Fermi and LSST Computing

R.Dubois
richard@slac.stanford.edu

Outline

- Intro to Fermi & Pass 8

- LSST
  - Camera Control System
  - Test Data Curation and Analysis
  - LSST Project Data Management
  - Dark Energy Science Collaboration (DESC)

Fermi Highlights and Discoveries

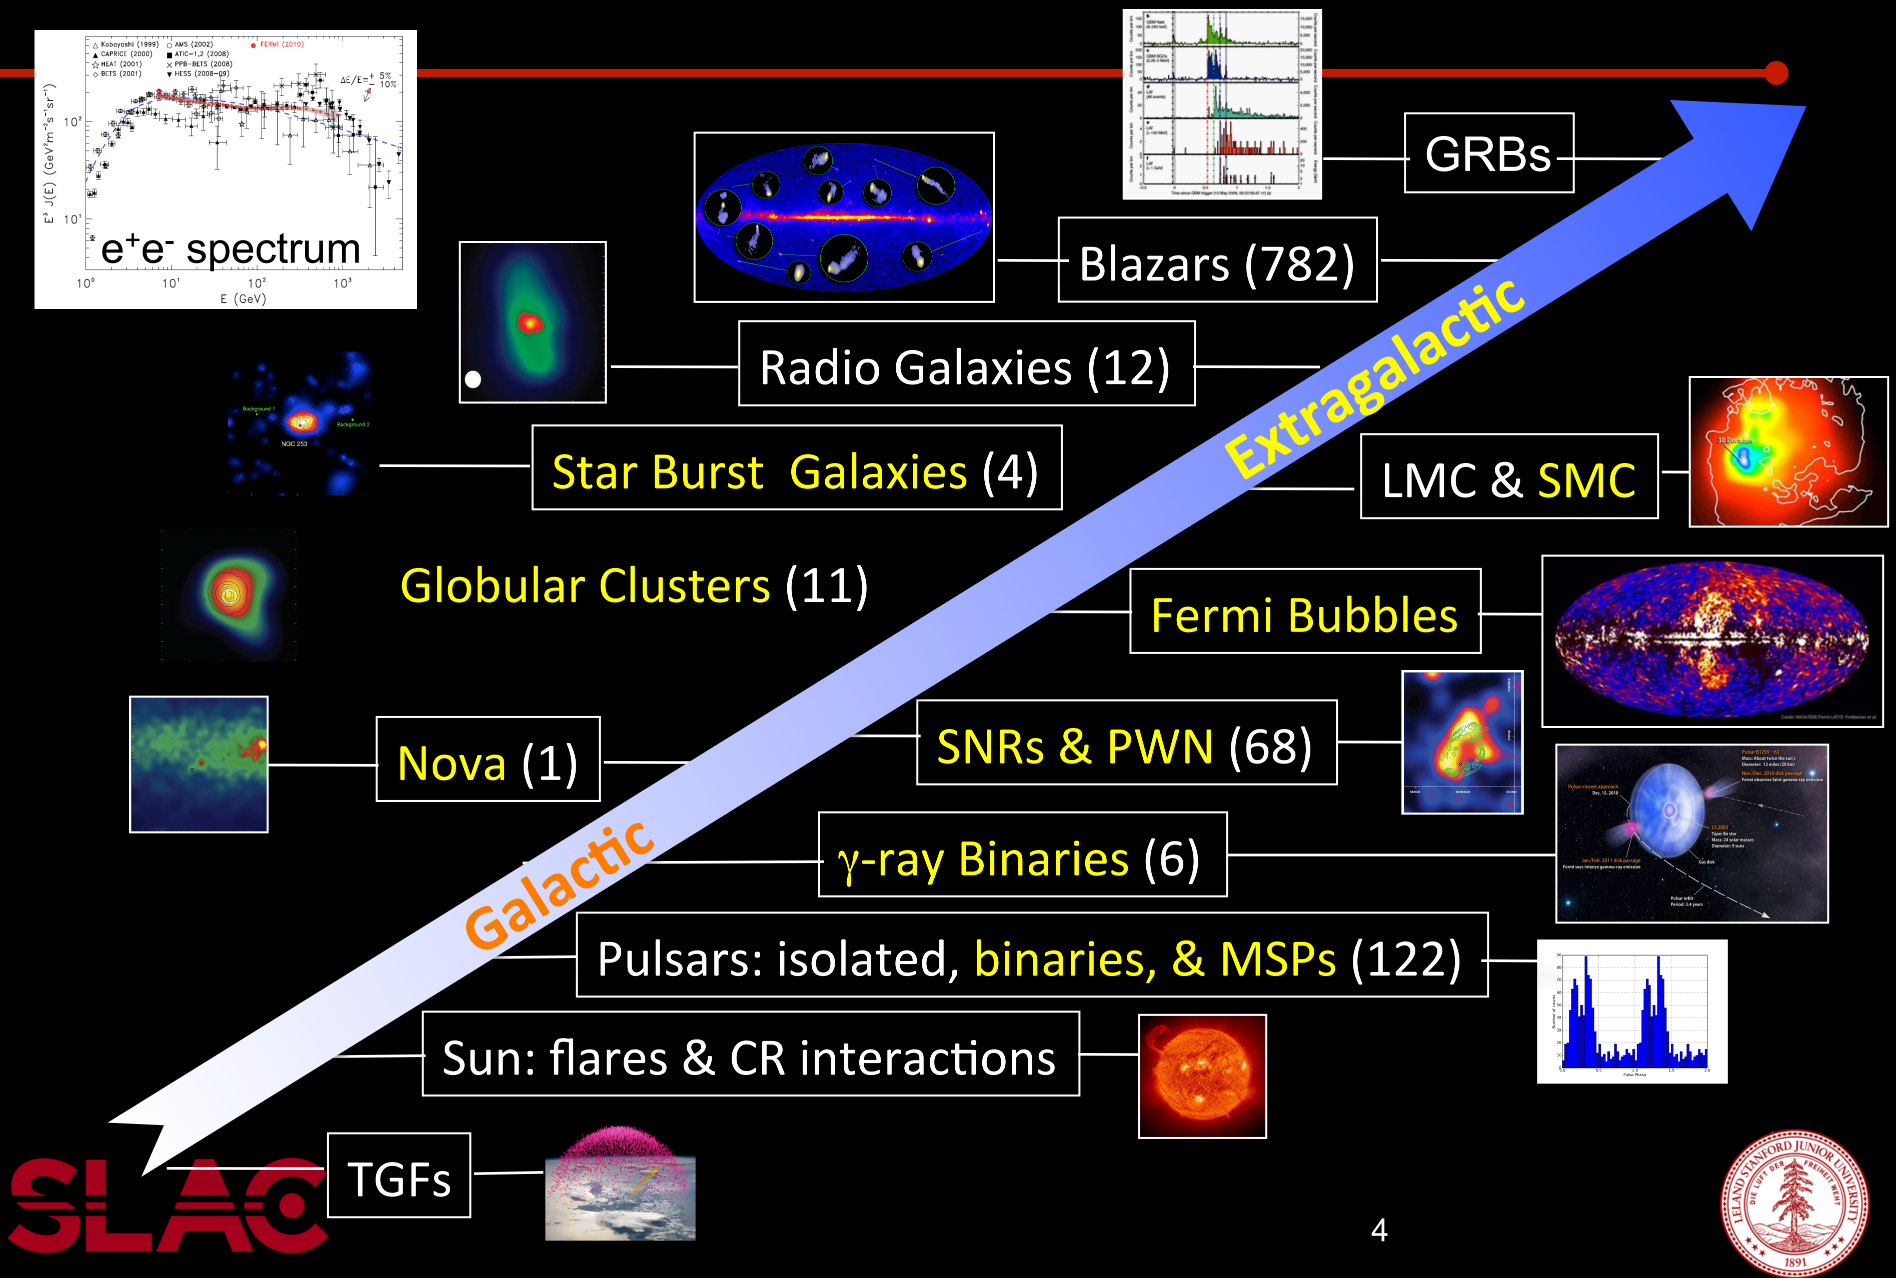

Fermi: Background Info

- Concept originated at SLAC & instrument integrated here
- Launched 2008-06
  - LAT data delivered to SLAC for processing. Turnaround time to be sent to public (via Goddard) is typically 90 mins.
- Mission goal is 10 years
  - NASA Senior review in 2012 secured ops through 2016
  - Next review in 2014 for “final” two years
  - We hope to go past 2018
- SLAC responsibilities: the LAT would not operate without SLAC
  - LAT Operations: Mission Planning, Flight software, LAT uploads
    - Maintenance of LAT testbed
  - Prompt Data processing: 90min turnaround to public
