The SDSS SkyServer and beyond

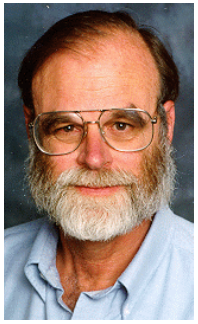

Alex Szalay

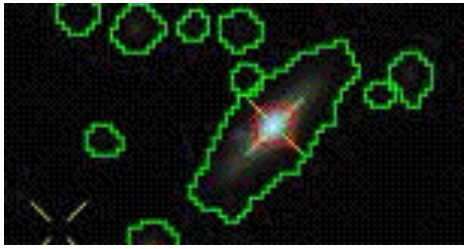

Historical Background

- The Sloan Digital Sky Survey (SDSS)
  The “Cosmic Genome Project”
  - *5 color images of ¼ of the sky*
  - *Pictures of 300 million celestial objects*
  - *Distances to the closest 1 million galaxies*
- JHU: build the public archive for the SDSS
- Lots of debate who the archive is for
  - *“power users”*
  - *“astronomers”*
  - *“students and amateurs”*
  - *“wide public”*
- Interesting challenge in digital publishing
  - *We have to publish first in order to analyze*

Sloan Digital Sky Survey

- **“The Cosmic Genome Project”**
- Started in 1992, finished in 2008
- 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
- Database and spectrograph built at JHU (SkyServer)
- Data served from FNAL
- Now SDSS-3, imaging completed
- SDSS-3 data served from JHU

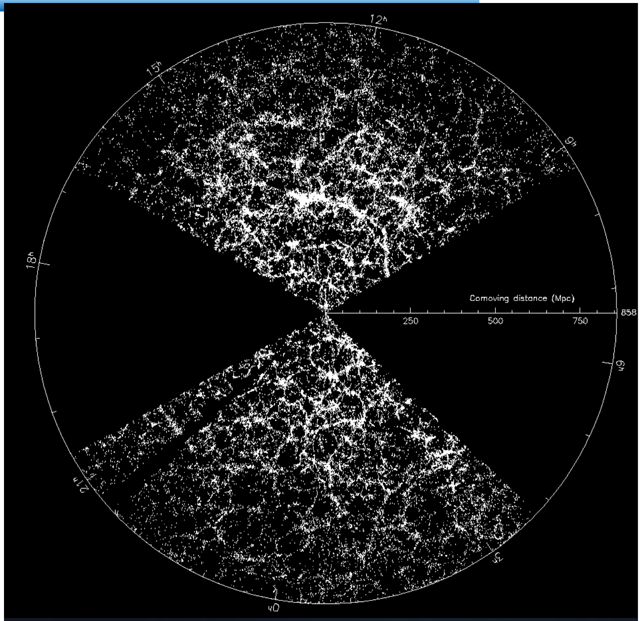

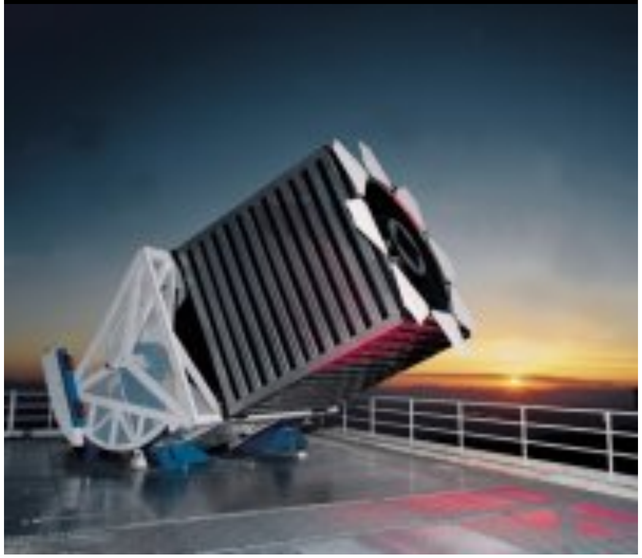

Skyserver

- Prototype in 21st Century data access
  - *1 billion web hits in 11 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 by 5K astronomers (30%)*

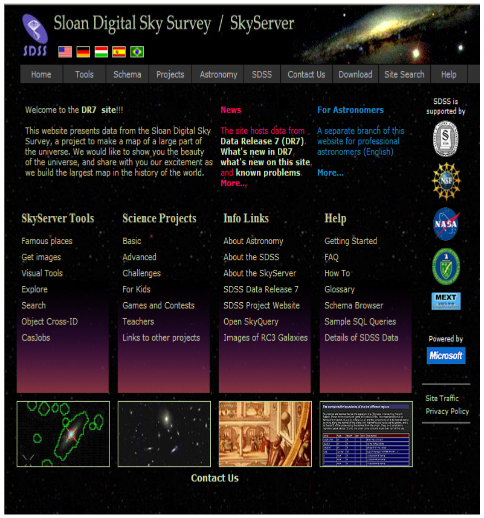

GalaxyZoo

- 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

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 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

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

SkyServer Goals

- Provide easy, visual access to exciting new data
  - *"hot off the press"*
- Illustrate that advanced content does not mean a cumbersome interface
- Understand new ways of publishing scientific data
- Demonstrate how to take analyses inside the DB
  - *Heavy use of user defined functions*
- Target audience
  - *Advanced high-school students, amateur astronomers, wide public*
- Multilingual capabilities built in from the start
  - *Heavy use of stylesheets, language branches*

DB Loading

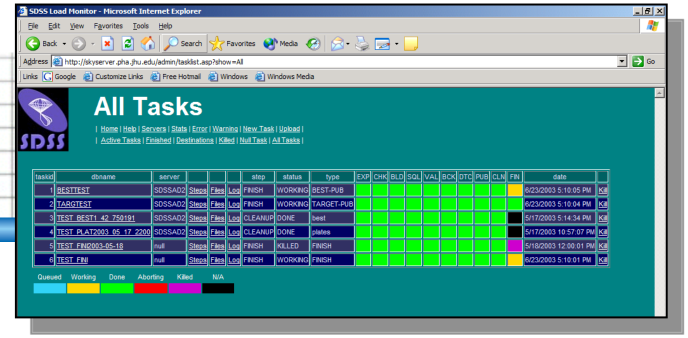

- Wrote automated table driven workflow system for loading
  - *Two-phase parallel load*
  - *Over 16K lines of SQL code, mostly data validation*
- Loading process was extremely painful
  - *Lack of systems engineering for the pipelines*
  - *Lots of foreign key mismatches*
  - *Fixing corrupted files (RAID5 disk errors)*
  - *Most of the time spent on scrubbing data*
- Once data is clean, everything loads in 1 week
- Reorganization of data is about 1 week

