Argonne NATIONAL LABORATORY

Introduction to vl3

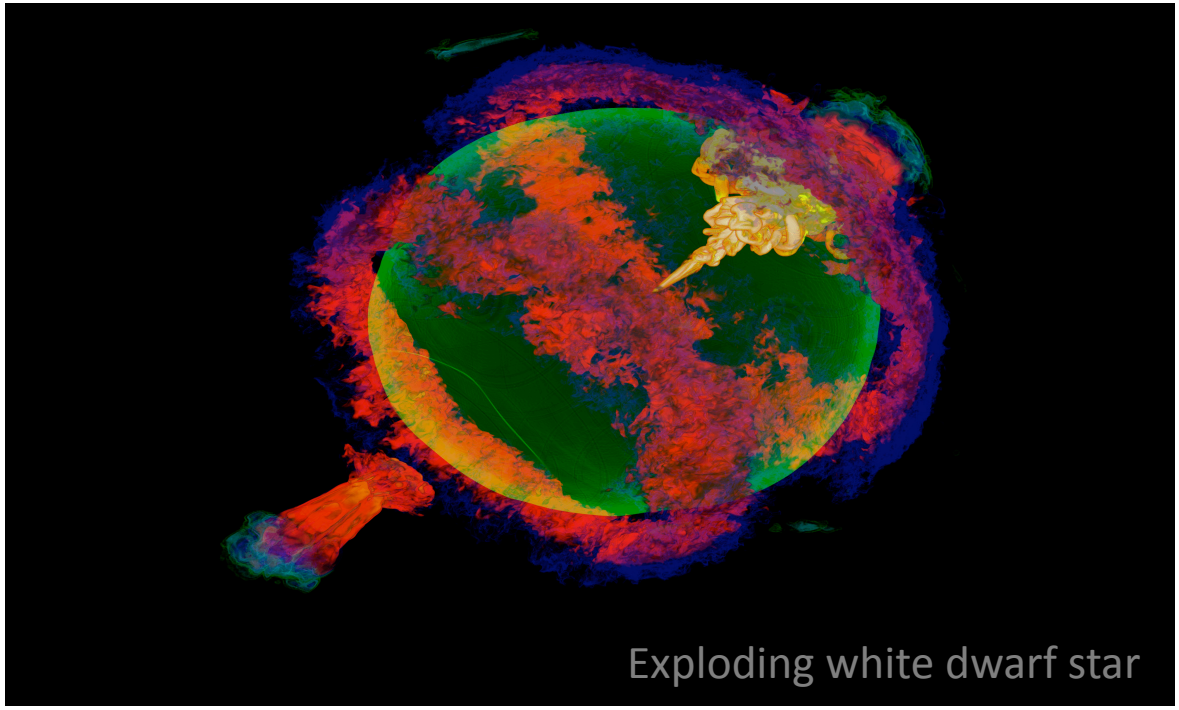

Exploding white dwarf star

Venkat Vishwanath

Argonne National Laboratory

University of Chicago

U.S. DEPARTMENT OF ENERGY

Volume Rendering Pipeline

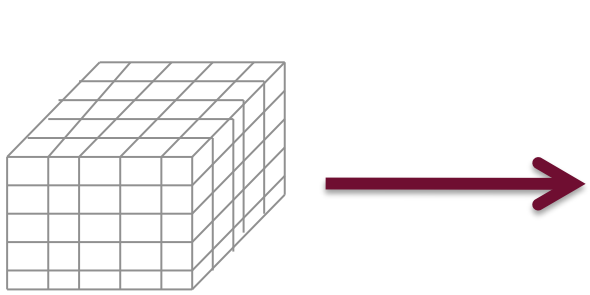

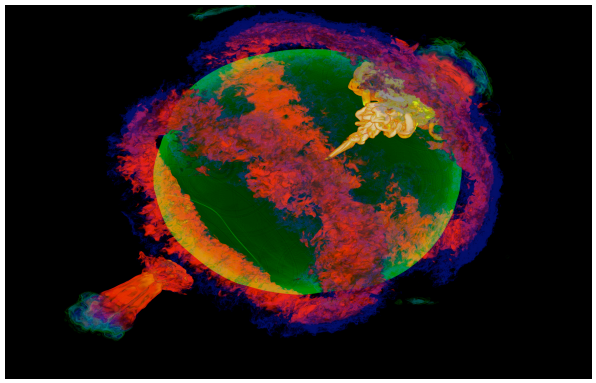

Data → Filter → Render → Composite → Image

- Reduces 2- and 3-dimensionsal datasets into 2D images
- Enables one to better understand scientific phenomena