    - raw data through to photons delivered to public (via Goddard)
    - Data quality monitoring from processing output
  - Data reprocessing
  - Collaboration support
    - Data access to LAT collaboration (more than is made public)
    - Sim/recon support for Pass 8
    - High level science tools
    - Supporting collaboration analysis on SLAC farm
    - Offline computing management, infrastructure (builds, release management, dev environment)

Pass 8 : Key Motivations

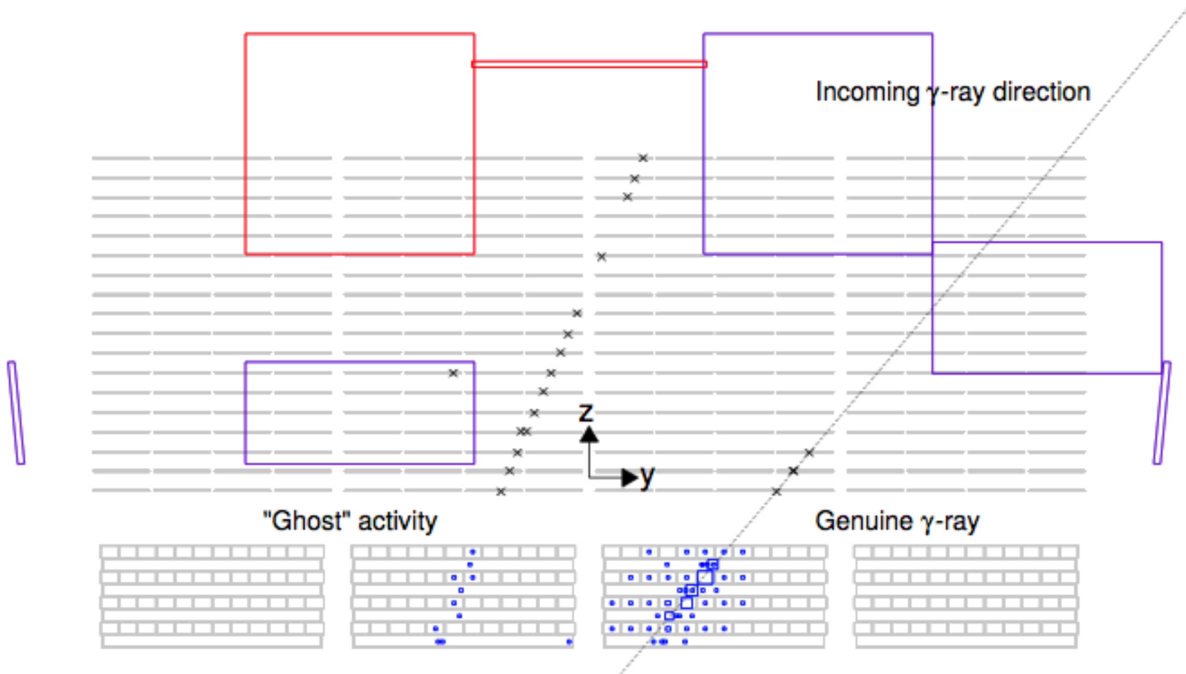

- The LAT event-level analysis was largely developed before launch.
- Considerable insight gained over the prime phase of the mission:
  - e.g., instrumental pile up is one of the original and main motivations for starting the Pass 8 development.
- Use this insight to maximize the instrument performance for science analysis.

Where we are now

Differential Point-Source Sensitivity (TS = 25; t=10^8 s; N > 10; b=90 deg)

