The Scalable Data Management, Analysis, and Visualization (SDAV) Institute

Key Technical Accomplishments and Major Work in Progress

(check out 13 posters) poster

http://sdav-scidac.org/

SDAV Institute

Arie Shoshani (PI)

Co-Principal Investigators from:

Laboratories	Universities
Argonne NL	Georgia Inst of Technology
Lawrence Berkeley NL	North Carolina State U
Lawrence Livermore NL	Northwestern U
Oak Ridge NL	Ohio State U
Los Alamos NL	U of California, Davis
Sandia NL	Rutgers U
Kitware (Industry)	U of Utah

22 people attending

SDAV Portfolio

(http://sdav-scidac.org/toolkit.html)

Data Management tools

I/O Frameworks

ADIOS

Darshan

Parallel netCDF

ROMIO

ViSUS/IDX

In Situ Processing

GLEAN

DIY

DataSpaces

EvPath

Indexing / Compression

FastBit

ISABELA

SDAV Portfolio

(http://sdav-scidac.org/toolkit.html)

Analysis and Visualization tools

Analysis and Visualization Frameworks

VisIt

ParaView

Analysis and Visualization Tools and Libraries

ExMage

Ultravis-P

IceT

VTK

TALASS

MSCEER

NDDAV

Multi-/Many-core Visualization Libraries

Dax

EAVL

PISTON

Statistics and Data Mining

NU-Minebench

STPMiner

Distance Field Computing

DATA MANAGEMENT

Area Co-Leads:
Rob Ross and Scott Klasky

SDAV Data Management Technologies

SDAV is a major provider of data management technologies for leadership computing platforms.

- **Parallel I/O and file formats** – efficiently moving data to/from storage and organizing that data
- **Data management frameworks** – retrieving data and handing data off to other services
- **Indexing and data compression** – tools for managing leadership-scale scientific data
- ***In situ* processing and code coupling** – integrating new services into I/O pathway

Parallel I/O, File Formats, and the Traditional I/O Model

Parallel I/O in the Traditional Model

SDAV technologies play central roles in the traditional post-processing I/O stack (and for checkpoint/restart).

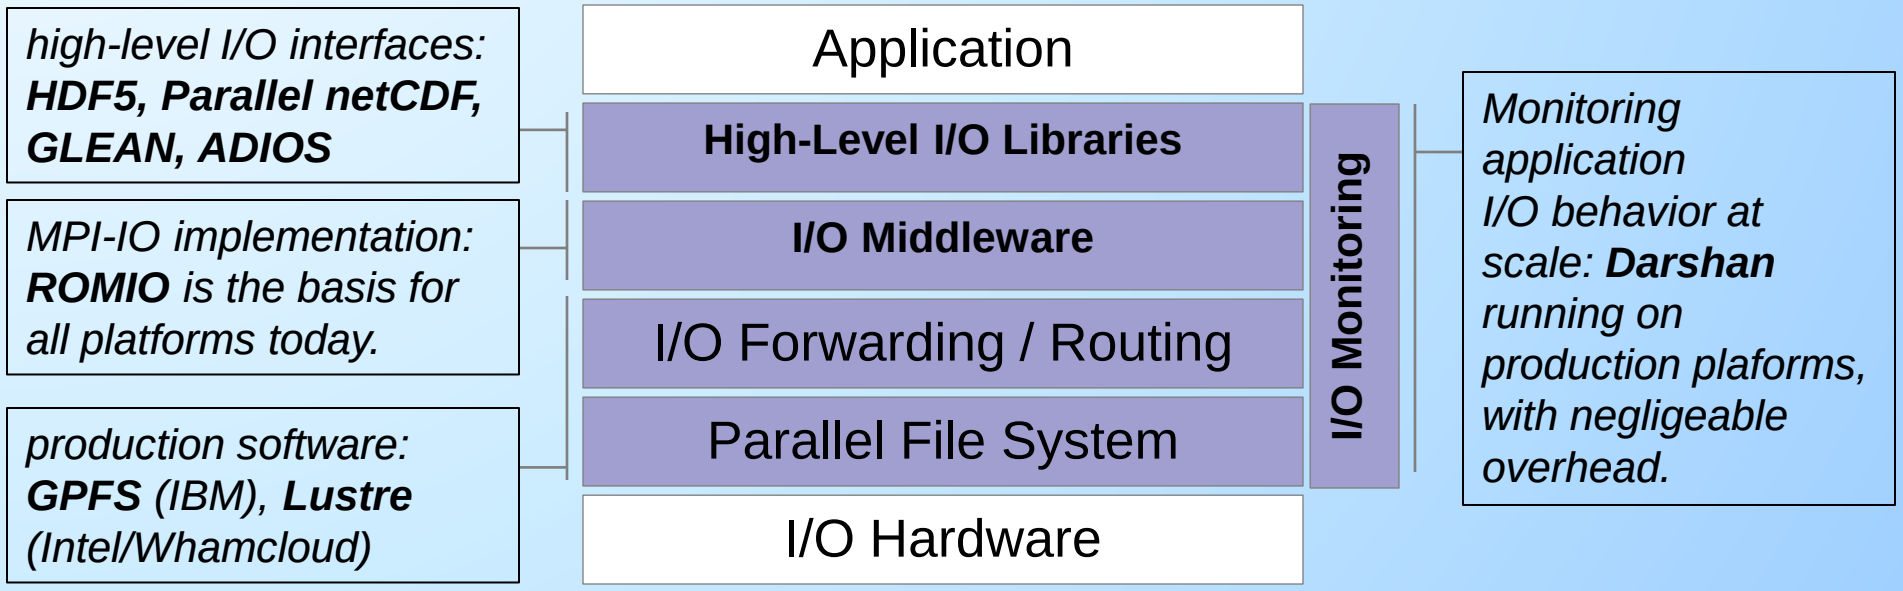

Assisting Scientists with Parallel I/O

- I/O performance is highly dependent on configuration settings
  - "Best" settings vary across workloads, scales, platforms
- e.g. HACC code often uses traditional I/O stack for checkpointing

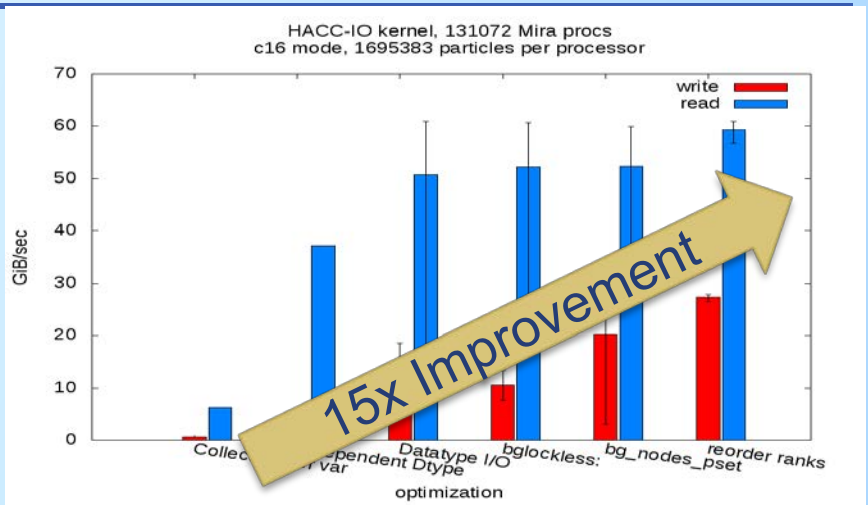

- I/O monitoring (**Darshan**) provided insight into performance problems
- Using **GLEAN** Small code tweaks plus configuration adjustments led to **15x** performance improvement for checkpoint operations.
- Achieved **160 GB/s** for HACC production simulations on Mira BG/Q system
- Currently working to automate selection of good configuration parameters (with SUPER)

poster

I/O Performance Improvement for Multiple Applications

ADIOS

Optimizations use:

- Synchronous I/O on leading HPC systems (Titan, Mira, Edison, Tiahne-1A, Garnet, …)
- Techniques are to optimize in-memory data copy, movement, and file system access
- Optimizations for GPFS, Lustre

