NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

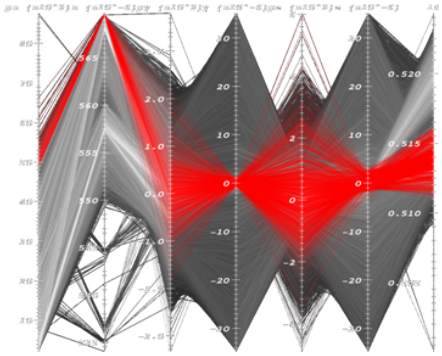

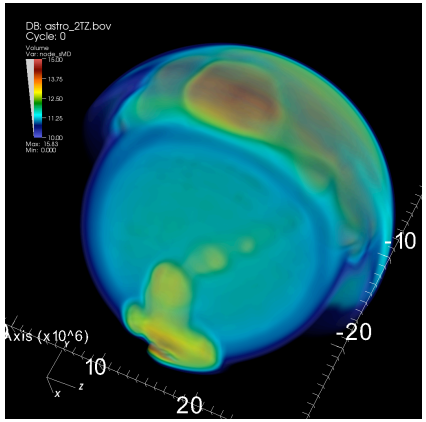

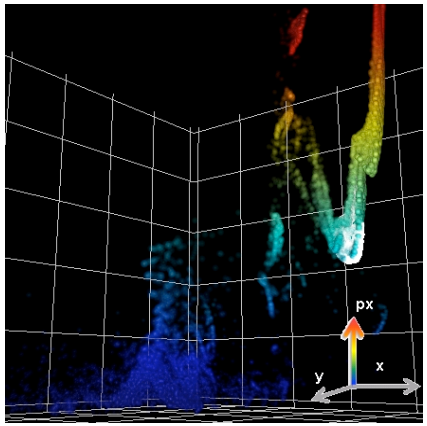

NERSC Analytics Accomplishments and Plans

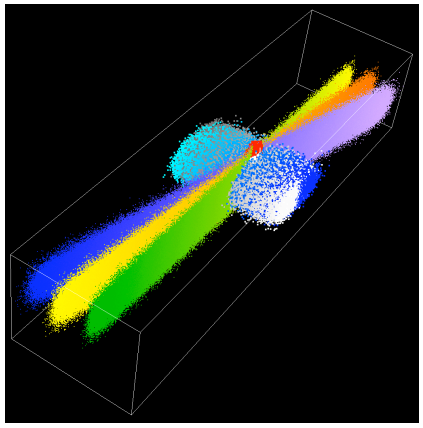

Hank Childs

NUG Meeting

October 7, 2009

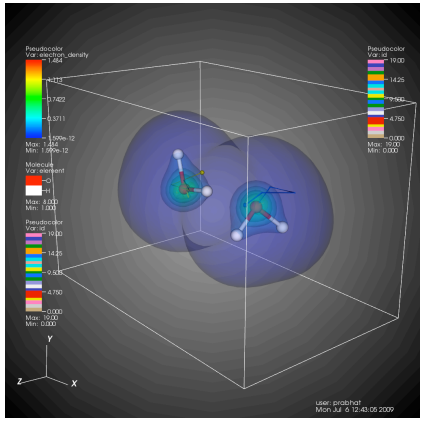

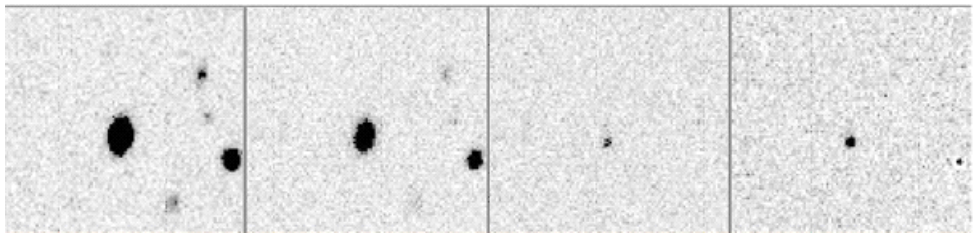

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Overview

- What do we mean by *analytics*?
- Mission statement
- Analytics resources
- Accomplishments: User projects
- Field-leading activities – VisIt hero run
- Contributions to infrastructure
- Analytics testbed: Tesla/Turing
- Plans for next year
- Members of the Analytics Group
- Contact information, Analytics Web page
- Analytics talks tomorrow
- Request for user feedback

What do we mean by *analytics?*

- **At NERSC, we include the following technologies in the Analytics Program:**
  - **visualization;**
  - **data analysis, data mining;**
  - **high performance I/O;**
  - **data management.**
- **We combine these technologies to enable *data exploration*.**

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Mission Statement

- **Our mission is to help NERSC users to analyze, visualize, and explore their data and simulation results. In some cases, we recommend software or algorithms to use; in other cases, we work collaboratively with users over a period of time. We use collaborations to augment and develop new tools and approaches for other users.**

Analytics Resources

- Hardware: daVinci is an SGI Altix with 32 Intel Itanium-2 processors and 192 GB of shared memory.
  www.nersc.gov/nusers/systems/davinci/index.php
  Tesla/Turing is a test analytics system, described later.
- Software: though there is some analytics and visualization software installed on all NERSC machines, the most analytics and visualization software is installed on daVinci. See the Software at NERSC pages for more information:
  www.nersc.gov/nusers/services/software/
  www.nersc.gov/nusers/resources/software/davinci/
- Staff for consulting and collaboration: the Analytics Program supports 3.5 FTEs at NERSC, who are shared with other programs (DOE Visualization and VACET). NERSC Analytics staff have expertise in visualization, analysis (data analysis and machine learning), and data management (data models, high performance I/O, and database systems).

