LA-UR-14-26861
Approved for public release; distribution is unlimited.

Title: Implications of Numerical and Data Intensive Technology Trends on Scientific Visualization

Author(s): Ahrens, James Paul

Intended for: Texas Advanced Computing Center Seminar, 2014-08-21 (Austin, Texas, United States)

Issued: 2014-09-02

Los Alamos
NATIONAL LABORATORY
EST. 1943

Implications of Numerical and Data Intensive Technology Trends on Scientific Visualization

James Ahrens

August 2014

Data Science at Scale Team and Collaborators

- Boonthanome Nouanesengsy, John Patchett, David Rogers, Patricia Fasel, Patrick O'Leary, Seb Jourdian, Chris Sewell, Jon Woodring, Christopher Mitchell, Ollie Lo, Kary Myers, Joanne Wendelberger, Curt Canada, Laura Monroe, Dave Modl, Bob Kares, Bob Greene, Paul Weber, Francesca Samsel, Garrett Aldrich, Ayan Biswas, Chris Bryan, Lalindra De Silva, Daniel Hill, Evgeni Makevnin, Denis Mosbach, Jesus Pulido, Andre Schmeisser, Wathsala Widanagamaachchi, Max Zeyen
- Marcus Daniels, Hilary Abhold, Gabe Rockefeller

Scope of this talk

- Technology trends in numerically and data intensive computing have the potential:
  - to reshape and significantly improve how we visualize and analyze scientific data

1. Numerically-intensive / exascale challenges and solutions

2. Data-intensive / cloud challenges and solutions

My background

- Originator and lead designer for ParaView
  - Open source large data visualization package
- Leadership roles in DOE
  - U.S. supercomputing planning
    - Exascale – 10^18 flops
  - White House briefings/white papers
    - On Big Data, High Performance Computing
    - J. Ahrens, B. Hendrickson, G. Long, S. Miller, R. Ross, D. Williams, "Data Intensive Science in the Department of Energy", A white paper for the Office of Science and Technology Policy.
    - On Big Data, High Performance Computing and Privacy
      - John Podesta led 90 day study

What are the challenges in your work?

A. Relentless increase in data sizes

- 3 orders of magnitude every ten years

B. Adapting to changing software & hardware infrastructure (shared memory, clusters, threading, data-intensive)

C. Advancing the fundamentals

- Improved end-to-end visualization and analysis workflow and cognitive understanding

Mega	Giga	Tera	Peta	Exa
10^6	10^9	10^12	10^15	10^18
Image speed	Storage & network speed	Operations speed	Operations speed	Operations speed

Supercomputing System Architecture Diagram

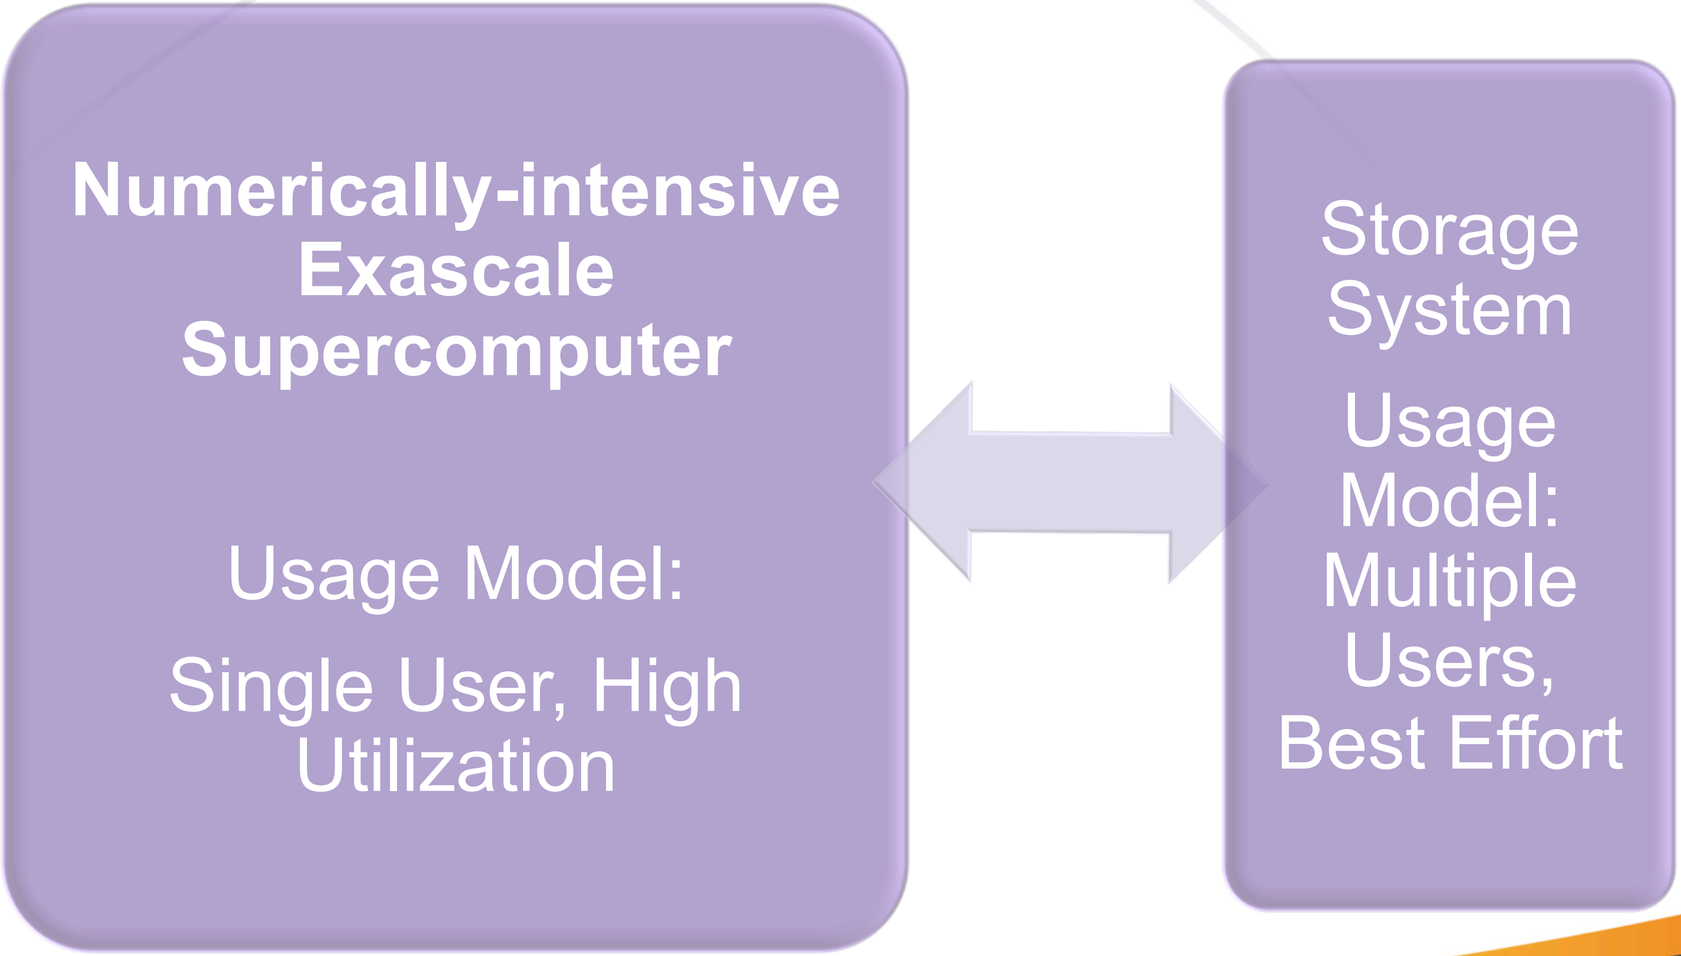

Numerically-Intensive / Exascale Computing
The Vision – Bill Herrod

Achieve order 10^{18} operations per second and order 10^{18} bytes of storage

– Address the next generation of scientific, engineering, and large-data workflows

- Deployed in early 2020s

Exascale Challenges

- Power is very costly
  - need 20 pJ per average operation
    - ~x 40 improvement over today's efficiency
  - Peak performance to increase 3 orders of magnitude -- System power to increase by a factor of 2
  - Most expensive power operation is data movement
- Achieve billion-way concurrency
- High reliability and resilience

Numerically-Intensive / Exascale Challenges for Visualization

- Traditional post-processing visualization approach is unworkable…
  - Temporal simulation snapshots are saved at intervals
    - Snapshots cannot answer arbitrary analysis questions
      - Intervals are getting longer
    - Full checkpoints are costly

Meta	Giga	Tera	Peta	Exa
10^6	10^9	10^12	10^15	10^18
Image size		Data size	Data size	Data size

Numerically-intensive / exascale solutions

- Solutions focus on power savings by reducing data movement
  - Also will improve analysis quality
- In situ analysis
  - Automatic
- Data reduction
  - Sampling
  - Compression

In Situ

- In situ saves reduced-sized data products during simulation run
  - Benefits:
    - Saving disk space
    - time in post-processing analysis
    - producing higher fidelity results
- Automatic visualization and analysis during the simulation run
  - Prioritized by scientist's importance metrics
  - Identify specific analysis questions
- Help manage cognitive and storage resource budget

