SciDAC Insti SAND2011-2710P alization

UltraVis Research Overview

DOECGF

April 29, 2011

Kenneth Moreland
Sandia National Laboratories

With contributions from:

Jian Huang, Kwan-Liu Ma, John Owens, Tom Peterka, Rob Ross, Han-Wei Shen, Debora Silver

Release Marking (e.g. Not Approved for Release, SAND XXXX, etc.)

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

UltraVis
SciDAC Institute for Ultrascale Visualization

Breadth of UltraVis Research

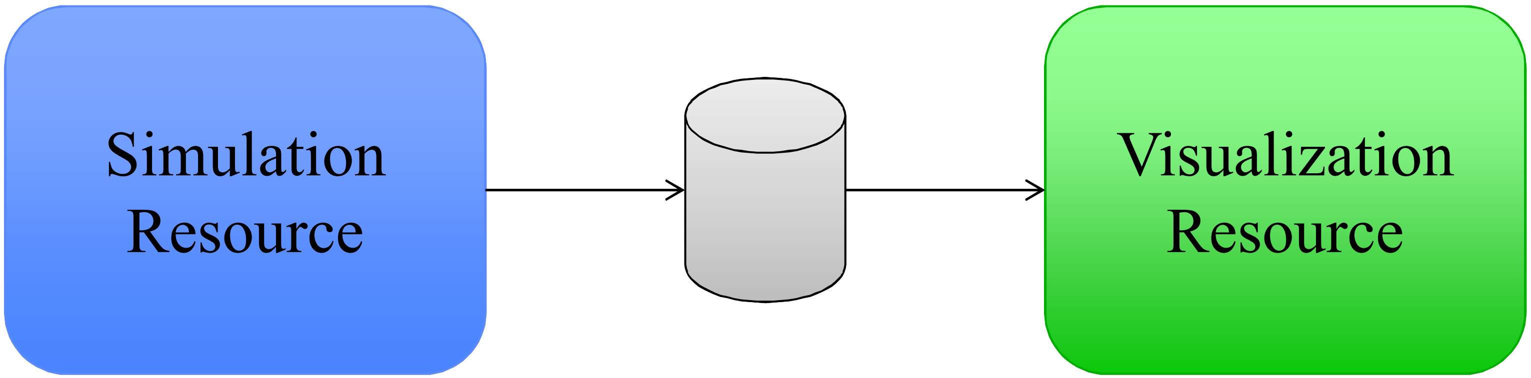

- In Situ Visualization
- Parallel Rendering
- Feature Tracking

- Parallel File I/O

- Particle Tracking
- Building Blocks
- GPGPU
- Query Interfaces

UltraVis
SciDAC Institute for Ultrascale Visualization

In Situ Visualization of Ultra-Scale Simulations

- When scientific supercomputing reaches exascale, new discoveries will become possible. However, scientists would not be able to study the full extent of the data generated by the high-resolution simulations with conventional post-processing data analysis methods.

- Computer scientists at the SciDAC Institute for Ultrascale Visualization have pioneered in situ data analysis and visualization of very large scale simulations.

- They have demonstrated in situ visualization of turbulent combustion simulations at the largest scale ever using the DOE Leadership Computing Facilities at ORNL, promising a solution to the upcoming exascale data analysis problems.

