



# VACEET

## Occam's Razor and Petascale Visual Data Analysis

*SciDAC 2009 Program Meeting*

*15 June 2009*

*San Diego, CA*

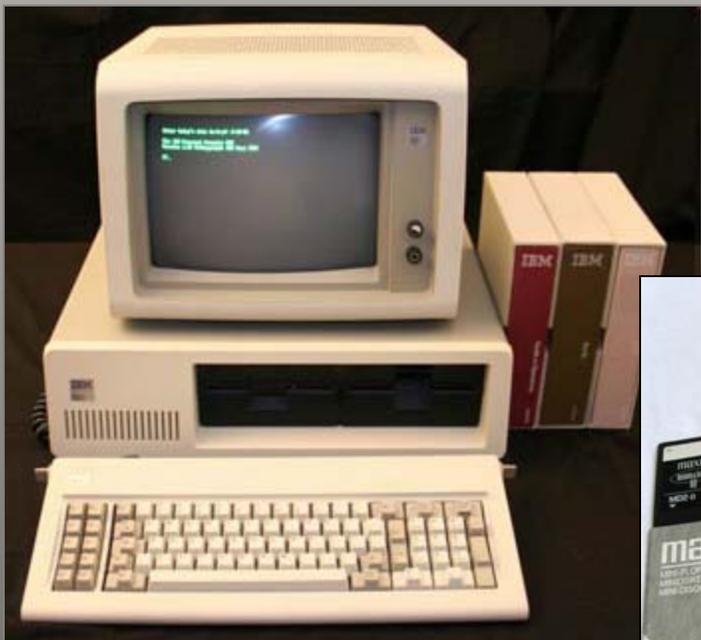
E. Wes Bethel

Lawrence Berkeley National Laboratory



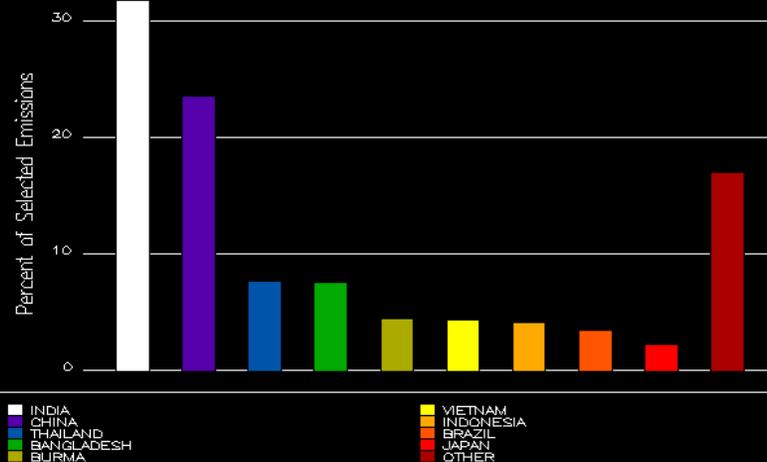
# VACET

## Remember When: 1981



Scenario Database Mapping Reports Exports Tools Exit

Methane Emissions from Rice Growing  
Tg/yr



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

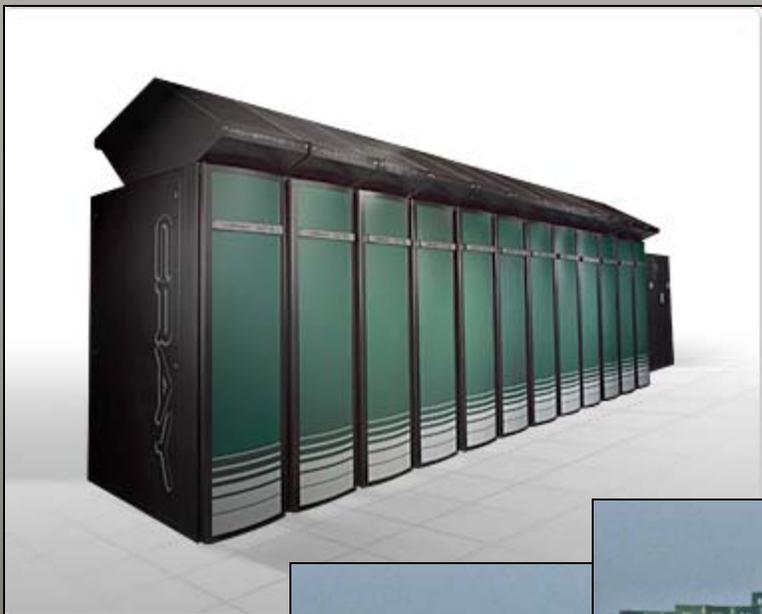


SciDAC  
Scientific Discovery through Advanced Computing



# VACET

This is no joke!



U.S. DEPARTMENT OF ENERGY  
**ENER**

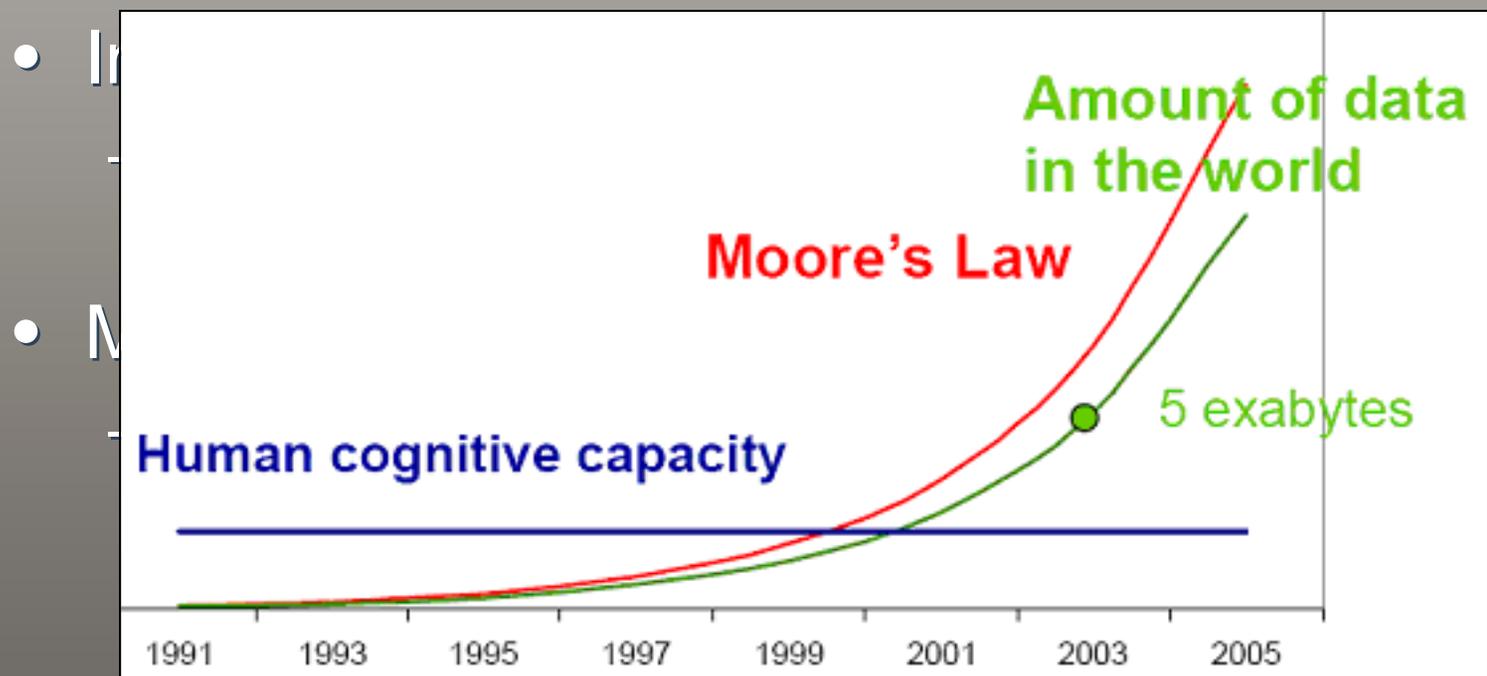


**SciDAC**  
Scientific Discovery through Advanced Computing



## Big Problem – Information Overload

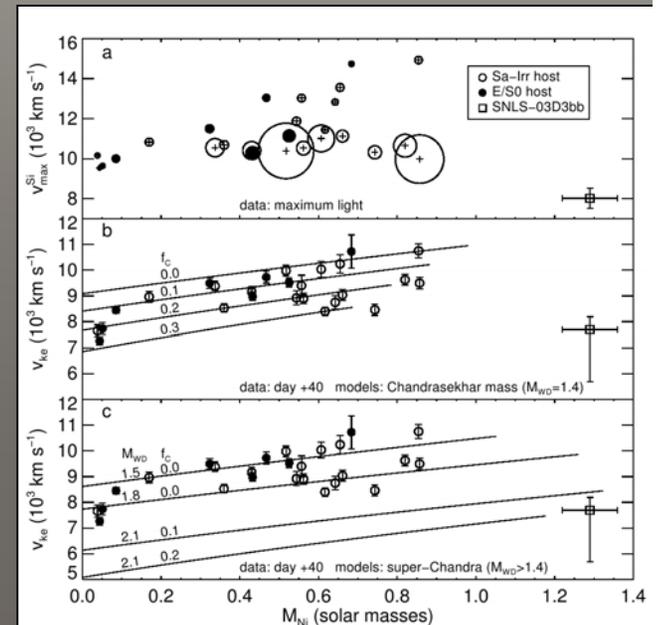
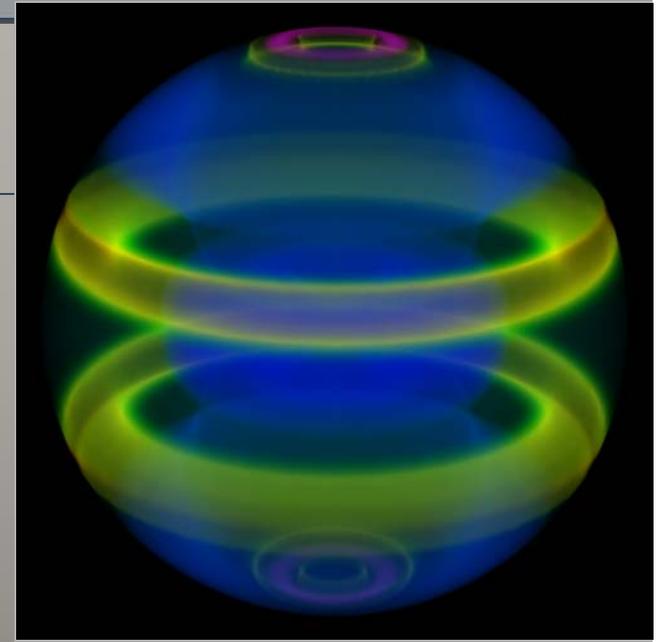
- Our ability to create and store information exceeds our capacity to understand it.





## Visualization Use Models

- Presentation visualization
  - You know what's there and want to show it to someone else
- Analytical Visualization
  - You know what you are looking for
- Discovery Visualization
  - You have no idea what you're looking for





# VACET

## VACET Mission and Vision

- SciDAC Visualization and Analytics Center for Enabling Technology (VACET)
- *Mission:* Leverage sci-vis and analytics software technology as an enabling technology for enabling scientific insight.
- *Vision:* adapt, extend, create, and deploy data understanding technologies for science stakeholders to enable petascale science.
- As a center, well positioned to respond to diverse needs/objectives through coordinated R&D, software engineering, outreach efforts.



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

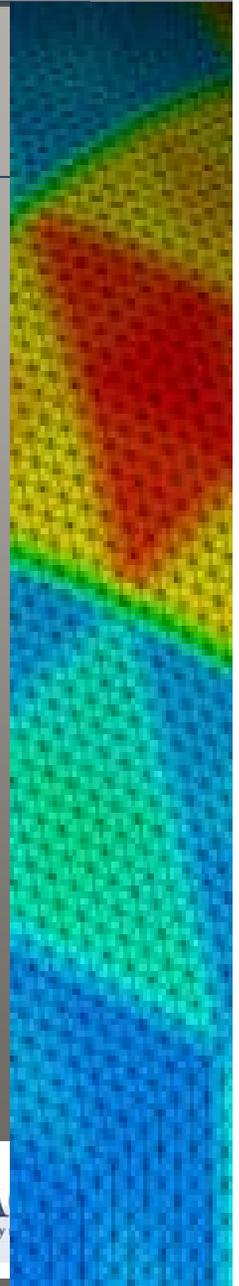


SciDAC  
Scientific Discovery through



## Discussion Topics

- Effective use of petascale platforms
  - First-ever runs on one- and two-trillion zone problems on petascale class machines.
- Effective knowledge discovery
  - Less is more: query-driven visualization and accelerator modeling
- VACET accomplishments: first 2.5 years





## Trillion-zone experiment: objectives

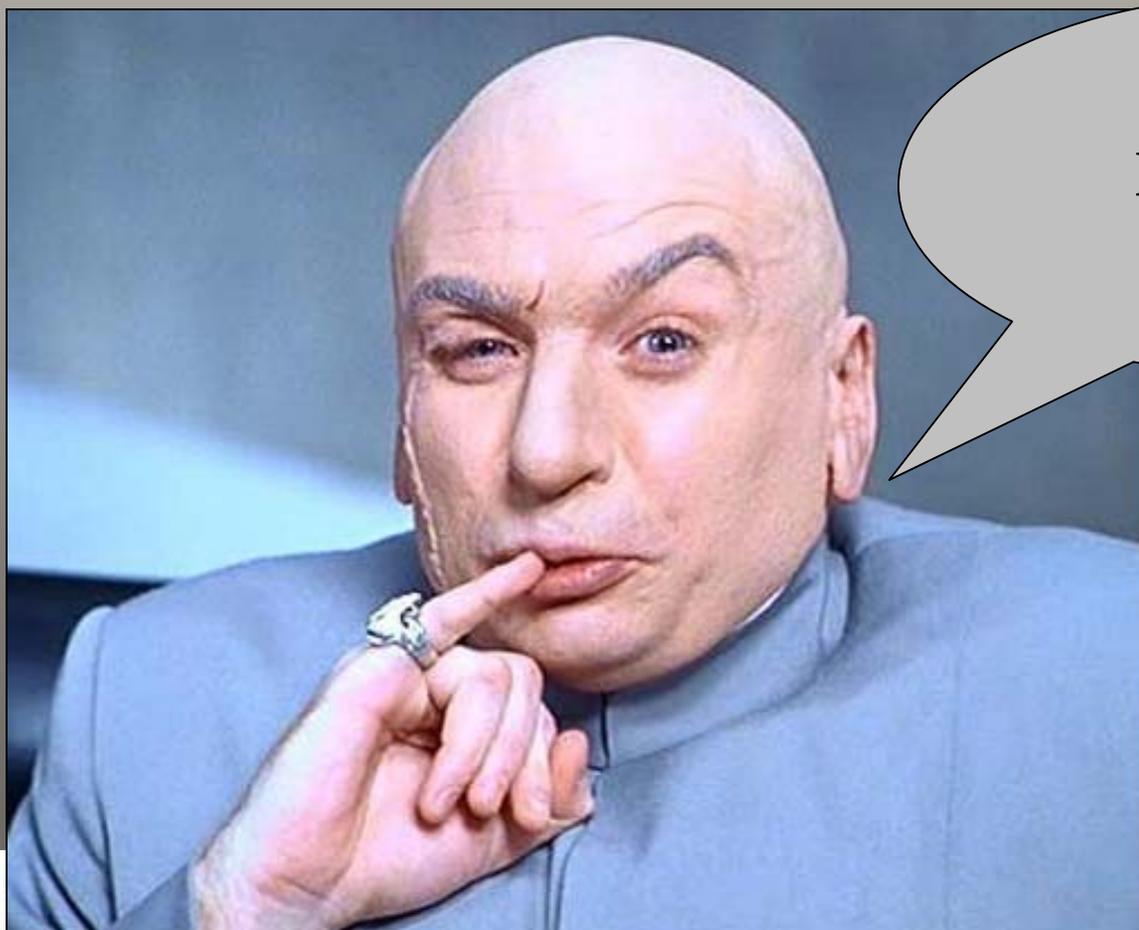
- Research questions:
  - Is it possible/feasible to run production-quality visual data analysis s/w on large machines and on large datasets?
  - What obstacles/bottlenecks do we encounter at this level of scale?
- Why?
  - Leverage I/O capacity, memory footprint of largest machines.
  - Share many CS research challenges with extreme scale computational science apps.