Supercomputing System Architecture Diagram

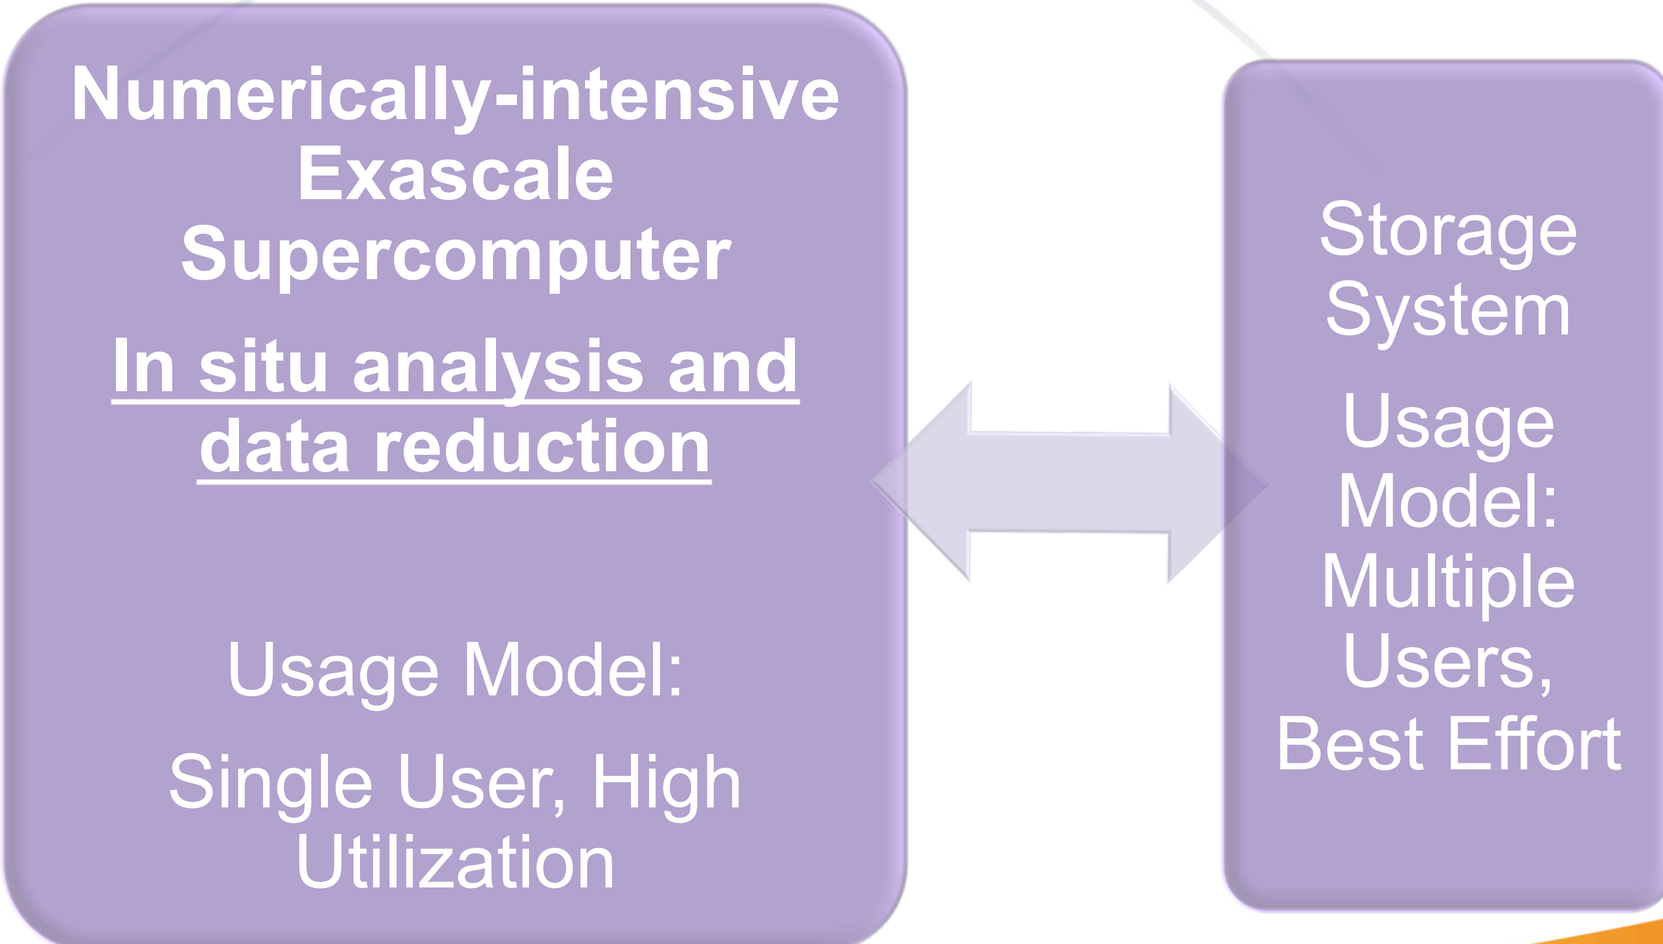

Demonstration of 3D *In Situ* in xRAGE

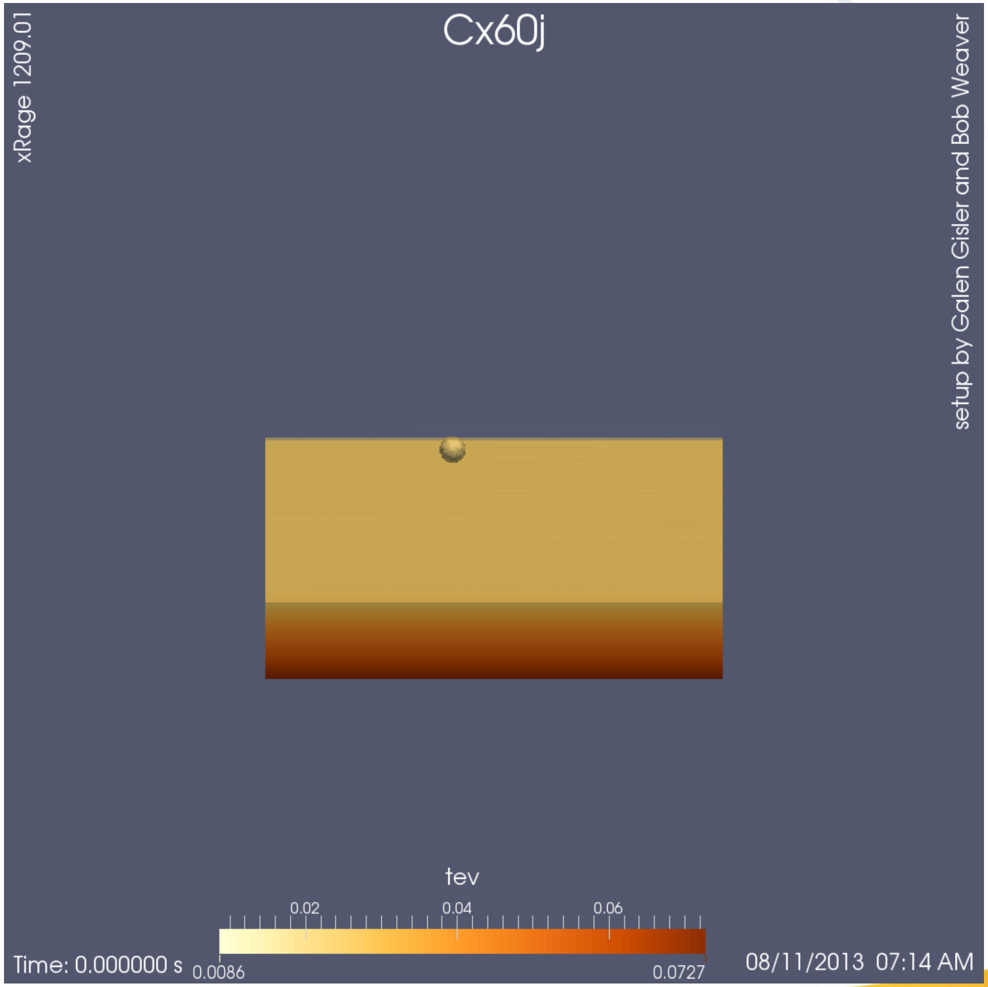

xRAGE *in situ* Weak Scaling on Cielo

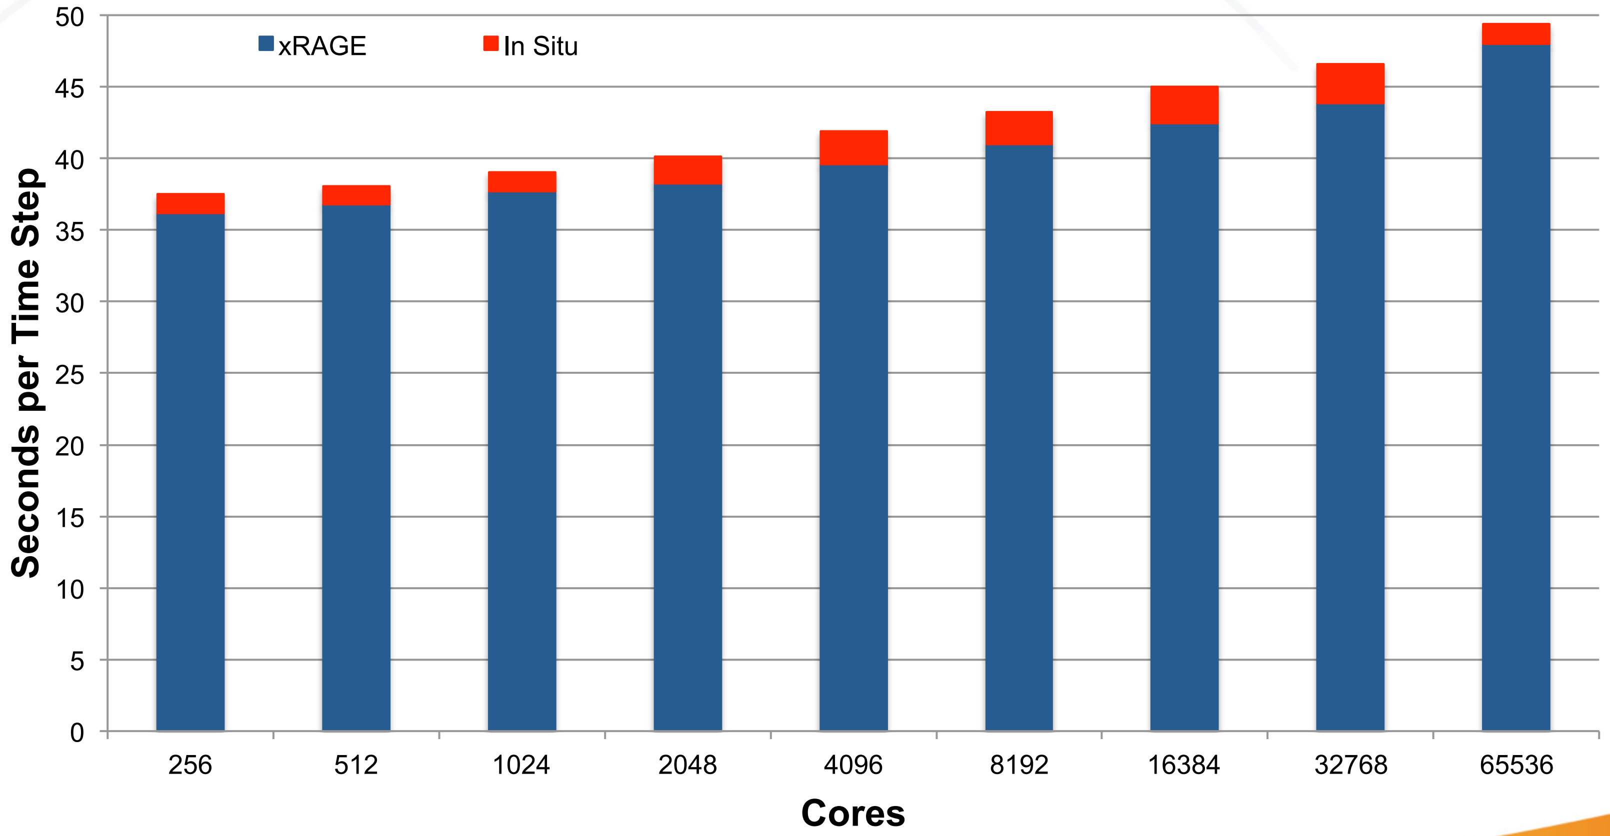

Costs of *in situ* xRAGE

Time of Processing	Type of File	Size per File	Size per 1000 time steps	Time per File to Write at Simulation
Post	Restart	1,300 MB	1,300,000 MB	1-20 seconds
Post	Ensight Dump	200 MB	200,000 MB	> 10 seconds
In Situ	PNG	.25 MB	250 MB	< 1 second

Decision Making When to Save and What

- Create a metric to indicate state of information
  - Trigger In Situ when the metric indicates sufficient change vs. last data product creation
  - Change Metric can be applied per variable

```
for each scalar field
   calculate a metric
   difference the metric from last save
   if difference > threshold
      produce data products defined for that scalar field
```

In situ driver: Adaptive focus based on selected scientific metrics

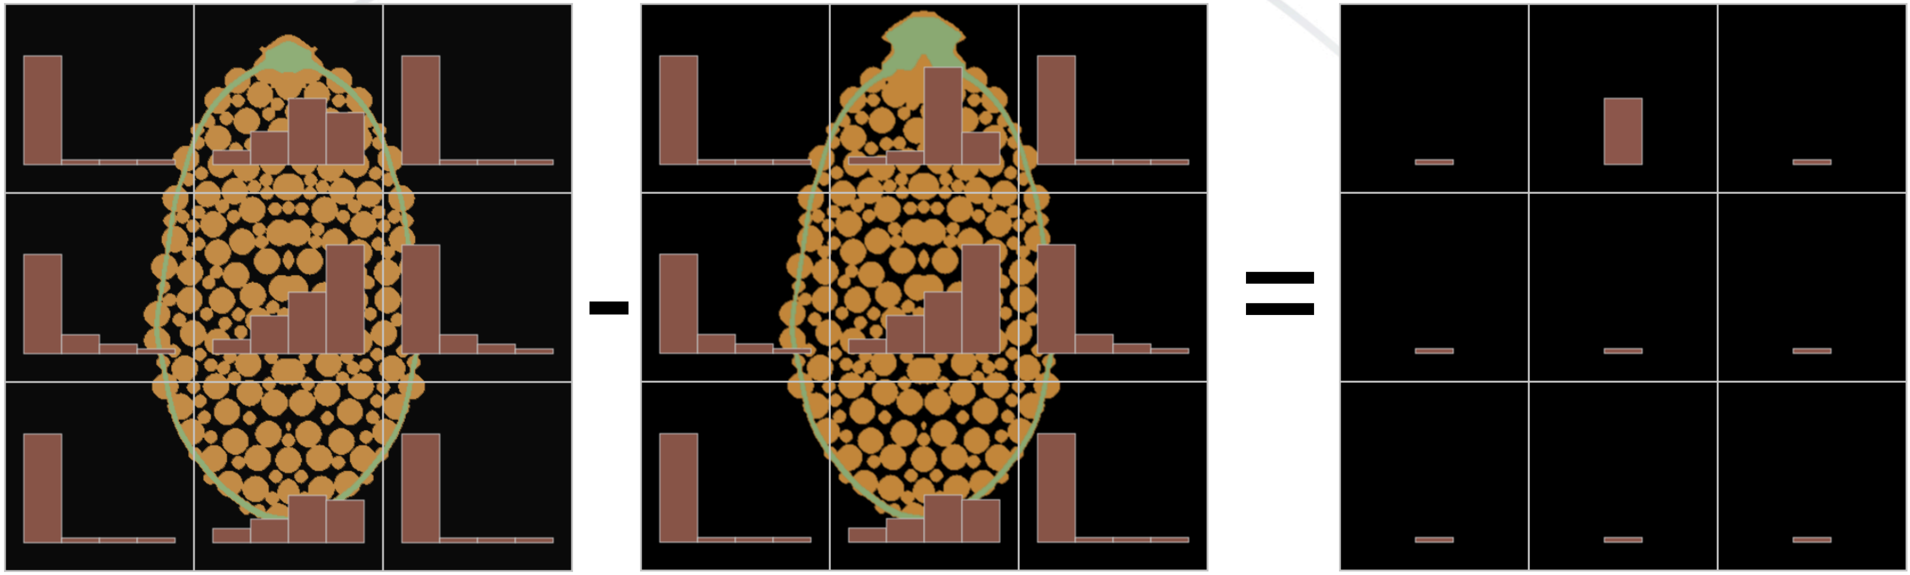

