NERSC Analytics Program Status and Update

DOE CGF

April 30, 2007

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

Today's Presentation

- Describe goals, components, resources, and value of NERSC Analytics Program
- Supported science domains
- Synergy with other analytics programs
- Examples of current technology applications – from scientific data management to data exploration

Goals of Analytics Program

Facilitate NERSC user knowledge discovery through the

use, adaptation, extension, creation, application, and deployment

of a diverse array of technologies spanning the domains of

high performance computing, data management, data analysis and visualization, and workflow management.

Components of Analytics Program

- **Scientific data management (SDM) – efficient access to data; data organization and security**
- **Workflow management – systematic approach to data processing & simulation pipelines**
- **Data analysis/data mining – broad category includes tools to processing data, finding features**
- **Visualization – visual data exploration; tools for communicating and sharing results**

Integration of these components facilitates data exploration that leads to knowledge discovery.

Resources of the Analytics Program

- Team of seven (three new hires late 2006-early 2007) with experience spanning all aspects of analytics, high performance computing, and many science domains.
- SGI Altix – 32 processors, 192GB RAM, 40TB attached FC storage
  - Architectural balance favors data intensive operations: large SMP memory, best I/O bandwidth on the floor at NERSC.
- Procurement process underway for new analytics machine.

Value of the Analytics Program

- For NERSC users, problem sizes are growing larger with larger computational platforms and higher-resolution instruments.
- All sciences need to store, access, process, explore and understand data and information.
  - Science is increasingly data intensive.
- Understanding and managing data is a limiting factor for a wide range of sciences (ASCR SDM Challenge workshop report).
  - There is no science without data analysis.

How We Work with Users

- **Generally, no off-the-shelf, general-purpose solutions for analytics exist.**
- **The Analytics Program adapts, extends, integrates, and sometimes creates technologies to meet user needs.**
- **Consulting and collaborative projects with users in all aspects of analytics technologies.**
- **Results and impact on science come through in-depth work with stakeholder/users.**

Supported Science Areas

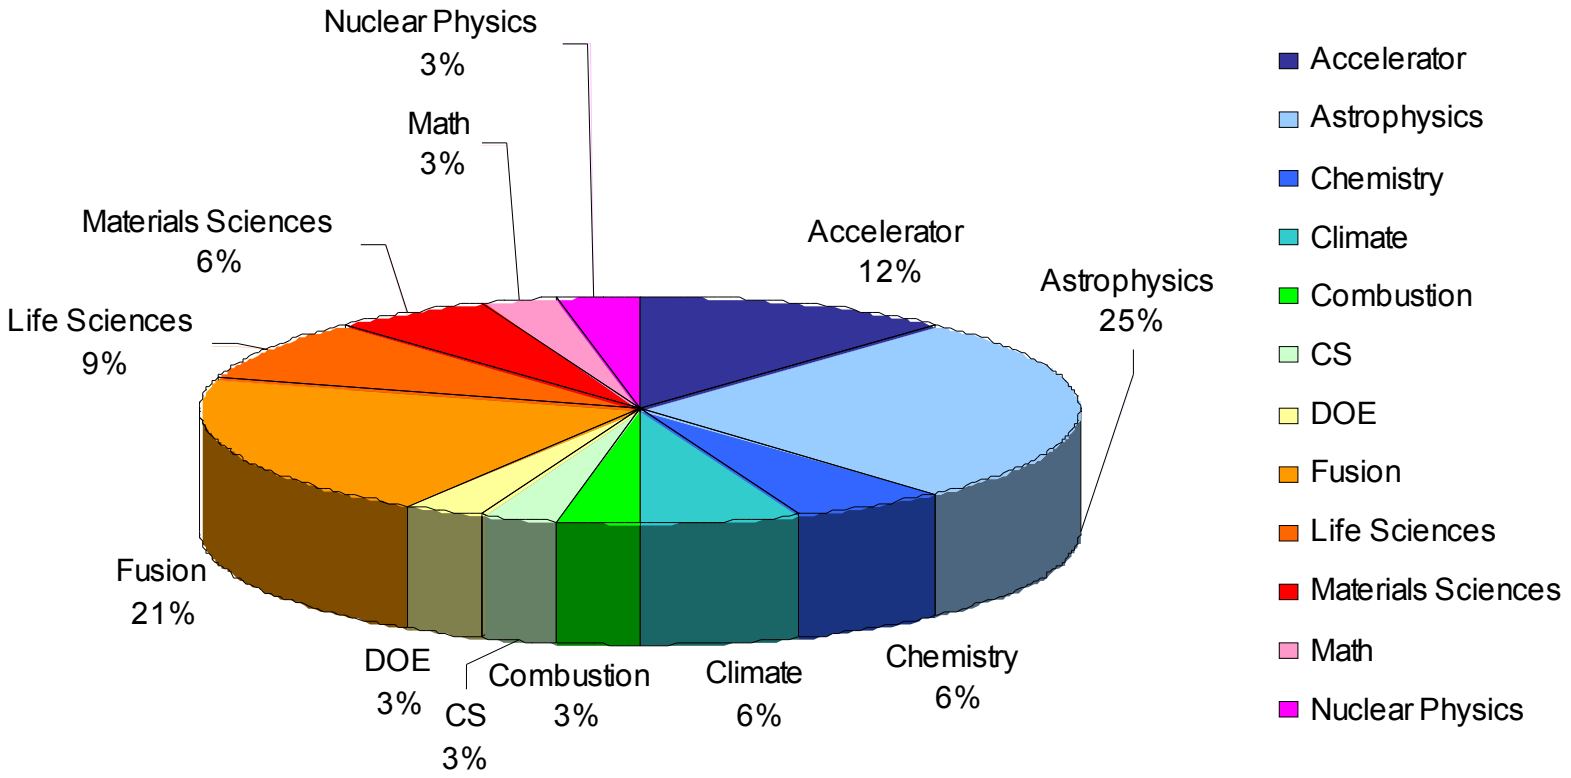

Supported Science Areas 2005-2007

Analytics Customer Technology Matrix

	SDM	Workflow	Analysis	Visualization
Accelerator	X	X	X	X
Astrophysics	X	X	X	X
Biology			X	X
Chemistry				X
Climate		X	X	X
Combustion			X	X
CS			X	X
Fusion	X		X	X
Math	X		X	X

Synergy with Other Analytics Programs

- SciDAC Visualization and Analytics Center for Enabling Technologies (VACET)
  - Collaboration of LBNL, LLNL, ORNL, Utah, UC Davis
  - Representatives from both NERSC Analytics and ORNL/CCS Visualization groups
  - Development of analytics infrastructure that will be deployed at DOE's open computing facilities
- DOE Base Program in Visualization
  - New technologies for high performance visualization
- SciDAC Scientific Data Management Center
  - FastBit index/query technology
  - Storage Resource Manger (SRM) technology
  - Expert assistance in applying technologies to user problems

Examples of Current Projects

- Scientific Data Management – Kurt Stockinger
  - SRM, H5Part, FastBit
- Workflow Management – Kurt Stockinger
  - Kepler
- Data Analysis and Data Mining – Peter Nugent, Raquel Romano
  - Climate modeling and astrophysics
- Visualization – Cristina Siegerist, Gunther Weber
  - Current projects and accelerating remote display performance
- Data Exploration – Cecilia Aragon
  - SUperNova Factory AssembLy Line

Scientific Data Management
Kurt Stockinger

- Storage Resource Manager (SRM) for distributed data management
  - integrated mechanism for transferring files from one location to another, uniform access to heterogeneous storage (disk, tape), fault tolerant
- H5Part is a simplified API on top of the scientific data format HDF5.
  - high-performance, parallel I/O
  - developed for accelerator modeling
- FastBit for efficient indexing and querying
- HDF5 FastQuery combines bitmap indices with HDF5 (no slides).