Data Delivery

- Small requests (<100MB)
  - *Anonymous, putting data on the stream*
- Medium requests (<1GB)
  - *Queues with resource limits*
- Large requests (>1GB)
  - *Save data in scratch area and use asynch delivery*
  - *Only practical for large/long queries*
- Iterative requests/workbench
  - *Save data in temp tables in user space*
  - *Let user manipulate via web browser*
- **Paradox:** if we use web browser to submit, users want immediate response even from large queries

CASJOBS/MyDB: Workbench

- Need to register 'power users', with their own DB
- Query output goes to 'MyDB'
- Can be joined with source database
- Results are materialized from MyDB upon request
- Users can do:
  - *Insert, Drop, Create, Select Into, Functions, Procedures*
  - *Publish their tables to a group area*
- Data delivery via the CASJobs (C# WS)
  - *Batch scheduler for large queries*

=> Sending analysis to the data!

MyDB

- Implemented by Nolan Li, from user feedback
- Results are materialized from MyDB upon request
- Users can collaborate!
  - *Insert, Drop, Create, Select Into, Functions*
  - ***Publish/share** their tables to a group area*
  - *Flexibility "at the edge"/ Read-only big DB*
- 6,800 registered users

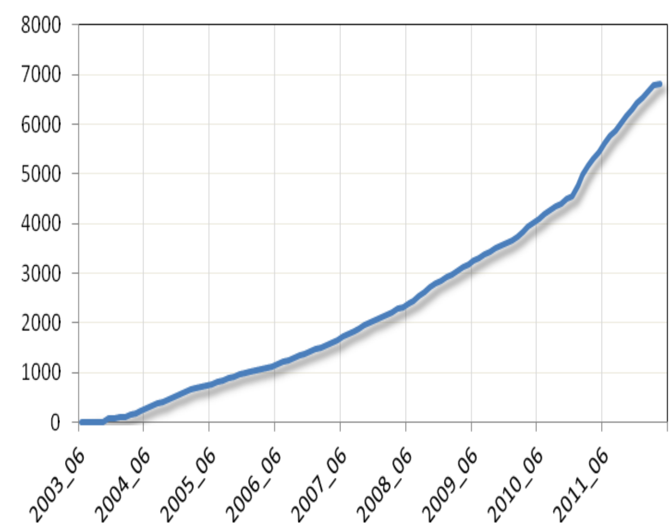

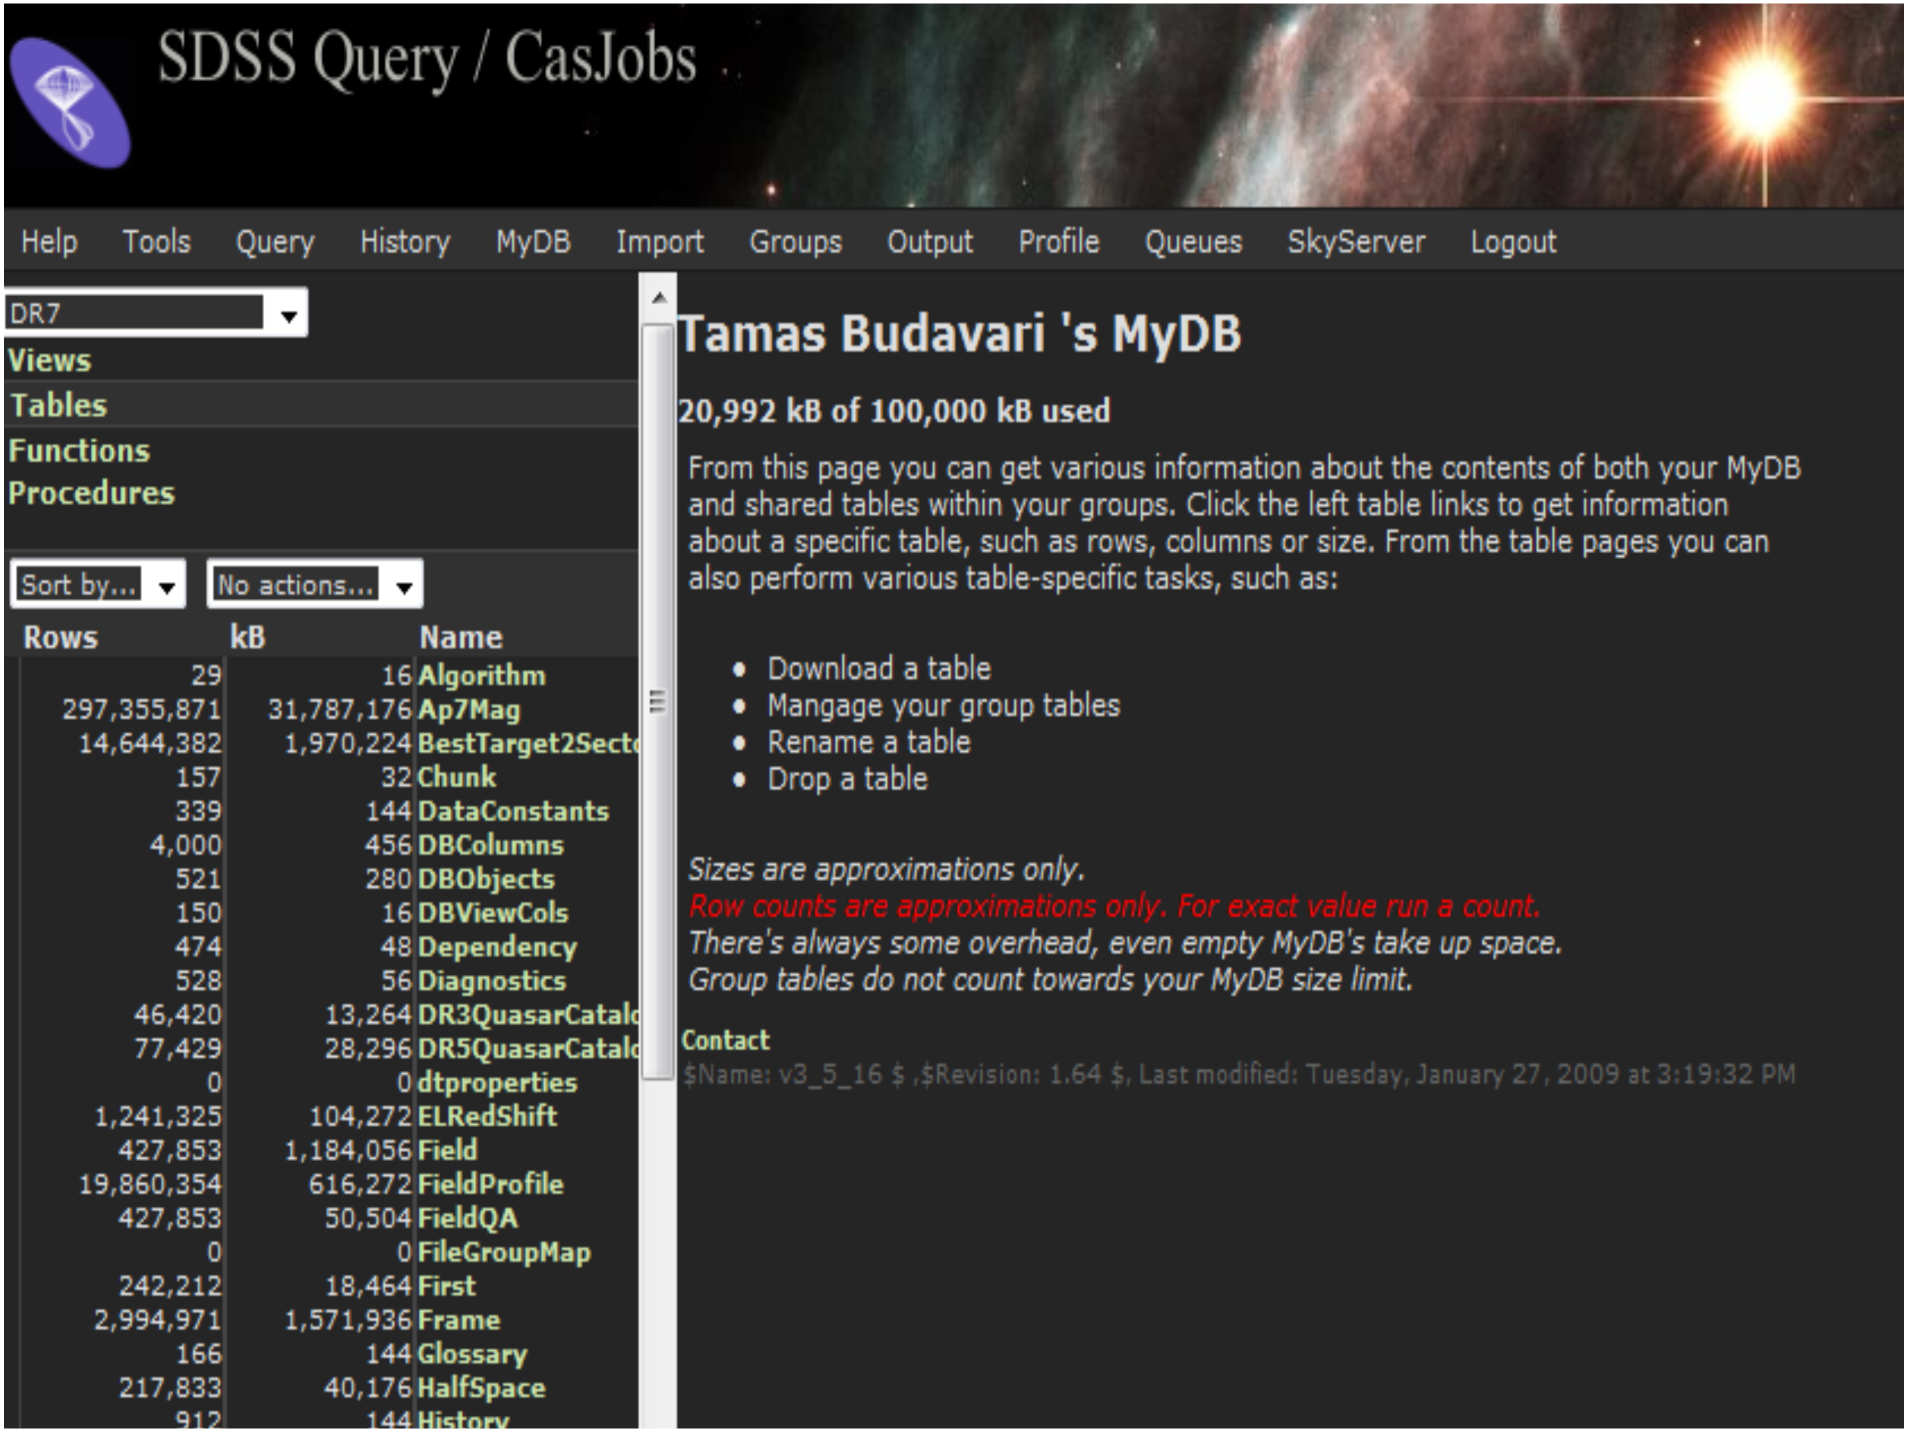
SDSS Query / CasJobs
Help Tools Query History MyDB Import Groups Output Profile Queues SkyServer Logout
