Argonne NATIONAL LABORATORY

... for a brighter future

Elements of Scalable Data Analysis and Visualization

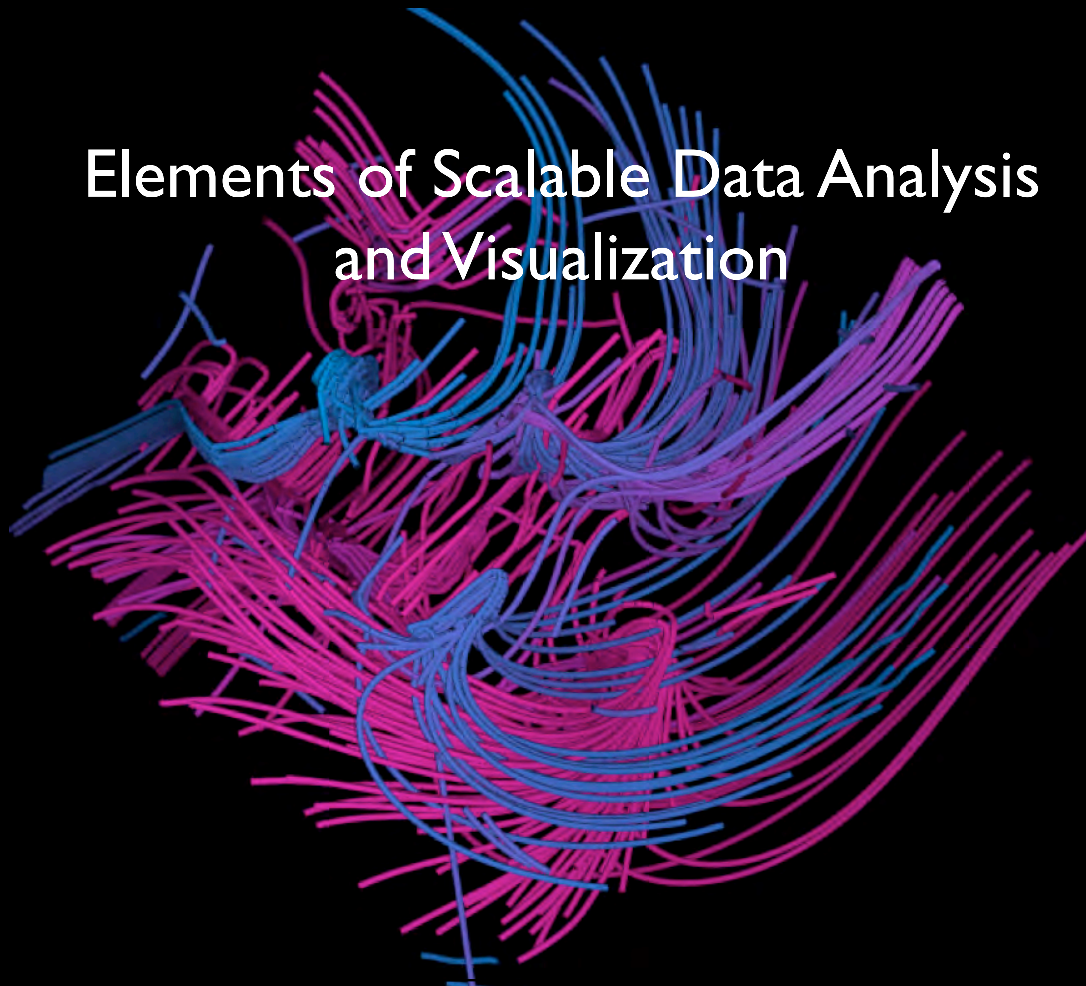

Tom Peterka

tpeterka@mcs.anl.gov

Mathematics and Computer Science Division

DOE CGF Petascale Computing Session

UltraVis SciDAC Institute for Ultrascale Visualization

www.ultravis.org

U.S. Department of Energy

UChicago ▶ Argonne LLC

Office of Science U.S. DEPARTMENT OF ENERGY

A U.S. Department of Energy laboratory managed by UChicago Argonne, LLC

0. Preface

- Science Applications
- Data Limitations

Science Applications

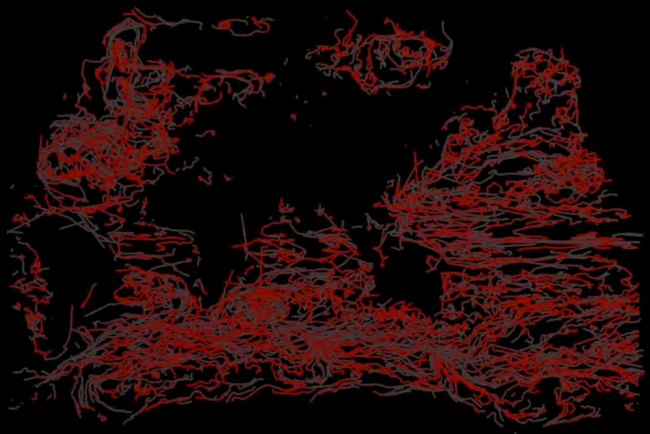

Ocean current data courtesy Rob Jacob, ANL

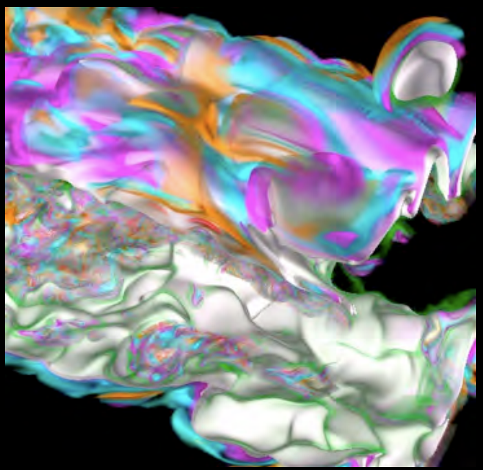

S3D combustion data courtesy Jackei Chen, SNL

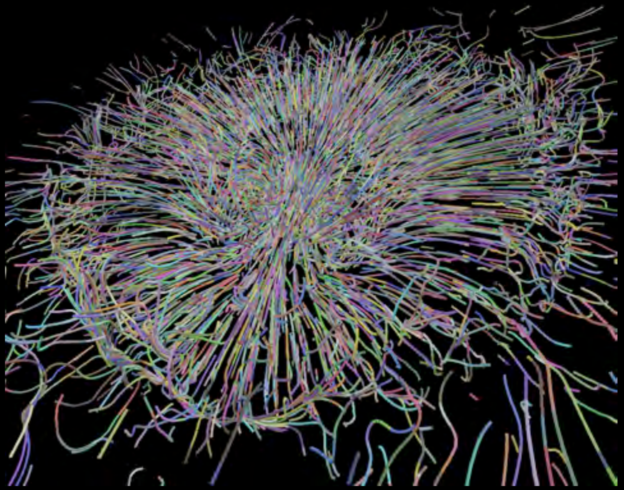

Type IA supernova data courtesy George Jordan, UofC FLASH Center

Core-collapse supernova data courtesy John Blondin, NCSU

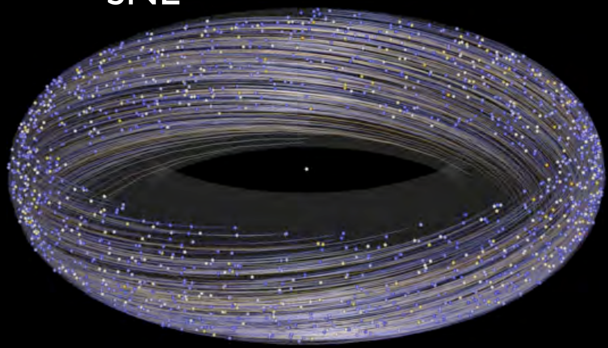