SRMs Used in Earth Science Grid for Robust Multi-File Replication

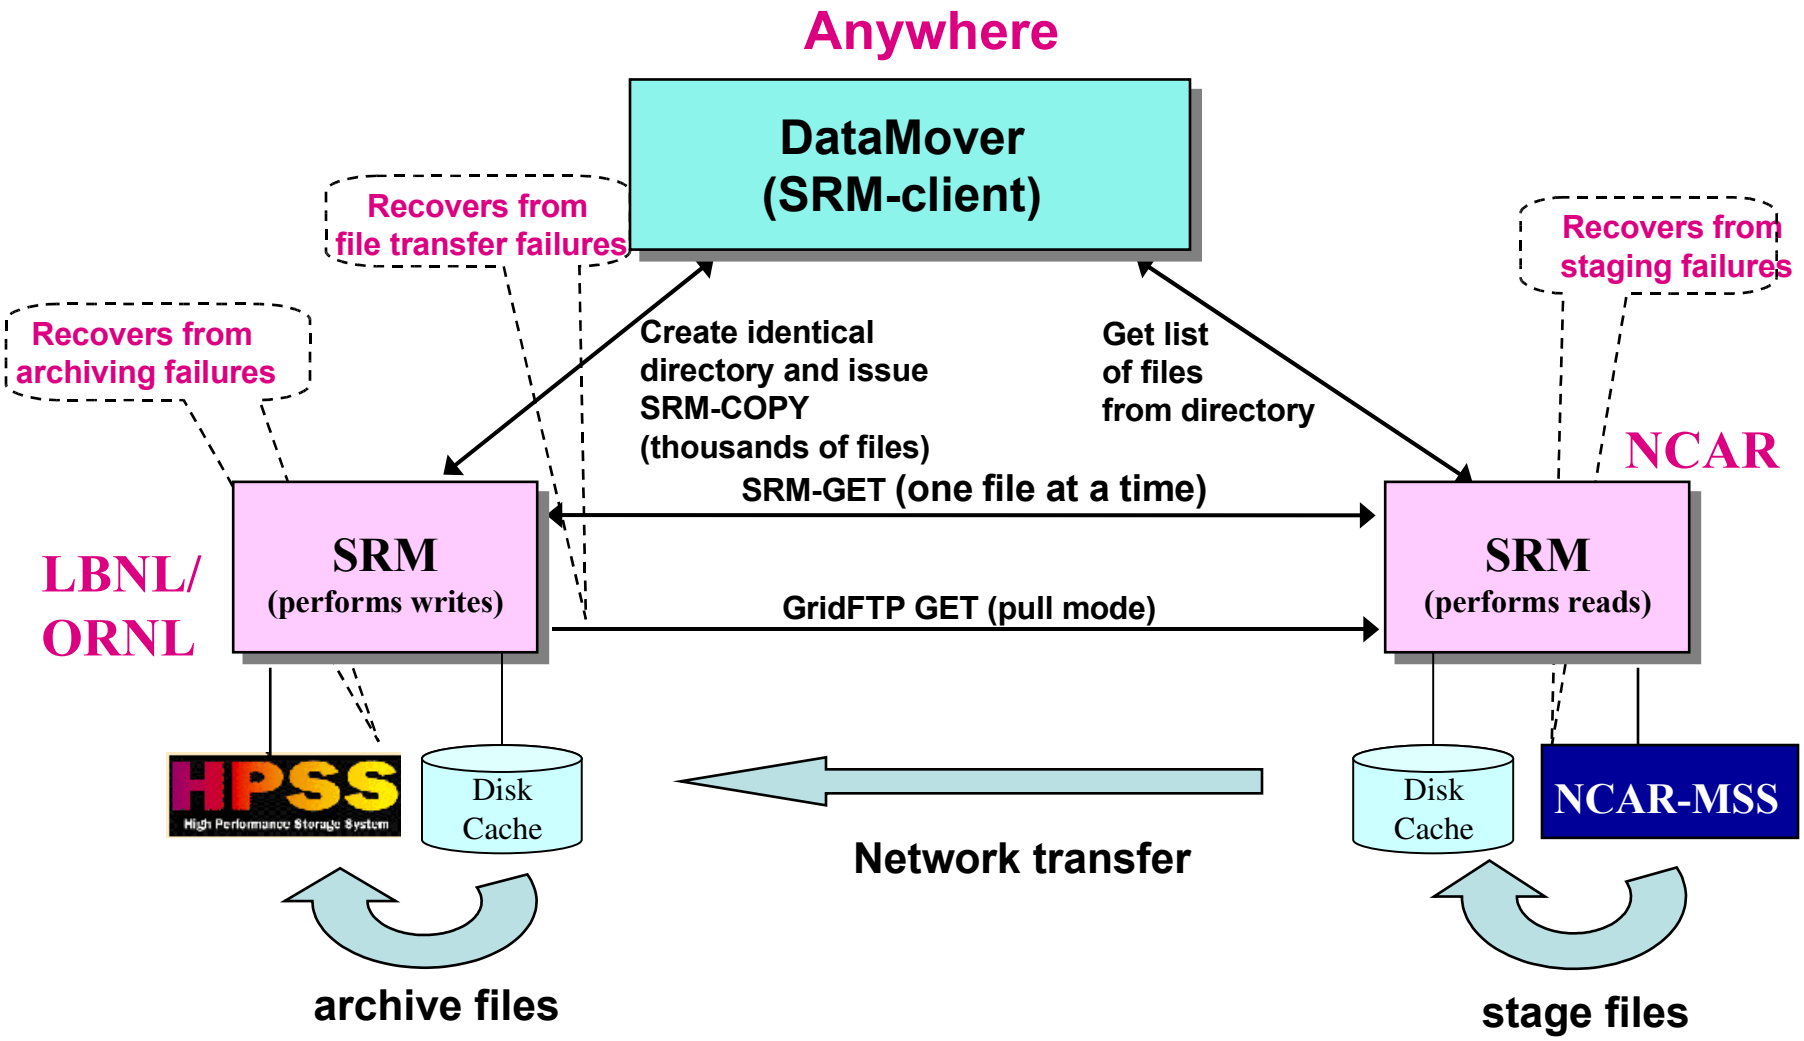

Courtesy: Arie Shoshani, Alex Sim (LBNL SDM Group)

Data Analysis with H5Part

H5Part

- provides high-performance parallel I/O
- integrated into codes:
  - BeamBeam3D, MaryLie/Impact (ongoing)
  - Used at LBNL and PSI (Switzerland)
- current focus on particles:
  - 6-vector of coordinates and momenta

VisIt plugin developed by Cristina Siegerist

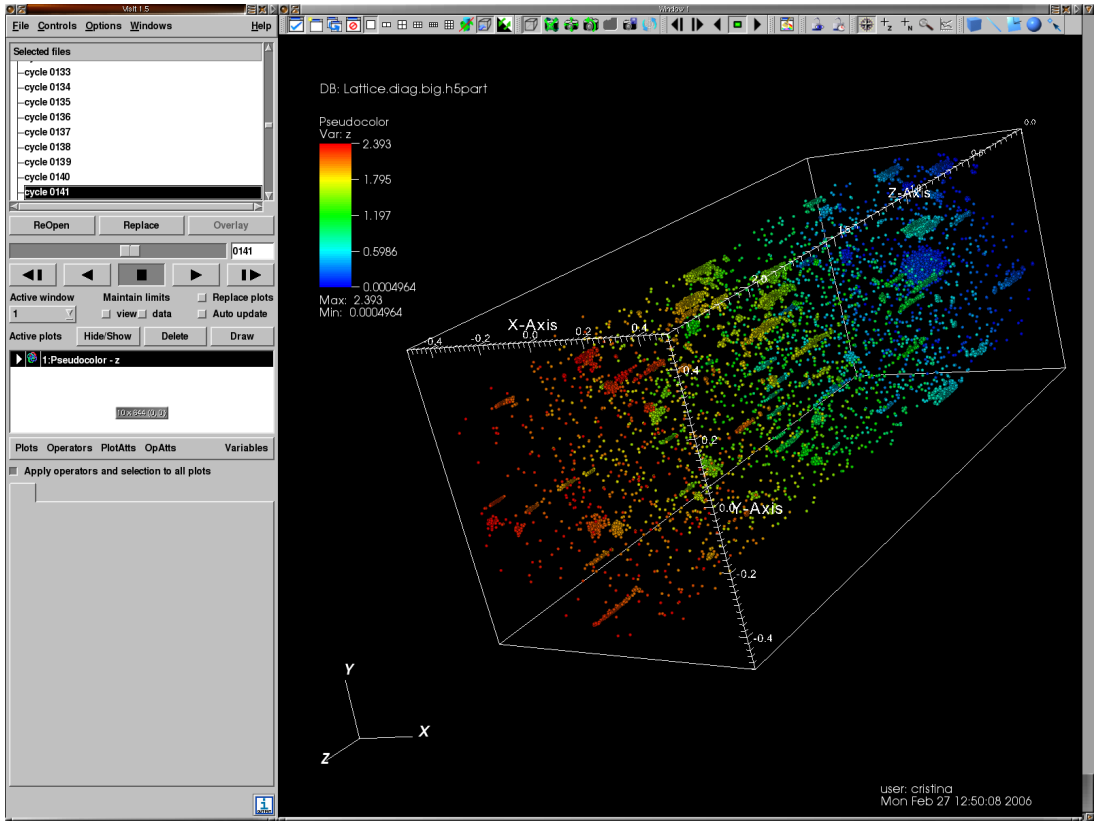

VisIt user interface showing an H5Part particle file rendered as spheres using the pseudocolor plot option. - visualization by Cristina Siegerist

Bitmap Indices in FastBit

base data bitmap index

		E^9	E^8	E^7	E^6	E^5	E^4	E^3	E^2	E^1	E^0
1	3	0	0	0	0	0	0	1	0	0	0
2	2	0	0	0	0	0	0	0	1	0	0
3	1	0	0	0	0	0	0	0	0	1	0
4	2	0	0	0	0	0	0	0	1	0	0
5	8	0	1	0	0	0	0	0	0	0	0
6	2	0	0	0	0	0	0	0	1	0	0
7	9	1	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0	1
9	7	0	0	1	0	0	0	0	0	0	0
10	5	0	0	0	0	1	0	0	0	0	0
11	6	0	0	0	1	0	0	0	0	0	0
12	4	0	0	0	0	0	1	0	0	0	0

Bitmap index compresses very well due to sparseness.
Queries can be answered very efficiently.

Workflow Management Using Kepler
Kurt Stockinger

Kepler (www.kepler-project.org)

- software application for building scientific workflows
- collaborative effort

Components consist of directors (scheduler), actors (perform a function), ports, relations, and parameters.

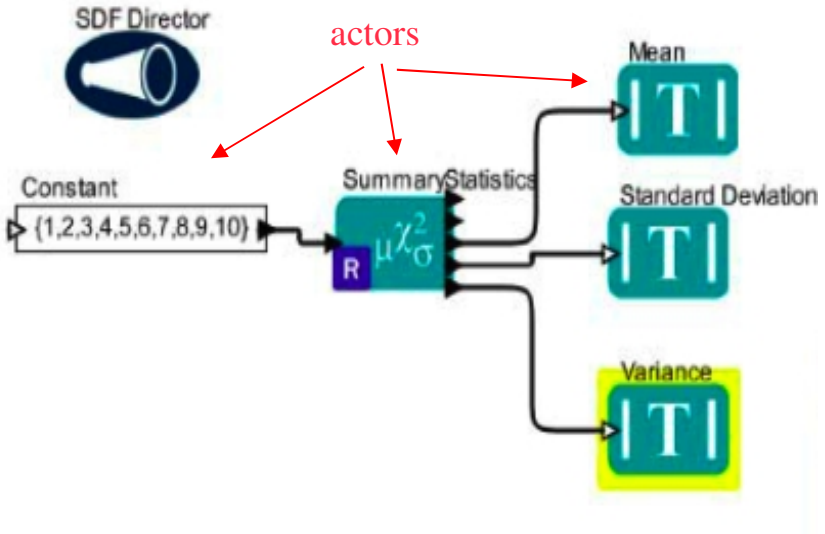

workflow to calculate simple statistics

Accelerator Simulation

- Iterate over input files in different directories
- Execute specific codes remotely
- Edit input files