DR7
Views
Tables
Functions
Procedures
Sort by...
No actions...
Rows kB Name
29 16 Algorithm
297,355,871 31,787,176 Ap7Mag
14,644,382 1,970,224 BestTarget2Secto
157 32 Chunk
339 144 DataConstants
4,000 456 DBColumns
521 280 DBObjects
150 16 DBViewCols
474 48 Dependency
528 56 Diagnostics
46,420 13,264 DR3QuasarCatalo
77,429 28,296 DR5QuasarCatalo
0 0 dtproperties
1,241,325 104,272 ELRedShift
427,853 1,184,056 Field
19,860,354 616,272 FieldProfile
427,853 50,504 FieldQA
0 0 FileGroupMap
242,212 18,464 First
2,994,971 1,571,936 Frame
166 144 Glossary
217,833 40,176 HalfSpace
912 144 History
Tamas Budavari 's MyDB
20,992 kB of 100,000 kB used
From this page you can get various information about the contents of both your MyDB and shared tables within your groups. Click the left table links to get information about a specific table, such as rows, columns or size. From the table pages you can also perform various table-specific tasks, such as:
Download a table
Mangage your group tables
Rename a table
Drop a table
Sizes are approximations only.
Row counts are approximations only. For exact value run a count.
There's always some overhead, even empty MyDB's take up space.
Group tables do not count towards your MyDB size limit.
Contact
$Name: v3_5_16 $,$Revision: 1.64 $, Last modified: Tuesday, January 27, 2009 at 3:19:32 PM