# VACET

## Is One Trillion Cells Big?



**One  
MILLION  
Cells!**

**1981: Cray-1,  
256MB RAM**



**ENERGY** | Science

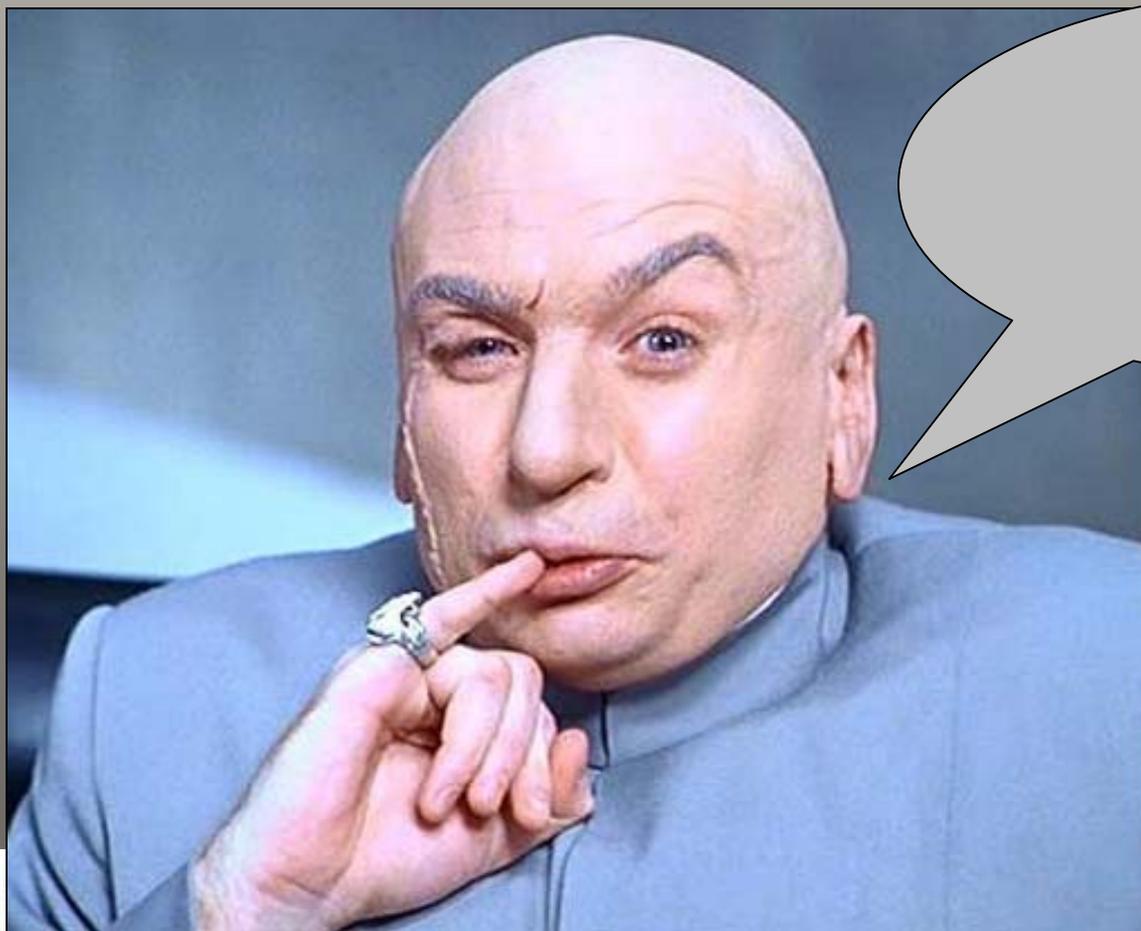


**SciDAC**  
Scientific Discovery through Advanced Computing



# VACET

## Is One Trillion Cells Big?



**One  
BILLION  
Cells!**

**1997: Cray T3E,  
128GB RAM**



**ENERGY** | Science

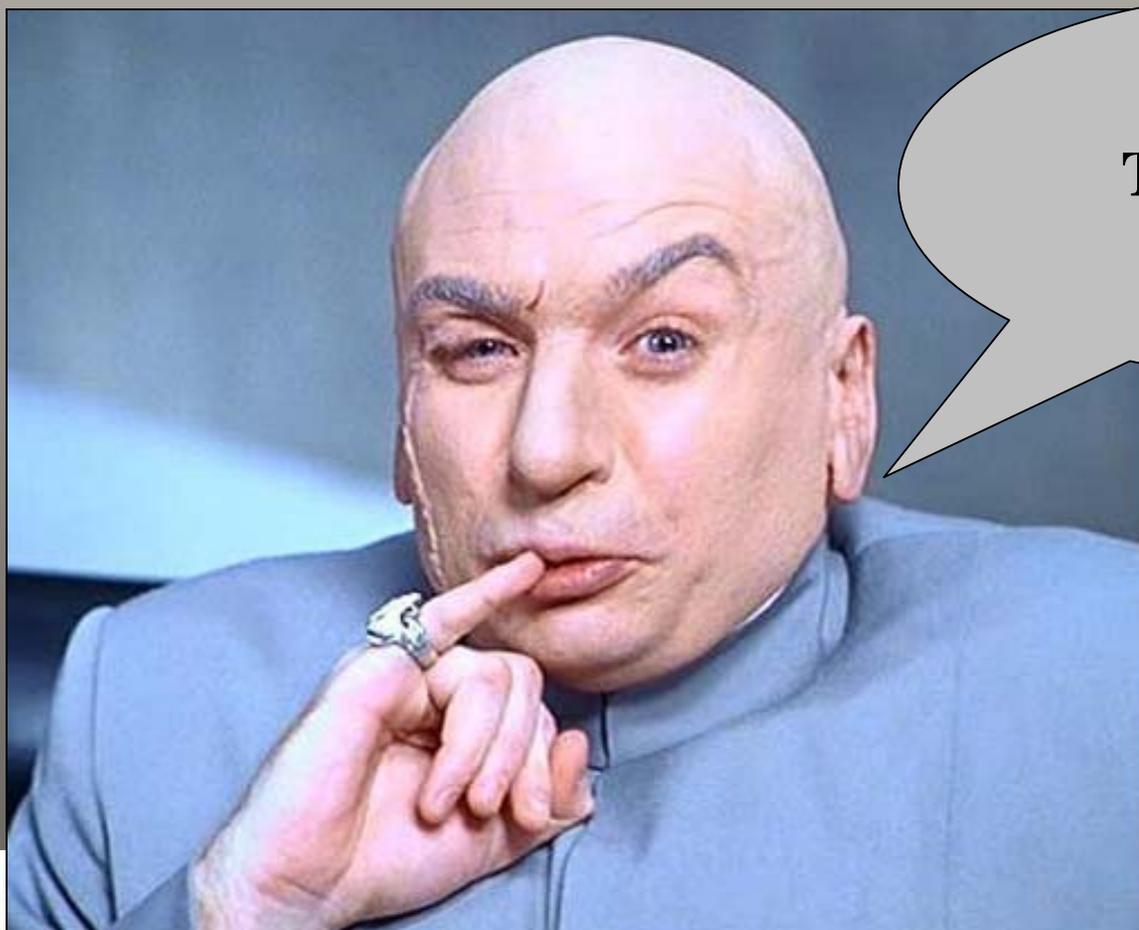


**SciDAC**  
Scientific Discovery through Advanced Computing



# VACET

## Is One Trillion Cells Big?



**One  
TRILLION  
Cells!**

**2009: Cray XT4:  
78 TB RAM**



**ENERGY** | Science



**SciDAC**  
Scientific Discovery through Advanced Computing



## Trillion-zone experiment: methodology

- Weak scaling study: ~62.5M zones/core
- Several machines:

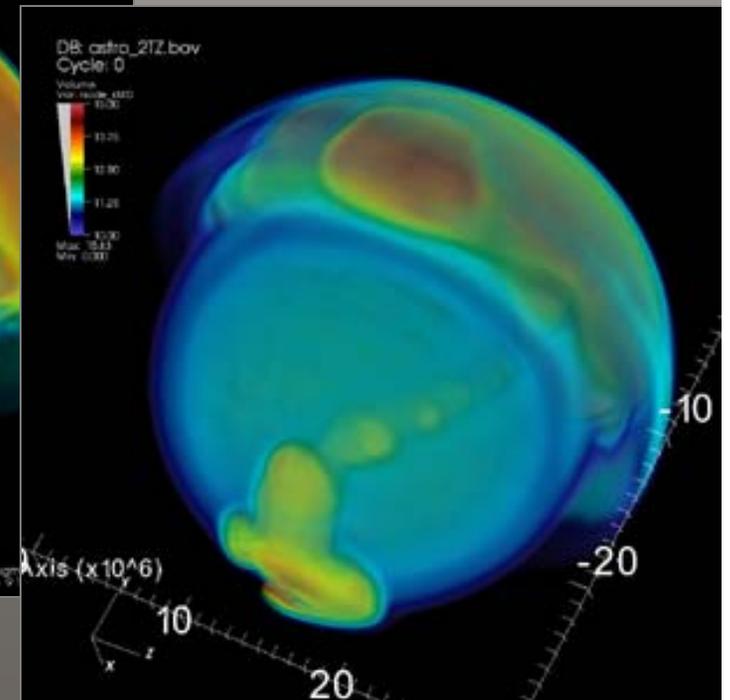
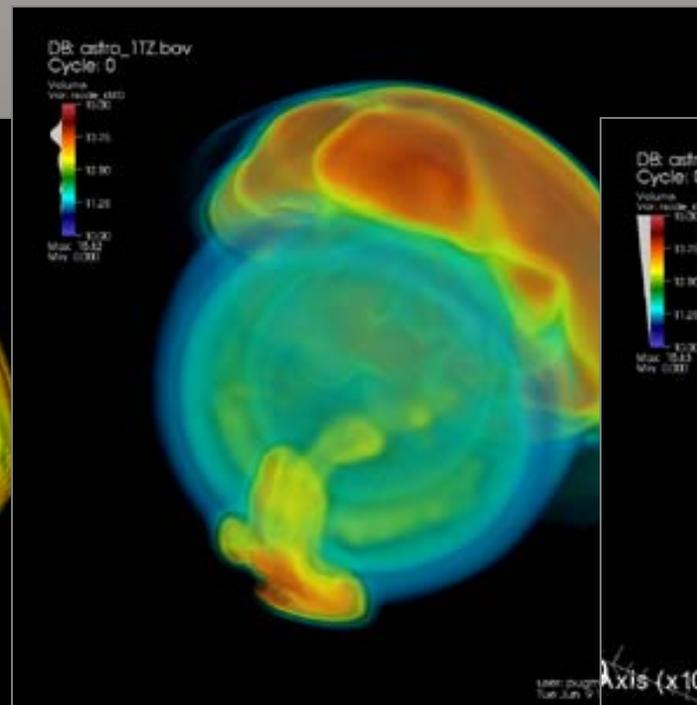
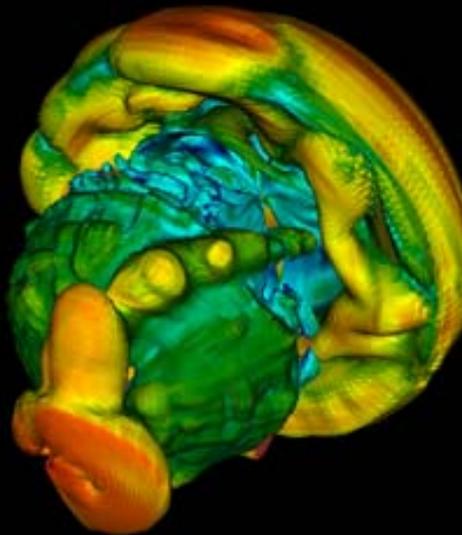
Machine	Model	Problem Size	#cores
Franklin	Cray XT4	1TZ, 2TZ	16K, 32K
Jaguar	Cray XT4	1TZ	16K
JaguarPF	Cray XT5	2TZ	32K
Juno	X86_64	1TZ	16K
Purple	IBM P5	0.5TZ	8K
Ranger	Sun	1TZ	16K





## TZ experiment methodology (ctd)

- Two common visualization techniques:
  - Isocontouring, volume rendering





## TZ experiment methodology (ctd)

- Application: VisIt 1.11.1 
  - VisIt is the first and only visualization application that is part of the “Joule” metric

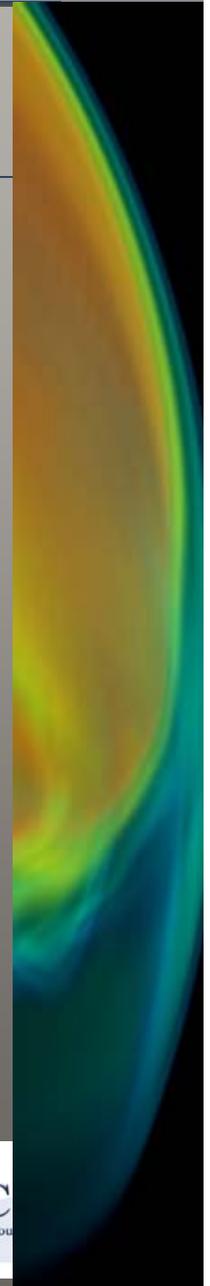
- Data:

- 512<sup>3</sup> sample data, astro simulation (courtesy T. Mezzacappa, B. Messer, S. Bruenn, R. Budjiara)



- Upsampled to target problem size

- Performance results largely independent of data content
- Experiment objectives met using this approach
- Representative of “tomorrow’s dataset sizes”





## TZ experiment “speed bumps”

- Application bottleneck: many-to-one status update
  - Workaround for study, fix to appear in future release of VisIt
- Application optimization: an NxN buffer for communication, which works well at small scale, proved problematic at large scale.
  - Workaround: remove optimization, evaluate best course of action for future VisIt release.
- System issues
  - Loading shared libraries takes a long time, as much as 5 minutes. Potential fix: implement VisIt plugins as static libraries.





## TZ experiment: Significance of Results

- VisIt runs effectively at high concurrency and on large problem sizes.
  - Test of visual data analysis infrastructure
  - Infrastructure is the underpinning of many different algorithms, not just isocontouring or volume rendering.
- Minor fixes to become part of production VisIt release (later this year).
- Successfully processed “tomorrow’s datasets today.”