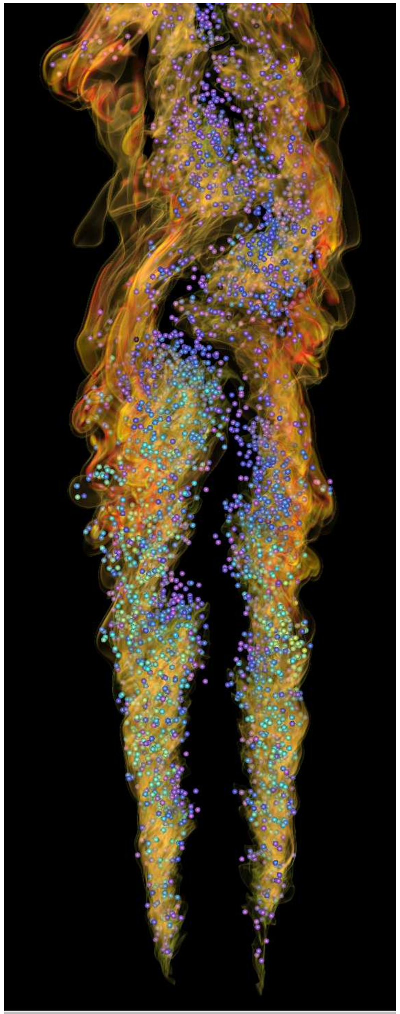

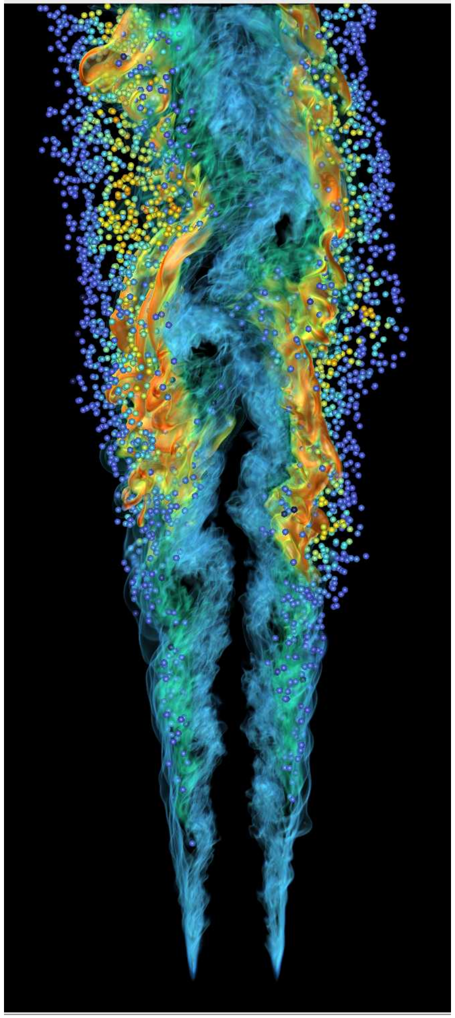

Turbulent Combustion Simulation of Lifted Flame

UltraVis
SciDAC Institute for Ultrascale Visualization

Parallel Image Composition With Radix-k

- **Problem**
  - Large-scale computational science requires parallel visualization to analyze data effectively. Image composition merges parallel results into one final image.
  - Composition algorithms were developed in the mid '90s when supercomputing architectures were very different than current and future ones.
- **Solution**
  - Take advantage of modern supercomputer network architectures with a better algorithm that adapts to largest HPC leadership machines and smaller graphics clusters alike
  - Implement in a popular image composition library called IceT
- **Impact**
  - Enables science results sooner; wall-size images can be composed in less than one millisecond.
  - Up to 5X faster than existing algorithms at full leadership scale
  - ParaView and VisIt can now take advantage of Radix-k via IceT

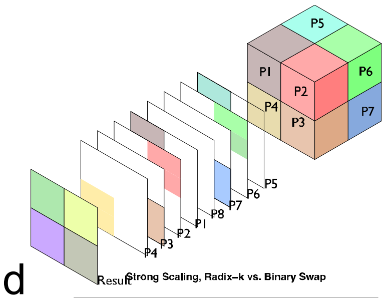

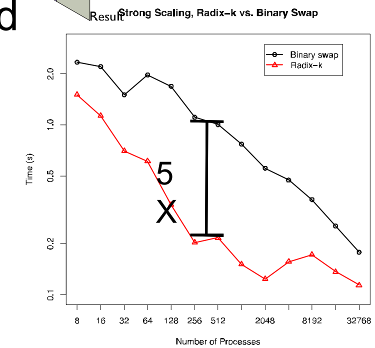

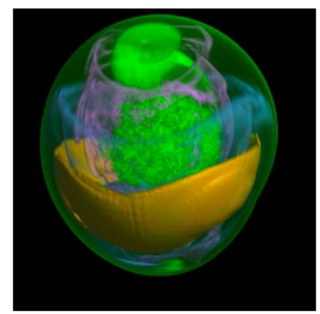

Binary swap | IceT | Radix-k | IceT-Radix-k

Direct-send | Compression | 2-3 swap | Optimized Radix-k

1993 | 1994 | 1998 | 2001 | 2007 | 2009 | 2010 | 2010

A Configurable Algorithm for Parallel Image-Compositing Applications. Peterka et al., SC'09

How the Radix-k Algorithm Works

SciDAC Institute for Ultrascale Visualization

Benefits

- Increase concurrency while managing contention
  2 <= k <= p
- Overlap communication with computation by nonblocking and careful ordering of operations
- No penalty for non-powers-of two numbers of processes is inherent in the algorithm design

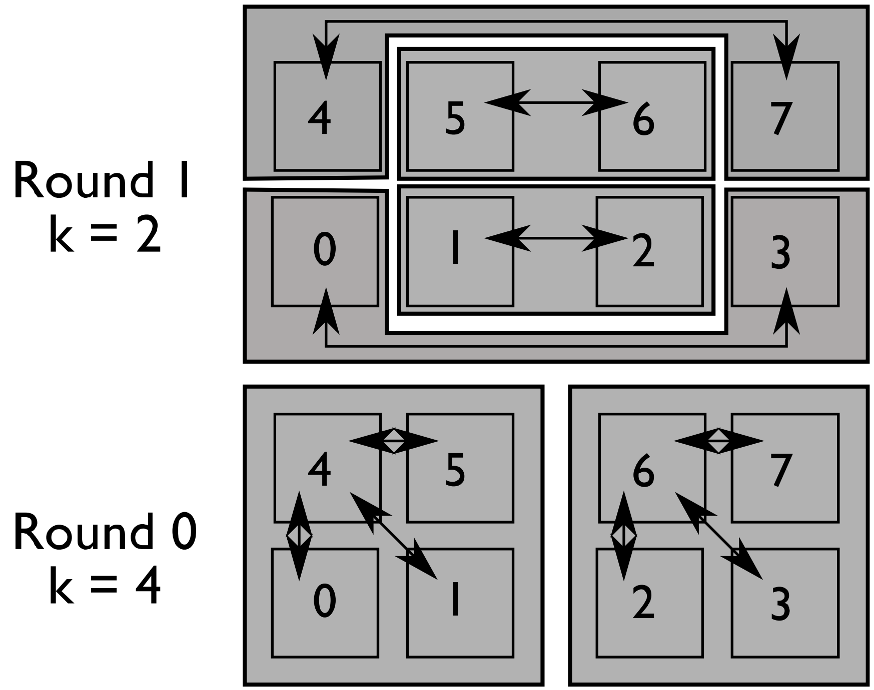

Example of 8 process compositing in 2 rounds of k=4 followed by k=2

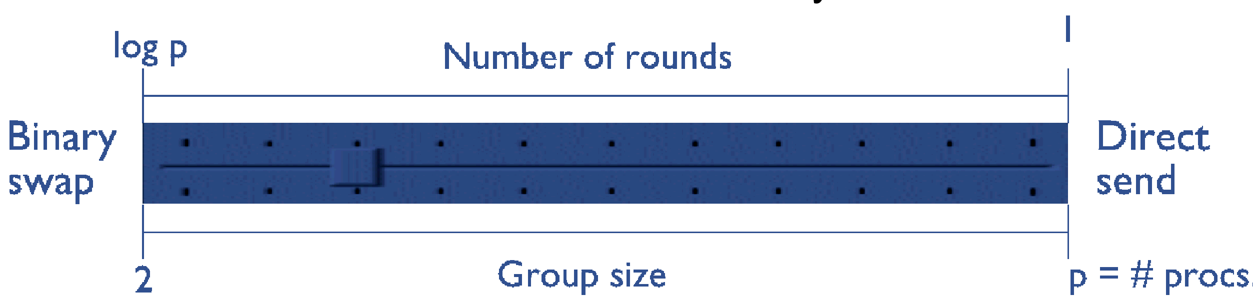

Radix-k covers binary swap, direct send, and many other configurations between these two extremes