Data Versions

- June 2001: EDR
- Now at DR5, with 2.4TB
- 3 versions of the data
  - *Target, Best, Runs*
  - *Total catalog volume 5TB*
- Data publishing: once published, must stay
- SDSS: DR1 is still used

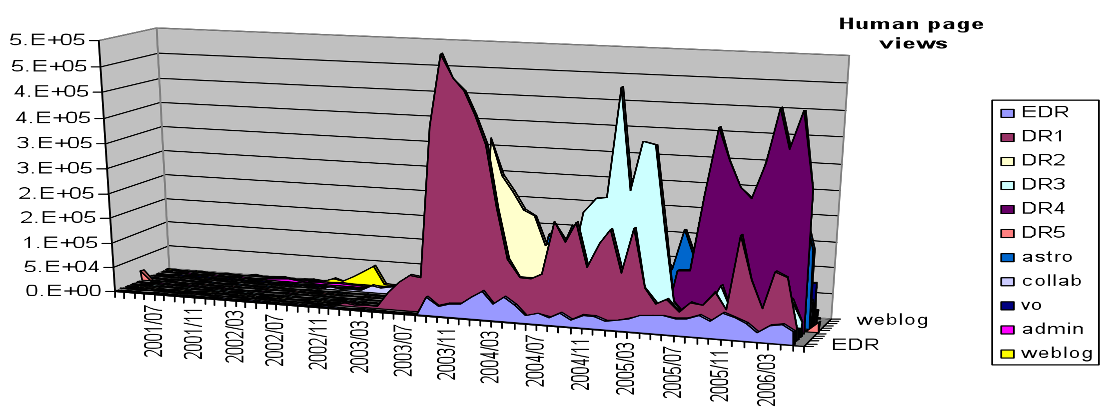

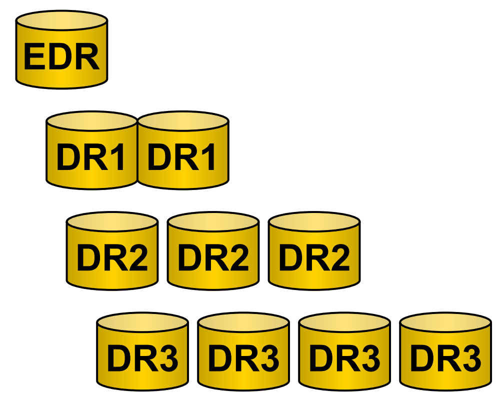

EDR: Early Data Release

- SDSS Early Data Release (June 6, 2001)
- 100 GB catalogs, few hundred square degrees
- SkyServer aimed solely at public outreach
- Built in 2 weeks by Szalay and Gray (20 hour days)
- Web site design by Szalay
- Images converted in PhotoShop scripts
- Content writing done by Stephen Landy
- Hardware donated by Compaq
- Highly interactive, using browser independent DHTML (“browser hell”)

DR1: Data Release 1

- The first main data release of SDSS (May 2003)
- 1.1TB of catalogs, linked to 6TB of low level data
- SkyServer has undergone a major facelift
  - *New graphic design by Curtis Wong, Asta Roseway (MS)*
  - *Modified stylesheets and embedded scripts only*
  - *Web site translated in 2 days*
- New visual tools using Web Services
  - *Szalay, Gray, Maria Nieto-SantiSteban*
- API's published
- Formal helpdesk in place
- Created MySkyServer
  - *0.65GB laptop version*

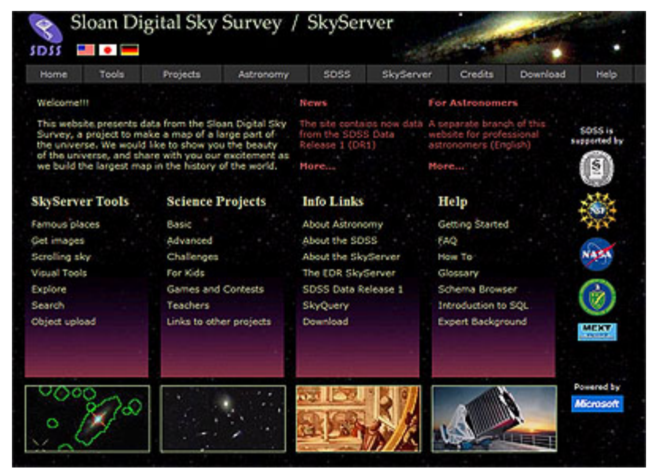

DR2: Data Release 2

- Live in March 15, 2004, with 2.2 TB of catalogs
- Only incremental changes in interface
- Web site under source control
- Color images dramatically improved
- New translations under way
  - *Japanese, French, German, Spanish, Hungarian*
- Tools overhauled
  - *now embraced by professional astonomers*
- Enormously increased traffic
- Moving to 3-way web front end + 3 DB servers
- Collaborative tools: MyDB with group access

Tutorials and Guides

- Developed by Jordan and Postdocs
  - *How to use Excel*
  - *How to use a database (guide to SQL)*
  - *Expert advice on SQL*
- Automated on-line documentation
  - *Ani Thakar, Roy Gal*
  - *Database information, Glossary, Algorithms*
  - *Searchable Help*
  - *All stored in the DB, and generated on the fly*

