

Visualizing your Data

Joseph Insley
Lead, Visualization & Data Analytics
Argonne Leadership Computing Facility

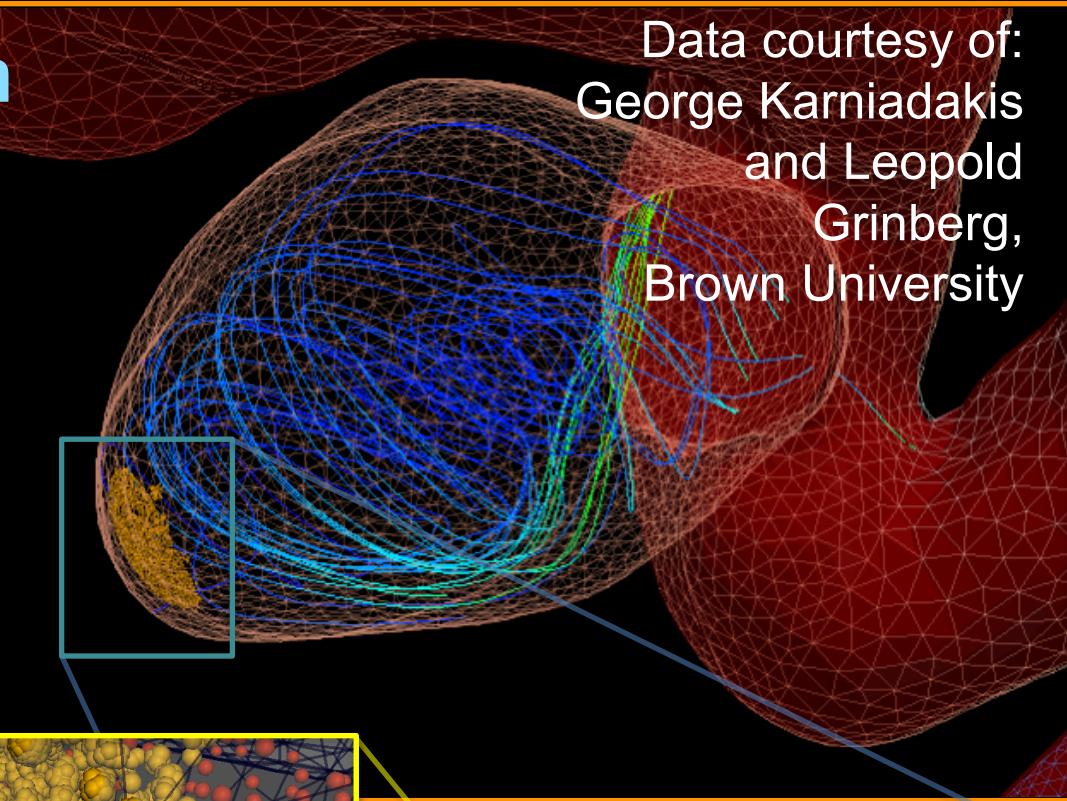
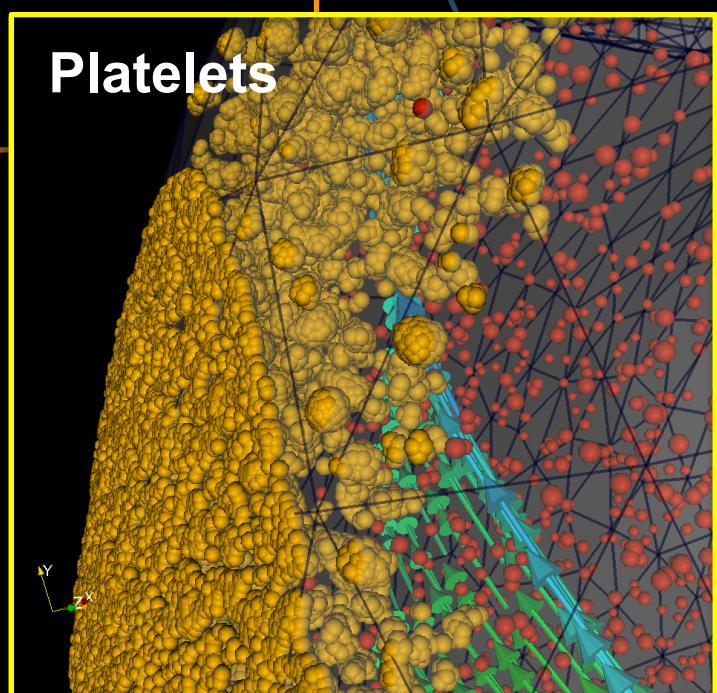
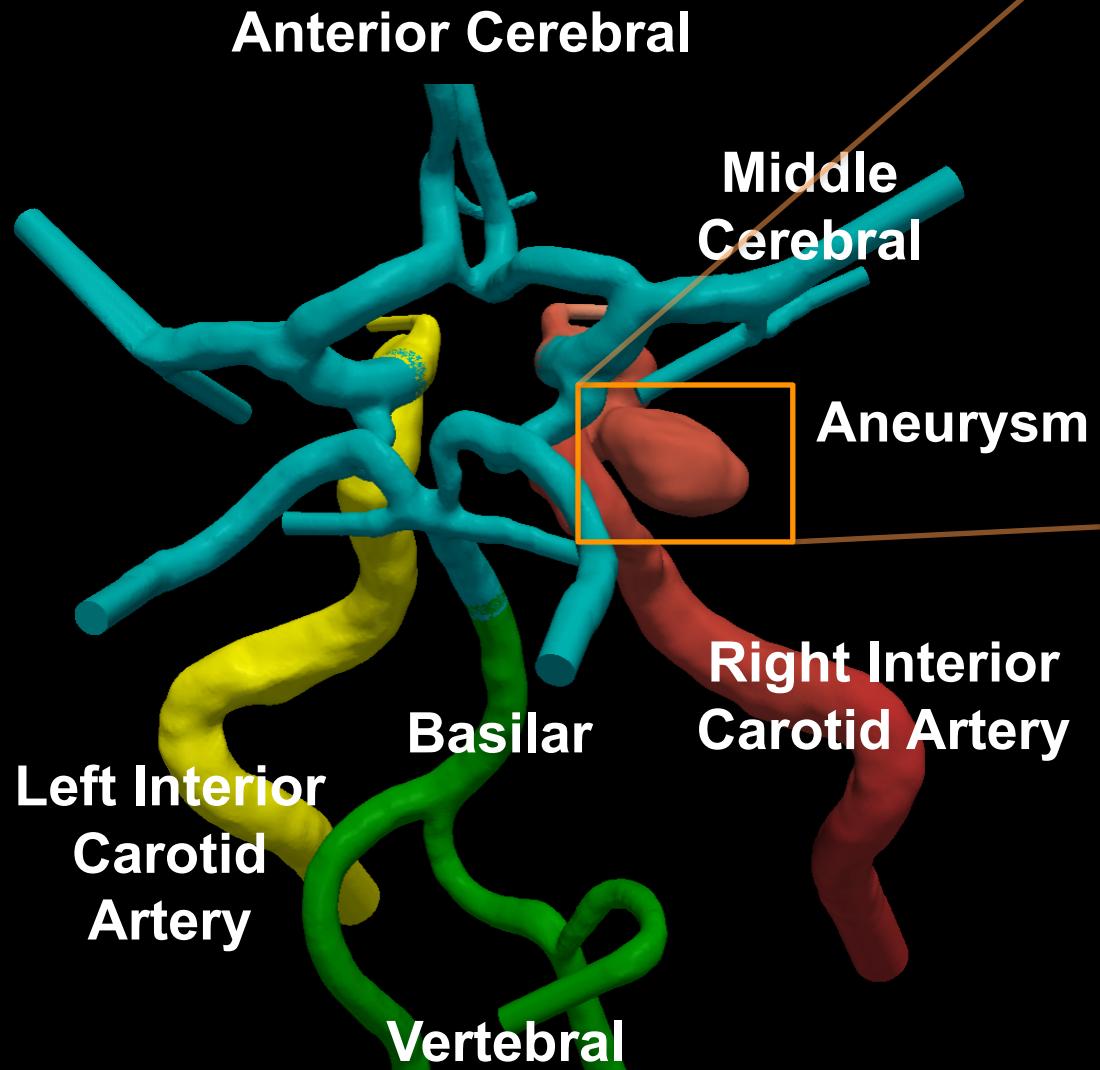
Silvio Rizzi
Assistant Computer Scientist
Argonne Leadership Computing Facility

Here's the plan...

- Examples of visualizations
- Visualization resources
- Visualization tools and formats
- Data representations
- Visualization for debugging
- In-Situ Visualization and Analysis