Accelerating and Benchmarking Radix-k Image Compositing at Large Scale. Kendall et al., EGPGV'10

SciDAC Institute for Ultrascale Visualization

Activity Recognition in Scientific Visualization

MOTIVATION: It is infeasible to visualize and search for events manually in ultrascale data → automated techniques to search and find features and events are necessary.

Activity recognition is the process that extracts specific temporal patterns (events) from a set of time varying datasets. It includes defining primitive events (actions) through tracking and complex events (activities) that are a specified sequence of primitive events.

- Activity recognition deals with the inference problem therefore it can answer a new set of questions for the scientists. (Such as how features evolve, whether particular events take place, etc.)
- Defining scientific events and cataloguing features is part of activity recognition

The logical relations between activity recognition, feature tracking and segmentation

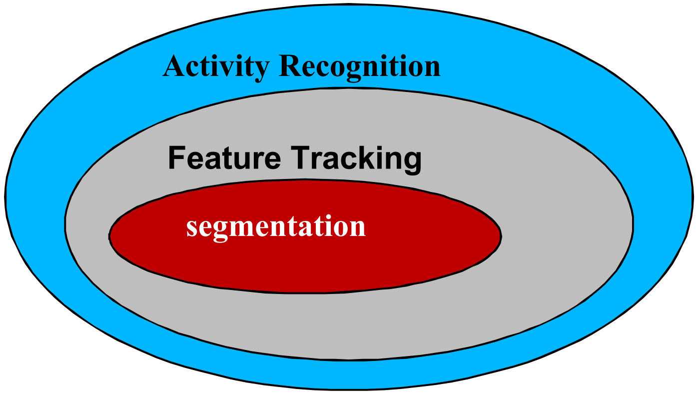

SciDAC Institute for Ultrascale Visualization

Activity Recognition

- An activity can be modeled as a sequence of atomic events (actions).
- Petri Nets can be used to define this activity
  - 5 tuple: {P,T,M,I,O}, where P=Place, T=Transition, M=Initial Marking (state),I defines the arcs and their weights from Places to transitions and O defines the arcs from transitions to Places. (where I and O can be defined as Matrices).
  - Places are the object states and transitions are the actions. Markings can be object IDs or the number of objects in each state.
  - This can then be used to search for activity within large datasets

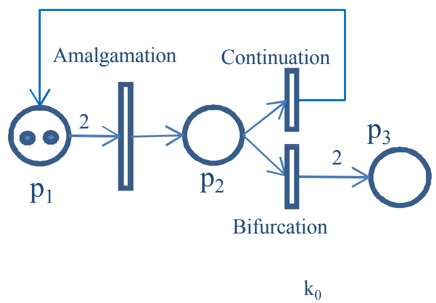

Example:
Petri Net for searching for a “Merge Split” activity

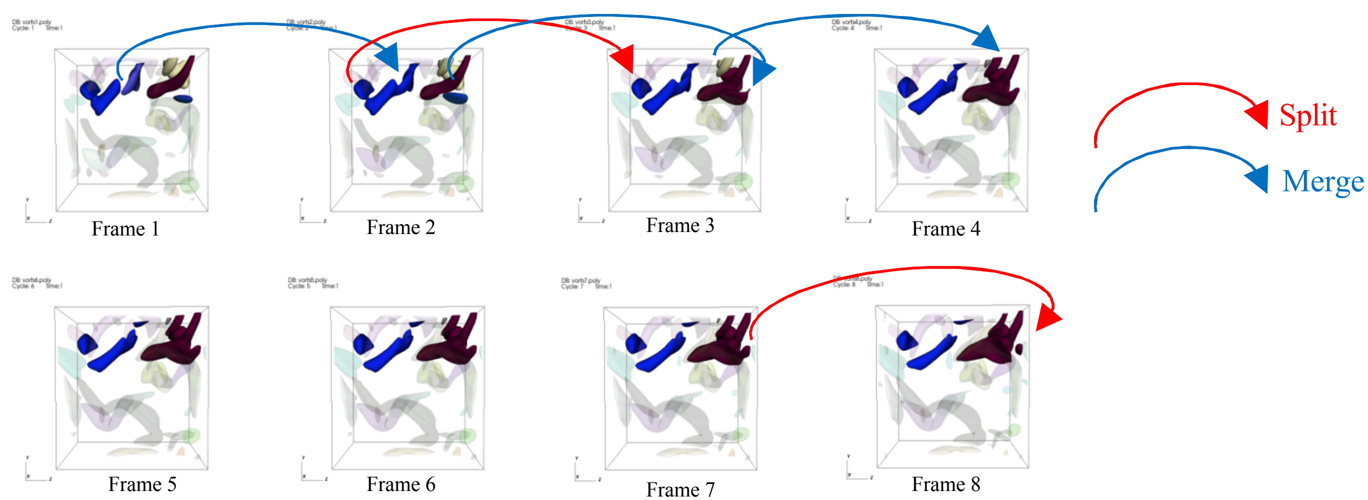

UltraVis
SciDAC Institute for Ultrascale Visualization

BIL - Block I/O Layer for Ultrascale Visualization and everything else!