Gyrokinetic data courtesy Stephane Ethier, PPPL

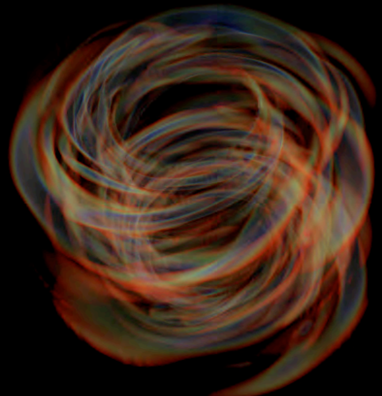

MRI data courtesy Aleks Obabko, ANL.

More than Peak FLOPS: data rate limits analysis capability. Data that is not stored can't be analyzed.

Normalized Storage / Compute Metrics

Machine	FLOPS (Pflop/s)	Storage B/W (GB/s)	Flops per byte stored	Bytes comp. per byte stored
LLNL BG/L	0.6	43	$O(10^4)$	$O(10^3)$
Jaguar XT4	0.3	42	$O(10^4)$	$O(10^3)$
Intrepid BG/P	0.6	50	$O(10^4)$	$O(10^3)$
Roadrunner	1.0	50	$O(10^5)$	$O(10^4)$
Jaguar XT5	1.4	42	$O(10^5)$	$O(10^4)$

-In 2001, Flops per bytes stored was approximately 500. Ref: John May, 2001.

-DOE science applications generate results at an average rate of 40 flops per byte of data. Ref: Murphy et al. ICS'05.

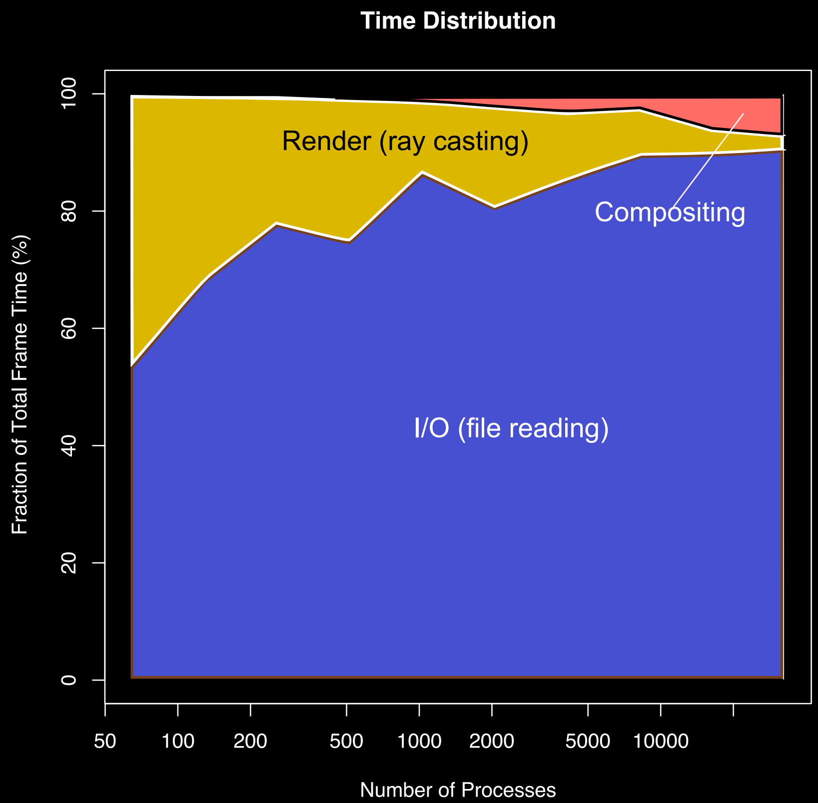

The relative percentage of time in the stages of volume rendering as a function of system size. Large visualization is primarily dominated by data movement: I/O and communication.