Multi-Scale Simulation / Visualization

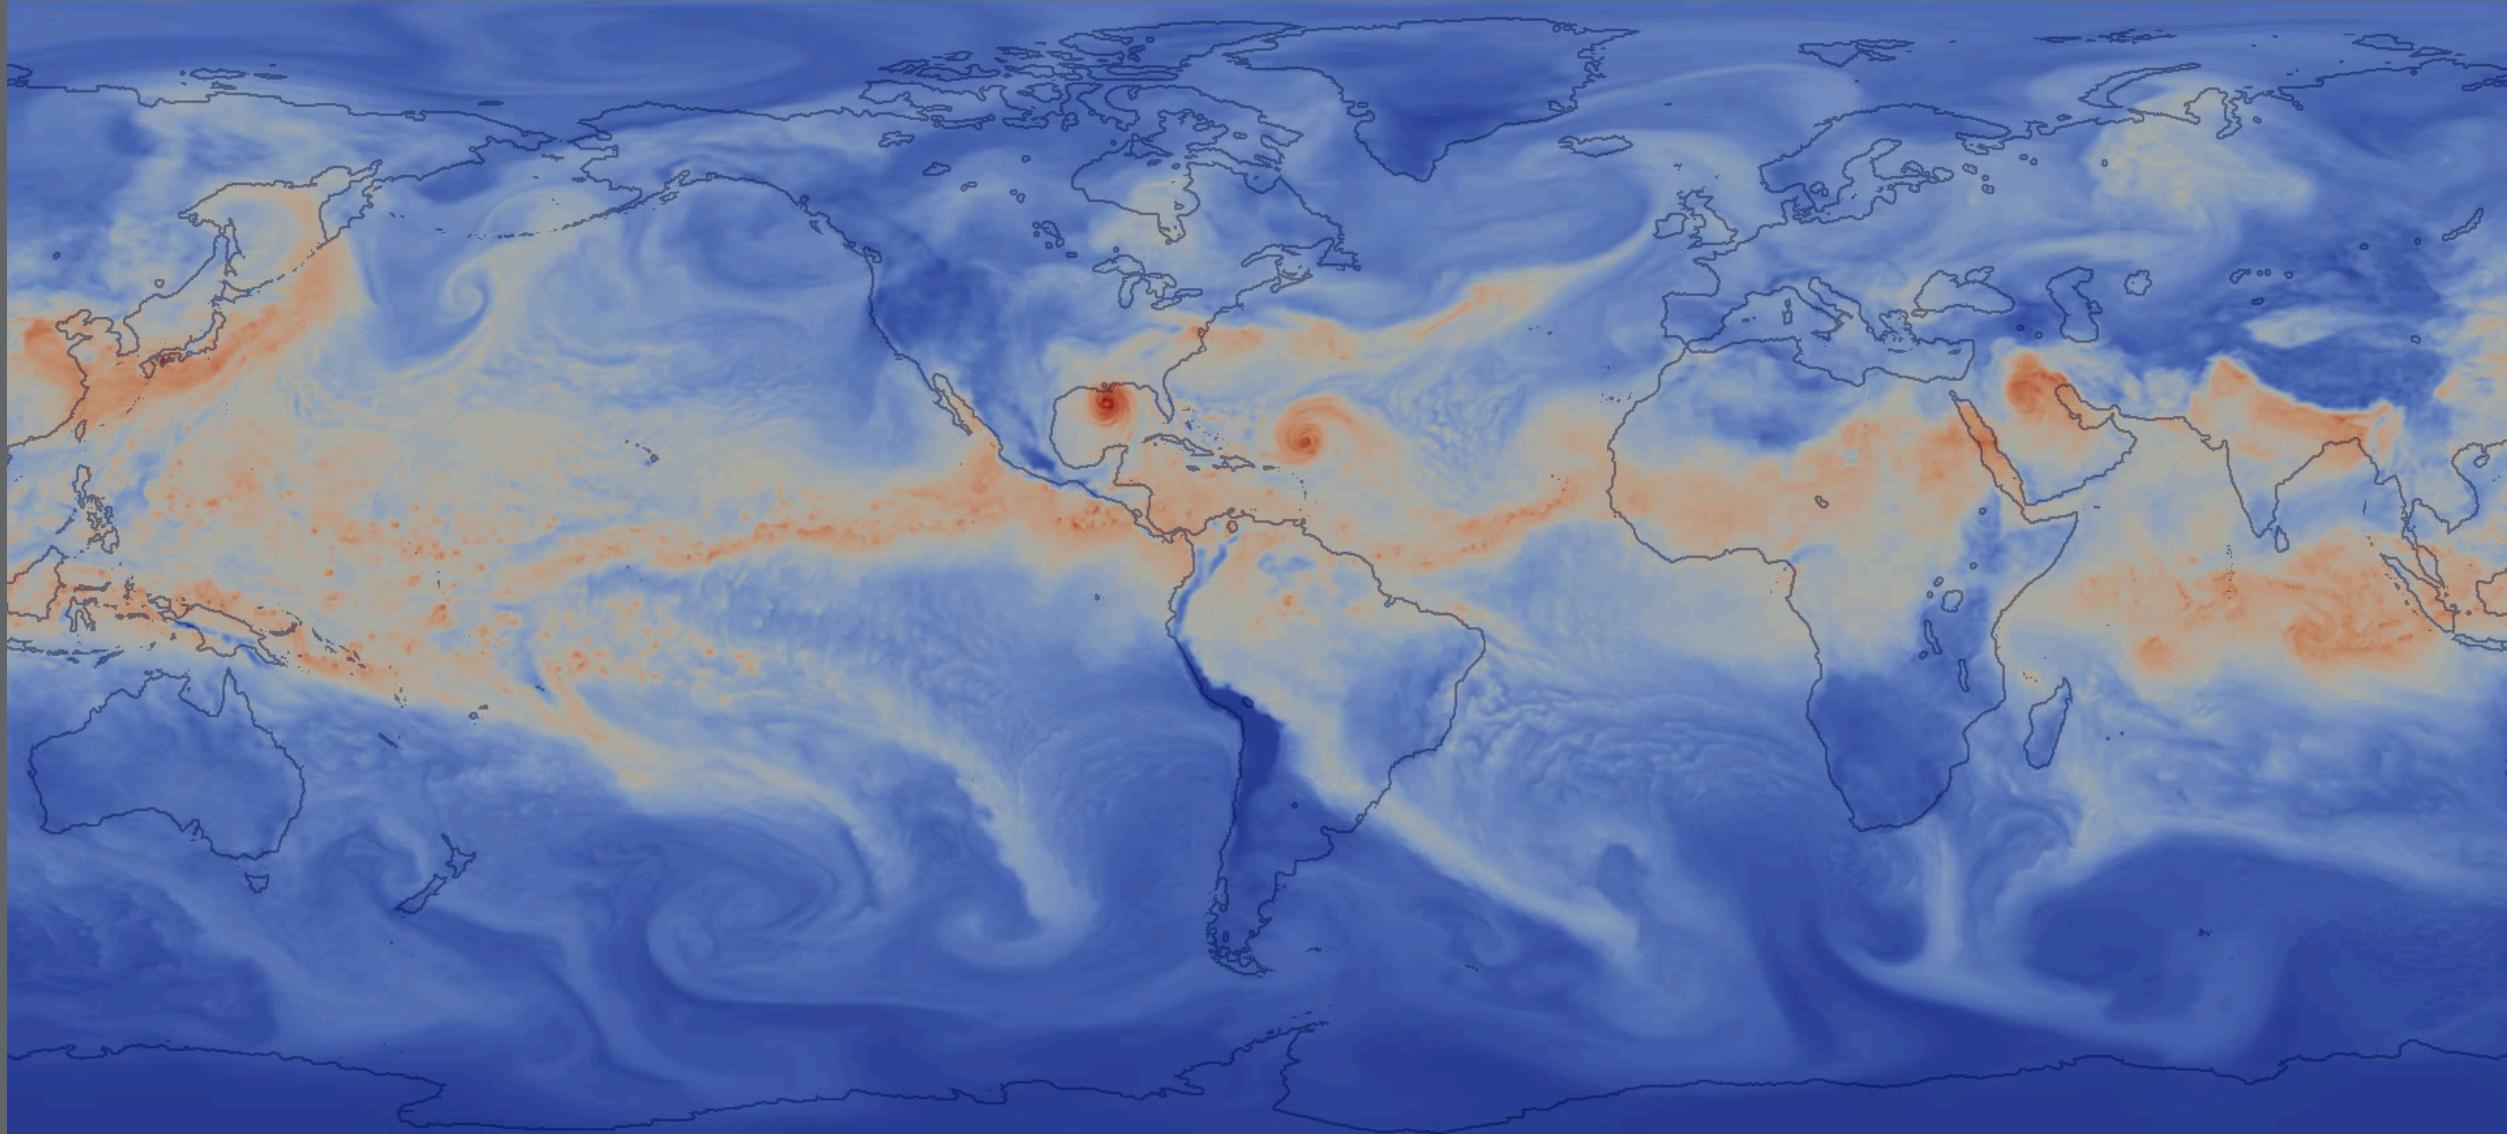
Arterial Blood Flow



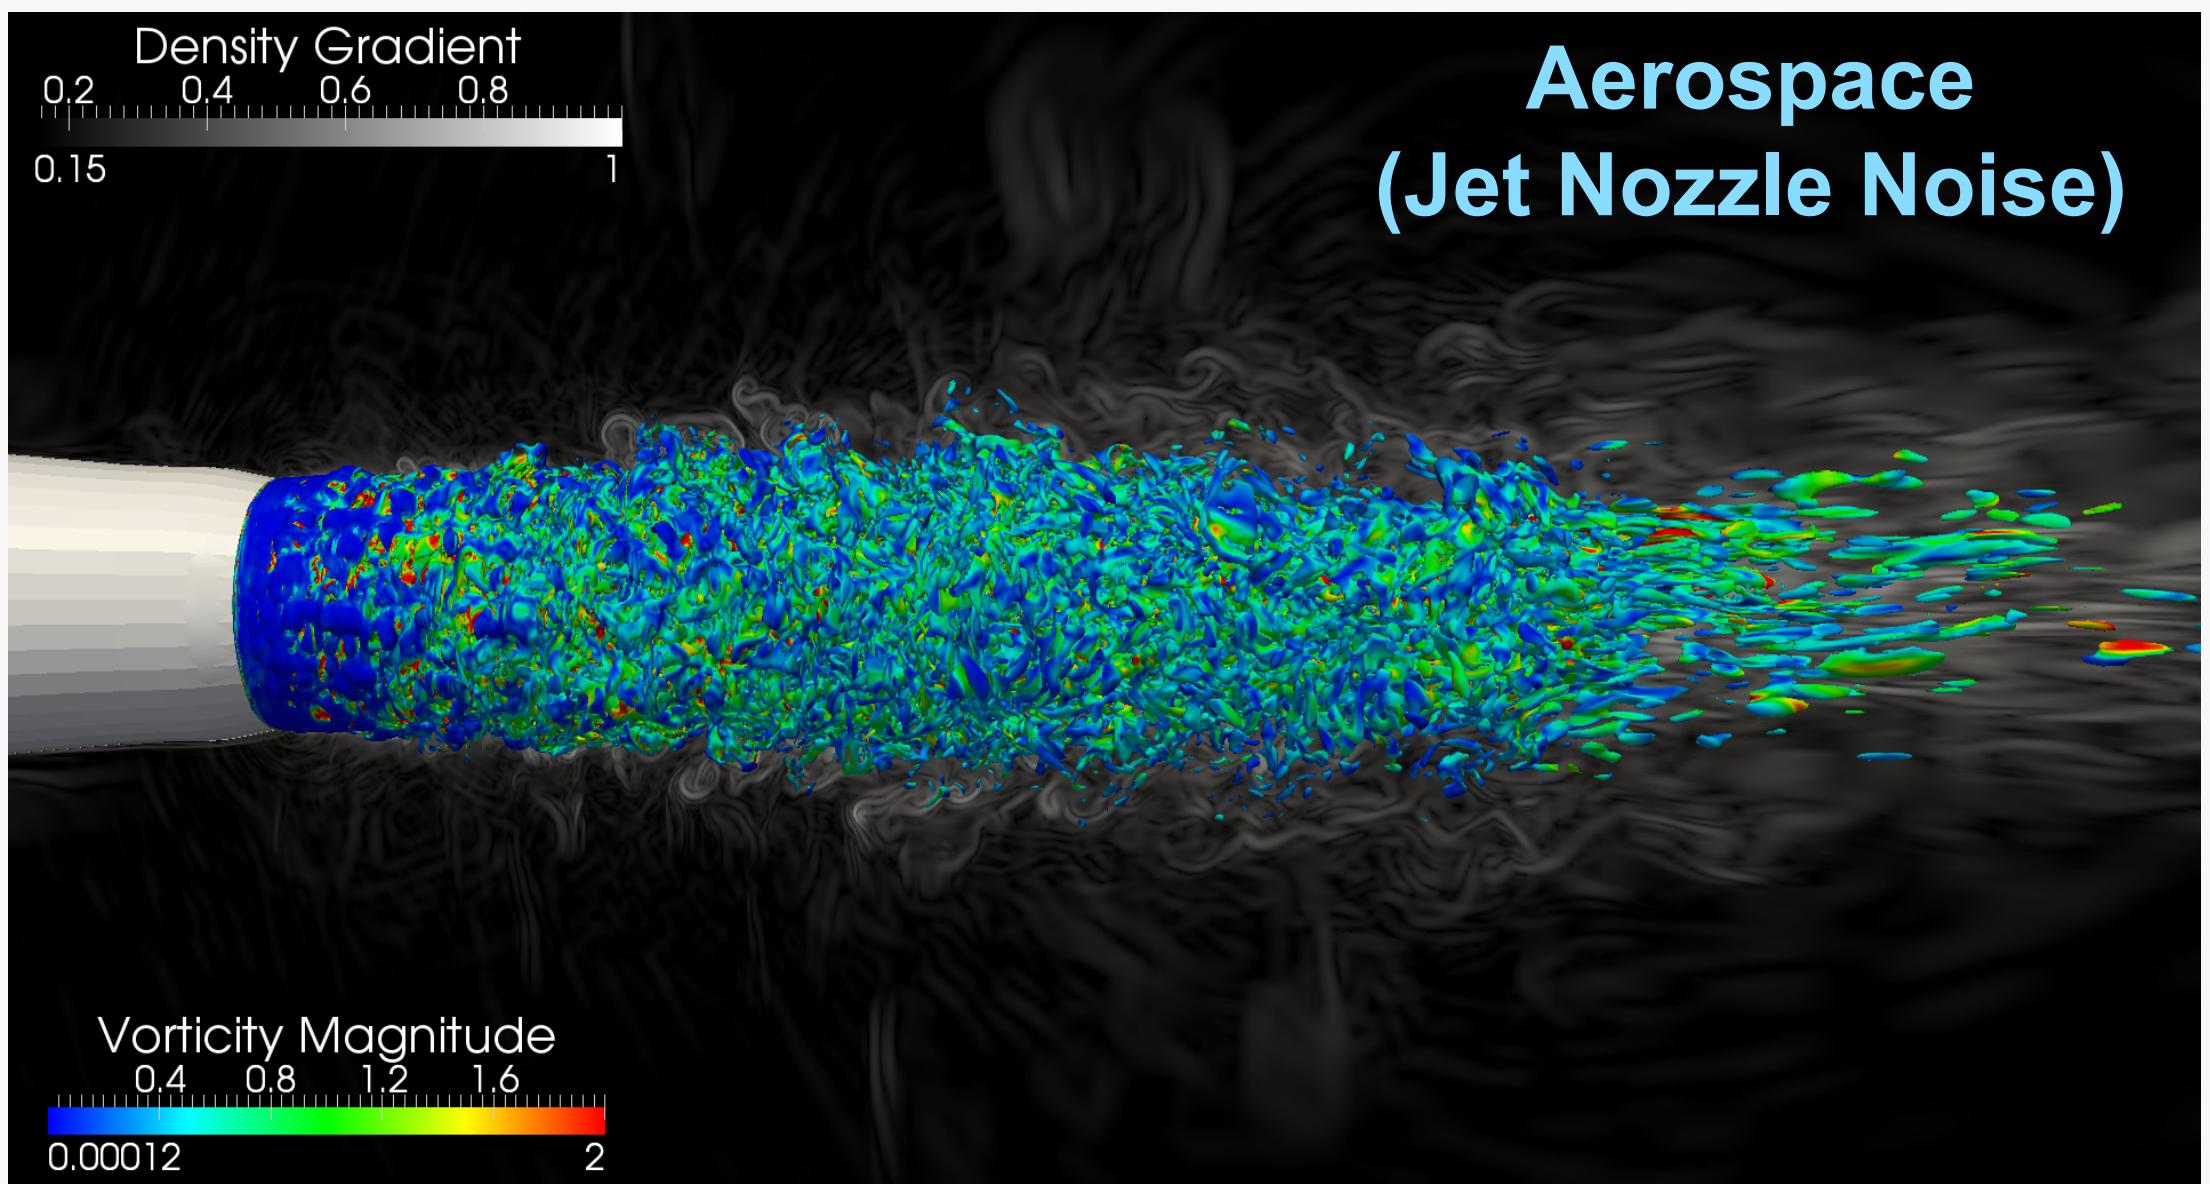
Data courtesy of:
George Karniadakis
and Leopold
Grinberg,
Brown University

Climate