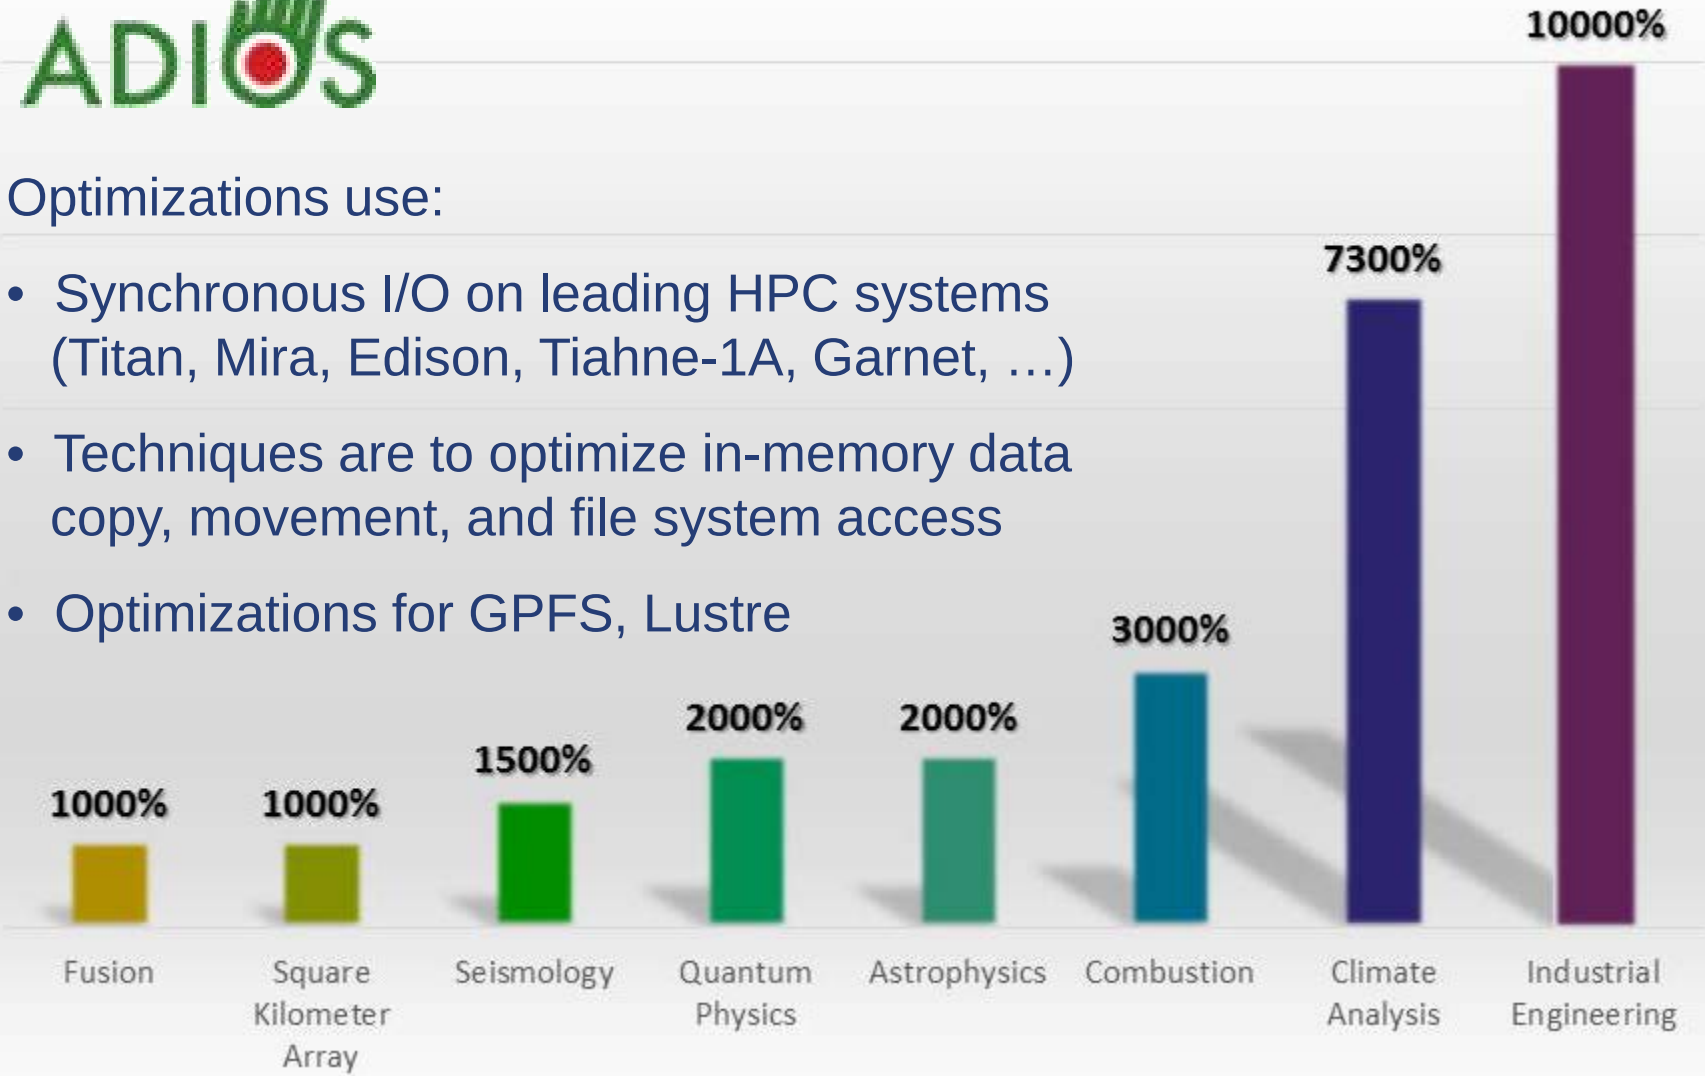

poster

Beyond the Traditional Model: Data Management Frameworks

Beyond Traditional I/O Capabilities

Data management frameworks provide a vehicle for deployment of *in situ* data analysis/transformations

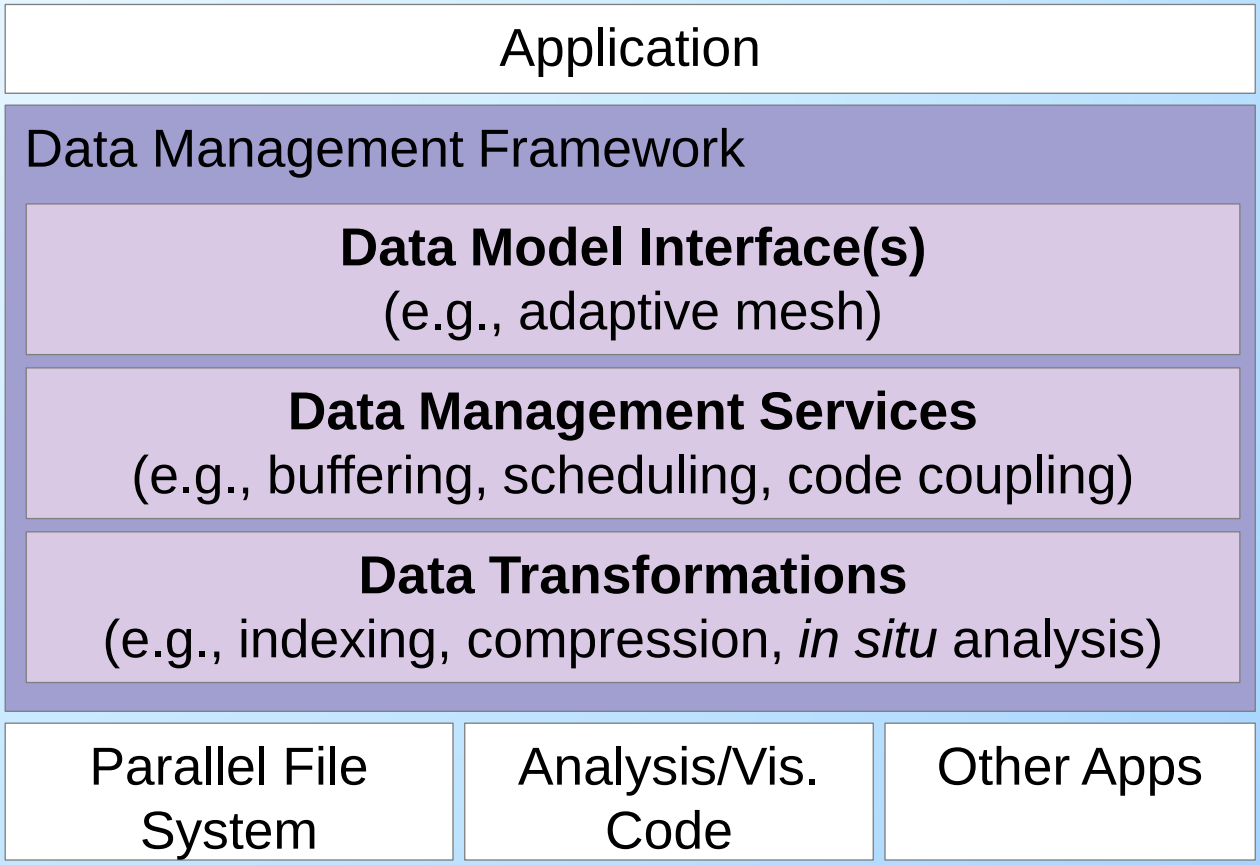

ADIOS

ADIOS is a framework for I/O and data management that facilitates high performance data movement at leadership scales.

- Mechanisms for changing I/O methods on the fly
- Plug-in architecture
- Specific component implementations optimized for key platforms and use cases
- Vehicle for deployment of successful I/O and data management technologies

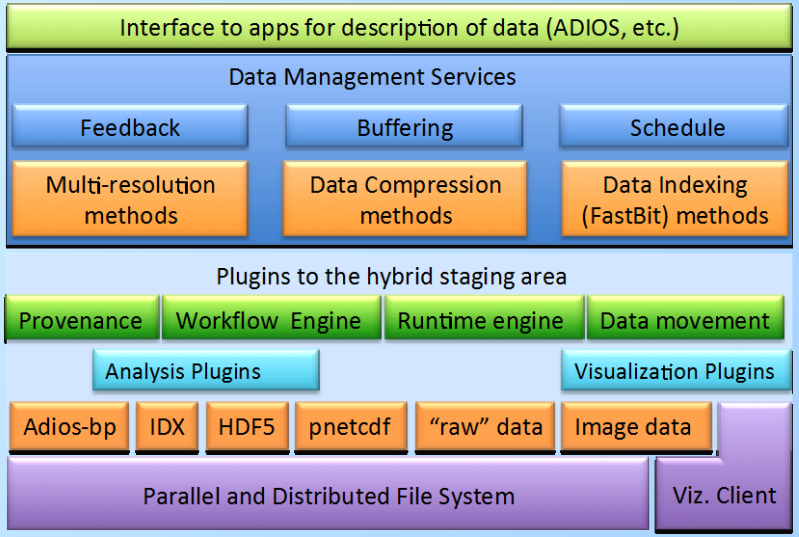

R&D 100

poster

Coupling Codes via Frameworks

Data management frameworks allow us to easily compose multi-scale / multi-physics applications along with analytic services.

- Significant challenge for science teams
  - Build huge application with all components?
  - Share data through files?
- **DataSpaces** provides a distributed, in-memory object store on HPC systems
- Integration into ADIOS eases use by science teams.