- Problem
  - I/O is the generally agreed villain to performance
  - Critical to visualization codes
  - Yet few in visualization implement leading edge I/O practices in code
  - Gap: too much I/O decisions in vis codes are due to file formats, too little are based on partitioning strategy
- Solution
  - Many visualization codes manages data on the granularity of blocks, including scalar vs. vector, static time vs. time-varying, single vs. multi/many-variable, single vs. ensemble output
  - For this parallel visualization design pattern, we created a BIL library (already under public release) where IOR benchmark performance can be consistently achieved in application code
- Impact
  - Happy users!

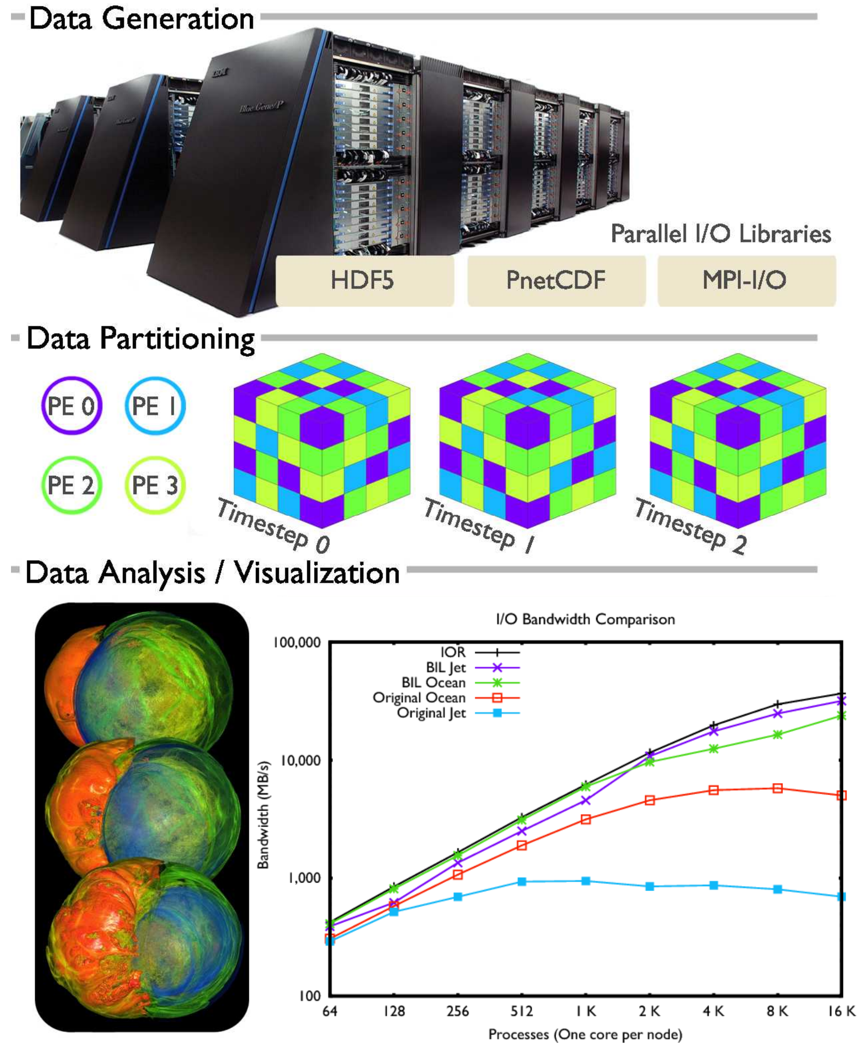

Parallel Particle Tracing for Large-scale Flow Fields

UltraVis
SciDAC Institute for Ultrascale Visualization

- Problem
  - Large-scale simulations (climate modeling, nuclear fusion, etc.) often attempt to model very complex fluid flow phenomena. Effective visualization of the resulting flow fields plays an important role in enabling validation of the simulations and scientific knowledge discovery.
  - To understanding of the complex flow fields, many visualization techniques based on computing a large number of particle tracing in space and time
  - Efficient parallel particle tracing is non-trivia due to irregular workload distribution and data access patterns.
- Solution
  - A scalable parallel flow visualization library is developed by IUSV researchers for DOE's Leadership Computing Facility.
  - The library contains optimized parallel load distribution, communication, I/O components with high quality rendering capability to support both postprocessing and in-situ visualization.
- Impact
  - The library exhibits strong scaling up to 16K processes.
  - It has been made freely downloadable.
  - IUSV is actively collaborating with scientists at Argonne National Lab (Nek5000 code) and climate scientists at PNNL to adopt this library.

