Exploring fusion science with SDAV's technologies

J-Y. Choi (ORNL), J. Dayal GA Tech), G. Eisenhauer (GA Tech), T. Jin (Rutgers), S. Klasky (ORNL), Q. Liu (ORNL), N. Podhorszki (ORNL), M. Parashar (Rutgers U.), D. Pugmire (ORNL), K-L. Ma (UC Davis), A. Sanderson (Univ. Utah), N. Samatova (NCSU), K. Schwan (GA Tech), A. Shoshani (LBNL), B. Whitlock (LLNL), M. Wolf (GA Tech/ORNL), J. Wu (LBNL), F. Zhang (Rutgers.Univ), Y. Zhang (UC Davis), F. Zheng (GA Tech), X. Zhang (GA Tech) and the SDAV team

Gyro-Particle Simulations

XGC, GTS, GTC-P

Large-scale I/O with ADIOS

Objectives:
- Fast, adaptable, and scalable I/O codes across all platforms
- Unified I/O framework to hide I/O complexity from users and manage individual I/O methods. Framework handles both file and stream processing.
- Non-contiguous parallel file format, highly scalable and resilient for large scale simulations.
- Support streaming methods: DataSpaces, DIMES, Flexpath

Impact:
- Checkpoint writing in XGC-1, GTC, and other large scale codes with low I/O overhead.
- E.g., XGC-1, 200k cores on Jaguar, 2TB checkpoint data to write: original parallel I/O >1 hour, ADIOS ~1 min.
- Code coupling with memory-to-memory data transfers & in situ visualization using staging methods (see other columns)

Scibox: Online Sharing of Scientific Data via Cloud

Objectives:
- A cloud-based write/read ADIOS method with data reduction
- Support partial object access in cloud
- Reduction functions are implemented using C-on-Demand, a dynamically-generated subset of C

Impact:
- Speed up sharing scientific data over the cloud.
- An order or magnitude faster read speed for analysis code of GTS output