Key Elements for Scalability

1. Moving Data
2. Transforming Data
3. Understanding Data

1. Moving Data

- Reorganize access pattern

- Reorganize data for faster access

Data Aggregation

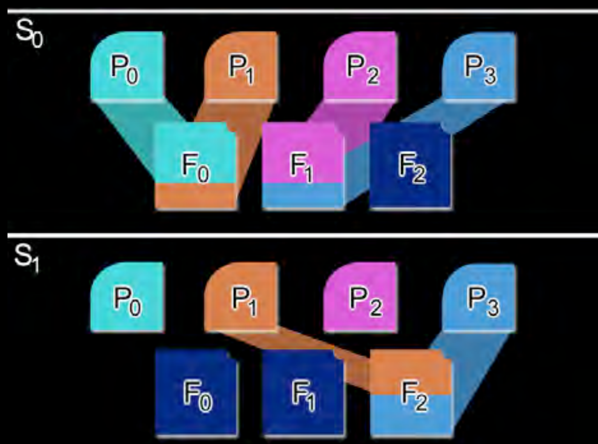

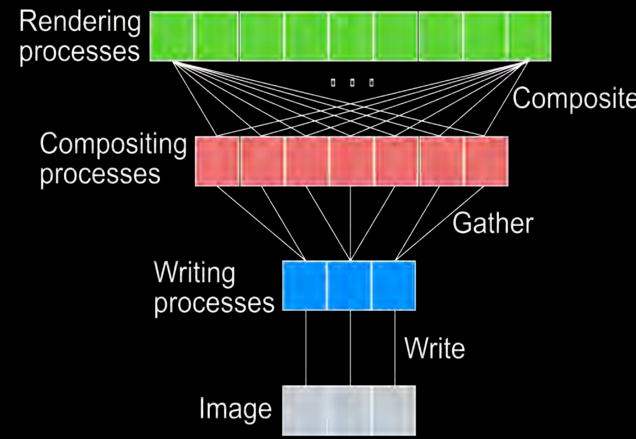

Aggregate I/O Bandwidth on Intrepid

