

**Microsoft**



Microsoft® Research

# Faculty Summit 2012

Riviera Maya, Mexico | May 23-25 | In partnership with CONACYT



# Data-Intensive Discoveries in Science: The Fourth Paradigm


Alexander Szalay

05/25/2012

# Gray's Laws of Data Engineering


Jim Gray:


- Scientific computing around data
- Need scale-out analysis
- Take the analysis
- Start with "20"
- Go from "work"




# Sloan Digital Sky Survey

- “The Cosmic Genome Project”
- Two surveys in one
  - Photometric survey in 5 bands
  - Spectroscopic redshift survey
- Data is public
  - 2.5 Terapixels of images => 5 Tpx
  - 10 TB of raw data => 120TB processed
  - 0.5 TB catalogs => 35TB in the end
- Started in 1992, finished in 2008
- Data volume enabled by Moore’s Law, Kryder’s Law

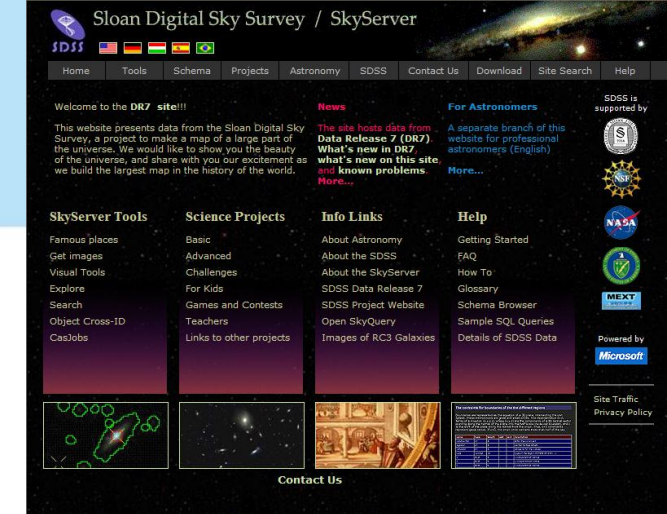








# SkyServer



- Prototype in 21st Century data access
  - One billion web hits in 10 years
  - 4,000,000 distinct users vs. 15,000 astronomers
  - The emergence of the “Internet scientist”
  - The world’s most used astronomy facility today
  - Collaborative server-side analysis done



# Impact of Sky Surveys

## Astronomy

## Sloan Digital Sky Survey tops astronomy citation list

NASA's Sloan Digital Sky Survey (SDSS) is the most significant astronomical facility, according to an analysis of the 200 most cited papers in astronomy published in 2006. The survey, carried out by Juan Madrid from McMaster University in Canada and Duccio Macchetto from the Space Telescope Science Institute in Baltimore, puts NASA's Swift satellite in second place, with the Hubble Space Telescope in third (arXiv:0901.4552).

Madrid and Macchetto carried out their analysis by looking at the top 200 papers using NASA's Astrophysics Data System (ADS), which charts how many times each paper has been cited by other research papers. If a paper contains data taken only from one observatory or satellite, then that facility is awarded all the citations given to that article. However, if a paper is judged to contain data from different facilities – say half from SDSS and half from Swift – then both

### Top 10 telescopes

| Rank | Telescope                       | Citations | Ranking in 2004 |
|------|---------------------------------|-----------|-----------------|
| 1    | Sloan Digital Sky Survey        | 1892      | 1               |
| 2    | Swift                           | 1523      | N/A             |
| 3    | Hubble Space Telescope          | 1078      | 3               |
| 4    | European Southern Observatory   | 813       | 2               |
| 5    | Keck                            | 572       | 5               |
| 6    | Canada–France–Hawaii Telescope  | 521       | N/A             |
| 7    | Spitzer                         | 469       | N/A             |
| 8    | Chandra                         | 381       | 7               |
| 9    | Boomerang                       | 376       | N/A             |
| 10   | High Energy Stereoscopic System | 297       | N/A             |

facilities are given 50% of the citations that paper received.

The researchers then totted up all the citations and produced a top 10 ranking (see table). Way out in front with 1892 citations is the SDSS, which has been