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Accomplishments: User Projects

What follows are descriptions of user projects that the Analytics Group contributed to or collaborated on this past year. The examples include:

- **Visualization, including developing data models and data import tools.**
- **Troubleshooting I/O performance problems.**
- **Database solution for image data that enables data mining.**
- **Analysis of particle data from a laser wakefield acceleration simulation.**

Visualization of Quantum Monte-Carlo Simulations

- Sinisa Vukovic (UC Berkeley) and Maciej Haranczyk (LBNL)
- Developed modules to import electron trajectory, density and orbital information into VisIt
- Developed custom visualizations to understand low-energy configurations

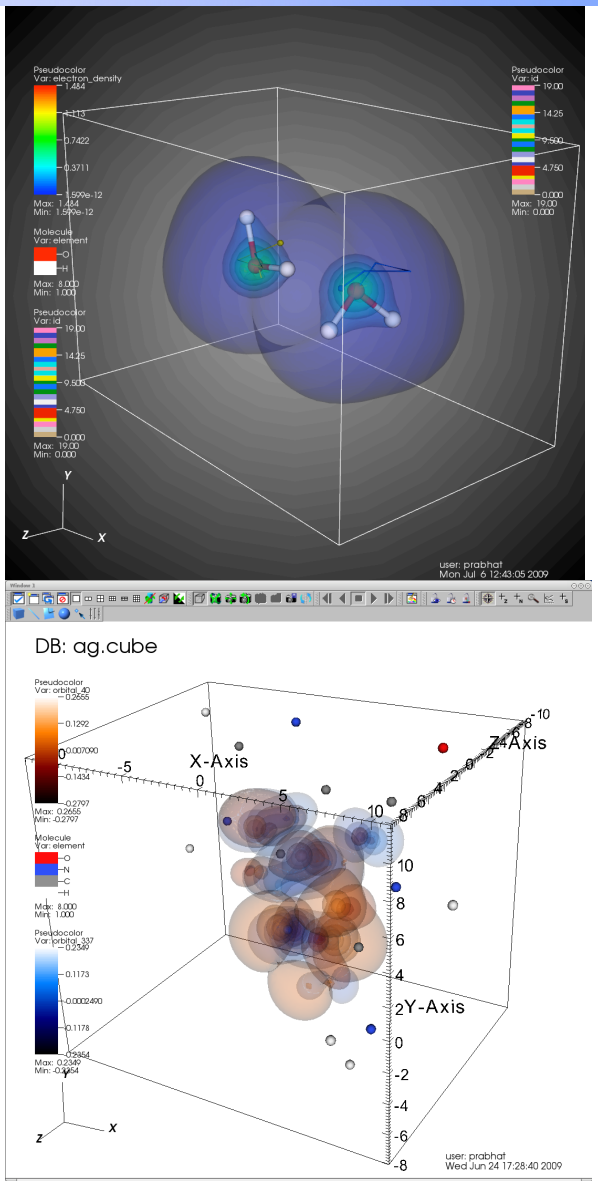

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Visualization of a Particle Beam

Fixing the view and transfer functions allowed the physicist to see how the beam changes as it travels through the chicane.