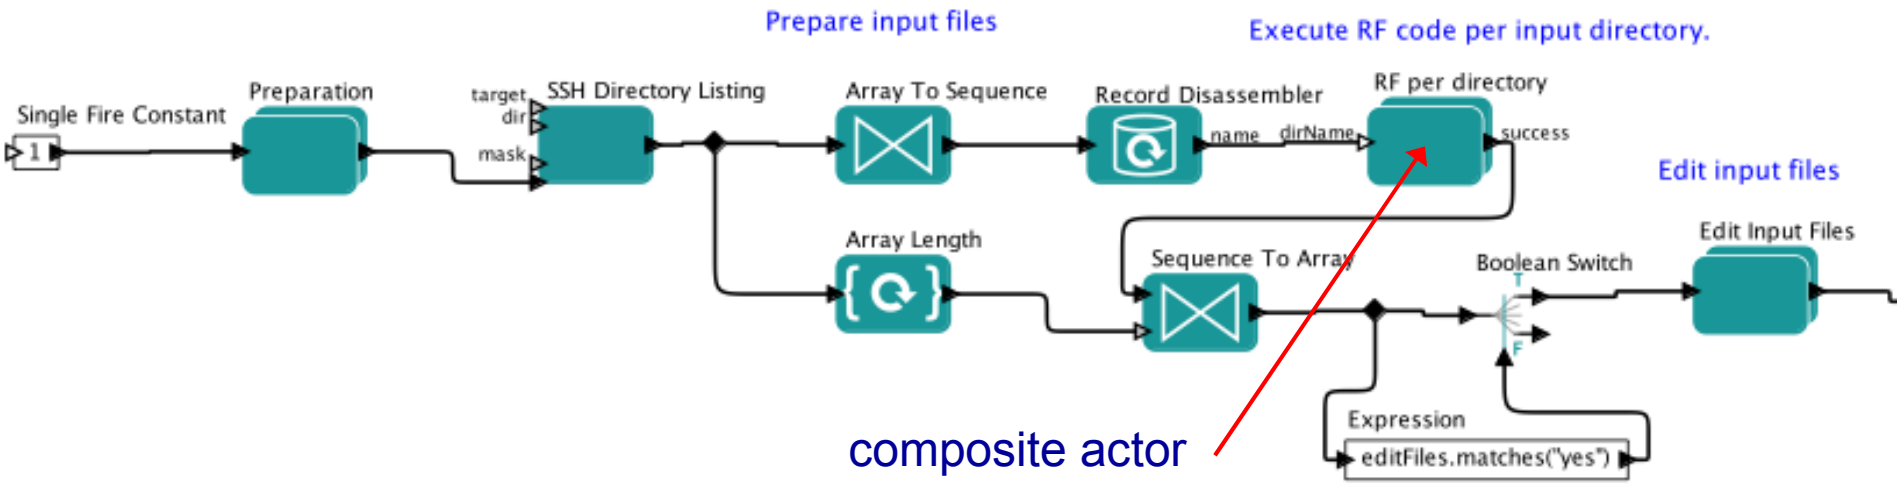

Benefits of Using Kepler Workflows

Script	Kepler Workflow
Needs to be written from scratch	Contains large number of re-usable actors
Mainly used for local execution	Allows parallel and distributed workflow execution
May be difficult to maintain	Easier to maintain
Often difficult to read ("write-only")	Visual representation of workflows: easier to understand

Analysis of Climate Modeling
Raquel Romano

Data analysis techniques used for feature detection and to remove 'model drift' from simulations.

- Blind source separation (BSS)
  - Can be used to find features (tropical storms) when no explicit criteria exist
- Multi-taper method
  - Used to separate model drift from natural variability

Automatic Feature Extraction by Blind Source Separation

- **Usual method of detecting tropical storms in GCM is to look for high wind, low SL pressure, high vorticity, *etc*.**
- **What if you don't know what explicit criteria to use to find features of interest?**
- **The BSS method (also known as blind deconvolution, independent component analysis (ICA)) is a statistical method for detecting unusual signals.**
- **Rachel applied the BSS method (using Matlab) to find features in climate models.**

Automatic Feature Extraction by Blind Source Separation

Extracted features can be used as templates for finding similar features.

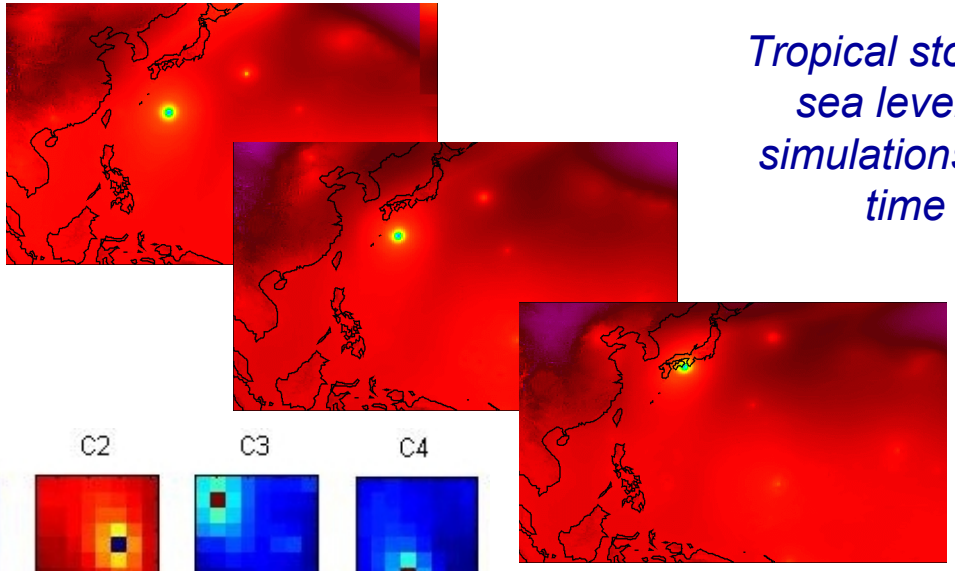

Tropical storm visible in sea level pressure simulations at multiple time steps.

In this case, the features were variations on rotating low-pressure systems. This was not assumed a priori.

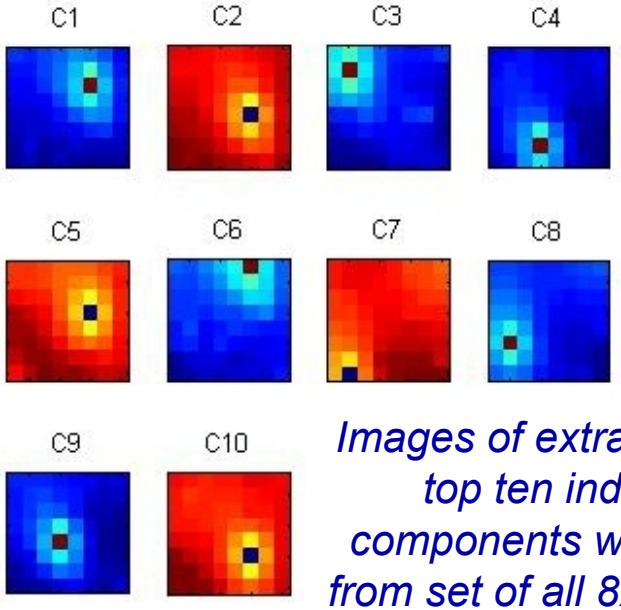

Images of extracted features: top ten independent components were extracted from set of all 8x8 subimages.

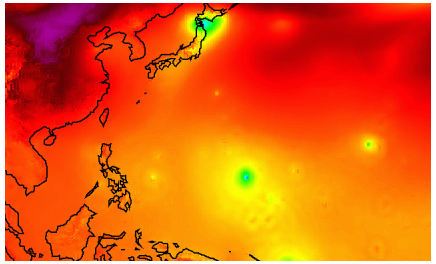

Nonlinear Trend Estimation in Long-Term Simulations of Surface Air Temperature

- Model-induced drift occurs in 1000-year surface air temperature simulations.
- Need to separate periodic oscillations due to natural variability from background trends (model drift).
- Multi-Taper Method (MTM): combines spectral frequency analysis with a statistical model.
- MTM used to find spatially varying nonlinear trends in order to detect spatio-temporal patterns of interest.

Statistical Detrending of Surface Air Temperature Using MTM

The red line is the trend that the modelers want wanted to remove. MTM was used to find trends both on a global scale, and on smaller spatial scales.

