U.S. DEPARTMENT OF ENERGY | Office of Science

SAND2017-2554PE

XVis: Visualization for the Extreme-Scale Scientific-Computation Ecosystem

ASCR CS PI Meeting

Kenneth Moreland Sandia National Laboratories

Berk Geveci Kitware, Inc.

David Pugmire Oak Ridge National Laboratory

David Rogers Los Alamos National Laboratory

Kwan-Liu Ma University of California at Davis

Hank Childs University of Oregon

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. SAND2017-XXXX O

XVis: Visualization for the Extreme-Scale Scientific-Computation Ecosystem

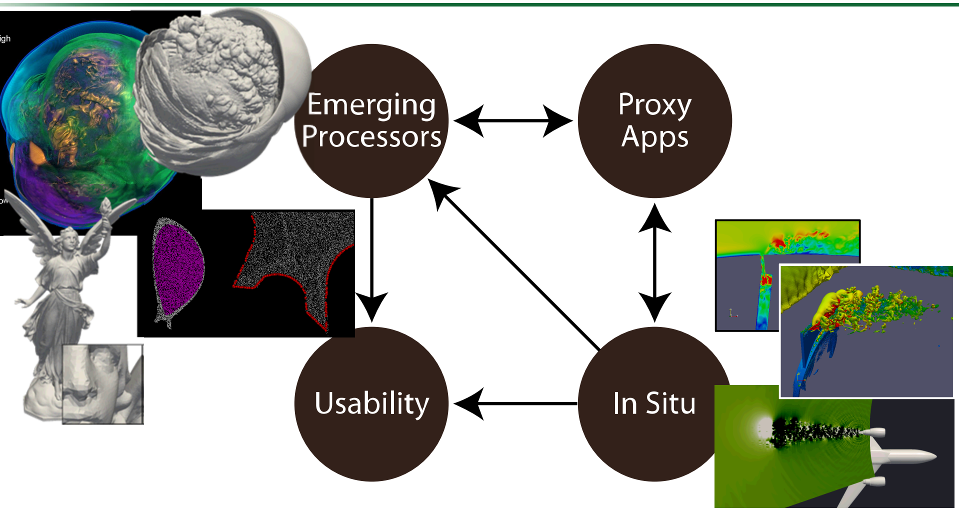

Novel Ideas

Enable visualization at extreme scale by addressing four interlocking challenges

- Leverage **Emerging Processor Technology** with the VTK-m toolkit for many-core visualization
- Improve ***In Situ* Integration** with lighter weight toolkits and better post hoc interaction
- Understand the **Usability** of new techniques with qualitative user studies
- Use **Proxy Analysis** to understand visualization behavior at extreme scale alone and coupled

Impact and Champions

The XVis project provides visualization research necessary for extreme-scale science in the DOE. Our studies on algorithms for emerging processor technology will update existing software like ParaView and VisIt to future architectures. The associated proxy analysis will guide our software on the most effective use of each platform and enable large-scale studies across numerous integrated software tools. Our *in situ* integration will make our software more accessible to simulation and the post hoc analysis more complete, as verified through user studies.

Principal Investigator(s): Kenneth Moreland (SNL), Berk Geveci (Kitware), David Pugmire (ORNL), David Rogers (LANL), Kwan-Liu Ma (UC Davis), Hank Childs (U Oregon)

Milestones/Dates/Status

	Scheduled	Actual
1.c Hybrid Parallel	FY16Q4	FY17Q1
1.d Additional Algorithms	FY17Q4	ongoing
1.e Function Characterization	FY17Q4	ongoing
2.c Flyweight In Situ	FY16Q4	ongoing
2.d Data Model Application	FY16Q3	FY16Q3
2.e Memory Hierarchy Streaming	FY16Q4	FY16Q4
2.f Interface for Post Hoc Interaction	FY16Q4	ongoing
3.c/3.d Start/Continue Usability Studies	FY17Q4	ongoing
3.e Apply Usability Studies	FY17Q4	ongoing
4.a/4.b Mini-App Impl/Characterization	FY16Q4	ongoing
4.c Architectural Studies	FY17Q4	ongoing