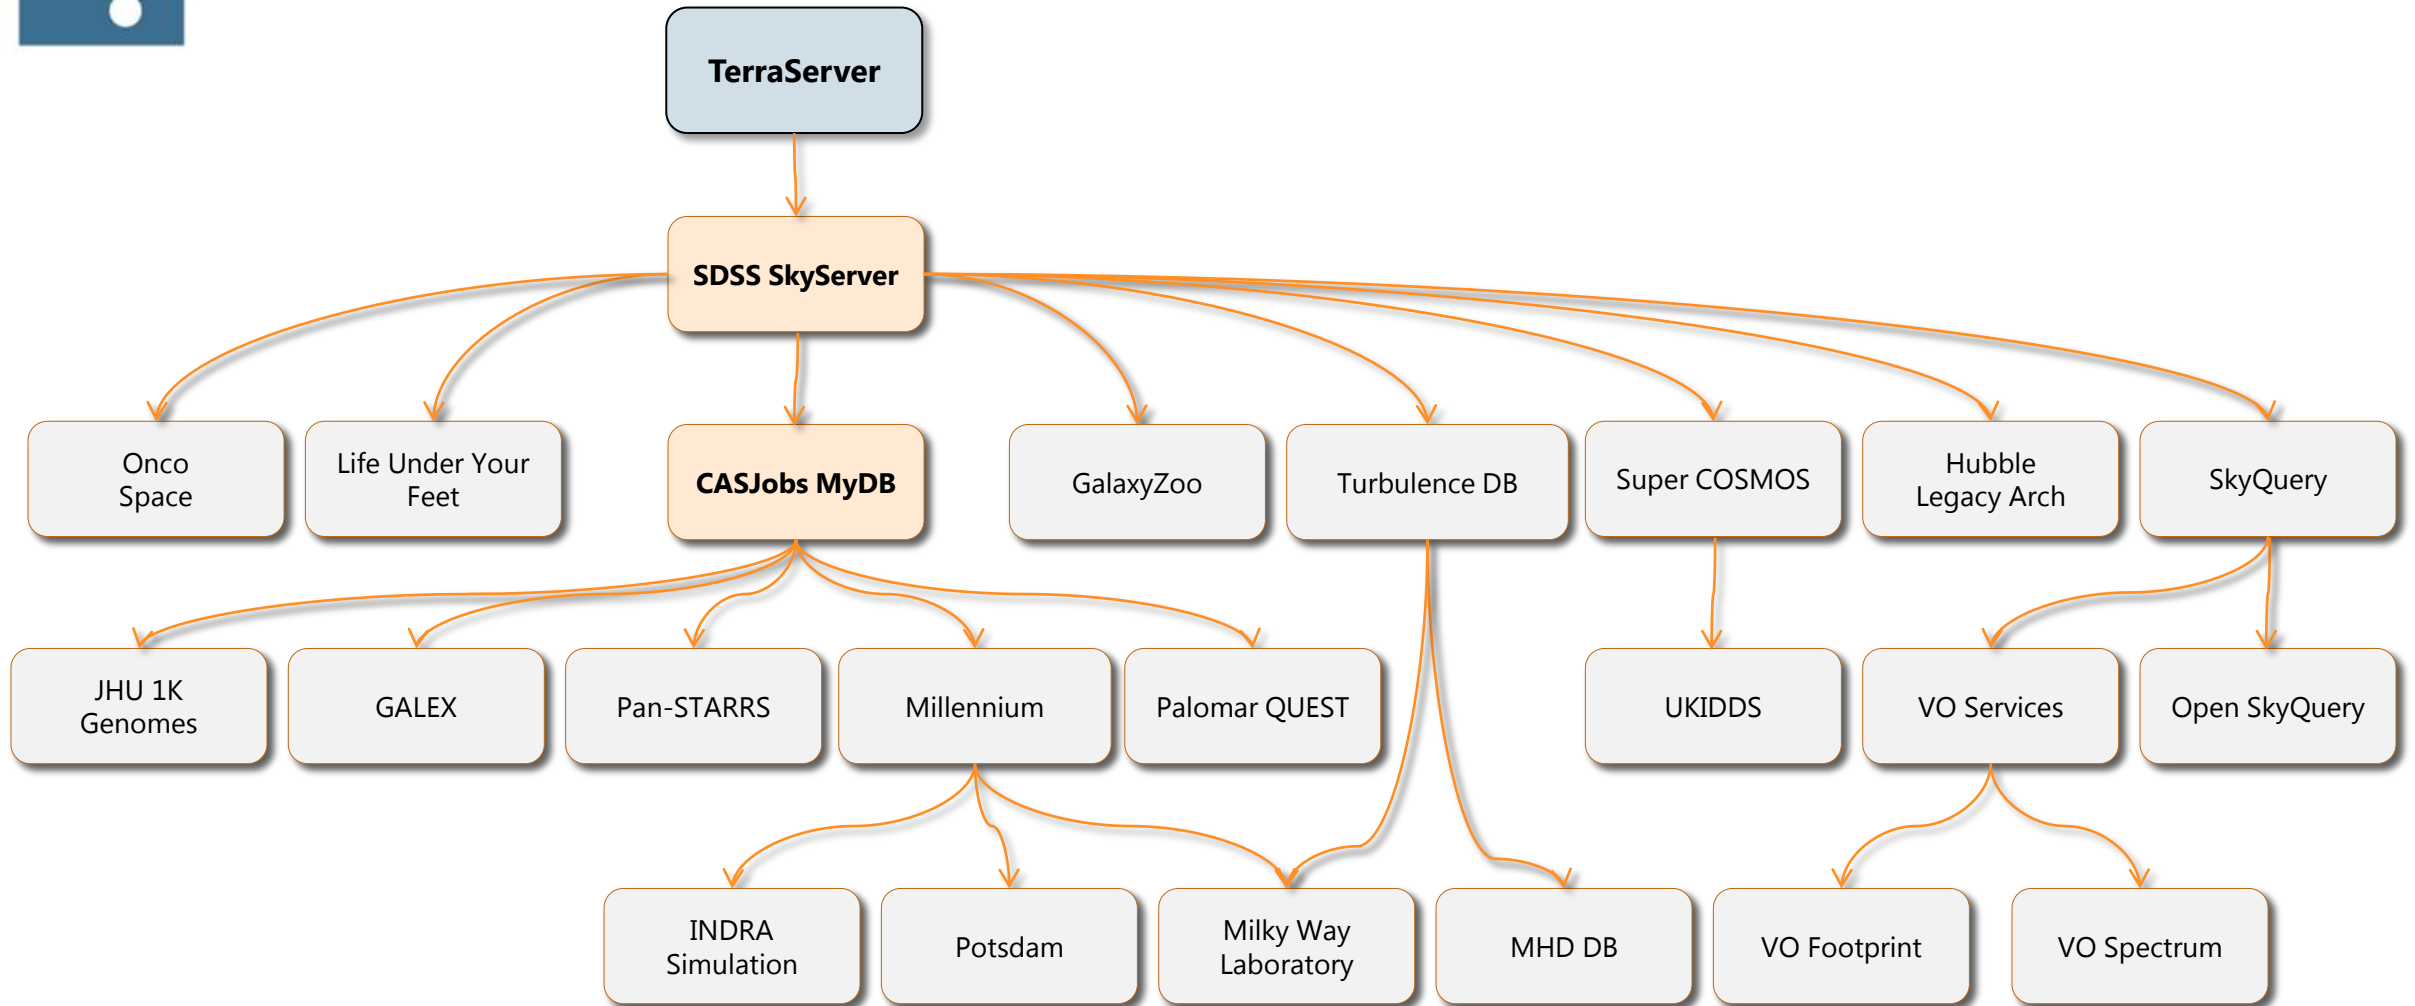
running since 2000 and uses the 2.5 m telescope at Apache Point in New Mexico to obtain images of more than a quarter of the sky. NASA's Swift satellite, which studies gamma-ray bursts, is second with 1523 citations, while the Hubble Space Telescope (1078 citations) is third.