Scalable Volume Rendering

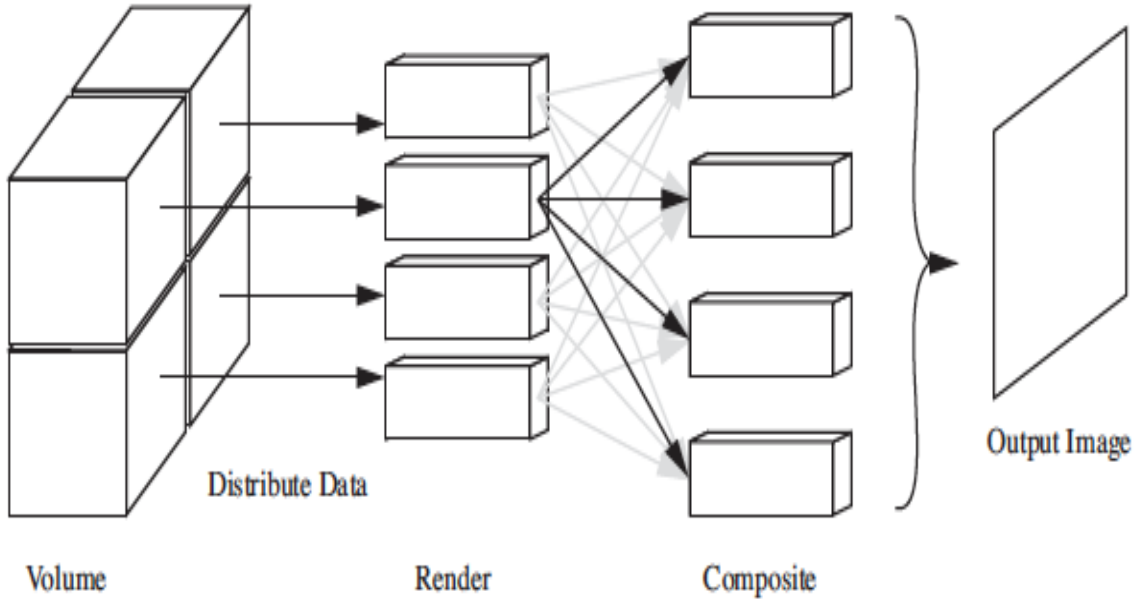

Leverage Parallelism in all stages

- Subdivision of data
- Parallel rendering
- Parallel image composition

Software Rendering vs Hardware Rendering

- Software rendering refers to the computations being performed on the CPU.
- Software rendering is extremely flexible and runs on any hardware.
- With the proliferation of commodity Graphical Processing Units (GPU), one can leverage GPUs for volume rendering. This is referred to as Hardware Rendering.
- Hardware Rendering requires a GPU and is portable across GPUs (AMD, NVidia) using the GLSL shading language. CUDA is not portable.

vl3 Architecture - Modular and Extensible

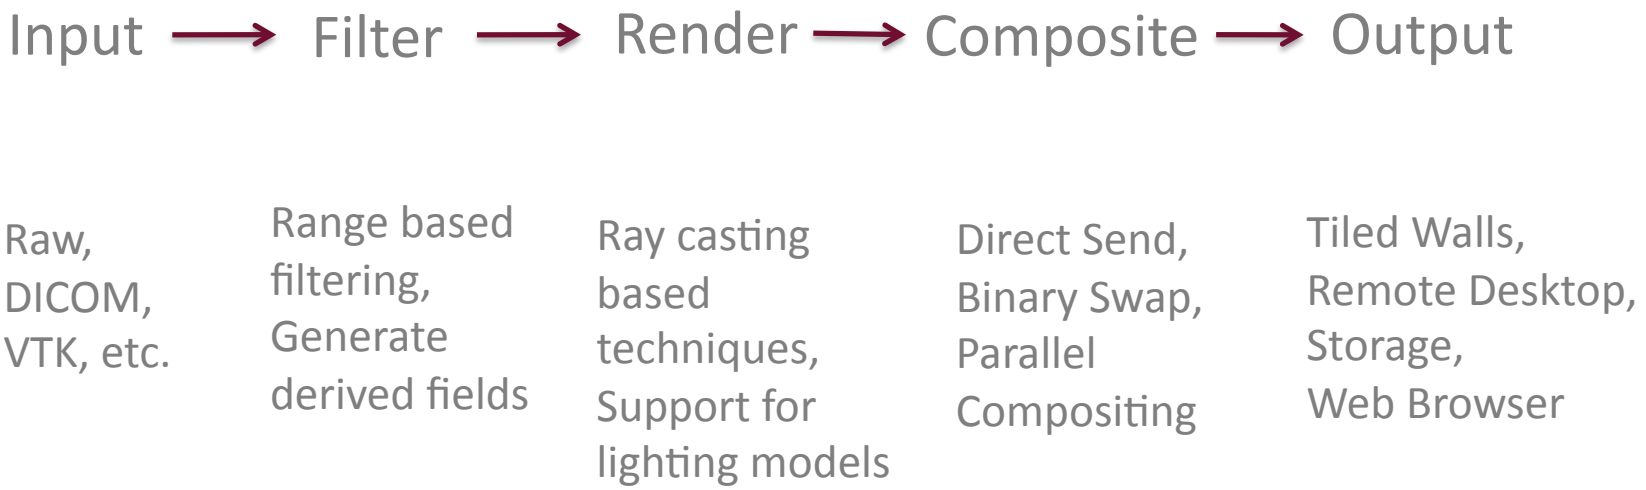

- vl3 runs on a desktop, GPU-based clusters and CPU-based clusters
- vl3 was used to generate the **top 2** visualizations at SciDAC 2010 and has won a total of 5 OASCRs in the past 3 years

vl3 Features

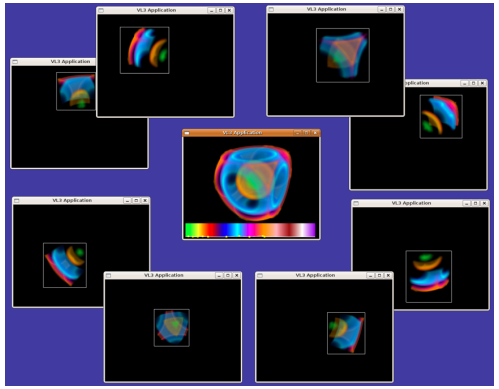

- Parallel Volume Visualization
- Supports both Software and GPU-based Hardware volume rendering and compositing
- Support for multiple data formats, filters, compositing algorithms via plugins
- Support for multiple output modes
- Support for multiple interaction devices
- Configurable design enables vl3 to be customized for the underlying architecture