## TZ experiment team



Sean Ahern, Dave Pugmire (ORNL)



Mark Howison, Prabhat (LBNL)



Hank Childs, LBNL (nee LLNL)



NERSC “special assistance”

– Katie Antypas





## TZ experiment: Acknowledgments

- Visit on Franklin (CNL): Dave Pugmire & Sean Ahern (ORNL), Gunther Weber & Janet Jacobsen & Prabhat (LBNL/NERSC Analytics) & NERSC Consultants
  - This pathfinding effort spanned many months last year.
- For access to facilities, help along the way:
  - LBNL/NERSC: Nick Cardo, Francesca Verdier, Howard Walter, Kathy Yelick
  - ORNL/OLCF: James Hack, Doug Kothe, Arthur Bland, Ricky Kendall
  - LLNL: David Fox, Debbie Santa Maria, Brian Carnes
  - UT Austin/TACC: Paul Navratil, Kelly Gaither, Karl Schulz





# VACET

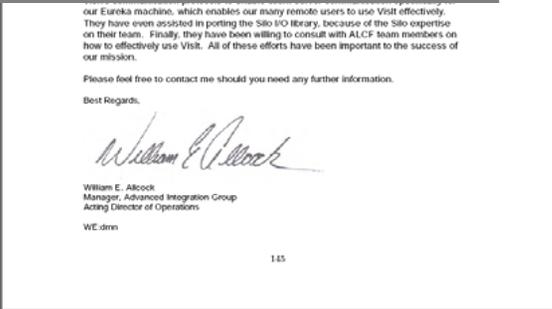
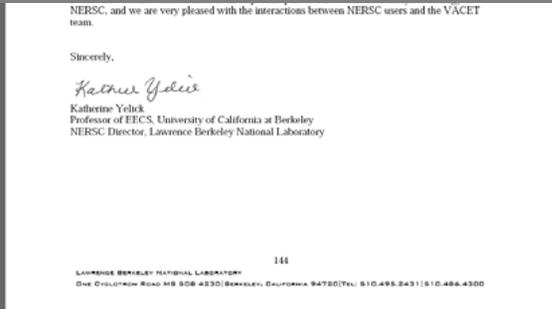
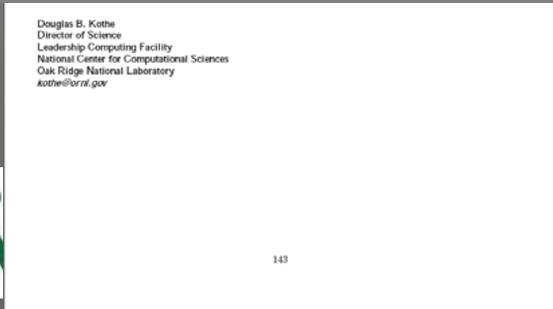
## We have ensured that our tools will run on the Office of Science's big machines



- Porting for Franklin (LBL), Jaguar (ORNL), and Eureka (Argonne, connected to BG/P)
- Changes to ensure job launching, client-server connections, I/O, etc run smoothly



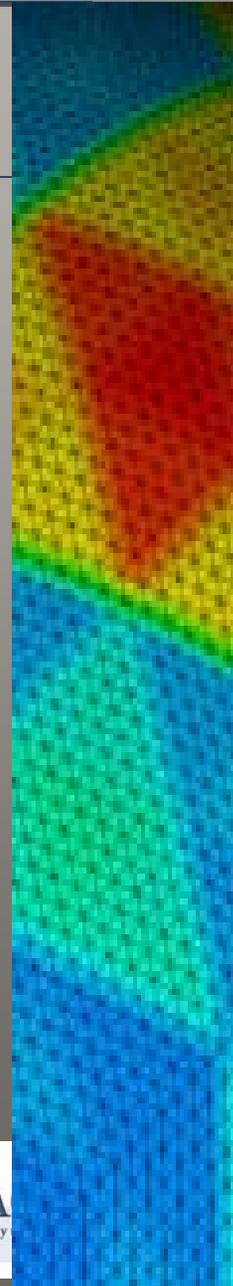
The SEG is ensuring that VACET tools are available to SciDAC stakeholders.





## Discussion Topics

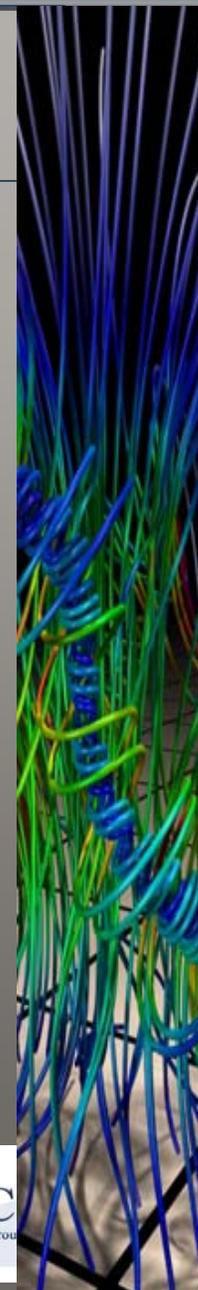
- Effective use of petascale platforms
  - First-ever runs on one- and two-trillion zone problems on petascale class machines.
- Effective knowledge discovery
  - Less is more: query-driven visualization and accelerator modeling
- VACET accomplishments





## Query-Driven Visualization

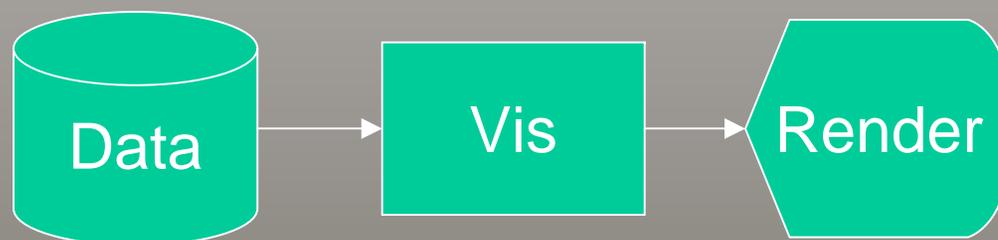
- What is Query-Driven Visualization?
  - Find “interesting data” and limit visualization, analysis, machine and cognitive processing to that subset.
- One way to define “interesting” is with compound boolean range queries.
  - E.g.,  $(CH_4 > 0.1)$  AND  $(T_1 < temp < T_2)$
- Quickly locate those data that are “interesting.”
- Pass results along to visualization and analysis pipeline.





# VACET

## Query-Driven Visualization



The Canonical Visualization Pipeline



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

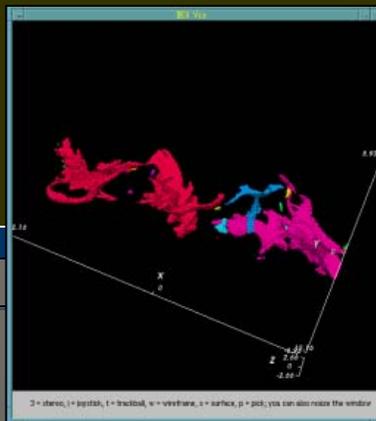
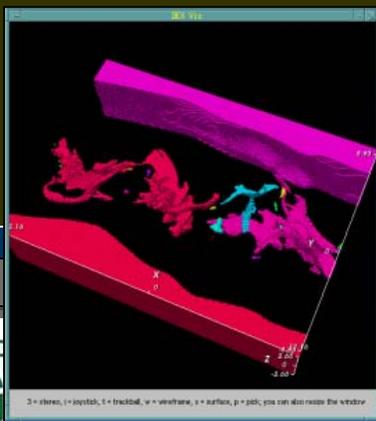
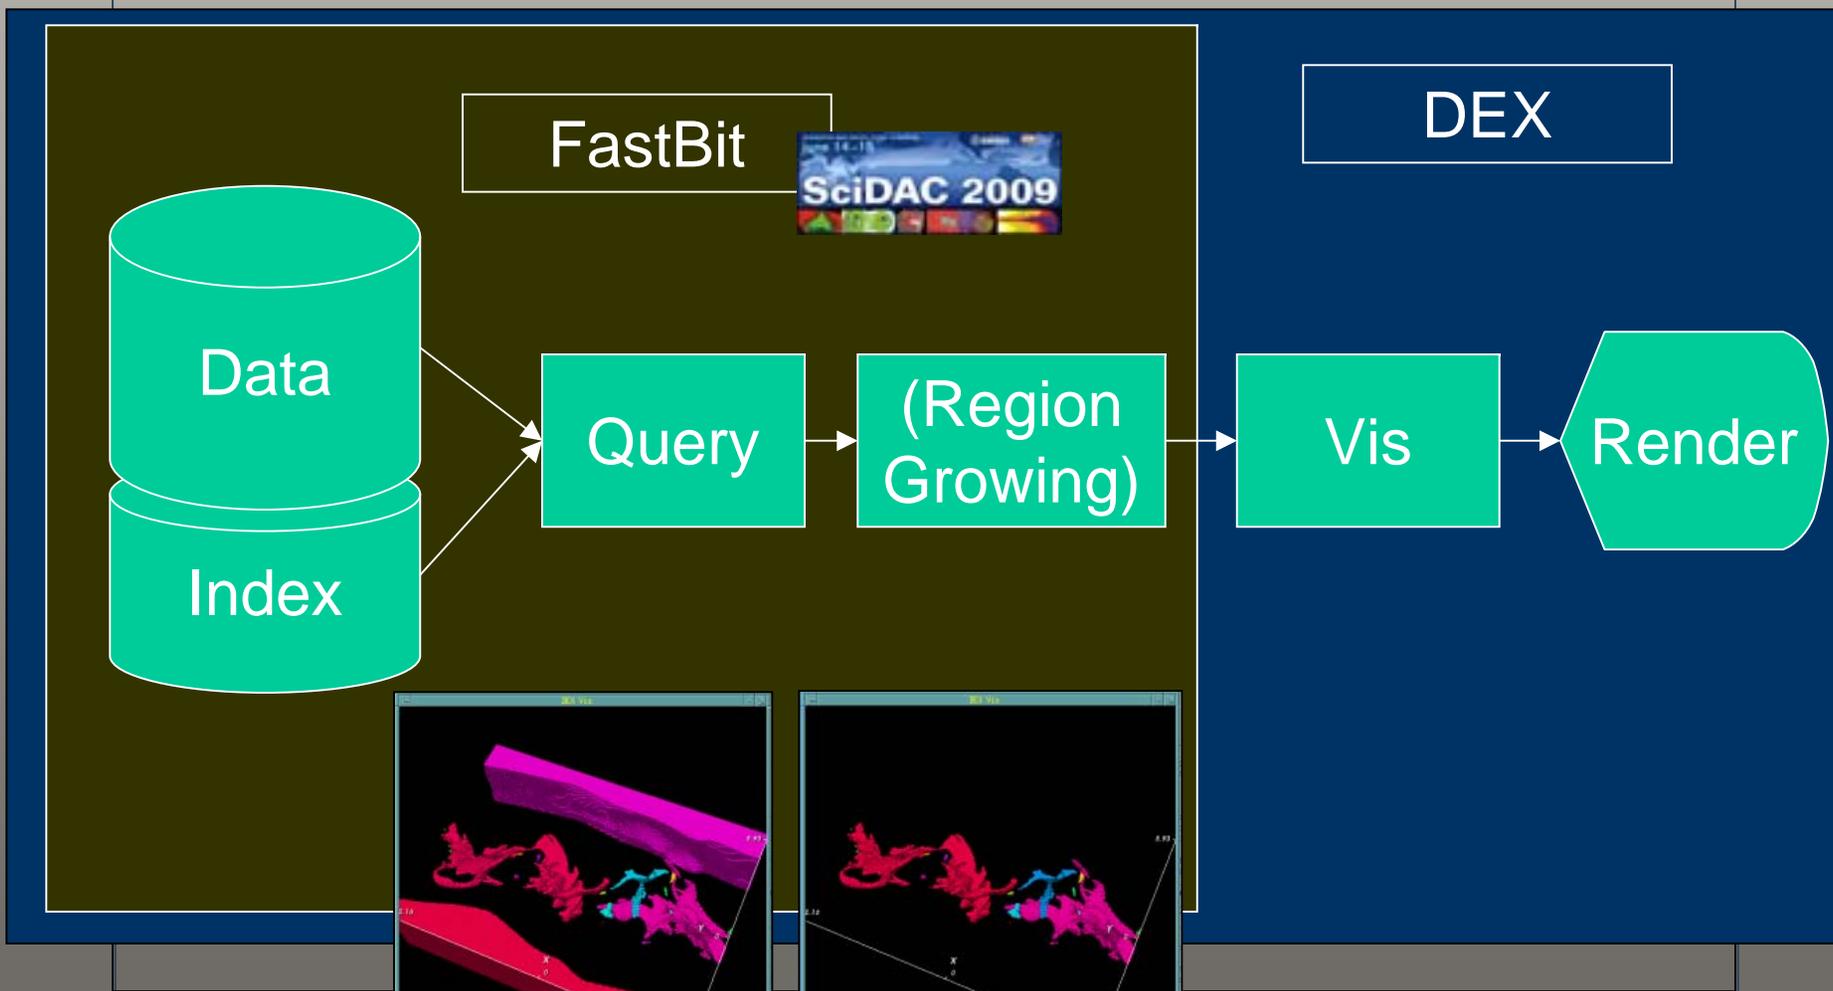


**SciDAC**

Scientific Discovery through Advanced Computing

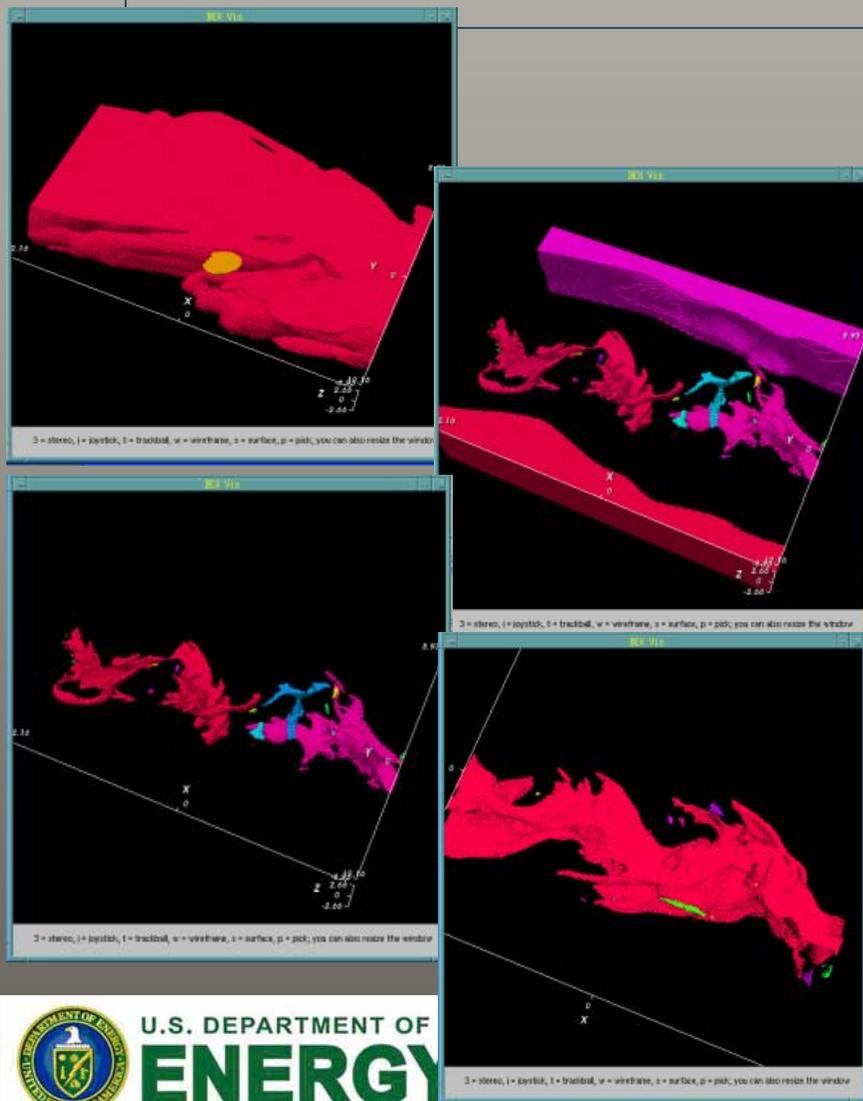


## Query-Driven Visualization





## Query-Driven Visualization



❖  $\text{CH}_4 > 0.3$

❖  $\text{Temp} < T_1$

❖  $\text{CH}_4 > 0.3$  AND  $\text{temp} < T_1$

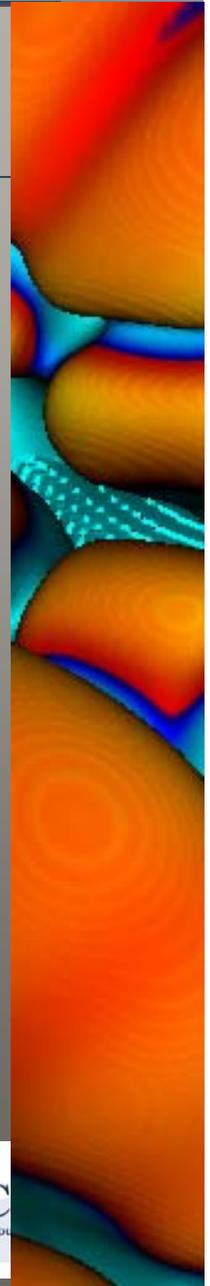
❖  $\text{CH}_4 > 0.3$  AND  $\text{temp} < T_2$   
▪  $T_1 < T_2$





## QDV Cybersecurity Case Study

- The next sequence of slides discusses application of the work to a cybersecurity application – evidence that the idea is generally applicable to large data visualization.
- The team:
  - NERSC Network Security
  - ESnet Network Engineers
  - Scientific Data Management Research
  - Visualization Research





## QDV – Detecting Distributed Scans

- The problem:
  - One day's worth of traffic consists of tens of millions of individual connections.
  - Traffic increasing by an order of magnitude every 48 months.
    - ESnet monthly traffic levels now exceed 1 PB.
  - The Internet is a hostile environment, and it will get worse.
  - Objective: enable rapid forensic data analysis (network flow records).