Data courtesy of: Mark Taylor, Sandia National Laboratory; Rob Jacob, Argonne National Laboratory; Warren Washington, National Center for Atmospheric Research

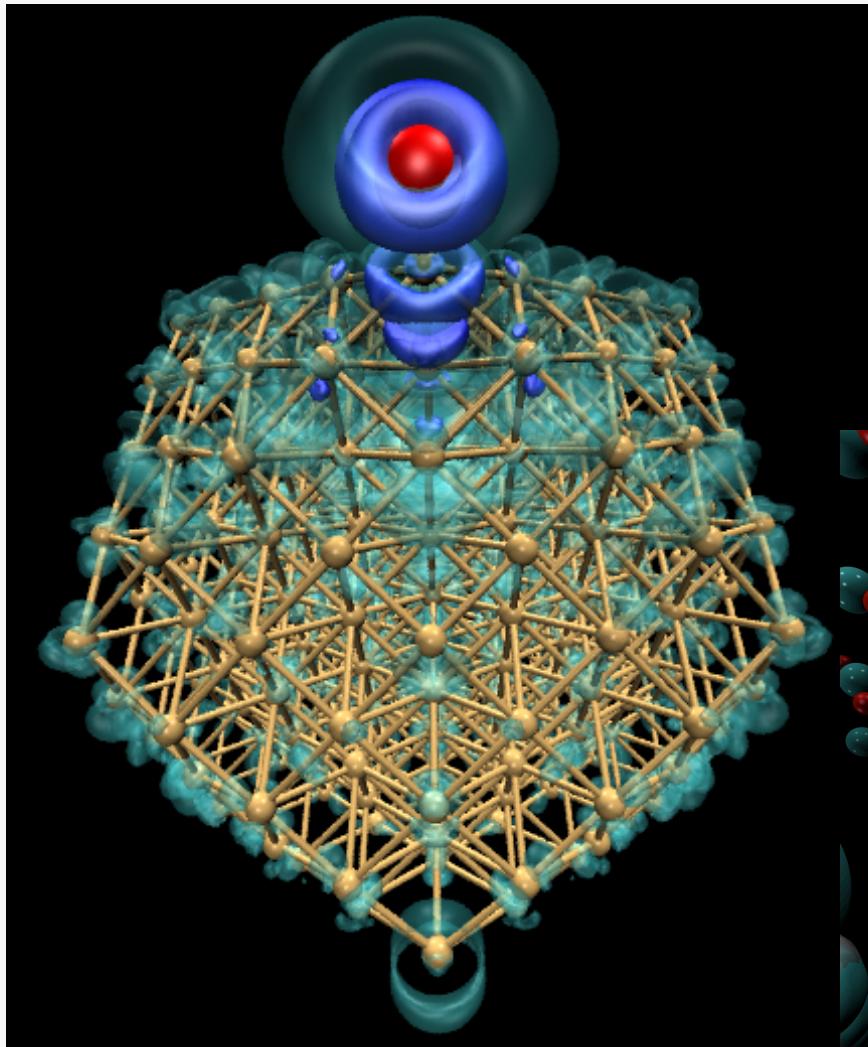


Aerospace (Jet Nozzle Noise)