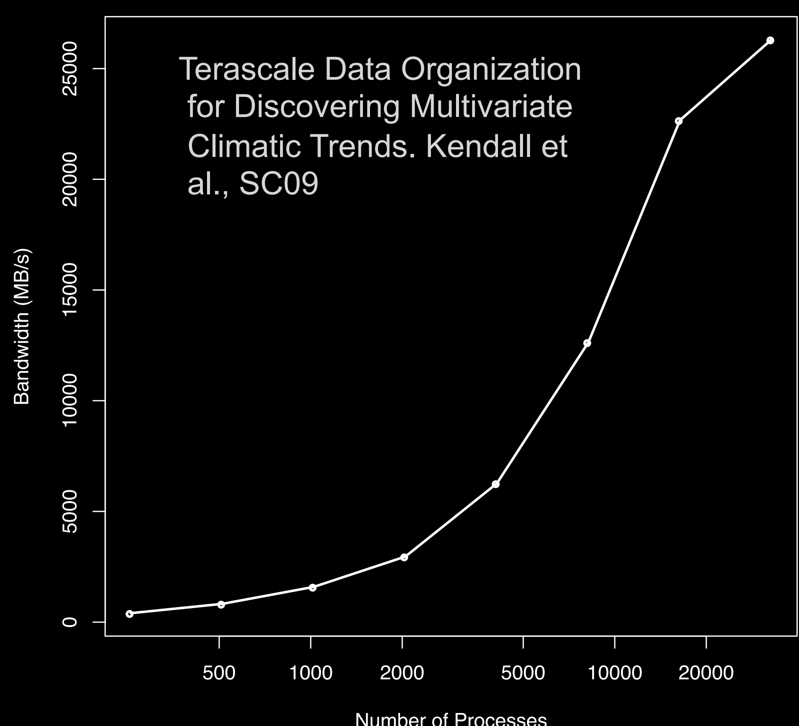

Output Time for Varying Numbers of Writing Processes

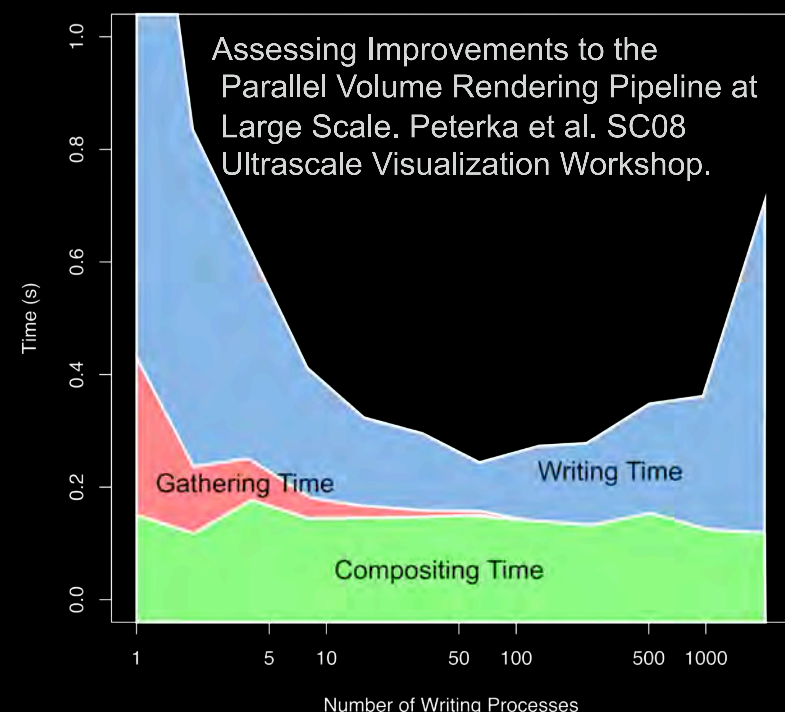

Data Reordering

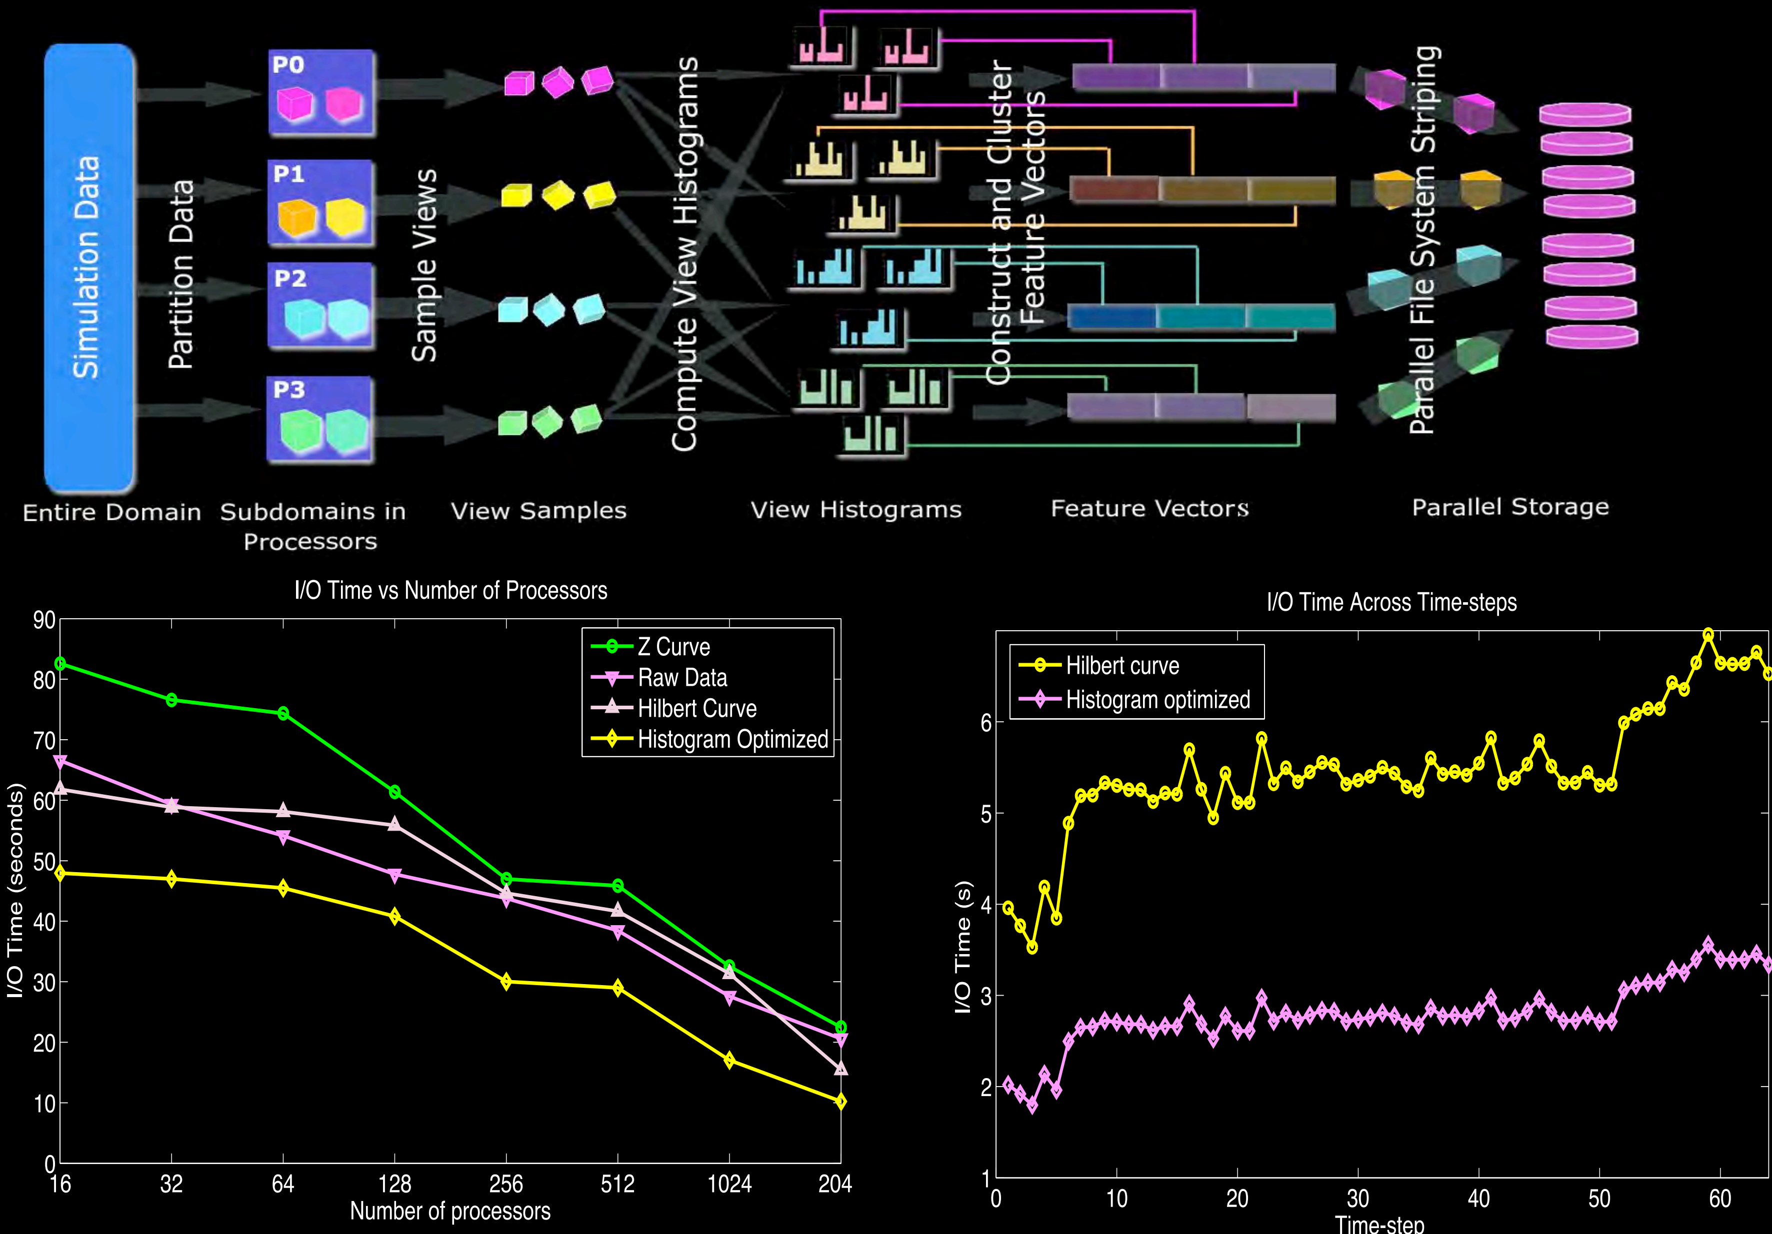

Histogram-based I/O Optimization for Visualizing Large-Scale Data. Hong et al. SC09 Ultrascale Visualization Workshop.