vl3 Examples

Computed Tomography

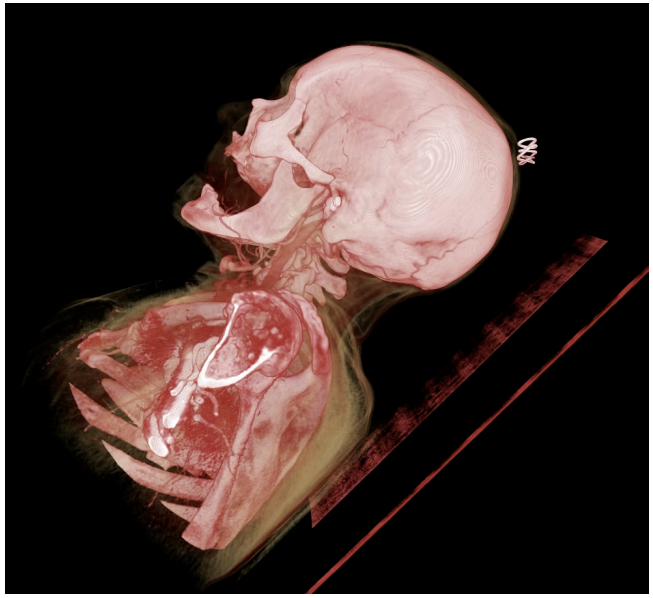

Microtomography

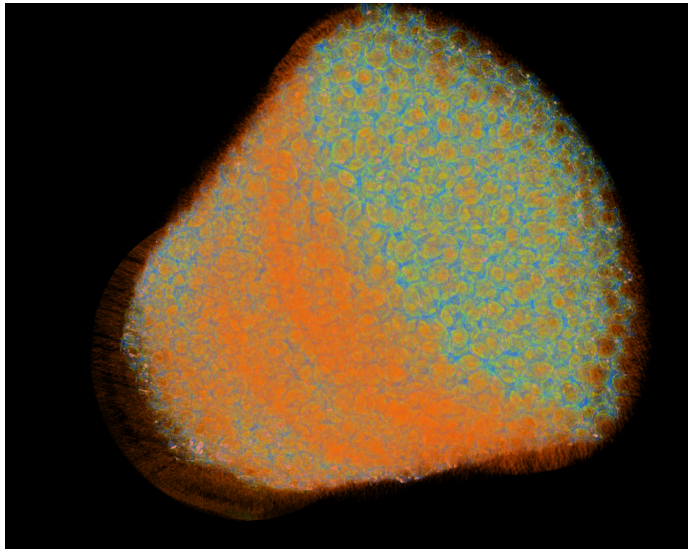

Cosmology

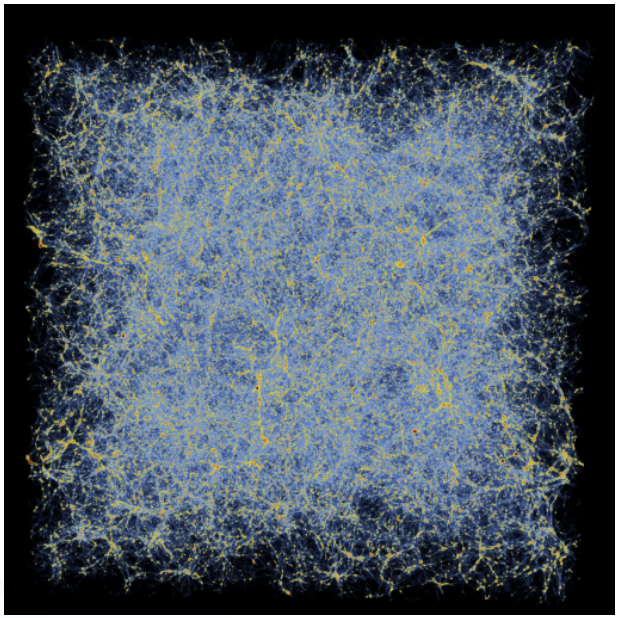

Bioengineering

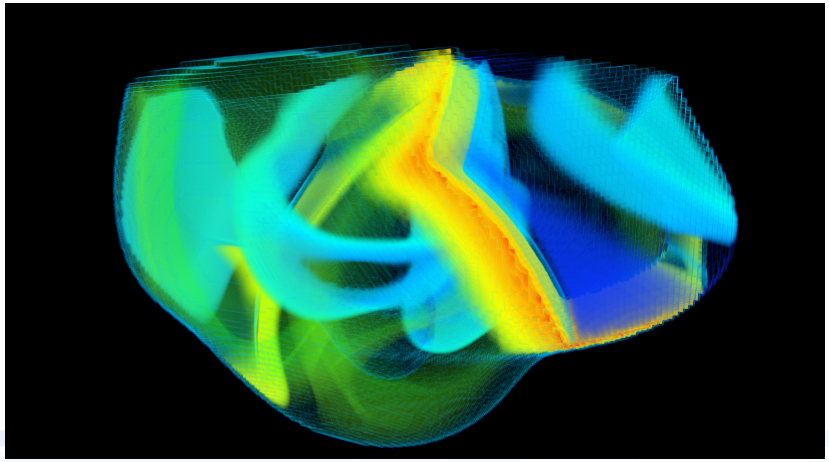

vl3- Enabling Domain-specific Visualization

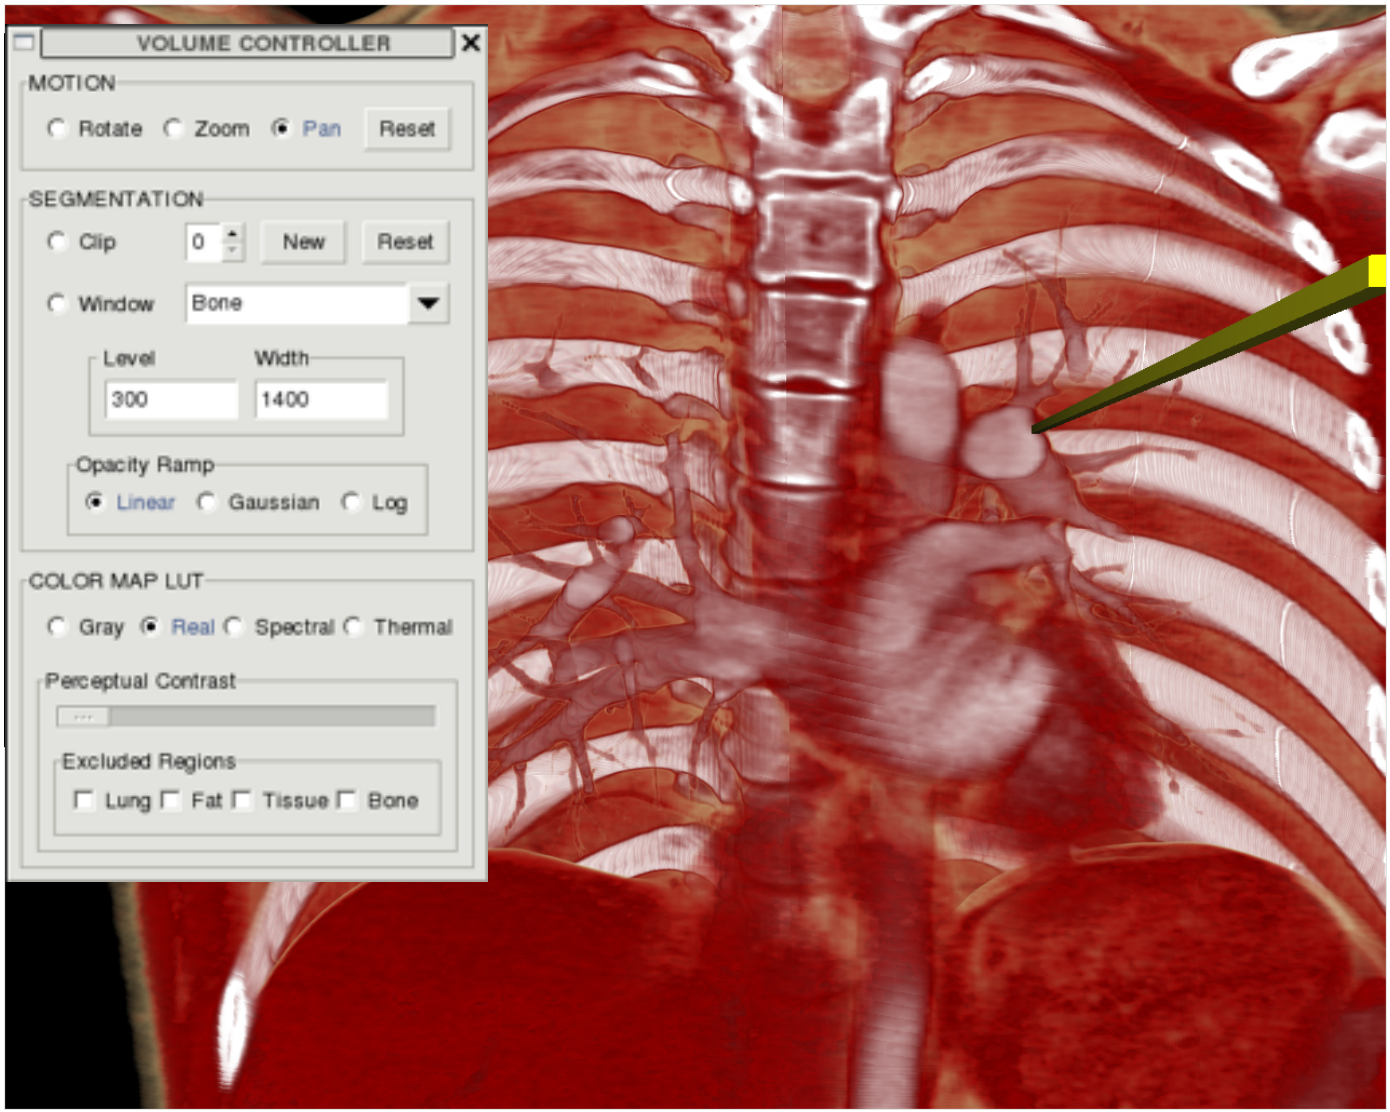

vl3- Enabling Domain-specific Visualization

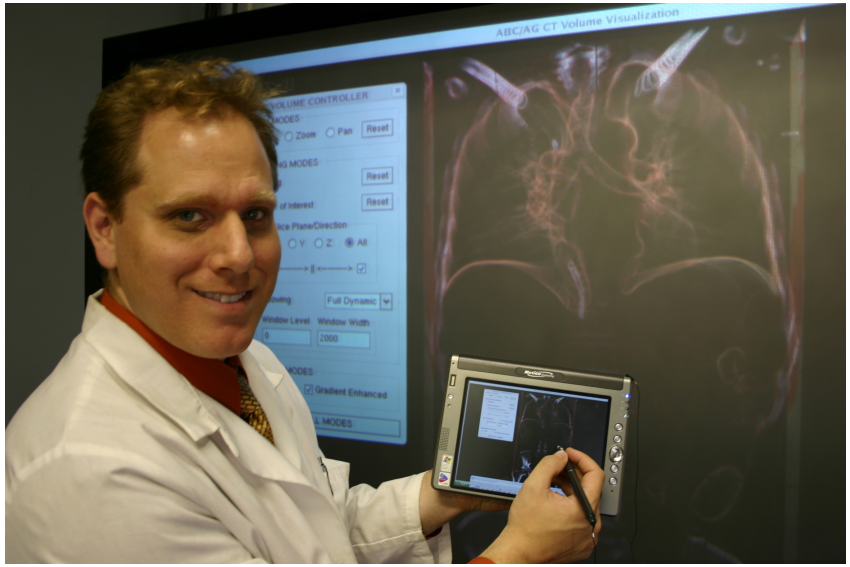

- vl3 is regularly used by surgeons at Univ. of Chicago and provides them with an intuitive user experience
- vl3 is used in teaching a virtual anatomy class at University of Chicago

FLASH Center

enzo

APS, Argonne

vl3 users

SCEC

an NSF + USGS center

THE UNIVERSITY OF CHICAGO

Department of Surgery

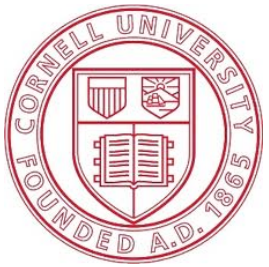

Argonne Leadership Class Facility (ALCF)