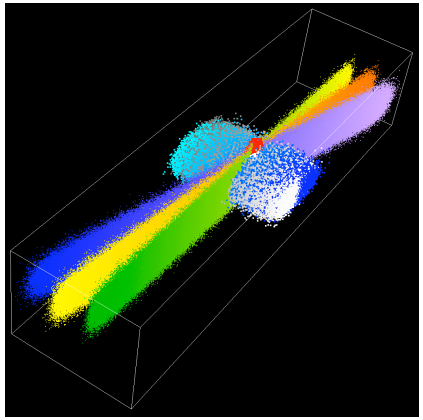

beams superimposed

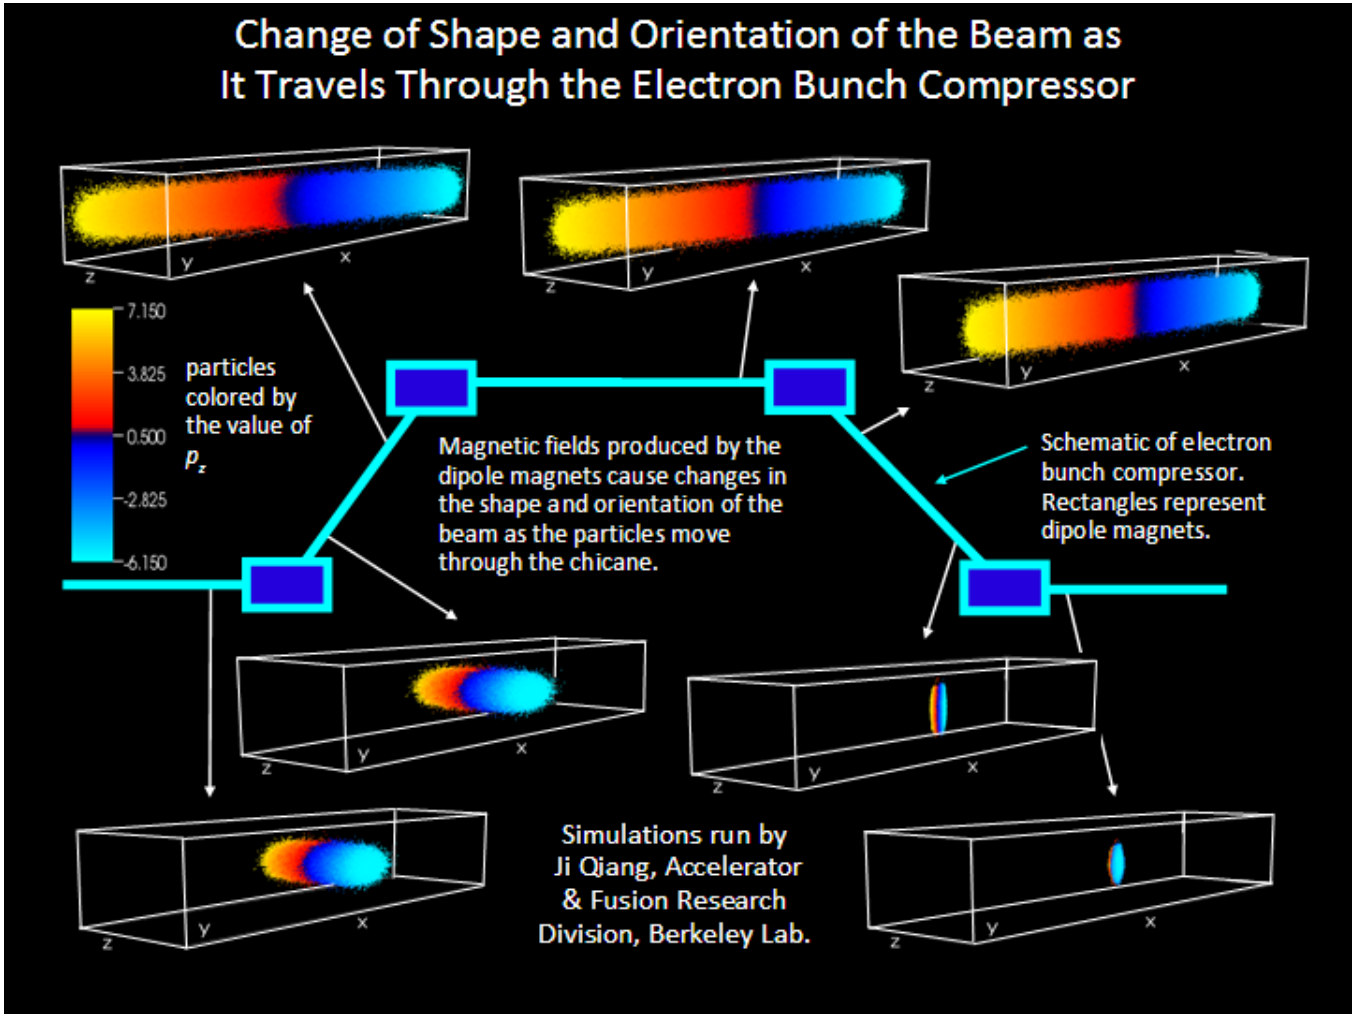

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

Visualization of a Groundwater Model

- Tim Schiebe and Bruce Palmer (PNNL)
- Hybrid Numerical Methods for Multiscale Simulations of Subsurface Biogeochemical Processes
- NERSC Analytics collaboration:
  - Data model for particles
  - Production deployment of parallel visualization capabilities (VisIt)

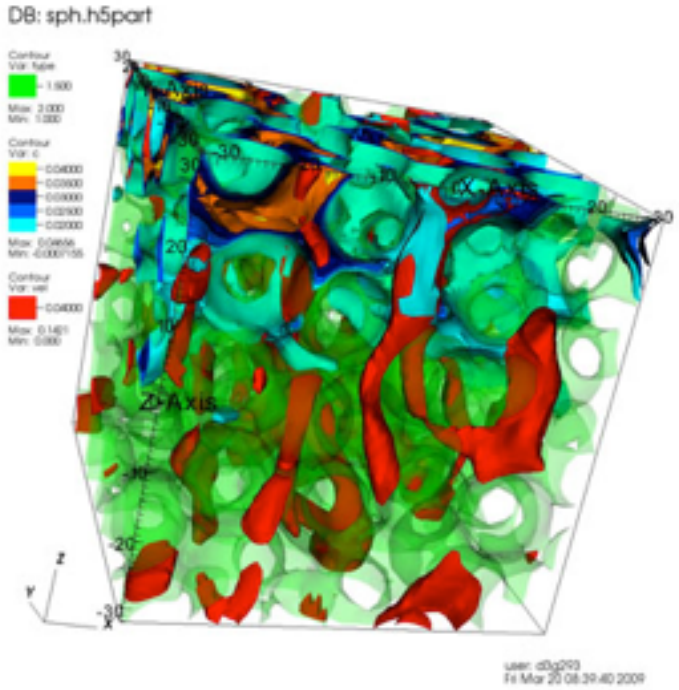

Visualization of a 3D pore-scale simulation of fluid flow and solute transport in a complex pore geometry. Red – velocity isosurfaces, blue to yellow – concentration isosurfaces, green – grains.