timestep$_n$ — timestep$_{n+1}$ = difference histogram

- Create adaptive analysis-based grid
  - Histogram at each grid element - across all axises (spatial, value, multivariate)
- Use for spatial, temporal selection
  - Cameras, storage, feature identification

Woodring et al, VisWeek – Large Data Analysis Visualization 2014

Demonstration of our *in situ* visualization and analysis capability applied to xRAGE

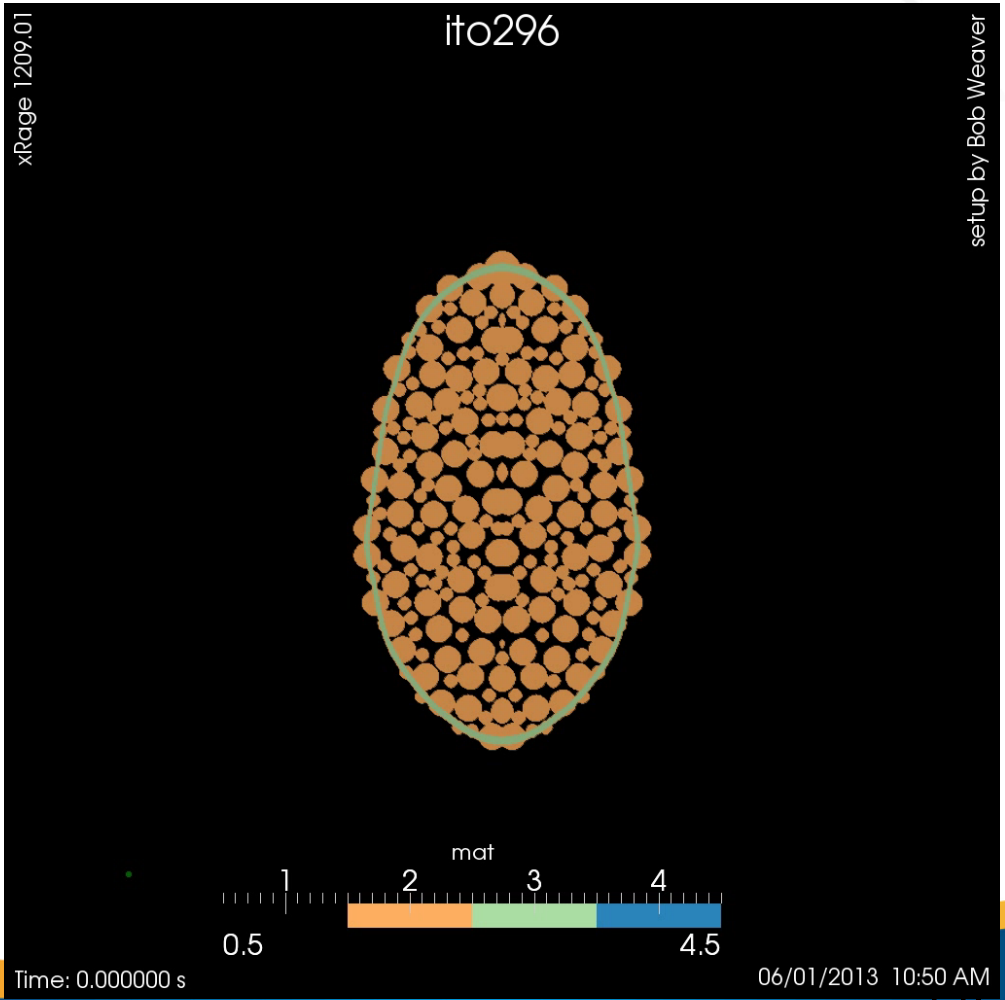

Solution: Statistical Sampling

- Random sampling provides a data representation that is unbiased for statistical estimators, e.g., mean and others
- Since the sampling algorithm is done in situ, we are able to measure the local differences between sample data and full resolution data
- (Simulation Data – Sampled Representation) provides an accuracy metric

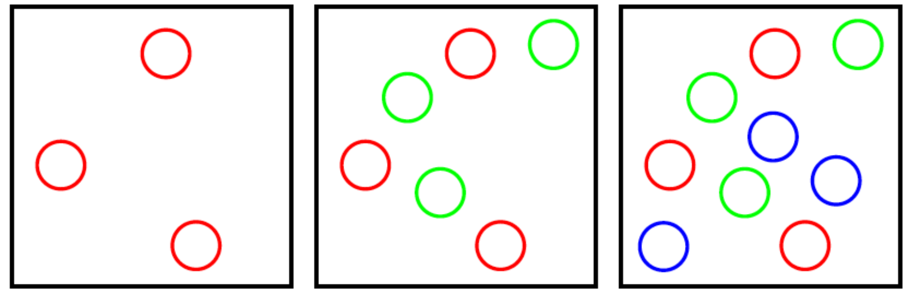

An abstract depiction of LOD particle data under increasing resolution with visual continuity. The particles in the lower resolution data are always present in the higher resolution data.

