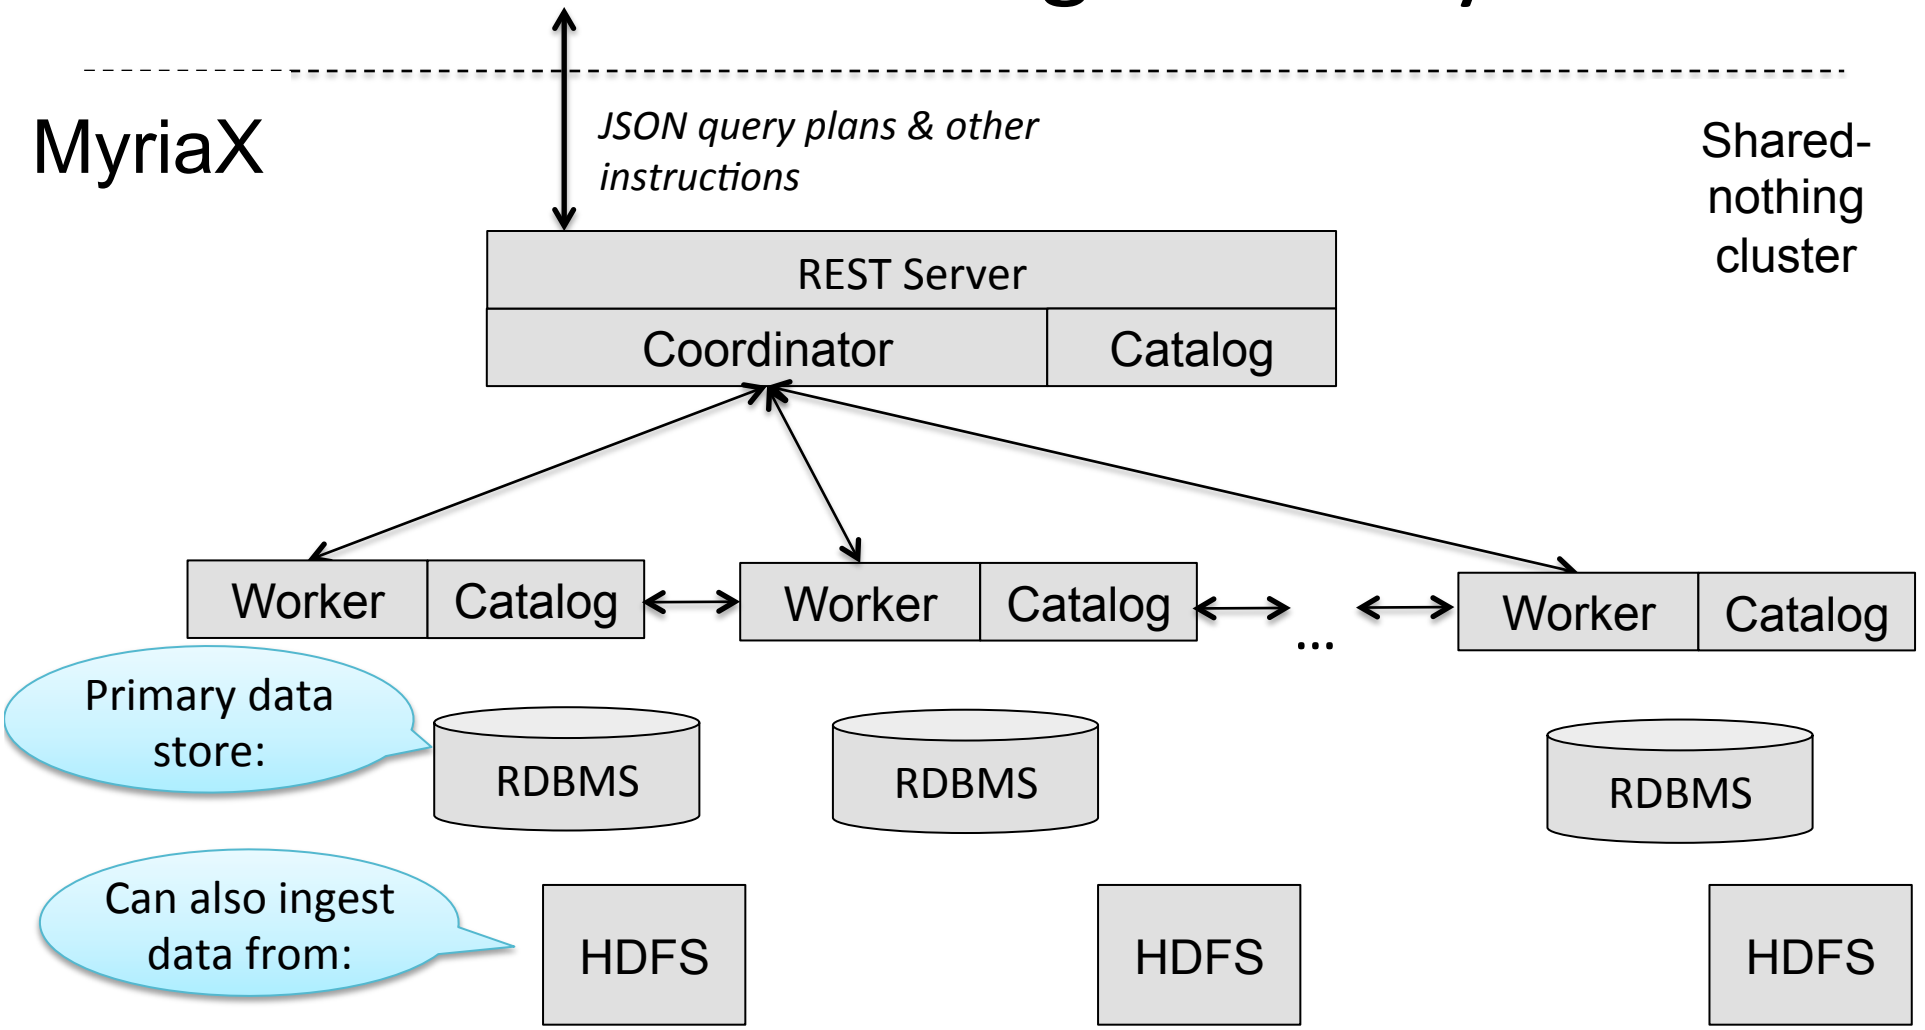

# MyriaX Query Execution Engine

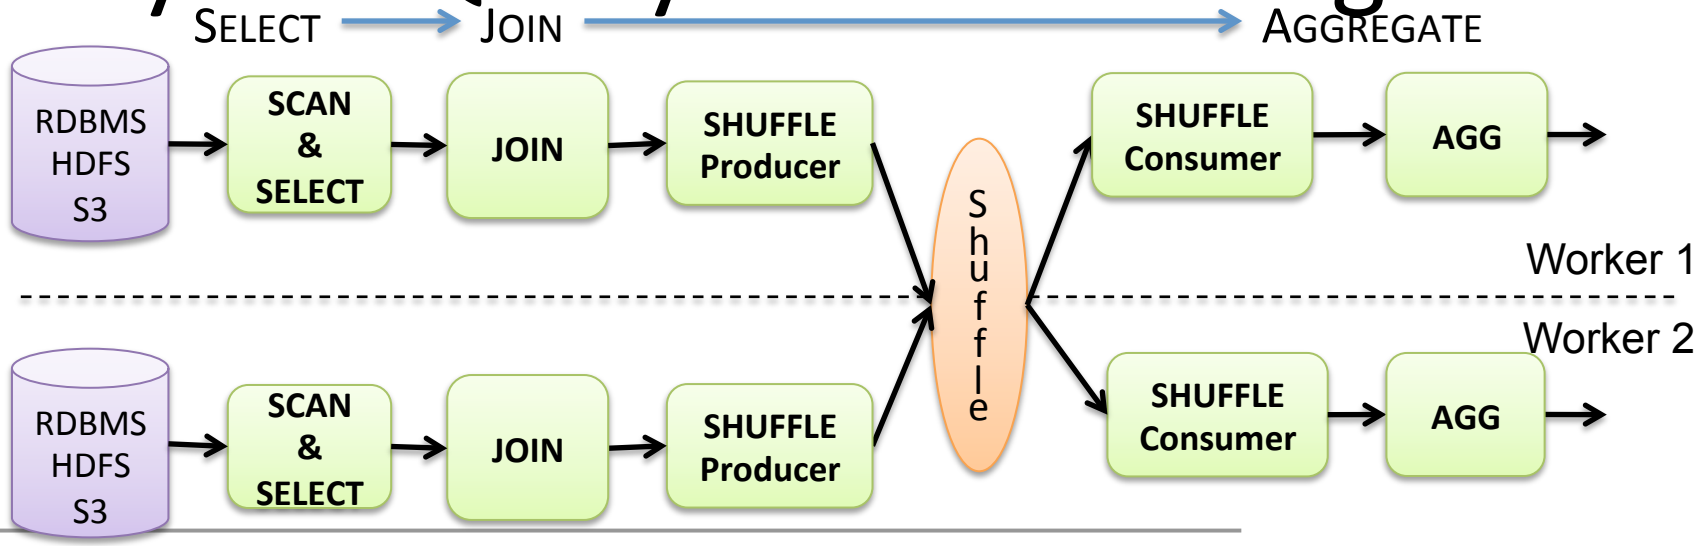


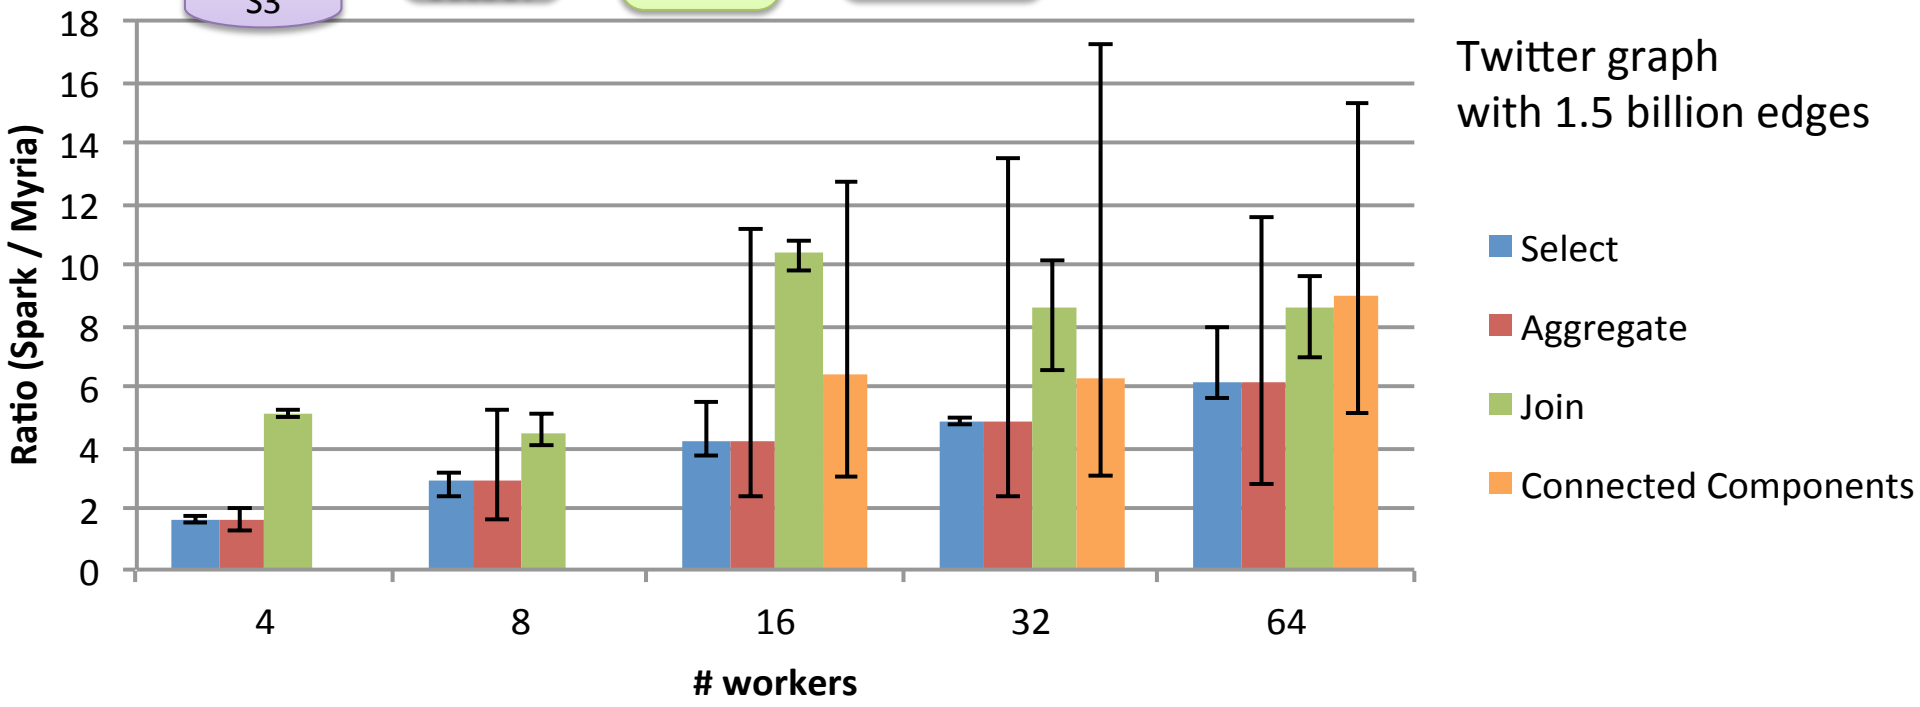

# Example Specific Research Problems

# Recursive Relational Query with Aggregations

- Positive Datalog
- Multiple relations with recursive dependencies

```
Edges(time,gid1,gid2,  :- Galaxies(time,gid1),
  $COUNT(*) as c)         Particles(time,gid1,pid),
                          Particles(time+1,gid2,pid)
Galaxies(time+1,gid2)  :- Galaxies(time,gid1),
                          Edges(time,gid1,gid2,c), c >= 2
```

- Aggregate functions:
  - Only in head
  - Bag-monotonic

```
Edges(time,gid1,gid2,$COUNT(*) as c)
Edges(time,gid1,gid2,c), c >= 2
```