## QDV – Detecting Distributed Scans

- The data:
  - 42 weeks' of connection records from Bro (NERSC).
  - 281GB for raw data, 78GB for compressed bitmap indices.
- “Hero-sized problem”
  - No previous network analysis work has ever attempted to perform interactive visual analytics on data of this scale (ca. 2006).
  - Result: what once took days (if at all possible) now takes seconds.





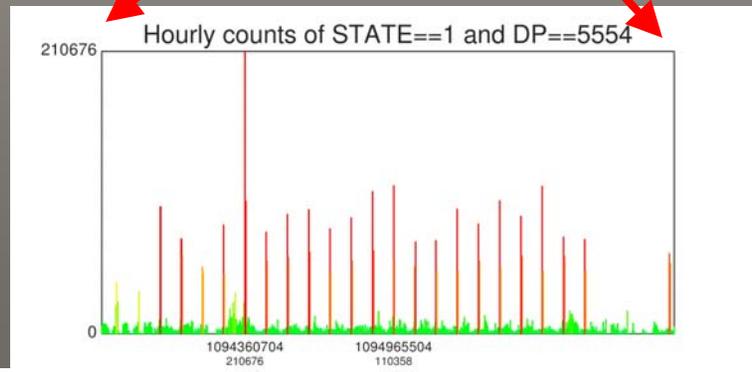
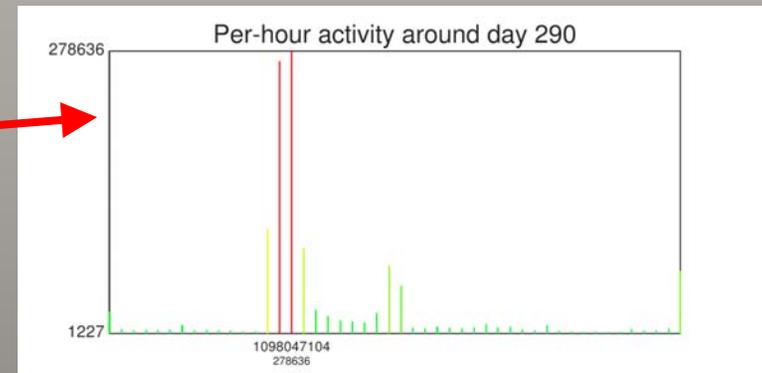
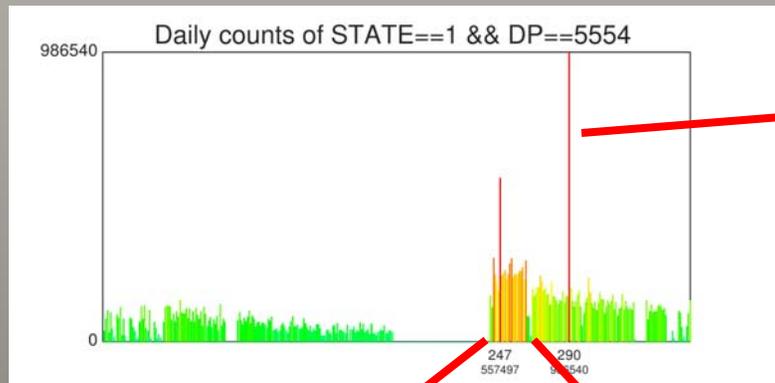
## QDV – Detecting Distributed Scans

- The starting point:
  - You are a network security analyst
  - Your beeper goes off [at lunch, in the shower...]
  - You receive an alert that “something odd is happening with the network...IDS showing unusual levels of activity on port 5554”
  - Your job – answer questions:
    - What’s going on now?
    - How long has this been happening?
    - Implications?





## QDV – Detecting Distributed Scans



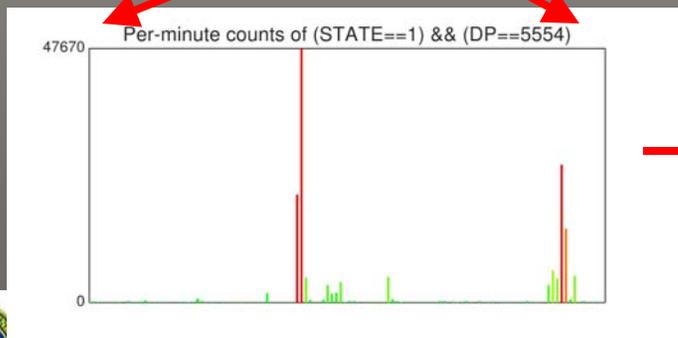
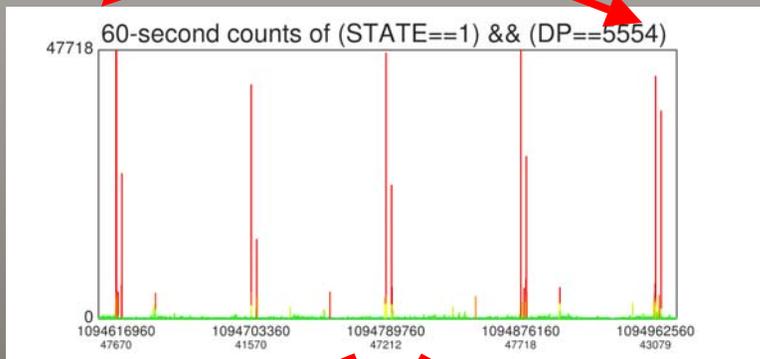
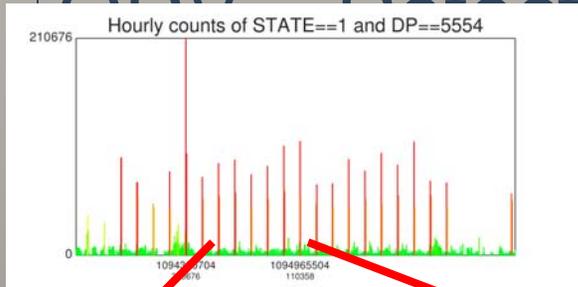
1. Query to produce a histogram of unsuccessful connection attempts over a 42-week period at one-day temporal resolution (upper left).
2. Drill into the data, query to produce a new histogram covering a four-week period at one-hour temporal resolution (lower left).
3. Generate a histogram of one-hour resolution over a two-day period around day 290 (upper right).





# VACET

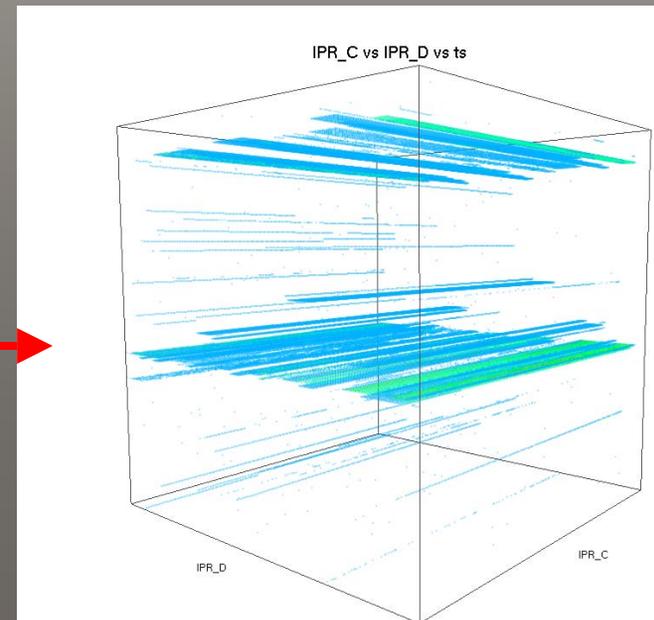
## DDV: Detecting Dist



5. Query to generate a histogram of unsuccessful connection attempts over a five-day period sampled at one-minute temporal resolution (middle, left). Regular attacks occur at 21:15L, followed by a second wave 50 minutes later.

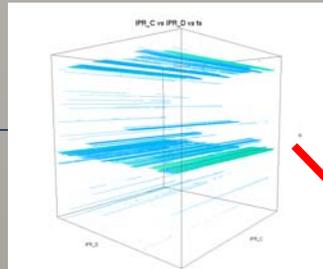
6. Query to generate histogram over a two-hour period at one-minute temporal resolution (lower left).

7. Query to generate a 3D histogram showing the coverage of attacks in destination address space (lower right).





## QDV –

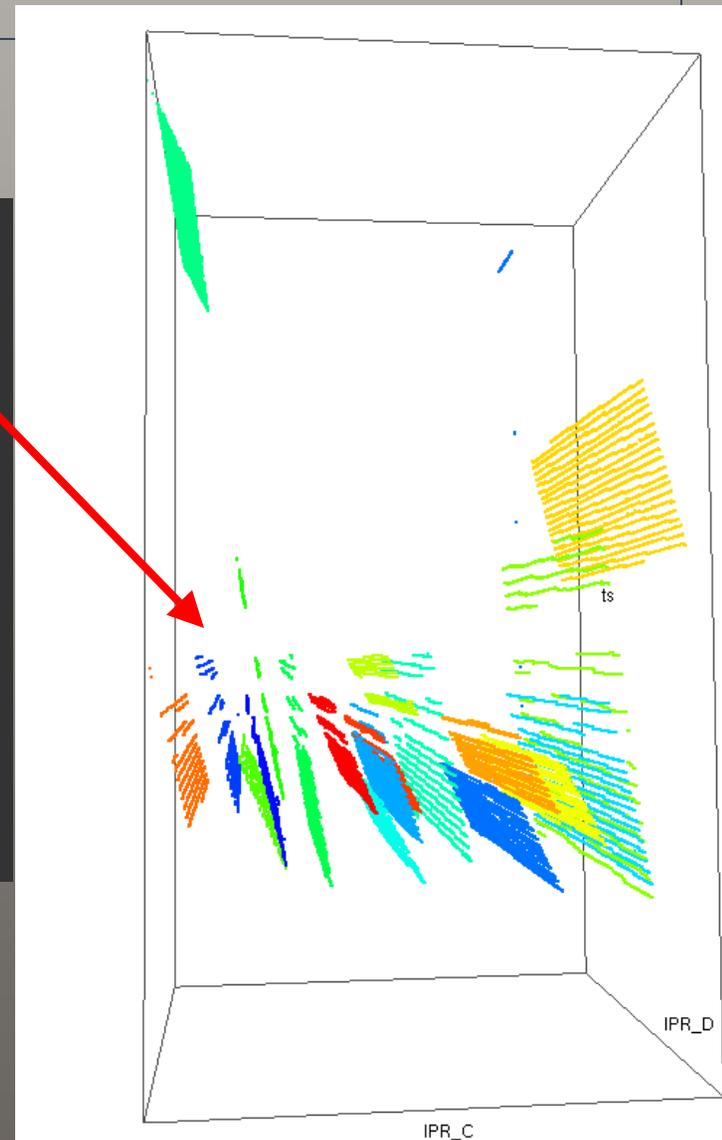


## ing Distributed Scans

After establishing that (1) a temporally regular activity is occurring, and (2) that it is in fact a systematic probe (scan) of entire blocks of network addresses, the next task is to determine the set of remote hosts participating in the attack.

Working backwards, we isolate the A, B, C and D address octets of the hosts participating in the attack.

This image shows a 3D histogram of the destination address space being attacked by each of 20 different hosts. The vertical axis is time – a seven-minute window at one-second temporal resolution.





## QDV – Detecting Distributed Scans

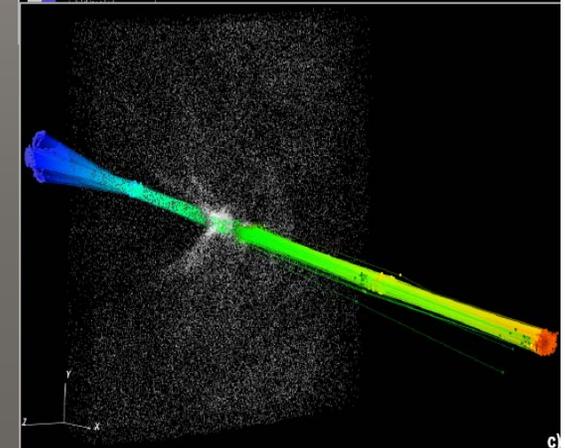
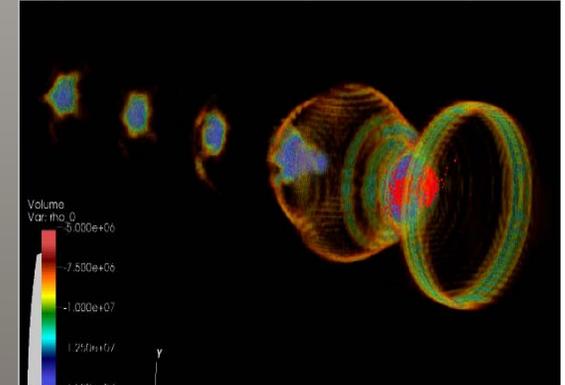
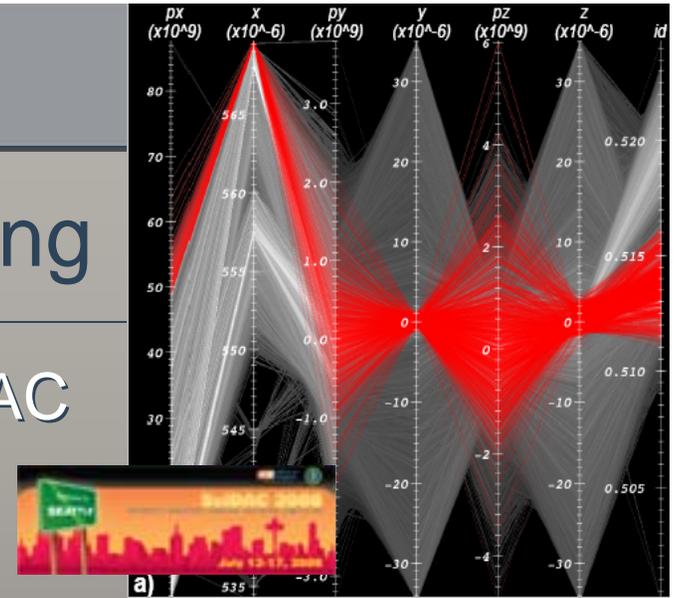
- Our analysis was performed in statistical space only.
  - We never accessed the raw data.
  - Our processing and visualization used only the index data.
  - Performance study focuses on parallel algorithms for multidimensional histogram computation from compressed bitmap indices.