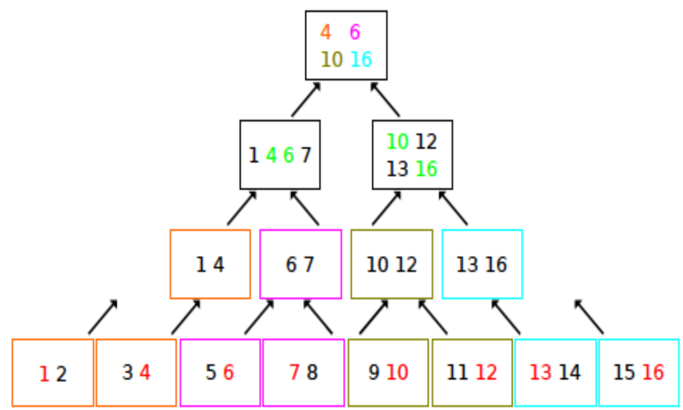

Empirically Comparing a 0.19% Sample compared to Full Resolution MC^3 Data

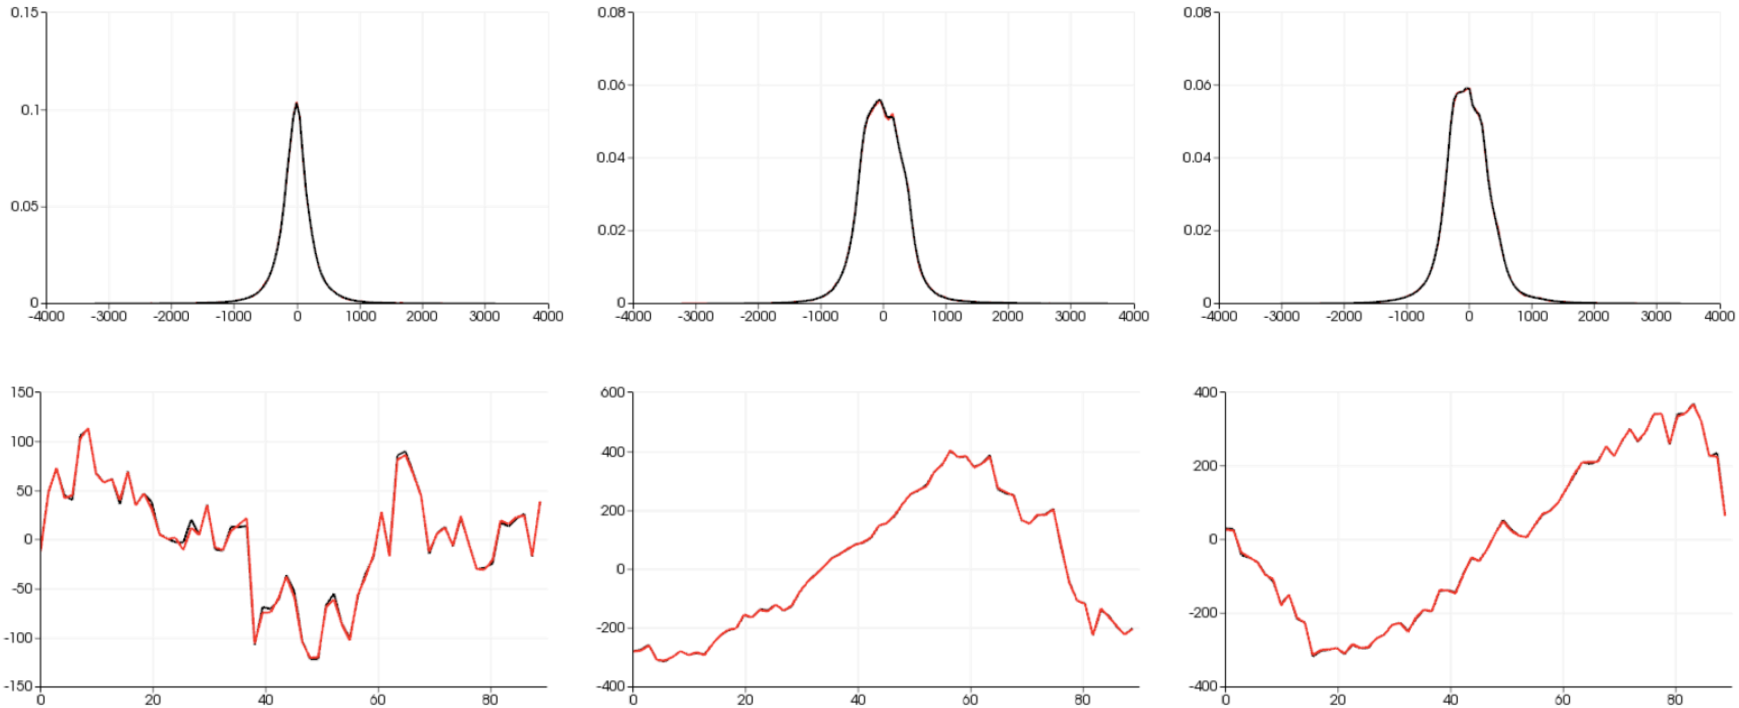

Red is 0.19% sample data, black is original simulation data. Both curves exist in all graphs, but the curve occlusion is reversed on top graphs compared to bottom graphs.

Effect of Sampling on Friend of Friends Algorithm

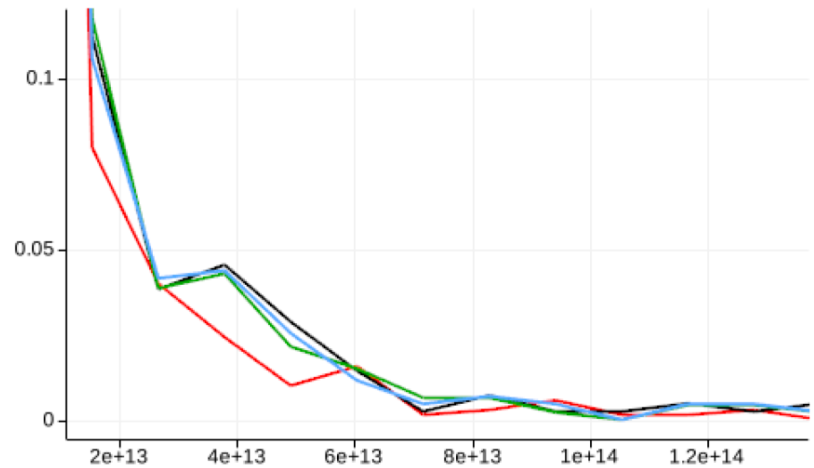

- The halo mass function for different sample sizes of 256^3 particles. The black curve is the original data. The red, green, and blue curves are 0.19%, 1.6%, and 12.5% samples, respectively.

Visual Downsampling

- Large-scale visualization tools (ParaView, VisIt, Ensight, etc.) have been effective, but render everything – a lack of display bandwidth compared to data sizes
- For the MC^3 data there is too much occlusion and clutter to see anything