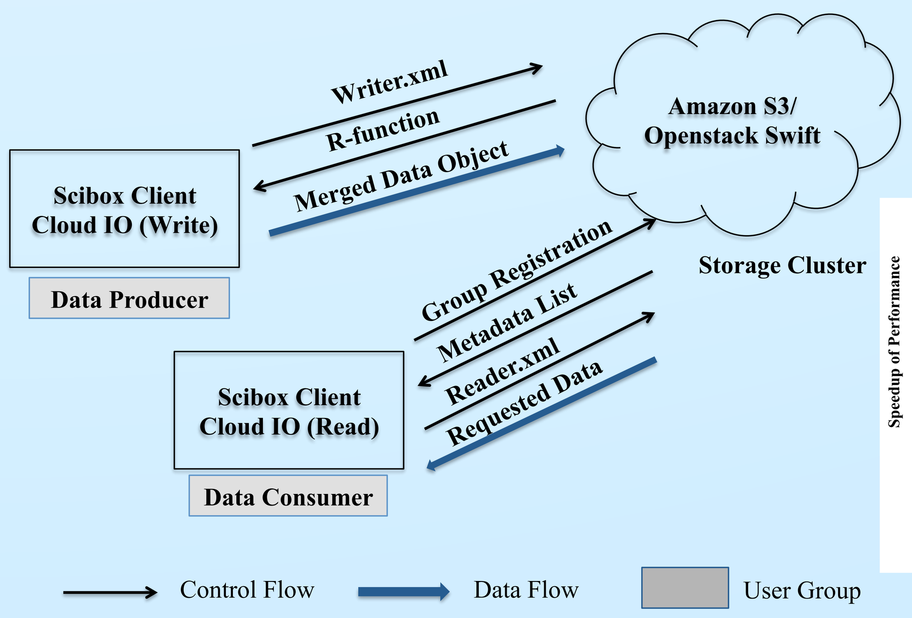

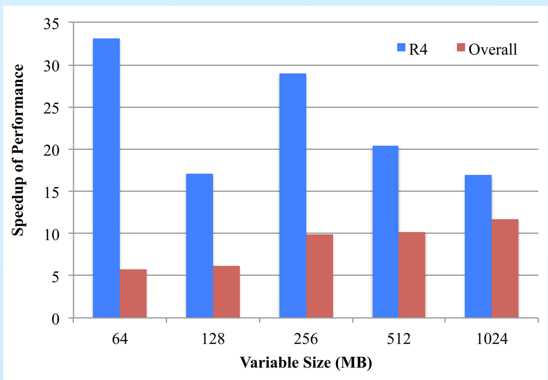

Type	Description	Example
R1	Max(variable)	Max(var_double_2Darray)
R2	Min(variable)	Min(var_double_2Darray)
R3	Mean(variable)	Mean(var_double_2Darray)
R4	Range(variable, dimensions, start_pos, end_pos)	Range(var_int_1Darray, 1, 100, 1000)
R5	Select(variable, threshold1, threshold2)	select var.value where var.value ∈ (threshold1, threshold2)
R6	Select(variable, R_Function1, R_Function2)	select var.value where var.value ≥ Mean(var)
R7	Select(variable1, variable2, threshold1, threshold2)	selcet var2.value where var1.value ∈ (threshold1, threshold2)
R8	Self defined function	double proc(cod_exec_context ec, input_type *input, int k, int m) { int i;int j; double sum = 0.0; double average = 0.0;for(i = 0; i<m; i= i+1) sum = sum + input.tmpbuf[i+k*m]; average = sum / m;return average; }

Energetic Particles

GTC, M3D-K

FastQuery: Finding Regions of Interest with Complex Mesh

Objective: locate regions of interest based on user specified range conditions.

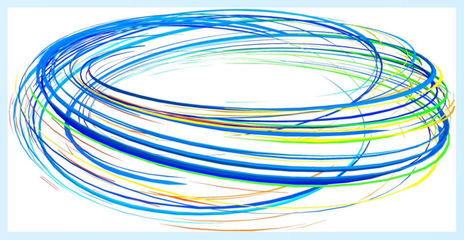

Example: a regions of special interest in a fusion simulation data might be regions of space with extreme high electromagnetic potential. The illustration shows each connected region in GTC with a different color.

Challenge: GTC and similar modeling code uses complex meshes that make identifying connected regions difficult. Approach: define connectivity based on the magnetic coordinate system instead of Cartesian coordinates.

Impact: improves speed of locating spatial features hundreds of times

A summary of identifying regions defined by two types of range query conditions. Our approach based on the magnetic coordinate system is hundreds of times faster than the common approach based on the Cartesian coordinate system.

			Labeling Time		
Query type	Avg # of Points	Avg # of regions	Cartesian Coord	Magnetic Coord	Speed up
1RQ	5.4×10^6	609	44,060	45.6	**966**
2RQ	2.6×10^6	17400	19,016	32.0	**546**

Range Based Queries

Objectives:
- Enable scientists to query large scale simulations to find "interesting" particles
- Develop a framework to utilize fast indexing schemes such as FastBit.

Impact:
- Production tool in VisIt providing an interactive environment.
- Temporal range queries on multivariate data for finding correlations between variables.
- Multidisciplinary usage, e.g., biomass, fusion.