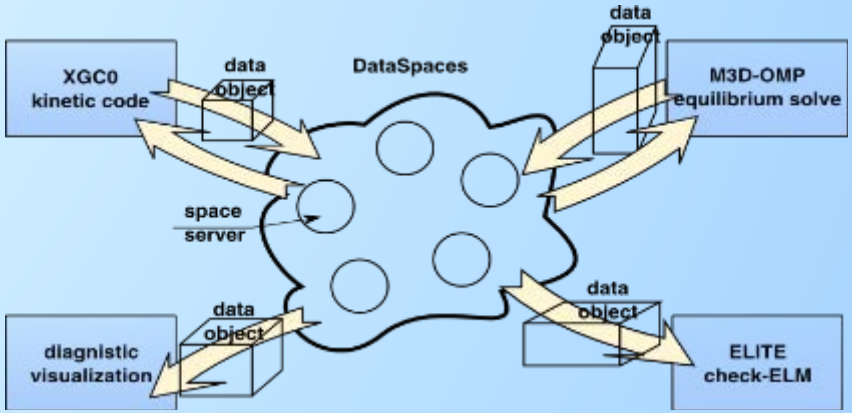

As part of our collaboration with the EPSI project, ADIOS/DataSpaces is being used to enable coupling of XGC0, M3D-OMP, and ELITE codes, as well as providing a means for runtime diagnostic visualization.

poster

Enabling Data Analysis: Indexing, and New Data Organization

Indexing Scientific Datasets

FastBit is a compressed bitmap indexing method specifically designed for scientific numerical data.

- Provides 10-100 fold query speedup as compared with existing methods
- Especially effective for in situ and real-time data analysis
- Has been applied to TB-size datasets in various domains (e.g., particle accelerators, combustion, fusion, biology)

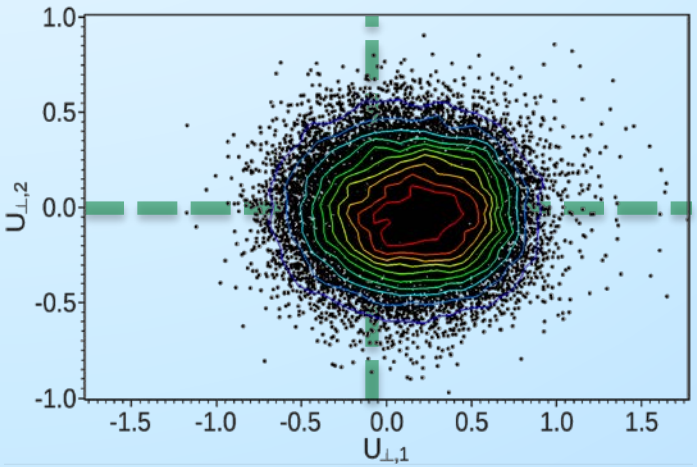

- Indexed and queried a trillion particle dataset for studying **magnetic reconnection**
- Located particles of interest (highly energetic) in 3 seconds using 1,250 cores
- The off-centered and oblong distribution of highly energetic particles confirms existence of asymmetry in electron behavior (agyrotropy)
- "This is the first time anyone has ever queried and visualized 3D particle datasets of this size." -- Homa Karimabadi, Physicist, UCSD

poster

Applying New Data Organizations

IDX is a library for storing multidimensional data in a multi-resolution format

- Enables fast browsing of very large datasets
- PIDX version enables writing of data directly into this format from simulation codes, streaming from simulation to analysis codes

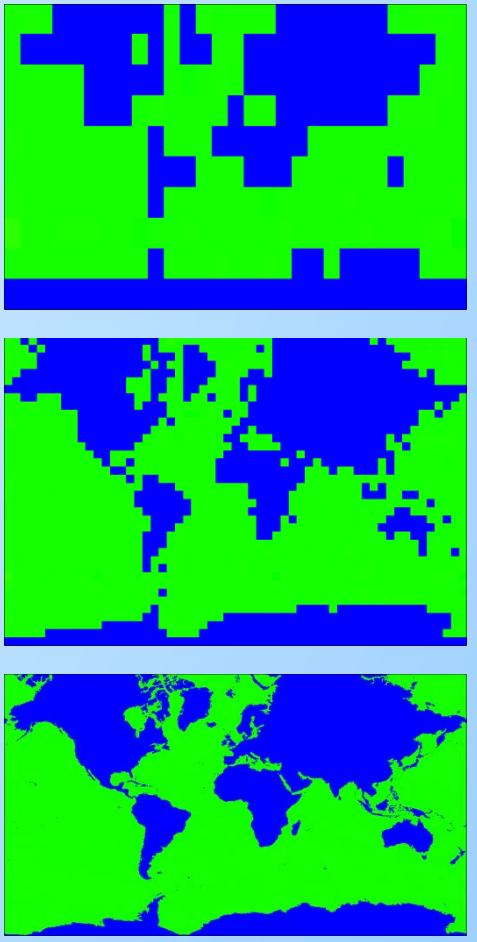

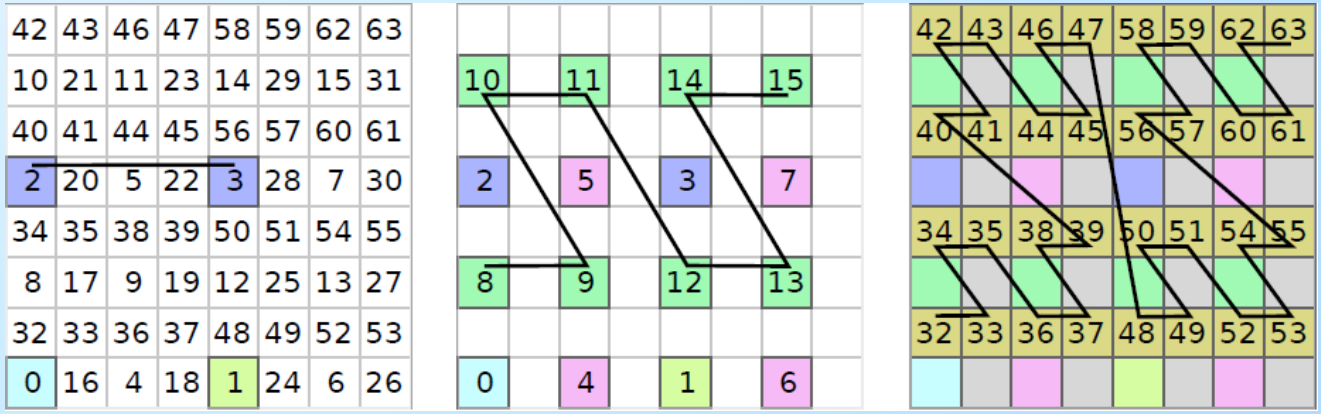

Hierarchical z-order data organization enables multi-resolution access while retaining spatial locality.

Understanding and Accelerating Science I/O

I/O Characterization with Darshan

Darshan is an open source, application-level instrumentation library that is used to capture patterns of storage access by DOE codes across multiple platforms.

- Approach
  - Limited, fixed memory footprint
  - No communication prior to MPI_Finalize (end of job)
  - Heavy use of compression
- Benefits
  - Aids in application performance tuning
  - Guides storage R&D and future procurements
- Darshan is deployed at NERSC and ALCF facilities

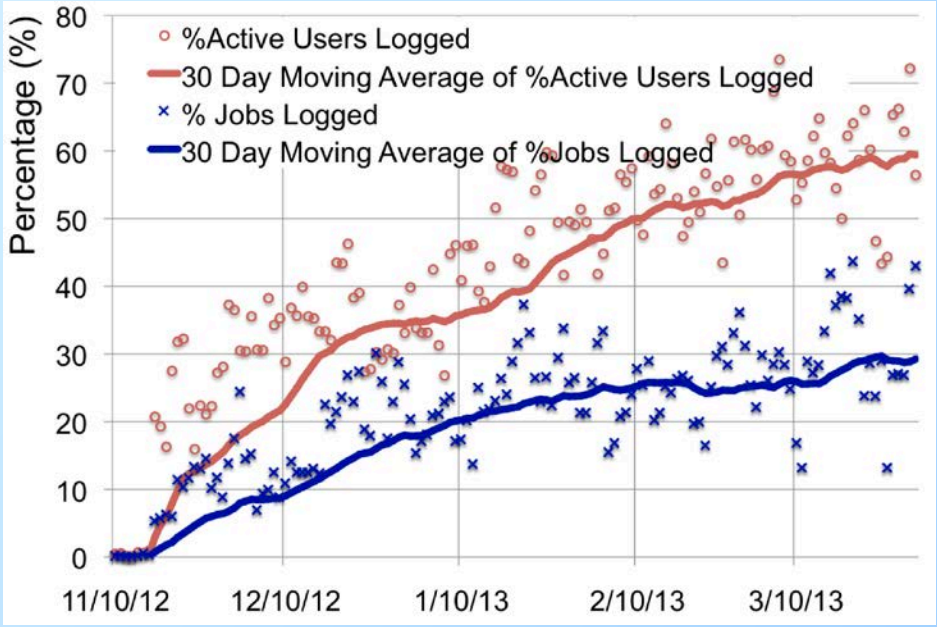

Coverage on NERSC Hopper platform as of March 2013.

poster

Data-Centric Science on DOE Leadership Computing Platforms

Matching large scale simulations of dense suspensions with empirical measurements *to better understand properties of complex materials such as concrete.*

Processing large-scale seismographic datasets *to develop a 3D velocity model used in developing earthquake hazard maps.*