## QDV and Accelerator Modeling

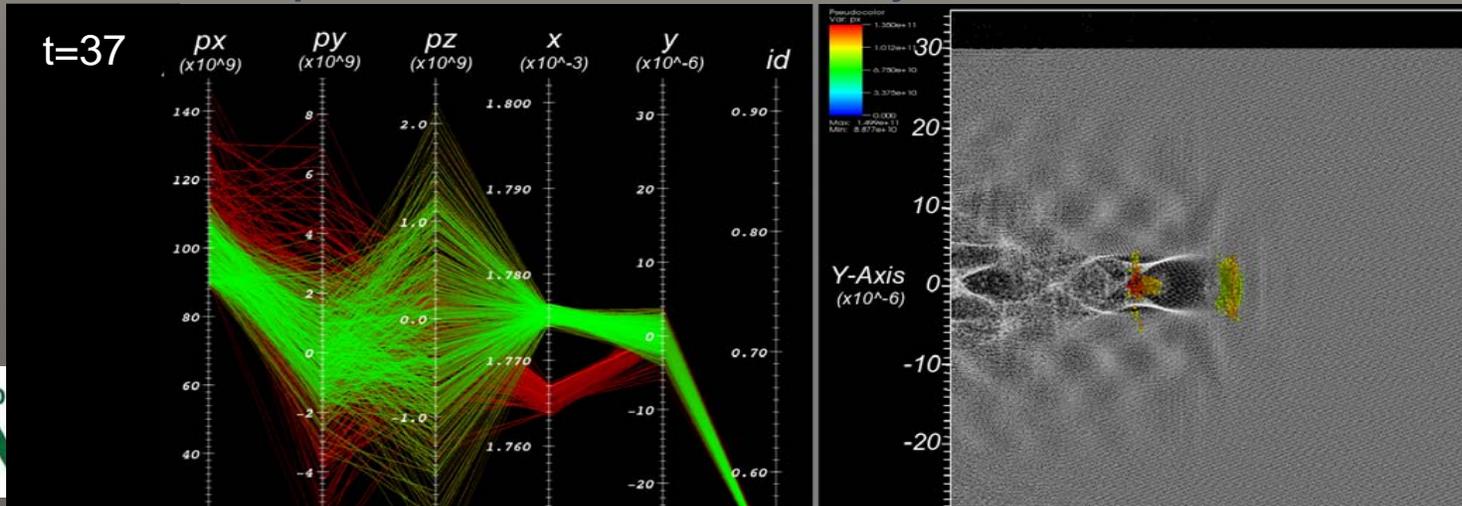
- PI: C. Geddes (LBNL), part of SciDAC COMPASS project, Incite awardee.
- Accomplishment:
  - Algorithms and production-quality s/w infrastructure to perform interactive visual data analysis (identify, track, analyze beam particles) in multi-TB simulation data.
- Science Impact:
  - Replace serial process that took hours with one that takes seconds.
  - New capability: rapid data exploration and analysis.
- Collaborators:
  - SciDAC SDM Center (FastBit)
  - Tech-X (Accelerator scientists)





## Analysis Task(s)

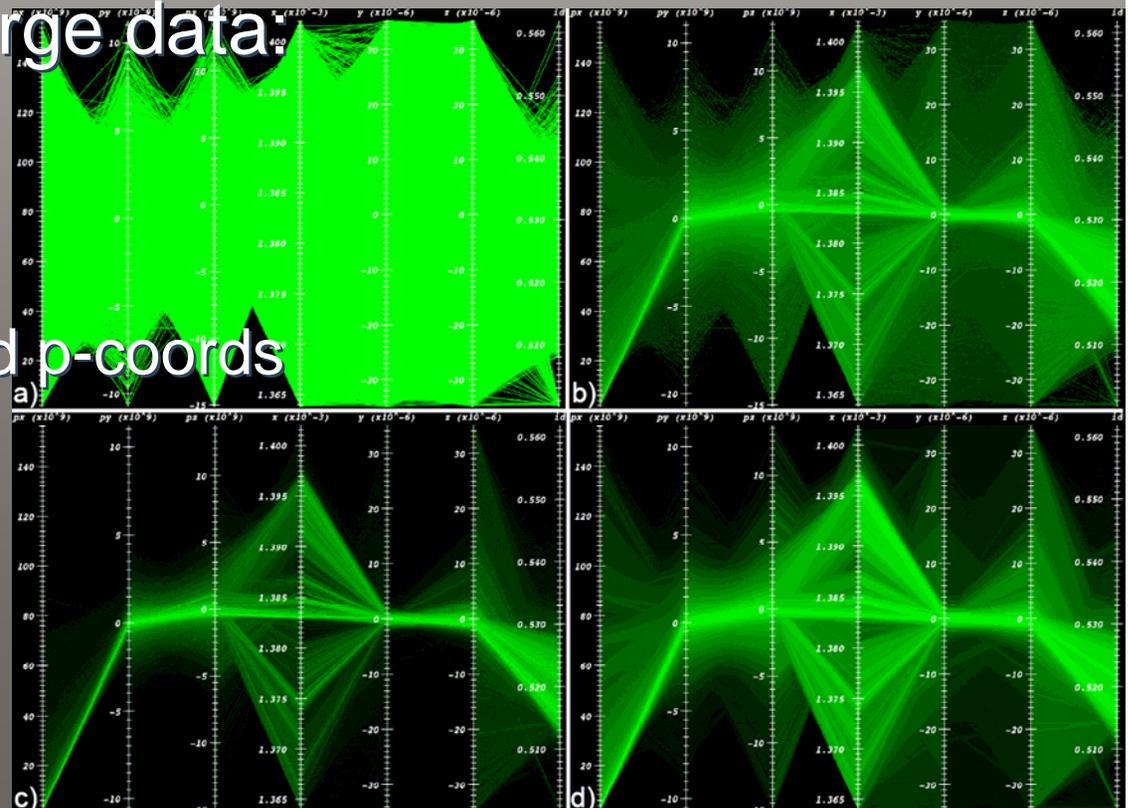
- Identify particles that form a beam
  - Interactive visual data exploration
  - Data subsetting
- Track them over time
  - Given particle ID's from a given time step,
  - Find all those particles in all time steps
  - Subsequent visual data analysis.





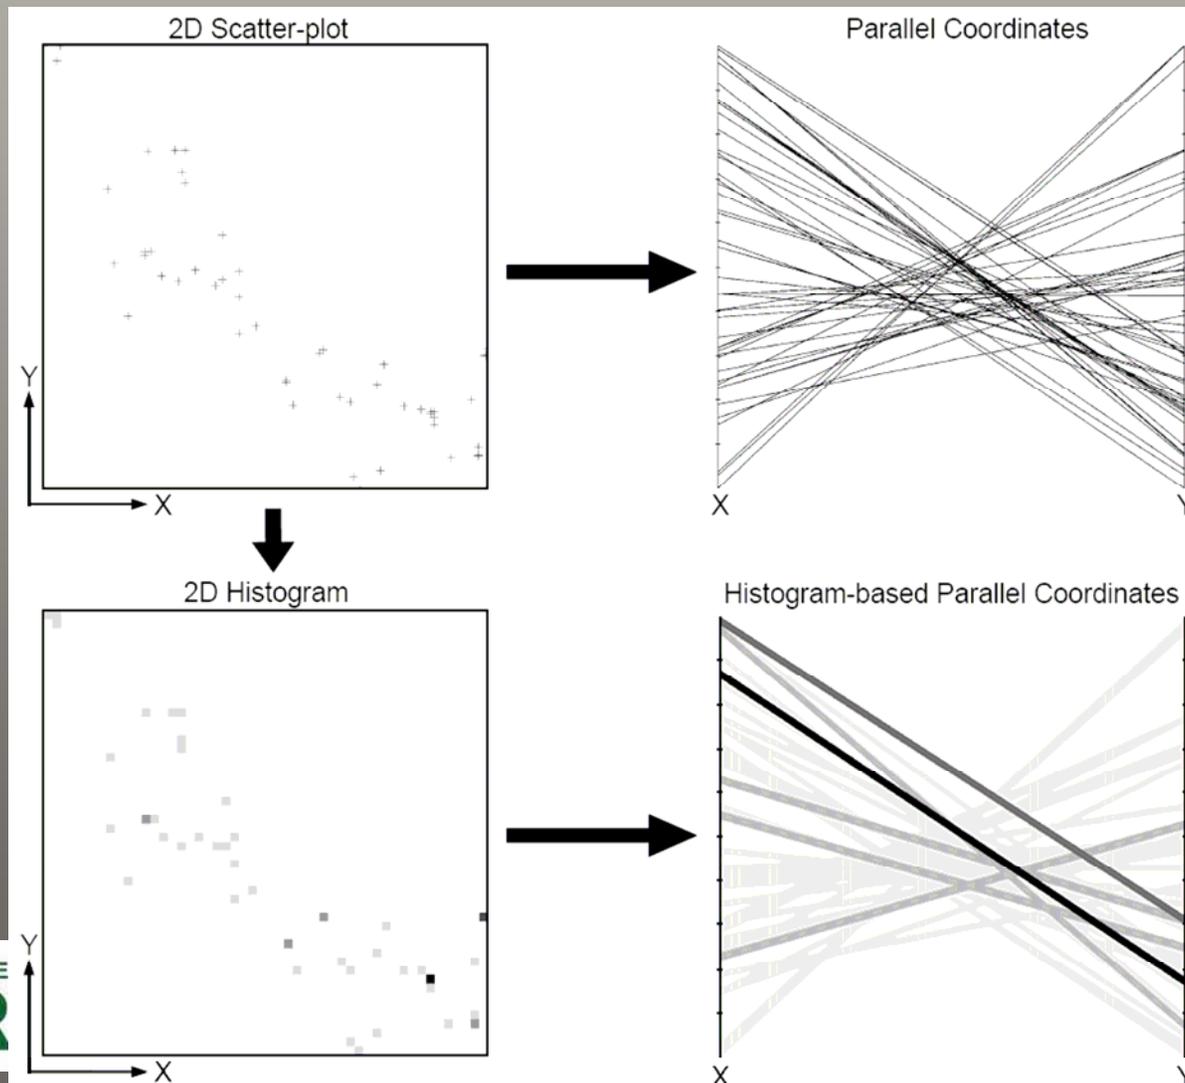
## Fundamental Problem #1 - Interface

- Parallel coordinates
  - An interface for subset selection.
  - A mechanism for displaying multivariate data.
- Problems with large data:
  - Visual clutter
  - $O(n)$  complexity
- Solution
  - Histogram-based p-coords



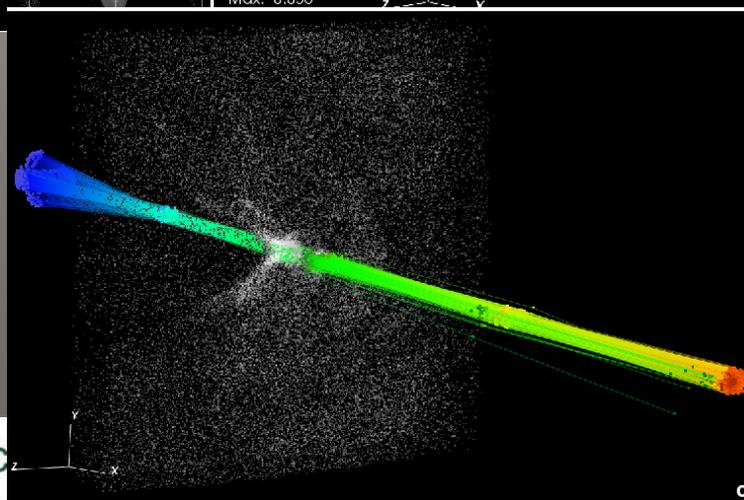
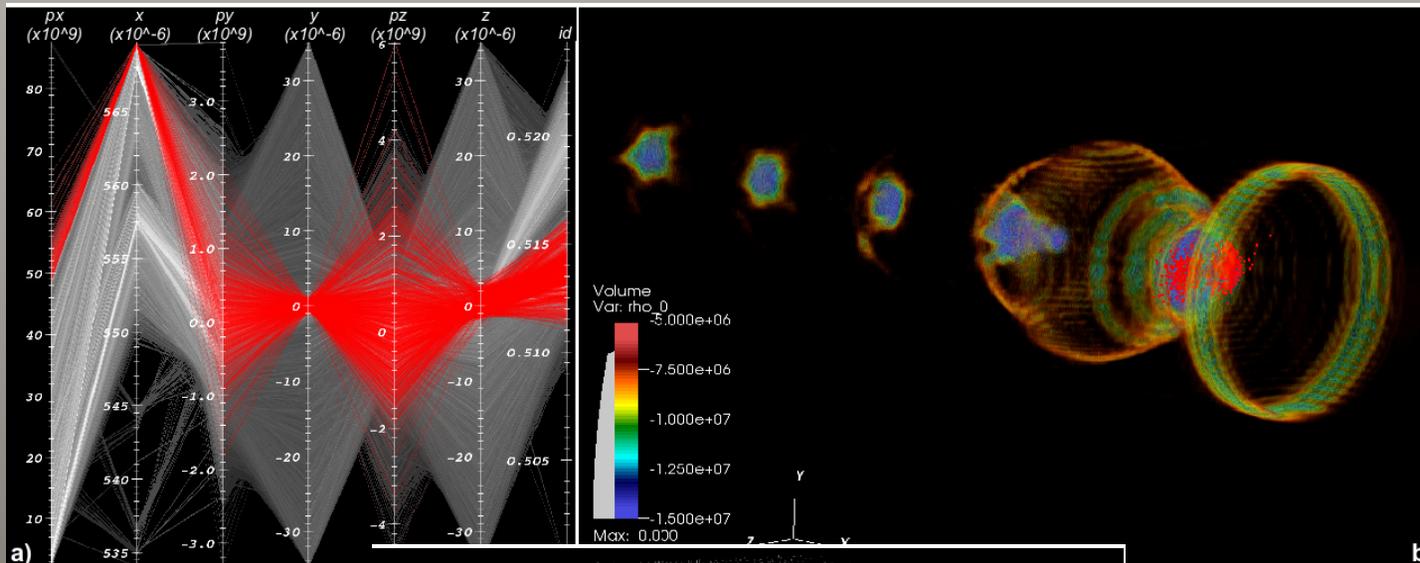


## Histogram-Based Parallel Coordinates





## 3D Example





## Fundamental Problem #2 – Performance

- How to efficiently construct a histogram?
  - Naïve approach:  $O(n)$
  - Better approach: “cheat” (use FastBit)
- How to efficiently do particle tracking?
  - Naïve approach:  $O(n^2)$
  - Better approach:  $O(H*t)$  (use FastBit)