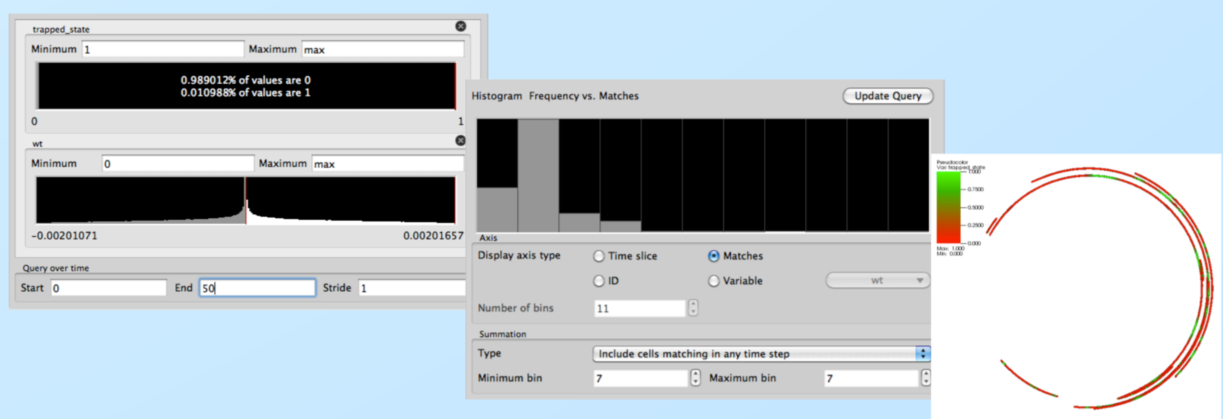

A range query over 50 time steps for 500K particles in GTC using two variables, resulting in 8 particles with their paths shown. Color of particles is based on being passing (green) or trapped (red).

Magnetohydrodynamic codes

M3D, Pixie3D, NIMROD, Siesta, VacField

Parallel vector field analysis in VisIt

Objectives:
- Enable scientists to understand the nature of velocity and magnetic fields in large scale simulations.
- Flexible, scalable framework for computing integral curves.
- Research algorithms for efficiently computing integral curves on large scale parallel computing platforms.

Impact:
- Production tools that provide scalable, parallel methods for computing integral curves which are the basis for fieldlines, streamlines, and pathlines.
- Multidisciplinary usage, e.g., astrophysics, fusion, radiation transport.
- Techniques deployed in VisIt and have become the foundation for building more complex analysis and visualization tools: Poincare plots and Finite Time Lyapunov Exponents (FTLE).

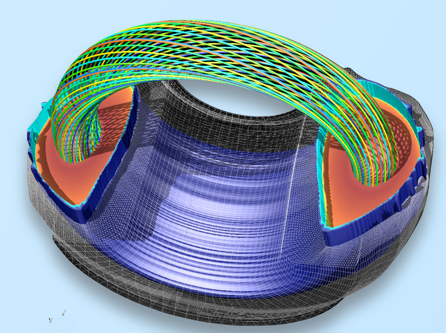

Fieldlines from a M3D simulation

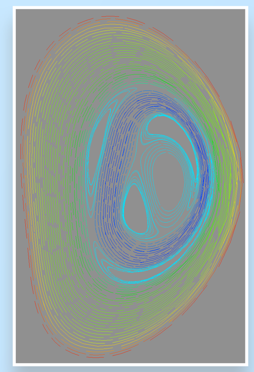

Poincare plot of a DIII-D disruption using NIMROD

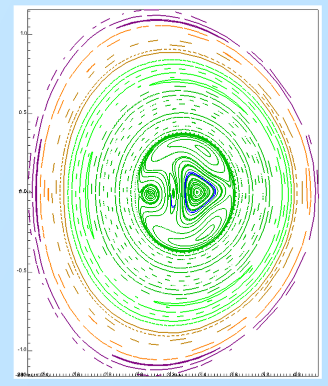

Poincaré plot of a sawtooth crash using M3D-C1

Advanced Visualization Techniques for Fusion Data

Objectives:
- High speed and high fidelity rendering design for visualizing 3D field data from tokamak simulations. (Left Image)
- Data decomposition techniques for addressing the complexity of data to study different aspects of the physics. (Rigth Image)

Impact:
- Enables scientists to better validate their simulations and to better communicate with others their data and findings.

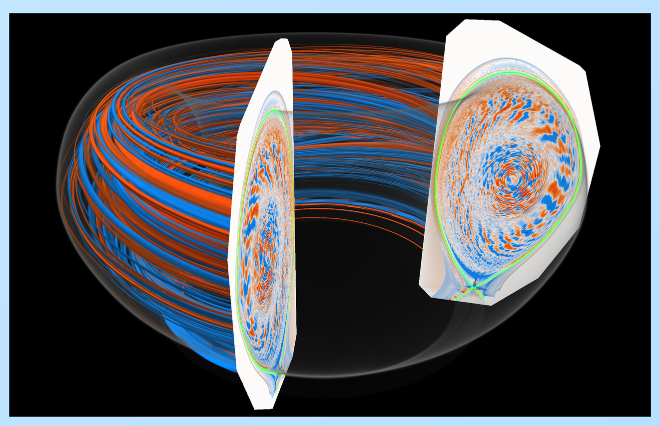

Electrostatic potential perturbation by Ion-Temperature-Gradient (ITG) driven plasma turbulence in a realistic diverted DIII-D tokamak geometry. XGC1 code.