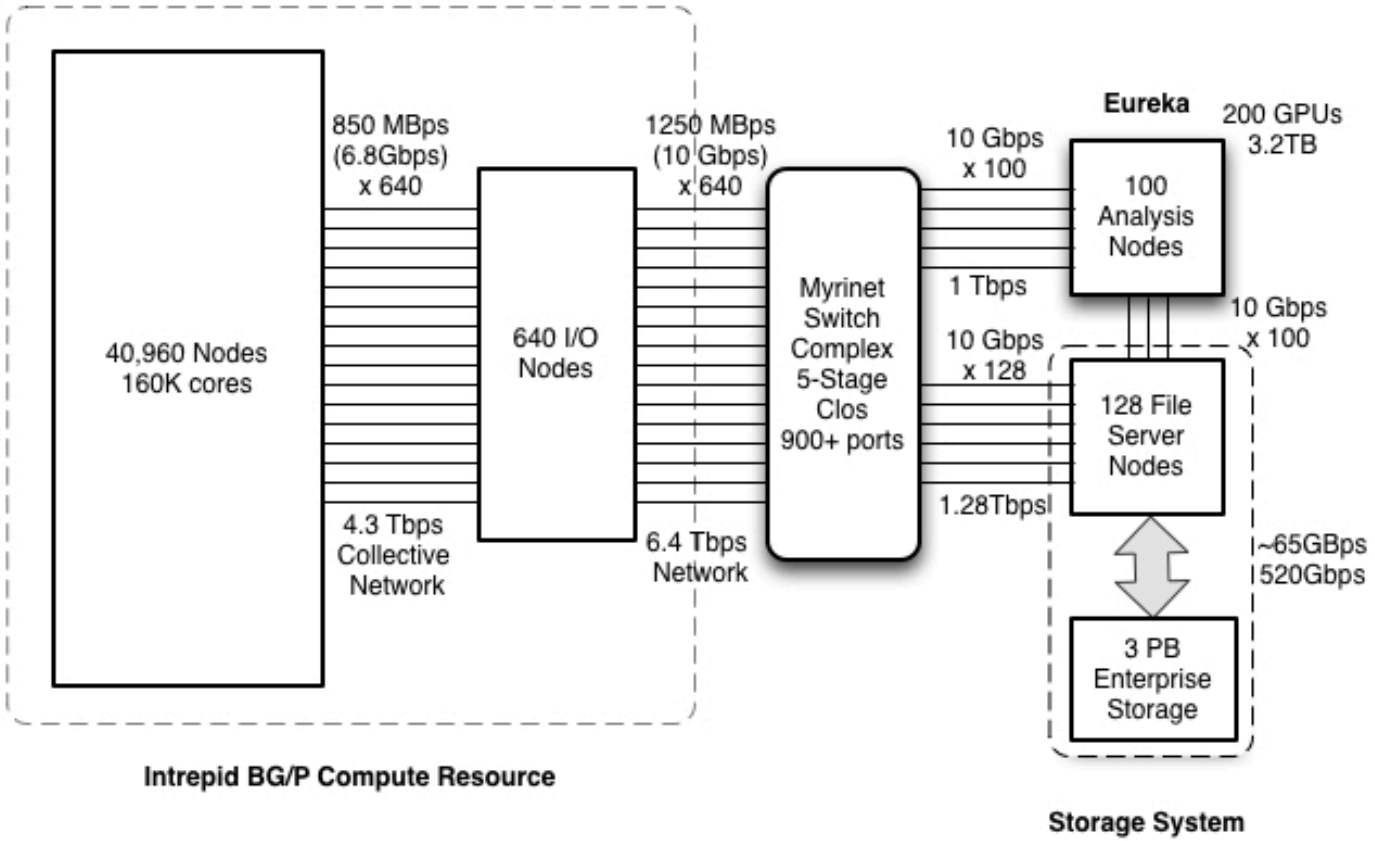

Eureka is a 110 Tflop GPU-based Data Analytics and Visualization Cluster facilitating major scientific discoveries and breakthroughs for programs including INCITE and ALCC

Exploring Cosmology With Supercomputers, Supernetworks, and Supervisualization

Intergalactic medium on 2 Glyr scale

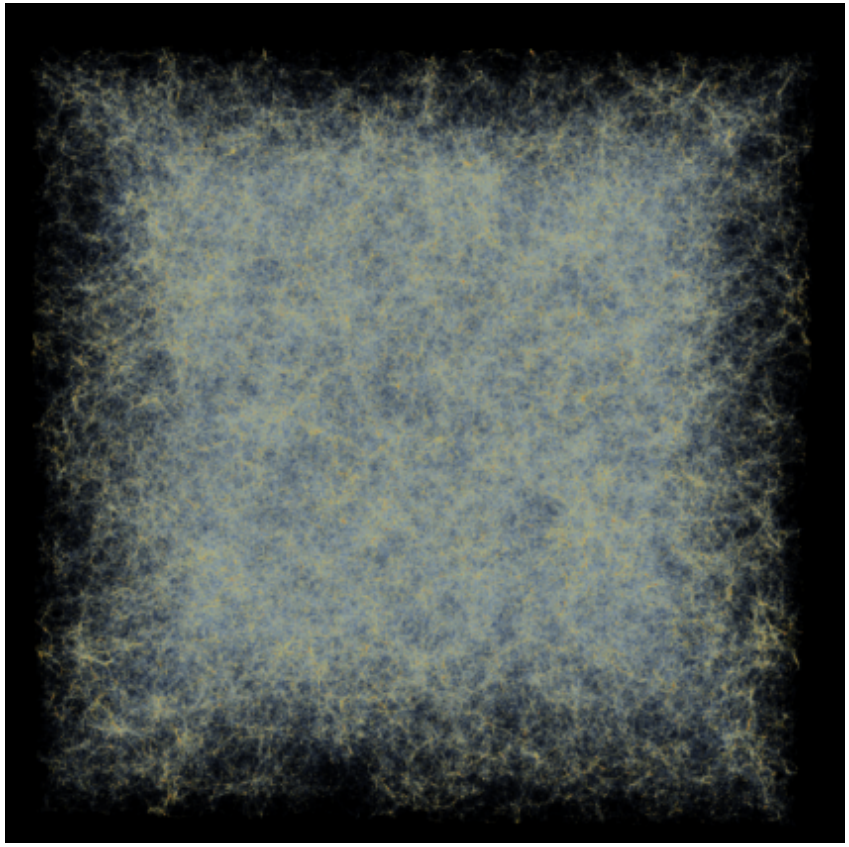

- 4096^3 particle/cell hydrodynamic cosmology simulation
- NICS Kraken (XT5)
  - 16,384 cores
- Output
  - 148 TB movie output
  - 80 TB diagnostic dumps