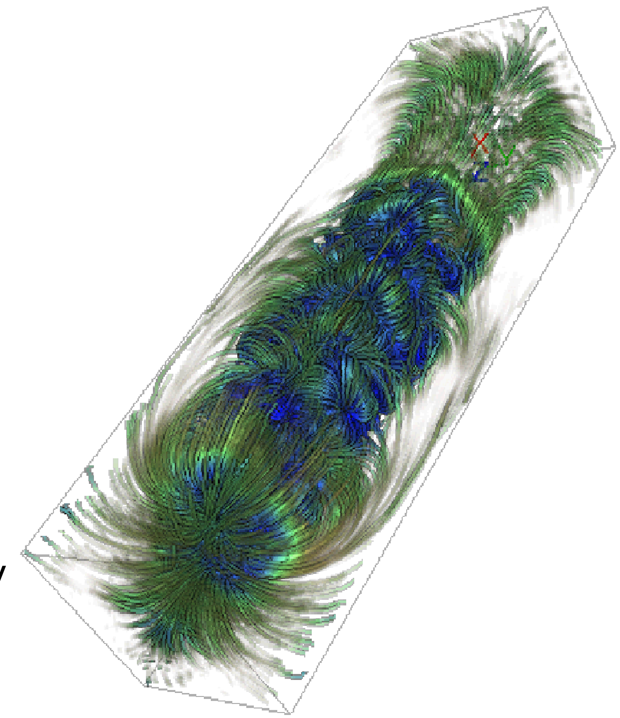

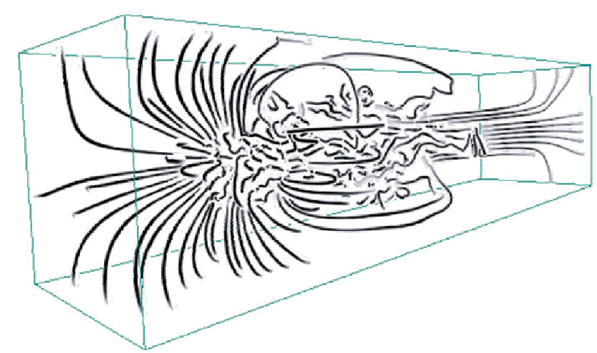

OSUFlow Library

- As a result of an NSF grant, OSU began to construct a flow visualization library since 2004
- In 2005, the core of OSUFlow Vis library was adopted by NCAR's visualization software VAPOR and released to the turbulence research community via SourceForge (more than 1000 downloads so far)
- In 2008, OSU and Argonne began to extend the library to run on DOE's leadership computing facility
- Currently more than 25,000 lines of code

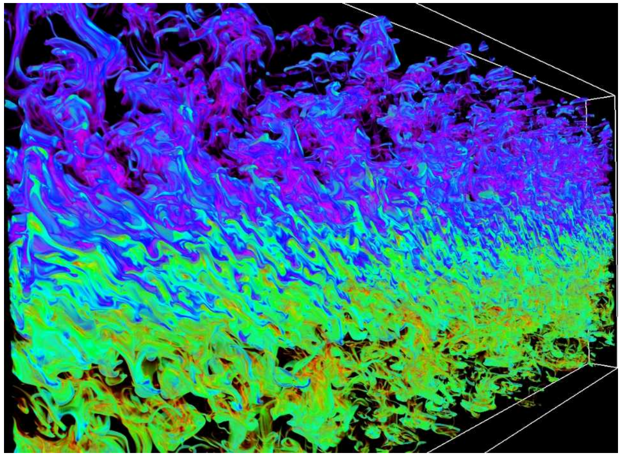

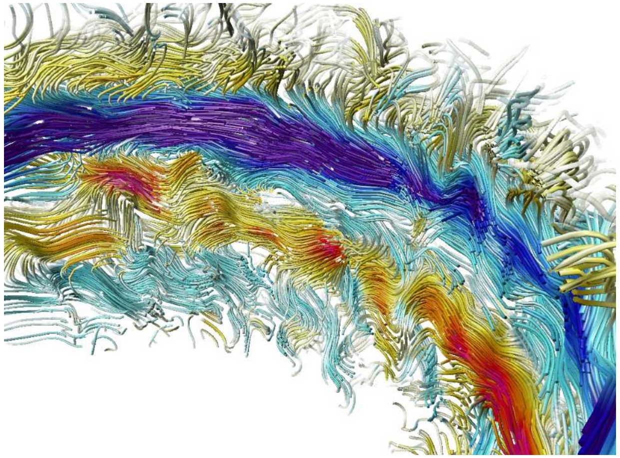

SciDAC Institute for Ultrascale Visualization

Large-Scale Parallel Particle Tracing

Benefits

- Large system size, up 32K cores
- Dense seeding, up to 128K particles
- Static and time-varying flows
- Downloadable library https://svn.mcs.anl.gov/repos/osuflow/trunk