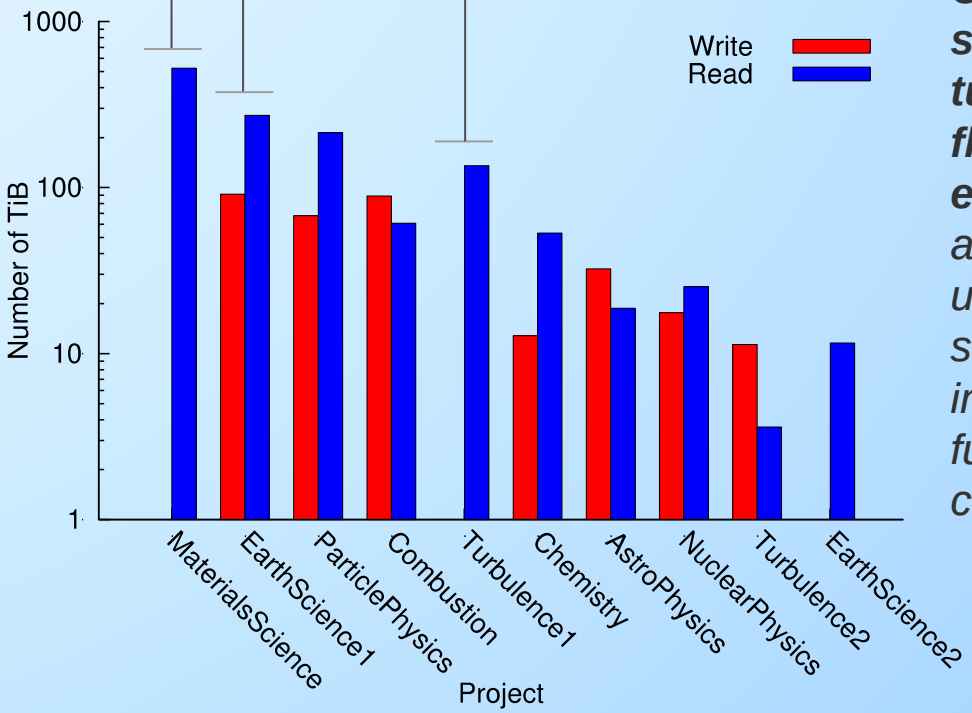

Comparing simulations of turbulent mixing of fluids with experimental data *to advance our understanding of supernovae explosions, inertial confinement fusion, and supersonic combustion.*

VISUALIZATION

Area co-leads

James Ahrens and E. Wes Bethel

SDAV Visualization Technologies

- **VisIt and ParaView**: ongoing deployment, support, evolve to new platforms, delivery vehicles for new technologies.
- **VTK-m**: our approach for realizing m-core across many technologies.
- **Flow visualization, Rendering, and Ensembles**: productize technology for enabling knowledge discovery.

VisIt and ParaView

VisIt and ParaView Progress

Visit and ParaView (Kitware, LANL, LBNL, LLNL, ORNL, SNL)

- Enhance Visit and Paraview to leverage multiple cores within a single MPI task
- Integrate VisIt and ParaView with ADIOS
- Demonstrate and evaluate *in situ* analysis methods with VisIt and ParaView

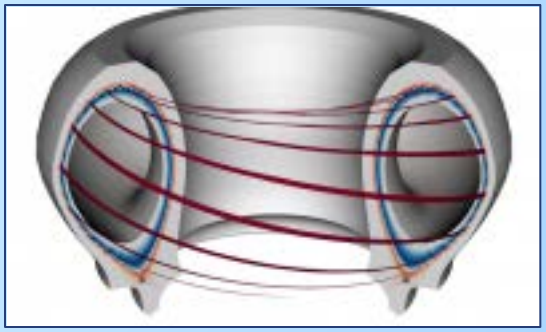

SDAV software installed and supported at major DOE SC facilities

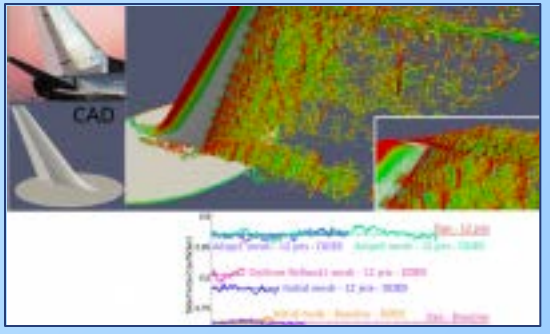

VisIt, ParaView and Multi-core

- Objective:
  - Want to be able to take advantage of multi-core platforms, don't want to run MPI task-per core.
- Approach:
  - Focus on key infrastructure in VTK: thread safety.
  - General-purpose threading interface to abstract back-end threads library (pthreads, Threading Building Blocks (TBB), etc.)
- New features:
  - Summer 2014: VTK 6.1, VisIt 2.8, ParaView (4.1.0)
  - Most embarrassingly parallel operators supported
  - Runs faster, uses less memory.
  - VTK 6.1: vtkSMP (symmetric multiprocessing) class

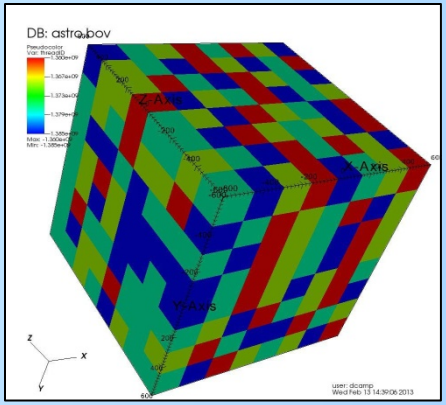

512-block astrophysics dataset colored by thread ID.

VisIt, ParaView, and ADIOS Integration

- Problems: increasingly intractable to do full-resolution I/O, science being lost due to analysis of partial results.
- Approach: integration, interoperation of SDAV technologies for *in situ* processing and vis/analysis; work with code teams to put this technology into practice.

New features:

- VisIt, PV ADIOS loaders.
- Code team interactions and deployment: SPECFEM3D (seismology, INCITE); XGC (fusion, EPSI-SciDAC), others.

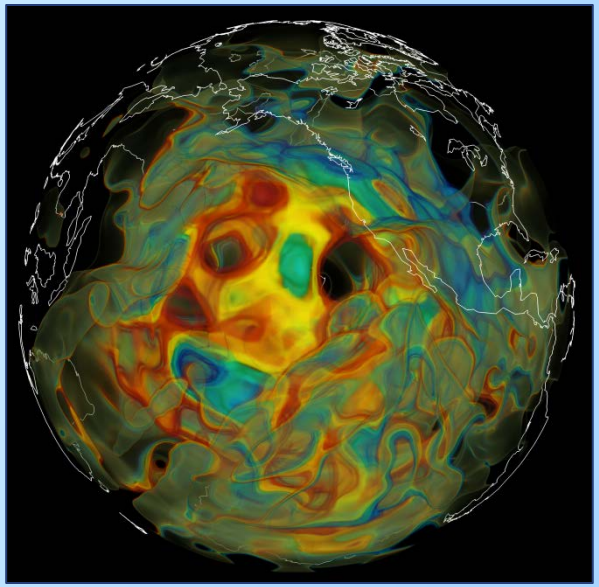

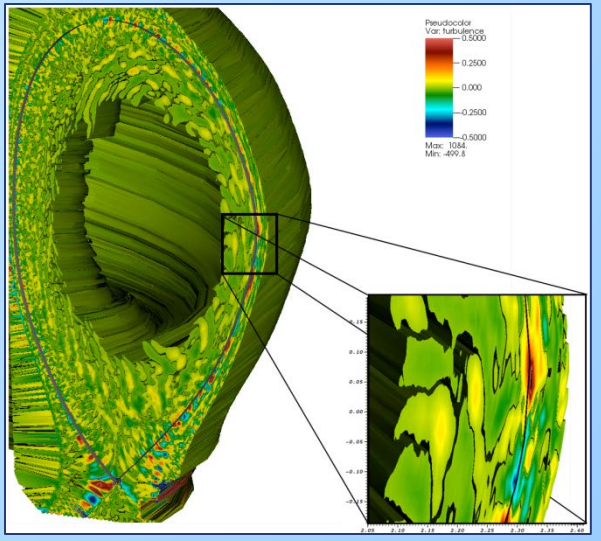

In Situ VisIt, ParaView Enables New Science

Problem:

science being lost due to infrequent temporal I/O, inability to do parallel *in situ*

New features:

- Release of **ParaView/Catalyst** library: open source library for coupling codes to VTK-, PV-based vis/analysis.
- **VisIt Libsim** engineering to optimize memory footprint, etc.
- Custom **PV+PISTON** adaptor to support interactions between science teams and *in situ* tools.

poster

poster

Applications:

- Worked with multiple applications/teams: VPIC and H3D (plasma physics), HACC (astro), Warp3D (accelerator), POP (climate), XGC (fusion), SPECFEM3D (seismology), many others.

Example:

- Warp3D+VisIt/Libsim working at 100s of cores on 2D configuation, plans for 10Ks of cores on 3D.
- *in situ* Data subsetting, analysis and visualization

poster

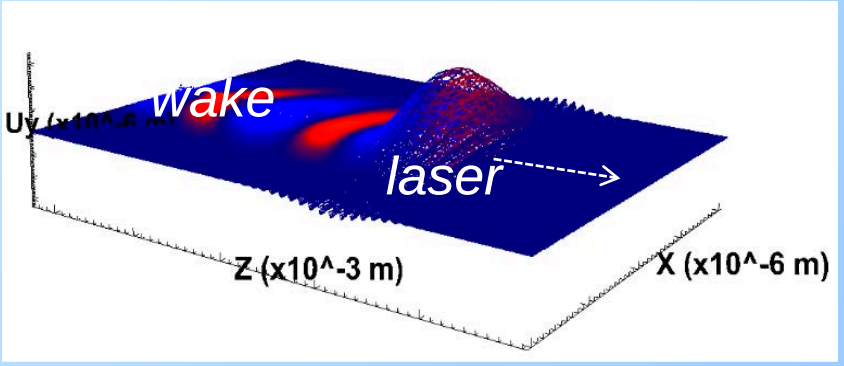

VTK-m: realizing m-core across technologies

VTK-m Background

- **VTK** is a serial, single-threaded class library (data structures and algorithms) used as the basis for important applications (VisIt, PV)
- **EAVL** emphasizes the development of a new data model,
- **Dax** emphasizes the development of a new execution model,
- **DIY** provides a lightweight toolkit of commonly-used distributed-memory parallel functionality,
- **PISTON** emphasizes portability and parallel algorithm development.