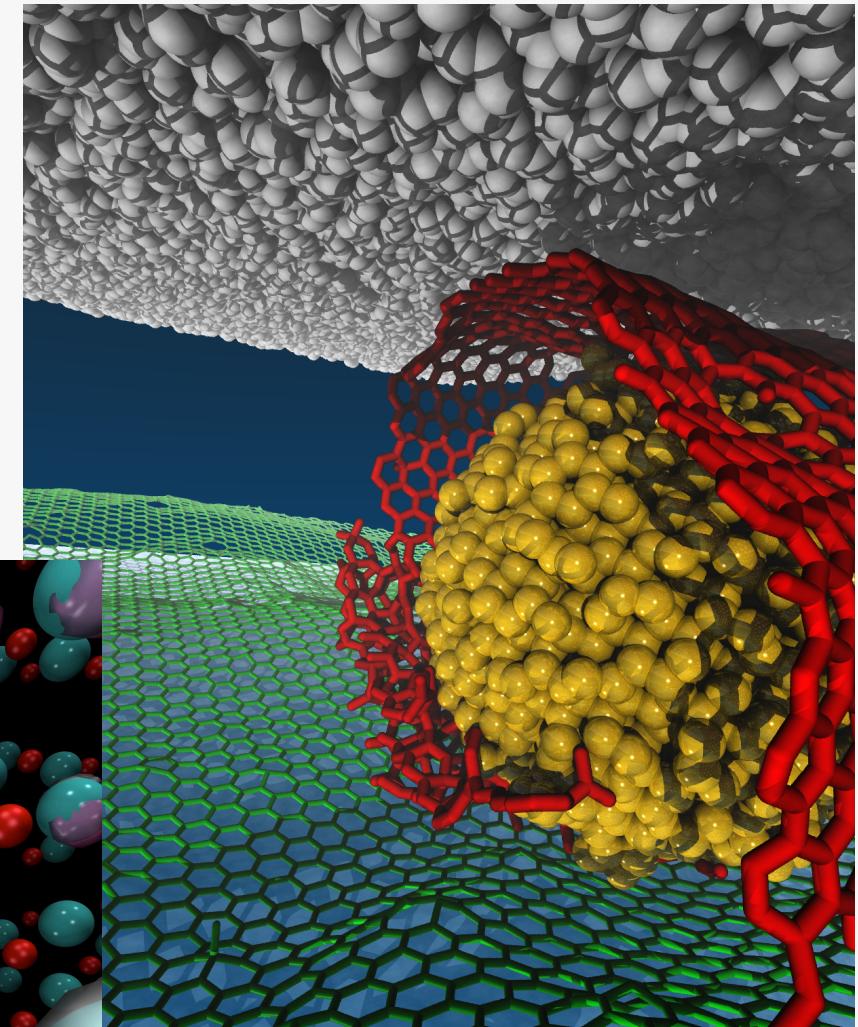
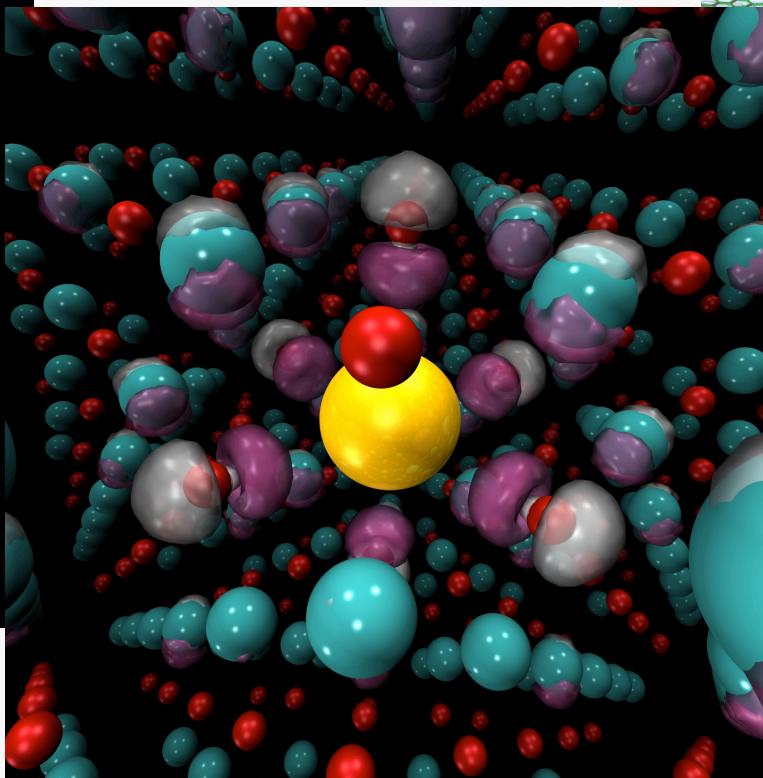
Data courtesy of: Anurag Gupta and Umesh Paliath, General Electric Global Research

Materials Science / Molecular



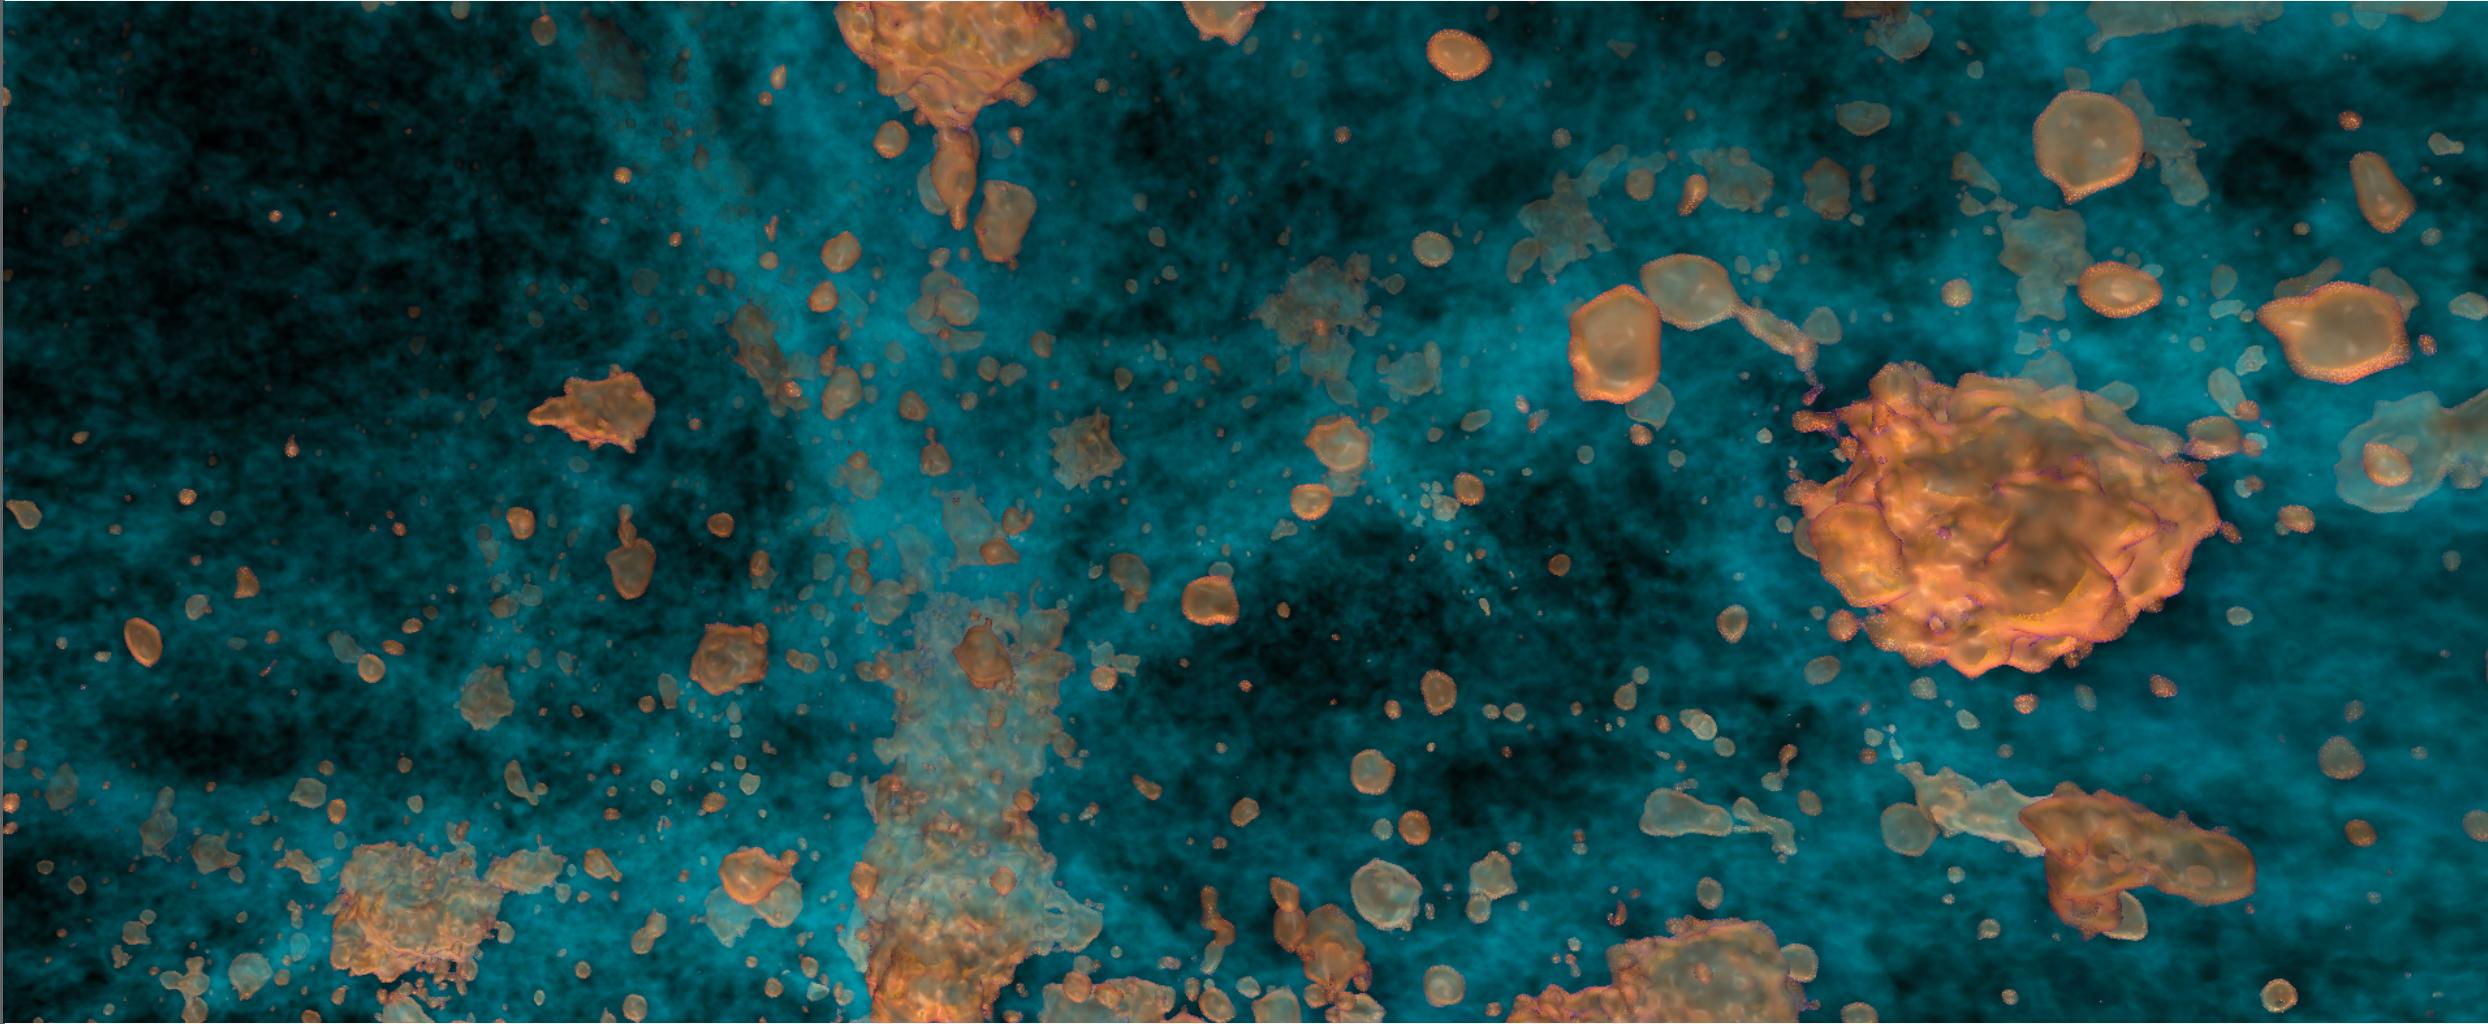
Data courtesy of: Jeff Greeley, Nichols Romero, Argonne National Laboratory