Science: Norman, Harkness,Paschos SDSC
Visualization: Insley, ANL; Wagner SDSC

StarGate Streaming Rendering

ESnet

ALCF Internal

3 A media bridge at the border provides secure access to the parallel rendering streams.

gs1.intrepid.alcf.anl.gov

SDSC

San Diego

10 Gbps

ALCF

Chicago

5 Updated instructions are sent back to the renderer to change views, or load a different dataset.

2

The full image is broken into subsets (tiles). The tiles are continuously encoded as a separate movies.

4 flPy, a parallel (MPI) tiled image/movie viewer composites the individual movies, and synchronizes the movie playback across the OptIPortal rendering nodes.

1 Simulation volume is rendered using vl3 , a parallel (MPI) volume renderer utilizing Eureka's GPUs. The rendering changes views steadily to highlight 3D structure.

- ANL * Calit2 * LBNL * NICS * ORNL * SDSC

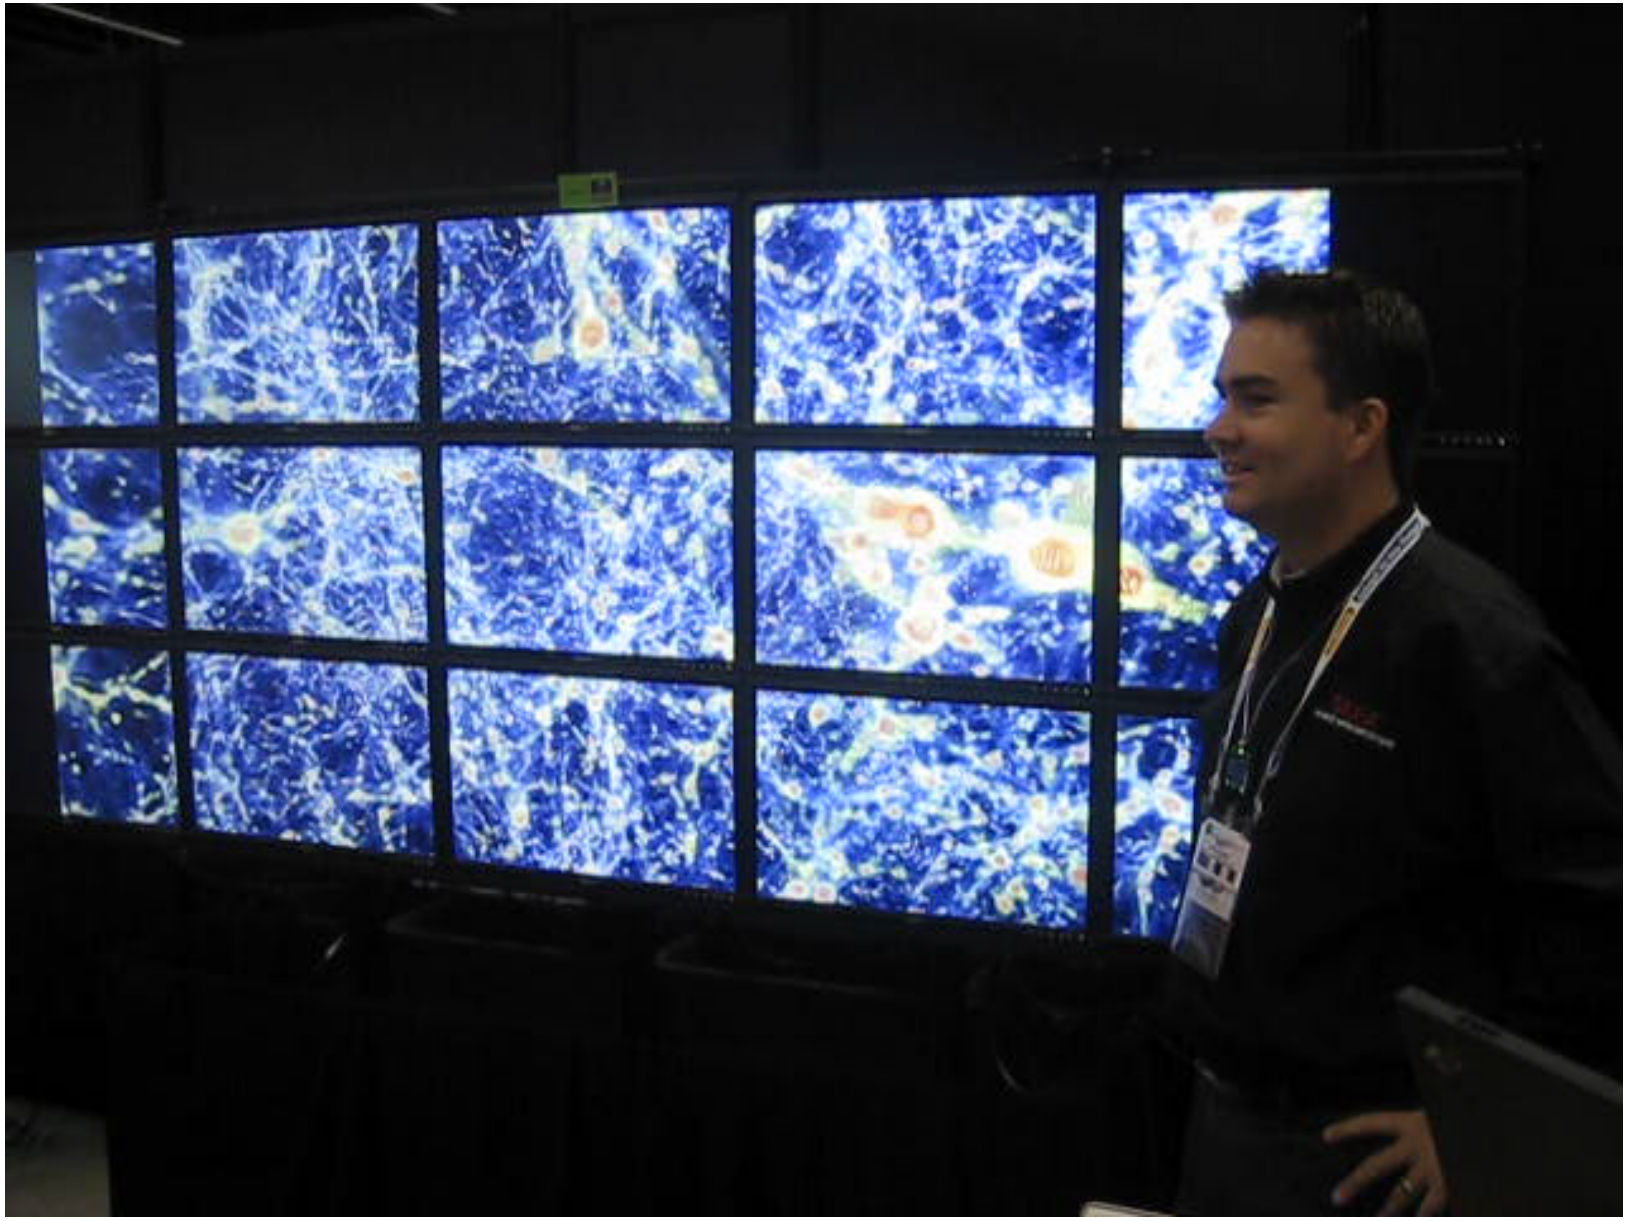

