Argonne NATIONAL LABORATORY

Argonne National Laboratory Site Report

Venkatram Vishwanath

Argonne National Laboratory

venkatv@mcs.anl.gov

U.S. DEPARTMENT OF ENERGY

Mission: Enhance scientific insight and productivity in computational and experimental sciences via research and development of scalable data analysis, visualization and user interfaces

Michael Papka*, Mark Hereld, Joseph Insley, Aaron Knoll, Eric Olson, Tom Peterka, Rob Ross, Tom Uram, Venkat Vishwanath

*Site representative: papka@anl.gov

Funding Sources

Argonne Leadership Computing Facility Resources

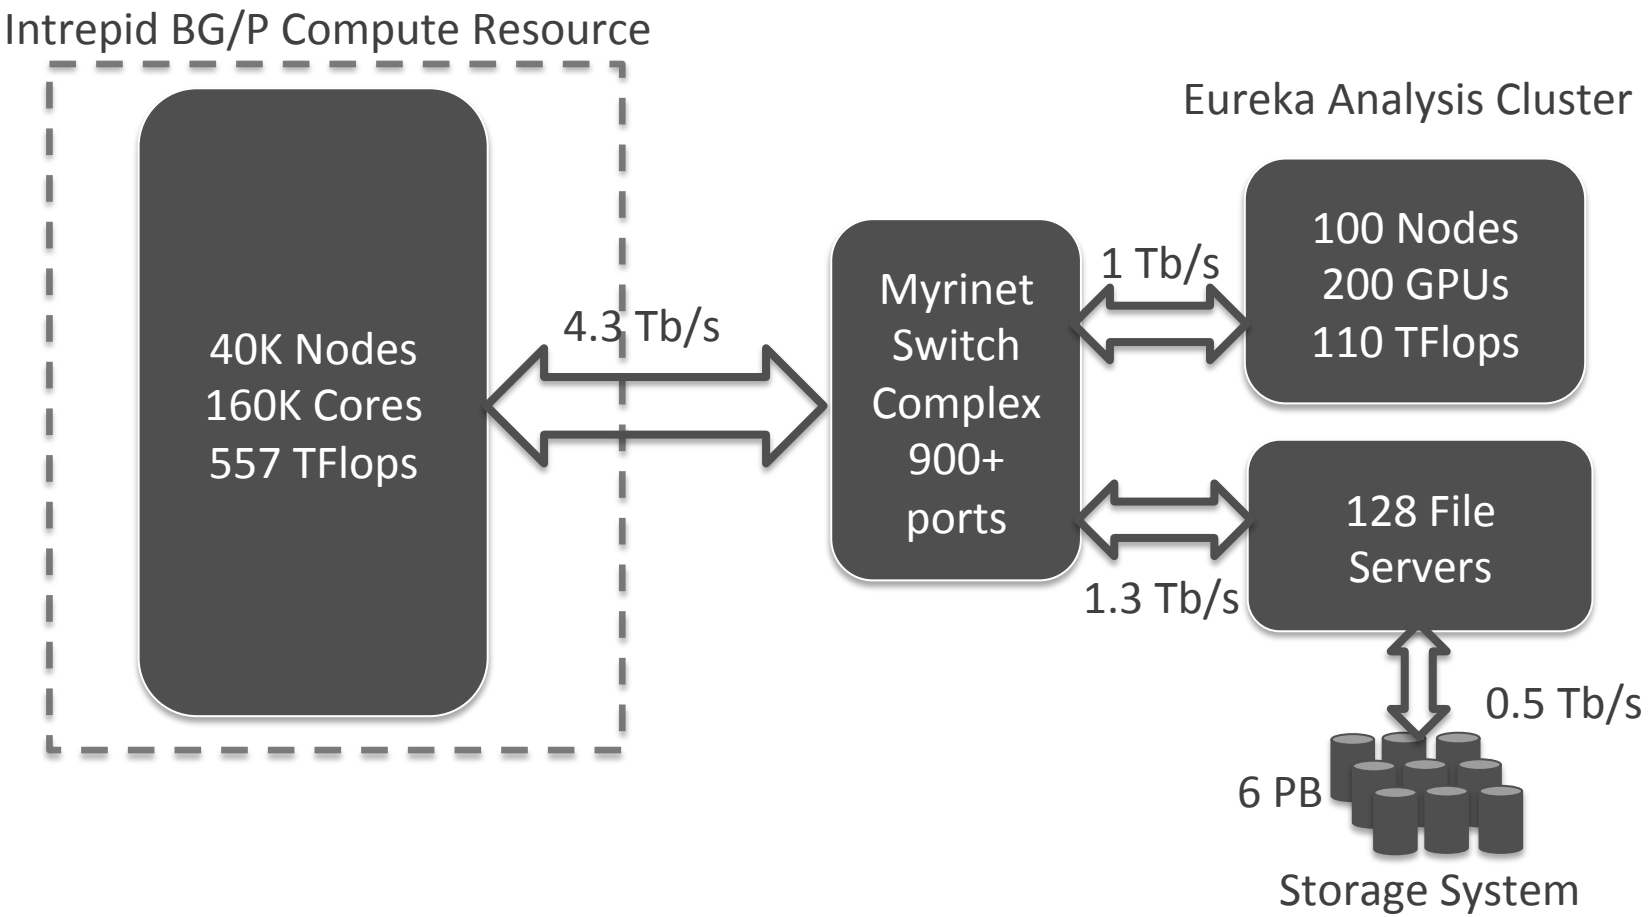

Argonne Magellan Cloud Hardware - Final Configuration

Compute Servers

504 Compute Servers
- Nehalem Dual quad-core 2.66GHz
- 24GB RAM, 500GB Disk
- QDR Infiniband

Totals
- 4032 Cores, 40TF Peak
- 12TB Memory, 250TB Disk

Active Storage Servers

200 Compute/Storage Nodes
40TB FLASH/SSD Storage
9.6TB Memory, 1.6PB Disk
QDR Infiniband

GPU Servers