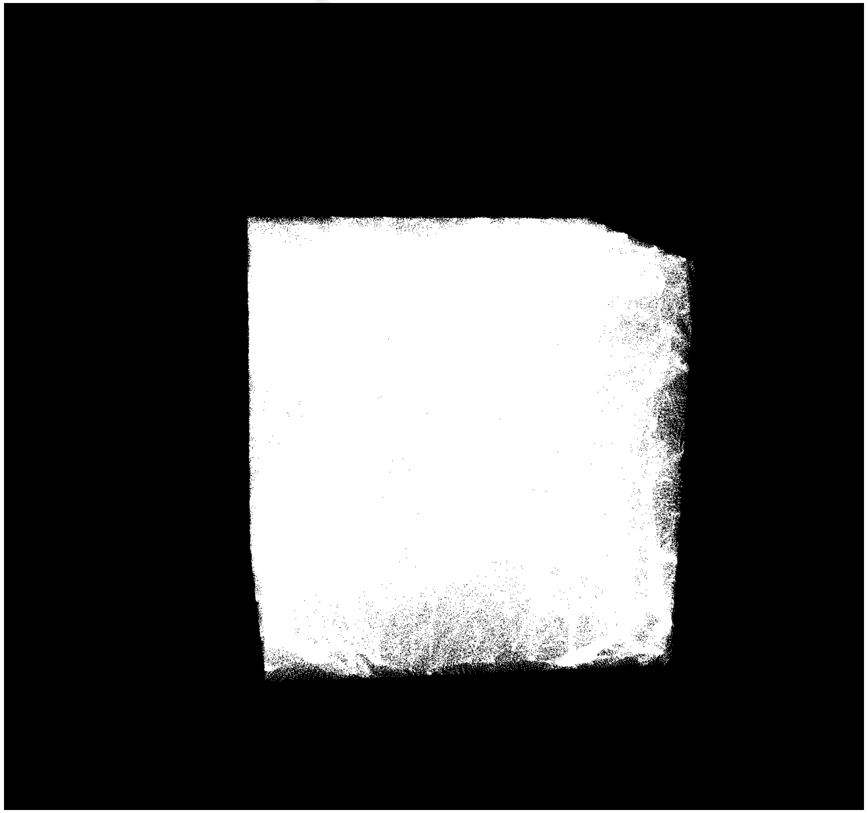

MC^3 visualization in ParaView

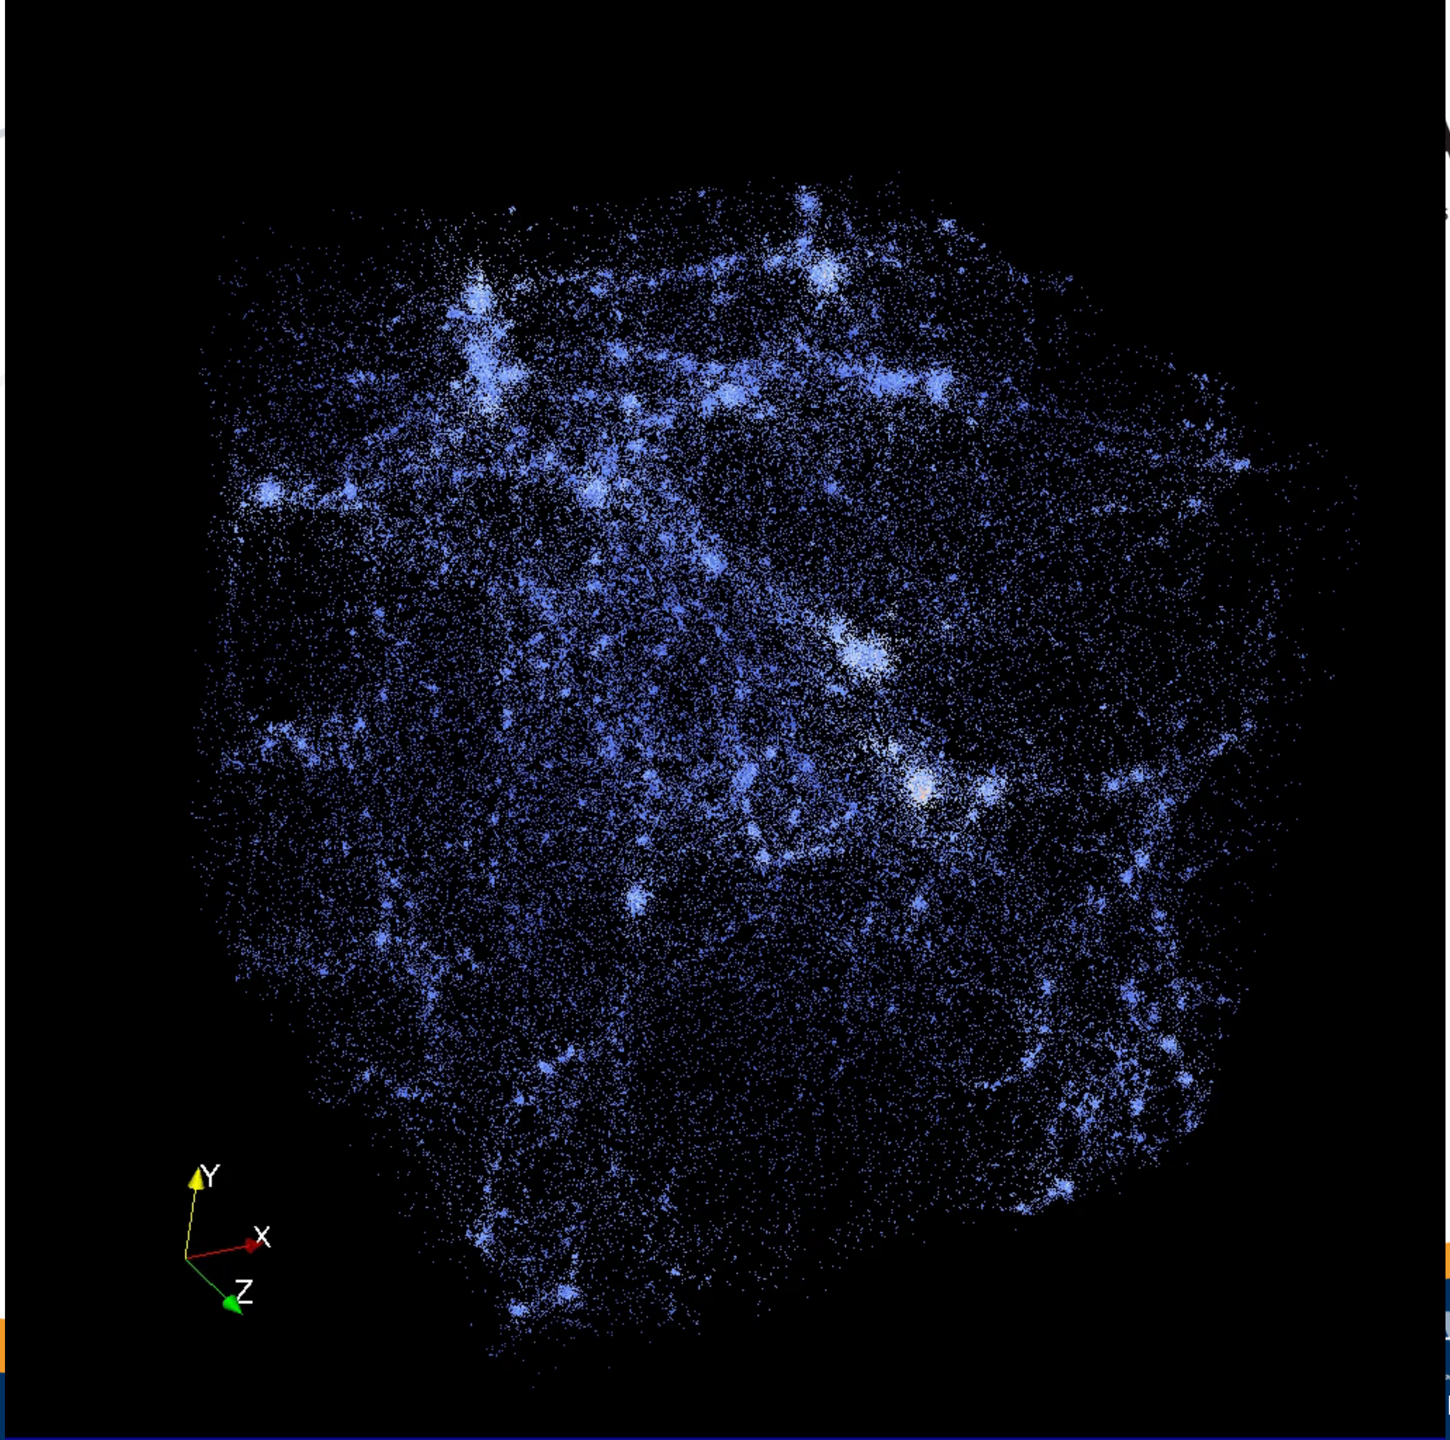
Y
X
Z

Solution: Compression

- (Simulation Data – Compressed Representation) provides an accuracy metric
- We measure the maximum point error so there is a guarantee that the data are accurate to *x* decimal places
- The user can trade read I/O time vs. data accuracy in a quantifiable manner

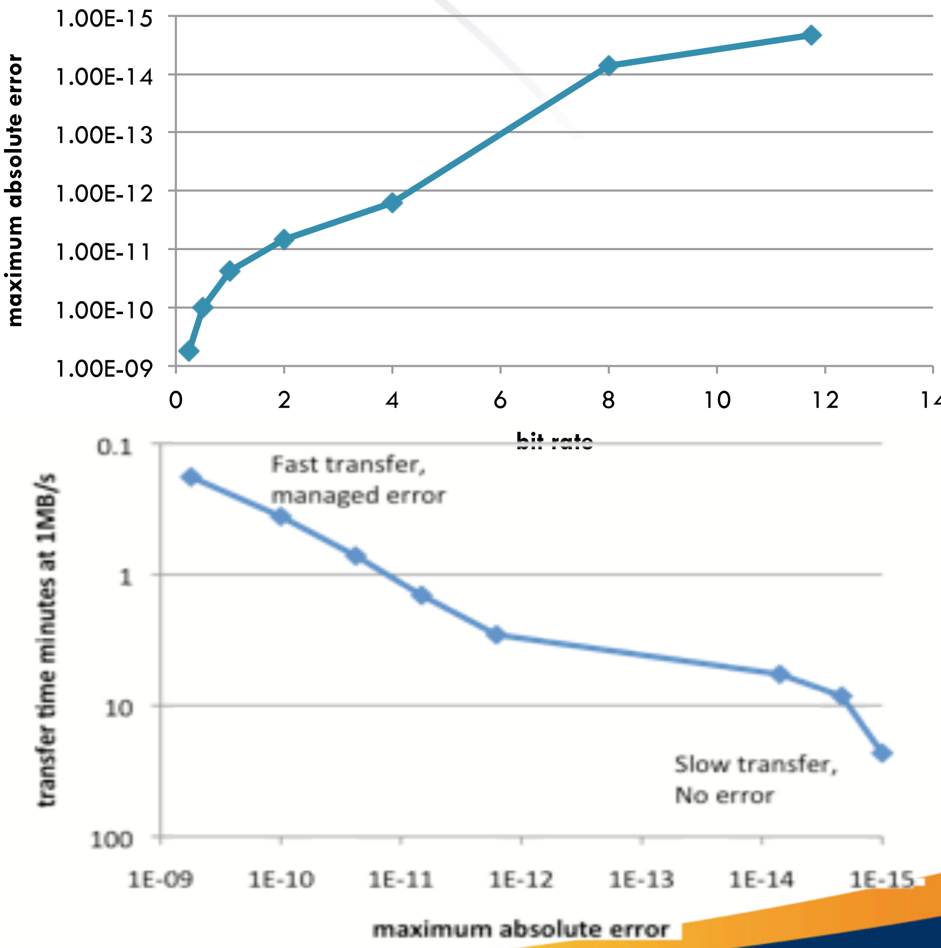

Isovalues on Compressed Simulation Data with Bounding Error - (32 bits, 3200x2400x42, 1.4 GB)