U.S. DEPARTMENT OF ENERGY | Office of Science

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. SAND2017-XXXX O

March, 2017

Scientific Visualization and Analysis Ecosystem for Large Data

(VTK, ParaView*, VisIt*, Cinema, SENSEI, and VTK-m*)

* ECP funded

- Problem
  - DOE simulation scientists generate petabytes of data that they need to understand
- Solution
  - Developed general-purpose, scientific visualization and analysis libraries and tools
    - designed with parallelism in mind to operate on world's largest data sets
    - collaborative open-source development model engaging multiple National Laboratories, universities, and industry
- Impact
  - Leading visualization and analysis solutions for Department of Energy scientists
  - Used on all ASCR supercomputing facilities, and worldwide inmost HPC facilities
  - Over 1 million downloads of software worldwide
  - Creates capability for DOE to look at data from the world's largest simulations
  - Steve Langer (LLNL) – "We rely on this software for gaining an understanding of complex and spatial variations in backscattered light simulations using pF3D, where data may consist of over 400 billion zones per time step."

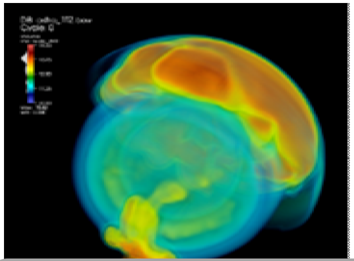

Core collapse supernova from GeneASIS.

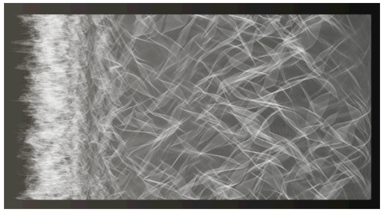

Novel visualization of ocean flow data, designed to show 'mixing barrier' (left)

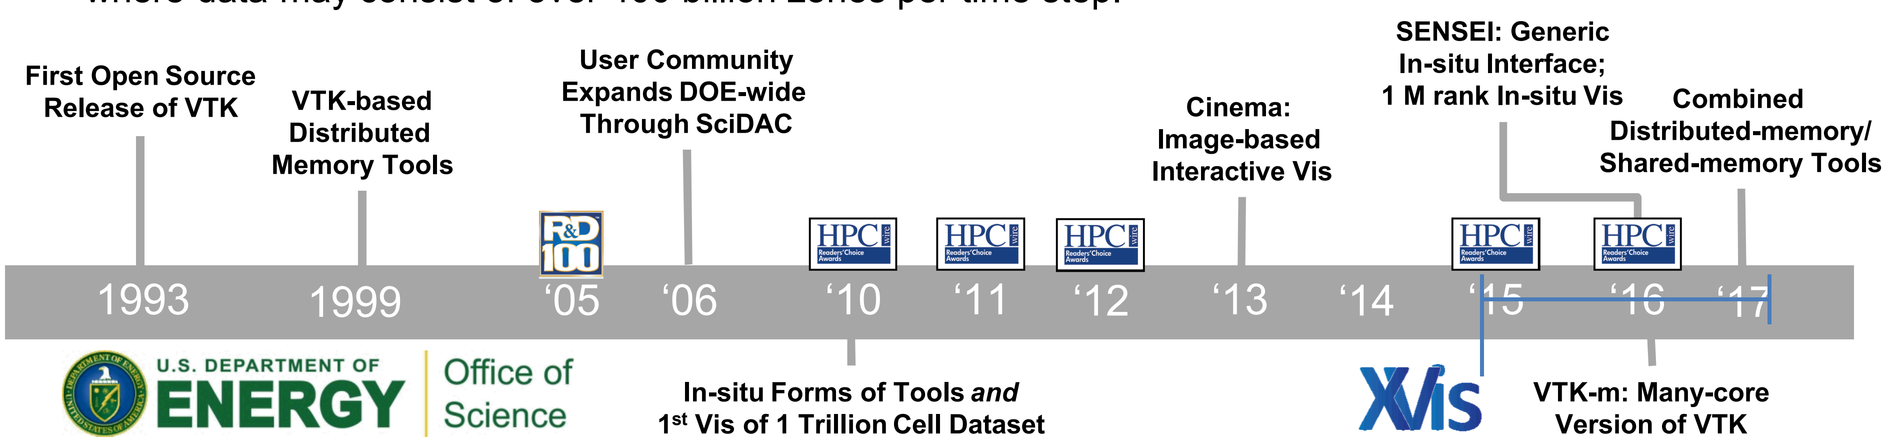

U.S. DEPARTMENT OF ENERGY | Office of Science

Enabling DOE Visualization Tools with VTK-m

Sandia National Laboratories
Kitware
Oak Ridge National Laboratory
Los Alamos National Laboratory EST. 1943