### Dataset:

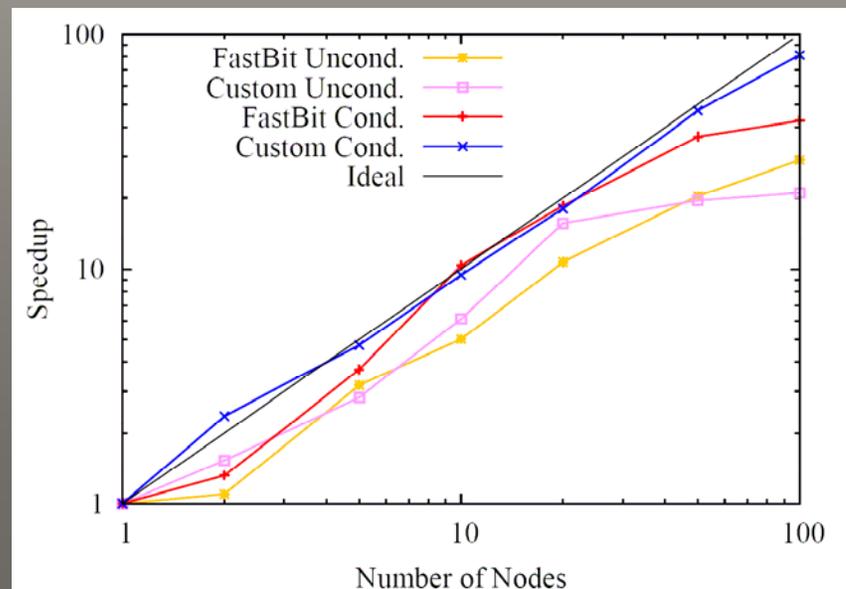
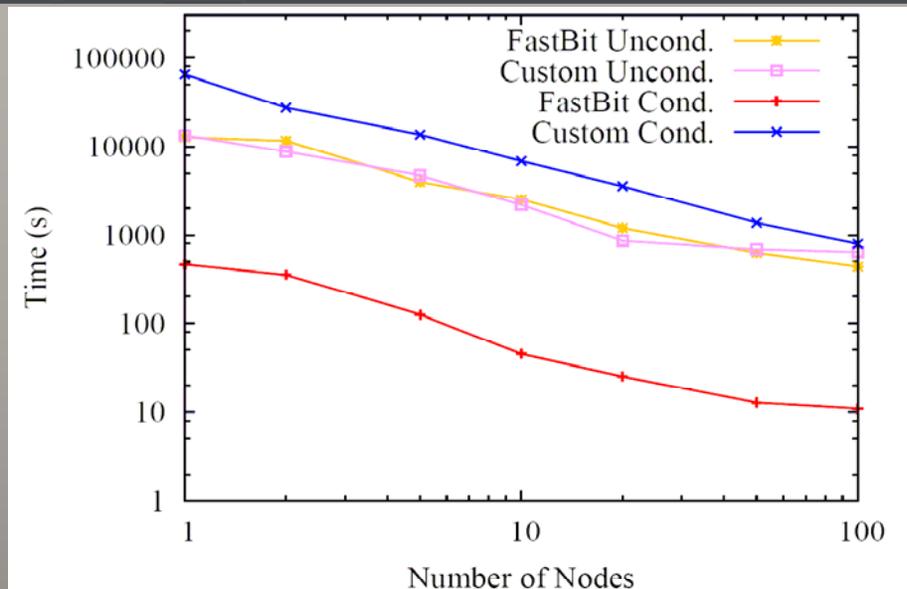
- 3D dataset consisting of 100 timesteps
- ~177 million particles per timestep
- ~10 GB per timestep
- ~1TB total size

### Test platform: (as of July.2008)

- franklin.nersc.gov
- 9,660 nodes, 19K cores Cray XT4 system
- Filesystem: Lustre Parallel Filesystem
- Each node consists of:
  - CPU: 2.6 GHz, dual-core AMD Opteron
  - Memory: 4GB
  - OS: Compute Node Linux

### Test setup:

- Restrict operations to a single core of each node to maximize I/O bandwidth available to each process
- Assign data subsets corresponding to individual timesteps to individual nodes for processing
- Generate five 1024x1024 histograms for position and momentum fields at each timestep
- Conditon:  $px > 7 * 10^{10}$
- Levels of parallelism: 1, 2, 5, 10, 20, 50, 100





### Test setup:

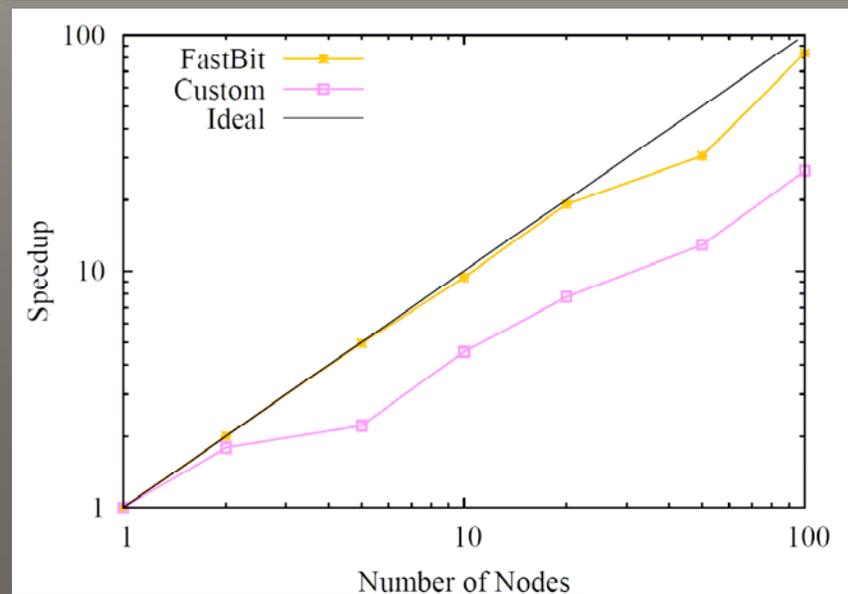
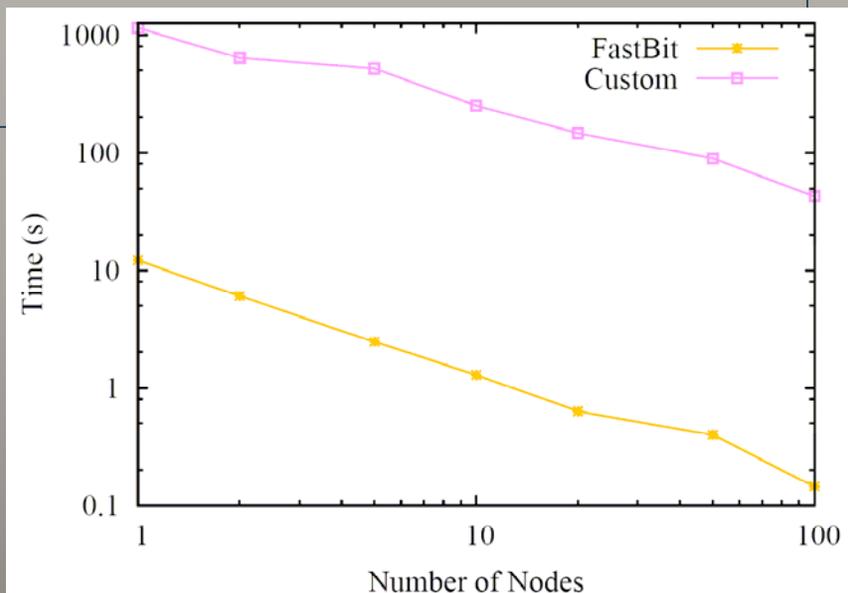
- Same as for histogram computation
- Track 500 particles (Condition:  $p_x > 10^{11}$ ) over 100 timesteps

### Results:

- FastBit is able to track 500 particles over 1.5TB of data in 0.15 seconds

### Performance of original IDL scripts:

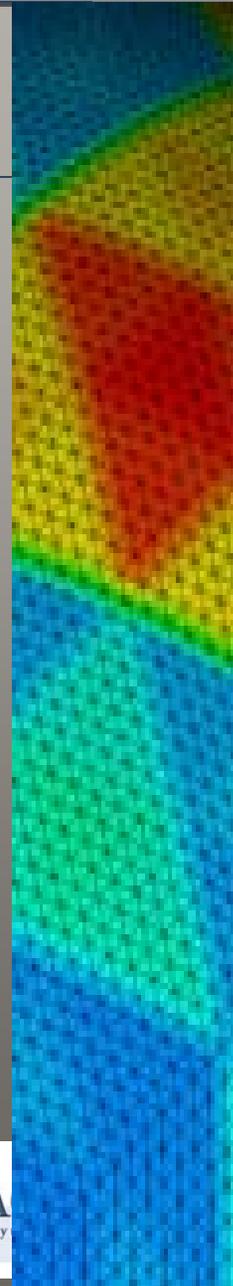
- ~2.5 hours to track 250 particles in small 5GB dataset





## Discussion Topics

- Effective use of petascale platforms
  - First-ever runs on one- and two-trillion zone problems on petascale class machines.
- Effective knowledge discovery
  - Less is more: query-driven visualization and accelerator modeling
- VACET accomplishments





# VACET

## VACET Accomplishments

- Flow visualization research
- AMR Visualization
- Topological analysis research
- Climate modeling visual data analysis
- High quality and parallel volume rendering



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science



**SciDAC**  
Scientific Discovery through

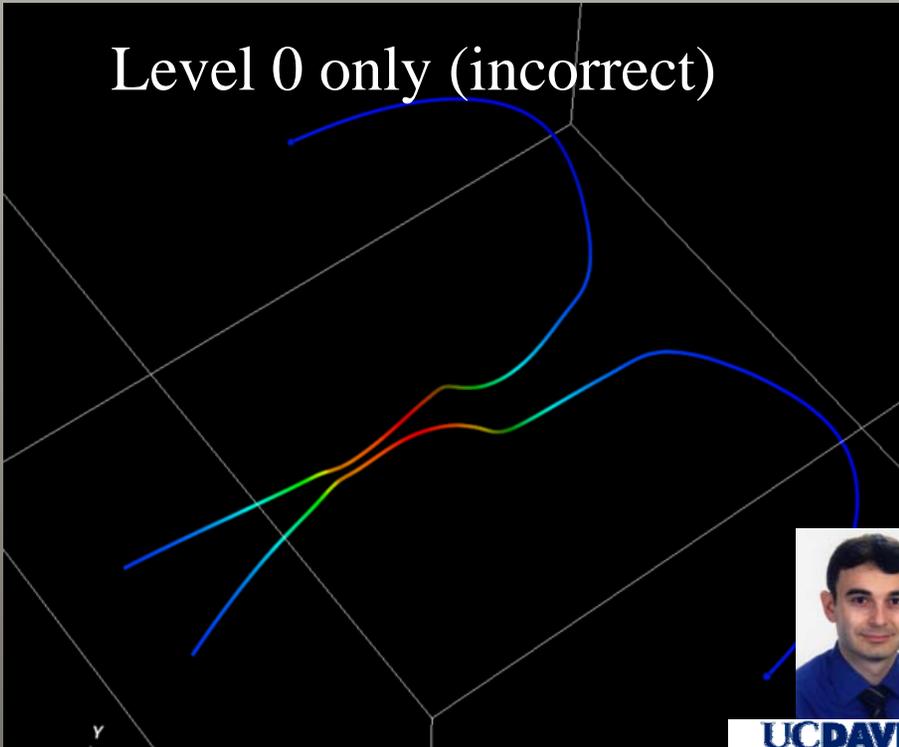




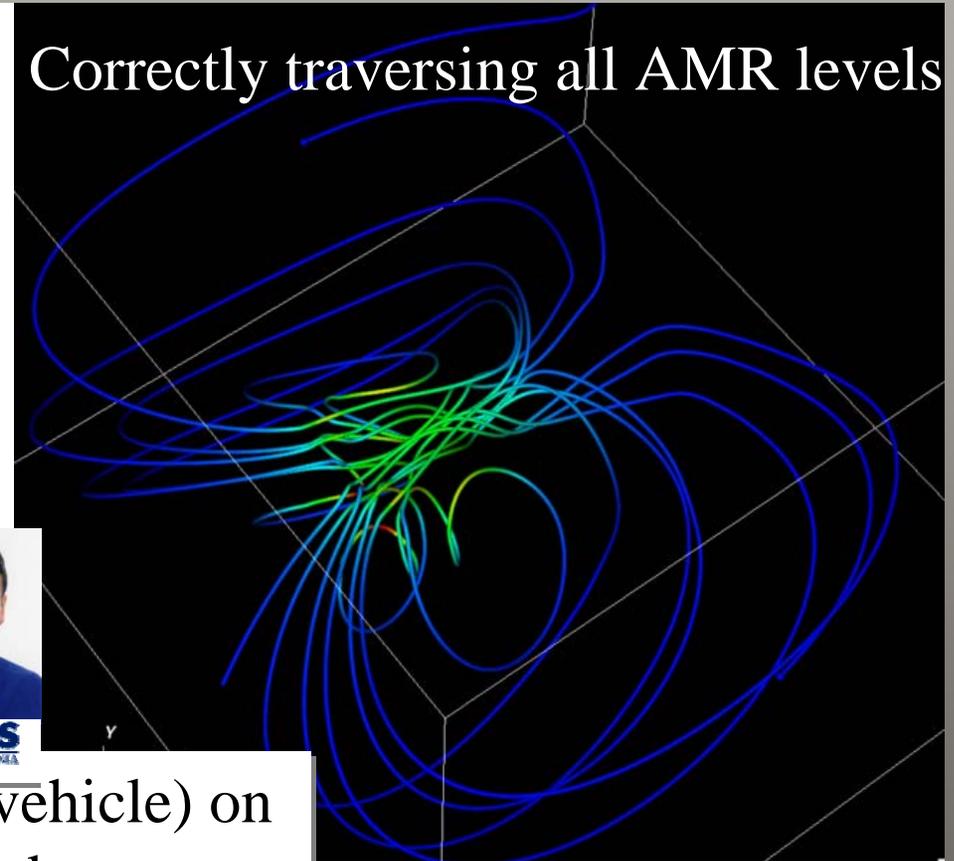


## Flow Visualization

Level 0 only (incorrect)



Correctly traversing all AMR levels



UC DAVIS  
UNIVERSITY OF CALIFORNIA

Research vehicle (and delivery vehicle) on streamline generation of AMR data sets.