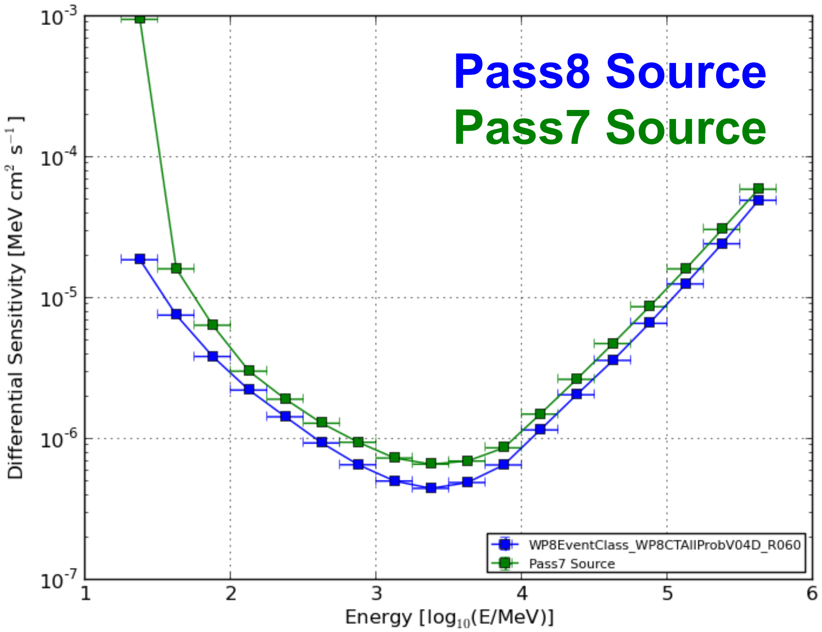

Ratio

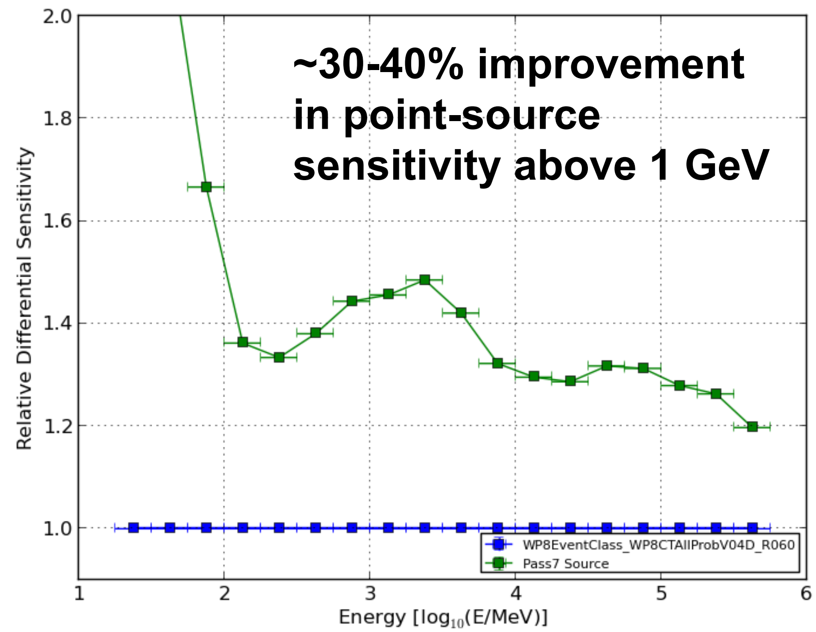

Pass8 demonstrates a clear performance improvement over Pass7!

Pass8 Timeline

- Recon Features Code Freeze Late March 2013
- GR Bug Fixes and Infrastructure Improvements March-April 2013
- Start of Pass8 Reprocessing Mid-April 2013
- Internal release of first Pass8 data sets (1+ yr) June 2013
- Pass8 Reprocessing Complete Sep/Oct 2013
- Final Validation Studies Oct-Dec 2013
- Operational Challenges:
  - Output file size has grown x3!
  - Execution time x2 slower

Contractually obligated to deliver 5 year Catalogue to NASA

Will be done using Pass 8.

LSST: A Deep, Wide, Fast, Optical Sky Survey

8.4m telescope 18000+ deg^2 10mas astrom. r<24.5 (<27.5@10yr)

ugrizy 0.5-1% photometry

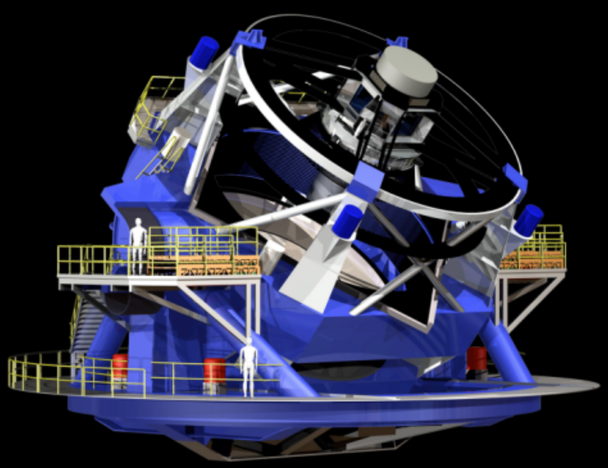

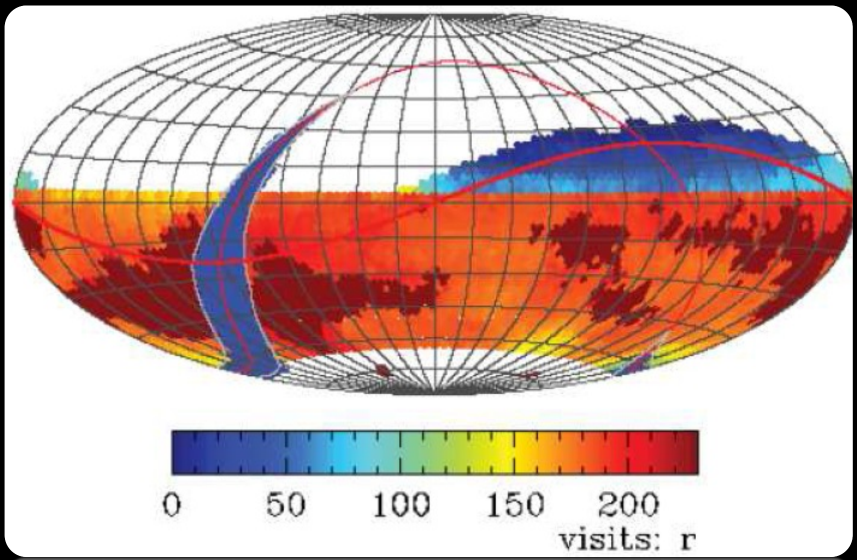

3.2Gpix camera 30sec exp/4sec rd 15TB/night 38 B objects

Imaging the visible sky, once every 3 days, for 10 years (825 revisits)

Three Aspects of LSST Program at SLAC

I. The LSST Camera Fabrication Project

- This is a major construction project: project management, design and development of major camera subsystems, and I&T of the entire camera assembly.
- SCA computing contributions to Camera Control System and Camera Test Data Curation and Analysis: 4 FTE on core funding.

II. The LSST Data Management Project

- SLAC selected because of the technical expertise present in the Laboratory. It is currently 50% externally supported, and 50% covered by core funding (1.5 FTE).
- SLAC has leadership in the database construction and management, and in various aspects of the middleware.

III. The LSST Dark Energy Science Collaboration

- DOE-OHEP will support the derivation of constraints on the nature of dark energy.
- We have formed the DESC to coordinate the planning and preparation, with SLAC as host Lab
- SLAC coordinating computing model and infrastructure; contribution to image simulation tools: 1 FTE on core funding.

LSST Camera Control System (CCS)

- Goals
  - Provide the distributed control and monitoring system for the camera
    - Provide framework for test benches used during acceptance testing and integration and testing (I&T)
    - Provide infrastructure for subsystem specific control functionality development
    - Provide Camera interface to the observatory control system (OCS) and once the camera is installed in the observatory

- CCS is in use now and will be continuously in future
  - Need to refine/iterate infrastructure design while supporting test benches/ subsystem developers
  - Need to work closely with subsystems (including test/database group) to understand requirements and provide support

Test Stand Consoles

- Refrigeration. This has been in continuous use at SLAC for several months.