- Objective
  - DOE's production visualization tools, representing 100's of man-years of investment, are not ready for exascale architectures.
  - Integrating VTK-m into our existing visualization tools provides functionality at exascale while leveraging previous investment.
- Technology
  - ParaView & VisIt: Interactive HPC visualization software leverages VTK and MPI for large scale jobs.
  - VTK: Toolkit for 3D computer graphics, image processing, and visualization. Uses MPI for parallelization.
  - VTK-m: Toolkit for the development and distribution of scientific visualization algorithms on multicore and accelerators.
- Impact
  - ParaView 5.3 and VisIt 2.11 offer VTK-m acceleration.
  - Upcoming VTK 8.0 will offer VTK-m accelerated filters.
  - In situ tools demonstrated on Titan with thousands of GPUs.

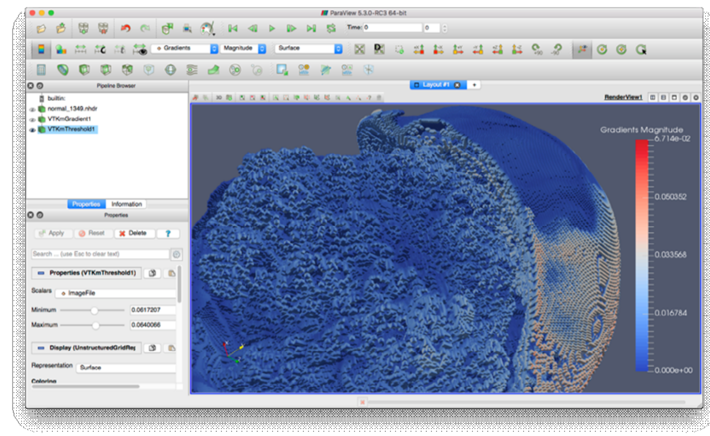

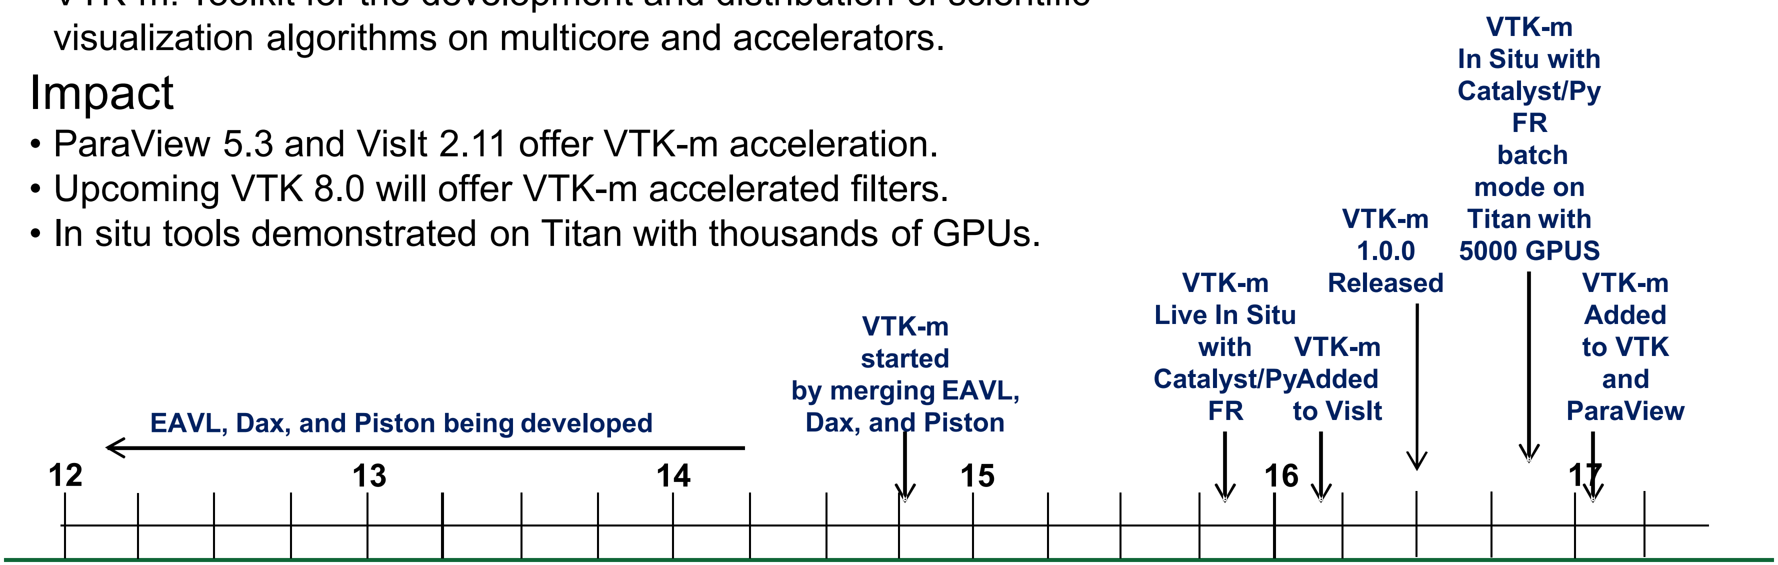

U.S. DEPARTMENT OF ENERGY | Office of Science

Live Demonstration of In Situ Visualization on Accelerator Processors

Sandia National Laboratories
Kitware
Oak Ridge
Los Alamos National Laboratory

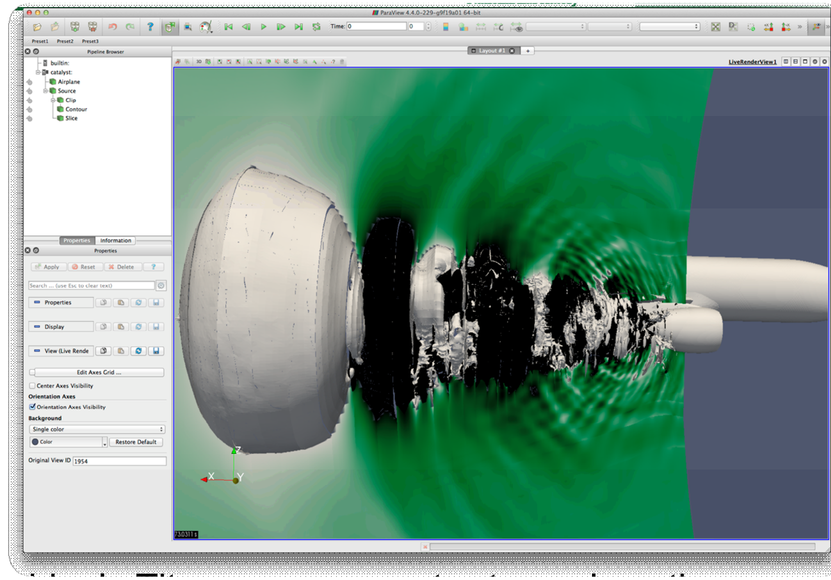

Objective

- Scientific discovery using 100 petascale to exascale supercomputers requires integrated visualization solutions on new computational hardware.
- Multiple technologies to be implemented and integrated.

Technology

- ParaView: existing HPC visualization infrastructure leveraging VTK over large MPI jobs with user interaction.
- Catalyst: in situ coupling of ParaView to simulations.
- VTK-m: toolkit for the development and distribution of visualization algorithms on multicore and accelerators.
- PyFR: CFD simulation running on 256 GPU devices of Oak Ridge's Titan supercomputer to analyze the turbulence behind a new serrated jet engine nozzle.

Impact

- A demonstration at SC 2015 shows the integration of these 4 key technologies for a live analysis of the simulation.
  - Top right: The pockets of air in the jet wake reduce noise.
  - Direct right: A live interactive visualization of a large simulation running on Titan viewed and controlled on the show floor in Austin.
- The entire process from simulation to analysis to rendering happens locally on the GPUs without memory transfer.
  - As increasing limitations in network, power, and storage limit the movement of data, sharing resources is critical for effective analysis.