Data courtesy of:
Subramanian
Sankaranarayanan,
Argonne National
Laboratory



Data courtesy of: Paul Kent, Oak Ridge National Laboratory, Anouar Benali, Argonne National Laboratory

Cosmology



Data courtesy of: Salman Habib, Katrin Heitmann, and the HACC team, Argonne National Laboratory

Cooley: Analytics/Visualization cluster

Peak 223 TF

126 nodes; each node has

- Two Intel Xeon E5-2620 Haswell 2.4 GHz 6-core processors
- NVIDIA Tesla K80 graphics processing unit (24GB)
- 384 GB of RAM

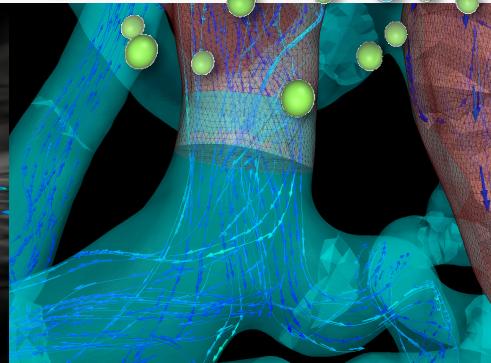
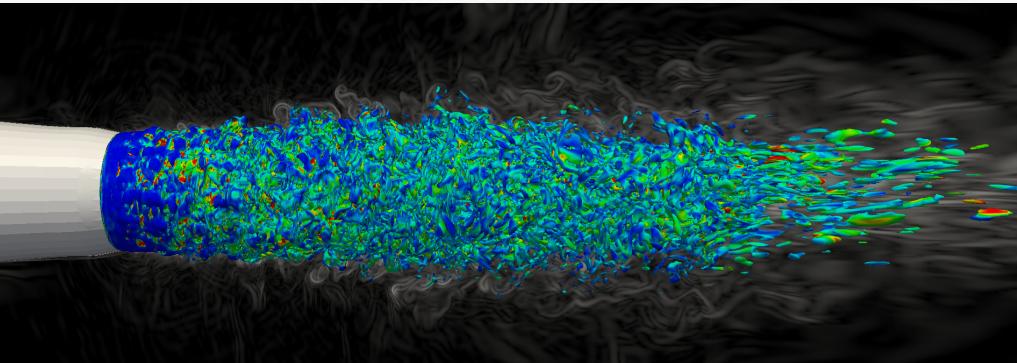
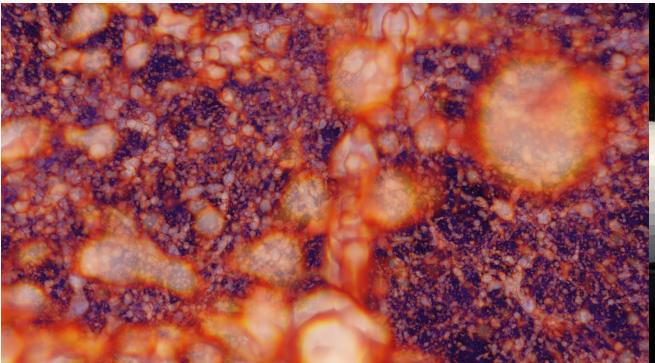
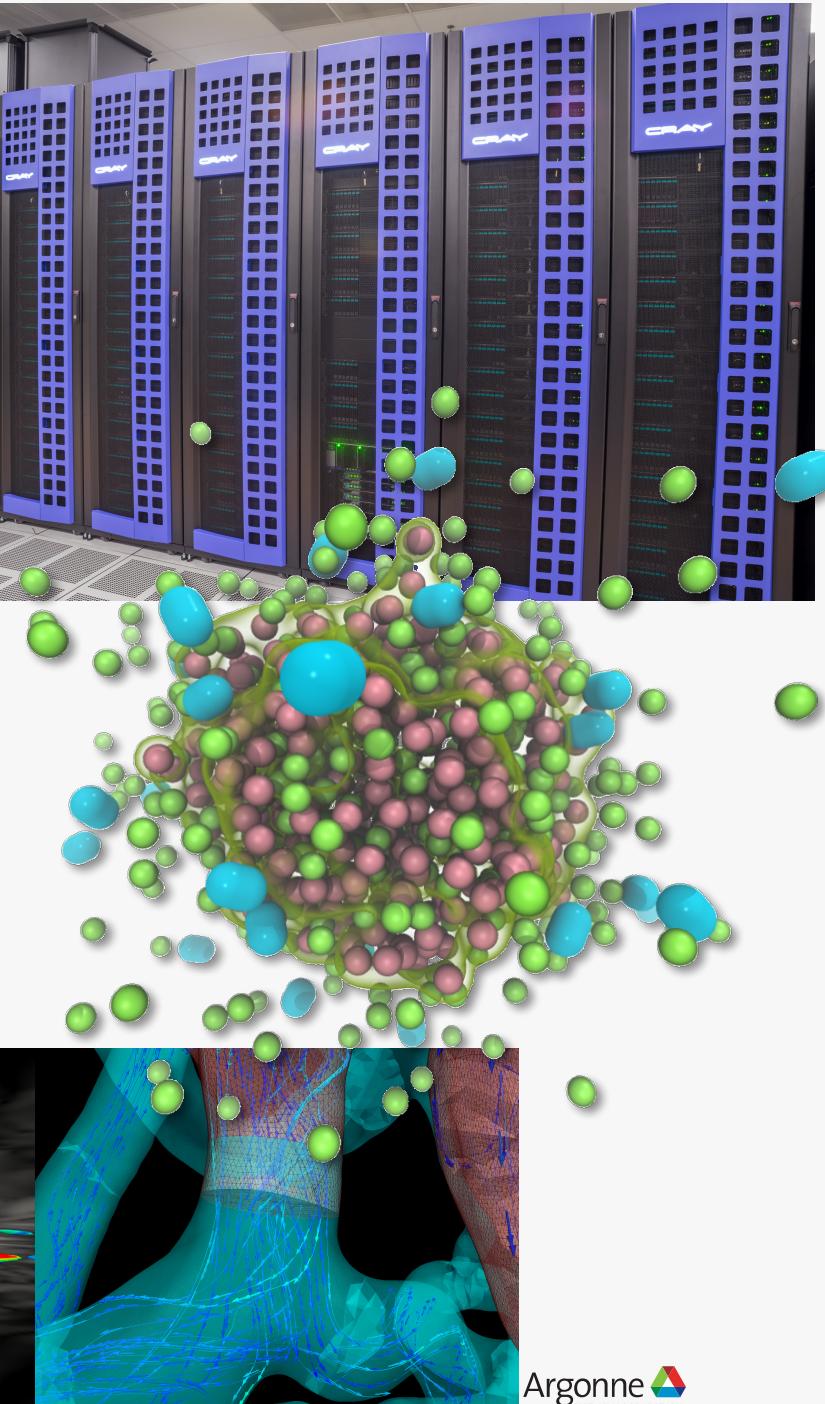
Aggregate RAM of 47 TB

Aggregate GPU memory of ~3TB

Cray CS System

216 port FDR IB switch with uplinks to our QDR infrastructure

Mounts the same GPFS file systems as Mira, Cetus (not Theta)



Moving Data between Theta and Cooley (Mira)

globus.org

- Select Institution: Argonne LCF
- Login using your ALCF username and CryptoCard PIN+password
- Select Endpoints
 - Source: alcf#dtn_theta
 - Destination: alcf#dtn_mira

The screenshot shows the Globus Transfer Files interface. At the top, there are tabs for Manage Data, Publish, Groups, Support, and Account. Below that is a navigation bar with Transfer Files, Activity, Endpoints, Bookmarks, and Console.