Visual Tools

- Goal:
  - *Connect pixel space to objects without typing queries*
  - *Browser interface, using common paradigm (MapQuest)*
- Challenge:
  - *Images: 200K x 2K x1.5K resolution x 5 colors = 3 Terapix*
  - *300M objects with complex properties*
  - *20K geometric boundaries and about 6M 'masks'*
  - *Need large dynamic range of scales (2^13)*
- Assembled from a few building blocks:
  - *Image Cutout Web Service*
  - *SQL query service + database*
  - *Images+overlays built on server side -> simple client*

User Level Services

- Three different applications on top of the same core
  - *Finding Chart (arbitrary size)*
  - *Navigate (fixed size, clickable navigation)*
  - *Image List (display many postage stamps on same page)*
- Linked to
  - *One another*
  - *Image Explorer (link to complex schema)*
  - *On-line documentation*

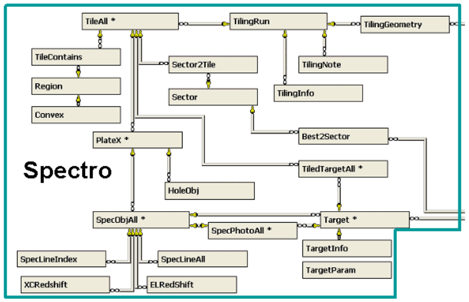

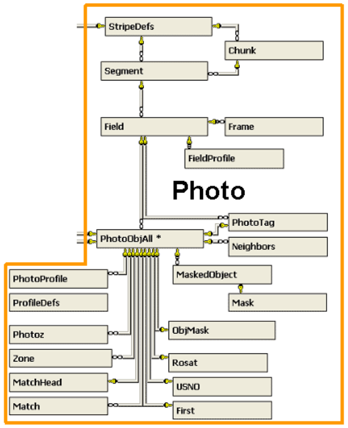

Images

- 5 bands, 2048x1489 resolution (u,g,r,i,z), 6MB each
  - *Raw size 200Kx6MB = 1.2TB*
  - *For quick access they must be stored in the DB*
  - *It has to show well on screens, remapping needed*
  - *Remapping must be uniform, due to image mosaicking*
- Built composite color, using lambda mapping
  - *(g->B, r->G, i->R), u,z was too noisy*
- Many experiments, discussions with Robert Lupton
  - *Asinh compression*
- Resulting image stored as JPEG
  - *From 30MB->300kB : a factor 100 compression*

Object Overlays

- Object positions stored in (ra,dec)
- At run time, convert (ra,dec)-> (screen_x, screen_y)
- Plotting pixel space quantities, like outlines:
  - *We could do (x,y)->(ra,dec)->(screen)*
  - *For each field we store local affine transformation matrix:*
    - (x,y) -> (screen)
- Apply local projection matrix and plot in pixel coordinates
  - *GDI plots correctly on the screen!*

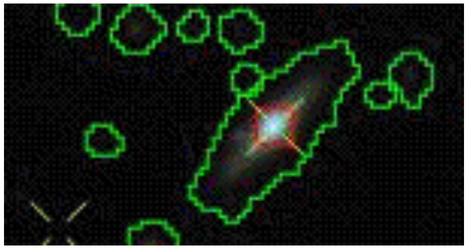

- Whole web service less than 1500 lines of C# code

Geometries

- SDSS has lots of complex boundaries
  - *60,000+ regions*
  - *6M masks, represented as spherical polygons*
- A GIS-like library built in C++ and SQL
- Now converted to C# for direct plugin into SQL Server2005 (17 times faster than C++)
- Precompute arcs and store in database for rendering
- Functions for point in polygon, intersecting polygons, polygons covering points, all points in polygon
- Using spherical quadtrees (HTM)

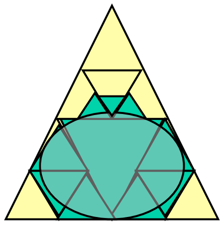

Things Can Get Complex

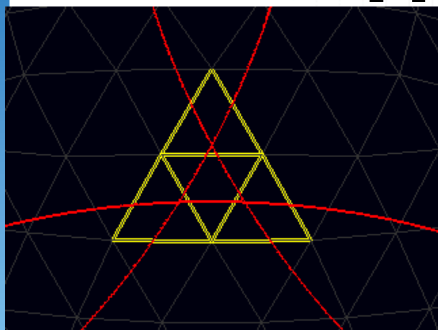

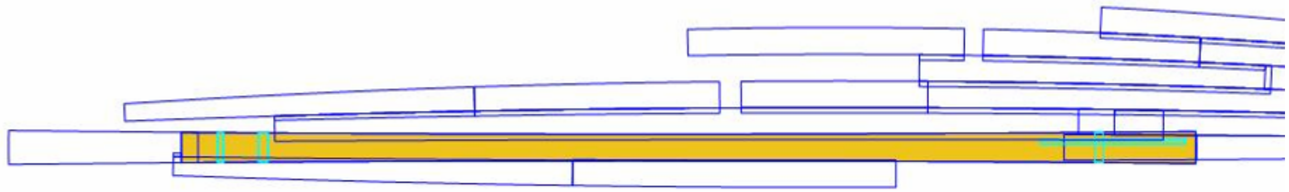

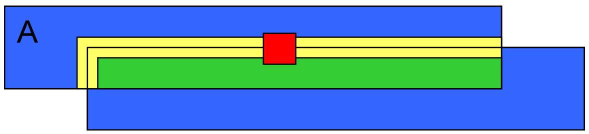

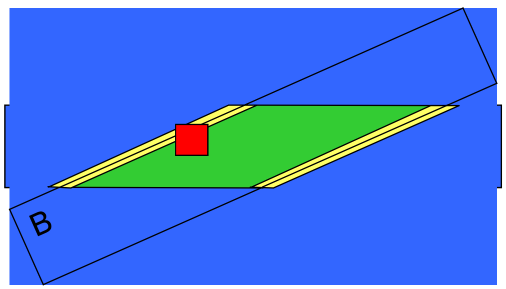

Green area: $A \cap (B-\varepsilon)$ should find B if it contains an A and not masked
Yellow area: $A \cap (B\pm\varepsilon)$ is an edge case may find B if it contains an A.