Although the 200 most cited papers make up only 0.2% of the references indexed by the ADS for papers published in 2006, those 200 papers account for 9.5% of the citations. Madrid and Macchetto also ignored theory papers on the basis that they do not directly use any telescope data. A similar study of papers published in 2004 also puts SDSS top with 1843 citations. This time, though, the European Southern Observatory, which has telescopes in Chile, comes second with 1365 citations and the Hubble Space Telescope takes third spot with 1124 citations.

**Michael Banks**



# The SDSS Genealogy



# GalaxyZoo

- 40 million visual galaxy classifications by the public
- Enormous publicity (CNN, Times, Washington Post, BBC)
- 300,000 people participating, blogs, poems...
- Original discoveries by the public (Voorwerp, Green Peas)
- Chris Lintott et al



# Virtual Observatory Web Services


- E.g. VOFootprint
- Describe footprints and exposures on the sky
- Point in region test
- Intersection, difference, unions
- Integrated with database (SQL Server)
- Now used for
  - SDSS
  - all HST exposures (Hubble Legacy Archive)
  - GALEX, Chandra
  - Pan-STARRS, LSST





# SphericalLib .NET

- 8,500 lines of C# code – 20k total, auto-ported also to C++
  - OS independent (Windows, Un\*x w/ Mono)
  - Documentation via Sandcastle
- Great performance
  - Fast Boolean operations, e.g., intersect, union
  - Correct handling of degenerate cases and other imprecision of representation
  - Fast point in shape searches
  - GDI+ based visualization
  - Custom binary I/O
  - Exact areas





# Oncospace

Todd McNutt,  
JHU Radiation  
Oncology

## OncoSpace: Adapting the SkyServer Approach

- **Active Databases**

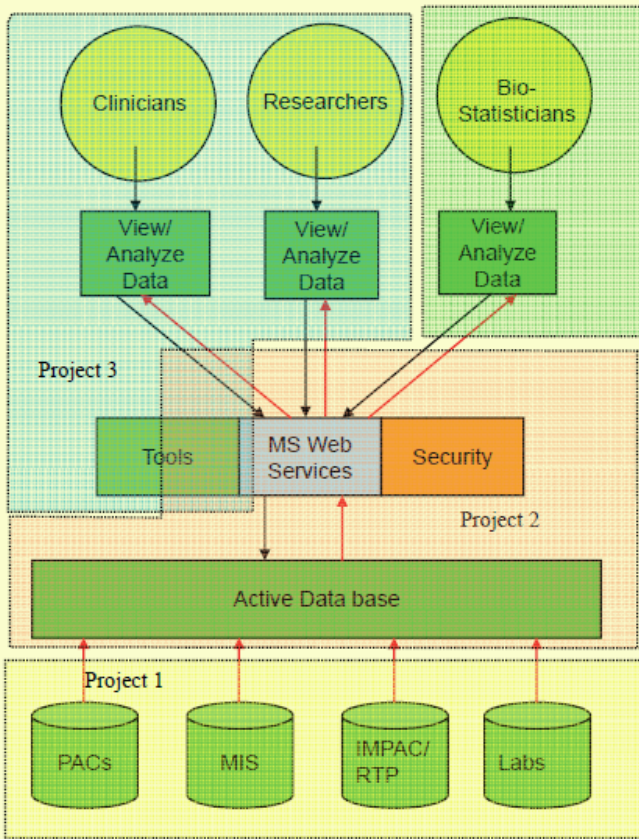
- There is too much data to move around,  
***take the analysis to the data!***