U.S. DEPARTMENT OF ENERGY | Office of Science

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. SAND 2016-0727 PE

NNSA National Nuclear Security Administration

Flyweight In Situ Analysis

- **Objective**
  - In situ analysis though tightly-coupled simulation and analysis codes.
  - Accelerator enabled scalable analysis and visualization though lightweight infrastructure
- **Technology**
  - VTK-m : toolkit for the development and distribution of visualization algorithms on multicore and accelerators.
  - Sensei : write-once, deploy-many in situ infrastructure with integration of Catalyst, Libsim, ADIOS and Glean.
  - Strawman (Alpine) : lightweight in situ visualization infrastructure for multi-physics HPC simulations.
- **Impact**
  - Initial demonstration of VTK-m integration into emerging in situ infrastructures.
  - Flyweight in situ analysis with minimal dependencies.
  - Upcoming scalability study with multiple mini-apps.

Usability of In Situ Generated PDFs for Post Hoc Analysis

Application and Background

- Large-scale scientific simulations produce too much data for effective storage and analysis
- Researchers need a way to quickly query large data subsets, sometimes using trial and error, to explore scientific phenomena
- Many methods currently in use by scientific teams inefficiently operate on the full data
- A study is being done with Dr. Jackie Chen, a combustion scientist at Sandia National Labs

Goals and Challenges

- Features of interest are not known *a priori* and must be explored during post hoc analysis
- Need to capture simulation trends in situ and store/represent them in an efficient format
- Must have little impact on the simulation and fit smoothly into the scientists' existing workflow
- An effective set of tools must be evaluated using a thorough usability study to determine their effectiveness and improve their functionality

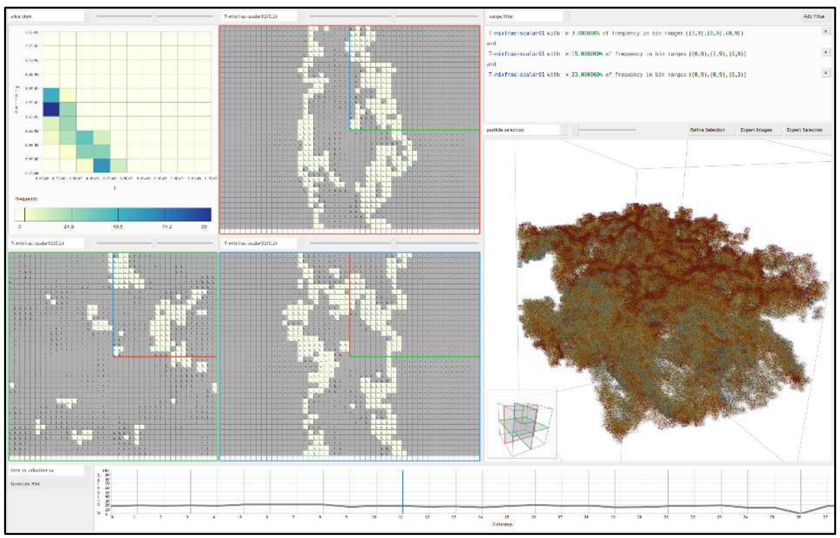

Results and Impact

- A set of in situ libraries with a post hoc visualization tool can capture trends efficiently using probability distribution functions (PDFs)
- Large-scale particle data is coupled with PDFs in situ to enable fast subset selection and analysis
- Extensive performance tests ensure insignificant simulation and storage overheads
- An ongoing usability study with expert users will evaluate both the in situ and post hoc tools

U.S. DEPARTMENT OF ENERGY | Office of Science

Kwan-Liu Ma

UC DAVIS UNIVERSITY OF CALIFORNIA

Parallel Peak Pruning: Scalable SMP Contour Tree Computation

Hamish A. Carr (University of Leeds), **Gunther H. Weber** (LBNL/Vis Base),
Christopher M. Sewell (LANL/XVis), James P. Ahrens (LANL/XVis)

Background and Problem

- *Contour trees* summarize topological properties of isosurfaces, making them a valuable tool for data analysis, such as identifying
  - Most "important" isosurfaces based on prominence; and
  - Features in wide range of applications, including burning regions (combustion simulations), halos (computational cosmology), atoms and bonds (computational chemistry), pores and pockets (materials).
- Optimal serial "sweep-and-merge" algorithm for computing contour trees based on inherently sequential metaphor, requires data processing in sorted order.