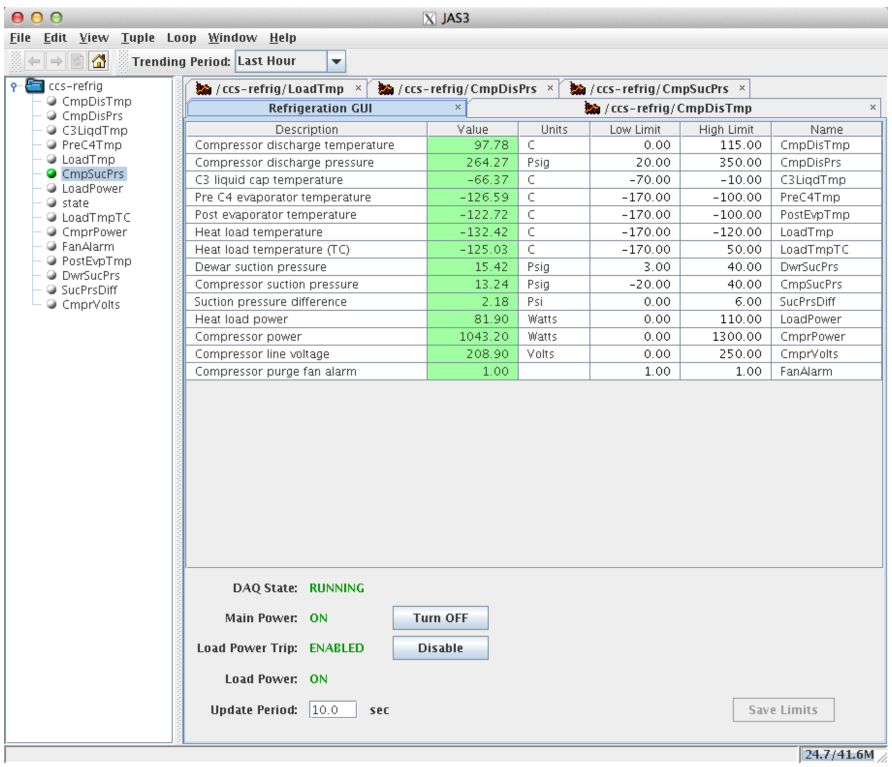

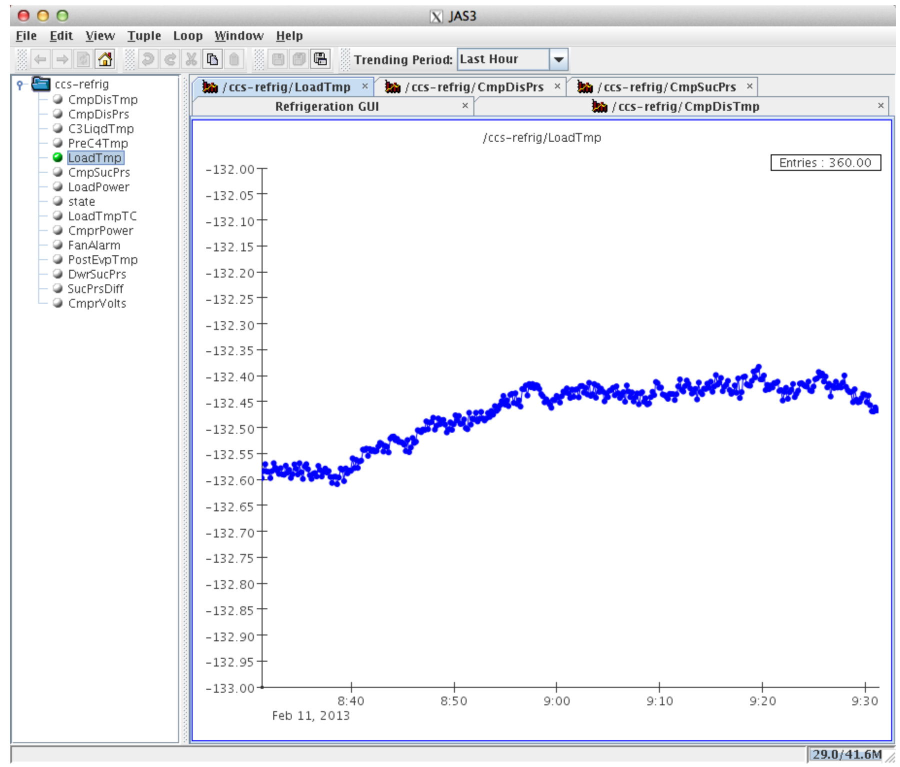

- CCD control (Brookhaven) and Filter Changer (Marseille) undergoing initial testing

SLAC NATIONAL ACCELERATOR LABORATORY

Manpower and Expertise

- Distributed core team at SLAC and in Paris
  - Paris: presently ~2.5 FTE
  - SLAC: presently ~3.5 FTE (2 FTE from SCA)
    - Expect manpower requirements to peak in ~2015

- Building on tools and experience of SCA applications group
  - Java expertise/support for collaborative development
    - Similar tools developed for BaBar, Fermi, EXO
  - Reuse of some existing SCA libraries/tools
    - Plotting package, trending (time-history) system, web interface, Java console framework
      - Risk: Aging and lacking clear long-term core support

Camera Test Data Curation and Analysis

- eManufacturing
  - Traveler
  - Procedures
  - Log Book
  - Datasheets
  - Calibration and test reports
  - Non-Compliance Reports
- Camera Data Management
  - Data Products
  - Data archive & Access tools
  - Algorithms
  - WorkFlow engine for bulk analysis

Goal

- **Develop integrated system for all:**
  - **Hardware types**
  - **Locations**
  - **Phases – test, production, operations**