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Climate: Enabling Global Cloud Resolving Model Simulations

- Dave Randall (CSU) and Karen Schuchardt (PNNL)
- Specification of data model for geodesic grid and variables
- Troubleshooting collective I/O on Franklin
  - Found issues in MPI-IO implementation
  - Improved performance from ~100MB/s to ~4GB/s
- Visualization of TBs of data
  - Developed modules to load data in VisIt
  - Customized 2D plots

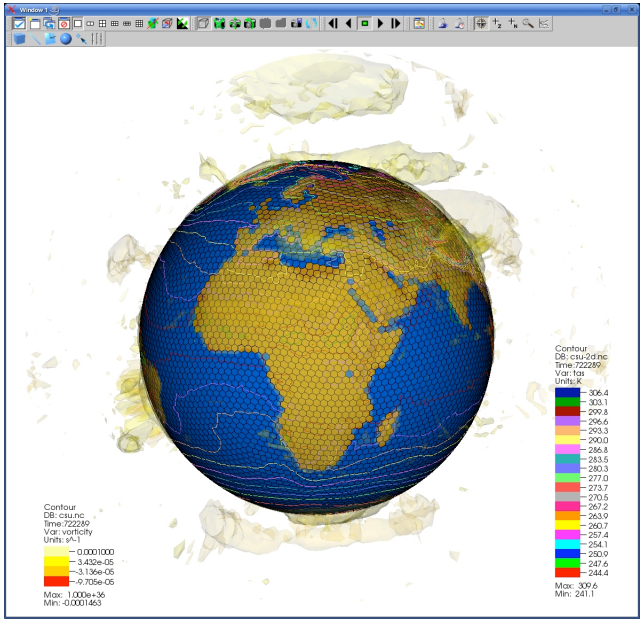

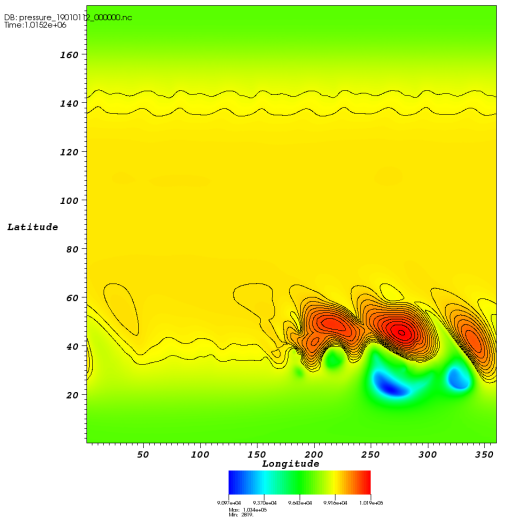

Collective I/O Performance

- Many users reporting 1GB/s write bandwidth on Franklin
- Profiled low-level I/O operations with IPM
- Discovered inefficiencies in Cray's MPI-IO implementation
- Cray revised MPI-IO to better interact with lustre filesystem
- Write bandwidth for several use cases rose to 6GB/s

IPM trace of old MPI-IO implementation writing 1 file at 1.2GB/s:

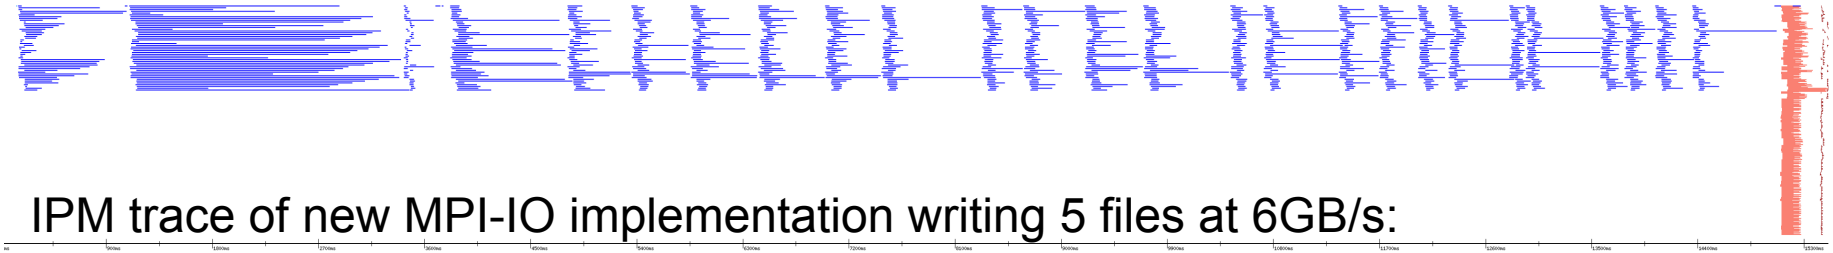

IPM trace of new MPI-IO implementation writing 5 files at 6GB/s:

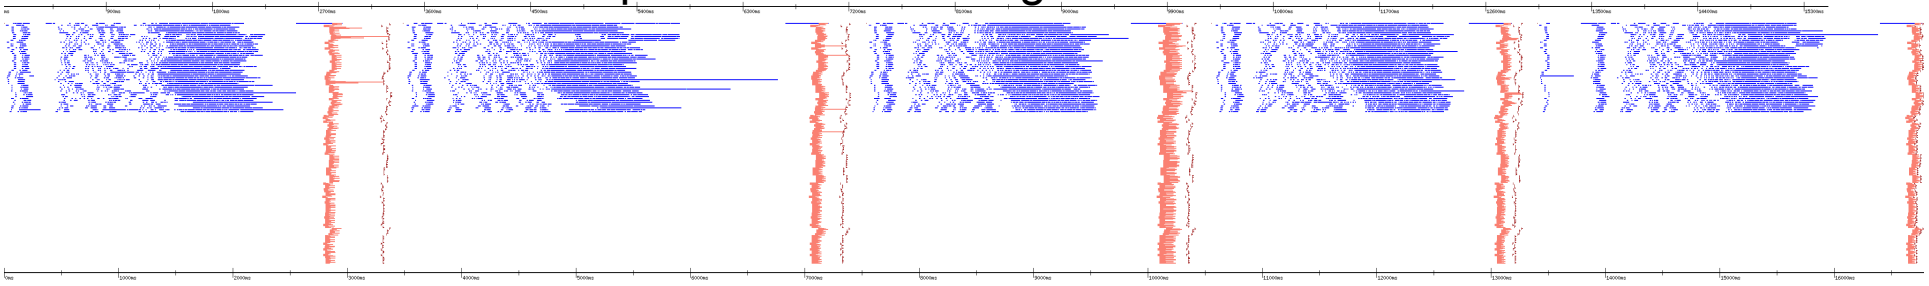

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Data Intensive Science: Deep Sky and the Palomar Transient Factory

- PI: Peter Nugent, CRD/LBNL
- Phase I: generate repository of deep reference images from historical data from Palomar QUEST sky survey.
  - Deep reference images used for transient surveys to discover supernovae. Web interface developed by Randy Kersnick (OSG) to serve the deep reference images to the public. Science gateway node used to host the database and serve the images.
- Phase II: apply infrastructure from Phase I to Palomar Transient Factory (PTF) sky survey to develop subtraction pipeline.
  - Database key to identifying candidate transient objects for follow-up telescope time.
  - 51 supernovae discovered by the PTF pipeline in four months.

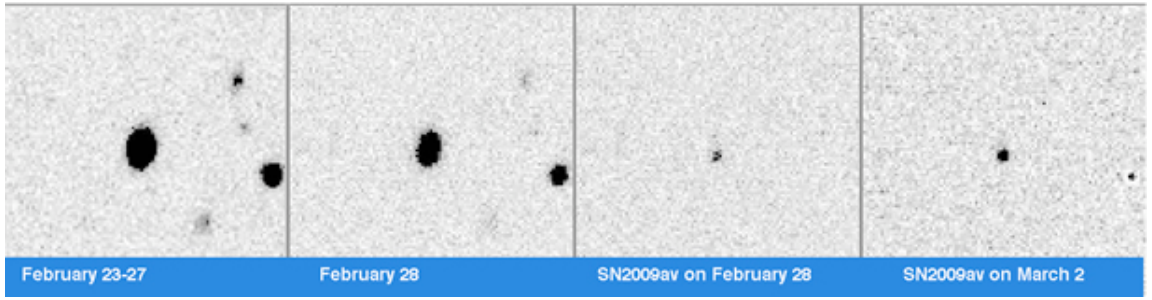

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Analysis: Automated Detection and Analysis of Particle Beams in Simulations of Laser Wakefield Acceleration

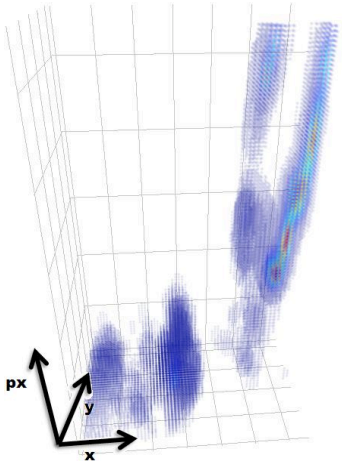

(a)

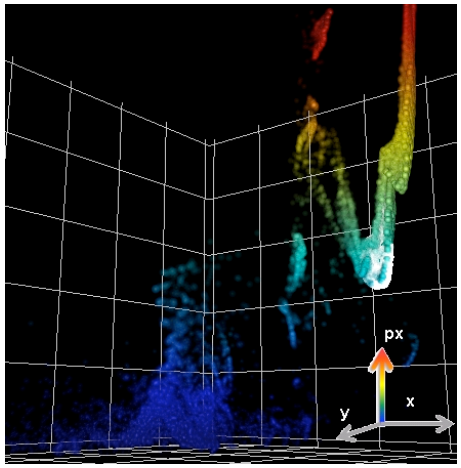

(b)

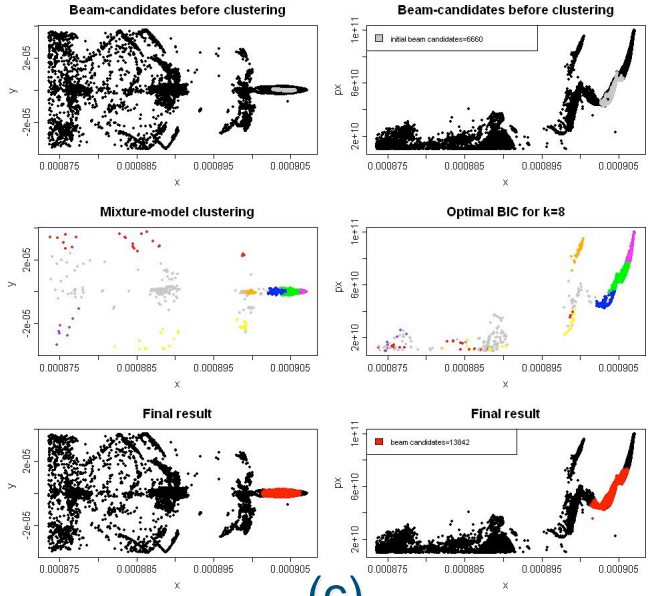

(c)