- A subset of Datalog-with-aggregation
  - Details in the paper

# Important Runtime Optimizations

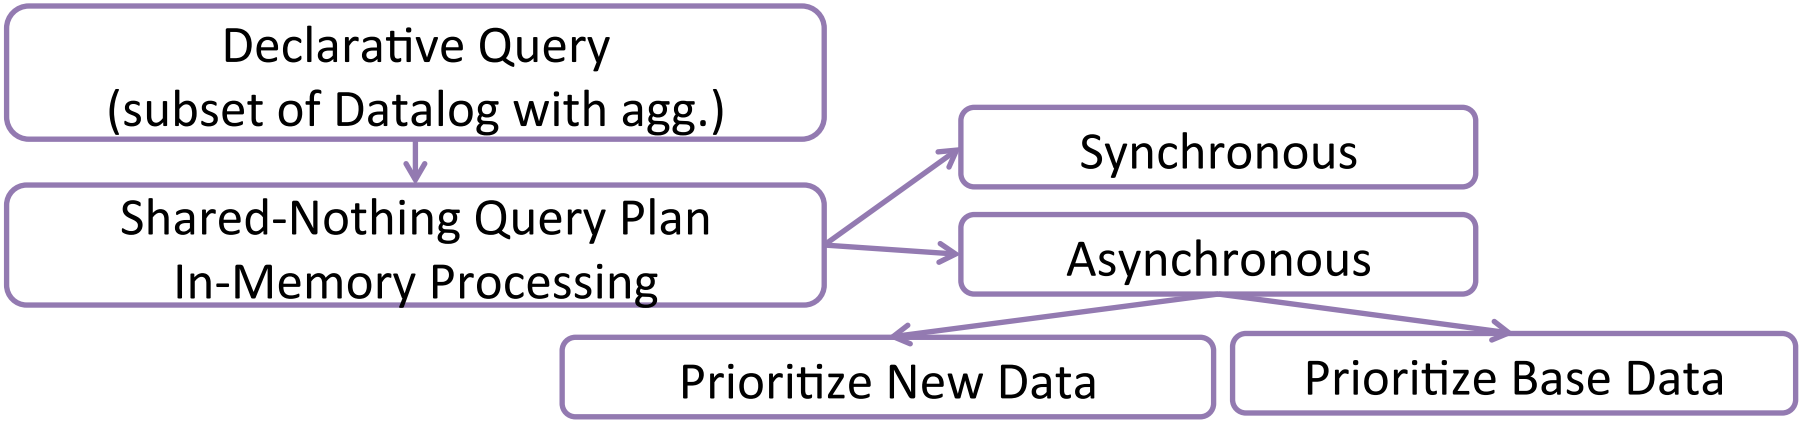


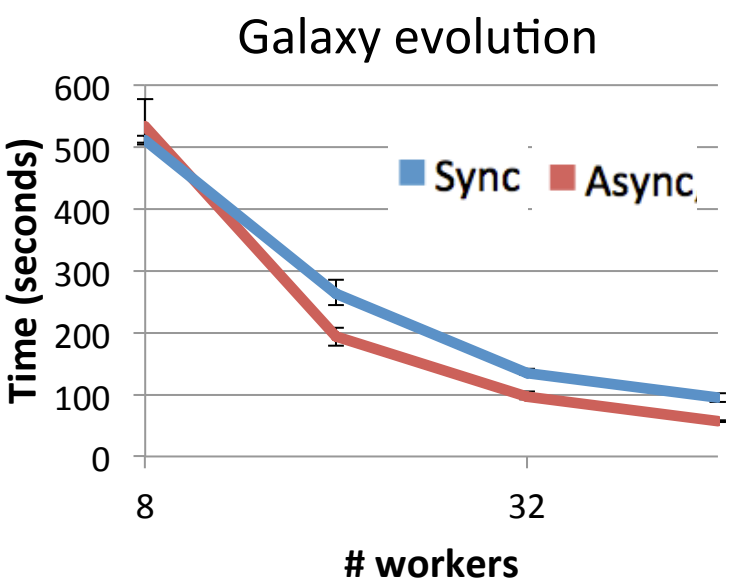