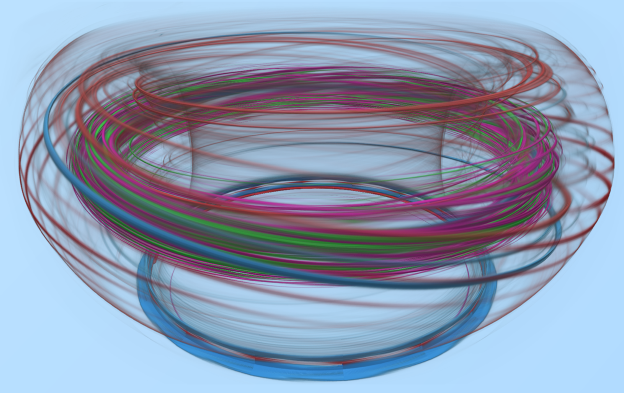

Visualization of different frequency bandpass data (f>30kHz and f<5kHz) shows different aspects of plasma and turbulence physics.

Exploring fusion science with SDAV's technologies

DataSpaces for In Situ / In Transit Data Staging and Analytics

Objectives:

- Combine in-situ/in-transit data analytics execution to move the analytics closer to the data, reduce data movement costs, mitigate disk IO costs, and accelerate the data to insights process.

DataSpaces

- Semantically-specialized virtual shared space abstraction with geometry-based queries.
- Asynchronous coordination and interaction.
- In-situ/in-transit mapping, scheduling and execution

Impact:

- A novel scalable approach to effectively transform large-scale simulations data into insights.
- ADIOS/DataSpaces is a production tool providing scalable memory-to-memory data transfers
- Used in EPSI project for coupling fusion codes.

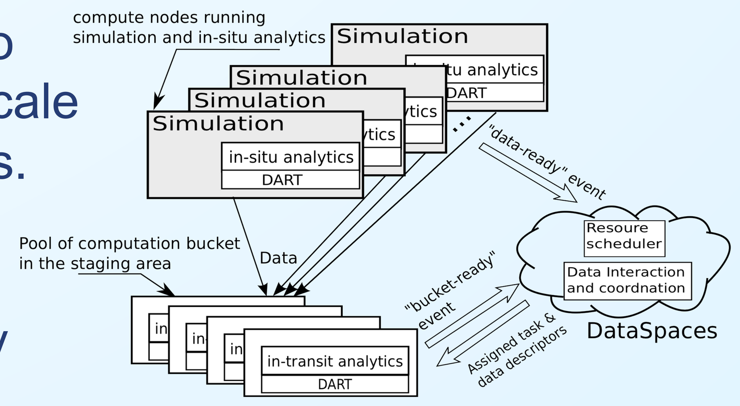

DataSpaces Architecture

Monitoring simulations with eSiMon

Objectives:

- The eSiMon dashboard:
- Uses data, web, and workflow service infrastructure for flexibility and portability.
- Keeps track of provenance information (complete data lineage) which is key to ease of use and efficiency.
- Updates output from a running simulation in real time.

Impact:

- Allows scientists with different backgrounds and levels of expertise to work together using one single online tool for analysis, visualization, and data movement.
- Allows collaborators to check/analyze a running simulation.