2. Transforming Data

- Parallelism is the single best tool we have
- Similar to and integrated into other parallel applications

Parallel Particle Tracing

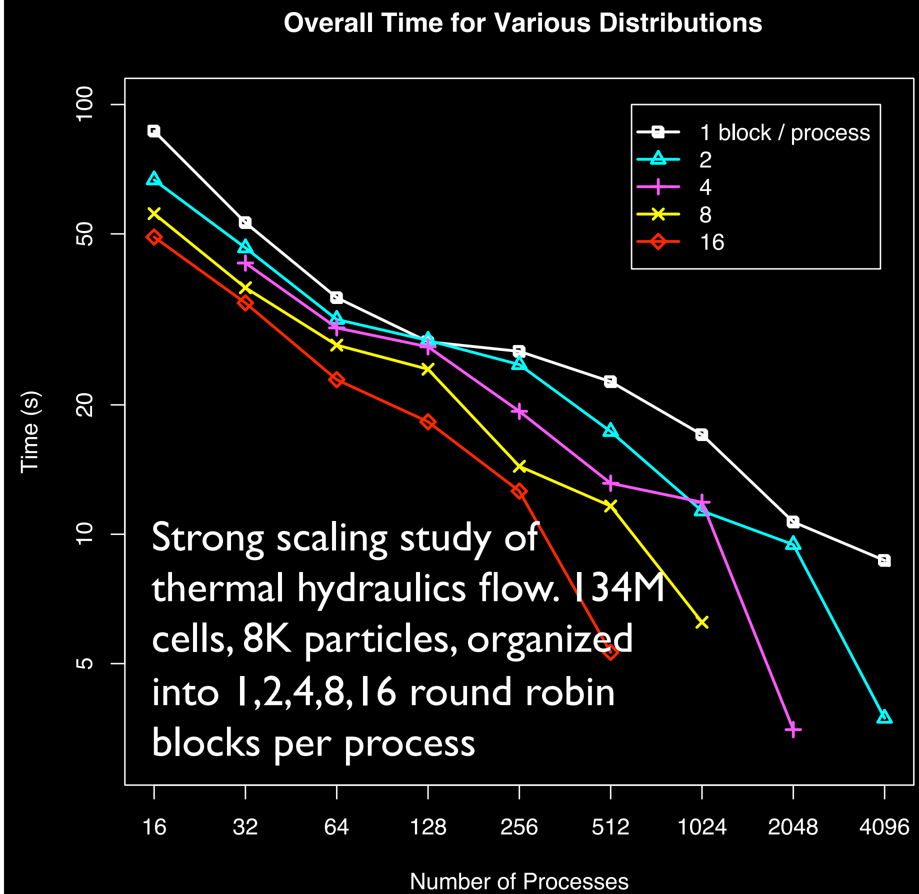

Strong scaling study of thermal hydraulics flow. 134M cells, 8K particles, organized into 1,2,4,8,16 round robin blocks per process

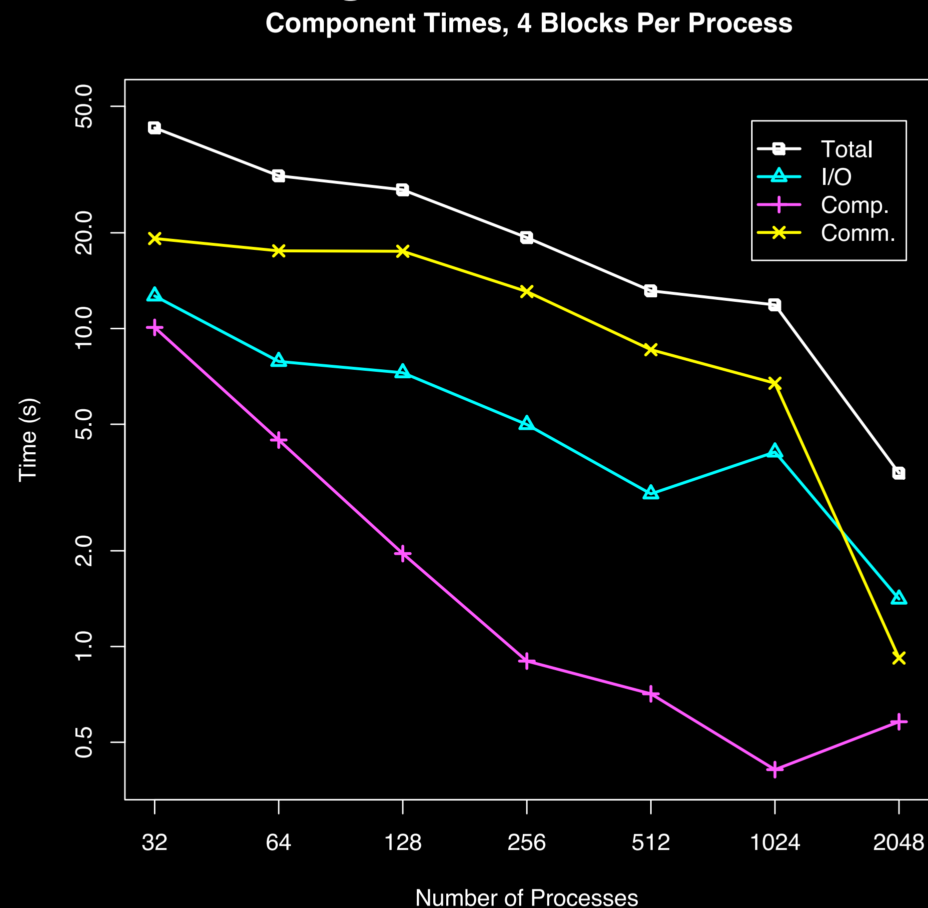

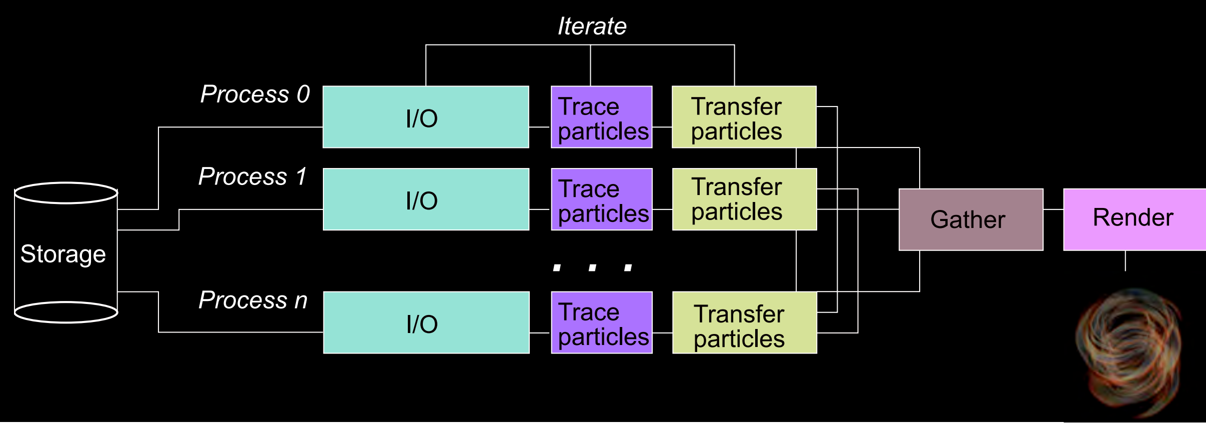

Parallel structure for flow visualization algorithm consists of iterations of particle tracing and transfer, followed by a rendering stage.