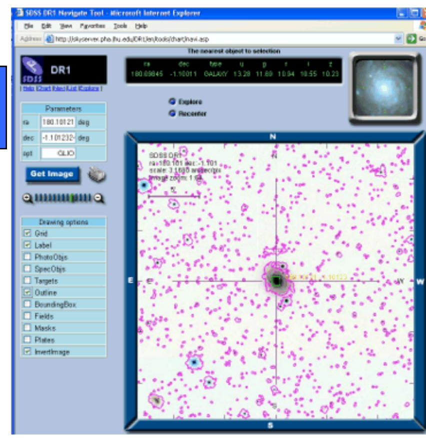

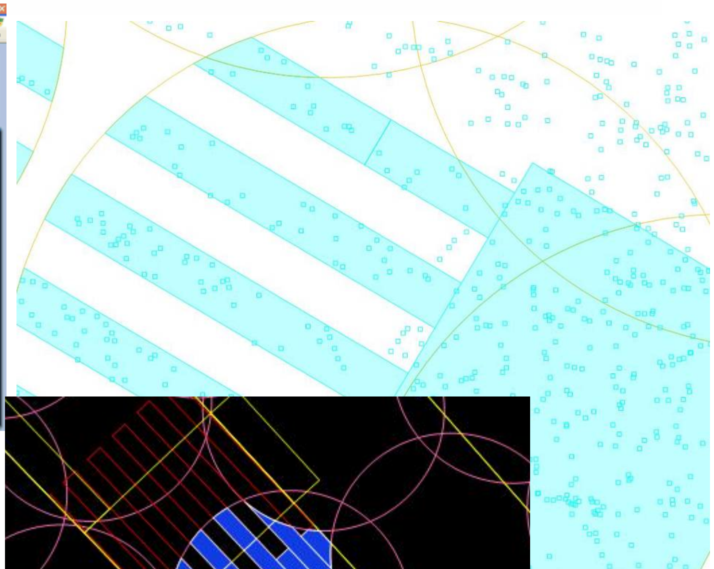

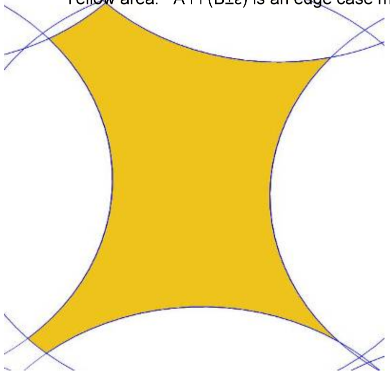

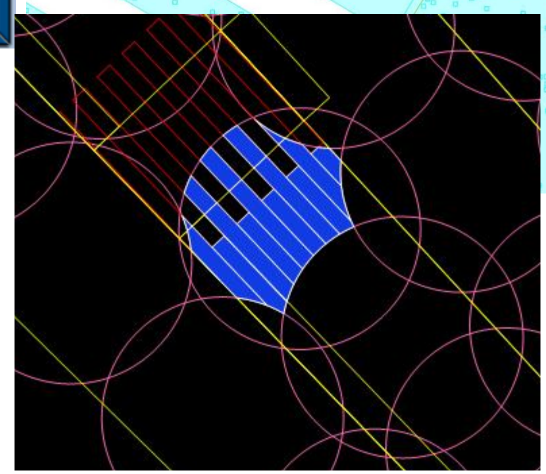

Trends

CMB Surveys		
1990	COBE	1000
2000	Boomerang	10,000
2002	CBI	50,000
2003	WMAP	1 Million
2008	Planck	10 Million

Angular Galaxy Surveys		
1970	Lick	1M
1990	APM	2M
2005	SDSS	200M
2008	VISTA	1000M
2012	LSST	3000M

Time Domain

- QUEST
- SDSS Extension survey
- Dark Energy Camera
- PanStarrs
- SNAP…
- LSST…

Galaxy Redshift Surveys		
1986	CfA	3500
1996	LCRS	23000
2003	2dF	250000
2005	SDSS	750000

Petabytes/year by the end of the decade…

Current Status

- SDSS-2 finished with DR7
  - *Database a bit over 10TB*
- SDSS-3
  - *One last run of imaging, completed area between Southern stripes, then turned off imaging camera*
  - *Rebuilt spectrographs, mostly LRG (BOSS)*
  - *DR8 in 2011, DR9 in end of July 2012*
  - *Database over 12TB*
- Planning started for AS3 (After SDSS 3)