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science



**SciDAC**  
Scientific Discovery through Advanced Computing



# VACET

## Flow Visualization, ctd.

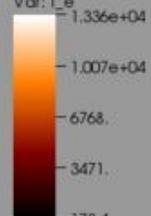




## Flow Visualization, ctd.

DB: phi\_BT.h5

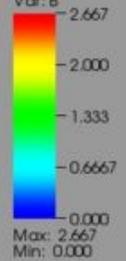
Pseudocolor  
Var: T\_e



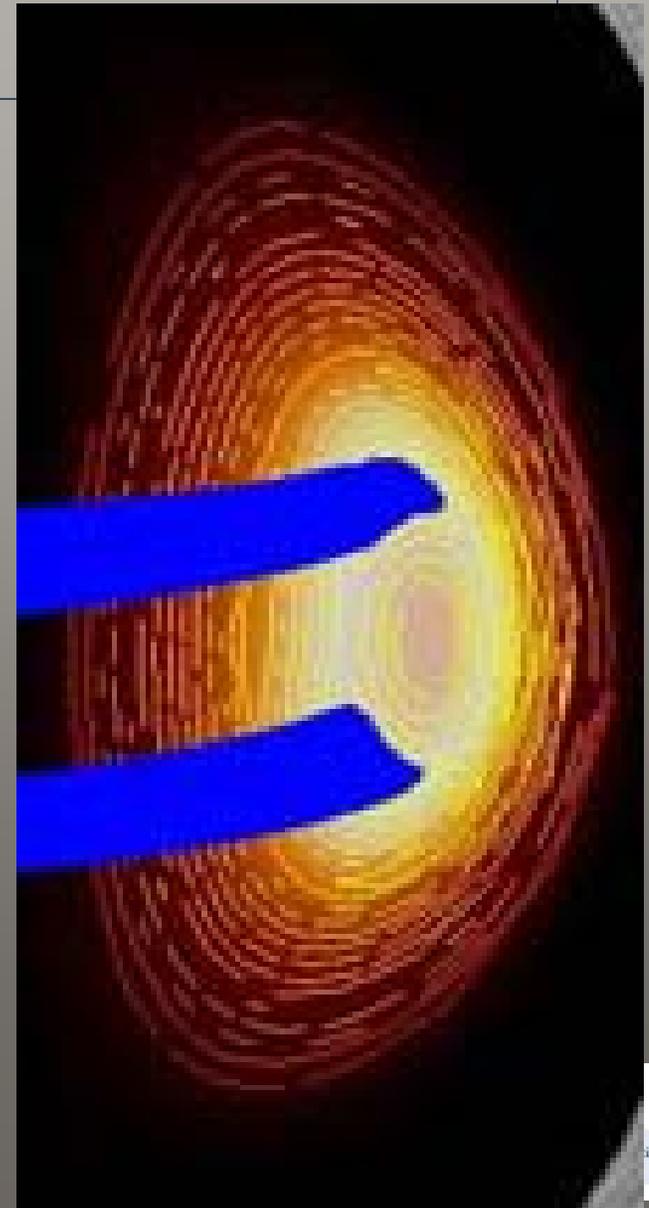
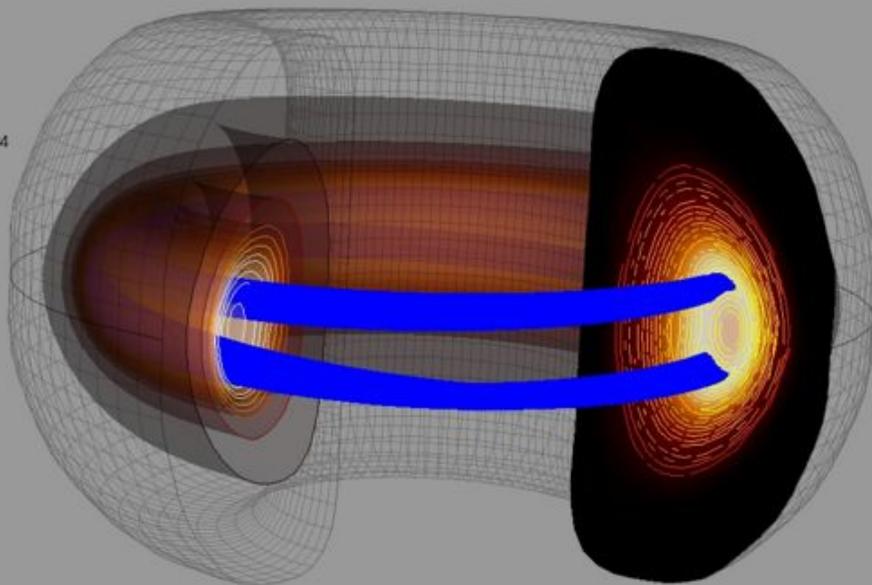
Max: 1.336e+04  
Min: 173.4

Mesh  
Var: Mesh

Poincare  
Var: B

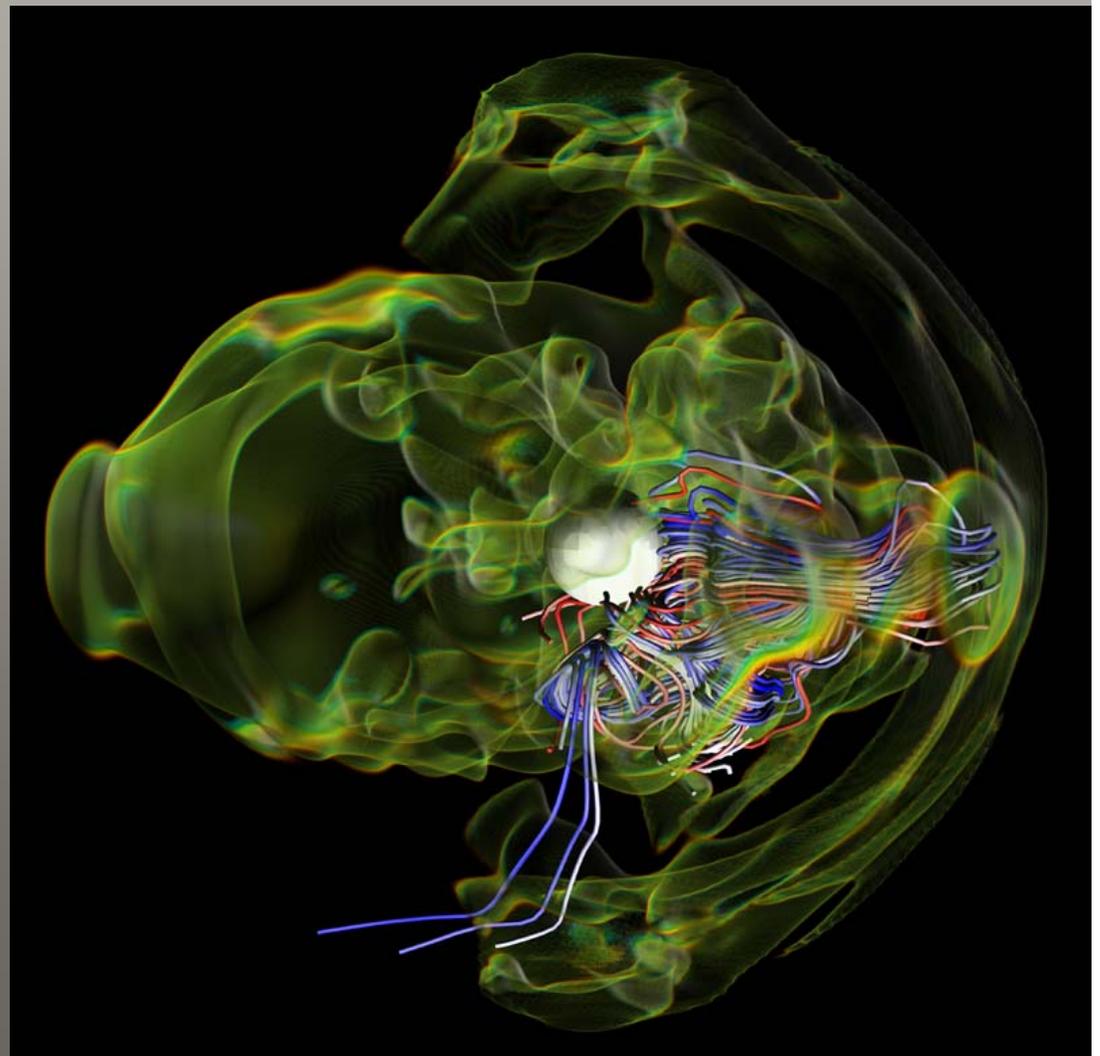
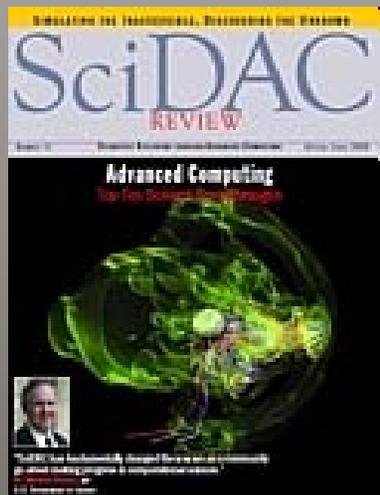


Max: 2.667  
Min: 0.000





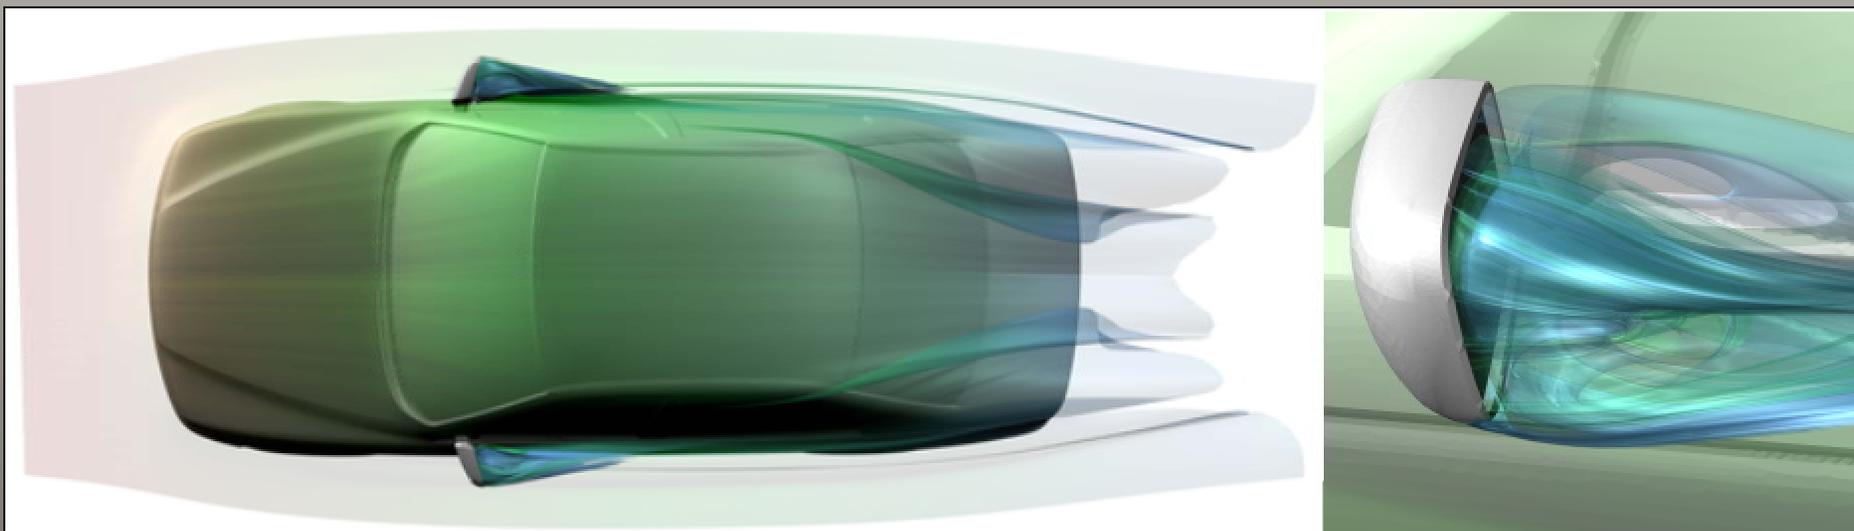
## Flow Visualization, ctd.





# VACET

## Flow Visualization, ctd.



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science



**SciDAC**

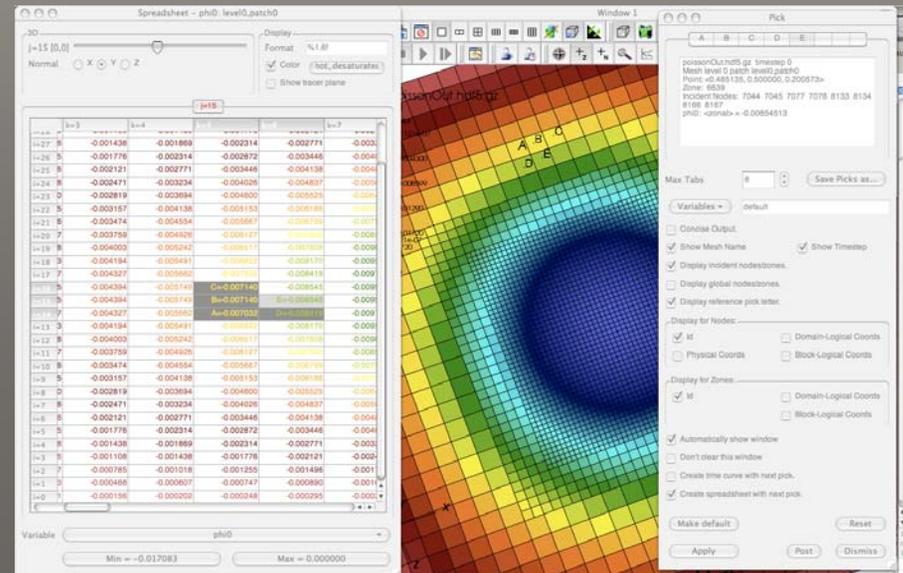
Scientific Discovery through Advanced Computing





## AMR Visualization

- Improvements:
  - Optimized many, many routines, resulting in end-to-end speedups of  $>10X$
  - Added critical debugging functionality, such as spreadsheets and the ability to connect with a debugger.
  - Mapped AMR grids
- Done for collaborators at APDEC & CAC.
  - Macros and python callbacks for customizability.
  - Responded to over 30 requests for interface changes and bug fixes.





## AMR Visualization

- APDEC adopts VisIt as their project-wide visual data analysis application.

**ChomboVis - Mozilla Firefox**

File Edit View History Bookmarks Tools Help

http://seesar.lbl.gov/anag/chombo/chombovis.html

Most Visited Getting Started Latest Headlines

### ChomboVis

[VisIt](#) has superseded ChomboVis as the visualization and analysis tool of choice to Chombo (AMR) HDF5 datasets. ChomboVis is no longer being maintained or developed. Chombo users are strongly encouraged to download, install, and use VisIt. Development of VisIt for Chombo AMR datasets is ongoing and future releases of VisIt should provide addition functionality for Chombo AMR datasets.

The development of [VisIt](#) for Chombo datasets is being done by [VACET](#) (Visualization and Analytics Center for Enabling Technologies).