The main area is titled "Transfer Files". It shows two file browser panes:

- Endpoint 1:** alcf#dtn_theta
Path: /projects/visualization/
Content: DATA, fujiwara, insley, paraview, preeti, soft, srizzi, turam, v13pkgs
- Endpoint 2:** alcf#dtn_mira
Path: /projects/visualization/
Content: ATPESC, DATA, KOVE, SC17, cam, cooley, esias, flash_data, fmiranda, ftessier, fujiwara, j, kasthuri, marrin, mira, mlewis, preeti, sanuar, srizzi, sshilpika

Below the panes are sections for "Label This Transfer" (with a text input field and placeholder "This will be displayed in your transfer activity.") and "Transfer Settings" (with checkboxes for sync, delete files, preserve times, verify integrity, and encrypt transfer). The "verify file integrity after transfer" checkbox is checked.

Visualization Tools and Data Formats

All Sorts of Tools

Visualization Applications

- [VisIt](#)
- [ParaView*](#)

EnSight

Domain Specific

- [VMD](#), PyMol, Ovito

APIs

- [VTK](#): visualization
- [ITK](#): segmentation & registration

GPU performance

- [vl3](#): shader-based volume and particle rendering

Analysis Environments

- [Matlab](#)
- Parallel R

Utilities

- [GnuPlot](#)
- [ImageMagick](#)



Available on Cooley



Available on Theta

ParaView & VisIt vs. vtk

ParaView & VisIt

- General purpose visualization applications
- GUI-based
- Client / Server model to support remote visualization
- Scriptable / Extendable
- Built on top of vtk (largely)
- *In situ* capabilities

vtk

- Programming environment / API
- Additional capabilities, finer control
- Smaller memory footprint
- Requires more expertise (build custom applications)



Data File Formats (ParaView & VisIt)

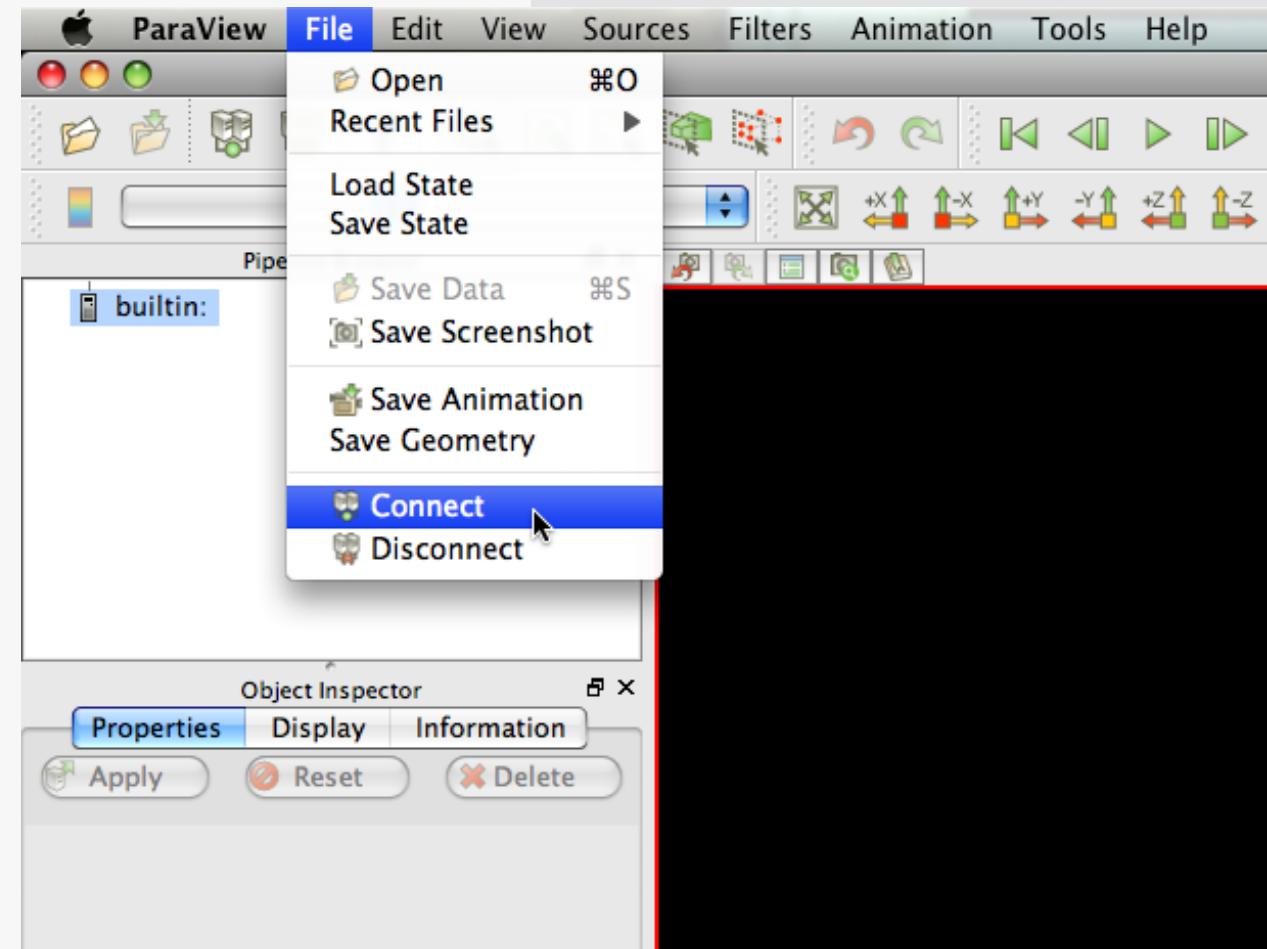
VTK	PLOT3D	Facet	Tetrad
Parallel (partitioned) VTK SpyPlot CTH	HDF5 raw image data	PNG	UNIC
VTK MultiBlock (MultiGroup, Hierarchical, Hierarchical Box)	DEM	SAF	VASP
Legacy VTK	VRML	LS-Dyna	ZeusMP
Parallel (partitioned) legacy VTK	PLY	Nek5000	ANALYZE
EnSight files	Polygonal Protein Data Bank	OVERFLOW	BOV
EnSight Master Server	XMol Molecule	paraDIS	GMV
Exodus	Stereo Lithography	PATRAN	Tecplot
BYU	Gaussian Cube	PFLOTRAN	Vis5D
XDMF	Raw (binary)	Pixie	Xmdv
PLOT2D	AVS	PuReMD	XSF
	Meta Image	S3D	
		SAS	

ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

After launching client locally

- Connect
- Fetch servers (first time only)
- Fetch Theta configuration
- Connect
- Configure server settings
- Connecting: Enter Password
- Open File

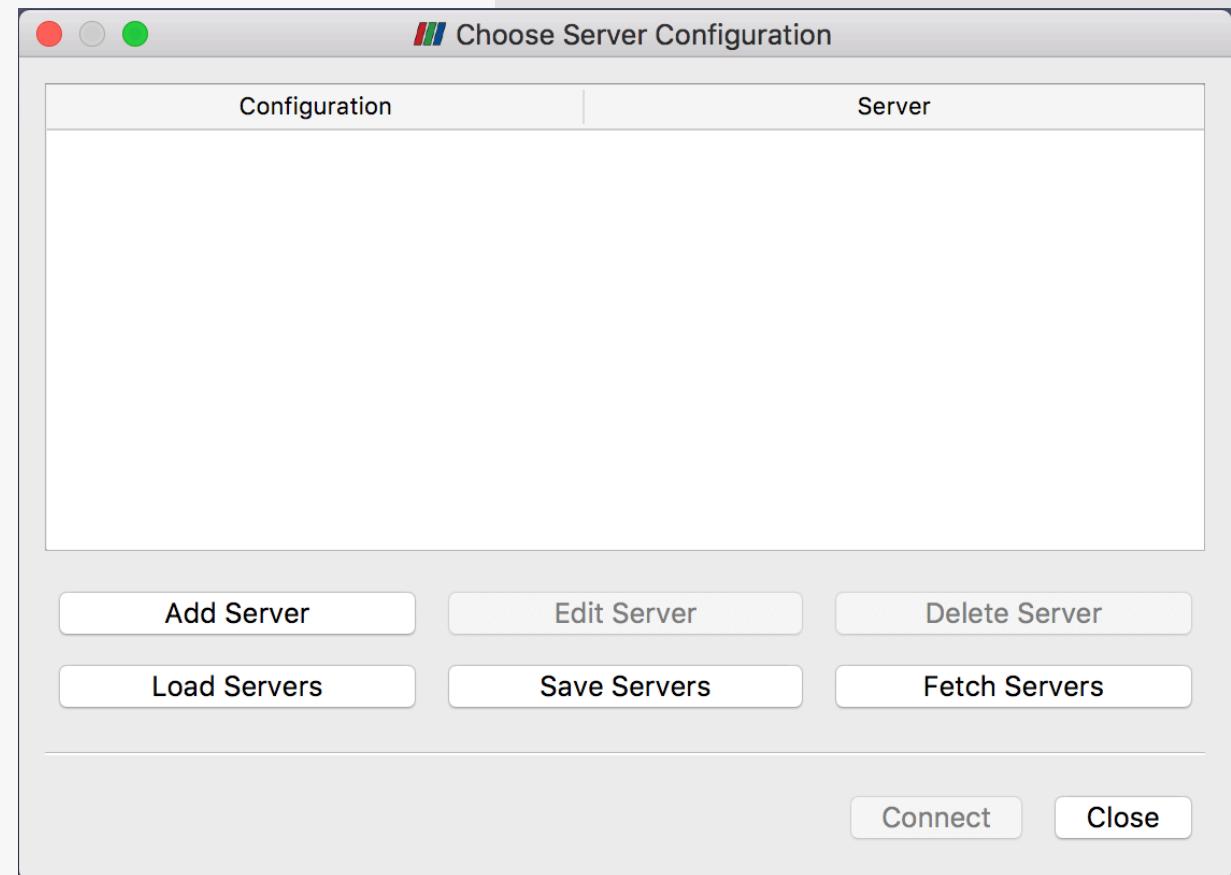


ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

After launching client locally

- Connect
- **Fetch servers (first time only)**
- Fetch Theta configuration
- Connect
- Configure server settings
- Connecting: Enter Password
- **Open File**