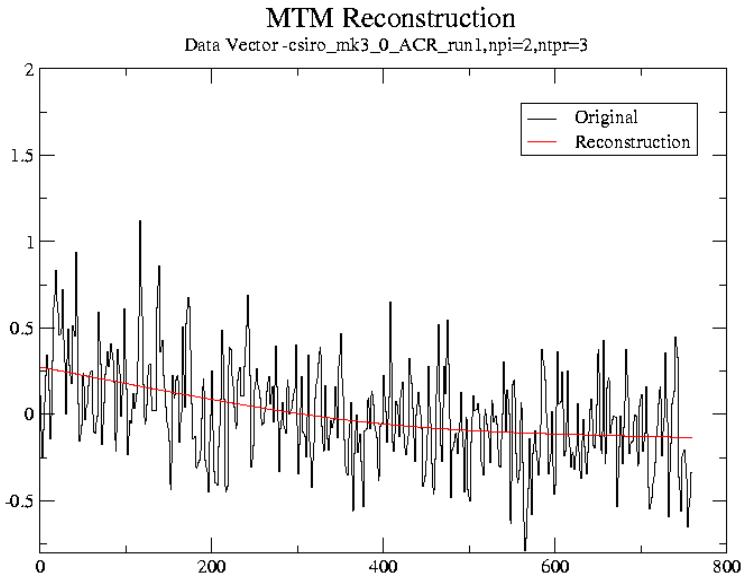

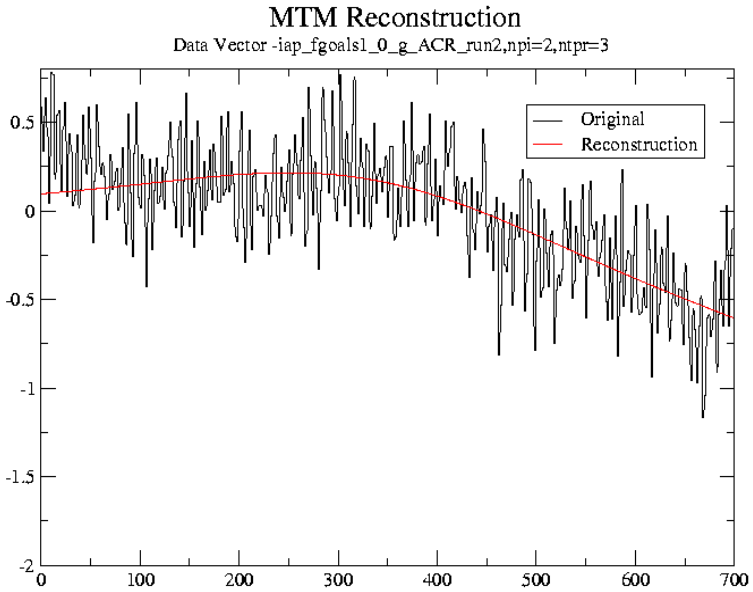

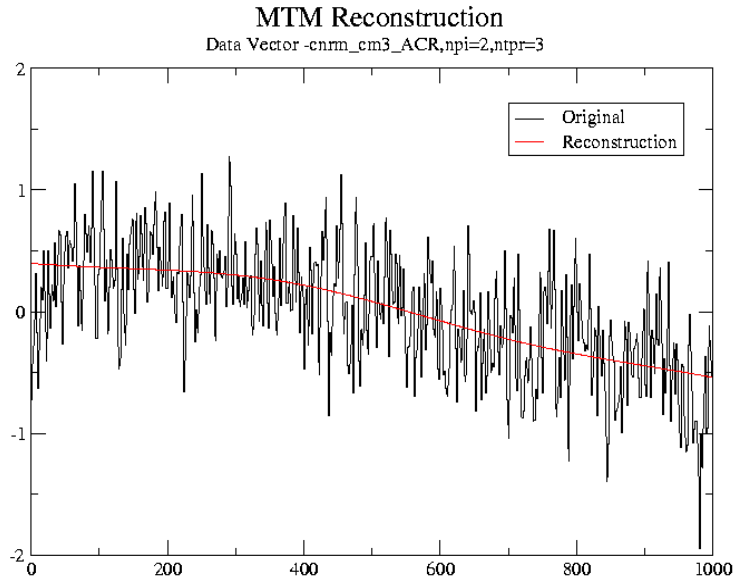

Analytics in Astrophysics
Peter Nugent

- Analysis of supernovae data led to the detection of the first Super Chandrasekhar Mass Type Ia supernova.
- Analysis of simulation results of supernova detonation facilitated comparisons with observations
- Workflow developed for processing 1 M 10 MB sky images to produce a library of reference images to use in investigating transient objects.

Confirmation of Type 1a SN from "Super-Chandrasekhar" Mass WD

- **Visualization and analysis used to test and confirm theory of new type of Type Ia SN – one having "Super-Chandrasekhar" mass.**
- **Top: brightness vs. velocity and amount of deceleration.**
- **Middle: velocity vs.mass and unburned carbon using a 1.4 solar mass model.**
- **Bottom: moving to a 2 solar mass model includes new observation.**

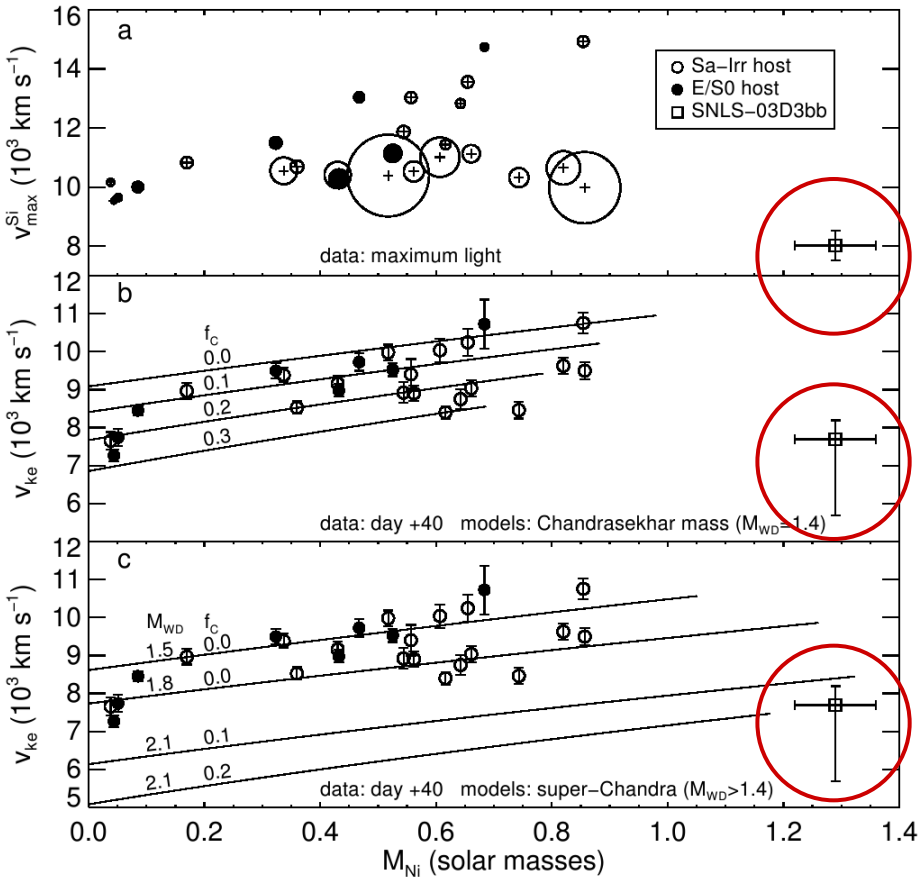

D.A. Howell, M. Sullivan, P.E. Nugent, R.S. Ellis, A.J. Conley, D. Le Borgne, R.G. Carlberg, J. Guy, D. Balam, S. Basa, D. Fouchez, I.M. Hook, E.Y. Hsiao, J.D. Neill, R. Pain, K.M. Perrett and C.J. Pritchet. The type Ia supernova SNLS-03D3bb from a super-Chandrasekher-mass white dwarf star. *Nature,* 443, 308-311, 21 September 2006.

Spectrum Synthesis

- The spectrum synthesis of the Incite model simulations compares favorably to observations.
- This may be the first explosion model that "naturally" explains the transition from deflagration to detonation in thermonuclear supernovae.

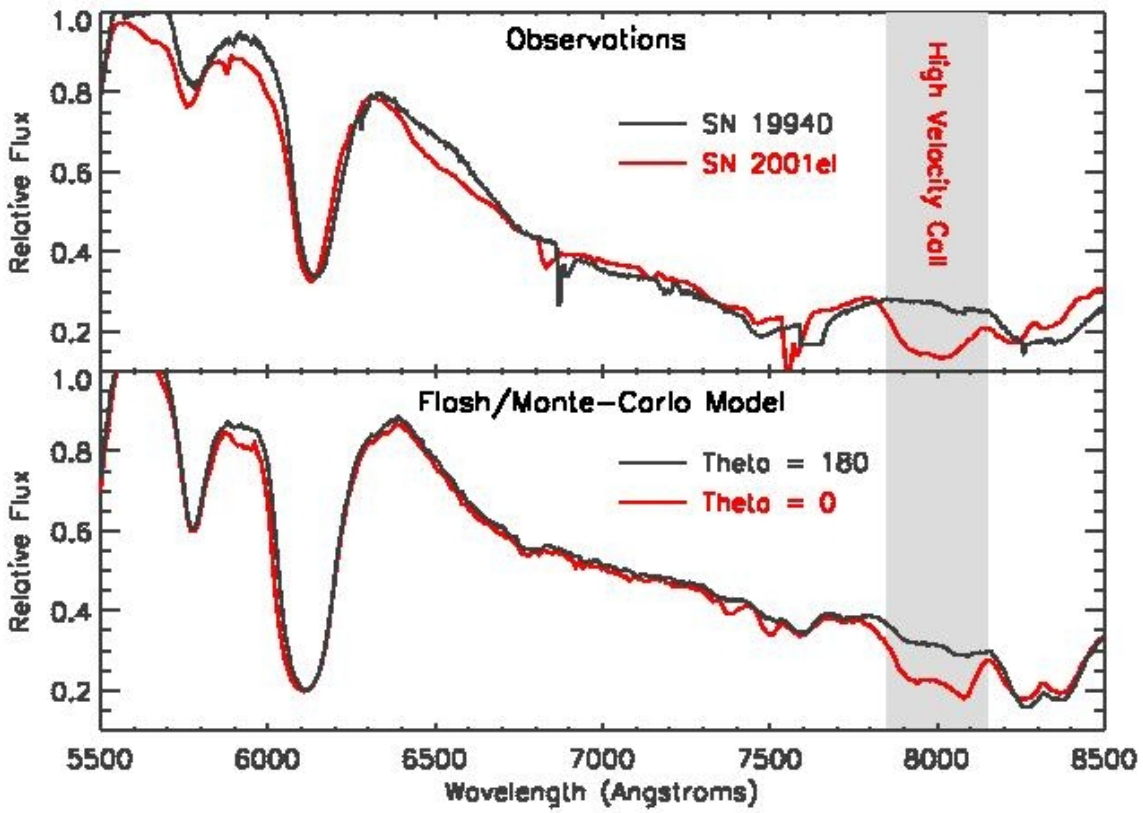

Workflow Management in the SN Factory

- Goal: make better reference images to use when investigating transient objects (*e.g.*, supernovae).
- About ~20,000 sq deg on the sky has been covered about 50 separate times: approximately 1 million, 10 Mb images resides on NGF.
- Images processed in 2 X 2 square degree chunks; half of the processing is a serial task and half is a parallel task.
- Serial task: Determine image "zero point" and astrometry.
- Parallel task: co-addition of 500-1000 images into one 250 Mb image of 2 X 2 square degrees.
- Each part takes roughly the same amount of time ~ 100 wallclock days.

Workflow Management in the SN Factory

- To handle the serial part in “parallel”, Peter wrote mpibatch, a script to perform a round-robin execution of a list of serial commands (or even scripts of several commands) through mpi.
- Other steps:
  - Perl script converts text list into Fortran namelist.
  - Fortran namelist read as character strings on the master processor, spit out to slaves to execute, who then ask for more when they are done.
  - The threaded code ‘swarp’ is run to co-add each of the sub-images into one large image.
- Example the Coma Cluster…

Workflow Management in SN Factory

500 images went into this incredibly deep co-addition of the Coma Cluster.