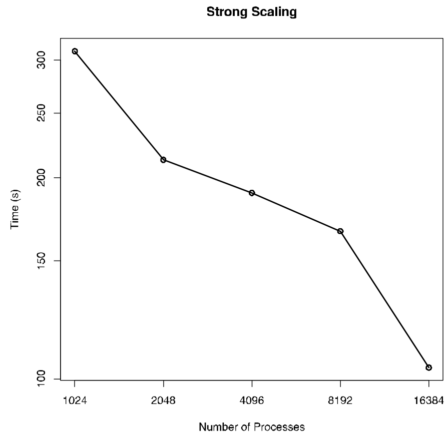

Strong scaling, 2048 x 2048 x 2048 data, 128K particles, 1 time-step

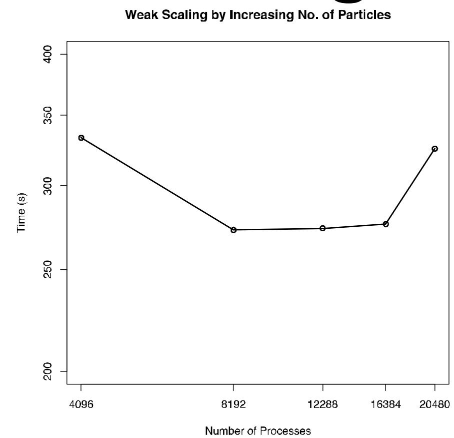

Weak scaling, 2304 x 4096 x 4096 data, 16K to 128K particles, 1 time-step

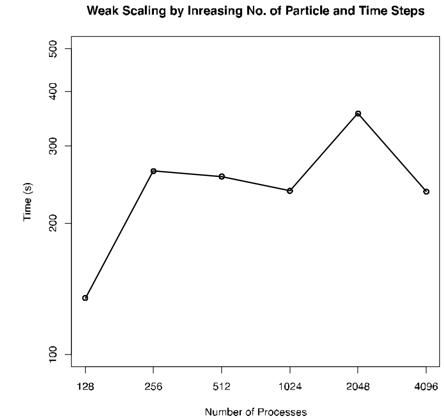

Weak scaling, 1408 x 1080 x 1100 data, 512 to 16K particles,1 to 32 time-steps

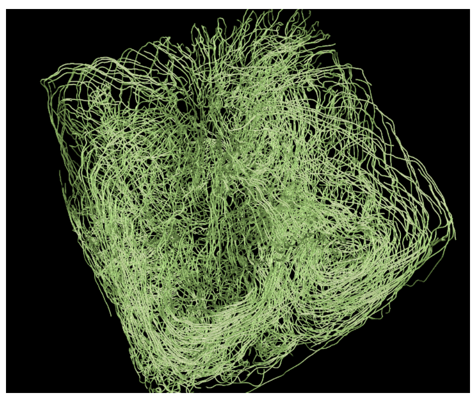

Thermal hydraulics data courtesy Aleks Obabko and Paul Fischer, ANL

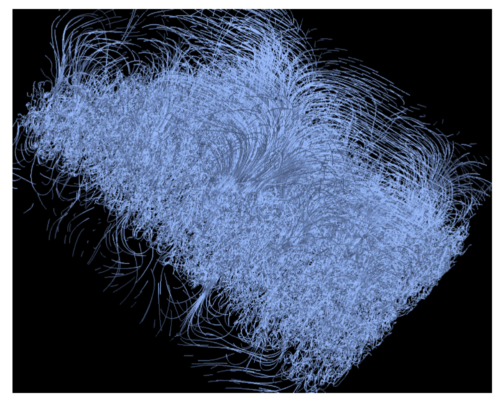

Rayleigh-Taylor instability data courtesy Mark Petersen and Daniel Livescu, LANL

Flame stabilization data courtesy Ray Grout, NREL and Jackie Chen, SNL

A Study of Parallel Particle Tracing for Steady-State and Time-Varying Flow Fields. Peterka et al., IPDPS’11

Building Blocks for Scalable Parallel Analysis

SciDAC Institute for Ultrascale Visualization

Goal:

- Achieve scalability through a common set of core data movement components
- Balance load (computation, communication)
- Minimize / optimize data movement (storage and network)
- Hide data movement (overlap with work)

Benefits:

- Researchers can study new algorithms
- Computer / computational scientists can build custom applications
- Reuse core components

Examples of components for communication:

- Global reduction (merging, compositing)
- Local nearest-neighbor exchange (particle tracing, ghost cell exchange, connected component labeling)

We are building a prototype library of these components.

Global Reduction

1. Round 1 exchange with k = 4, eg.
2. Round 2 exchange with k = 2, eg.
3. Repeat

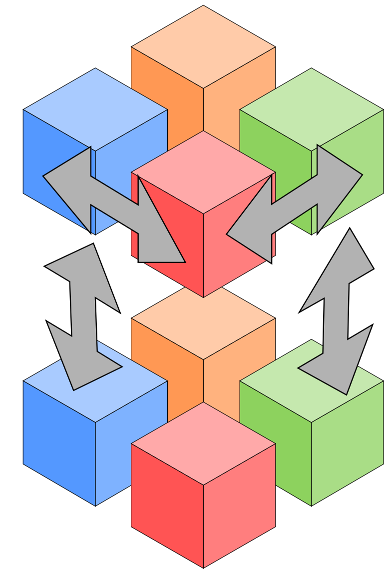

Local Nearest Neighbor Exchange

1. Perform local computations on blocks

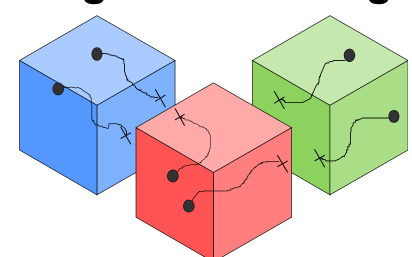

2. Exchange objects among neighbors

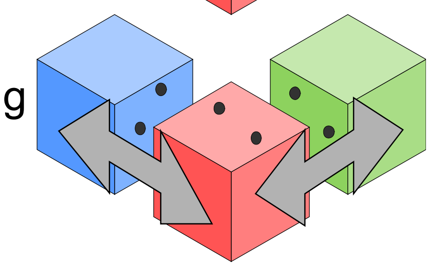

3. Repeat

Building Blocks for GPU Computation

- CUDPP ("CUDA Data Parallel Primitives"): open-source, popular library for GPU computing primitives
- DCGN ("Distributed CPU-GPU Network"): open-source MPI-like message-passing library for GPU clusters
- GPMR ("GPU MapReduce"): open-source high-level MapReduce library for GPU clusters
  - Implemented volume renderer within MapReduce framework (with K.-L. Ma, C.-K. Chen, Ultravis)

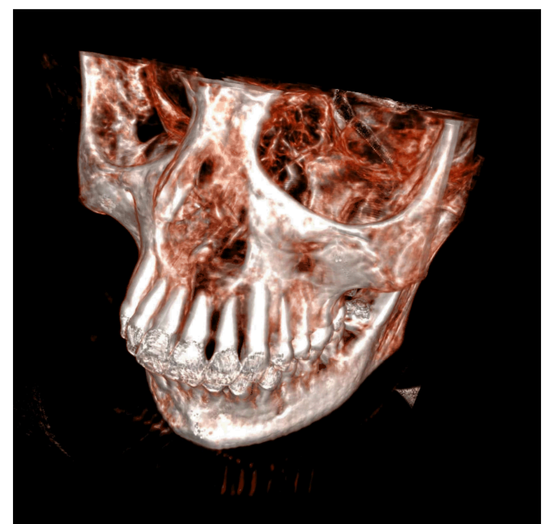

- Project lead: John Owens, UC Davis; students: Jeff Stuart, Shubhabrata Sengupta, Andrew Davidson, Ritesh Patel

UltraVis
SciDAC Institute for Ultrascale Visualization

Complex Feature Visualization by Programming Language Interface

- Problem
  - Today multi-scale/physics systems are modeled with unprecedented complexity
  - A daunting task to manually explore all aspects of such complex systems
  - Infeasible to always require users to fully substantiate a feature specification in all aspects
  - In result, traditional visualizations could be unreliable, inconclusive or incomprehensive
- Solution
  - Enabling versatile exploration and visualization of complex features by using supercomputers to automatically explore a feature specification's un-specified aspects.
  - Domain scientists control this automation via light weight programming languages (regex).
- Impact
  - Reliably capture and visualize ocean structures at fine precision on global scales
  - Useful to all climate modelers needing to understand or improve ocean models