0.25 bits
10.8 MB

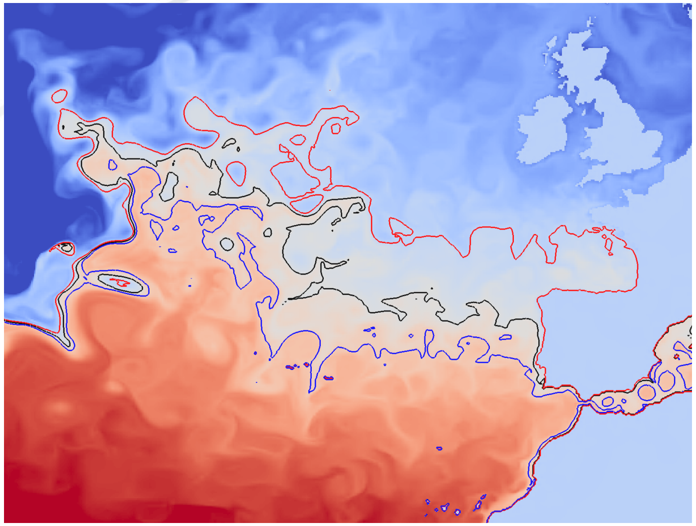

0.5 bits
21.6 MB

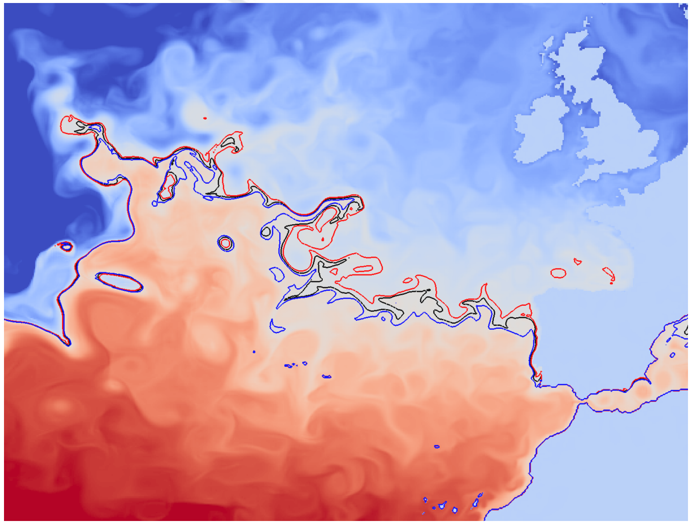

1.0 bits
43.3 MB

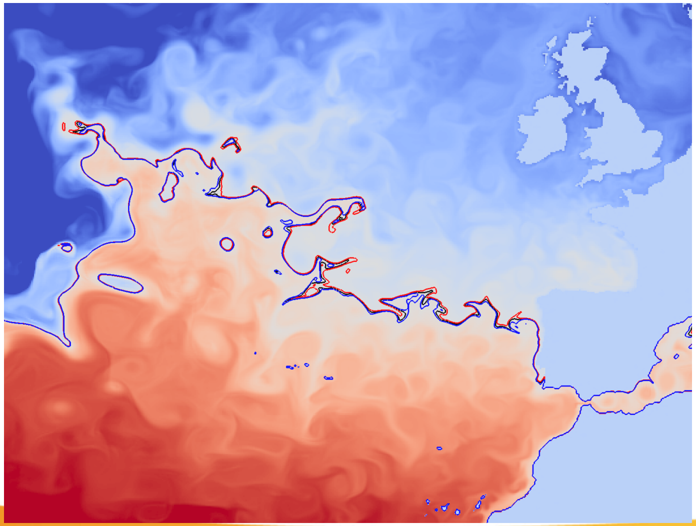

2.0 bits
86.5 MB

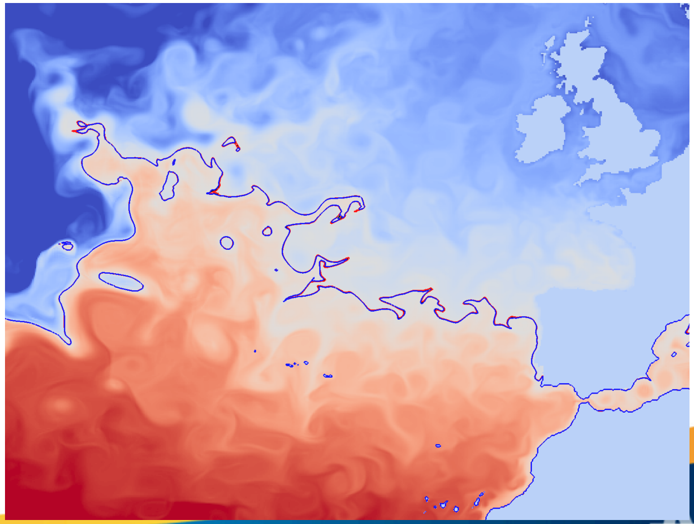

Data-Intensive Technology

- Server infrastructure
  - Clusters of “a few thousand processors,” with disk storage associated with each processor from pool of 10^6 processors
- Large database community driver
  - Success of this approach
    - Thousands of processors, terabytes of data, tenth of second response time
- Commodity parallel computing for data analysis
  - Available to scientific community
  - How do we leverage this?

The NIST Definition of Cloud Computing

- Cloud computing is a model for enabling:
  - ubiquitous, convenient, on-demand network access to
  - a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services)
  - rapidly provisioned and released with minimal management effort or service provider interaction

The NIST Definition of Cloud Computing - Essential Characteristics:

- **On-demand self-service**
  - automatic
- **Broad network access**
- **Resource pooling**
  - multi-tenant model
  - different physical and virtual resources dynamically assigned and reassigned according to demand

The NIST Definition of Cloud Computing - Essential Characteristics:

- **Rapid elasticity**
  - Scale rapidly commensurate with demand
  - Capabilities appear to be unlimited and appropriated in any quantity at any time
- **Measured service (*cost model*)**
  - Automatically control and optimize resource use by leveraging a metering capability
  - Resource usage is monitored, controlled, and reported, providing transparency

Service and Infrastructure Models (for clarity)

- **Infrastructure as a Service (IaaS)** - control over operating systems, storage, and deployed applications

- **Private cloud** - Provisioned for exclusive use by a single organization comprising multiple consumers

Axis	Sub-axis	Numerically Intensive	Data Intensive
Hardware	Nodes and Interconnect	High performance and power	Lower performance and power
	Storage	Separate, independent	Integrated
SW	Synchronization	Tightly coupled	Loosely coupled
	Reliability	Checkpoint restart	Replication
Workload	Number of Users	**Single per node**	**Multiple per node**
	Data	Dynamic, heterogeneous (unstructured grid)	Static, homogeneous (text, images)
	Algorithms	Global	Distributed
	User Interface	**Complex Application**	**Simple Web**
	Data Model	**Files**	**Database**
Workflow	Scheduling	Batch	Interactive
	Analysis	Offline post-processing	Online
	I/O	Bulk parallel writes	Streaming writes