Development of novel unsupervised machine learning techniques to automatically locate and analyze beam particles:

- **Kernel-density estimation to represent particle densities;**
- **Detection of a high density bunches of particles;**
- **Model-based clustering of particles to determine a compact electron bunch, followed by time series analysis: moving averages to identify periods of bunch stability and dispersion measures to characterize beam compactness and quality.**

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Analysis: Interactive Exploration of Laser Wakefield Accelerator Simulation Data

- **Integrated FastBit and VisIt**
  - Provide new capability for efficient rendering of parallel coordinates.
  - Support interactive selection of particle beams in extremly large datasets.
  - Support fast tracing of particles over time.
  - New capability for temporal persistent selection of particles.
- **Science Impact:**
  - Enabled a more interactive and accurate analysis process.
  - Reduced the time for tracing of particle beams from hours to seconds.

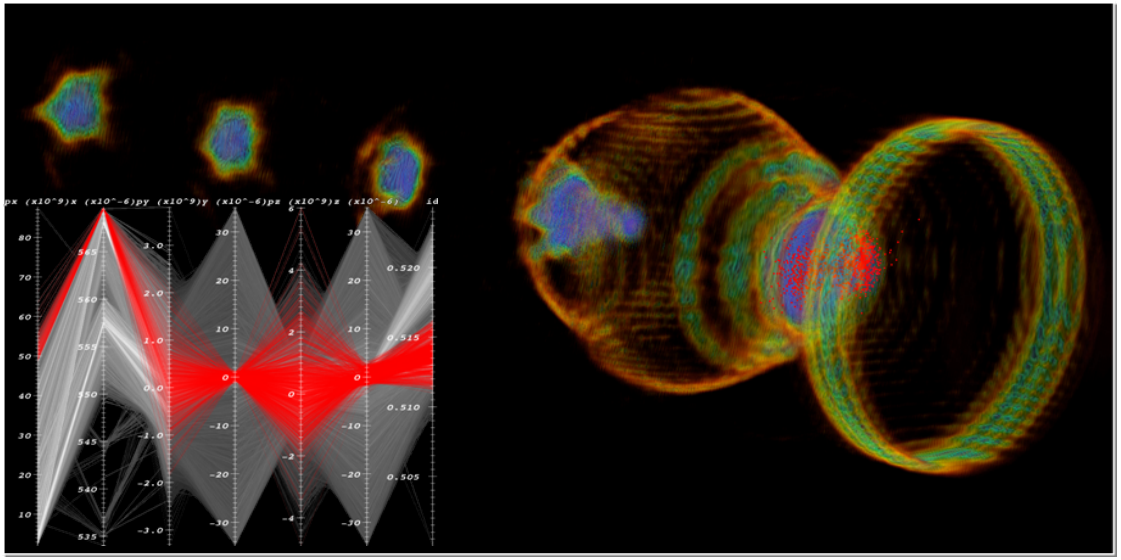

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

Analysis: Automatic Beam Path Analysis for Simulations of Laser Wakefield Acceleration

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Development of efficient method for automatic temporal classification of multiple particle bunches.

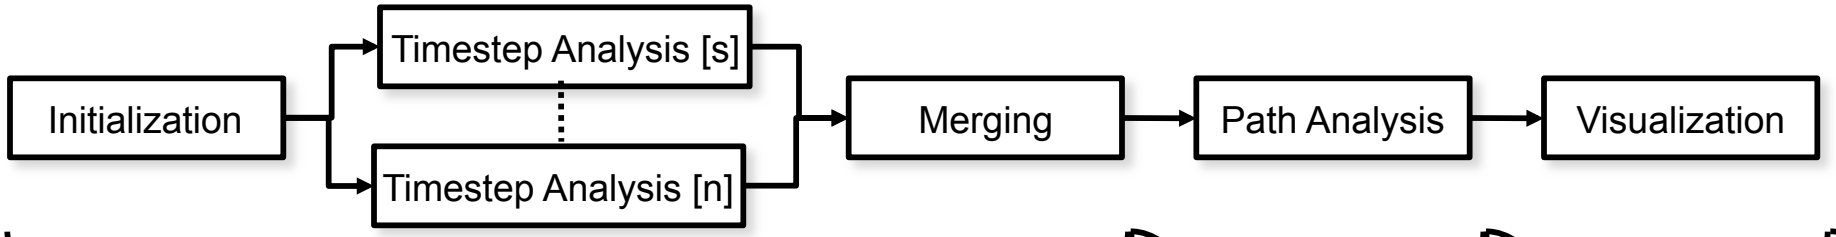

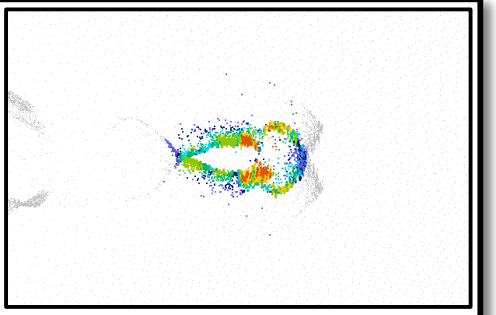

Detect the particle bunches and define for each bunch a set of reference particles.