Visualization
Cristina Siegerist

- Particle-in-Cell Simulation of Laser Wakefield Particle Acceleration
- Accretion-Induced Collapse of White Dwarfs
- VisIt plugin for H5Part – shown earlier
- VisIt plugin for NIMROD output in HDF5
- Journal covers

PIC Simulation of Laser Wakefield Particle Acceleration

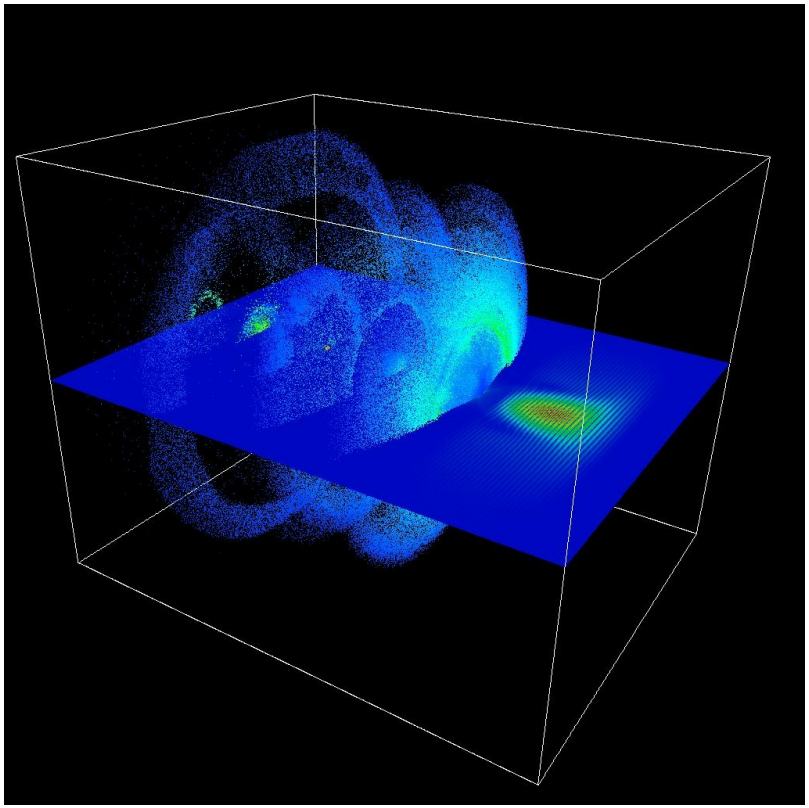

This image shows a horizontal slice through the electric field; the electrons are colored by the magnitude of the momentum. *(AVS/Express)*

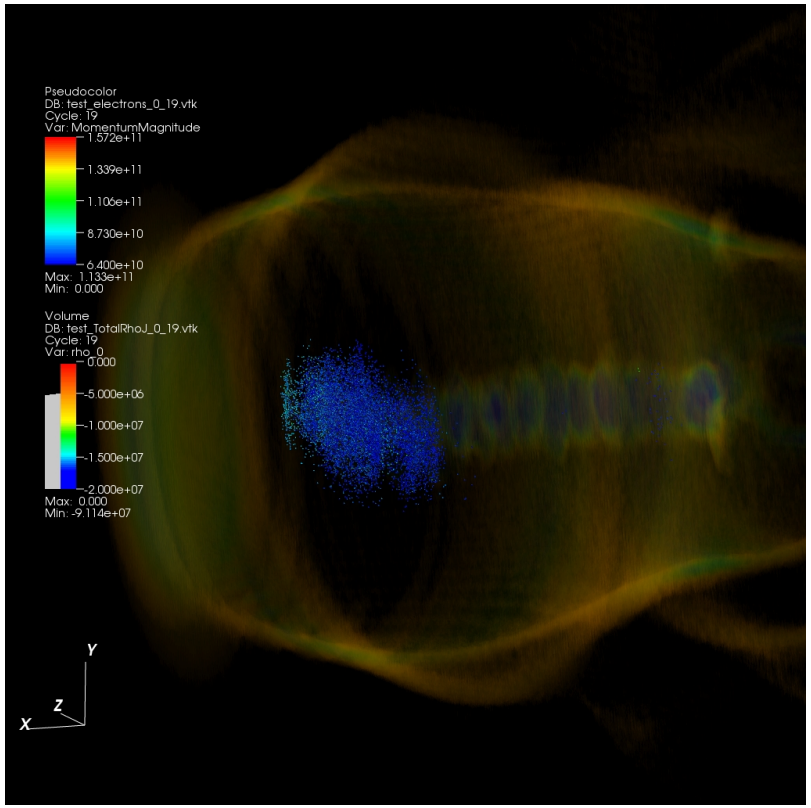

This image uses volume rendering to show the plasma density field. *(VisIt)*

2006 INCITE Project. PI: Cameron Geddes, LBNL

Accretion-Induced Collapse of White Dwarfs

- Data from 2D radiation-hydrodynamics simulations visualized using both AVS/Express and VisIt.
- Data include 68 scalar fields and 25 vector fields.
- Cristina wrote a reader for AVS/Express and converted the data to vtk format for VisIt.
- Future simulations will be 3D.

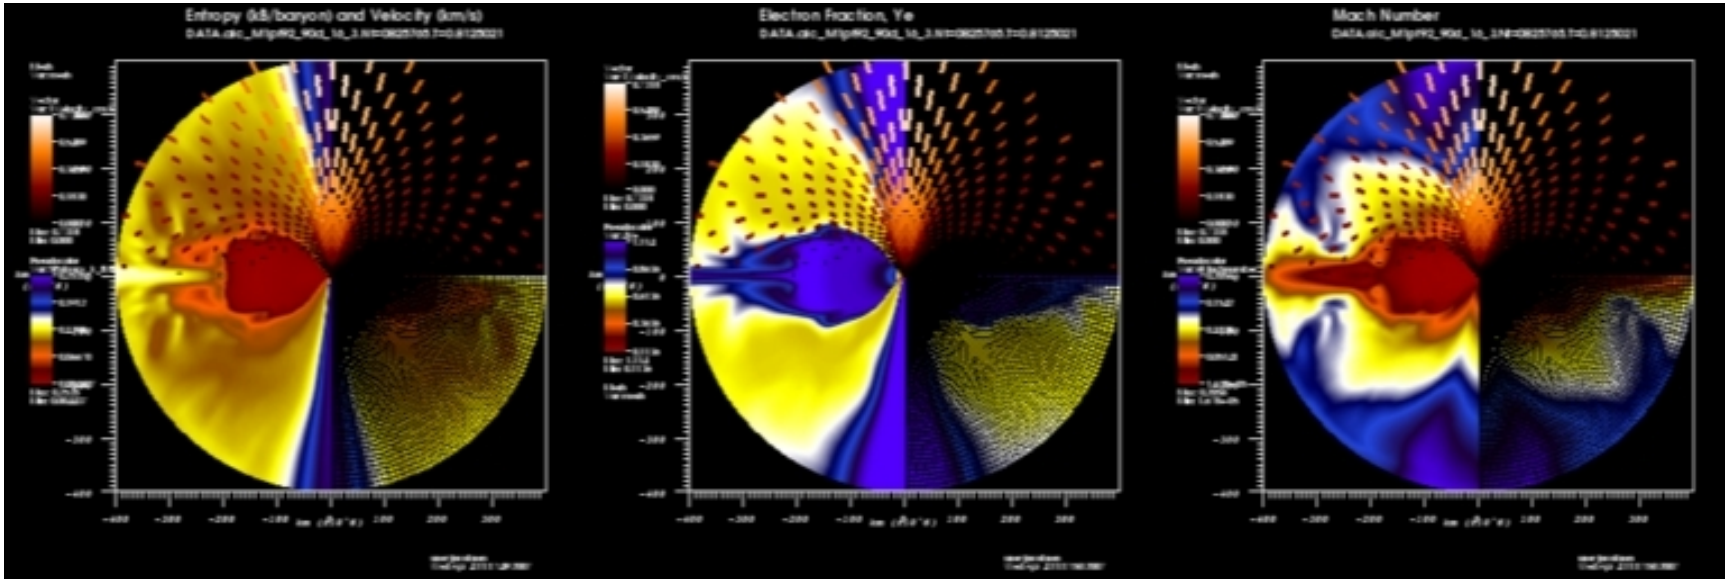

Accretion-Induced Collapse of White Dwarfs

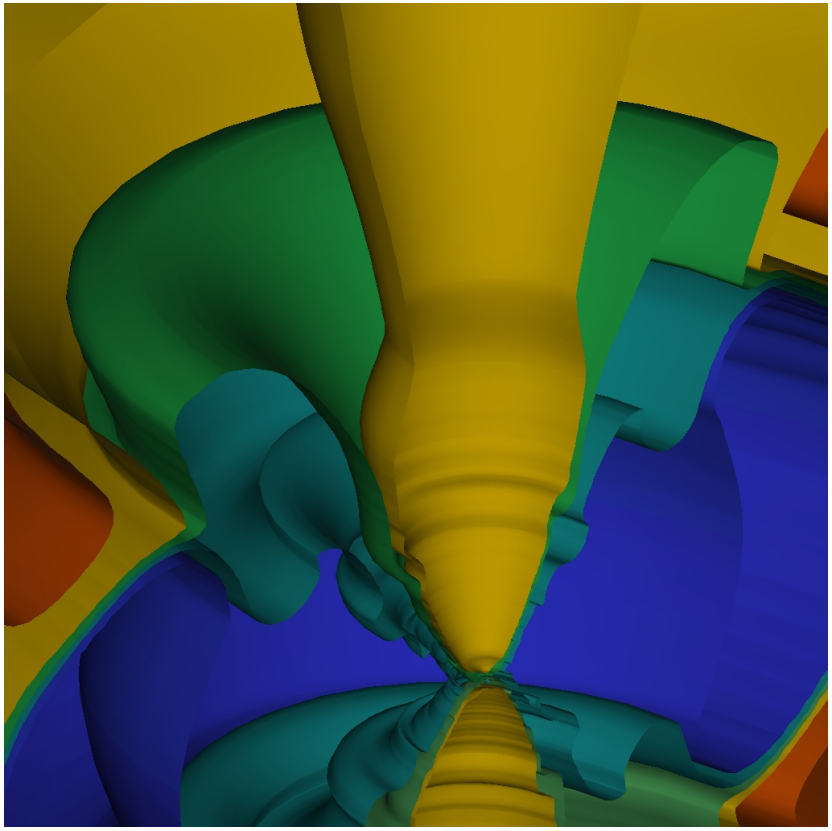

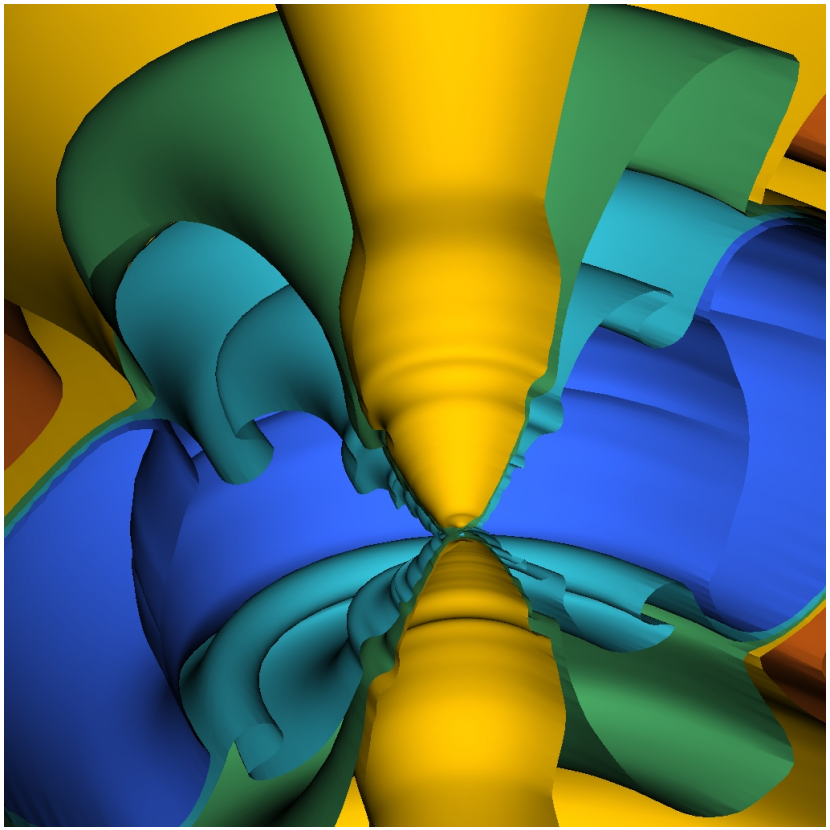

These are images of isosurfaces of entropy. One of these images was made with AVS/Express, and the other with VisIt.