133 GPU Servers
8.5TB Memory, 133TB Disk
266 Nvidia 2070 GPU cards
QDR Infiniband

Big Memory Servers

15 Compute Servers
15TB Memory, 15TB Disk
QDR Infiniband

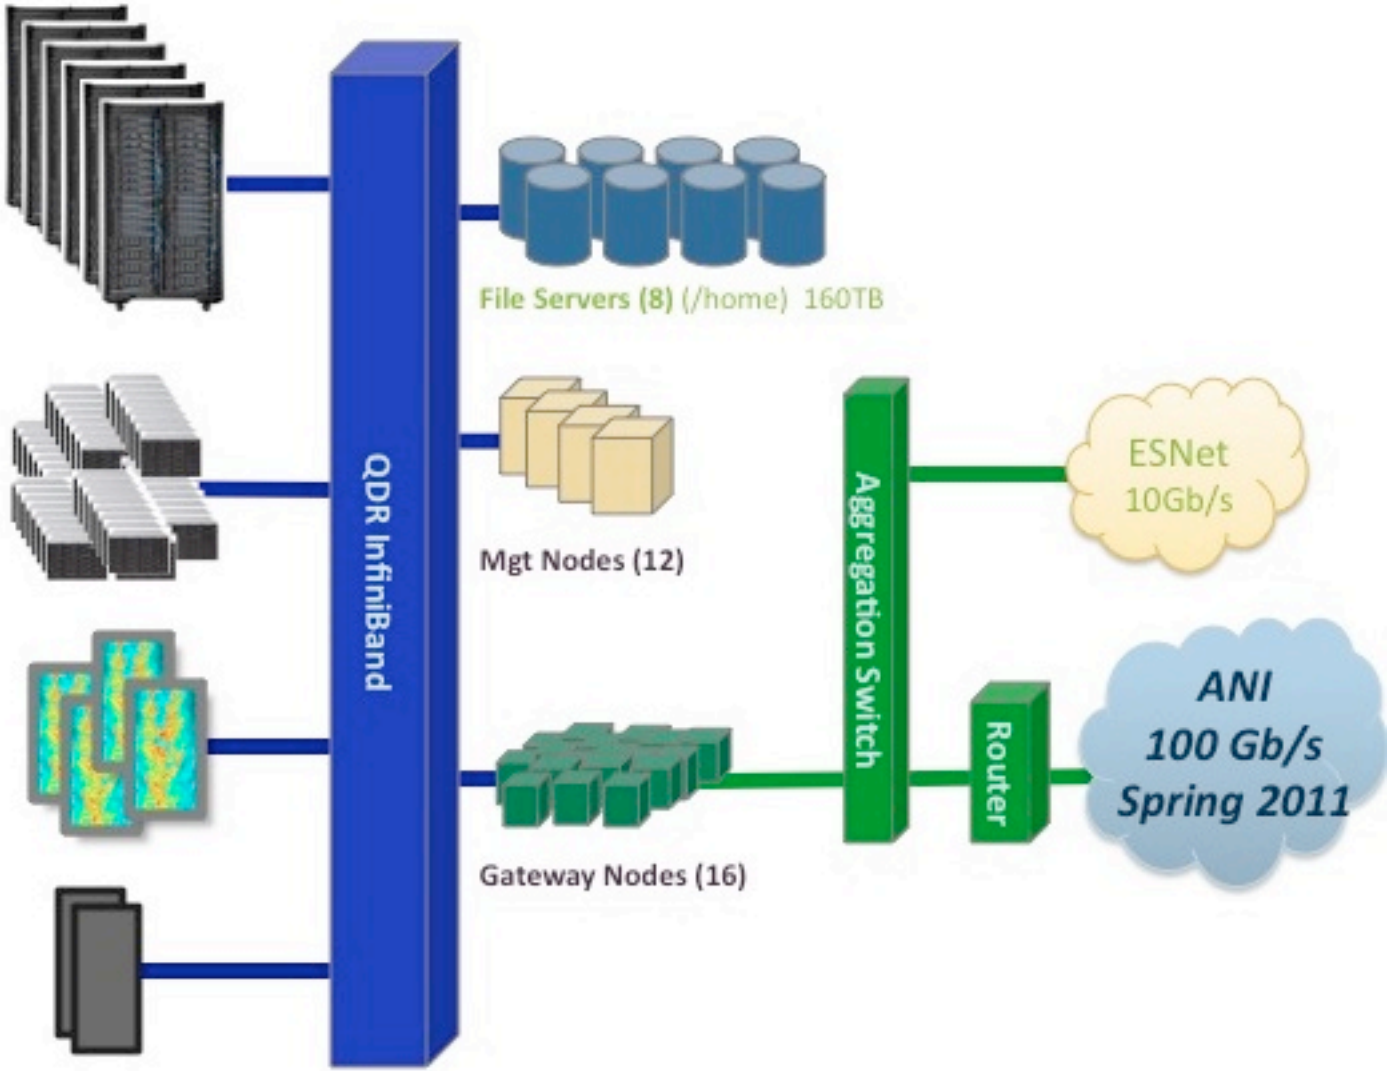

Planned Growth

- Hardware
  - Mira 10 PF BG/Q
  - Data Analysis Cluster for Mira
- Software
  - Updated versions of various visualization and analysis software
- Personnel
  - Depending on funding

Smaash – Environment for Simulation Capture, Curation and Analysis

- Problem: management of simulations is haphazard
- Smaash makes the following easier:
  - Starting, watching & controlling
  - Data analysis & visualization
  - Cataloging results & preparing publications
  - Tracking provenance
- How?
  - Automation, uniform practices
  - One repository of meta-data
  - Generic utility (other sciences)
  - Security, extensibility

Experiences Using Smaash to Manage Data-Intensive Simulations, Hudson et al, HPDC 2011, San Jose, CA

Multiscale Problem:
From O(10cm) to O(1nm)