- Do all data manipulations at database
  - ***Build custom procedures and functions in the database***

- Established Web-service for broad access
  - Query across multiple databases

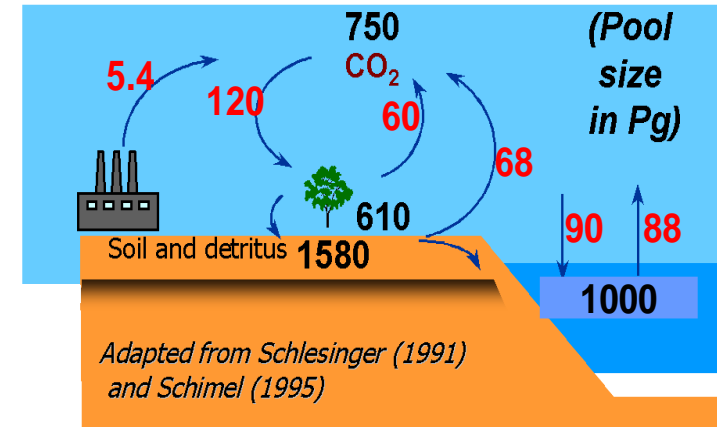

- Automatic parallelism guaranteed

Hopkins OncoSpace



# Life Under Your Feet

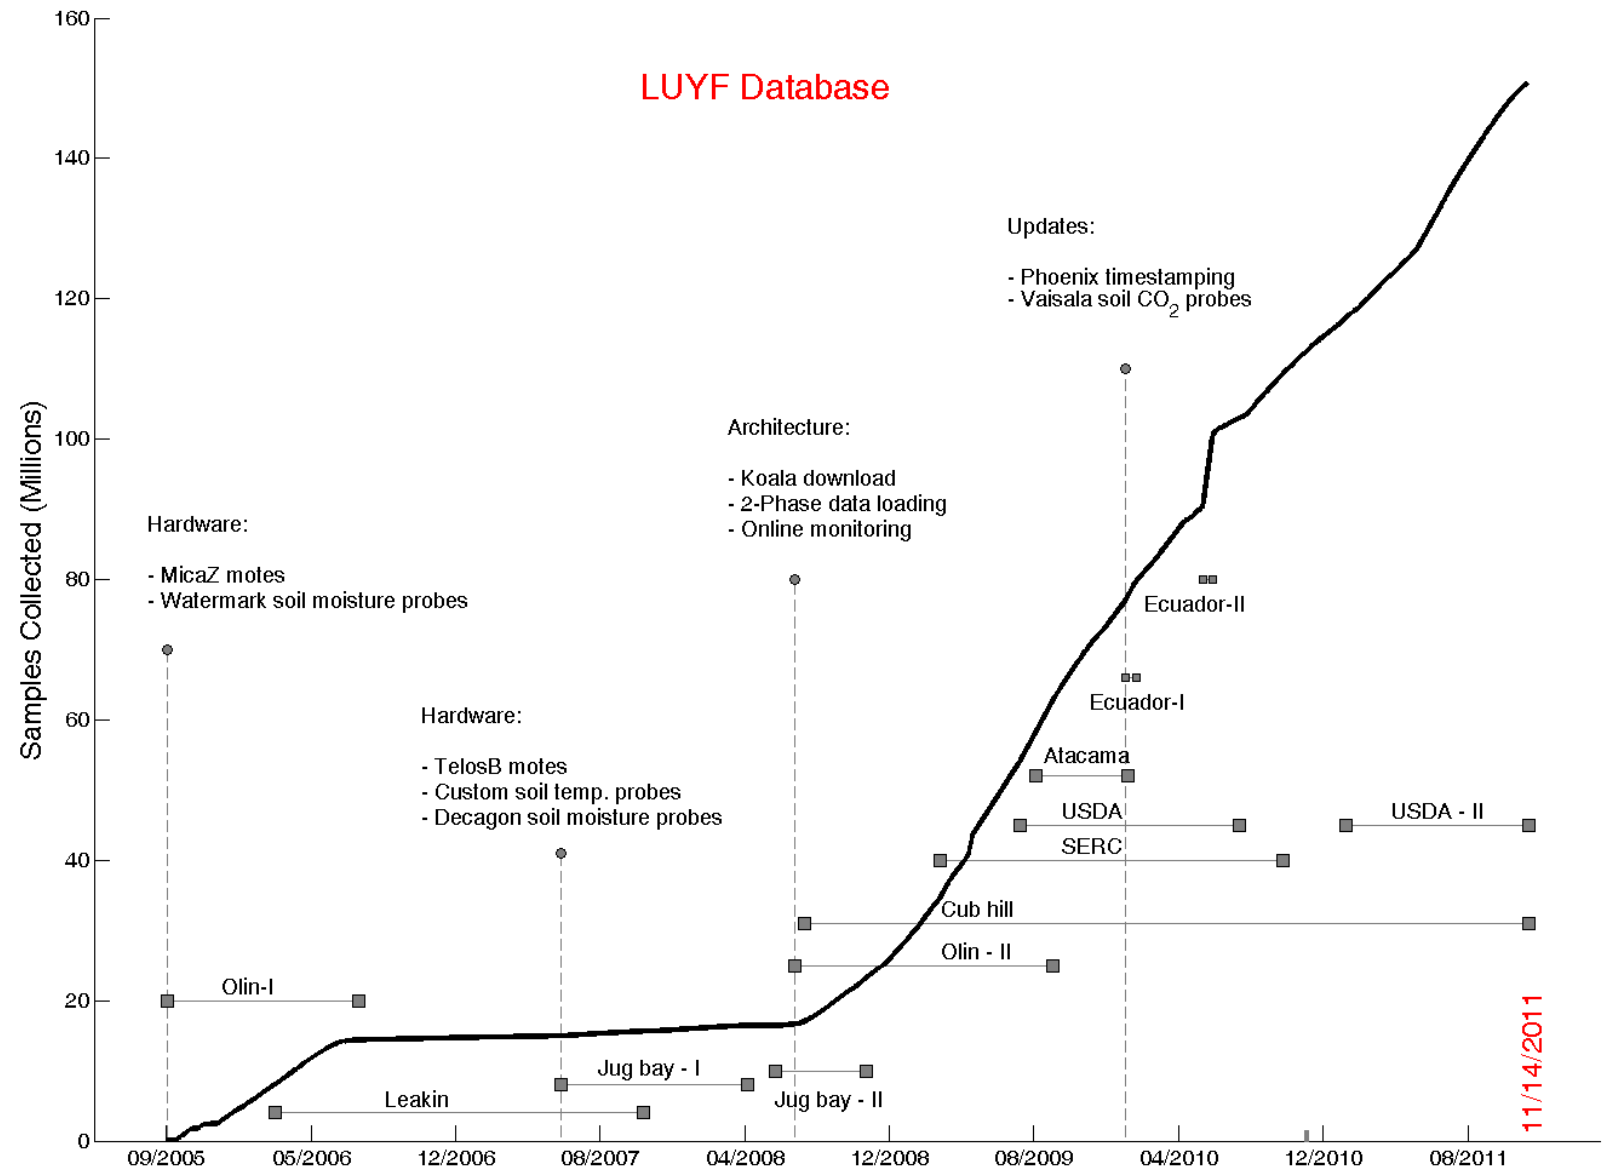
- Role of the soil in Global Change
  - Soil CO<sub>2</sub> emission thought to be **>15 times** of anthropogenic
  - Using sensors we can measure it directly, in situ, over a large area
- Wireless sensor network
  - Use 100+ wireless computers (motes), with 10 sensors each, monitoring
    - Air +soil temperature, soil moisture, ...
    - Few sensors measure CO<sub>2</sub> concentration
  - Long-term continuous data, 180K sensor days, 30M samples
  - Complex database of sensor data, built from the SkyServer
  - End-to-end data system, with inventory and calibration databases with K.Szlavec (Earth and Planetary), A. Terzis (CS)