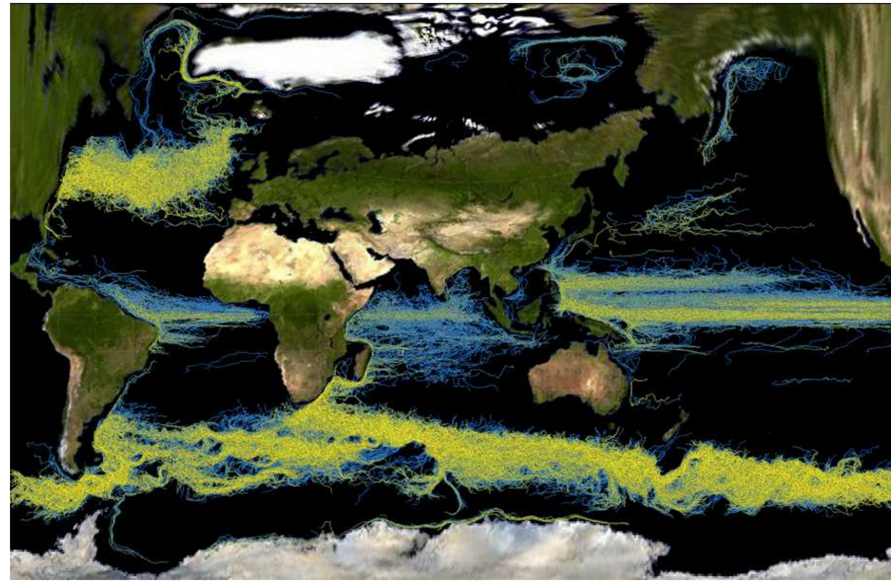

```
/* floating point numbers */
\-?[0-9]+
| \-?"."([0-9])+
| \-?([0-9])+"."([0-9])*
| \-?([0-9])+\.?([0-9])*[eE]\-?([0-9])+
        return VALUE;

\[ return RBEGIN;
\] return REND;
\? return RWILDCARD;
\* return MREXPANDER;
\- return RTO;
T return TIMEMARKER;
[ \t\n] ;
. ;
```

```
statement:
        | statement range
        | statement multirange
        | statement TIMEMARKER range
        | statement TIMEMARKER
        multirange
        ;
range: RBEGIN values REND
        | RWILDCARD
        ;
values: VALUE RTO VALUE
        | VALUE
        | RWILDCARD
        ;
multirange: range MREXPANDER
        ;
```

UltraVis
SciDAC Institute for Ultrascale Visualization

Visualization of Particle-based Groundwater Simulations

- Problem
  - Scientists at PNL have been developing parallel codes for multiscale particle-based simulations of subsurface biogeochemical processes. They simulate fluid flow, solute transport, and biogeochemical reactions in porous media at the pore scale for fundamental groundwater research.
  - Advanced visualization techniques are needed to understand the essential character of pore-scale dispersion processes as well as the the effects of localized zones of relatively low velocity within the pore-scale flow field, and their effect on transport phenomena at larger scales.
- Solution
  - To enable visualization of specific particles of interest in context and better understanding of particle trajectory behavior, IUSV researchers have invented a hybrid rendering technique based on advanced volume and surface construction methods.
- Impact

  The new visualization capability has helped scientists at PNL to better
  - Conduct visual checking of output for physical consistency
  - Synthesize large datasets in a visual manner to gain understanding of key processes
  - Communicate the results of complex simulations to scientific colleagues and non-scientific audiences