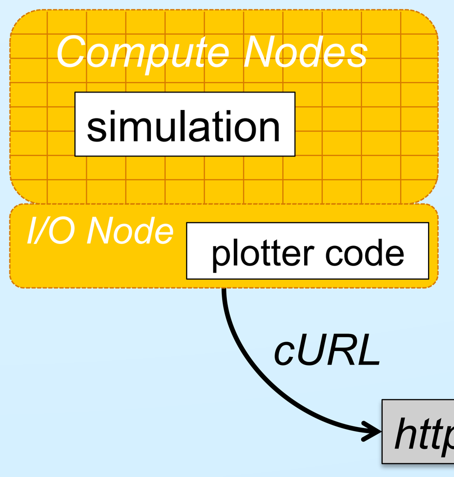

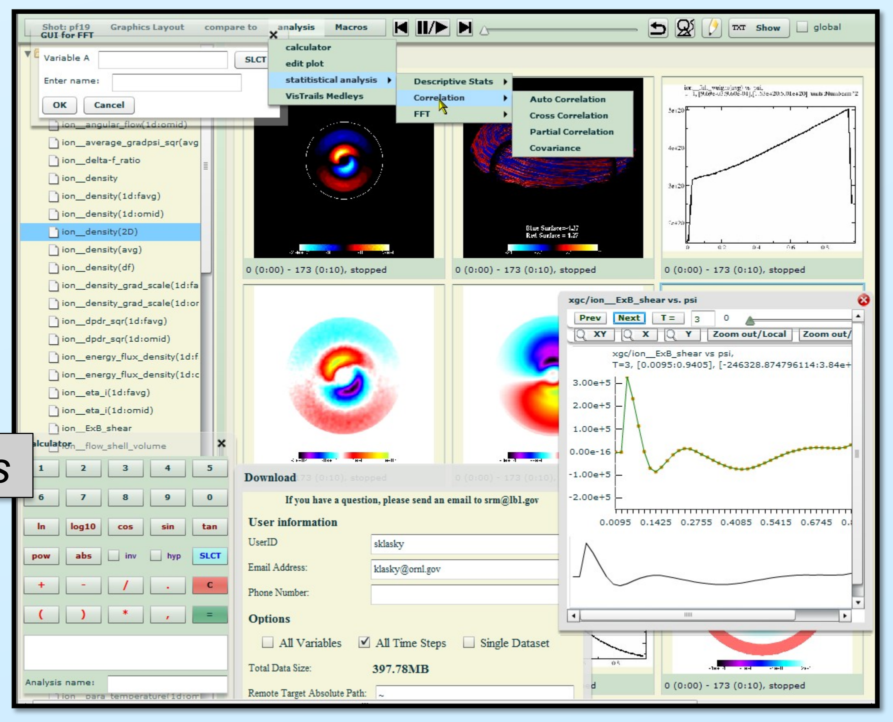

An XGC1 run in eSiMon. Movies generated by a workflow are shown as well as interactive vector plotting

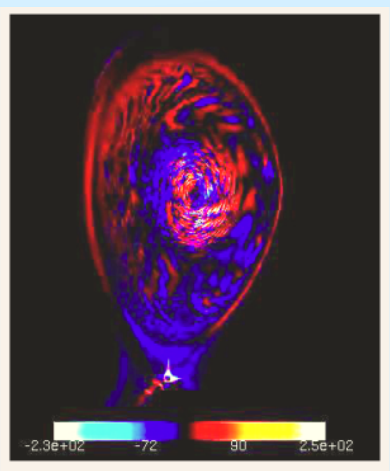

Movie frame of XGC1 potential field

FlexPath Messaging

Objectives:

- Efficient execution of parallel I/O pipelines
- Data staging methods for running analytics and visualization
- Data streaming and the online QoS control of such data streams
- Aggressive use of source-based data reduction and filtering
- convenient ways to carry out remote data visualization

FlexPath

- An event-based messaging middleware
- Active messaging provides a publish/subscribe paradigm
- Built on top of the EVPath event-based messaging layer
- C-on-Demand: users can create and deploy data conditioning plug-ins that can perform filtering and transformation operations on live data streams

Impact:

- Allows for decoupling of interacting components
- Ability to customize data streams
- Improved fault management
- Enable in situ analytics (histograms, sorting, range queries) for GTC and GTS.