Objective

- Effective utilization of multi-core architectures (per compute-node resources), such multi-core CPUs, GPUs
- Need efficient data-parallel contour tree computation algorithm.

Approach

- New "parallel peak pruning" method finds peaks and their "governing saddles," creating superarcs in the join/split tree by recursive pruning of peak/saddle pairs and their corresponding regions.
- Can prune many peak/saddle pairs concurrently, effectively using data parallel resources.
- Algorithm has formal guarantees of O(log(n) log(t)) parallel steps and O($n$log(n)) work–i.e., number of primitive operations–for grid with n vertices and t critical points, i.e., depends data on size *and* complexity.
- Native OpenMP and Nvidia Thrust implementations enable effective data-parallel contour computation on multi-core CPUs and GPUs.
- Recently finalized VTK-m port (https://gitlab.kitware.com/vtk/vtk-m) makes new algorithm available to DOE and HPC stakeholders.

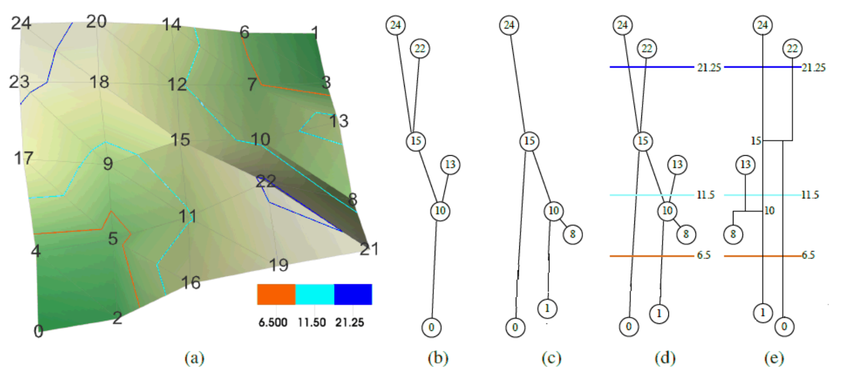

(a) Terrain and select contour lines. (b) Join and (c) split tree record where maxima and minima respectively "meet." (d) The contour tree combines join and spit tree into a representation of the full connectivity of contour lines. (e) The branch decomposition orders features of the contour tree hierarchically based on a simplification measure, such as prominence (persistence).

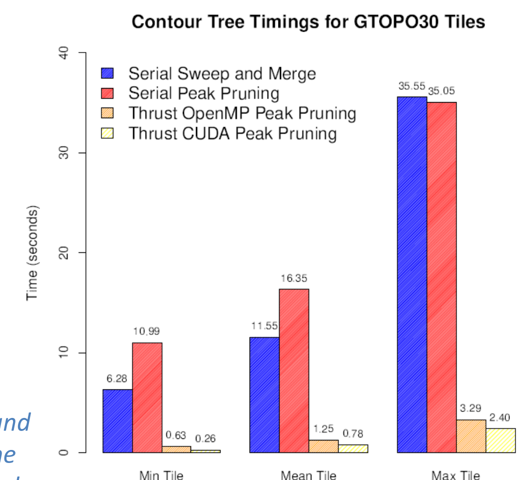

Results and Impact

- **Significant parallel performance improvement:** Up to **10x** parallel speed-up on multi-core CPUs and up to **50x** speed up on Nvidia GPUs, compared to optimal serial sweep-and-merge algorithm. (Tested on USGS GTOPO30 elevation maps.)
- Important building block for analysis on future **exascale platforms**.
- **Publication** "Parallel Peak Pruning for Scalable SMP Contour Tree Computation" by H. Carr, G. Weber, C. Sewell, J. Ahrens won **best paper award** at IEEE LDAV 2016.