Supercomputing System Architecture Diagram

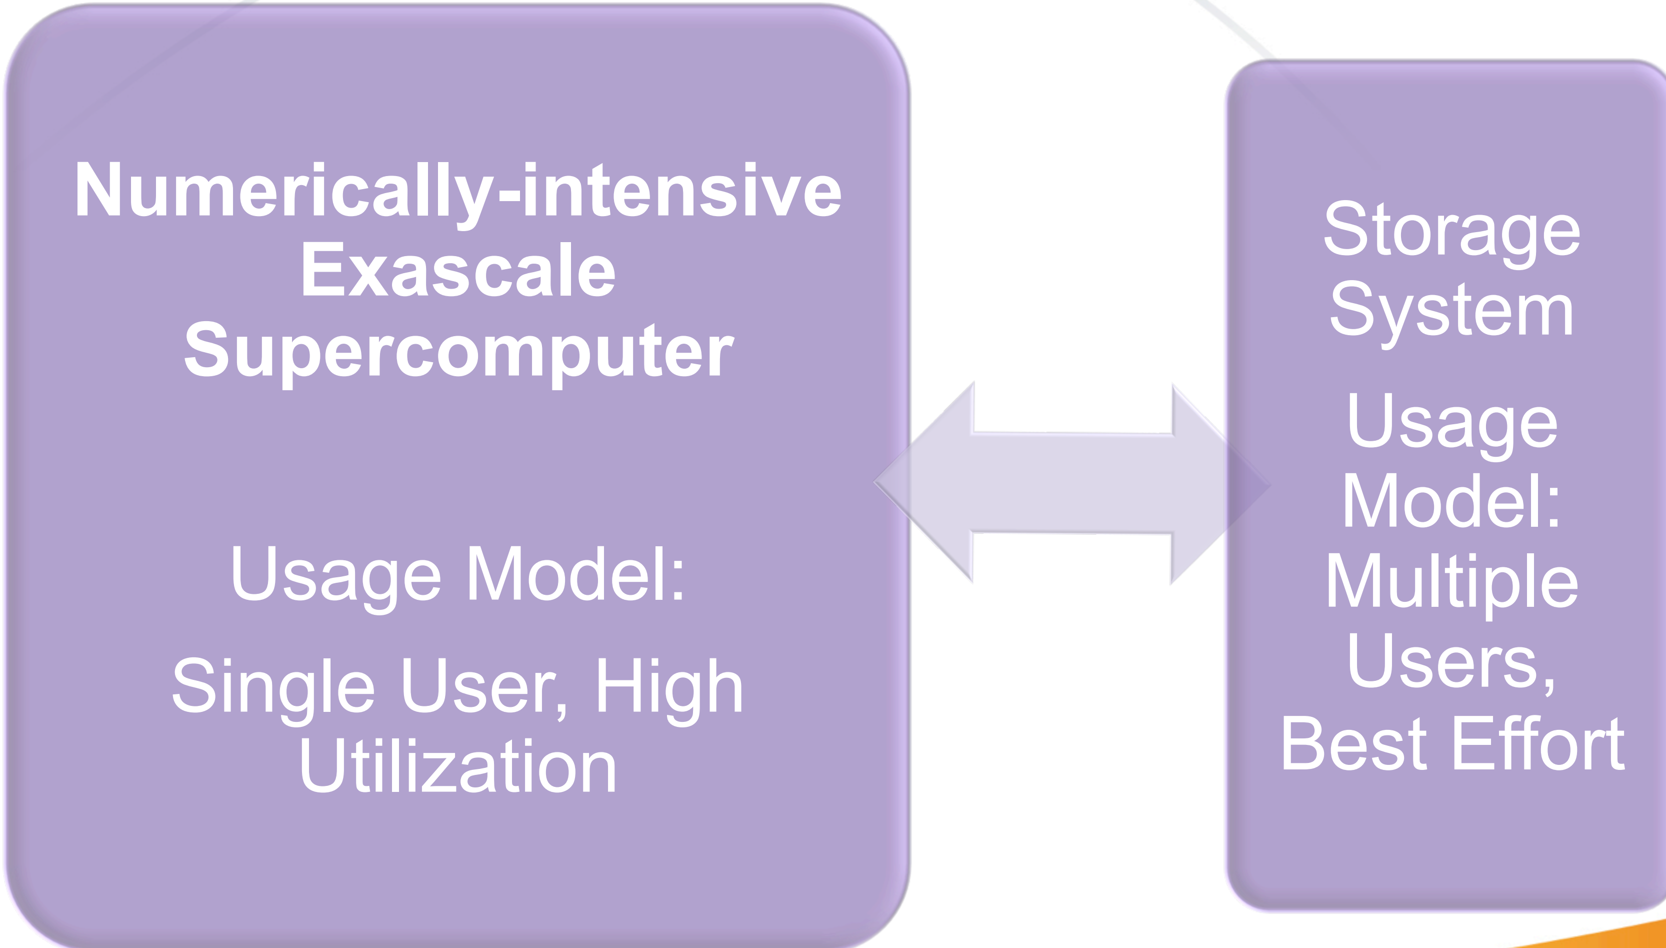

Supercomputing System Architecture Diagram

Los Alamos NATIONAL LABORATORY EST. 1943

Numerically-intensive Exascale Supercomputer

Usage Model:

Single User, High Utilization

Intelligent Storage System Impl. w/ Cloud Technology

Usage Model: Multiple Users, Best Effort

Solution: First steps – ParaView with HDFS

- Use Hadoop Distributed File System (HDFS) instead of Lustre
  - with ParaView visualization application
  - *3x improvement and reduced variance* in read times
- Compose relevant parts of each ecosystem
  - Did not use map reduce scheduler
- TACC/LANL success!
  - Drawback: Still used HPC model… copy experimental data onto Longhorn for each run

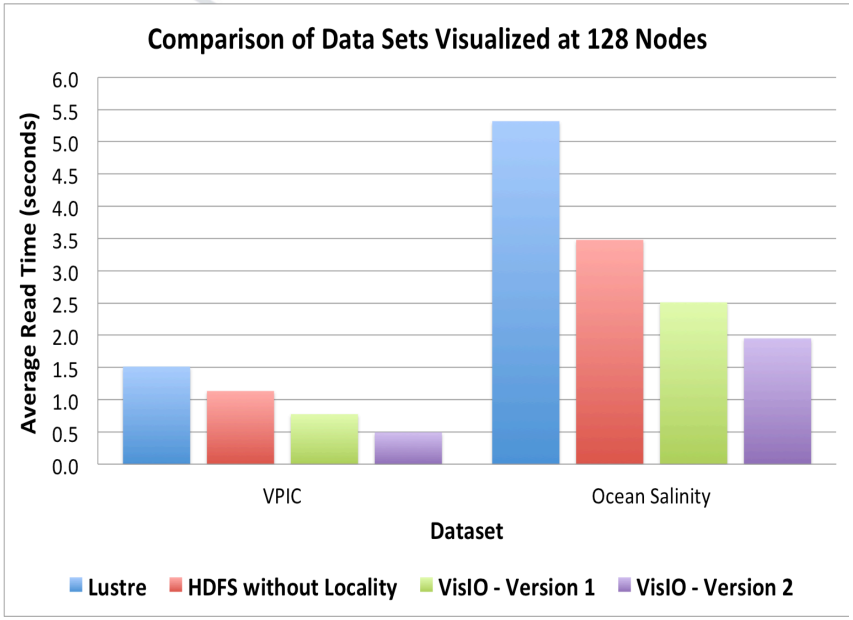

C. Mitchell, J. Ahrens, and J. Wang. "VisIO: Enabling Interactive Visualization of Ultra-Scale, Time Series Data via High-Bandwidth Distributed I/O Systems". IEEE International Parallel and Distributed Processing Symposium, May 2011.

Solution: Image Database
Embrace cloud approach

1. <u>During in situ analysis save out massive image **database**</u>
   – Images contain values and depth
   – Cartesian product of camera positions, operators, variables, timesteps
   – Guided by budget and analysis questions
2. <u>During post-processing</u>

- Visualization process = exploration of image **database**
  – Interactive **web interface**
    - Camera
    - Visualization operators
      - Through image compositing and rendering techniques
  – Search **web interface**
    - Explore and search from simple web based interface
- Supercomputing 2014

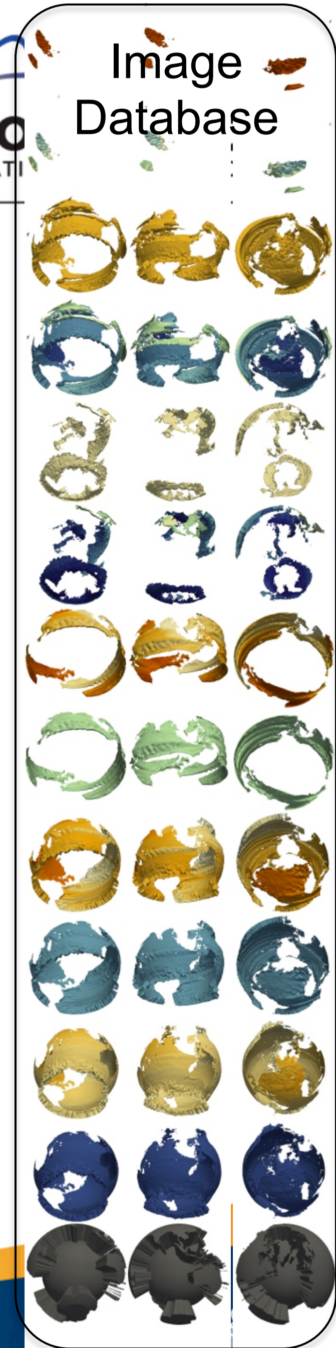

Mega	Giga	Tera	Peta	Exa
10^6	10^9	10^12	10^15	10^18
Image size		Data size	Data size	Data size