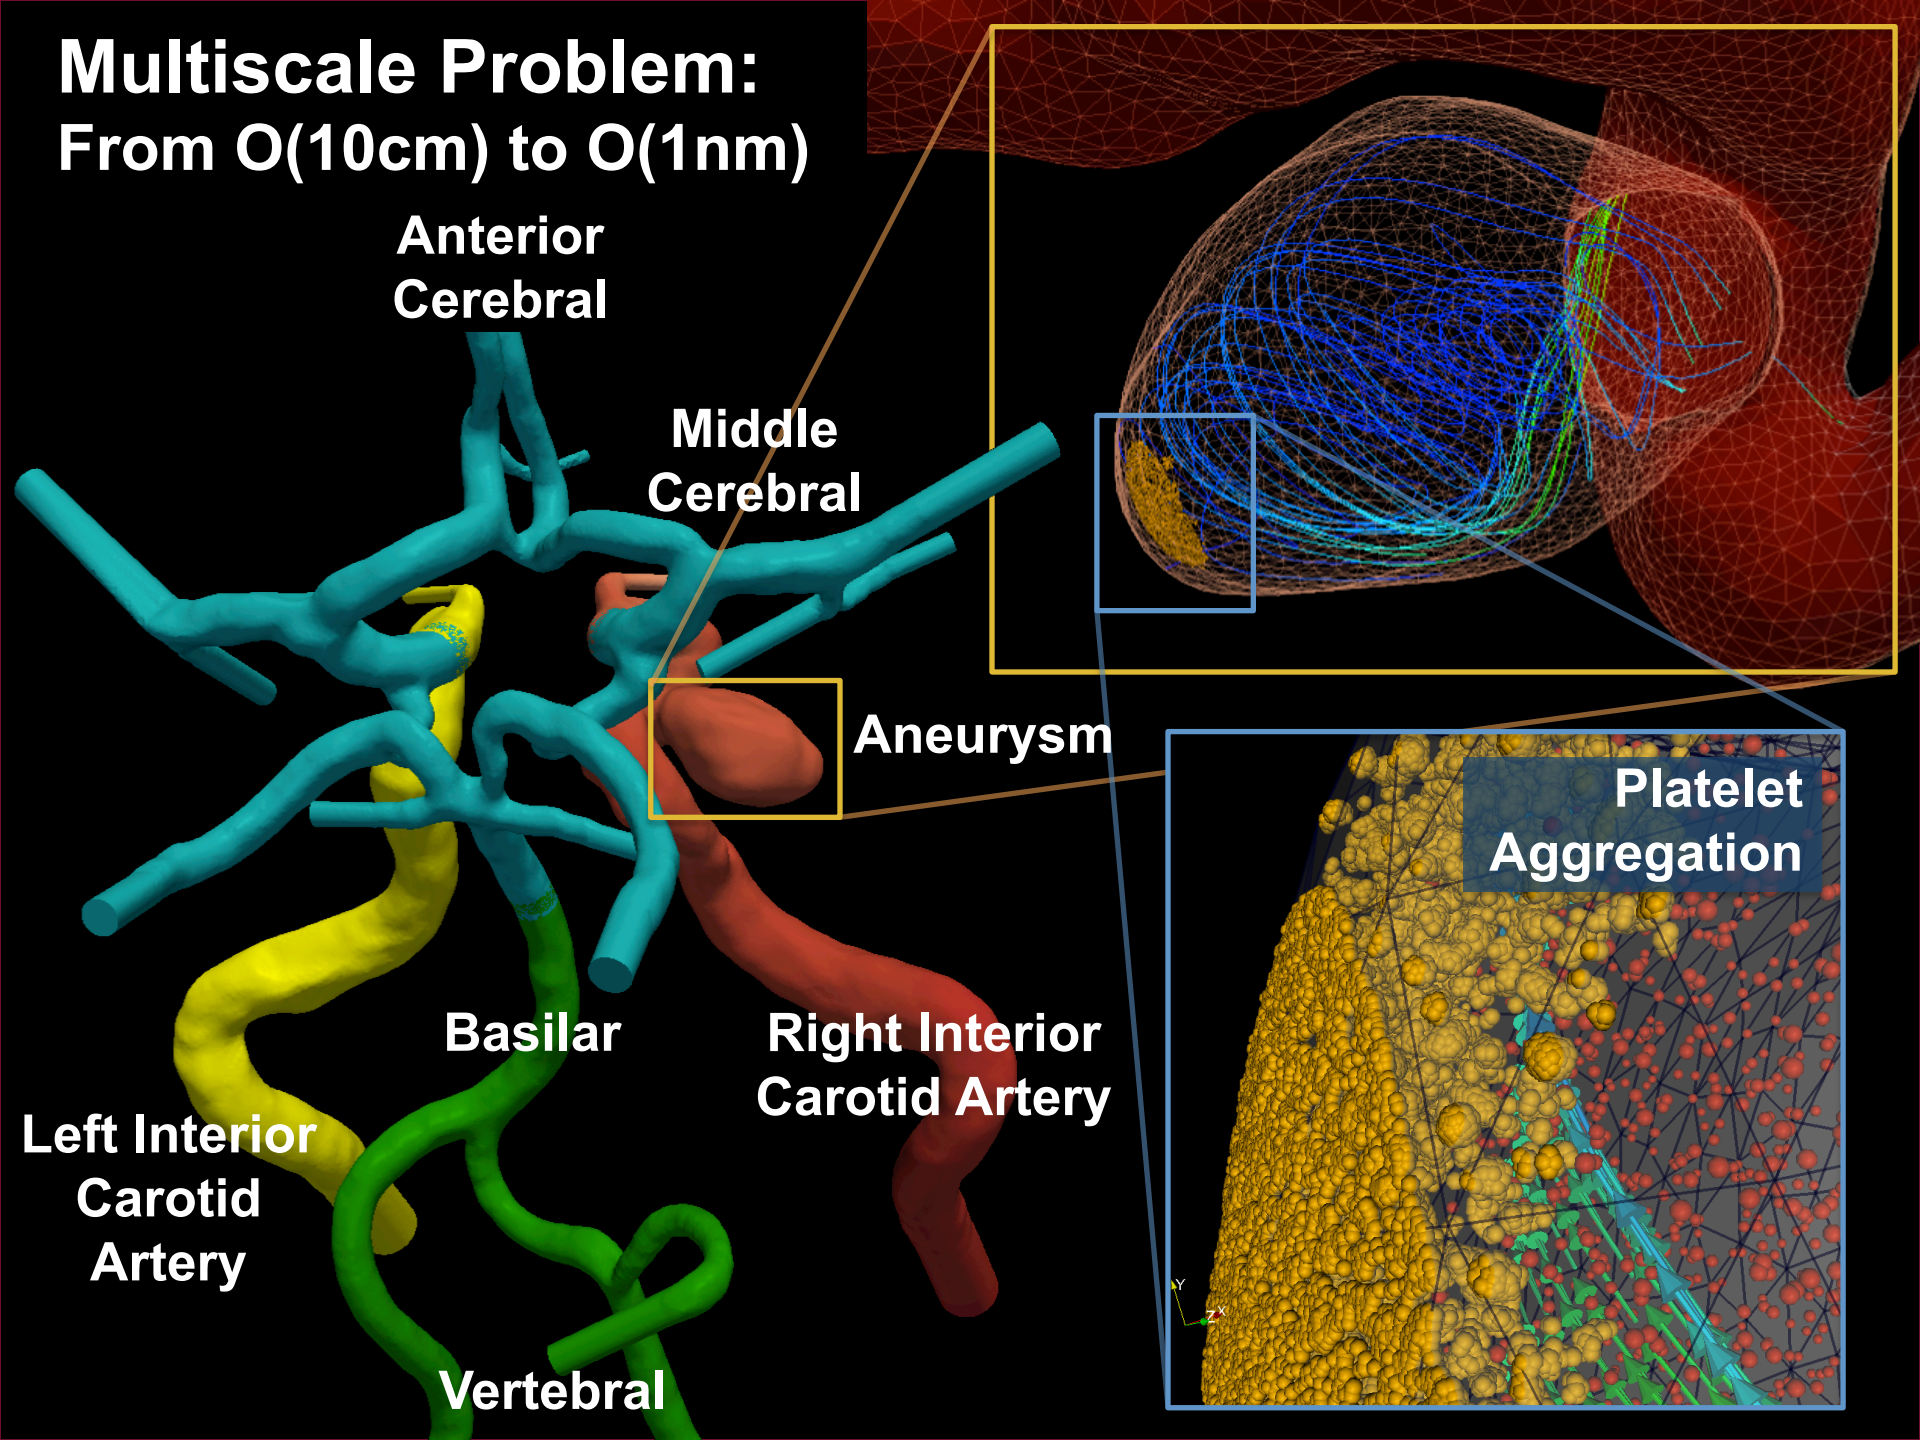

Materials Science Visualization

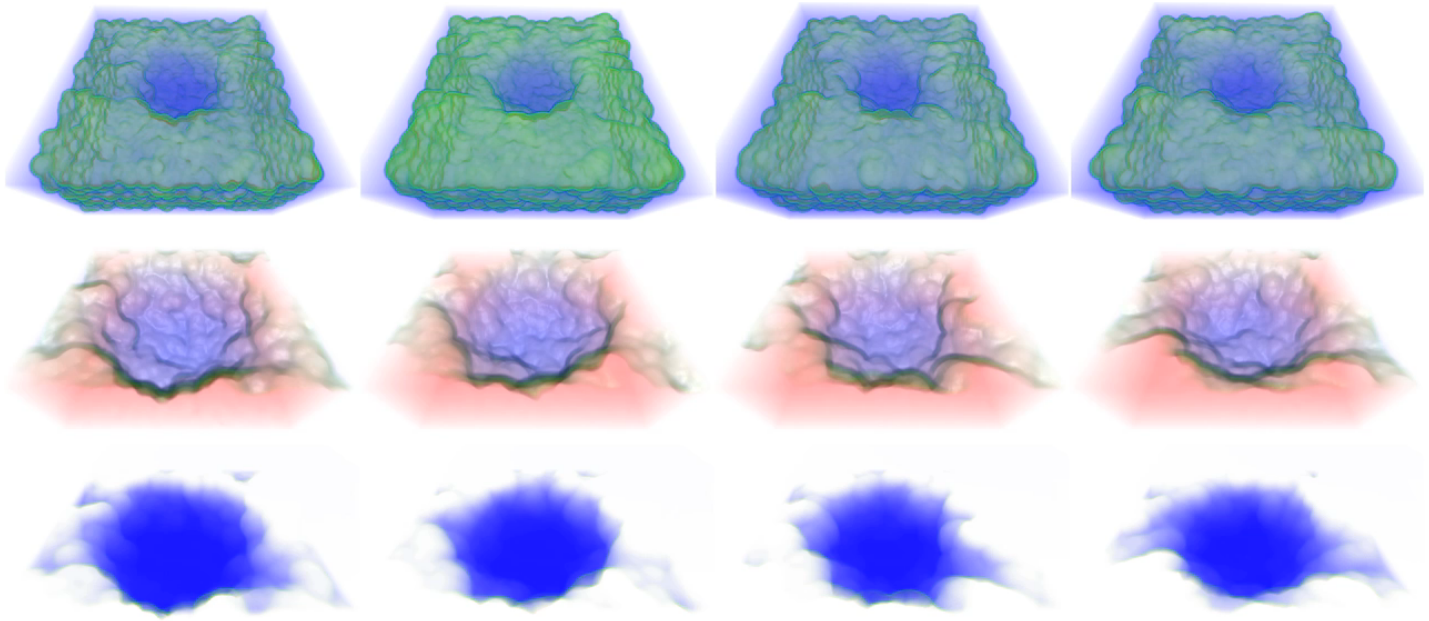

- Nanobowls are synthetic structures designed to trap and hold catalysts, fuel molecules, or other reaction participants
- The goal is to understand effects of temperature on nanobowls
- Each column is a different temperature (1000K - 1350K), and we perform automatic registration and classification of the bowl region