Rick Wagner in front of OptiPortal on SC09 show floor in Portland, OR, with vl3 streaming from Eureka at Argonne

vl3 Rendering Performance on Eureka

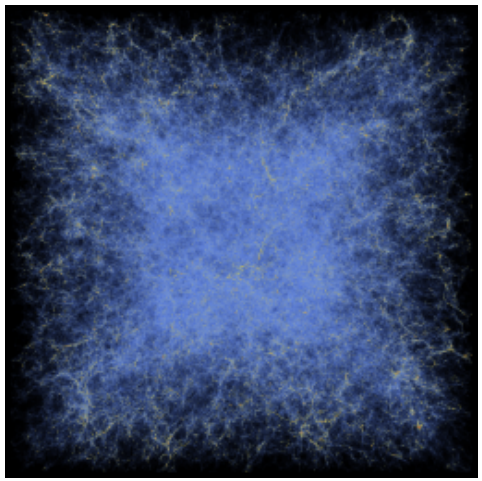

Data Size	Number of Processors/ Graphics Cards	Render + Composite Time
1024^3	8	0.9 sec
4096^3	64	1.3 sec

vl3 has currently scaled to $(6400)^3$ data sizes and the entire Eureka cluster

Acknowledgements

This work was supported was in part by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357, by the National Library of Medicine as part of contract #N01-LM-3-3508, and by the National Science Foundation as part of contract OCI-0504086, and in part by the DOE-supported ASC / Alliance Center for Astrophysical Thermonuclear Flashes at the University of Chicago.

- Eric Olson, Joseph Insley, Michael Papka, Mark Hereld, Tom Uram (ANL); Brad Gallagher (UChicago)
- Matt McCrory (NWU) and Joe Paris (NWU)