# Cumulative Sensor Data in the LUYF DB

## Deployments:

- Baltimore/Washington
- Chesapeake Bay
- Brazil
- Ecuador
- Chile (Atacama)





# Deployments








# Data in HPC Simulations

- HPC is an instrument in its own right
- Largest simulations approach petabytes
  - from supernovae to turbulence, biology and brain modeling
- Need public access to the best and latest through interactive numerical laboratories
- Creates new challenges in
  - how to move the petabytes of data (high speed networking)
  - How to interface (virtual sensors, immersive analysis)
  - How to look at it (render on top of the data, drive remotely)
  - How to analyze (algorithms, scalable analytics)

# Silver River Transfer


- 150TB in less than 10 days from Oak Ridge to JHU using a dedicated 10G connection



# Immersive Turbulence

*"... the last unsolved problem of classical physics..." Feynman*

- Understand the nature of turbulence
  - Consecutive snapshots of a large simulation of turbulence: now 30 Terabytes
  - Treat it as an experiment, **play** with the database!
  - **Shoot test particles** (sensors) from your laptop into the simulation, like in the movie Twister
  - Next: 70TB MHD simulation
- New paradigm for analyzing simulations!




with C. Meneveau (Mech. E), G. Eyink (Applied Math), R. Burns (CS)



# Typical Daily Usage

**Turbulence Database Usage by Day**





*2011: exceeded 100B points publicly delivered*





# Spatial queries, random samples



- Spatial queries require multi-dimensional indexes.
- $(x,y,z)$  does not work: need discretisation
  - index on  $(ix,iy,iz)$  with  $ix=\text{floor}(x/8)$  etc
- More sophisticated: space filling curves
  - bit-interleaving/octtree/Z-Index
  - Peano-Hilbert curve
  - Need custom functions for range queries
  - Plug in modular space filling library (Budavari)