The SDSS Genealogy

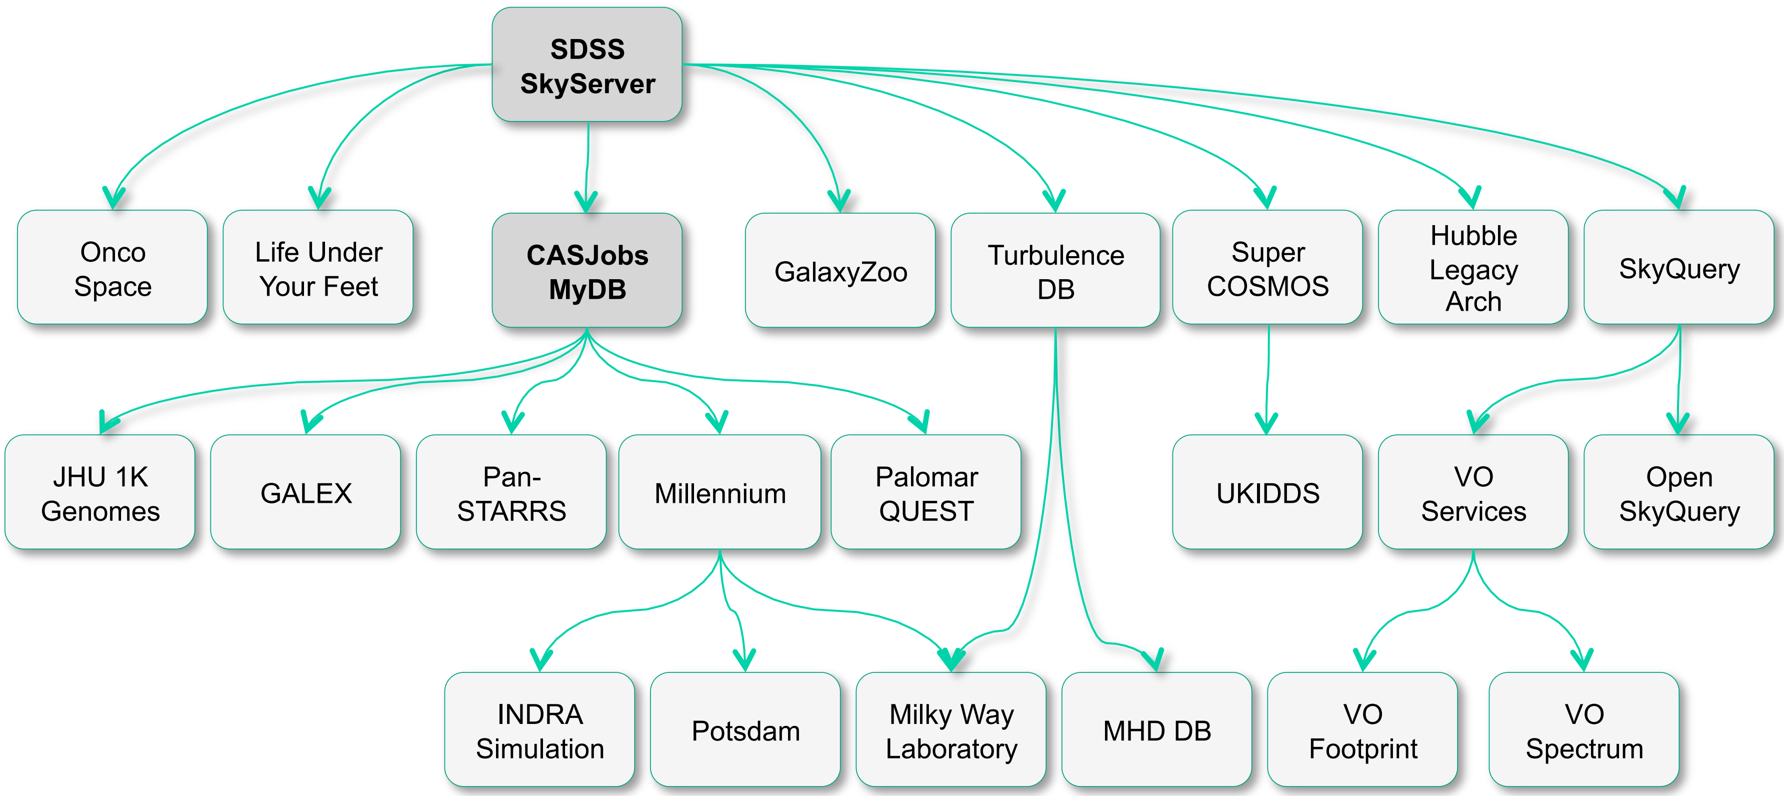

Virtual Observatory

- Started with NSF ITR project, “Building the Framework for the National Virtual Observatory”, collaboration of 20 groups
  - *Astronomy data centers*
  - *National observatories*
  - *Supercomputer centers*
  - *University departments*
  - *Computer science/information technology specialists*
- Similar projects now in 15 countries world-wide

⇒ International Virtual Observatory Alliance

NSF+NASA=>

VO Services

- Simple services to find data resources (VORegistry)
- SIAP - Simple Image Access Protocol
- TAP – Table Access Protocol
- VOTable
- VOTheory – Simulations
- VOFootprint – Sky Footprints
- VOSpectrum
- ….

Virtual Observatory Challenges

- Most challenges are sociological, not technical
- **Trust**: scientists want trustworthy, calibrated data with occasional access to low-level raw data
- Career rewards for young people still not there
- Threshold for publishing data is still too high
- Robust applications are hard to build (factor of 3…)
- Archives (and data) on all scales, all over the world

- Astronomy has successfully passed the first hurdles… but it is a long journey… no instant gratification