Choices…

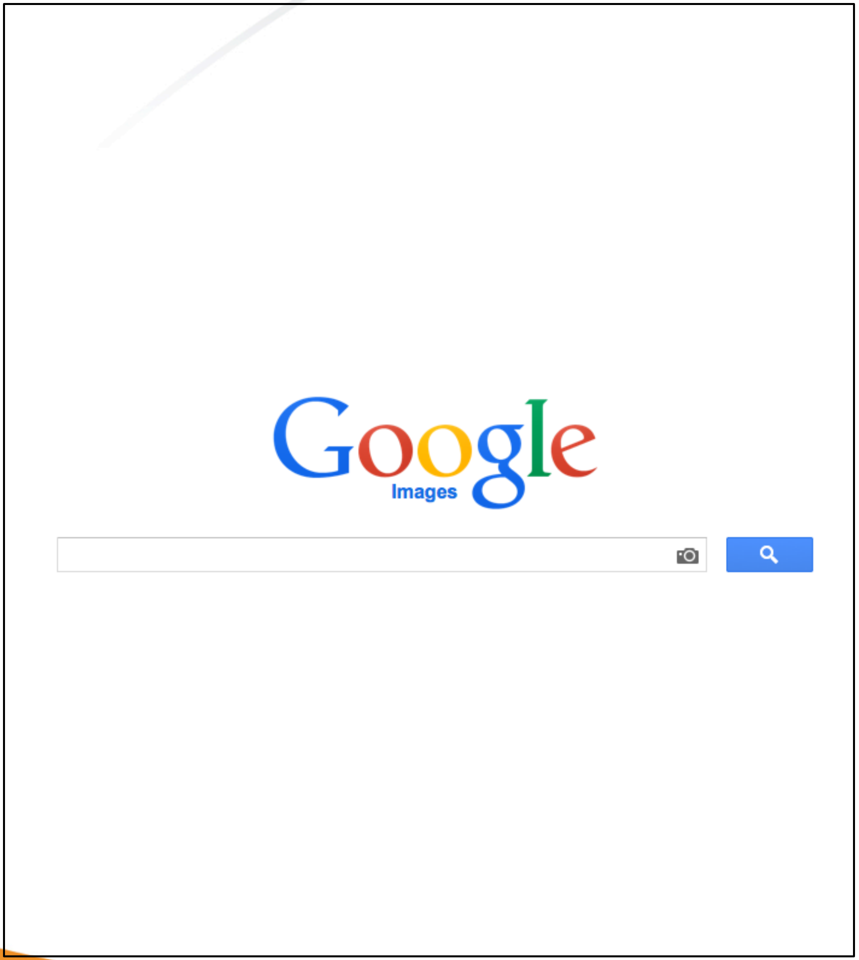
Google
Images

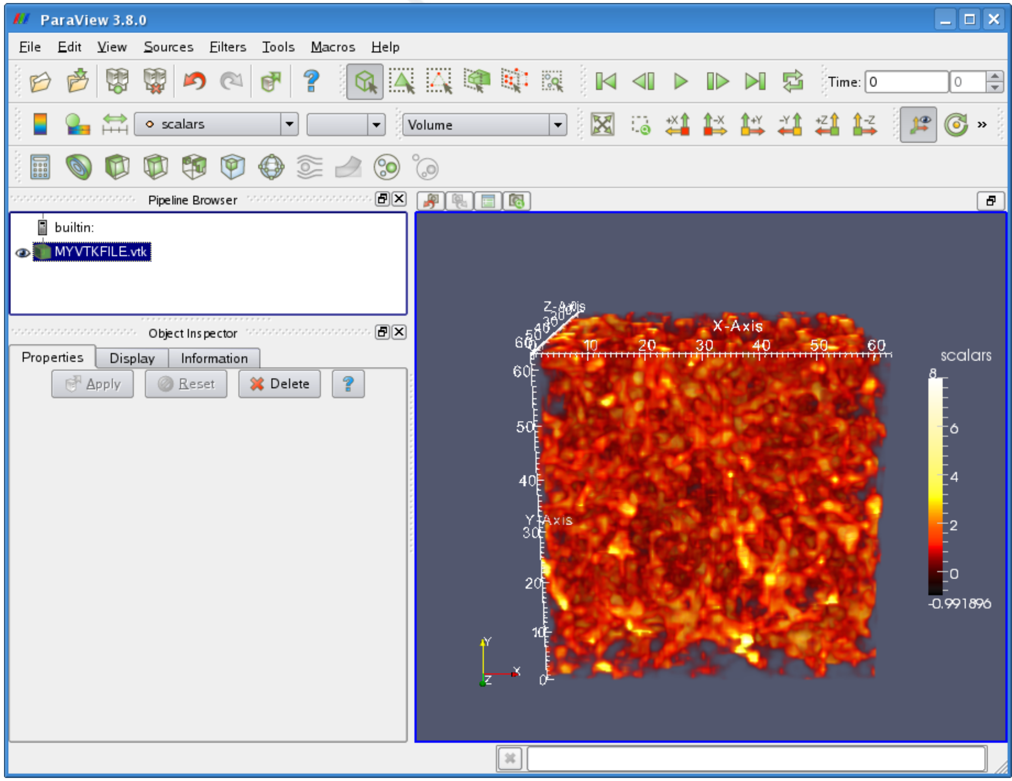
ParaView 3.8.0
File Edit View Sources Filters Tools Macros Help
Time: 0
scalars
Volume
Pipeline Browser
builtin:
MYVTKFILE.vtk
Object Inspector
Properties Display Information
Apply Reset Delete
X-Axis
Y-Axis
Z-Axis
scalars
-0.991896

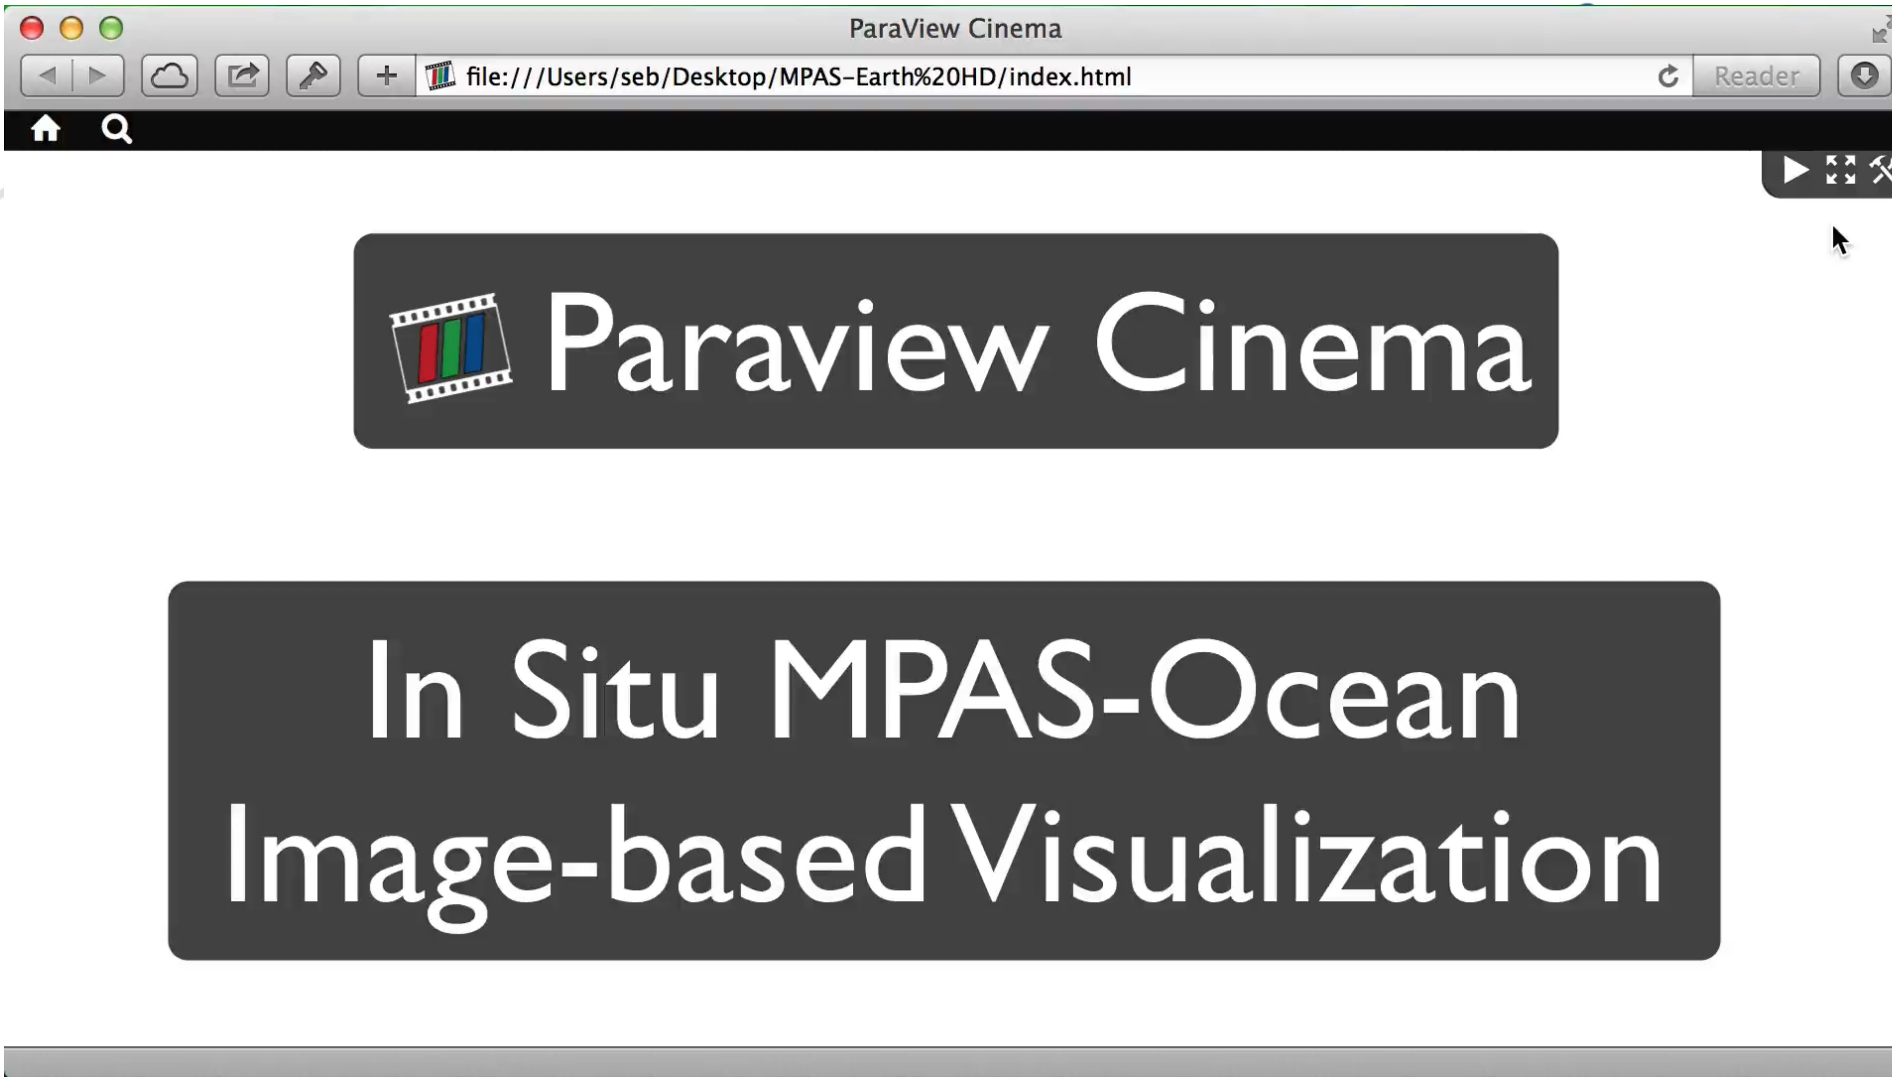
ParaView Cinema
file:///Users/seb/Desktop/MPAS-Earth%20HD/index.html
Reader
Paraview Cinema
In Situ MPAS-Ocean
Image-based Visualization

Ongoing experimentation and evaluation of data intensive technology for visualization

- Leveraging the rapid evolution of **commercial big data analytics**
- Advanced Database Technology
  - Relational database technology (**Microsoft SQL Server**/Johns Hopkins University) applied to turbulence simulation data
  - Query engine developed for **Spark** applied for particle-based cosmological simulation analysis
- Fast/Local Access Storage Technology
  - **Hadoop** distributed filesystem inspired by **Google** applied to scientific visualization and analysis (ParaView)
  - Christopher Mitchell, James Ahrens, and Jun Wang. "VisIO: Enabling Interactive Visualization of Ultra-Scale, Time Series Data via High-Bandwidth Distributed I/O Systems". IEEE International Parallel and Distributed Processing Symposium, May 2011.
- Cloud Technology
  - Virtualization approaches (**VMWare**/OpenStack) applied to managing data-intensive scientific supercomputing

Summary - Major technology changes forcing changes in standard approaches

- Numerically intensive/ Exascale: reduce power/data movement
  - Automated in situ data reduction will be the norm
  - Help users to understand the analysis choices they will have to make
    -
- Data Intensive/Cloud:
  - Creative energy and massive industry resources in this area
    - This is now commodity and standard
      - Shared memory to clusters
  - Adapt, leverage, succeed
    - Example - Image database

Conclusions

- How do these changes affect your research, development, deployment?
  - In situ?
    - What can be automated?
  - Data intensive?
    - Embrace usage model!
    - What flexible data products can be saved?
    - What are the new approaches enabled by this technology?
      - Databases

END