# Important Runtime Optimizations

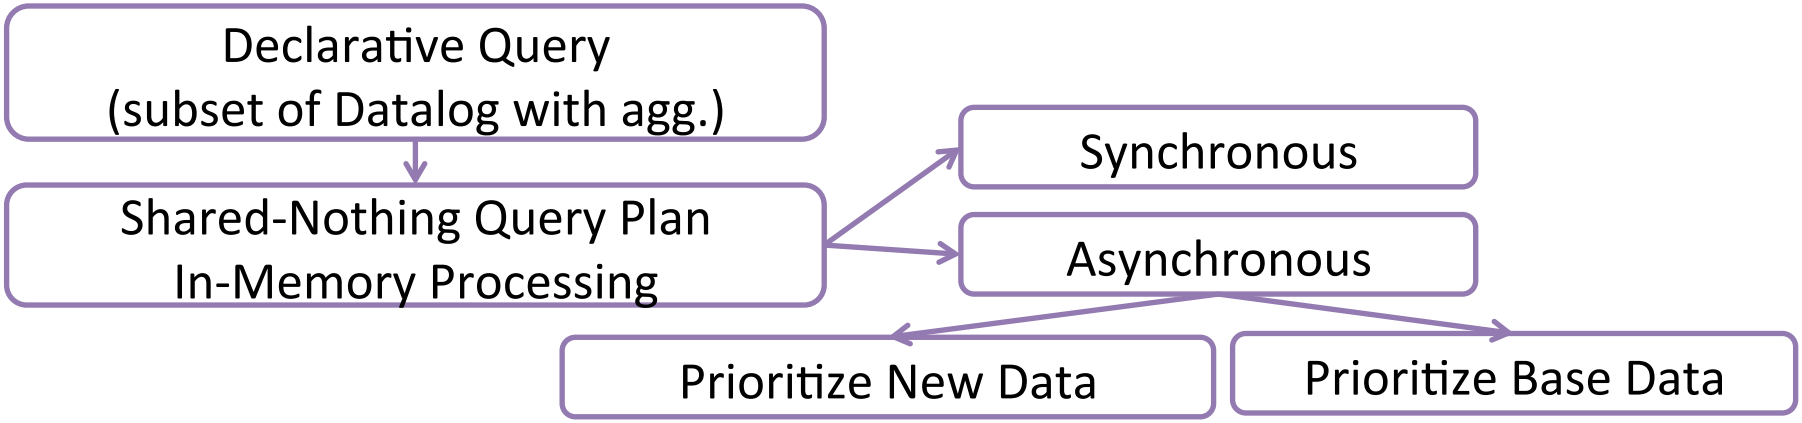


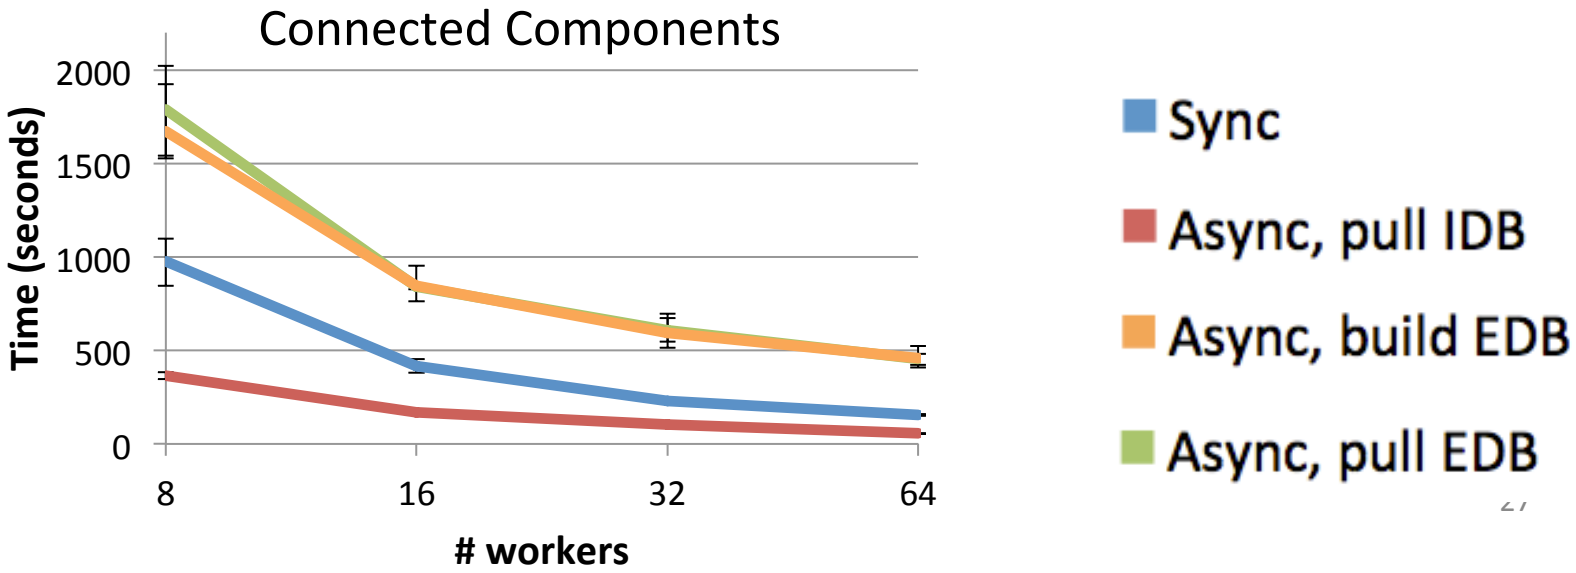

# Myria Changes the Face of Cloud Services

Today:

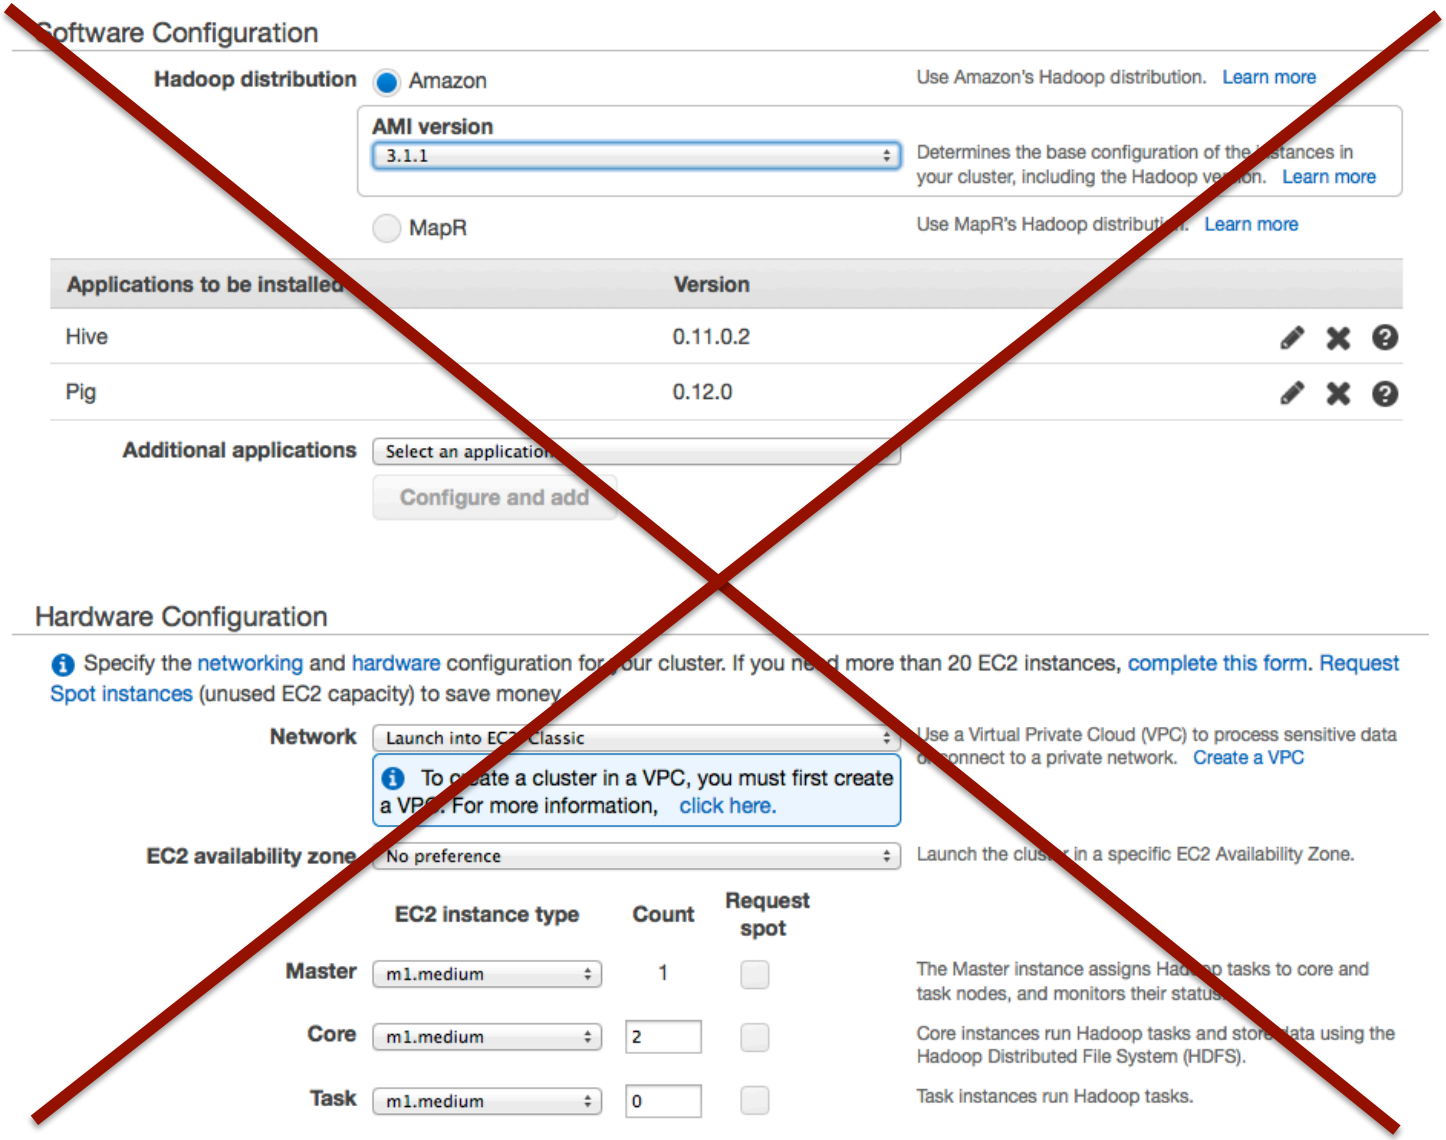

# Myria Changes the Face of Cloud Services

## With Myria: Personalized Service Level Agreements

Available SLAs

Tier #1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (22 ATTR.) FROM (JOIN 4 Tables) WHERE '<10% CONDITION'<br>SELECT (22 ATTR.) FROM (JOIN 5 Tables) WHERE '<1% CONDITION'<br>SELECT (27 ATTR.) FROM (JOIN 5 Tables) WHERE '<0.1% CONDITION' | 60 |
| SELECT (9 ATTR.) FROM (JOIN 4 Tables)<br>SELECT (27 ATTR.) FROM (JOIN 5 Tables) WHERE '<10% CONDITION' | 120 |
| SELECT (27 ATTR.) FROM (JOIN 5 Tables) | 180 |

Purchase @ $0.35/hour

Tier #2, benefits in addition to Tier#1

| Query Template | Runtime (seconds) |
|---|---|
| SELECT (27 ATTR.) FROM (JOIN 5 Tables) WHERE '<10% CONDITION' | 60 |
| SELECT (27 ATTR.) FROM (JOIN 5 Tables) | 120 |

Purchase @ $0.89/hour

# Conclusion

- We live in the “big data era”
- We need to build new data management and analytics systems to handle the data
- We work with domain scientists and companies to use these tools and enable new discoveries and new services

# Acknowledgments


- The University of Washington (UW) Database Group
- UW eScience
- UW astronomy
- And many others

- Our sponsors: NSF, ISTC Big Data, HP Labs, Microsoft Research, Yahoo!, Google, Intel, Amazon, Facebook, UW eScience