vl3 - A scalable framework for parallel volume rendering on leadership systems

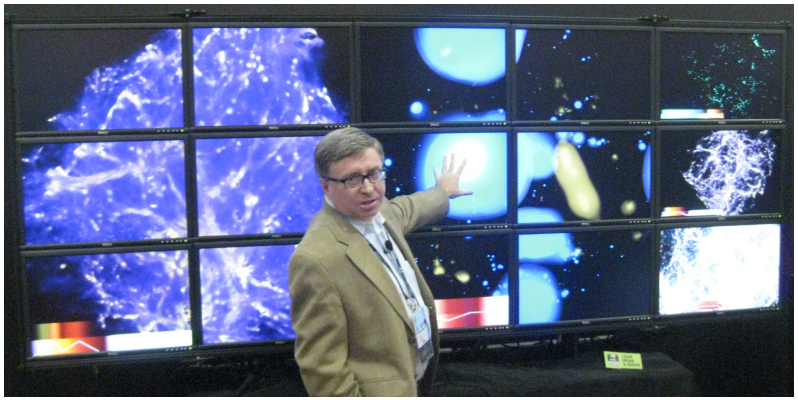

Volume rendering of ENZO using vl3
Courtesy: Joe Insley, Rick Wagner

- Supports laptops, Blue Gene/P supercomputer and GPU-based clusters
- Modular architecture
- Scalable performance to **16384 GPU cores**
- Renders a $(4096)^3$dataset at 2 fps

- Uses a hierarchical parallel compositing mechanism fully exploiting GPUs as well as CPUs
- Supports multiple data formats and has achieved 38 GiBps to read 256 GB ENZO time-steps from storage