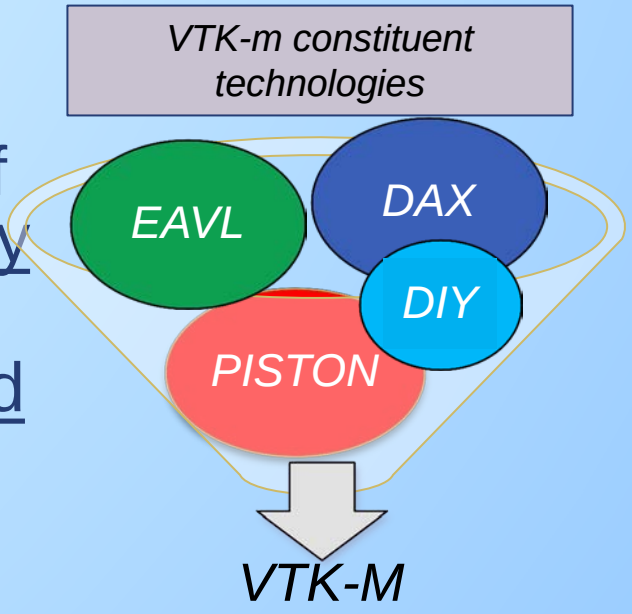

poster

Enhancements for LCF Codes, Prototype Implementations

EAVL and XGC, loosely coupled *in situ* with ADIOS

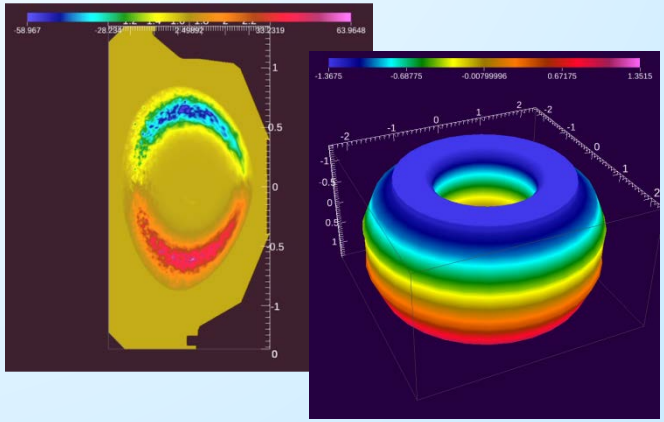

PISTON-PV enables in situ computation of halos on 8192^3 dataset in 16K Titan nodes (not previously possible)

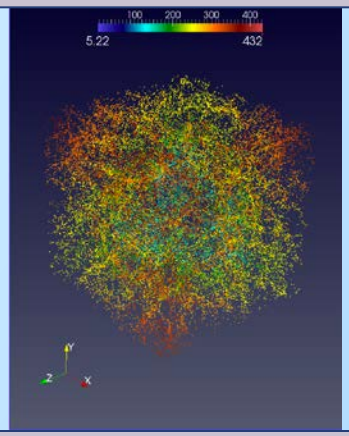

poster

DIY enables parallelization of *in situ* computational geometry capability

poster

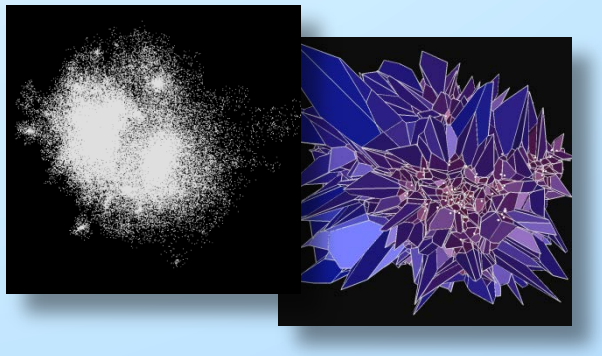

Dax-PV used for contour with subsequent vertex welding, coarsening, subdivision, and curvature estimation

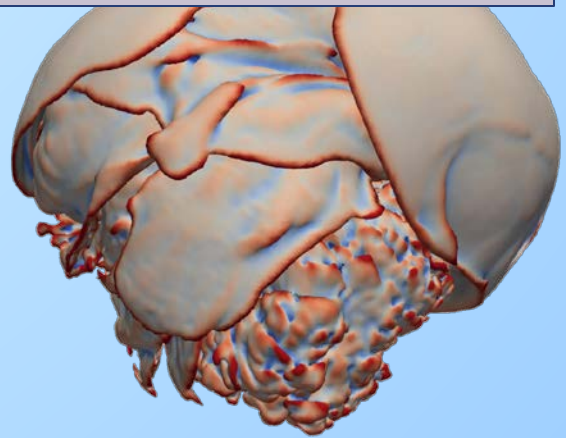

Early VTK-m prototype: vtkSMP

- *Problem:* VTK operators are single-threaded, want to support parallel operation in a way that is transparent to the application.
- *Approach:* implement a lightweight threading abstraction (vtkSMP) that draws ideas from Dax, EAVL, PISTON, and that supports various backends (platform portability).
- *Results:* vtkSMP provides means for parallelizing many filters/operators, and runs using multiple backends (Threading Building Blocks, OpenMP, Kaapi)

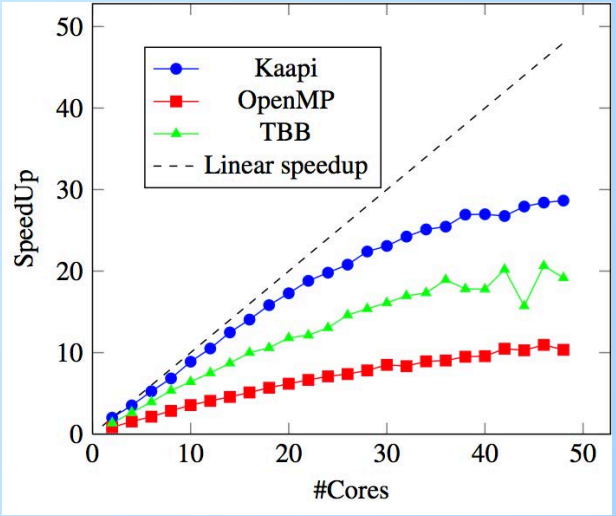

Contour Filter: not trivially parallel due to need to merge triangles. Runs on CPU and GPU. In VTK 6.1.0 (Jan 2014)

Flow visualization

Flow Visualization and Analysis

- *Problem*: the means to understand complex phenomena in flow fields; parallelizing flow analysis codes is difficult.
- New: optimization of scalable particle advection infrastructure (VisIt)
- New: OSUFlow+VTK integration
  - capabilities: flow line characterization for load balancing, analysis, out-of-core integral curve calculation.
  - Parallelization via DIY, deployment productization in VTK as vtkOSUFlow, accessible via applications (ParaView).

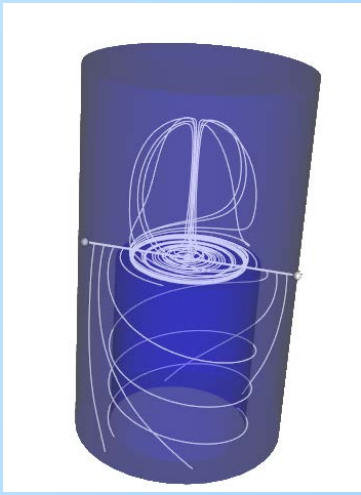

Flowlines in unstructured grid

poster

Sample Streamline

Extracted Streamlines

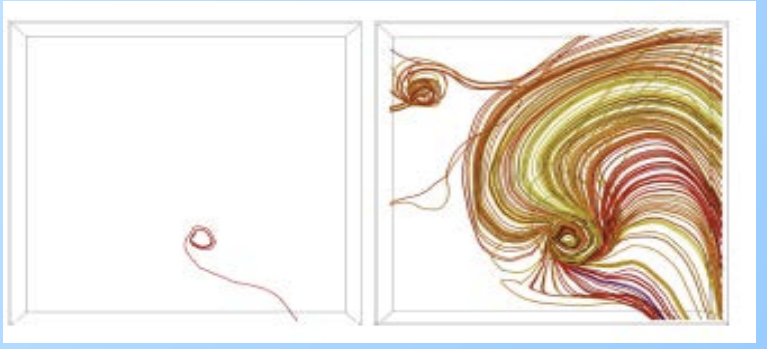

Data Analysis

Area co-leads
Valerio Pascucci, Kwan-Liu Ma

Analysis Tools Applied to New Applications

- **Feature-driven analysis**
  - *Dynamic tracking graphs and multi-core exploration, e.g.,*

  Localized thresholds for vortex detection
- **Importance driven analysis**
  - *In-situ data triage/reduction, e.g.*,

  Explorable images
- **Statistical analysis and data mining**
  - *Feature-based algorithms in MPI, e.g.,*

  Memory efficient graph analytics
- **Topological techniques**
  - *Integrate higher order feature descriptors, e.g.,*

  Abstract representation of material structure to investigate Lithium diffusion
- **Vector field analysis**
  - *Incorporate advanced statistical methods, e.g.,*