U.S. DEPARTMENT OF ENERGY | Office of Science

UNIVERSITY OF LEEDS · BERKELEY LAB · Los Alamos NATIONAL LABORATORY EST. 1943

VTK-m Support for ECP Projects

VTK-m is being used for visualization and analysis in the "Coupled Monte Carlo Neutronics and Fluid Flow Simulation of Small Modular Reactors" ECP project.

Direct visualization support for Constructive Solid Geometry (CSG) allows scientists to perform more accurate analysis

- **Conversion to mesh-based representations not needed**
  - Expensive (1000s of rods!) and introduces errors
- **In situ analysis and visualization possible without the need to copy and convert data**

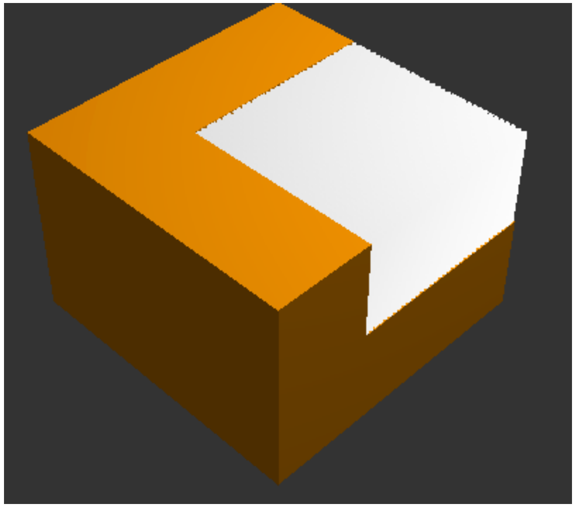

A reactor pin assembly rendered directly from the native CSG representation.

Detail view of CSG representation of a single rod in the pin assembly.

Wavelet Compression in VTK-m

Background

- **Supercomputing trends have the ability to generate data going up faster than the ability to store data, which will deter visualization.**
- **One solution is to transform and reduce simulation data to a small enough form that post hoc analysis is possible.**
- **Wavelets are an excellent technology for this reduction, and are used widely in other fields**

Technique

- **Implemented wavelet compression in VTK-m**
  - Enables portable performance over many architectures
  - Also soon to be available in many end-user products, due to VTK-m being adopted (VisIt, ParaView, ADIOS)
- ***Impact*: wavelet compression "is now available" to simulation codes.**
  - Level of data reduction can be controlled by user

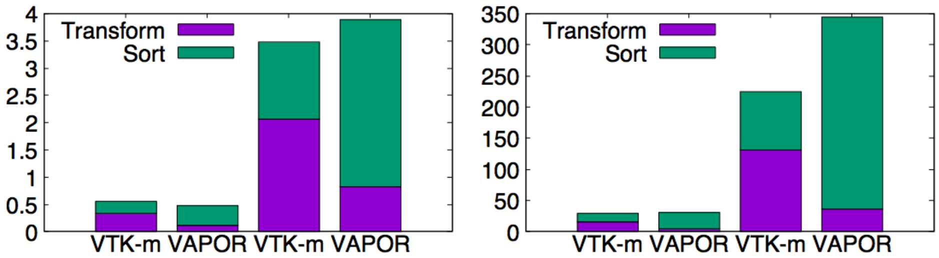

(a) *Left two bars:* 256^3. *Right two bars:* 512^3.

(b) *Left two bars:* $1{,}024^3$. *Right two bars:* $2{,}048^3$.

Performance study comparing our hardware-agnostic approach with a hardware-specific solution.

Progress

- **Study demonstrates:**
  - **VTK-m approach is comparable to reference CPU implementation**
  - **Also performs well on GPU**
- **"Achieving Portable Performance For Wavelet Compression Using VTK-m" by Li, Sewell, Clyne, and Childs. In submission to EuroGraphics Symposium on Parallel Graphics and Visualization (EGPGV) 2017.**

External Facelist Calculation in VTK-m

Background

- External facelist calculation is an essential algorithm for visualization
- Algorithm produces renderable surfaces from three-dimensional volumes.
- Needed in conjunction with popular algorithms like clipping, material interface reconstruction, and interval volumes.