NIMROD

- NIMROD: Fusion Magnetohydrodynamics Simulation Project
- Cristina wrote a VisIt plugin for NIMROD output that has been converted to HDF5.

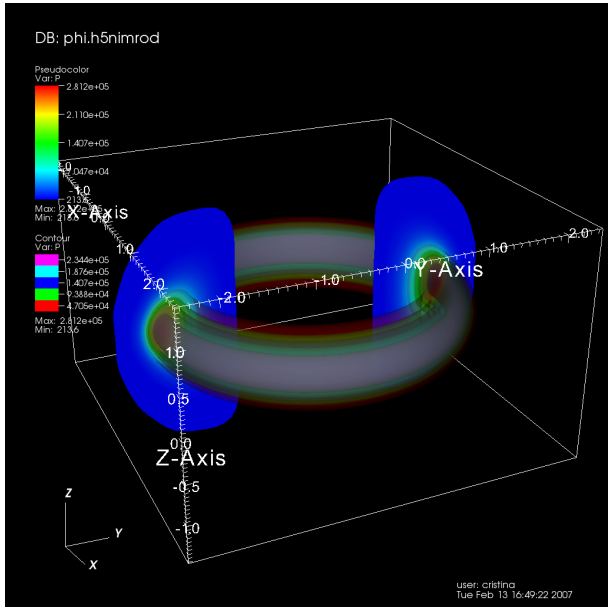

Pressure field visualized with cut planes and volume rendering.

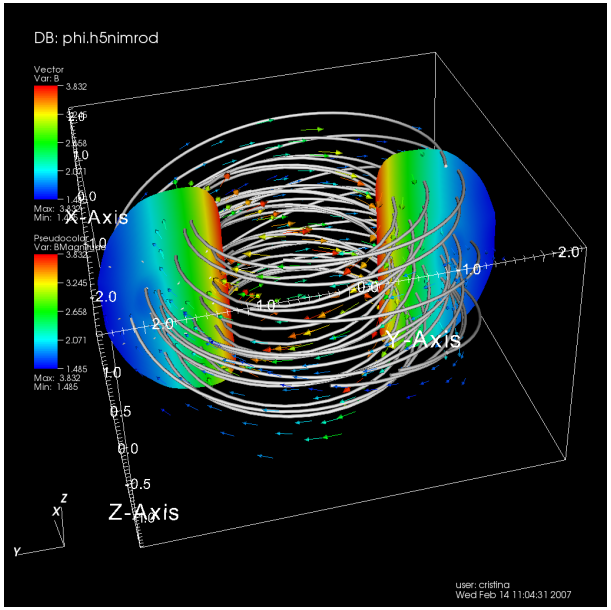

Magnetic field vectors and cut planes showing magnitude of field.

Journal Covers
Cristina Siegerist

- Journal of Physical Chemistry/Chemical Physics, March 2006
- Journal of Microscopy, July 2006
- SciDAC Review, Spring 2007

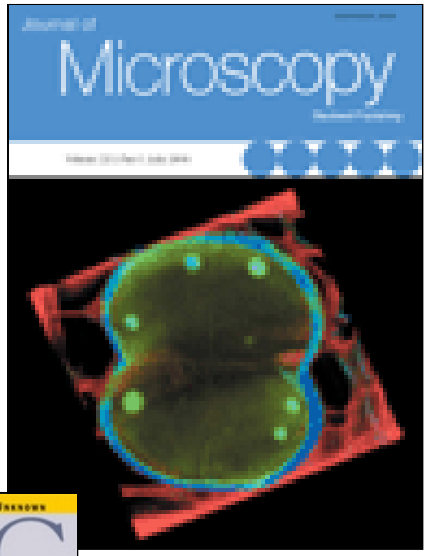

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

Accelerating Remote Display Performance - Gunther Weber

- Poor X11 network performance makes it difficult for remote users to use NERSC resources.
- Identify and evaluate alternatives to X11.
- Two alternatives under consideration: VNC (different vendors) and (Free)NX
- Evaluation criteria: speed/interactivity (usable over high latency link? frames/sec? time to finish update), deployability, security
- Must provide OpenGL/GLX functionality.

Simulating the Network

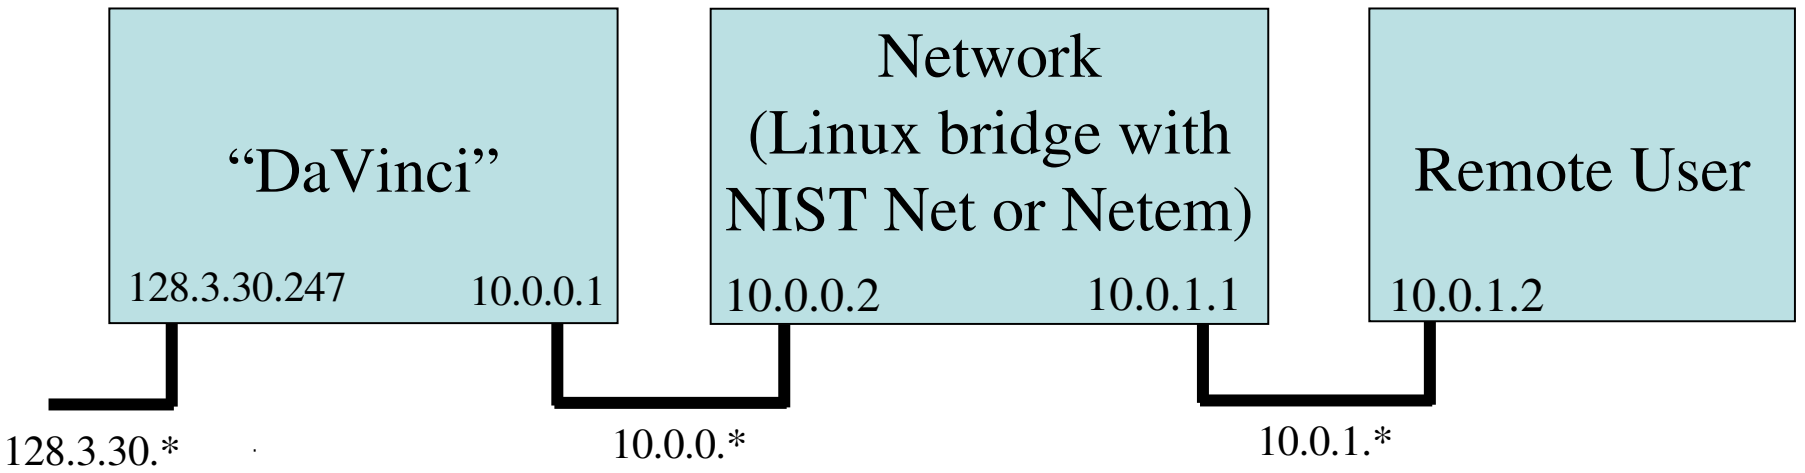

- **Run most test applications locally on fake "DaVinci"**
- **Use LBNL/NERSC network to test some applications (Matlab, …) on real DaVinci**

Measuring Performance

- Test with existing applications (Matlab, Maple, IDL, …)
  - Subjective speed experience
  - Time between "mouse click" and "last screen update" (using XTrap and XDamage extensions)
  - Frame per second updates
    - Draw alternating frames and measure frame change rate at remote client
    - Aliasing?
- Mock-up application
  - Need to be careful to ensure that application uses X11 in the same way as Matlab, Maple, IDL, *etc.*

Data Exploration using SUNFALL
Cecilia Aragon

- SNfactory data pipeline in 2005 involved many manual, repetitive tasks.
- The SUperNova Factory AssembLy Line (SUNFALL) is a software framework for data capture, data management, data analysis, data and information visualization, and workflow management.
- The SUNFALL project has
  - automated – and is continuing to automate – many repetitive tasks,
  - developed a suite of tools for data management (Data Forklift),
  - developed a Web-based monitor for workflow management,
  - developed improved image processing algorithms,
  - salvaged legacy code.