  Flowline indexing via geometric signatures

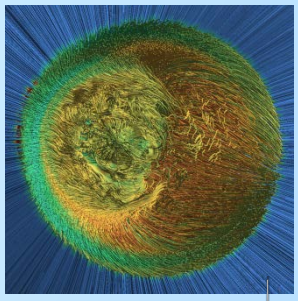

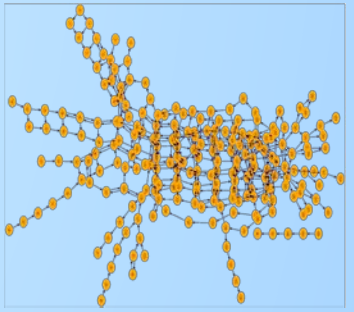

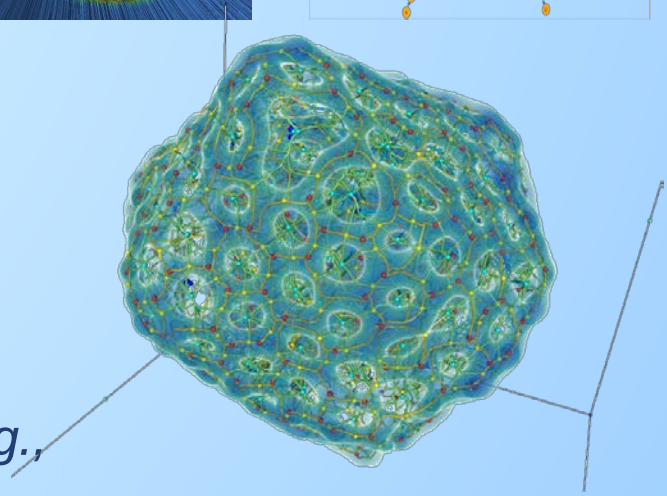

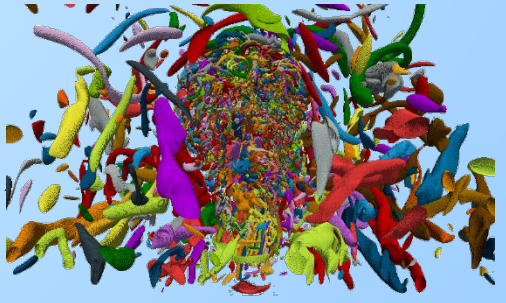

Feature-Driven Analysis:

Interactive exploration of families of features and their evolution in dynamic simulations

ExMage: Explorable Images for In Situ Computation of Vis & Analysis that is Editable in Postprocessing

- ExMage: provides in situ pathtube generation and visualization
- An *in situ* method that exploits the high temporal data resolution available at simulation time and generates images that are *explorable*.
- The resulting visualizations are usable for visual monitoring, debugging, analysis, and tuning of the simulation.
- A new data reduction solution

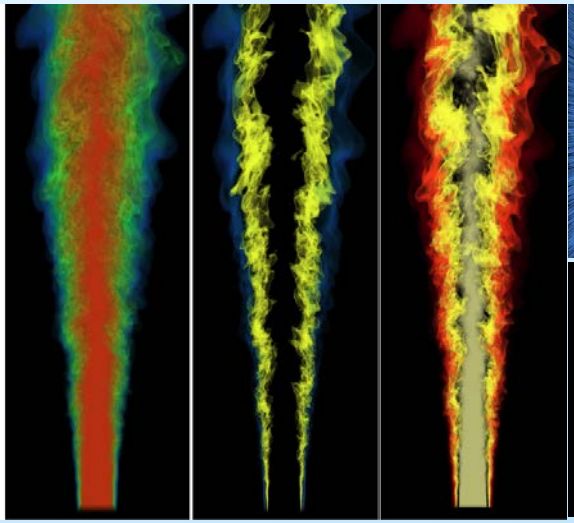

Combustion scalar field exploration

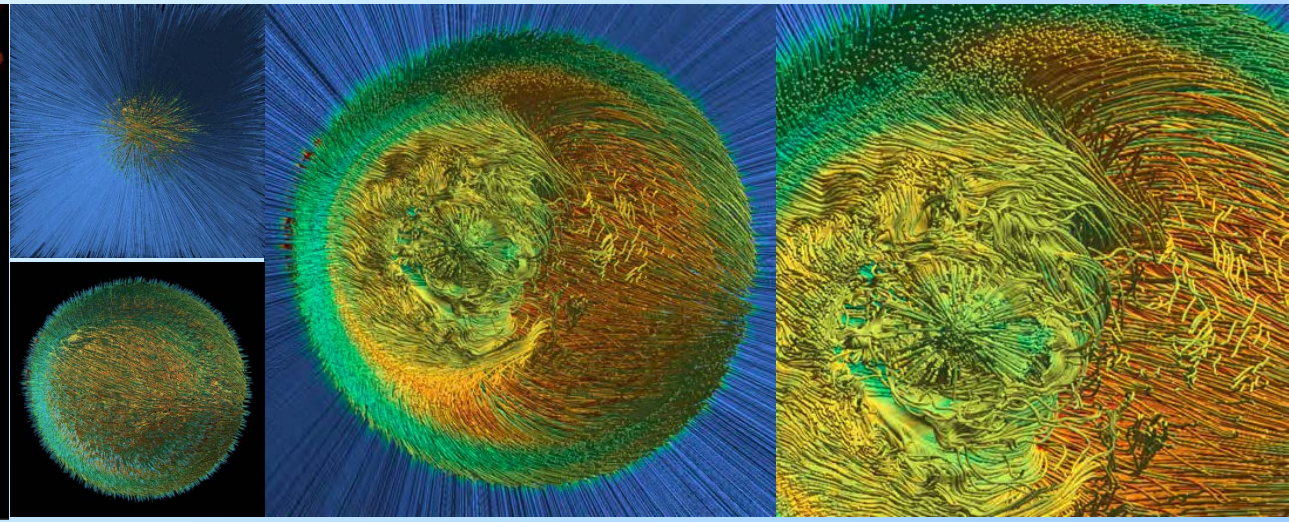

Supernova vector field exploration

Importance Driven Analysis:

Processing of Data Based on User Driven Metrics of Relevance

Parallel Computation of Distance Field for In-Situ Data Reduction and Rendering

- Distance fields can be used as importance fields to guide rendering, data compression, sampling, and feature-based optimizations.
- This project creates a highly scalable parallel implementation to support in situ processing and data reduction.
- The resulting technology will benefit many application areas from combustion, fusion, to climate and astrophysics simulations.
- Results on Intrepid show 86% efficiency at 32K cores

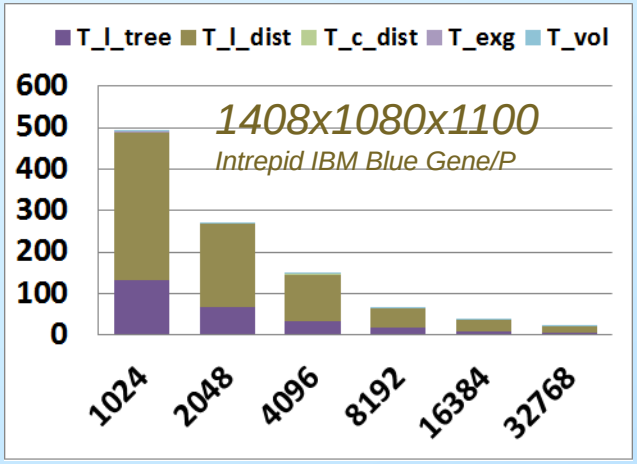

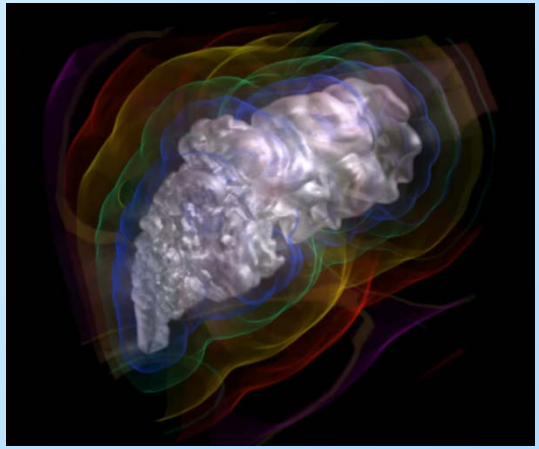

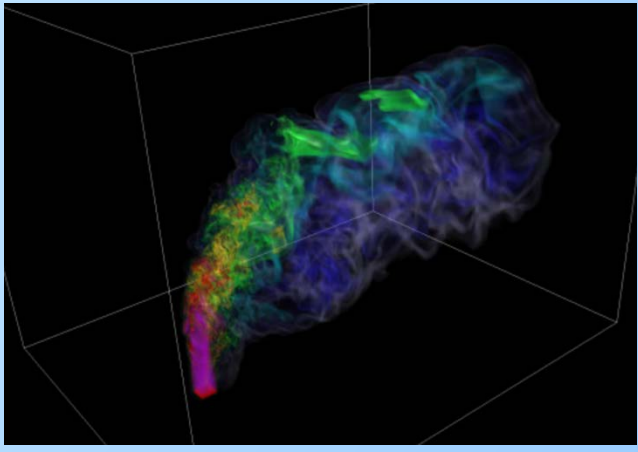

Topological Techniques:

Extraction of Abstract Data Representation Based on Complete Topological Representations

Investigate Advanced Battery Materials Through Structural Analysis of Carbon Defects

- **MSCEER:** is a set of tools and libraries for feature extraction and exploration in scalar fields
- Structural analysis of the density fields in carbon nano-spheres provides insight into Lithium Ion diffusion and storage
- Use MSCEER to generate Morse-Smale complexes and explore connectivity of void structures
- Model defects in the connectivity structure of carbon bonds to understand potential Lithium diffusion paths