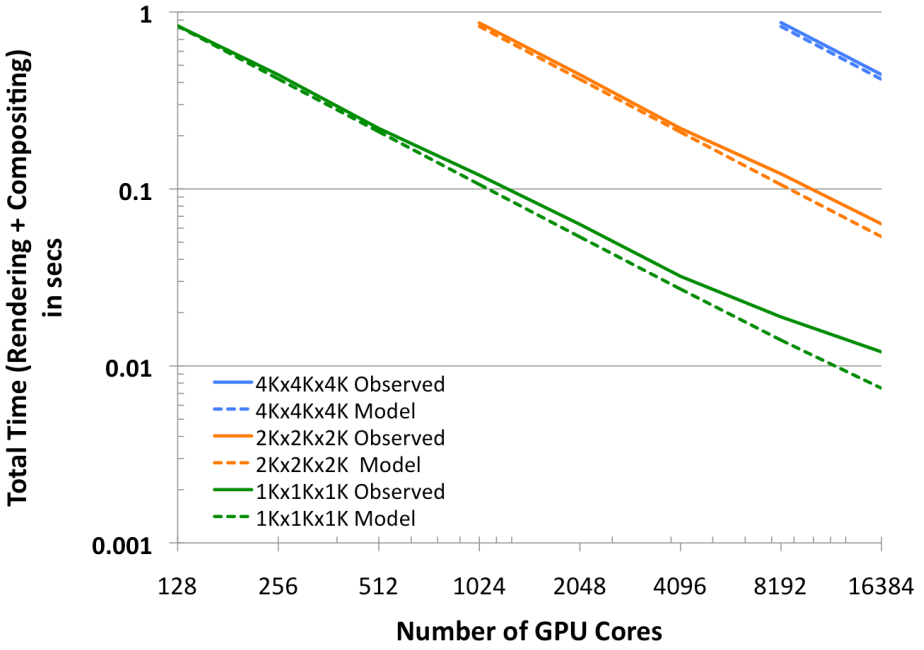

Enzo Astrophysics Simulations

- Four variable Radiation-Hydrodynamics

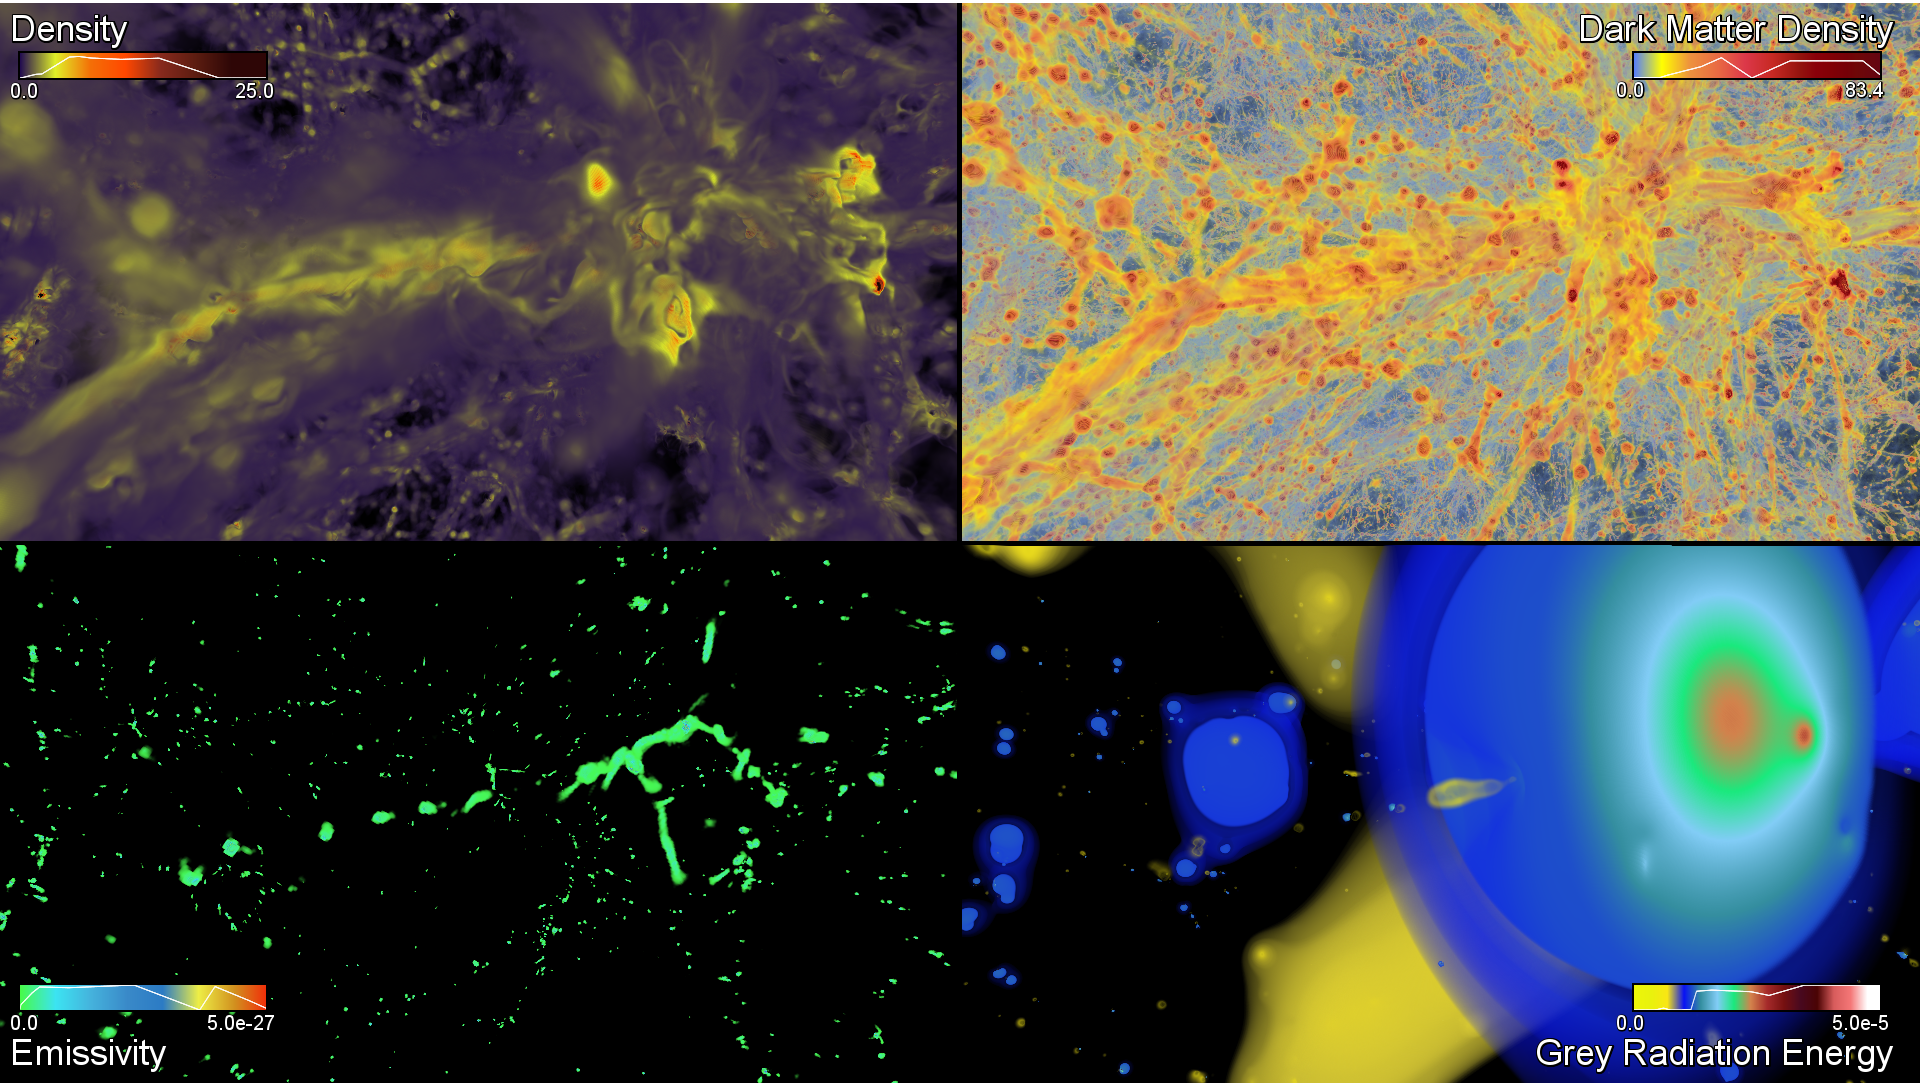

Parallel Topological Analysis

Collaboration with SCI Institute, University of Utah

-Transform discrete scalar field into Morse-Smale Complex
-Nodes are minima, maxima, saddle points of scalar values
-Arcs represent constant-sign gradient flow
-Used to quickly see topological structure
-Never parallelized before