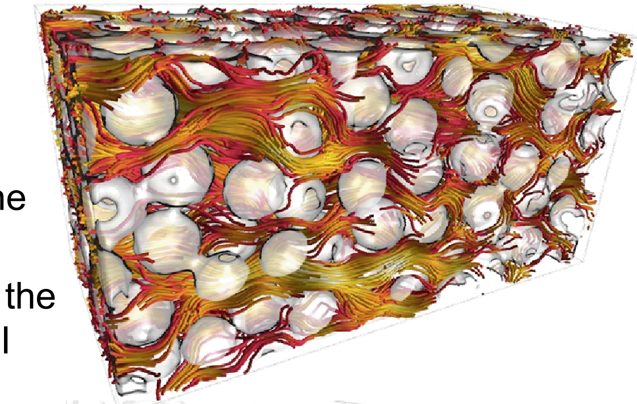

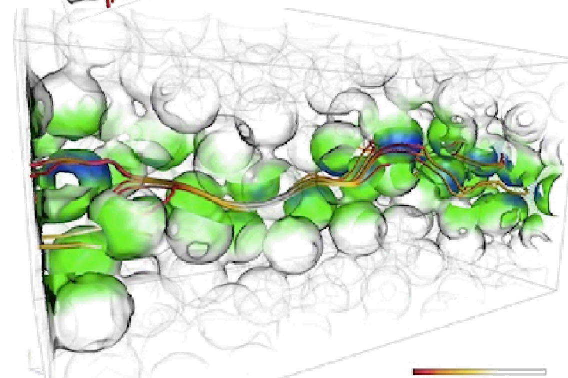

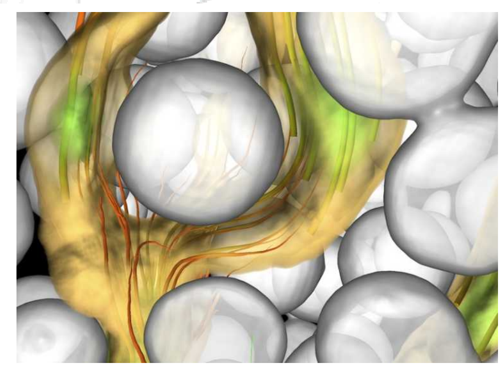

Simulation and Visualization of Potential Multipacting

- Problem
  - Multipacting is a phenomenon of resonant electron multiplication exponentially in the particle accelerators, leading to an electron avalanche.
  - Multipacting can cause significant power losses and heating of the wall, low achievable field gradient and thermal breakdown in superconducting structures.
- Solution
  - Using DOE's supercomputing facilities, ComPASS accelerator simulation codes and UltraVis Institute's visual analysis tools, multipacting behavior can be simulated, visualized, and analyzed.
- Impact
  - By comparing simulation results to damage in existing accelerator components, SLAC validated the simulation codes.
  - Analysis of new cavity designs has enabled that costly multipacting will not occur.
  - This important step was made possible by a multidisciplinary collaboration between scientists at SLAC and the Institute for UltraScale Visualization.

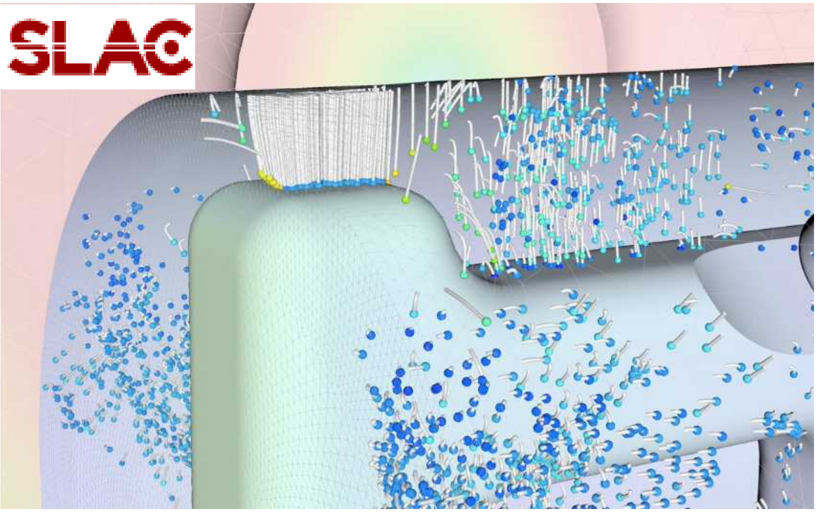

UltraVis
SciDAC Institute for Ultrascale Visualization

Full Range Analysis of Ultrascale Climate Data Sets

- As high quality data sets become available through NASA MODIS satellite programs and IPCC simulation ensembles, the critical link is to develop extreme scale visualization methods and systems that allow discoveries of multi-scale, multi-physics long-term trends in climate change.

- Computer scientists at the SciDAC Institute for Ultrascale Visualization have developed novel feature characterization and visualization methods that are scalable to tens of thousands of processors, promising a solution to the upcoming exascale data analysis problem.

- These new methods enable ORNL climate scientists with the SciDAC CCSM Consortium, DOE C-LAMP and International C^4MIP projects to analyze the full extent of their data.

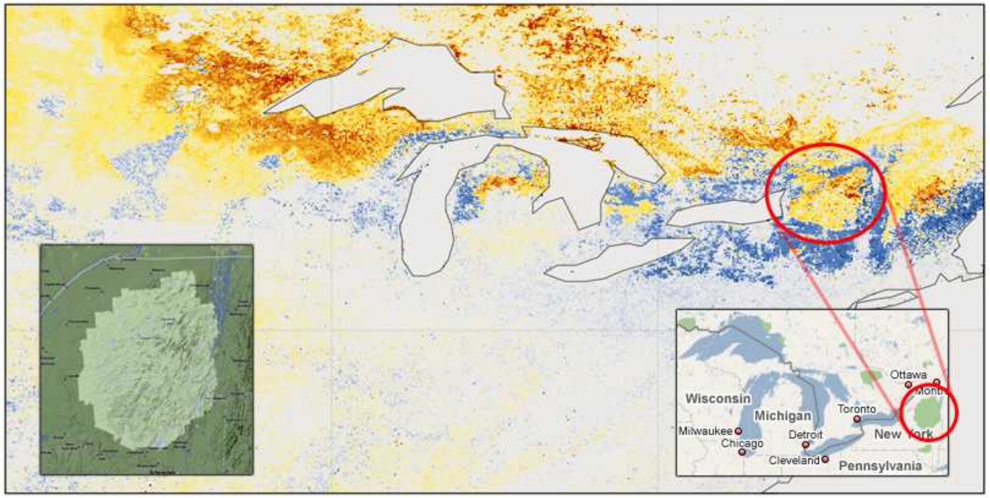

Impact:

This is the first time climate scientists can study a full terabyte of climate-related satellite monitoring data (NASA MODIS) and IPCC simulation data sets, and have discovered, in one single analysis, changes in patterns of phenology in undistributed vegetated ecosystems suggestive of climate change.

– Forrest Hoffman, ORNL

UltraVis
SciDAC Institute for Ultrascale Visualization

Visualization of the Structures and Dynamics of Turbulent Eddies

- Scientists at Sandia National Laboratories utilize DOE's supercomputing facilities to simulate combustion processes at the finest scales where chemical reactions occur, essential to understanding the process of clean and efficient combustion of alternative fuels.
- Computer scientists at the SciDAC Institute for Ultrascale Visualization have developed a novel method for the visualization and understanding of complex multi-scale features in the data output by such simulations.
- Impact

  This data reduction and visualization method enables us to zoom in and **see for the first time** the interaction of small turbulent eddies with the preheat layer of a turbulent flame, a region that was previously obscured by the multi-scale nature of turbulence.

  Jackie Chen, SNL-CA