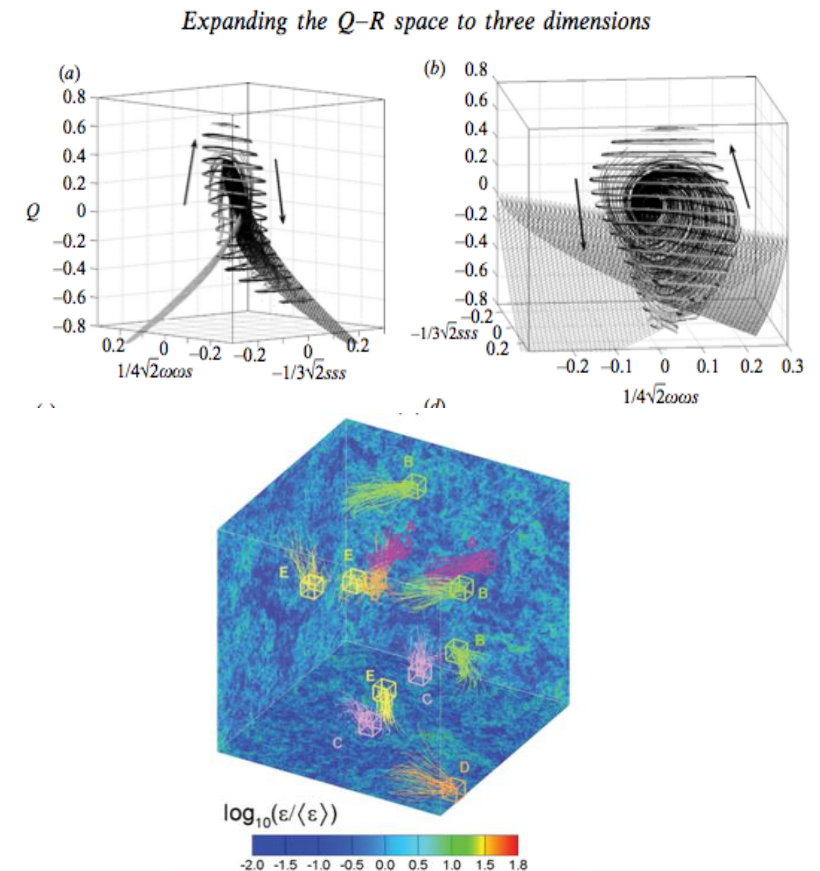
# Turbulence Research with the Database

Experimentalists testing PIV-based pressure-gradient measurement  
(X. Liu & Katz, 61 APS-DFD meeting, November 2008)

Measuring velocity gradient using a new set  
of 3 invariants,  
Luethi, Holzner & Tsinober,  
J. Fluid Mechanics 641, pp. 497-507 (2010)

Lagrangian time correlation in turbulence  
Yu & Meneveau,  
Phys. Rev. Lett. 104, 084502 (2010)

Now about 20 papers published




# Cosmological Simulations

In 2000 cosmological simulations had  $10^{10}$  particles and produced over 30TB of data (Millennium)

- Build up dark matter halos
- Track merging history of halos
- Use it to assign star formation history
- Combination with spectral synthesis
- Realistic distribution of galaxy types
- **More than 1,000 CASJobs/MyDB users**


Today: simulations with  $10^{12}$  particles and PB of output are under way (MillenniumXXL, Silver River, Exascale Sky)



# The Milky Way Laboratory

- Cosmology simulations as immersive laboratory for general users
- Via Lactea-II (20TB) as prototype, then Silver River (50B particles) as production (50M CPU hours)
- 800+ hi-rez snapshots (2.6PB) => 800TB in DB
- Users can insert test particles (dwarf galaxies) into the system and follow trajectories in pre-computed simulation
- Compute dark matter annihilation maps interactively
- Users will interact remotely with a PB in 'real time'

Madau, Rockosi, Szalay, Wyse, Silk,  
Kuhlen, Lemson, Westermann, Blakeley







# Visualizing Petabytes

- Send the rendering to the data ...
- It is easier to send a HD 3D video stream to the user than all the data
- Interactive visualizations driven remotely
- Visualizations are becoming IO limited
- It is possible to build individual servers with extreme data rates
- Prototype on turbulence simulation already works: data streaming directly from DB to GPU
- N-body simulations next





# Streaming Visualization of Turbulence






2000: 1TB

2008: 1PB


2012: 8PB





# Current Data-Intensive Projects at JHU

| Discipline        | data [TB] |
|-------------------|-----------|
| Astrophysics      | 930       |
| HEP/Material Sci. | 394       |
| CFD               | 425       |
| BioInformatics    | 414       |
| Environmental     | 660       |
| Total             | 2823      |



19 projects total proposed for the Data-Scope, more coming,  
data lifetimes between 3 mo and 3 yrs