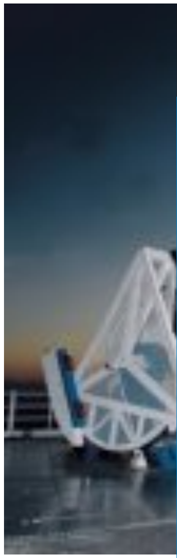

SD
2.4

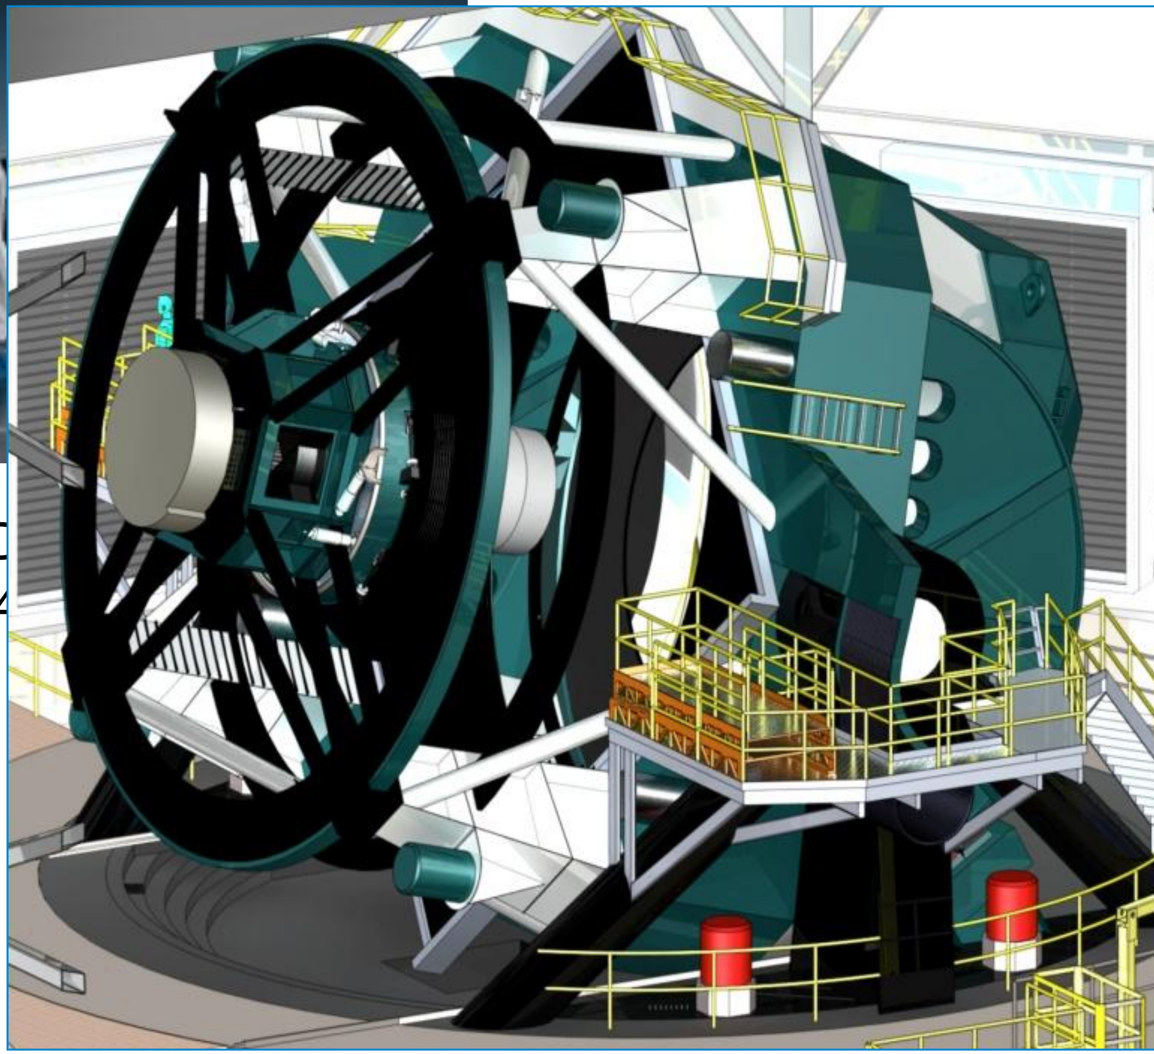

LSST
8.4m 3.2Gpixel

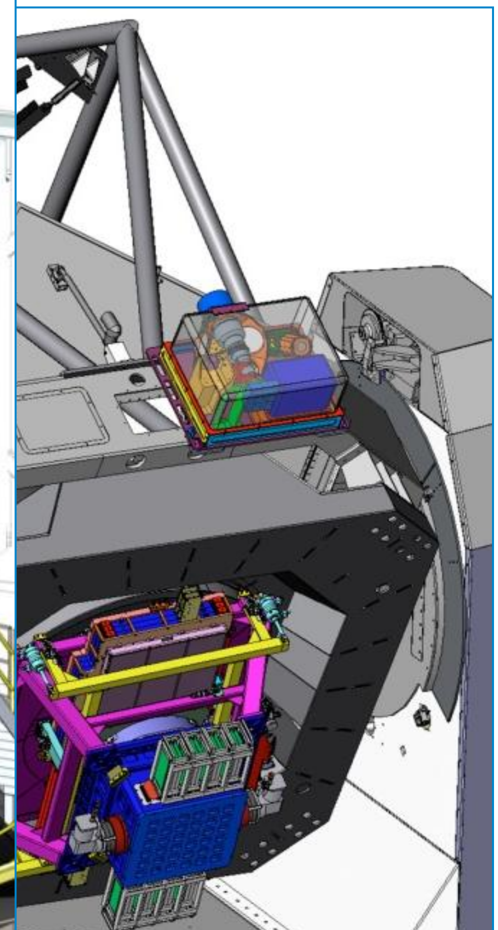

PanSTARRS
1.8m 1.4Gpixel

Continuing Growth

How long does the data growth continue?

- High end always linear
- Exponential comes from technology + economics
  - *rapidly changing generations*
  - *like CCD's replacing plates, and become ever cheaper*
- How many generations of instruments are left?
- Are there new growth areas emerging?
- **Software is becoming a new kind of instrument**
  - *Value added data*
  - *Hierarchical data replication*
  - ***Large and complex simulations***

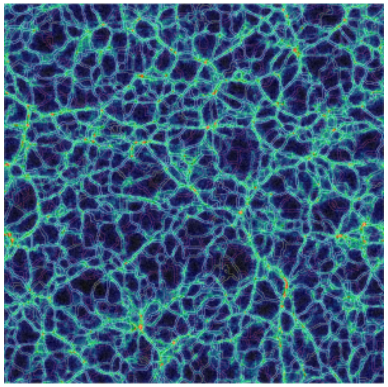

Cosmological Simulations

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

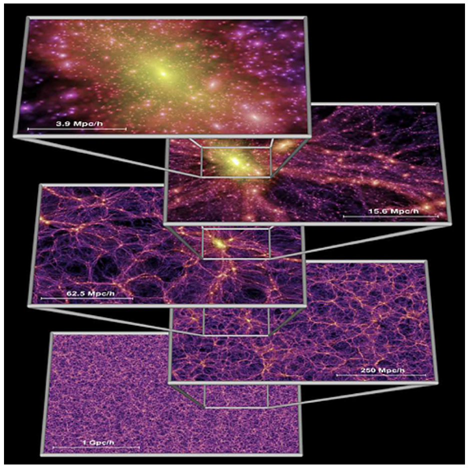

- Today: simulations with 10^{12} particles and PB of output are under way (MillenniumXXL, Silver River, etc)
- Hard to analyze the data afterwards -> need DB
- What is the best way to compare to real data?

Non-Incremental Changes

- Science is moving from hypothesis-driven to data-driven discoveries
- Need new randomized, incremental algorithms
  - *Best result in 1 min, 1 hour, 1 day, 1 week*
- New computational tools and strategies

… not just statistics, not just computer science,
not just astronomy…

Astronomy has always been data-driven….
now becoming more generally accepted