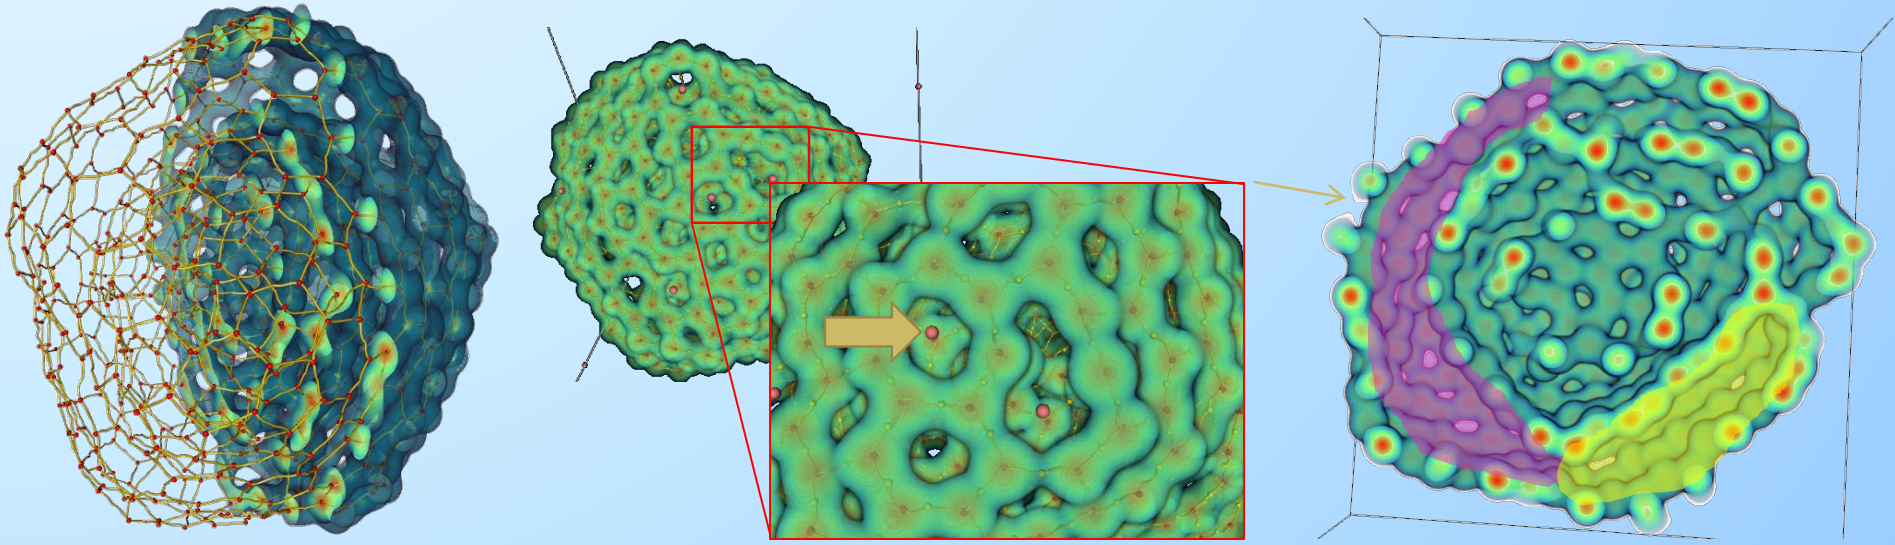

Vector Field Analysis:

Characterization of local and global trends in complex flows

Automatic Detection and Classification of Particle Beams and their Temporal Sub-Features

- Large-scale simulations of laser-plasma particle accelerators in support of ComPASS project
- Novel algorithms for fast and accurate classification of particle acceleration features
- New methods for feature-based data exploration
- First quantitative analysis of beam substructures and transverse particle loss

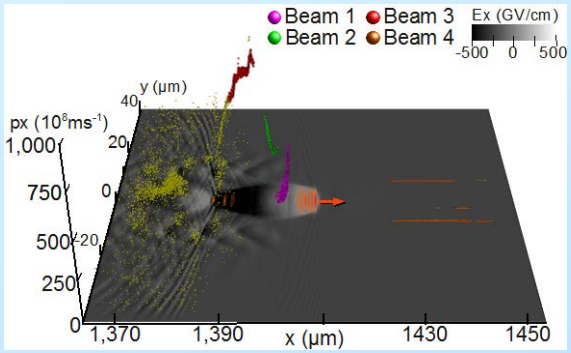

Single laser pulse

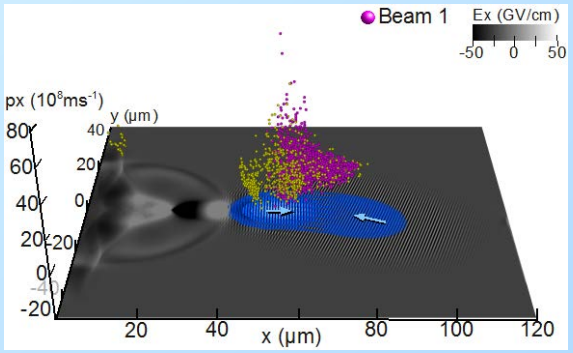

Dual colliding laser pulse

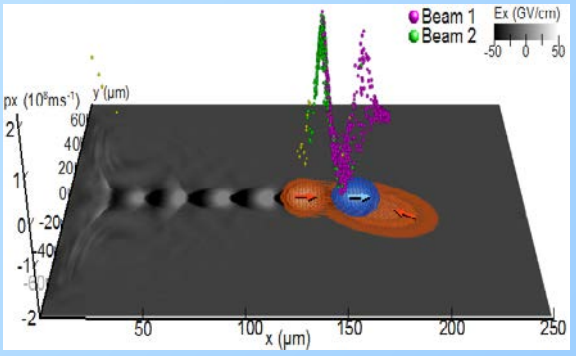

Triple colliding laser pulse

Tools and Libraries

TALASS: Addresses Analysis Challenges Through Theoretical and Practical Advances

- New data representations and techniques:
  - Parameter independent feature representation
  - Novel data transforms
- Efficient in-situ algorithms to allow
  - High frequency analysis
- Interactive tools to explore results
  - Spatio-temporal exploration of the solution space
  - Integrated per-feature statistics

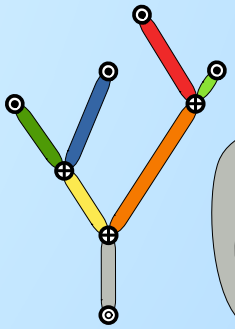

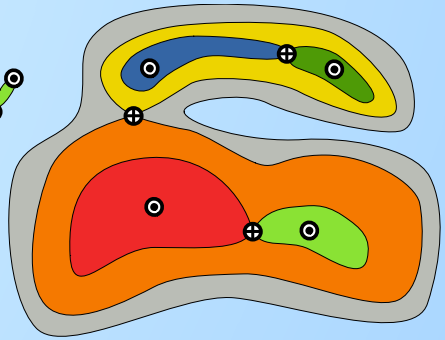

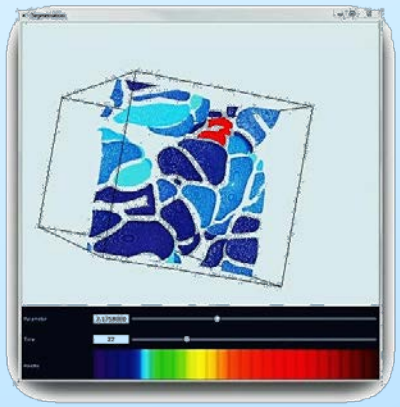

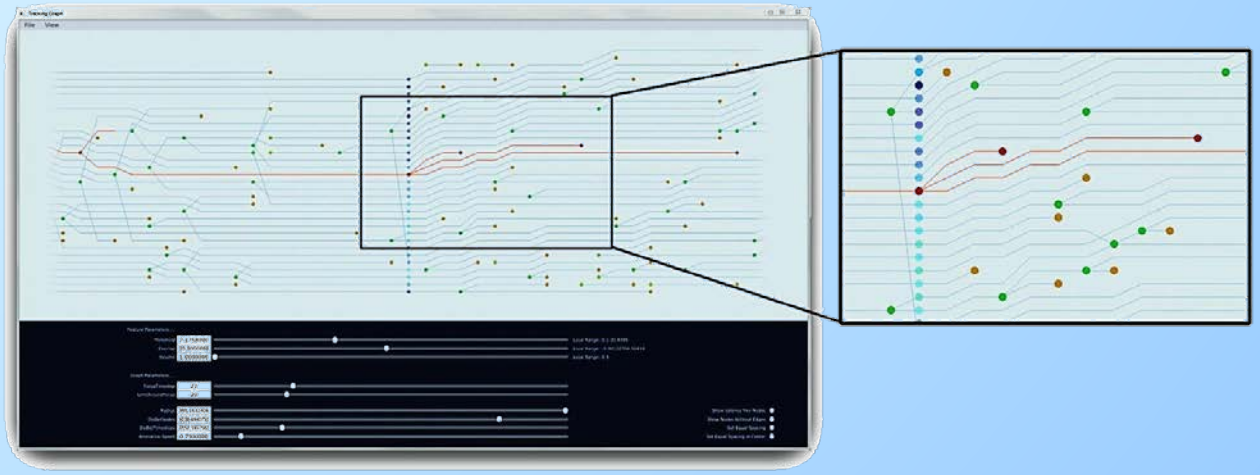

ND²AV, HDViz: High-Dimensional Analysis Tool Chains to Support New Use Cases

- Provide intuitive and theoretical insight into high-dimensional parameter spaces for a wide variety of use cases
  - Design of catalytic materials with PNNL
  - Nuclear Reactor Safety with INL
  - Climate and High-energy physics with LLNL
- Combine topology, geometry, machine learning, and statistics into interactive tools for abstract problems