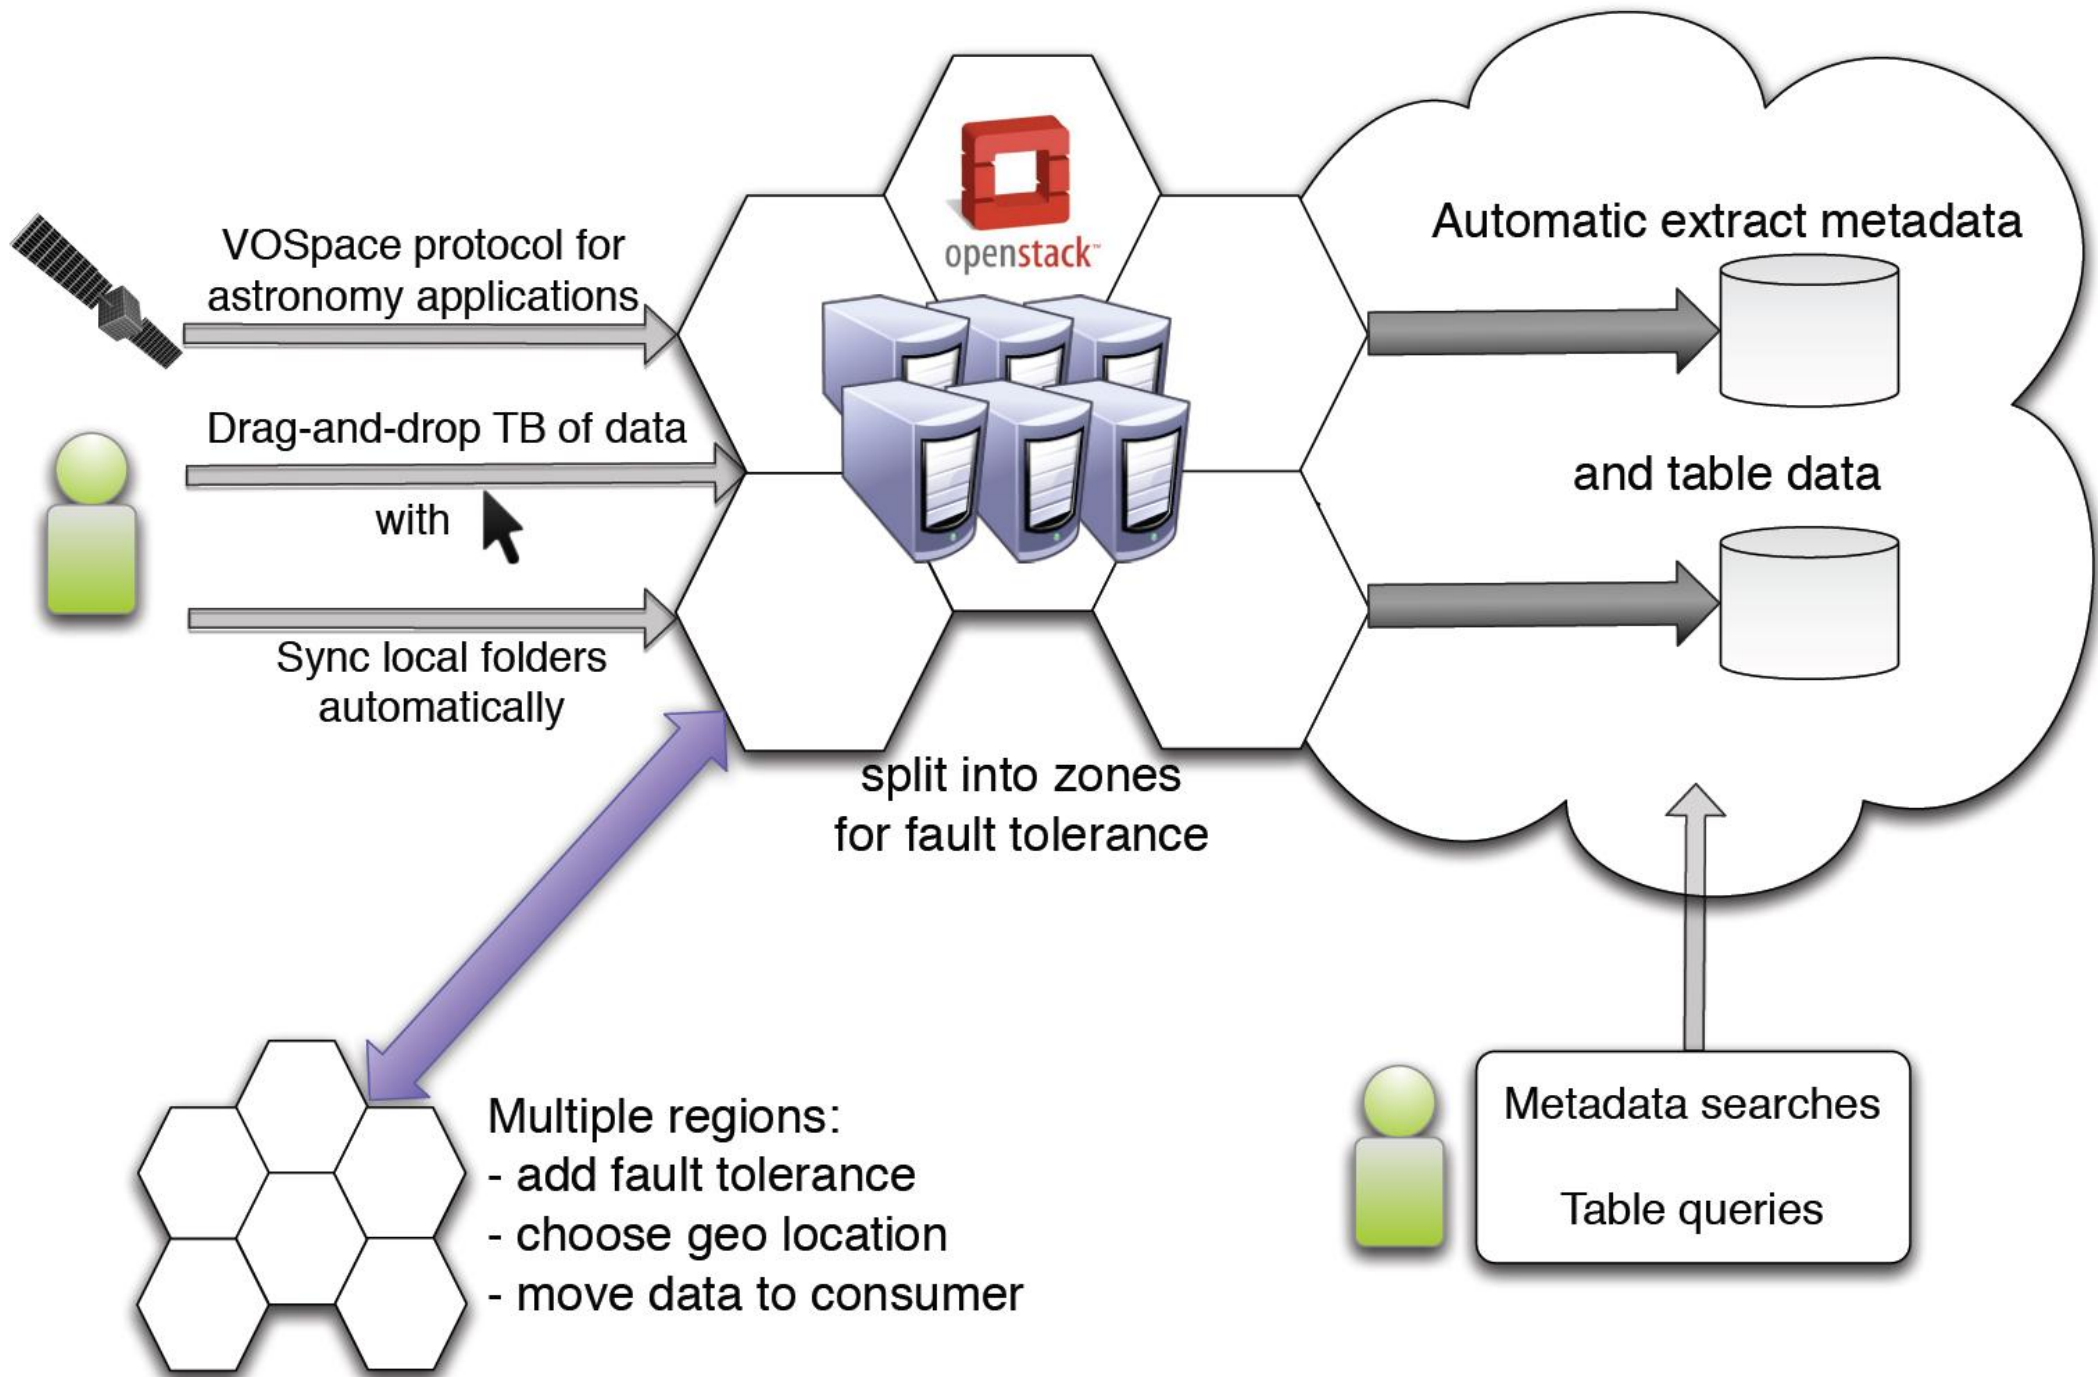
# The Long Tail

- The “Long Tail” of a huge number of small data sets
  - The integral of the “long tail” is big!
- Facebook: bring many small, seemingly unrelated data to a single cloud and new value emerges
  - What is the science equivalent?
- The DropBox lesson
  - Simple interfaces are much more powerful than complex ones
  - API public




# DropBox for Science

- DropBox+MyDB+VOspace
  - Dmitry Mishin (JHU)
- Free cloud storage for the VO Community
- Automatic harvesting of metadata upon “drop”
- Metadata and tables go into MyDB, can be joined with contexts
- Metadata immediately available for services (e.g. VOFootprint)
  - After a certain amount of data new connections appear
- Generalize to other areas of science (environmental science)



# Sociology


- **Broad sociological changes**
  - Convergence of Physical and Life Sciences
  - Data collection in ever larger collaborations
  - Virtual Observatories: CERN, VAO, NCBI, NEON, OOI,...
  - Analysis decoupled, off archived data by smaller groups
  - Emergence of the citizen/internet scientist
  - Impact of demographic changes in science
- **Need to start training the next generations**
  - $\Pi$ -shaped vs I-shaped people
  - Early involvement in "Computational thinking"





# Summary

- Science is increasingly driven by data (large and small)
- Large data sets are here, COTS solutions are not
- Changing sociology
- From hypothesis-driven to data-driven science
- We need new instruments: “microscopes” and “telescopes” for data
- There is also a problem on the “long tail”
- Same problems present in business and society
- Data changes not only science, but society
- A new, Fourth Paradigm of Science is emerging...





***Microsoft***<sup>®</sup>

© 2012 Microsoft Corporation. All rights reserved. Microsoft, Windows, and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.

The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.