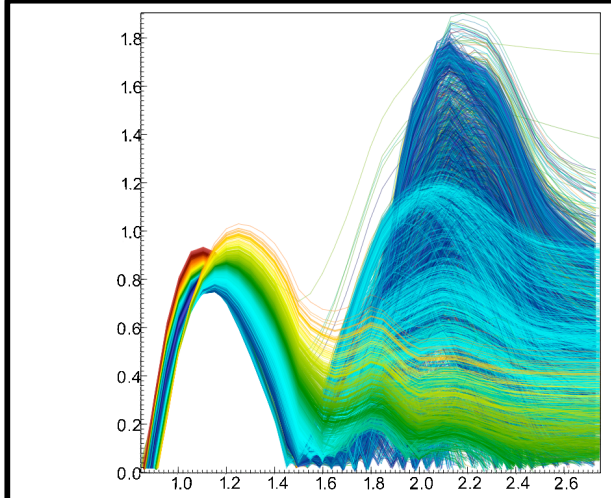

Compute distances of particle traces to bunch reference paths.

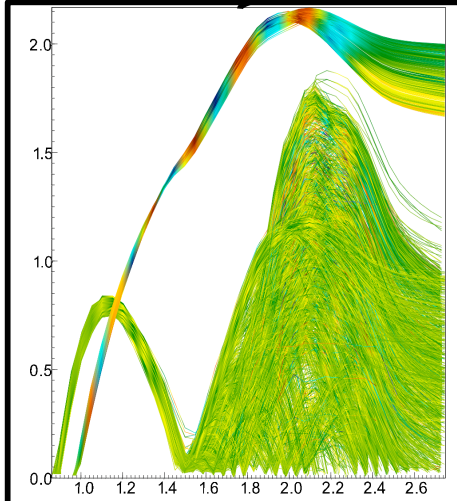

Analyze traces and compare bunches.

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

Contributions to Infrastructure

- Installed software, primarily on daVinci. New this year: R, R packages, more Python libraries, ImageJ.
- Installed VisIt on Franklin.
- Installed NX on Franklin.
- Collaborated on or developed extensions to VisIt including improving streamlines and enabling FastBit.
- Standing up analytics testbed – *next set of slides*.

Analytics Testbed: Tesla/Turing

- **Small-scale testbed platform with shared-memory, multi-core architecture and multiple GPUs.**
- **Gain experience with user applications and administration of such a platform.**
  - **Batch queues and GPUs (GPU/CUDA, OpenGL/visualization)**
  - **Remote delivery of hardware-accelerated graphics/visualization**
- **Service some user analytics workload: large memory footprint, shared-memory parallel, graphics intensive.**

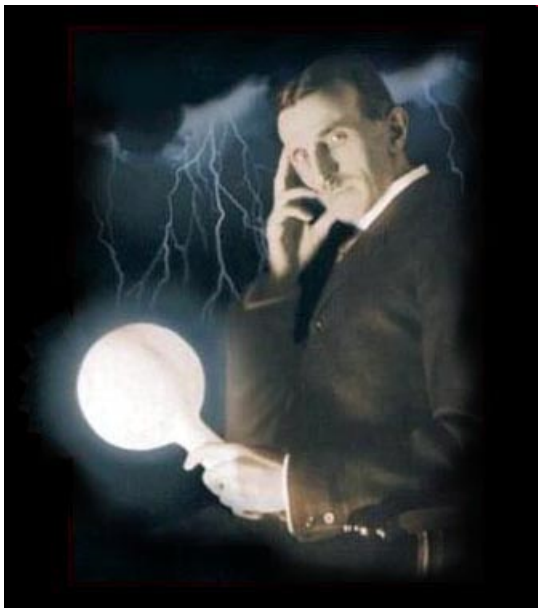

Nikolai Tesla, inventor of alternating current

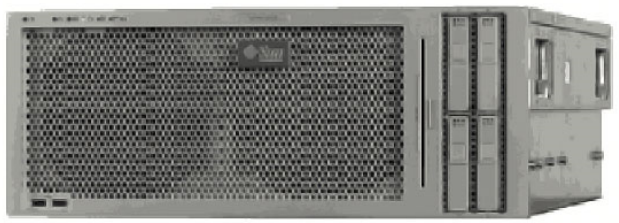

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Tesla/Turing System Architecture

- Tesla and Turing are identically configured SUN Sun Fire X4600-M2s
- Hardware Specifications (per system):
  - 8 AMD 2.8 GHz Quad Core Processors - 32 cores (Shanghai 75W)
  - 256GB of shared Memory - DDR2-667 (4GB DIMMS)
  - 10Gb Ethernet connection to NERSC network
  - 2 Dual-Port 4Gb FC HBA connections to NGF (/project, /global)
  - InfiniBand DDR HBA (private Tesla-to-Turing network)
  - Local disk for OS file systems only, no local /scratch
  - NGF file systems for /global/homes, /project and /global/common
  - Operating System - Scientific Linux (RHEL 5.3 derivative)
- Graphics Specifications
  - Nvidia QuadroPlex 2200 s4 (1u enclosure - shared between Tesla/ Turing)
  - Four FX-5800 Quadro GPUs - two per host via PCIe-x8 adapters
  - 16GB Memory total - 4GB per GPU
  - 960 CUDA Parallel Processor Cores total - 240 per GPU

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Tesla/Turing – Access and Use Policies

- Users may request access to Tesla/Turing to test:
  - CUDA/GPGPU, shared-memory multi-core tests, hardware-accelerated graphics/visualization with remote delivery, etc.
  - When? 1 November 2009.
- How to gain access?
  - Short term: contact David Turner in User Services.
  - Longer term: Web form.
  - Users must declare intended use of machine and must provide a brief report of testing results.
- Duration of access:
  - 90-day access period, may be renewed for additional 90-day units.
- Service level:
  - Best effort: Tesla/Turing is a test machine, not production.
  - It may be taken down without much advance notice

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Field Leading Activities: Visualization at Unprecedented Scale and Size

Performance experiment aimed at understanding the limits of scalability of visualization apps on DOE's largest machines.