Operations Overview: Database, Web Interface, Job Harness, eLog

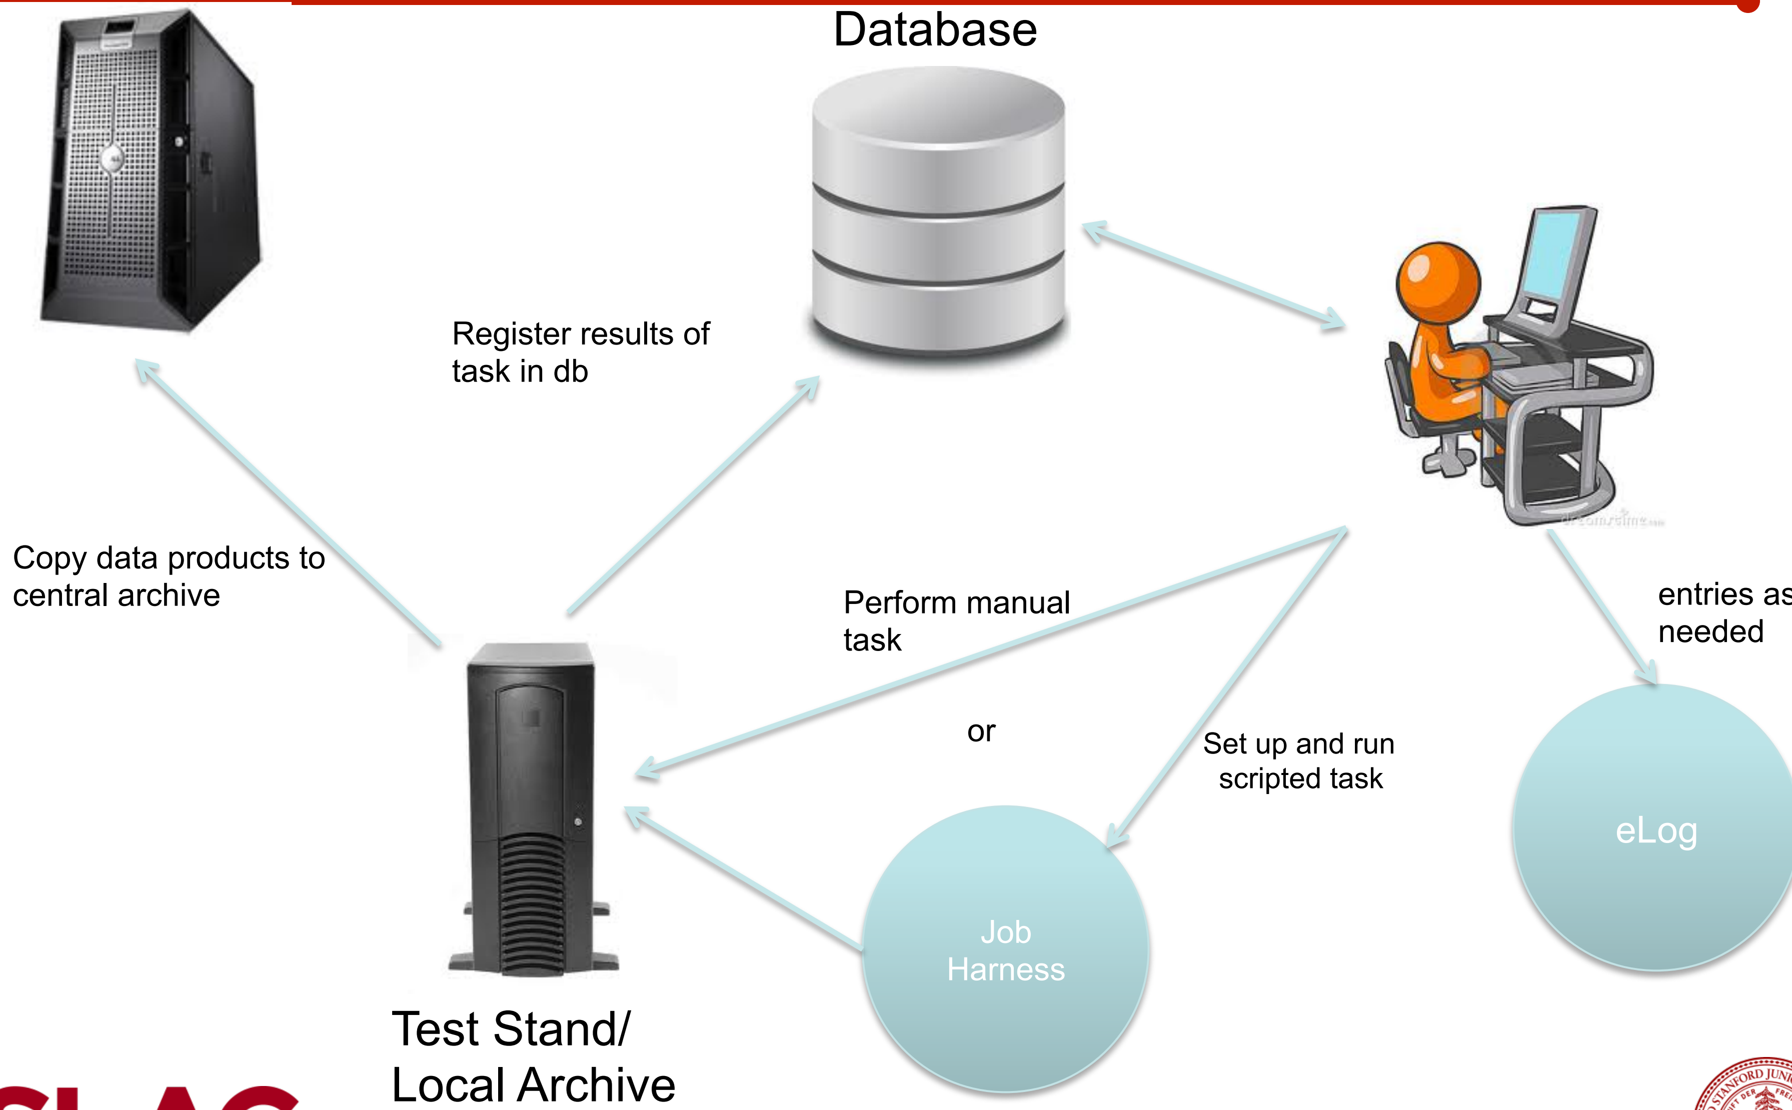

Building on SCA tools

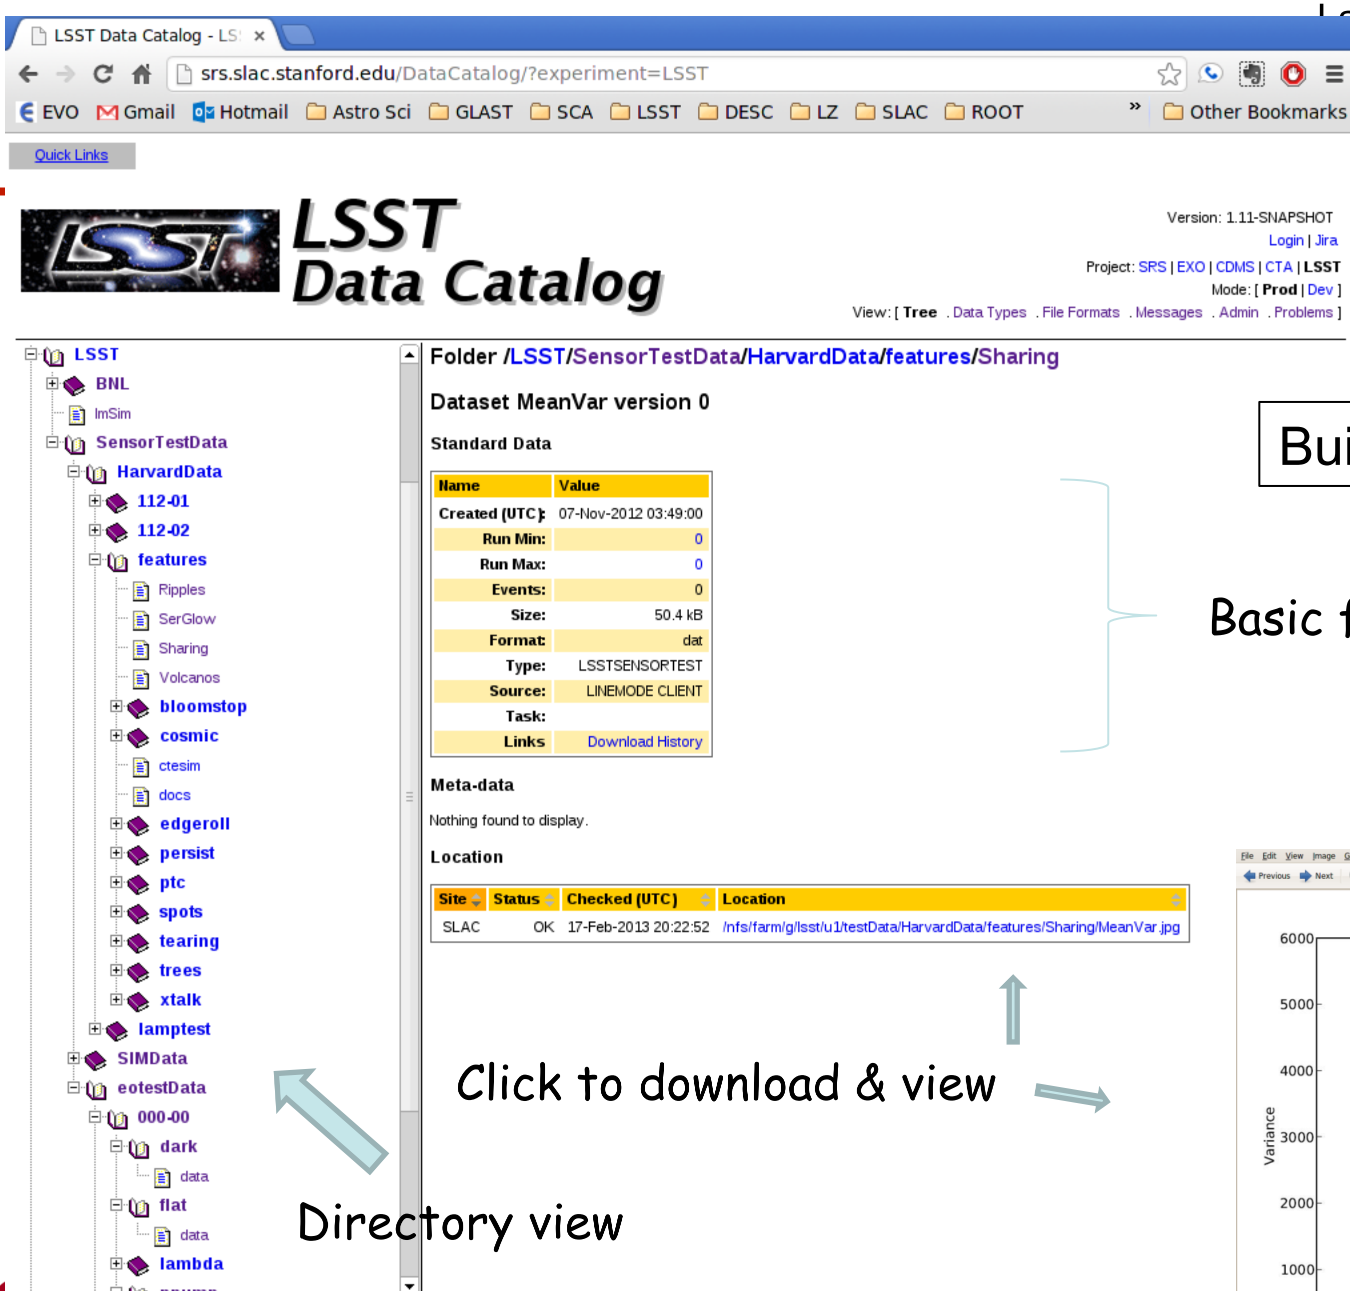

Basic file information

Click to download & view

Directory view

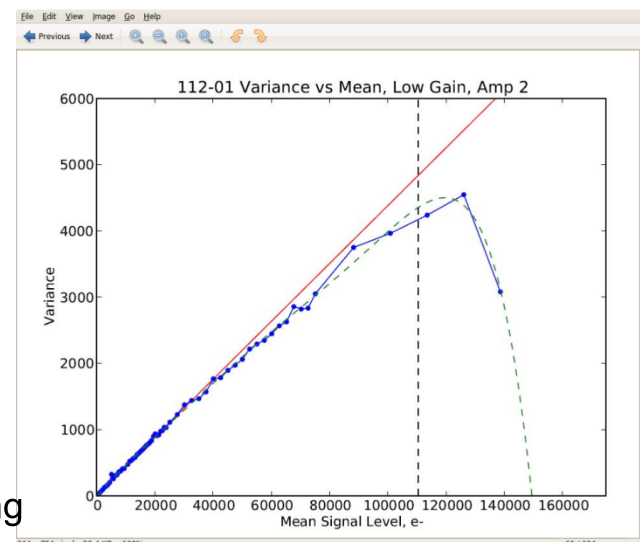

http://srs.slac.stanford.edu/DataCatalog/?experiment=LSST

NSF: Data Management Activities

- LSST DM delivered its 8th consecutive data challenge release in January, 2013
  - Twice-yearly release cycle, identical to construction/commissioning phase
- The latest release processed 10.4 TB in ~25 days (50,000 CPU hours) on the XSEDE Gordon cluster at SDSC:
  - 298 SDSS Stripe 82 runs (2 million fields)
  - Created a deep co-add covering -40 deg < R.A. < 55 deg, -1.25 < Dec < 1.25 (237 deg2)
  - Co-adds were used to detect 14.7 million sources, most of which would otherwise fall below the faint limit of individual exposures
  - Initial data quality analysis indicates that the software is already achieving data quality equal to or surpassing current fully operational surveys (e.g. SDSS. Pan-STARRS)
- The plan for achieving SRD-level requirements in construction/ commissioning is defined for all critical areas except one
  - Exception is Object Characterization, which is the focus of the Summer 2013 data challenge release

DM Data Releases and Algorithm Development in 2012

- Demonstrated single-frame extended source photometry (model fitting)
- Demonstrated co-addition by making a deep co-add of SDSS Stripe 82
- Demonstrated point-source forced-photometry
- Prototyped the deblender
- Opening up the codebase
  - Source installer
  - Binary installer
- High interest from scientists and EPO