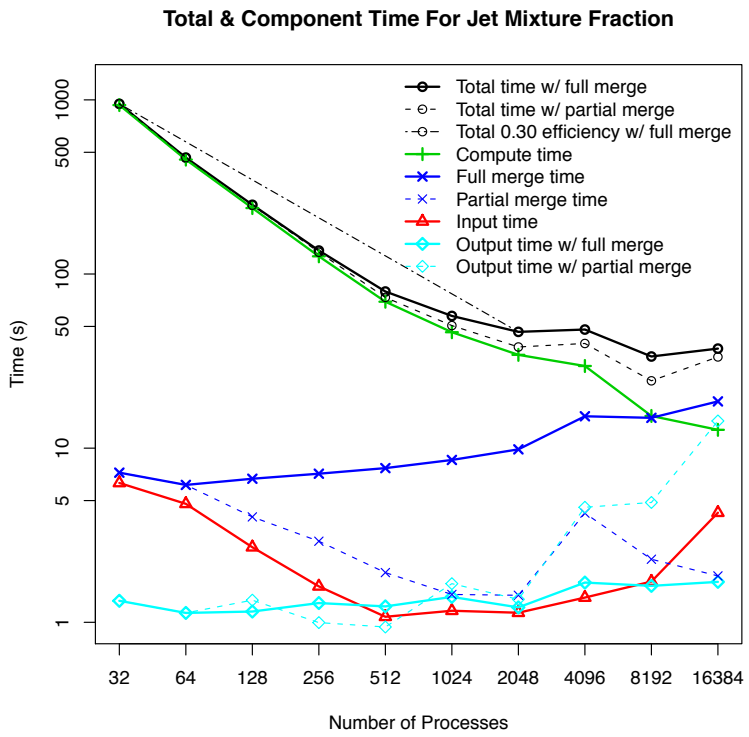

Benchmarked scalability to 16 K processes

Two levels of simplification of the Morse-Smale Complex for jet mixture fraction.

Jet mixture fraction data courtesy Jackie Chen (SNL) and Ray Grout (NREL). Generated by the S3D combustion code. Image courtesy Attila Gyulassy, University of Utah

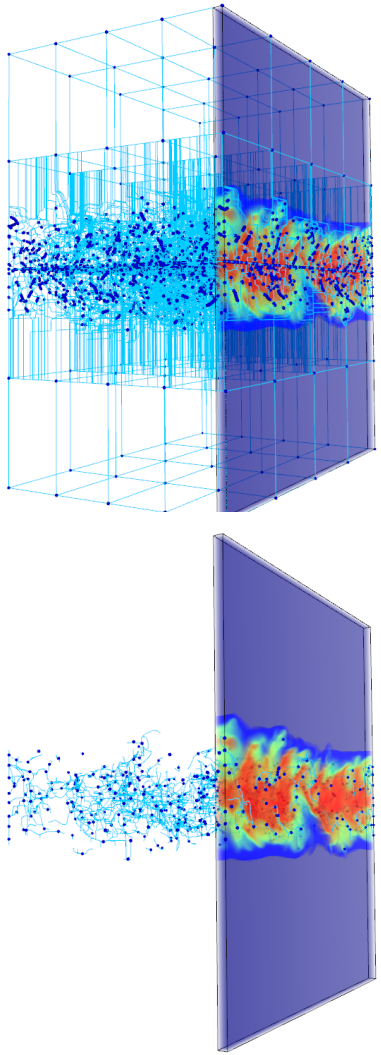

Gyulassy et al. Characterizing the Parallel Computation of Morse-Smale Complexes. Submitted to SC'11

GLEAN- Enabling simulation-time data analysis and I/O acceleration

- A flexible and extensible data analysis framework taking into account application, analytics and system characteristics to perform **the right analysis at the right place and time**
- Provides I/O acceleration by asynchronous data staging
- Scaled to 160K cores and a generic design to enable wide deployment

Infrastructure	Simulation	Mode
Co-analysis	PHASTA	Visualization using Paraview
Staging	FLASH, S3D	I/O Acceleration
In situ	FLASH	Fractal Dimension
In flight	MADBench2	Histogram

Leverages **data models** of applications including adaptive mesh refinement for FLASH and unstructured meshes for PHASTA

Mission: Enhance scientific insight and productivity in computational and experimental sciences via research and development of scalable data analysis, visualization and user interfaces

Michael Papka*, Mark Hereld, Joseph Insley, Aaron Knoll, Eric Olson, Tom Peterka, Rob Ross, Tom Uram, Venkat Vishwanath

*Site representative: papka@anl.gov