Prototyping In Situ Visualization

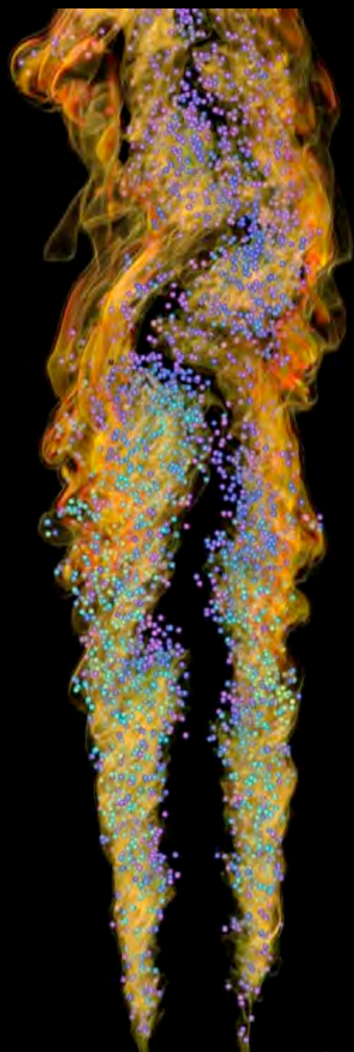

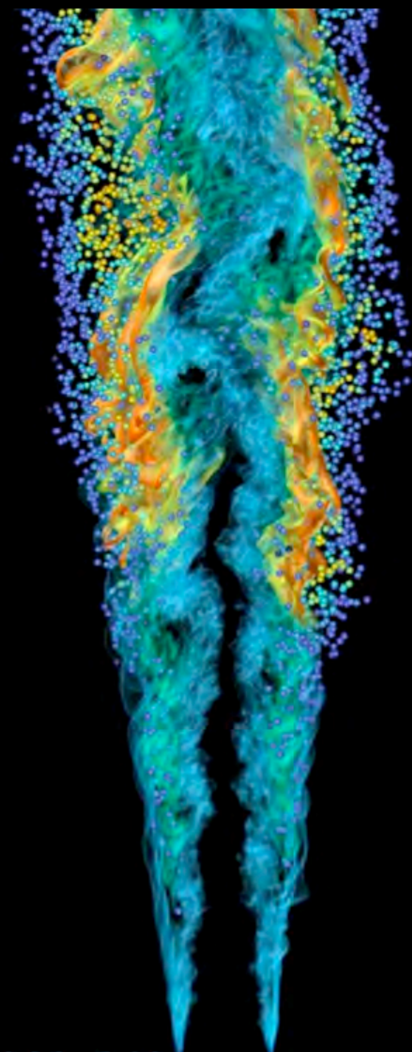

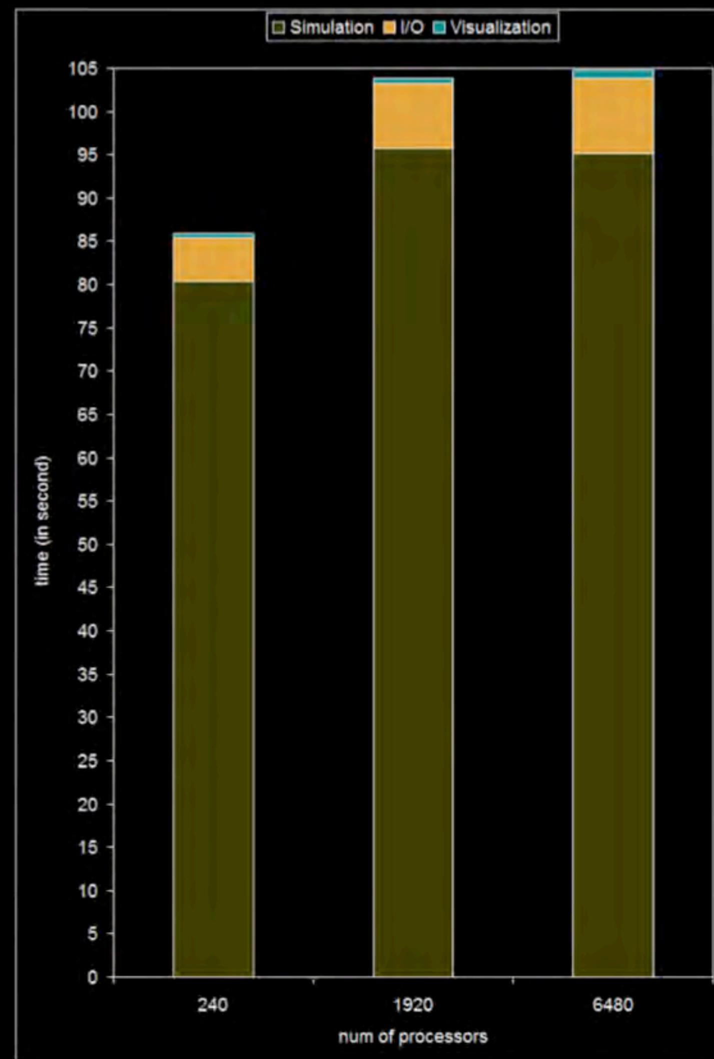

Simulation of S3D combustion courtesy of Jackie Chen with in situ visualization of volume and particle data that requires 7% of overall computing time at up to 6480 processors on Cray XT5 at NCCS. (for 1024^2 images)

In Situ Visualization for Direct Numerical Simulation of Turbulent Combustion. Yu et al., SC09 Ultrascale Visualization Workshop.