SNfactory Supernova Search Data Flow

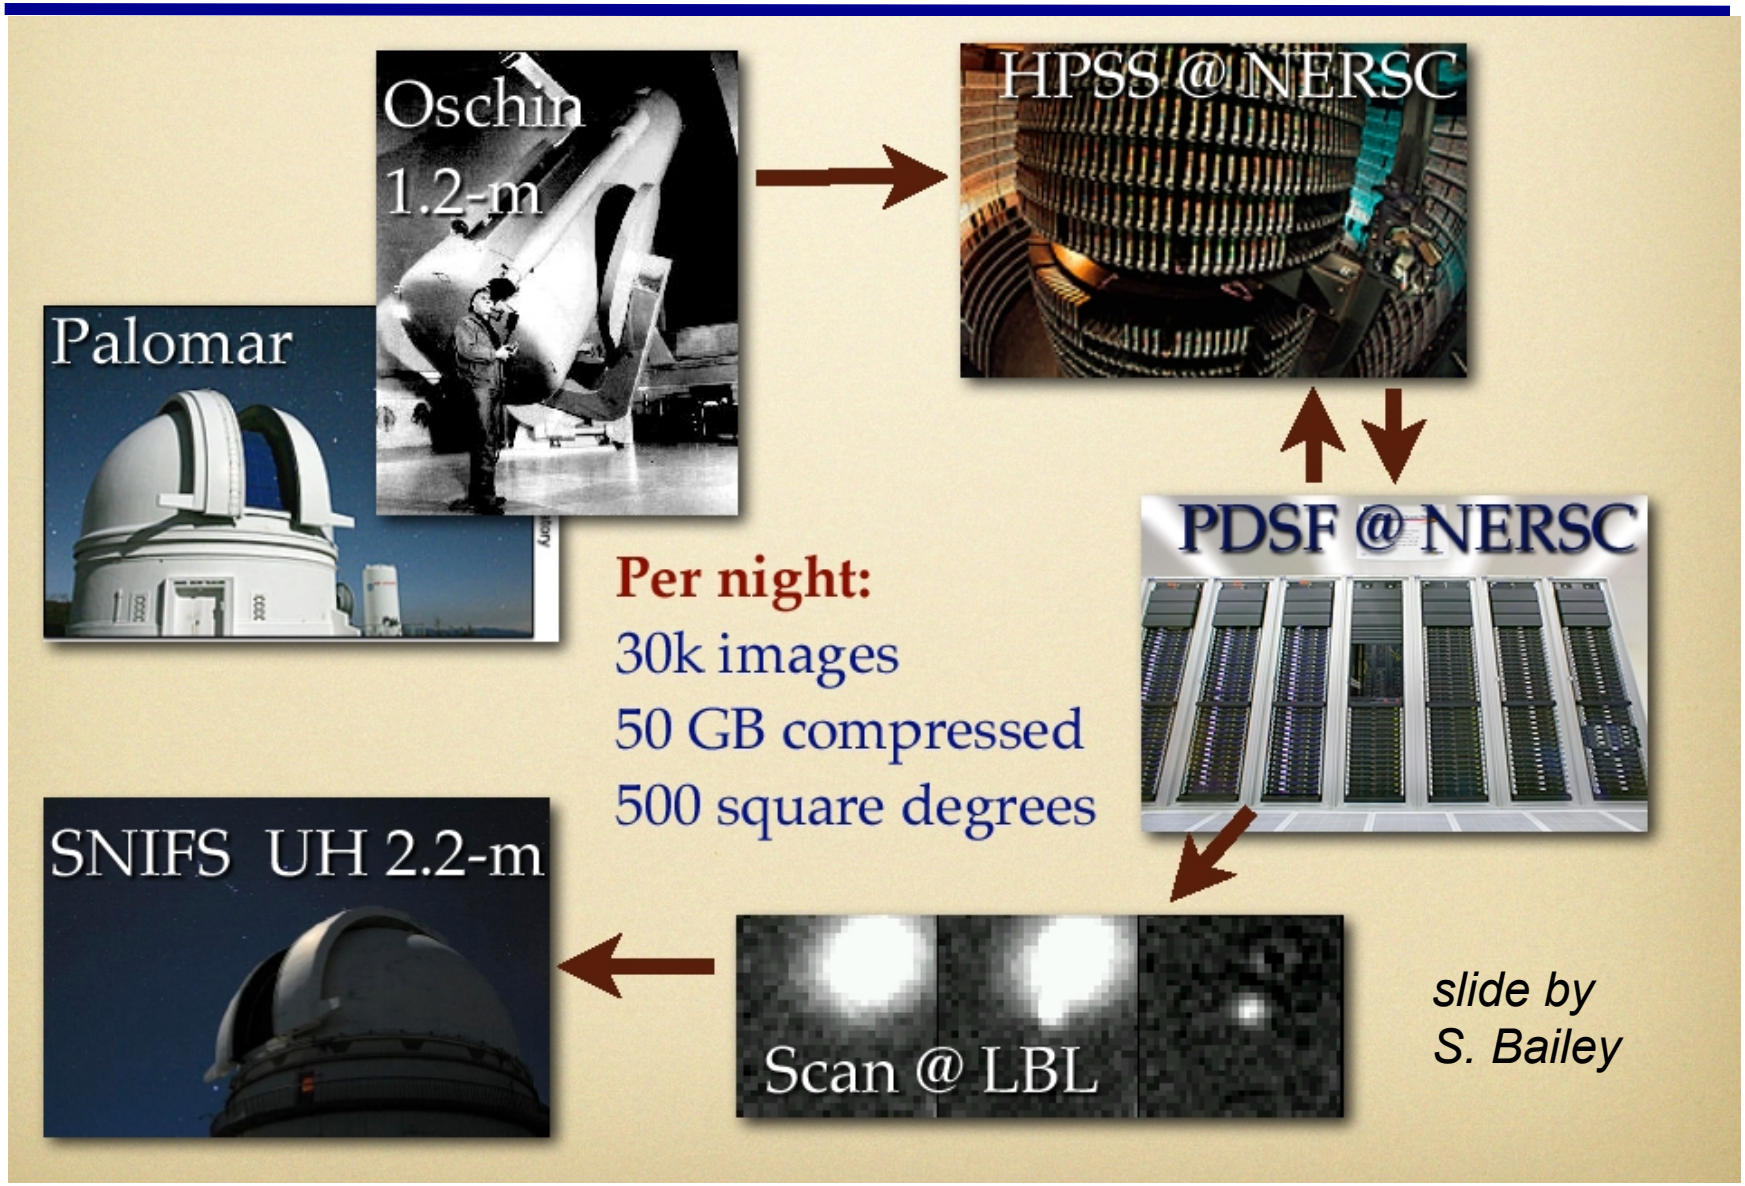

SNfactory Data Pipeline – Today

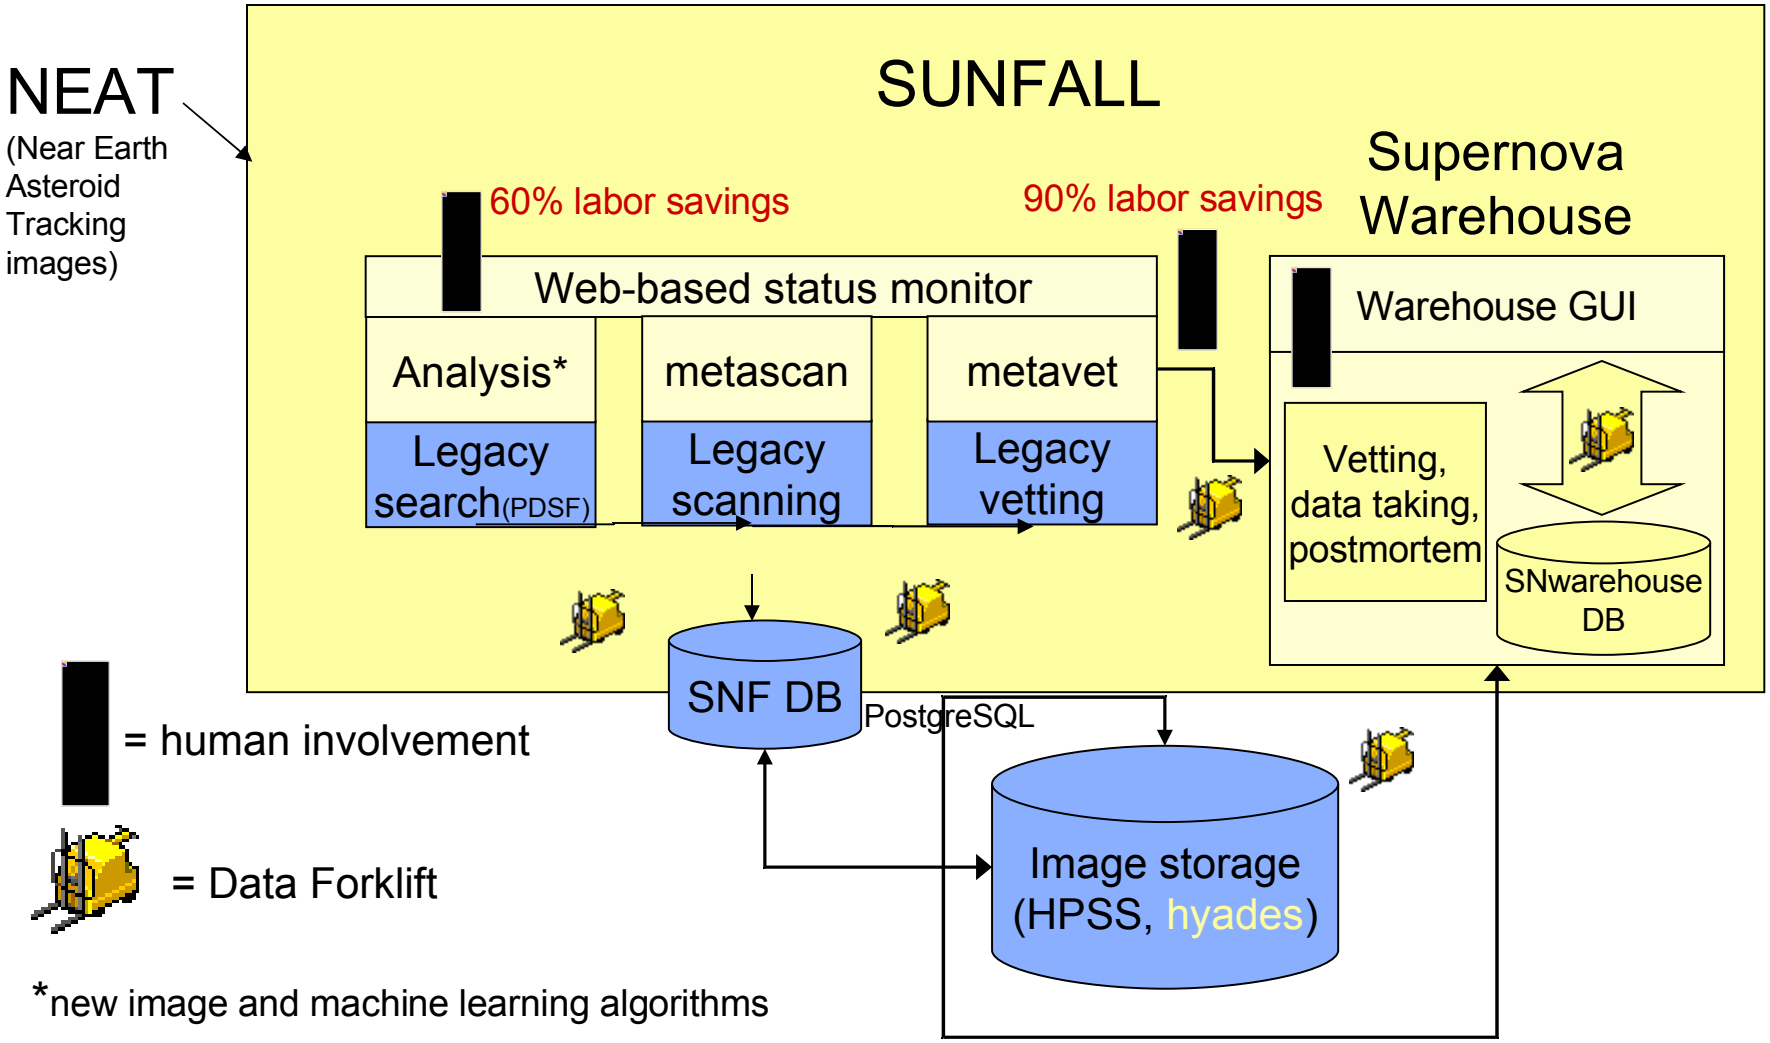

SUNFALL

- SUNFALL went into production in the fall of 2006. New capabilities continue to be incorporated.
- Number of false positive candidate supernovae reduced by 80%.
- Labor costs reduced dramatically: from 6 to 8 people working 4 hr/day to one person working 2 hr/day.
- Less time spent doing repetitive tasks means more time to do science, publish results.