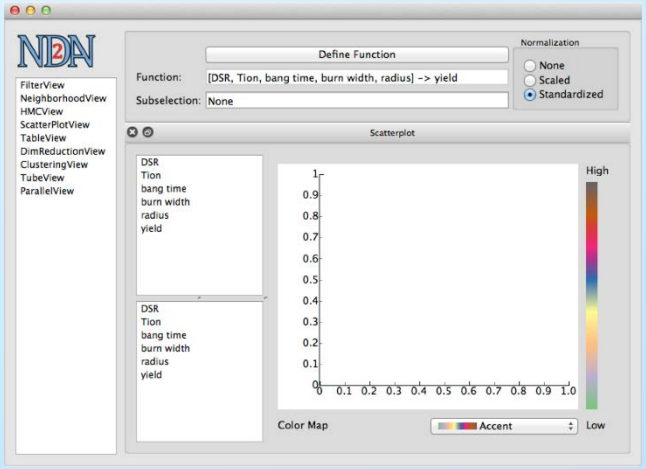

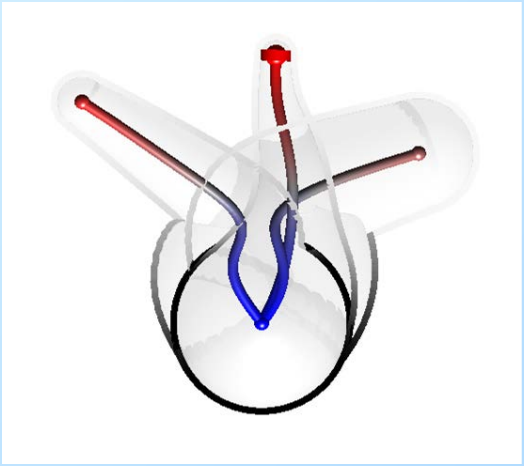

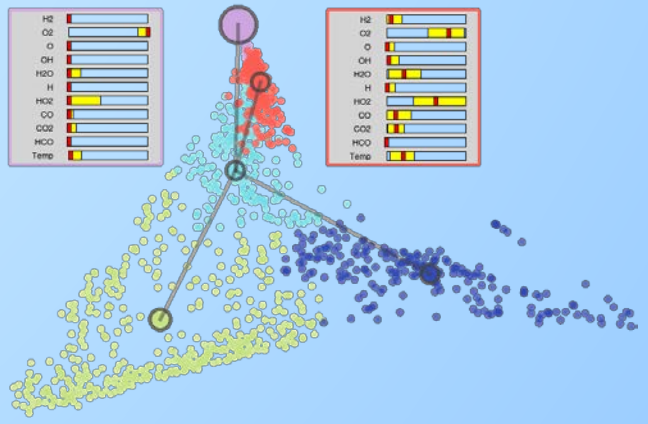

CosmoTools: In Situ Analysis for Cosmological Simulations

- CosmoTools provide a tightly integrated analysis environment for HACC Code both in-situ and stand-alone
- Enable high fidelity cross-correlation between large-scale simulations and digital sky surveys at extreme scales
- CosmoTools relies on a small, extensible API designed to integrate new analysis algorithms off-line or in-situ

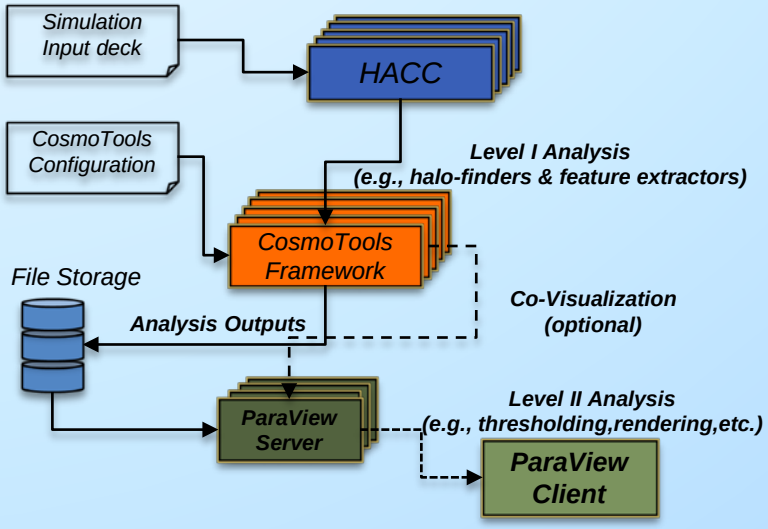

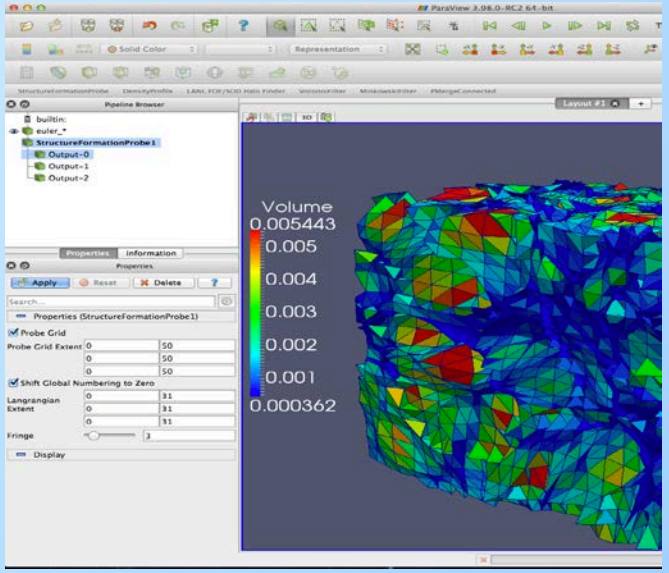

Collaborations with Domain Sciences (1)

	Project/Code	Program/ Domain	SDAV Technologies
1	Optimizing SuperConductor Transport Properties (OSCon)	BES	DIY
2	Advanced Structural Characterization using Experimental Scattering Data from Multiple Facilities	BES	ADIOS
3	High Resolution Coherent Imaging of Energy Materials	BES	DIY, PtychoLib
4	Synchrotron Data Pilot	BES	VisIt, HDF5/H5hut
5	Ice Sheet and Climate Evolution	BER	ParaView, ADIOS, VisIt
6	Center for Edge Physics Simulation (EPSI)	FES	DataSpaces, DIMES, ADIOS, VTK, VisIt, EAVL
7	Plasma Surface Interactions (PSI)	FES	EAVL, ParaView
8	Accelerator Simulation (ComPASS)	HEP	VisIt, FastBit
9	Computation-Driven Discovery for the Dark Universe	HEP	DIY, tess, ParaView, ROMIO, VTK-m (PISTON), CosmoTools, VisIt, Parallel DBSCAN, Parallel NetCDF, GLEAN
10	Nuclear Computational Low Energy Initiative (NUCLEI)	NP	ADIOS

Collaborations with Domain Sciences (2)

11	Nuclear Energy (NEAMS)	NE	VisIt
12	Explosive Hazard Predictions with the Uintah Framework	combustion	PIDX
13	Fuel-Flexible Gas Turbines and IC Engines	combustion	TALASS, MSCEER, PIDX, ADIOS
14	Global Seismic Tomography for Understanding Earth's Mantle (SpecFM3D)	geoscience	ADIOS, VisIt
15	Regional Scale Earthquake Simulation	geoscience	ADIOS
16	QMC Glue project	BES	ADIOS, VisIt
17	Magnetic Plasma Confinement in Tokamak Reactors	fusion	ADIOS
18	Relativistic Kelvin-Helmholtz Instability (PIConGPU)	physics	VisIt
19	Center for Exascale Simulation of Advanced Reactors (CESAR)	reactors	DIY, ROMIO
20	Climate Data Analysis (UV-CDAT)	climate	VisIt
21	Parallel Climate Analysis Library (ParCAL)	climate	ROMIO, PnetCDF, Ultravis-V
22	Climate Model Intercomparison (PCMDI)	climate	VisIt
23	Plasma Simulation (VPIC)	plasma	ParaView, FastQuery
24	Fluid/Plasma Simulation (BOUT++)	fusion	VisIt

Extra slides

Data Management Tools

- **I/O Frameworks**
  - **ADIOS**: provides a simple, flexible way for scientists to describe the data in their code, and based on that to provide efficient I/O, and in situ data processing
  - **Darshan**: captures an accurate picture of application I/O behavior
  - **Parallel netCDF**: a library providing high-performance I/O while still maintaining file-format compatibility with Unidata's NetCDF
  - **ROMIO**: is a high-performance, portable implementation of MPI-IO
  - **ViSUS/IDX**: Provides data streaming techniques for progressive processing and visualization of large surface and volume meshes
- ***In Situ* Processing**
  - **GLEAN:** is an extensible framework that takes system characteristics into account in order to facilitate simulation-time data analysis and I/O acceleration
  - **DIY**: Provides scalable building blocks for data movement tailored to the needs of large-scale parallel analysis workloads
  - **DataSpaces**: facilitates *in situ* code coupling using a shared-space abstraction
  - **EvPath**: is an event transport middleware layer providing processing over virtual data paths
- **Indexing / Compression**
  - **FastBit:** A very fast indexing method based on compressed bitmap representation specially suitable for scientific data
  - **ISABELA:** a tools for lossy but highly accurate (> .99 correlation) compression of spatio-temporal scientific data

Analysis and Visualization tools

- **Analysis and Visualization Frameworks**
  - **VisIt:** is an Open Source, interactive, scalable, visualization, animation and analysis tool.
  - **ParaView**: is an open-source, multi-platform data analysis and visualization application.
- **Analysis and Visualization Libraries**
  - **ExMage**: provides in situ pathtube generation and visualization.
  - **TALASS**: is a collection of routines for parallel and distributed processing of particle data.
  - **Ultravis-P**: is a collection of routines for parallel and distributed processing of particle data.
  - **MSCEER:** is a set of tools and libraries for feature extraction and exploration in scalar fields.
  - IceT: is a high-performance sort-last parallel rendering library that provides the unique ability to generate images for tiled displays.
  - **NDDAV**: is an interactive analysis framework for high-dimensional data.
  - VTK: is an open-source system for 3D computer graphics, image processing and visualization.

Analysis and Visualization tools

- **Multi-/Many-core Visualization Libraries**
  - **Dax:** The Dax Toolkit supports the fine-grained concurrency for data analysis and visualization algorithms required to drive exascale computing.
  - **EAVL:** is the Extreme-scale Analysis and Visualization Library that expands traditional data models to support current and forthcoming scientific data sets.
  - **PISTON:** is a cross-platform software library providing frequently used operations for scientific visualization and analysis.
- **Statistics and Data Mining**
  - **NU-Minebench**: is a data mining benchmark suite containing a mix of several representative data mining applications from different application domains.
  - **STPMiner:** is a collection of spatiotemporal data mining algorithms including: outliers/anomalies, colocation patterns, change patterns, clustering, classification, and prediction algorithms.
  - **Importance-Driven Analysis**: is a tool that uses a newly-designed spatial data structure, named parallel distance tree, to enable highly scalable parallel distance field computing.