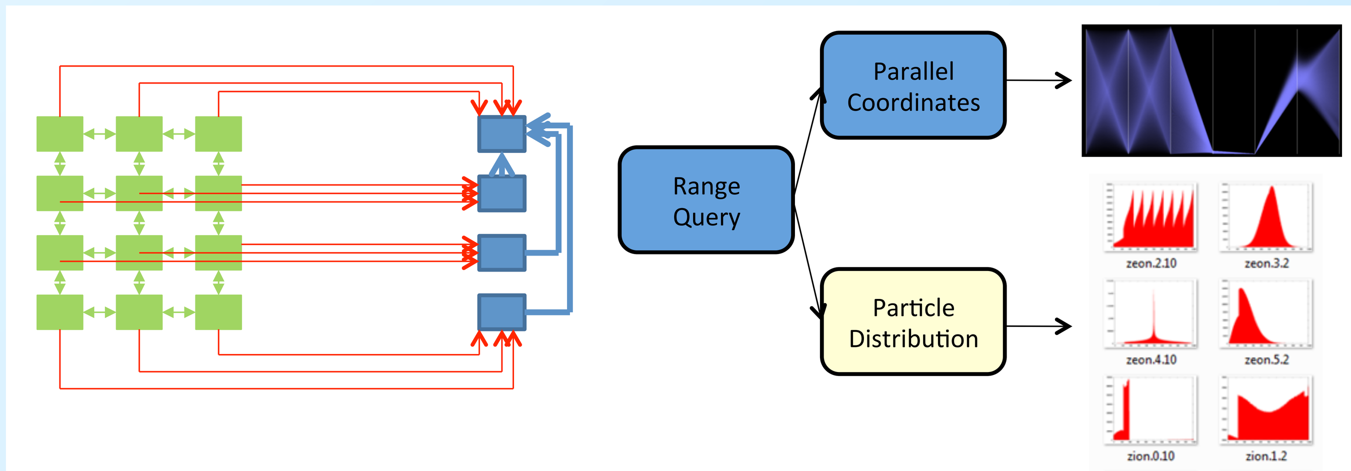

Reduction tree that implements the range query over GTS particles, resulting in the visualizations at right

In transit Analysis and Visualization

Objectives:

- Allow for concurrent analysis and visualization tasks
- Memory-to-memory transfers
- Asynchronous stepping (simulation is running free from analysis)

Impact:

- ADIOS/DataSpaces is a production tool that provides scalable memory-to-memory data transfers using file I/O
- Used for Pixie3D' s non-scalable analysis (PixPlot) and interactive visualization with ParaView

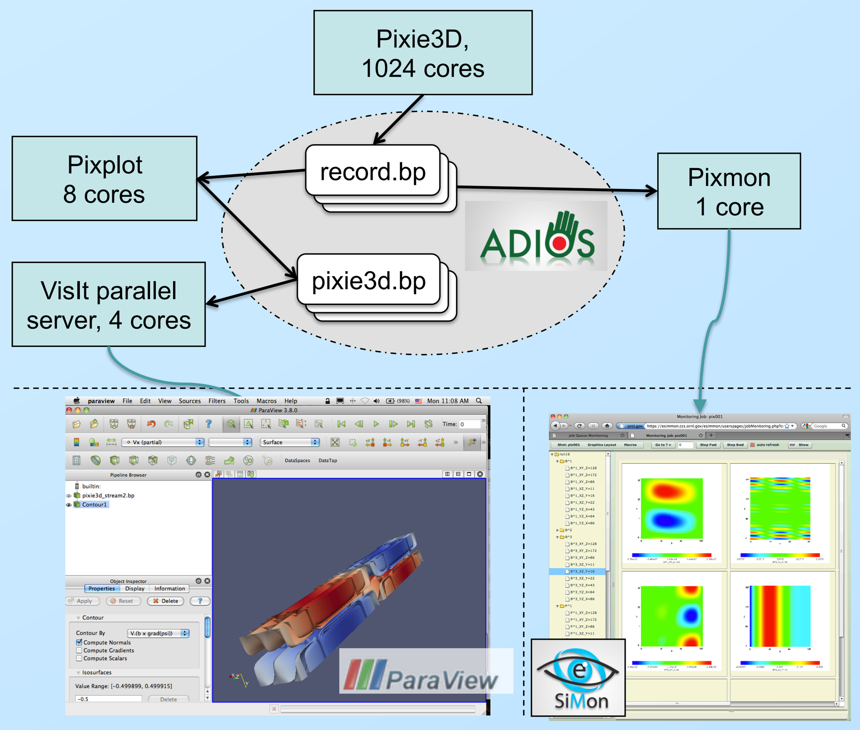

In situ Pixie3D fusion simulation analysis/viz pipeline

Data Compression

Objectives:

- Robust and parallelizable compression library that improves over widely used classical wavelet transform.