3. Understanding Data

- Libraries, specific- and general-purpose tools

- Environments for exploration

Interfaces for a Variety of Applications and Users

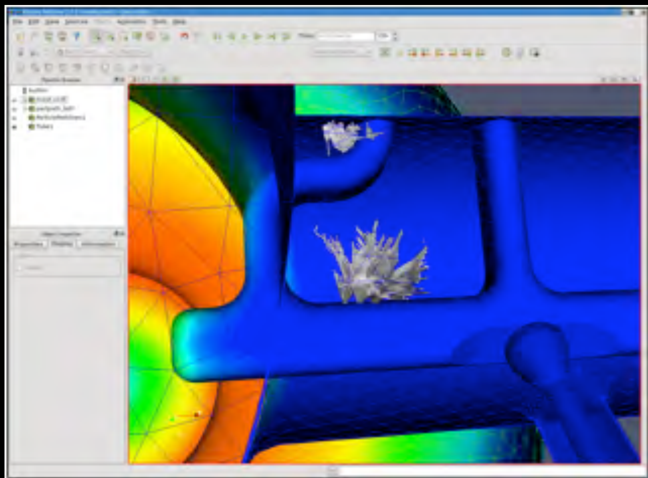

Sometimes a full-featured tool is appropriate

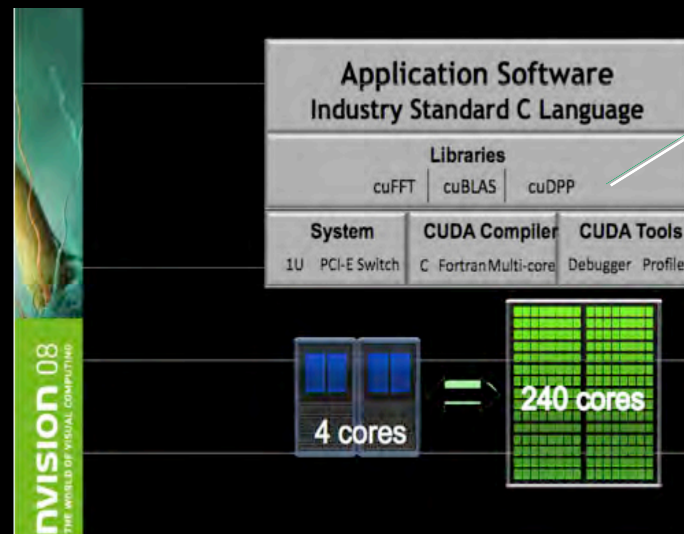

While other times, low level libraries such as cuDPP enable operations such as prefix sums and stream compactions to be performed efficiently on GPUs and can be more scalable.

Spatial Domain

Variable Domain

Temporal Domain

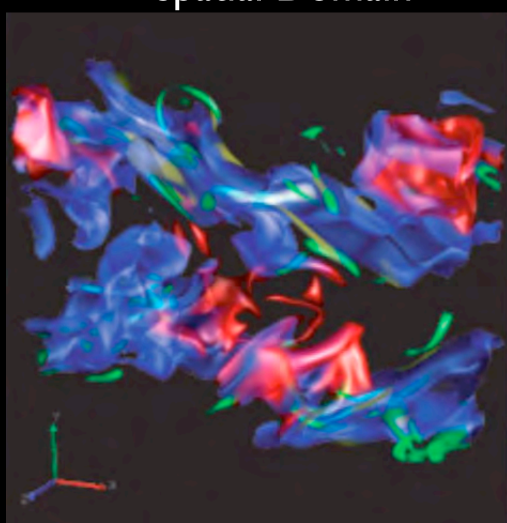

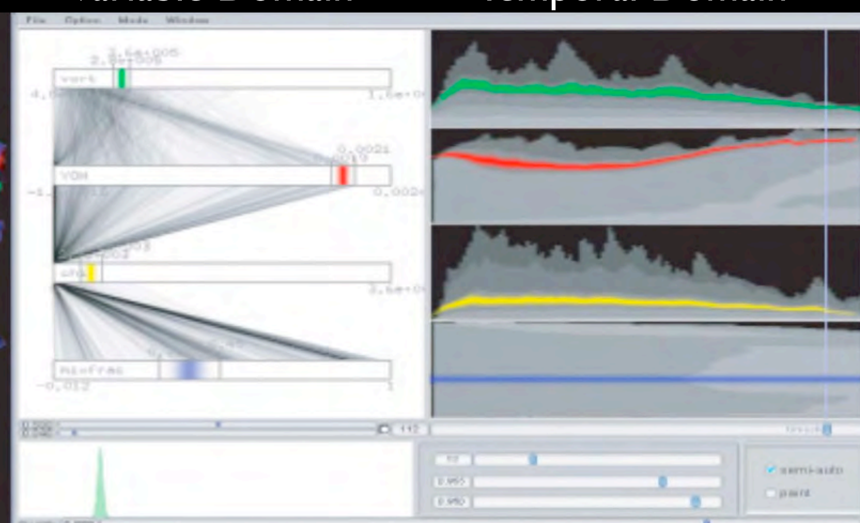

Interfaces designed to perform tasks tailored to the scientist's workflow are another option.

Immersive Workspaces

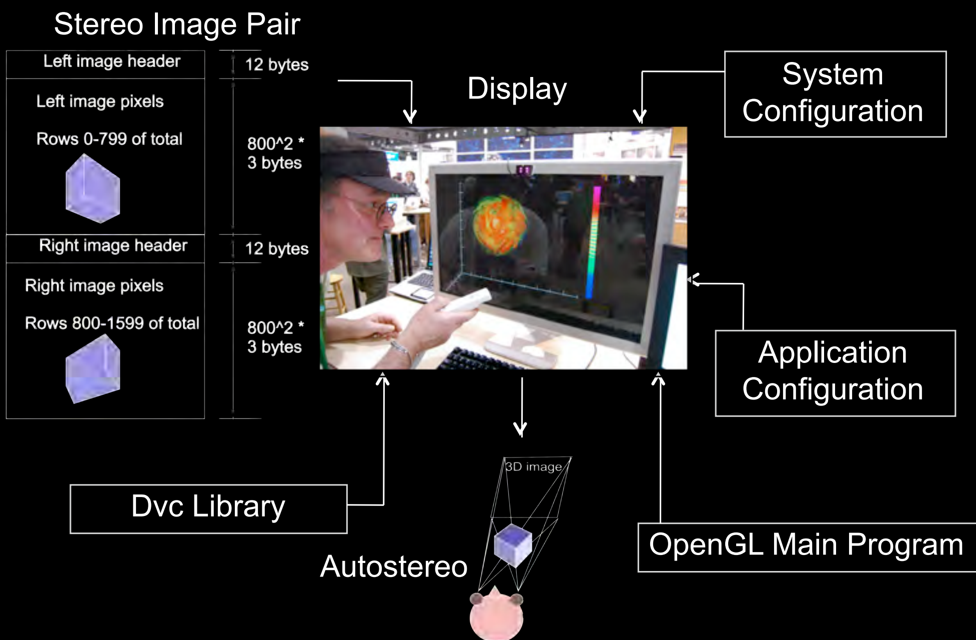

Stereo parallel volume rendering: The server (BG/P) computes stereo pairs of volume-rendered images and streams them to the client, which runs the dvc library to display them remotely in autostereo.