SUNFALL Workflow Manager

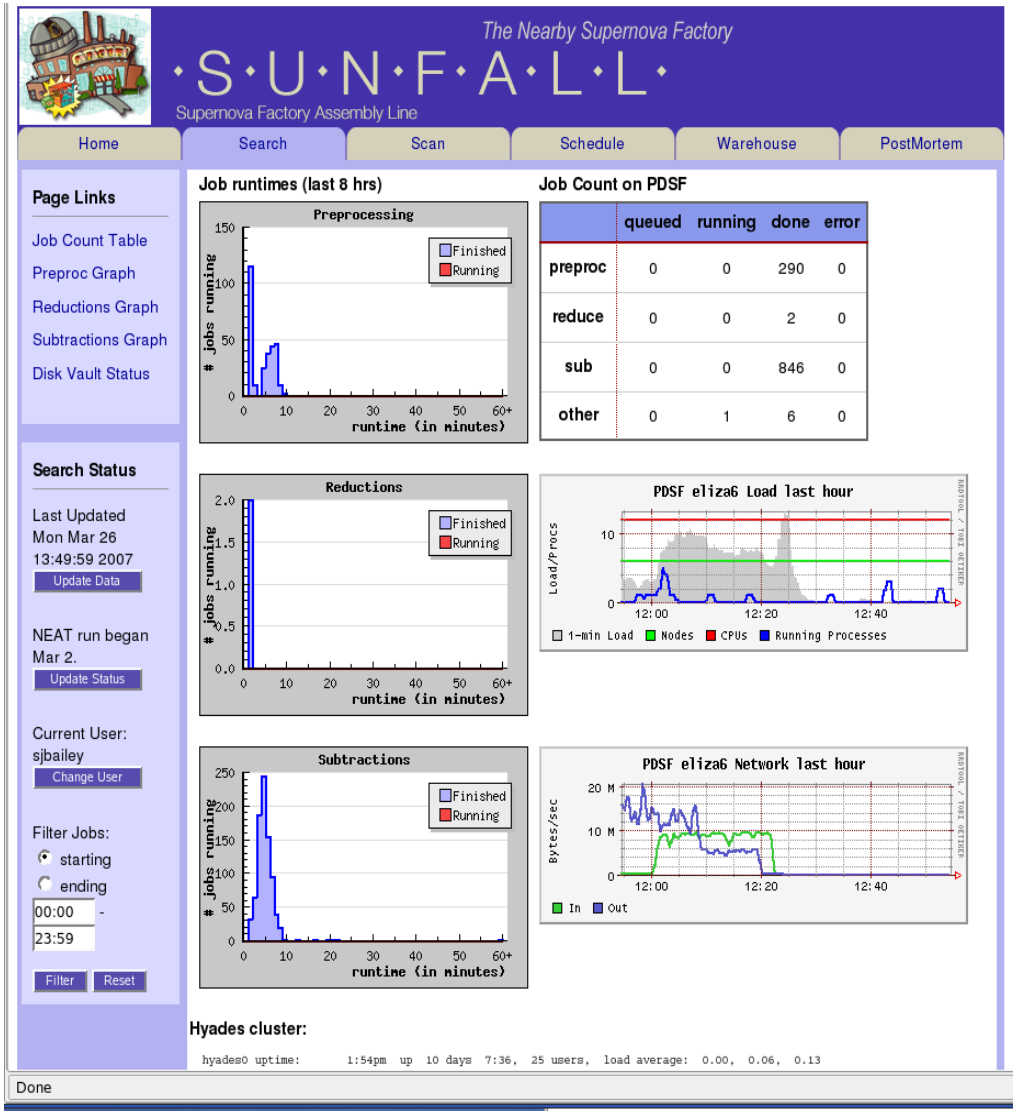

Supernova Warehouse

Provides tools for data management, data analysis, and data, information, and workflow visualization.

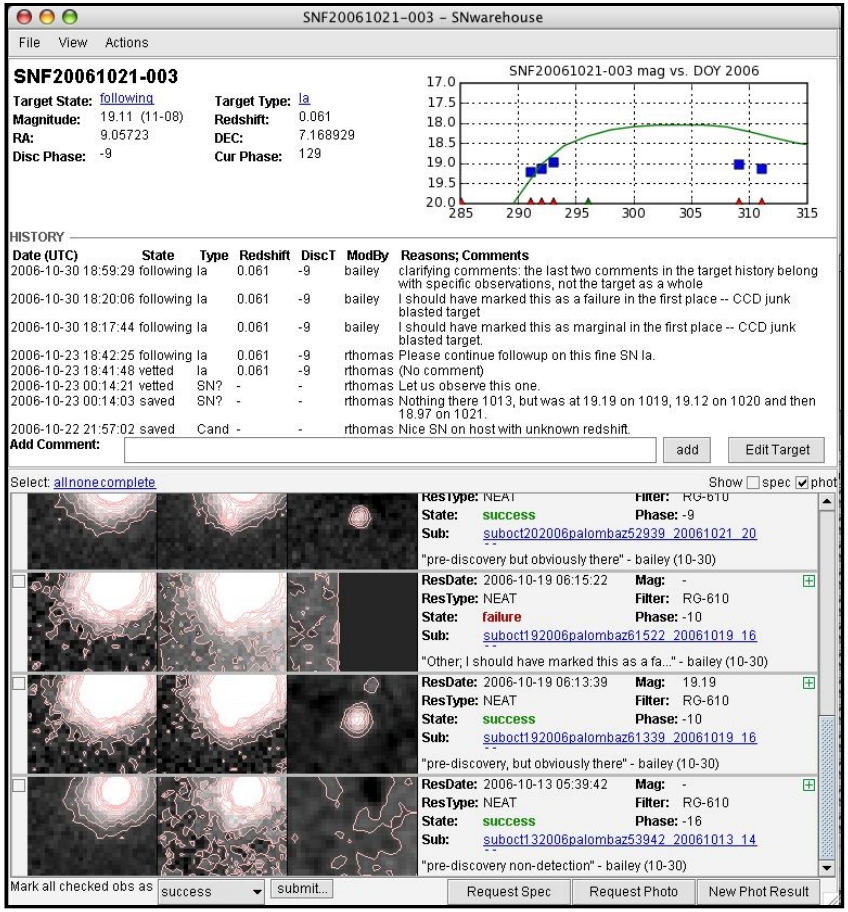

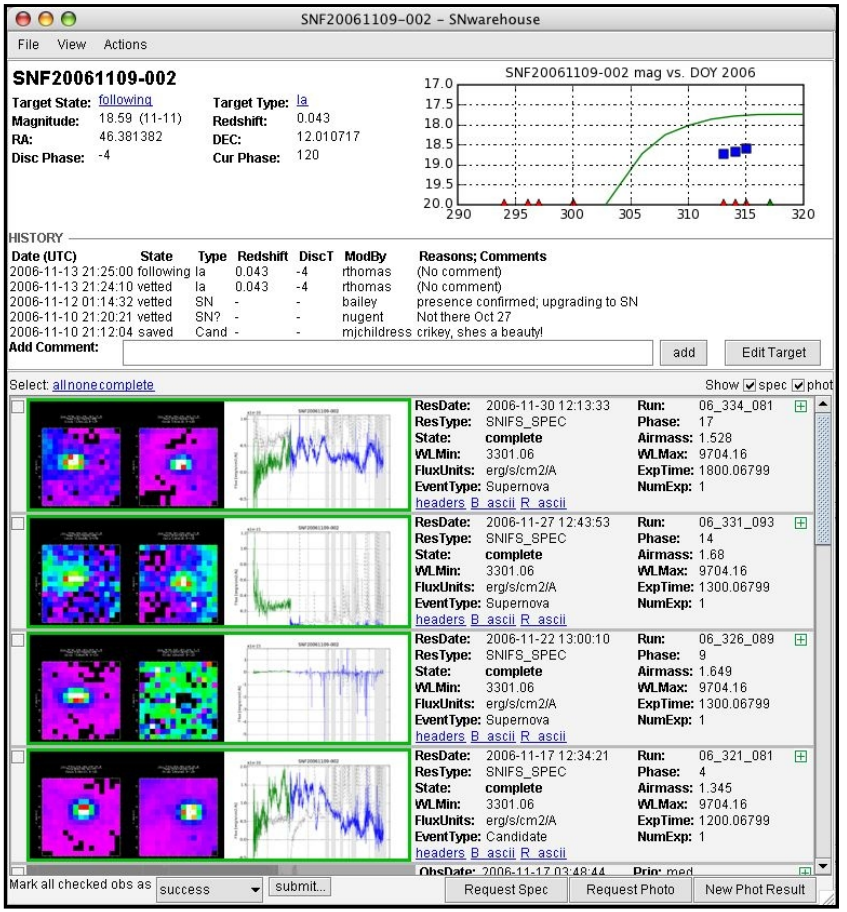

Acknowledgments

Members of Vis Group/
NERSC Analytics

- Wes Bethel (GL)
- Cecilia Aragon
- Peter Nugent
- Cristina Siegerist
- Kurt Stockinger
- Gunther Weber
- Raquel Romano (now at Google)