Franklin runs:

16K cores, 1 trillion zones
32K cores, 2 trillion zones

Largest-ever visualization runs in DOE community.

Lessons learned benefit all NERSC users who run at scale.

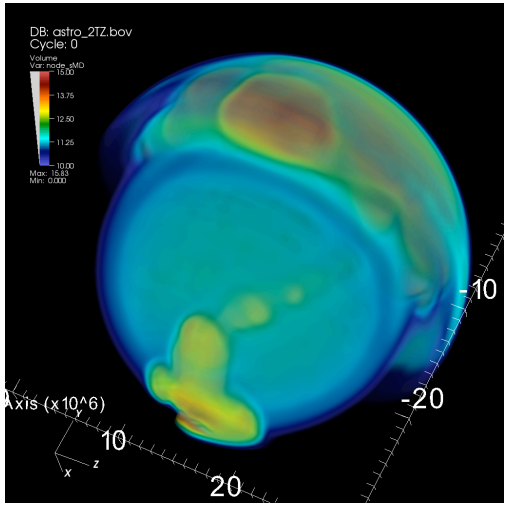

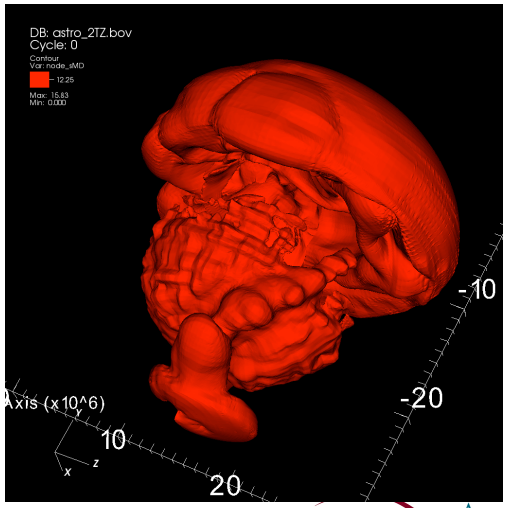

Visualizaton of CHIMERA supernova simulation. Top: Isocontouring of two trillion zones on 32,000 Opteron cores of Franklin. Bottom: Volume rendering of two trillion zones on 32,000 cores. *Simulation: Tony Mezzacappa ORNL, Visualization: NERSC Analytics.*

Office of Science
U.S. DEPARTMENT OF ENERGY

BERKELEY LAB

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Plans for Next Year

- Contact users who noted dissatisfaction with analytics in the NERSC User Survey.
- Contact top daVinci users to assess how their work will be impacted when daVinci is turned off and what analytics/visualization challenges they anticipate in 2012-2013.
- Contact users who ask for help with analytics/ visualization in their ERCAPs.
- Improve documentation for analytics/visualization software at NERSC, including developing new tutorials.
- Perform testing using Tesla/Turing as noted earlier.

Members of the Analytics Group

Wes Bethel, Group Lead

Hank Childs*, Mark Howison*, Janet Jacobsen, Prabhat, Oliver Ruebel, Daniela Ushizima*, Gunther Weber

*** Giving talks tomorrow**

Contact Information

The Analytics Group is part of the User Services Department.

If you have an analytics questions, send email to consult@nersc.gov. USG will forward your email to us.

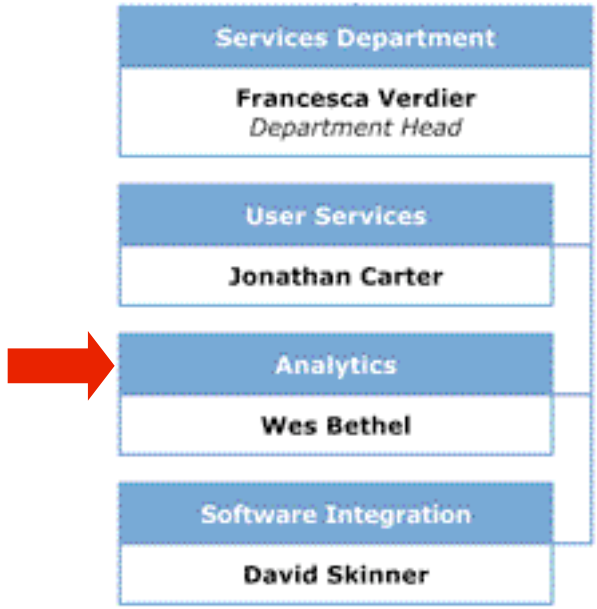

www.nersc.gov/nusers/analytics

Analytics Talks Tomorrow

- I/O Using H5part – Mark Howison
- Scientific Visualization using VisIt
  - Hank Childs
- Data Analysis Using the R Project for Statistical Computing
  - Daniela Ushizima

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

Request for User Feedback

- Why don't more users do visualization and data analysis/data exploration at NERSC?
- What are the obstacles you encounter in trying to do visualization and data analysis at NERSC?
- As your datasets grow in size and complexity, what resources will you need at NERSC to do visualization and data analysis/data mining/data exploration?
- Other?