Technique

- Implemented two variants of algorithm in VTK-m – variants based on hashing and sorting, respectively
- Both required advanced usage of VTK-m features.
- ***Impact***: by implementing the algorithm in VTK-m, it is now available on multiple architectures.
  - This is the first-ever many-core implementation of this algorithm.

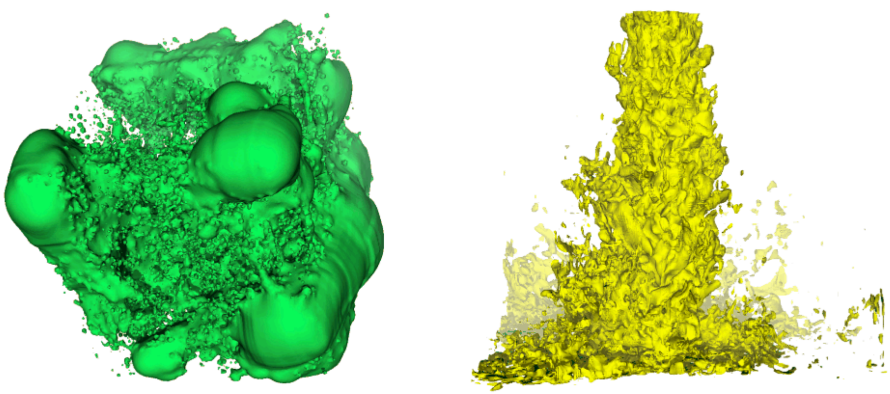

Visualization of two data sets used in study

Progress

- **Study demonstrates:**
  - **serial performance comparable to existing serial implementation**
  - **good parallel performance**
  - **hashing-based variant is fastest technique**
- **Algorithm contributed back to VTK-m repo**
- **"External Facelist Calculation with Data-Parallel Primitives" by Lessley, Binyahib, Maynard, and Childs, EuroGraphics Symposium on Parallel Graphics and Visualization (EGPGV), 2016.**

Rendering in VTK-m

Background

- **Rendering is an essential operation for visualization**
- **Many types of rendering:**
  - rasterization, ray-tracing, volume rendering
- **Also key for CINEMA-style data compression, which addresses I/O gap by saving many, many images instead of simulation data**

Progress

- **Implemented rendering techniques in VTK-m**
- **Studies show excellent performance on GPU, CPU, and Xeon Phi**
- **Comparison with industry standards (Intel OSPRay, NVIDIA OptiX) show our hardware-agnostic code has comparable performance**
- **Code available in VTK-m branch**

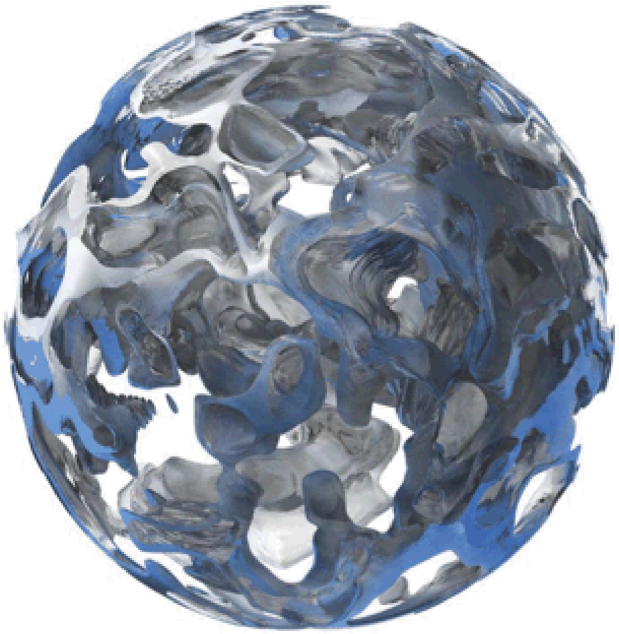

Ray-traced rendering in VTK-m.

Impact

- **For VTK-m to be successful, it is critical that it has performant, reliable rendering infrastructure**
- **With this effort, we now have "future proofed" rendering and can avoid external dependencies with respect to rendering**
  - **This will improve in situ integration, due to smaller binaries and less code complexity**
- **Publications:**
  - **"VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Architectures", Moreland et al. CG&A**