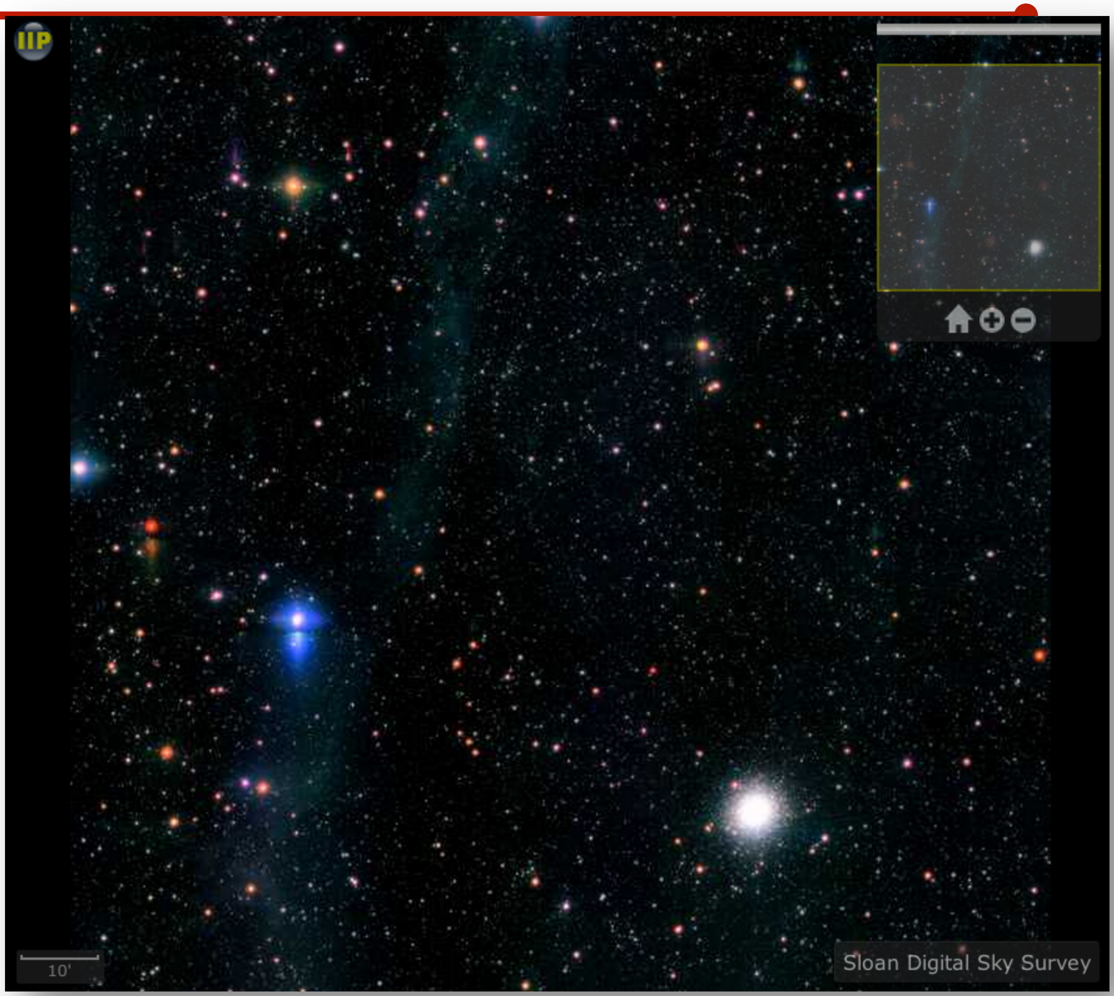

Background-matched co-add of SDSS Stripe 82 in the vicinity of M2

Broad uses for LSST data storage techniques

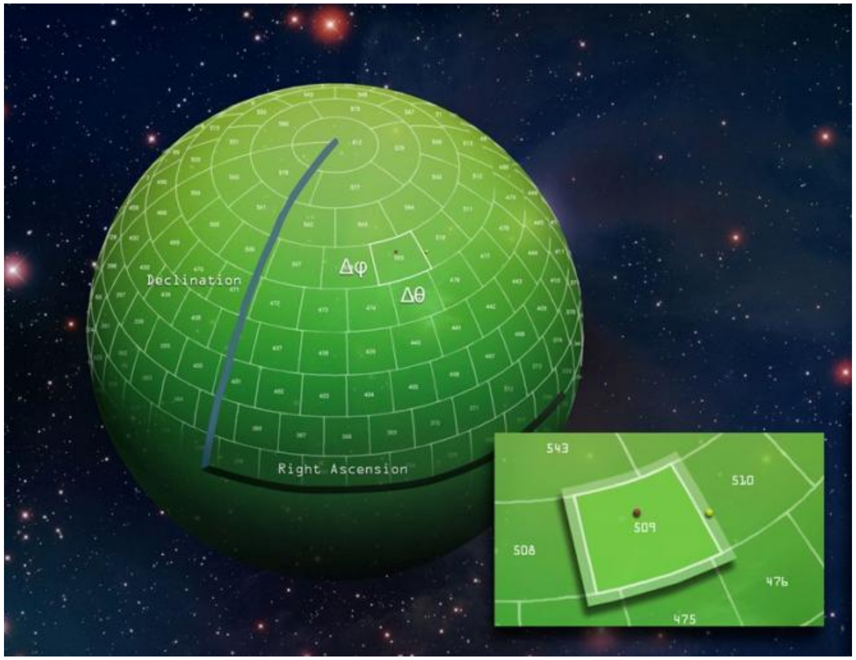

- Spherical partitioning with overlapping edges developed by LSST database team for efficient searching of enormous databases
- Technique shown to be linearly scalable without degrading system performance
- Useful in many fields that store spatial information (maps) and information that changes with time
  - Financial sector
  - Geosciences; Climate modeling
  - Fraud detection; internet usage behavior
  - Medical imaging; Drug discovery
  - Oil and gas exploration
- Featured as Research.gov highlight

Dark Energy Science Collaboration

- Collaboration formed June 2012
- Steve Kahn & Bhuv Jain co-spokesmen
- White Paper written: arXiv:1211.0310
- Very interested in Curation of camera test data for simulations realism and for understanding systematics during Operations

- Analysis Working Groups *Jeff Newman*
  1. Weak Lensing — Michael Jarvis, Rachel Mandelbaum
  2. Large Scale Structure — Eric Gawiser, Shirley Ho
  3. Supernovae — Alex Kim, Michael Wood-Vasey
  4. Clusters — Steve Allen, Ian Dell'Antonio
  5. Strong Lensing — Phil Marshall
  6. Combined Probes, Theory — Rachel Bean, Hu Zhan
  7. Photo-z Calibration — Jeff Newman (acting)
  8. Analysis-Computing Liaison — Rick Kessler
- Computing and Simulation Working Groups *Andy Connolly*
  1. Cosmological Simulations — Katrin Heitmann
  2. Photon Simulator — John Peterson
  3. Computing Infrastructure — Richard Dubois
  4. Software — Scott Dodelson
- Technical Working Groups *Chris Stubbs*
  1. System Throughput — Andrew Rasmussen
  2. Image Processing Algorithms — Robert Lupton
  3. Image Quality — Chuck Claver
  4. Science Operations and Calibration — Zeljko Ivezic

The image simulation workflow

- A representative view of the universe and its properties
- Reflects the current LSST designs
- Scales from the individual images to Terabytes
- Fidelity of the simulations is designed to match the current needs of the project
- Flexible to respond to as-delivered components
- Images complement the capabilities of the analytic models, catalog simulations and real data