Back up Slides

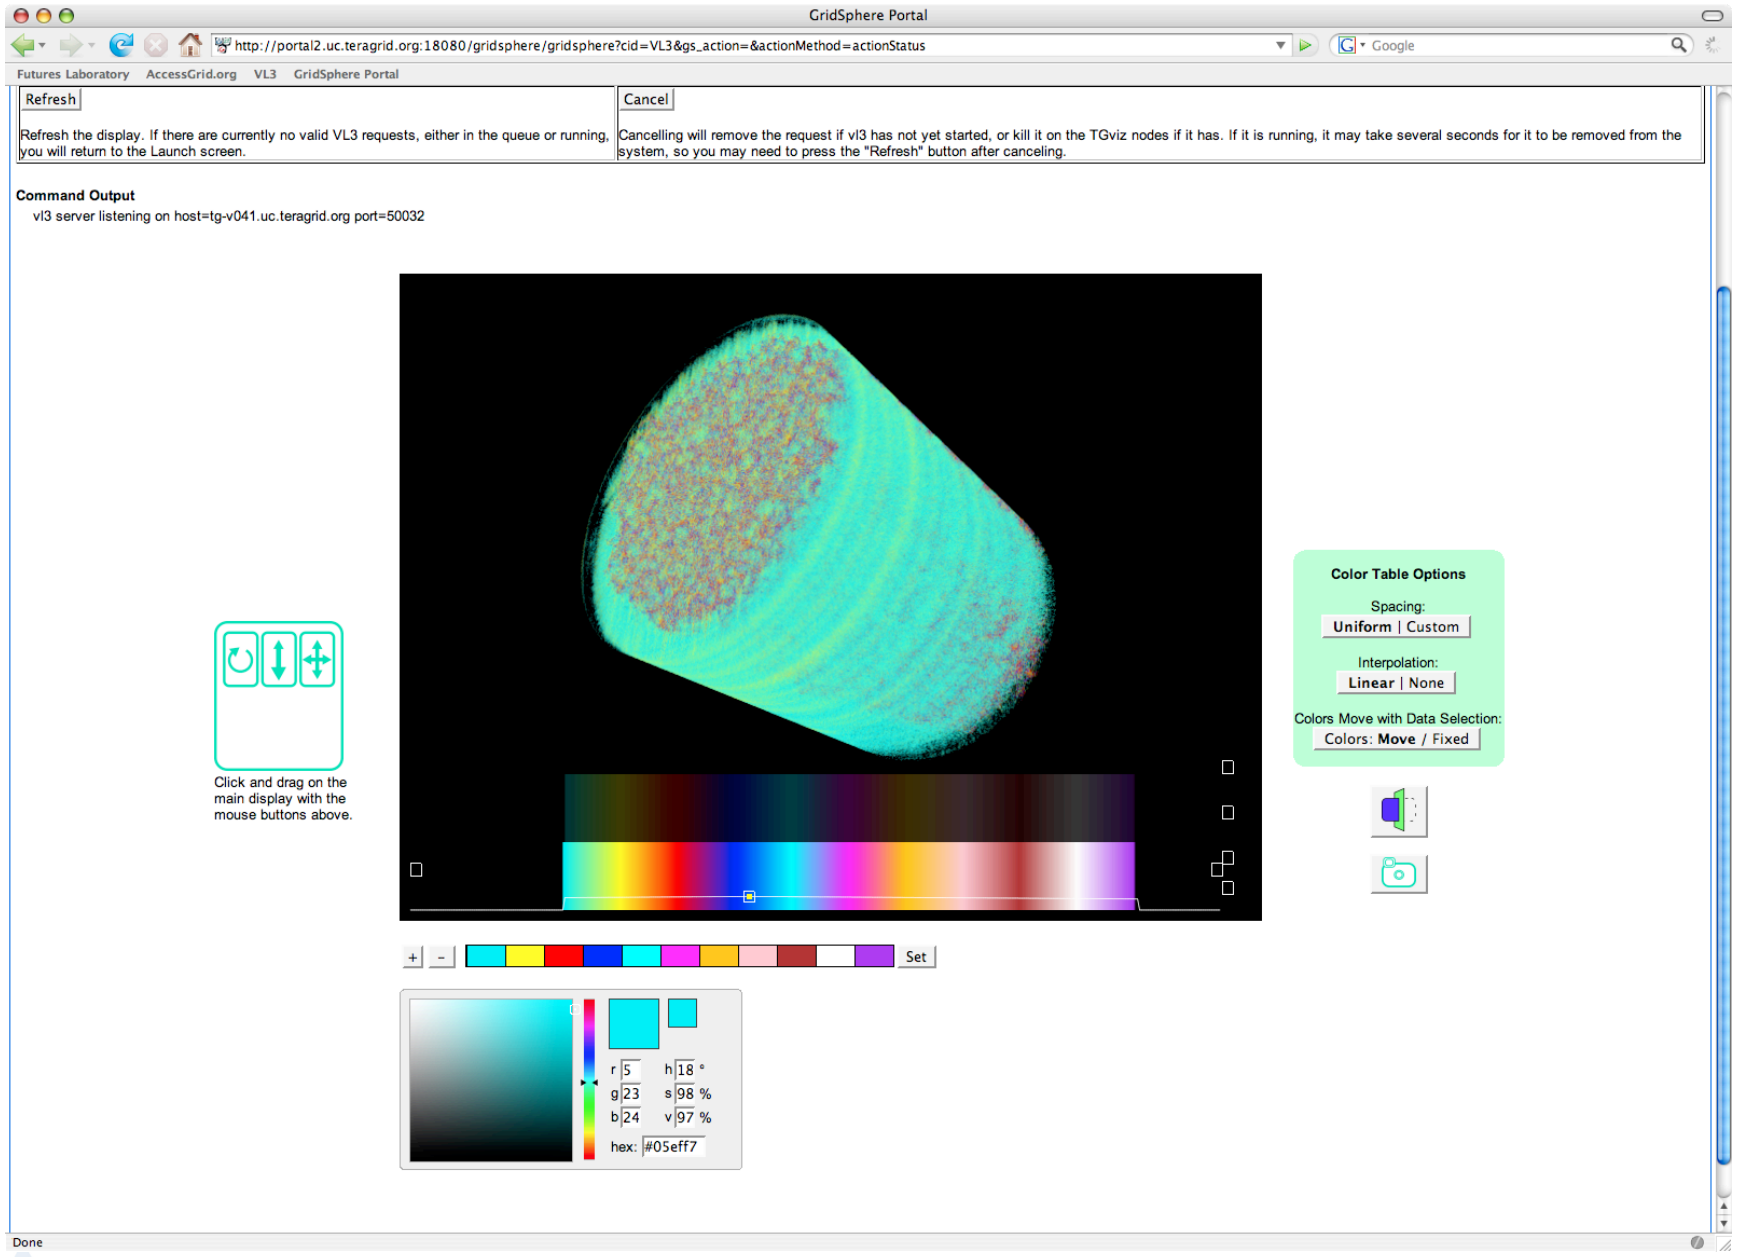
GridSphere Portal
http://portal2.uc.teragrid.org:18080/gridsphere/gridsphere?cid=VL3&gs_action=&actionMethod=actionStatus
Google
Futures Laboratory
AccessGrid.org
VL3
GridSphere Portal
Refresh
Refresh the display. If there are currently no valid VL3 requests, either in the queue or running, you will return to the Launch screen.
Cancel
Cancelling will remove the request if vl3 has not yet started, or kill it on the TGviz nodes if it has. If it is running, it may take several seconds for it to be removed from the system, so you may need to press the "Refresh" button after canceling.
Command Output
vl3 server listening on host=tg-v041.uc.teragrid.org port=50032
Click and drag on the main display with the mouse buttons above.
Color Table Options
Spacing:
Uniform | Custom
Interpolation:
Linear | None
Colors Move with Data Selection:
Colors: Move / Fixed
+
-
Set
r 5
g 23
b 24
h 18 °
s 98 %
v 97 %
hex: #05eff7
Done