Impact:

- Four-to-five times storage reduction on over 20 different datasets from applications (XGC, S3D, GTS, GroundWater, GCRM) with 0.99+ correlation and 0.01-normalized point-by-point error.

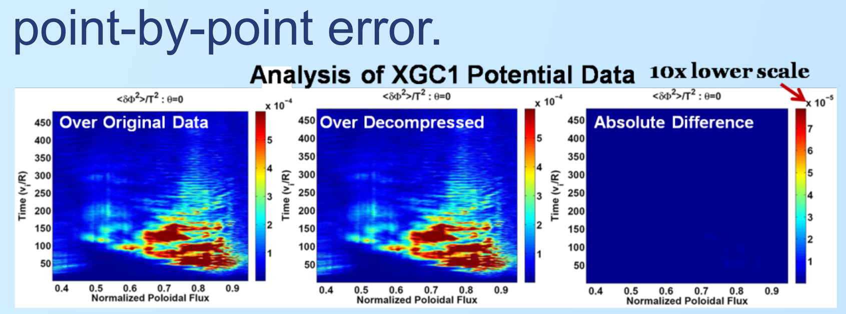

XGC-1 Potential data after ISABELA compression shows no differences from the original at applicable resolutions

Tight-coupling of Fusion Codes for Exascale Computing

Problem:

- Goal is a mathematically tight kinetic-kinetic coupling in XGC, to extend the first-principles full-f simulation to experimental time scale

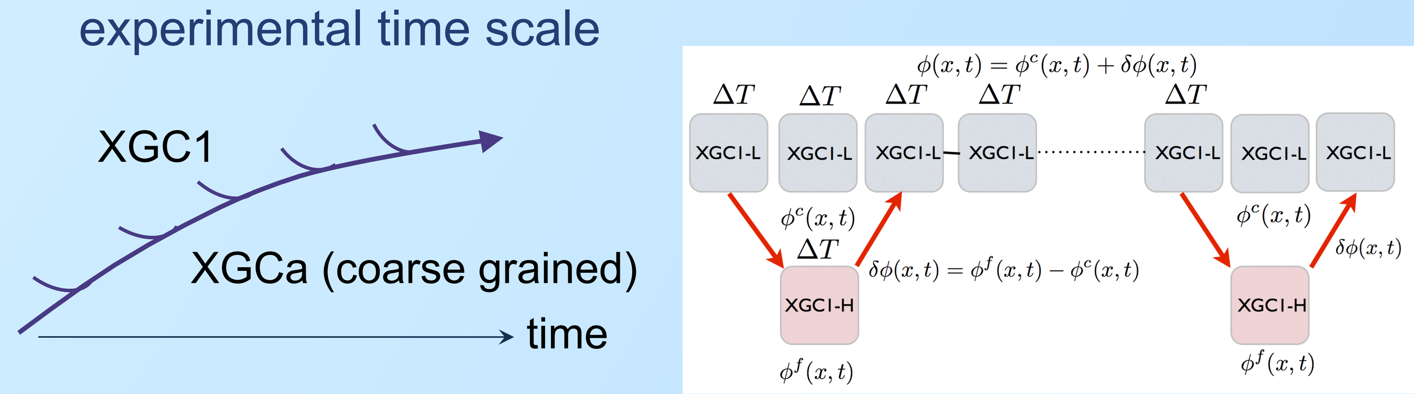

- Development on extreme scale in-memory data management, analysis, and visualization (as well as the state-of-the-art applied math) is desperately in need

Objectives:

- Locality-aware process placement, moving interacting processes closer to take advantage of data locality
- Reduce network data traffic by transferring only the analytic kernels and retrieving the results
- Reduce application execution time by offloading and executing data computations in parallel

Impact:

- Enables predictive capabilities using codes from first-principles coupled to faster (time-scale) codes

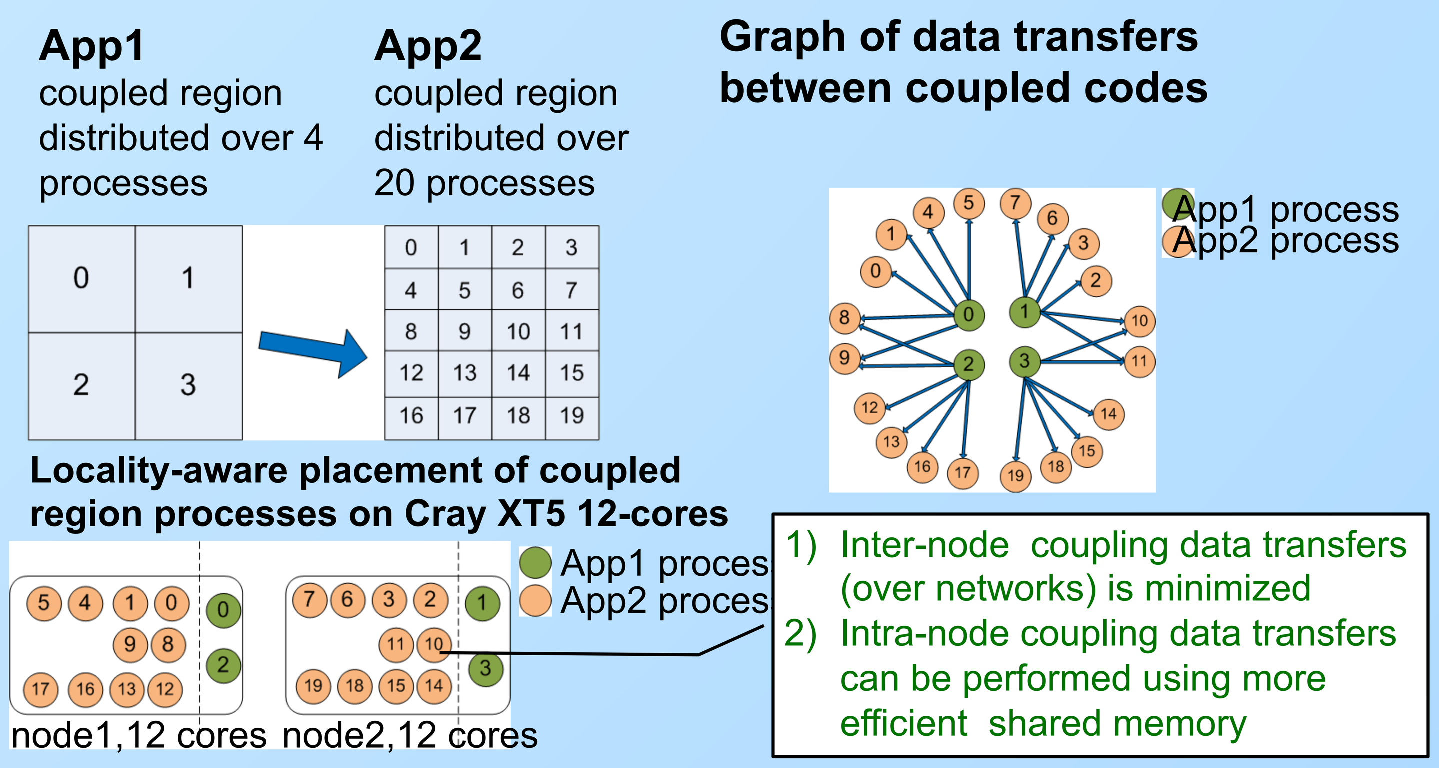

Locality-aware placement of coupled region processes on Cray XT5 12-cores using DataSpaces (IPDPS 2012)

SDAV Scalable Data Management, Analysis, and Visualization

U.S. DEPARTMENT OF ENERGY | Office of Science