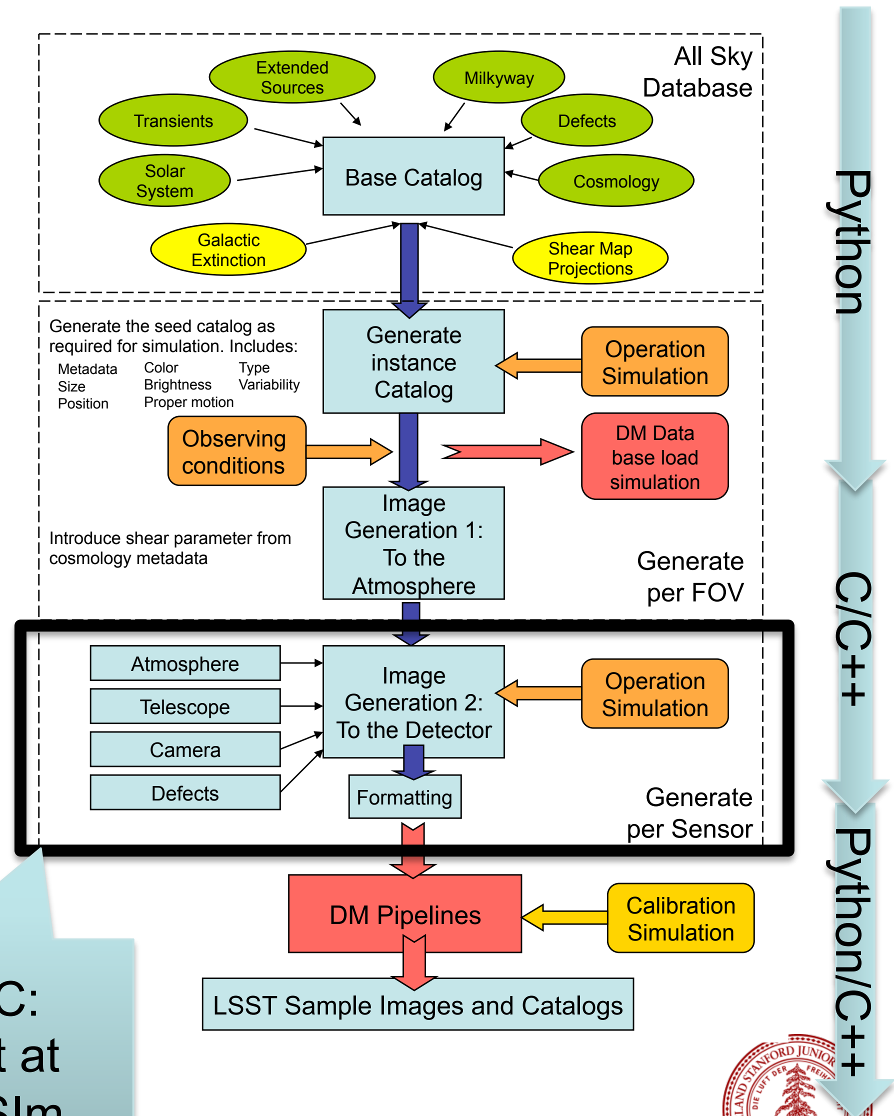

Key to DESC: specific effort at SLAC on ImSIm

NATIONAL ACCELERATOR LABORATORY

ImSim Computational Flow

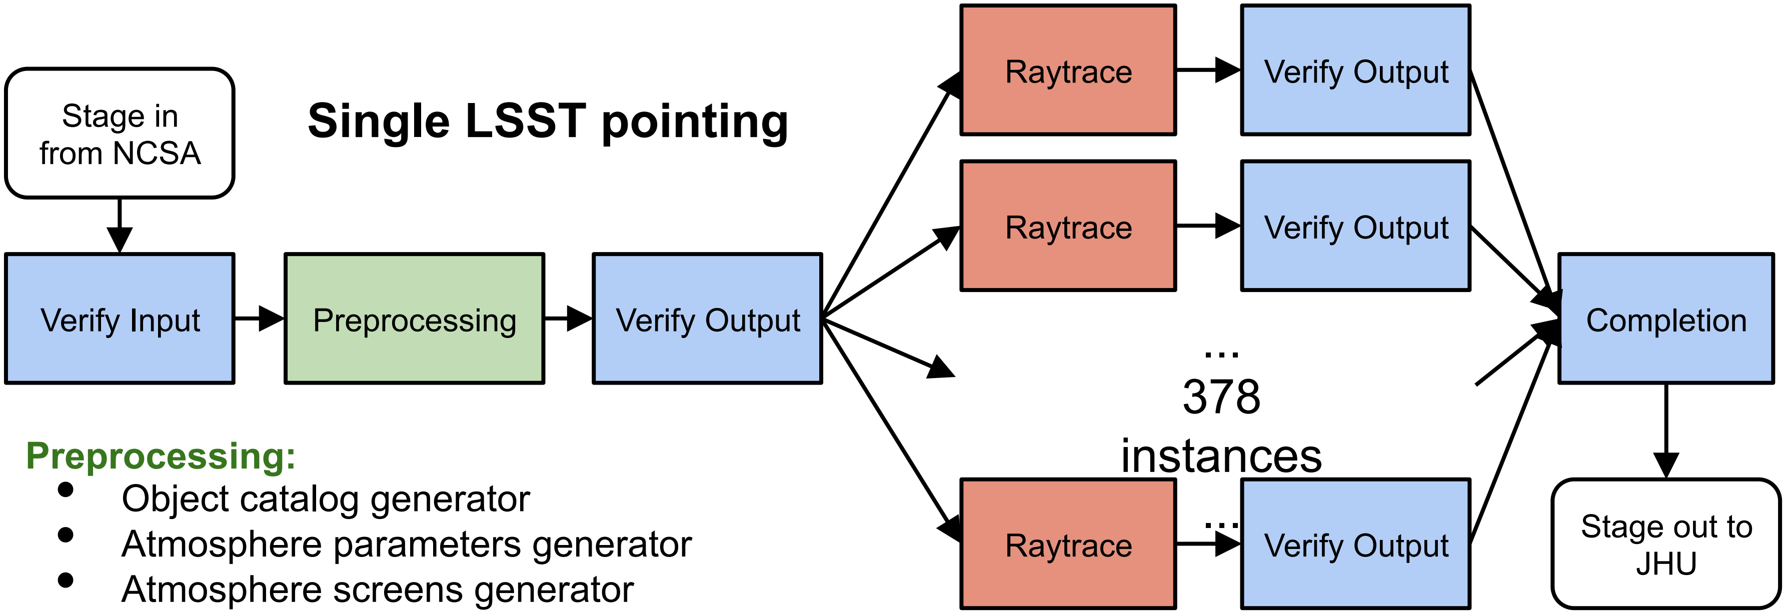

Preprocessing:

- Object catalog generator
- Atmosphere parameters generator
- Atmosphere screens generator
- Cloud screens generator
- Telescope optics parameters

Raytracing:

- Photon raytracer
- Cosmic ray adder
- Background radiation adder
- CCD readout modeler

- 1000 CPU-hrs per pointing
- Each CCD and exposure can be treated independently
- Therefore, we have *378-way parallelism per pointing.*

SLAC NATIONAL ACCELERATOR LABORATORY

DESC Computing Directions

- Need for data intensive computing
  - many runs of about 1000-10000 cores with fast IO and machines with at least 1GB per core – worry about future big scale machines!
  - 2.5 M-CPU-hrs allocation at NERSC for 2013
    - Work still needed to be able to run our code effectively on MPI machine
    - likely want to double our allocation on NERSC each year for the next 5 years
  - Working with Torre Wenaus on interfacing workflow engine to Panda
- Provision disk space to support our big runs
  - 20-40 TB but likely doubling with the compute resources
  - Distribution/access to collaborators
- Need DM stack installed (and phosim etc) on NERSC and kept current
- Planning workshop at FNAL in June to develop use cases with Science Working Groups
  - Better understand modes of working
  - Start scoping out the framework for DESC computing and analysis
  - Computing Model proposal due at October collaboration meeting