Display devices and interaction techniques bring virtual environments to scientific visualization.

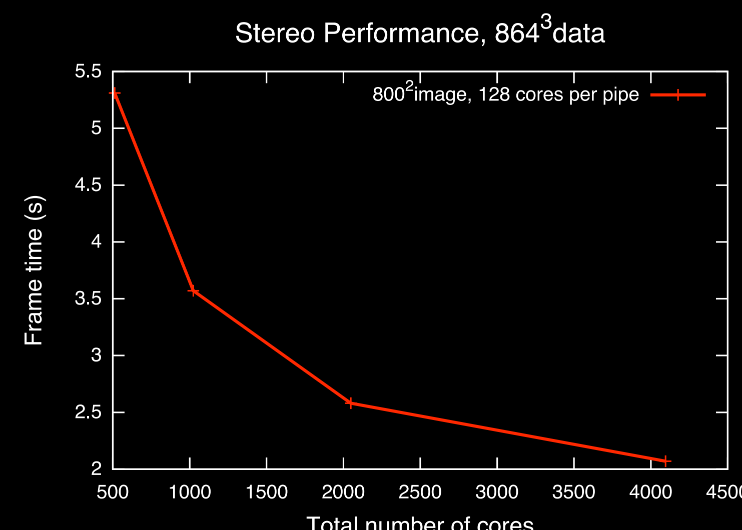

Scalability in end-to-end pipeline from BG/P to autostereo display

Display of Large-Scale Scientific Visualization. Peterka et al. SPIE'09

The Computational Science Interface as a Whole

A linear, sequential pipeline where tasks mapped to architectures in fixed fashion

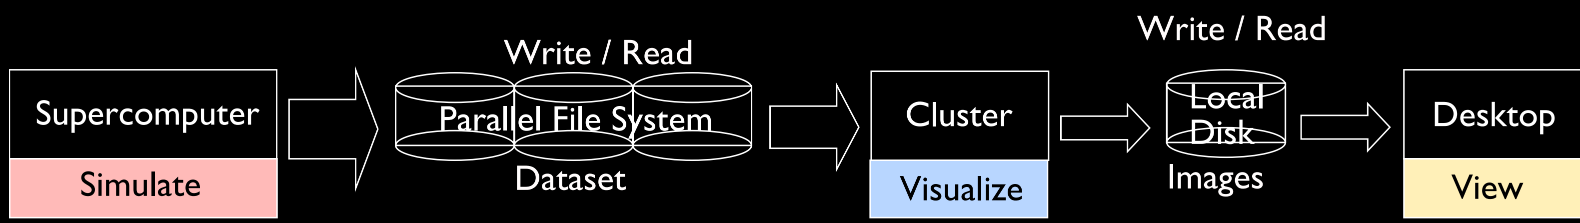

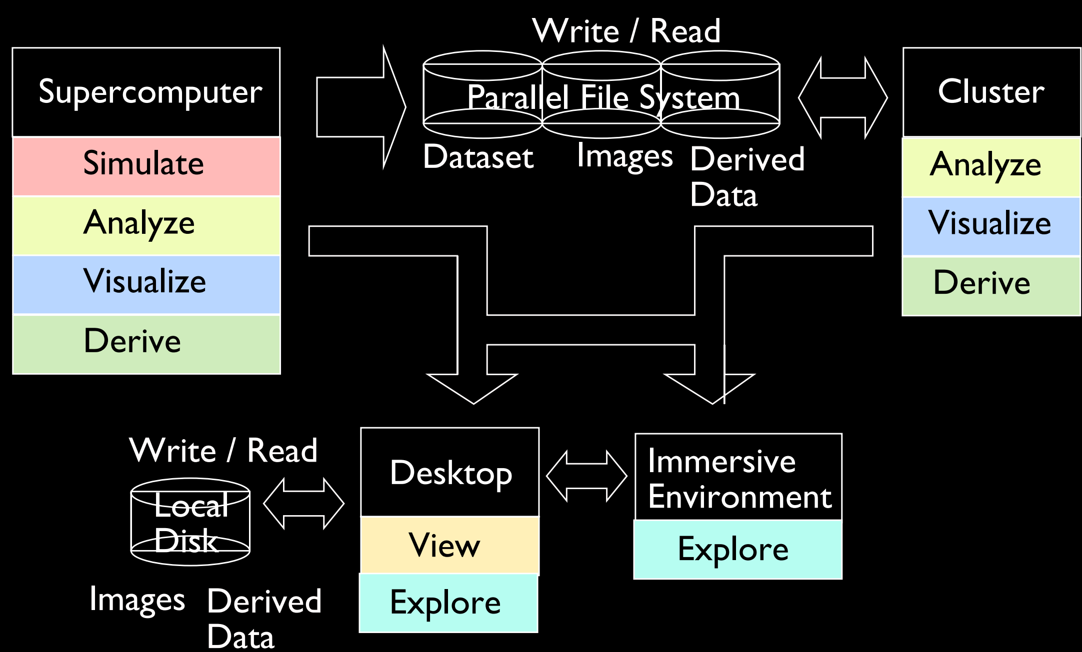

A complex network with flexible mappings of tasks to architectures, remote collaborators, and exploratory science environments

Elements of Scalable Data Analysis and Visualization

Argonne NATIONAL LABORATORY

... for a brighter future

Thank you

Acknowledgments:

Facilities

Argonne Leadership Computing Facility (ALCF)
Oak Ridge National Center for Computational Sciences (NCCS)

Funding

US DOE SciDAC UltraVis Institute

People

Kwan-Liu Ma, Rob Ross, Han-Wei Shen, Jian Huang, Hongfeng Yu, Wes Kendall, Jackie Chen, John Blondin, Rob Jacob, Aleks Obabko, and many, many others

Tom Peterka

tpeterka@mcs.anl.gov

Mathematics and Computer Science Division

DOE CGF Petascale Computing Session

UltraVis SciDAC Institute for Ultrascale Visualization

www.ultravis.org

U.S. Department of Energy

UChicago Argonne LLC

Office of Science U.S. DEPARTMENT OF ENERGY

A U.S. Department of Energy laboratory managed by UChicago Argonne, LLC