**STATUS:** No longer being maintained or developed (see above).

Contact [ChomboVis Development Team](#)

---

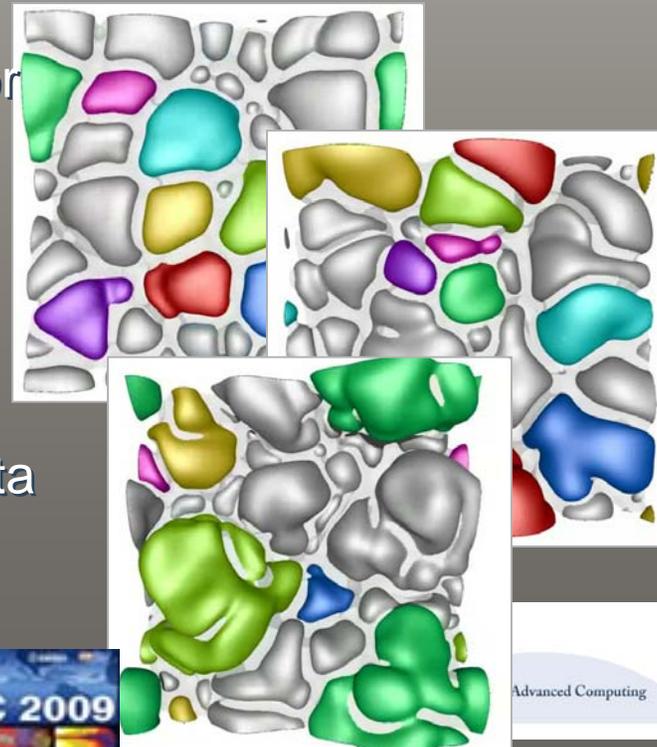
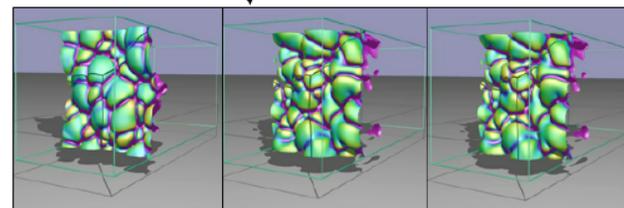
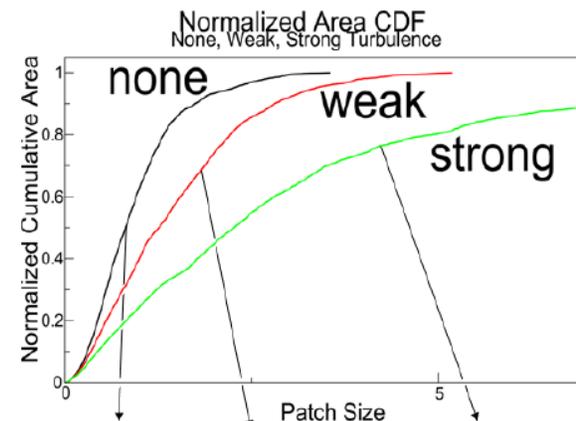
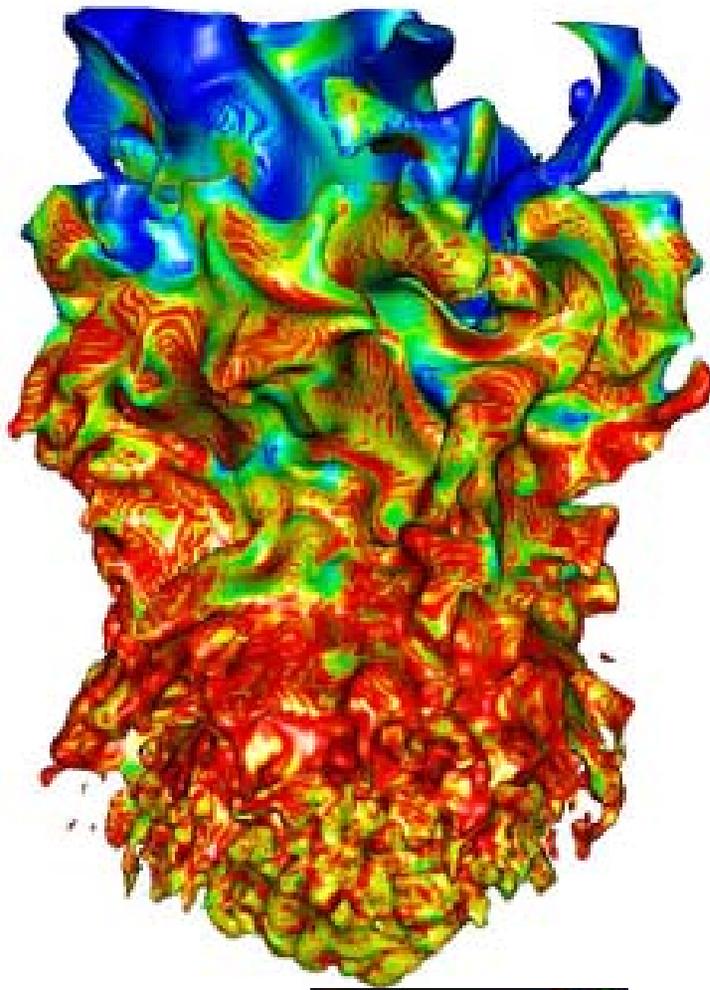
Return to [ANAG Home Page](#)





## Combustion

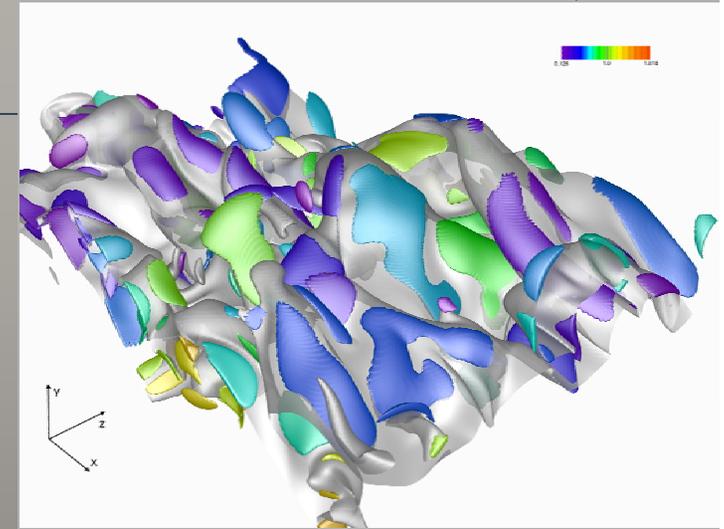
- PI: Jo...
- Com...
- Cons...
- Awar...
- Acco...
- Ne...
- stu...
- pa...
- Jo...
- Scier...
- Fir...
- tim...
- to...
- siz...
- lea...



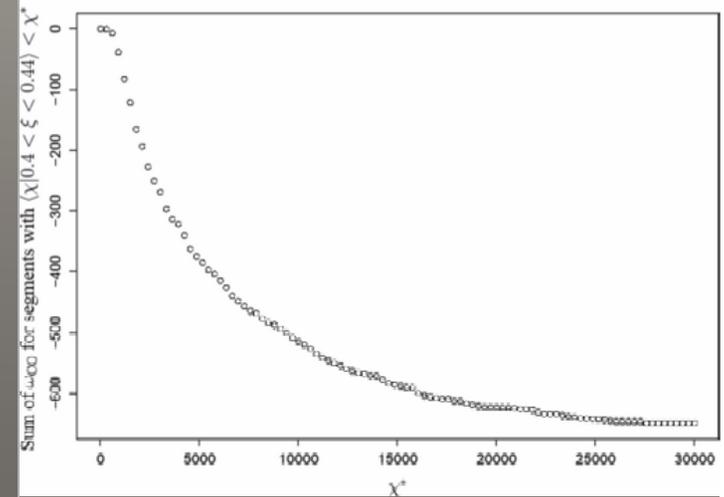


## Combustion

- PI: Jacqueline Chen (SNL-CA), Incite awardee.
- Accomplishment(s)
  - Algorithms for feature segmentation, tracking, and analysis.
  - Co-authors on multiple papers.
- Science Impact(s)
  - New capability: first-ever ability to see relationship between simulation parameters (e.g., level of turbulence) and scalar dissipation rate.



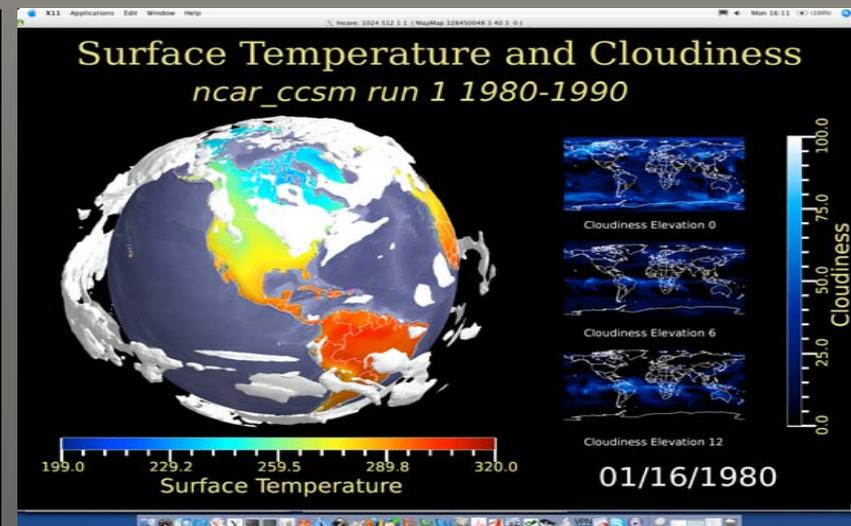
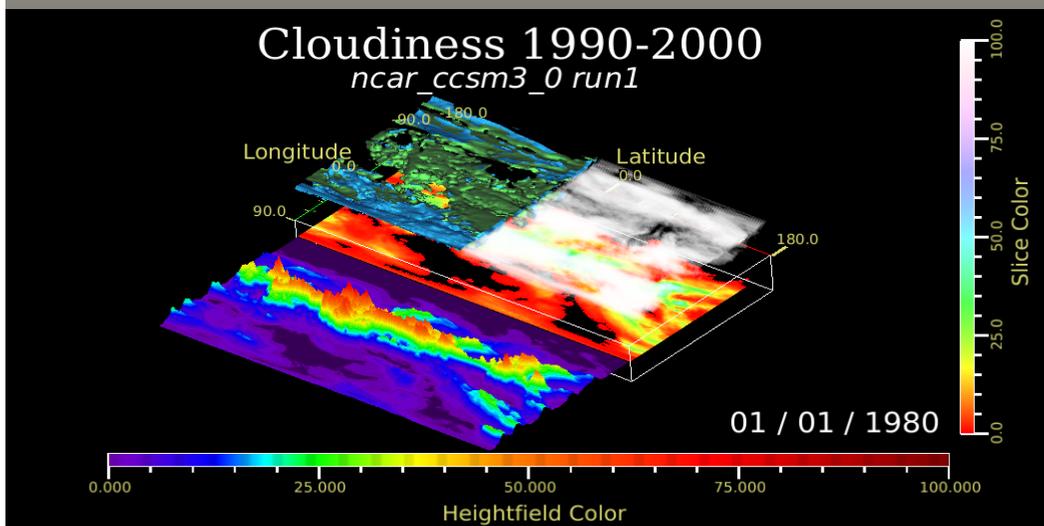
Cumulative carbon monoxide source term inside segments





## Climate Modeling

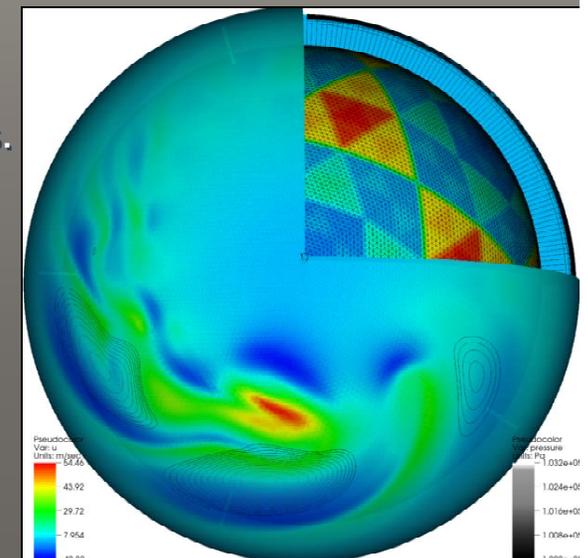
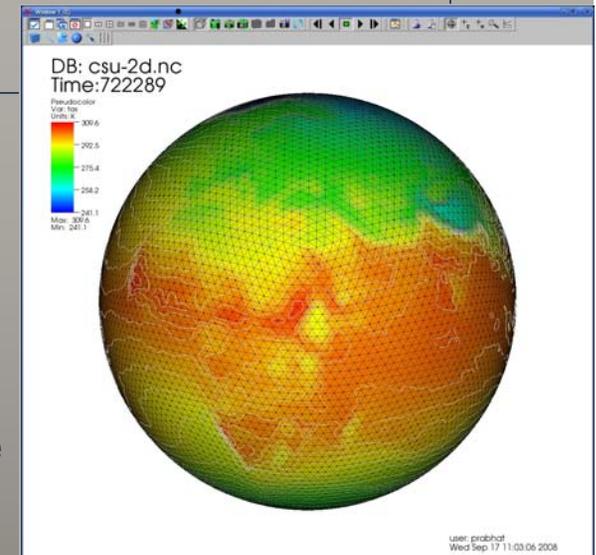
- CDAT: dominant, highly domain specific visualization application for climate data
  - But CDAT lacked 3D capabilities.
  - Need: very lightweight library for tight integration.
- The SEG developed SFSG (Simple and Flexible Scene Graph) library and integrated and deployed it in CDAT.





## Climate Modeling, ctd.

- PI: Dave Randall (CSU). SciDAC Application: Role of Clouds in Global Climate
- Accomplishments
  - Debug and optimize parallel I/O to meet performance objectives.
  - New visualization infrastructure for icosahedral grid.
- Science Impact
  - Enable effective use of INCITE allocation at NERSC
  - Critical s/w infrastructure to enable visualization and analysis of ensemble runs of new global cloud models.
- Other Collaborators:
  - NERSC Center staff
  - Karen Schuchardt (PNNL)

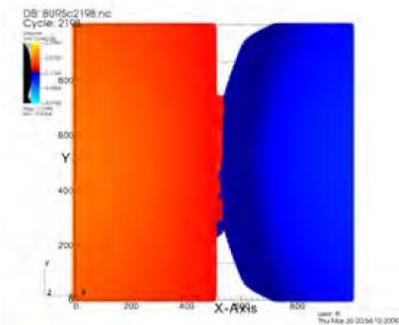




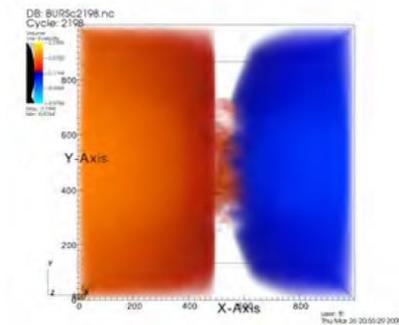
## Tech transfer: the SEG is productizing the University of Utah's volume renderer.

- Utah's pioneering work in multi-dimensional transfer functions is being productized to reach a wider audience.
  - SLIVR library and its next generation replacement, Tuvok.
  - SLIVR is integrated into VisIt; Tuvok integration planned.

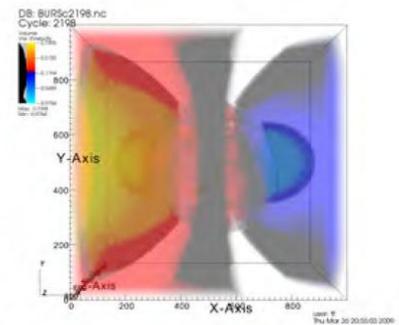
### Tuvok's Clearview mode



(a) Splat-based volume rendering.



(b) 3D texture-based volume rendering.



(c) SLIVR volume rendering.

Widgets

rect0001
tri0002
paraboloid0003

Widget attributes

Name	paraboloid0003
Size/Location	0.2 0.778095 0.5
Color	<input type="color" value="#00FFFF"/> 80%

Rendering method: SCI, University of Utah (SLIVR)



# VACET

## VACET Accomplishments Meet SciDAC Program Objectives

- Science Impact
- Production-Quality, Petascale capable software infrastructure
- Field-leading, award-winning research
- Progress towards petascale
- Effective use of ASCR Computing resources
- Service and Outreach

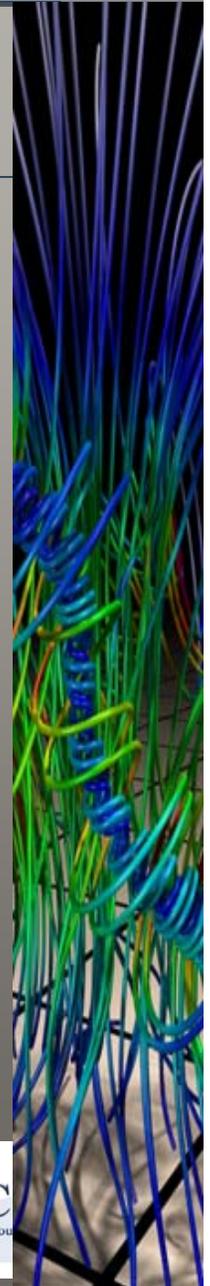


U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science



SciDAC  
Scientific Discovery through





# VACET

The End



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science



**SciDAC**  
Scientific Discovery through Advanced Computing