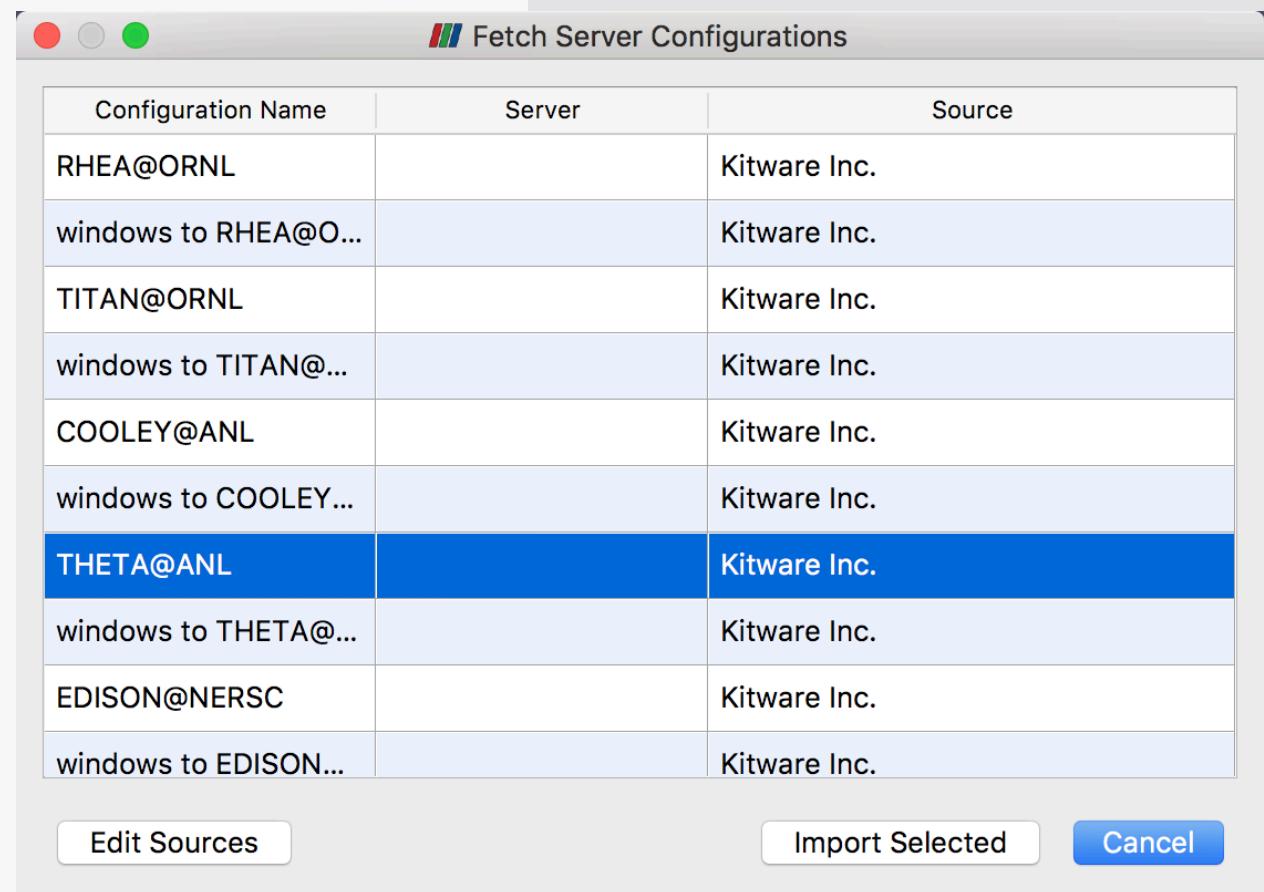


ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

After launching client locally

- Connect
- Fetch servers (first time only)
- **Fetch Theta configuration**
- Connect
- Configure server settings
- Connecting: Enter Password
- **Open File**



The screenshot shows a Mac OS X style dialog box titled "Fetch Server Configurations". It contains a table with three columns: "Configuration Name", "Server", and "Source". The "Source" column for all entries is "Kitware Inc.". The "Configuration Name" column lists several server configurations: RHEA@ORNL, windows to RHEA@O..., TITAN@ORNL, windows to TITAN@..., COOLEY@ANL, windows to COOLEY..., THETA@ANL, windows to THETA@..., EDISON@NERSC, and windows to EDISON... . The row for "THETA@ANL" is highlighted with a blue background.

Configuration Name	Server	Source
RHEA@ORNL		Kitware Inc.
windows to RHEA@O...		Kitware Inc.
TITAN@ORNL		Kitware Inc.
windows to TITAN@...		Kitware Inc.
COOLEY@ANL		Kitware Inc.
windows to COOLEY...		Kitware Inc.
THETA@ANL		Kitware Inc.
windows to THETA@...		Kitware Inc.
EDISON@NERSC		Kitware Inc.
windows to EDISON...		Kitware Inc.

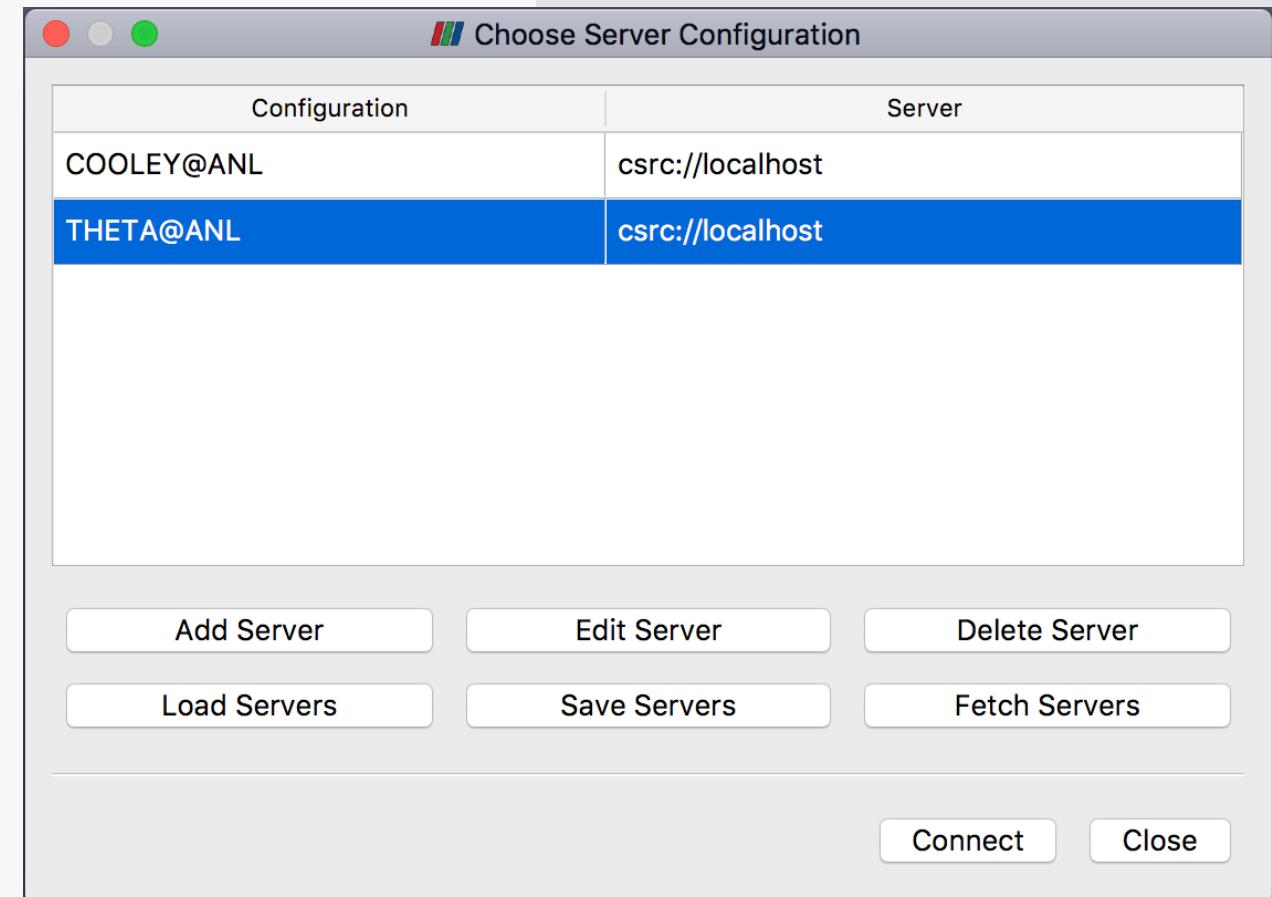
Buttons at the bottom of the dialog include "Edit Sources", "Import Selected", and "Cancel".

ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

After launching client locally

- Connect
- Fetch servers (first time only)
- Fetch Theta configuration
- **Connect**
- Configure server settings
- Connecting: Enter Password
- **Open File**

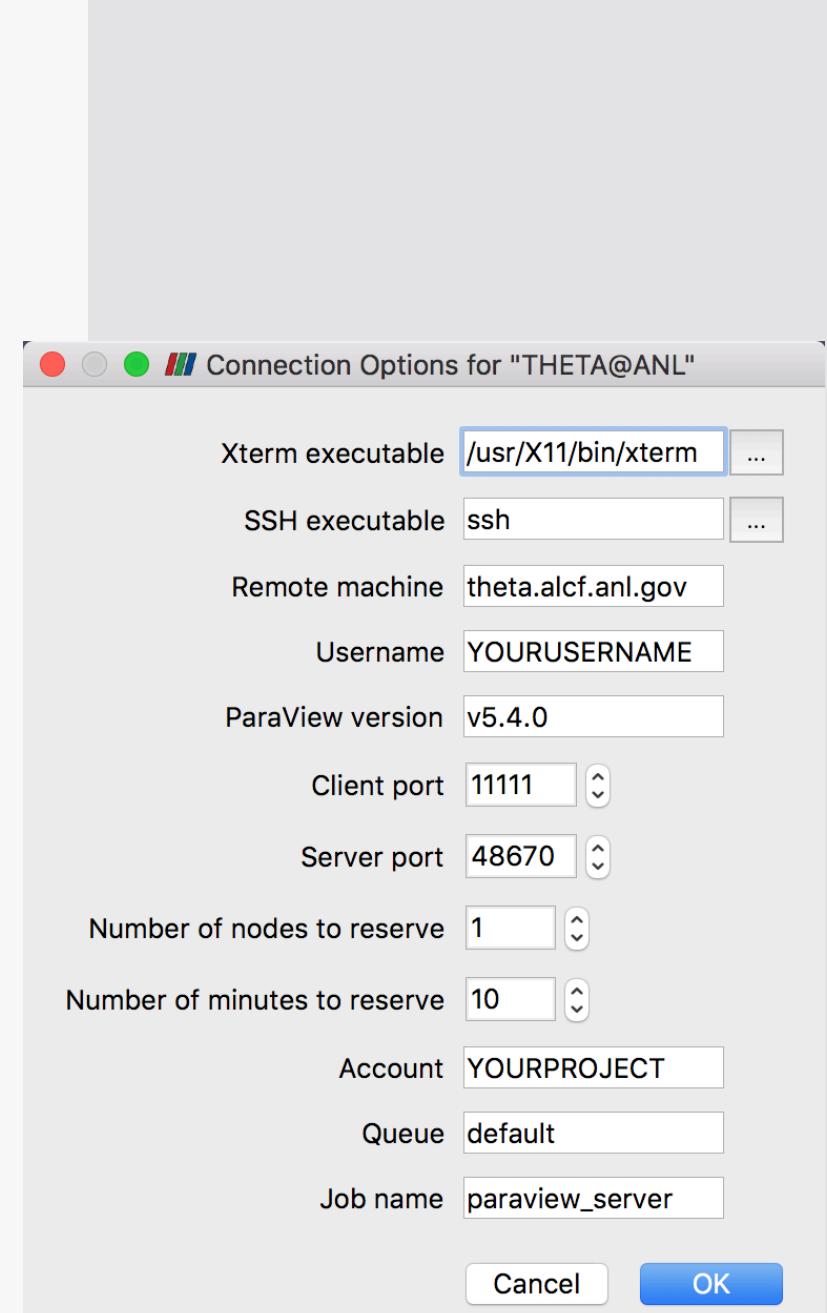


ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

After launching client locally

- Connect
- Fetch servers (first time only)
- Fetch Theta configuration
- Connect
- **Configure server settings**
- Connecting: Enter Password
- **Open File**

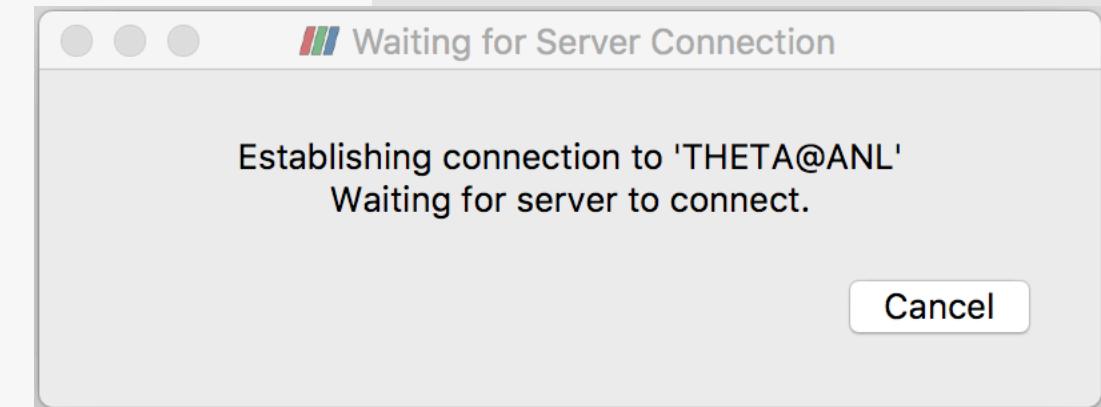


ParaView on Theta

Version 5.4.1 (Client and Server versions must match)

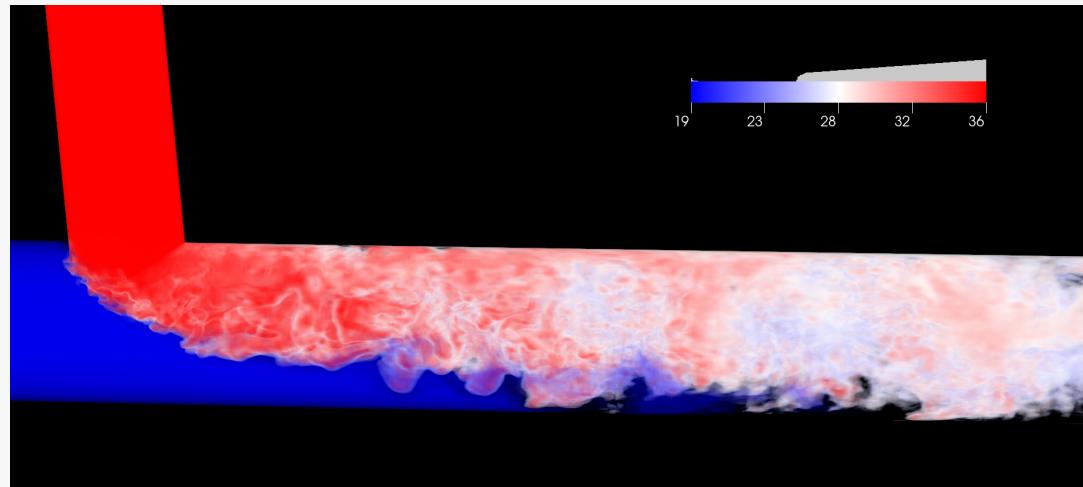
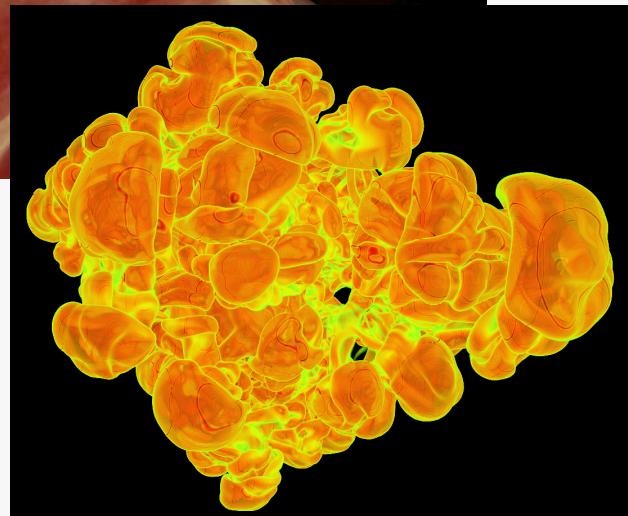
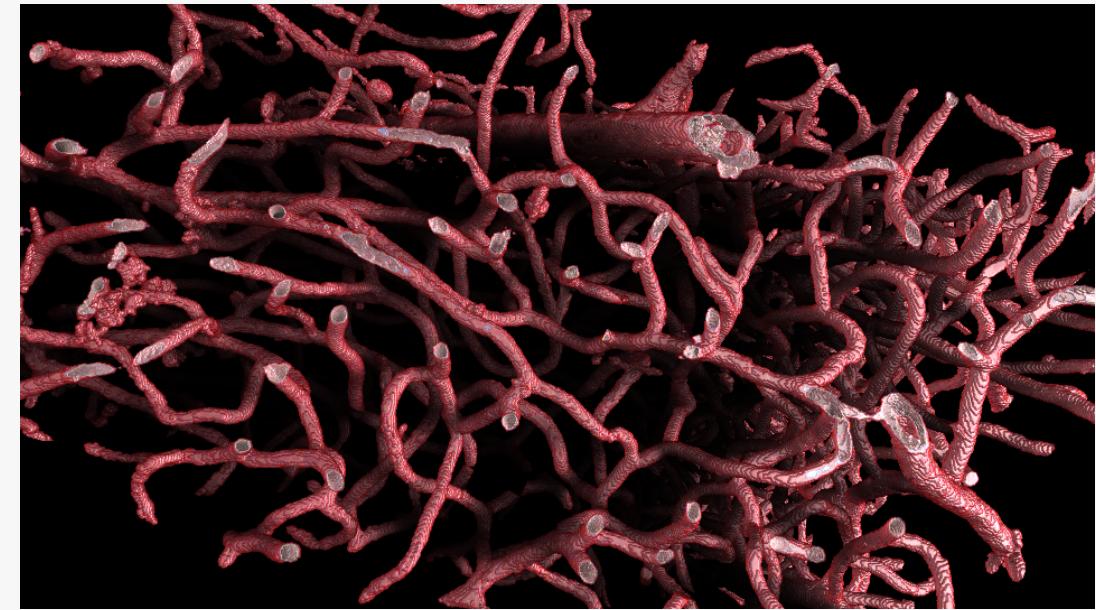
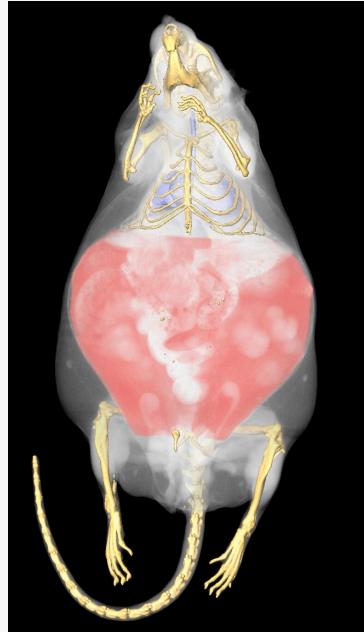
After launching client locally

- Connect
- Fetch servers (first time only)
- Fetch Theta configuration
- Connect
- Configure server settings
- **Connecting: Enter Password**
- **Open File**



Data Representations

Data Representations: Volume Rendering



Data Representations: Glyphs

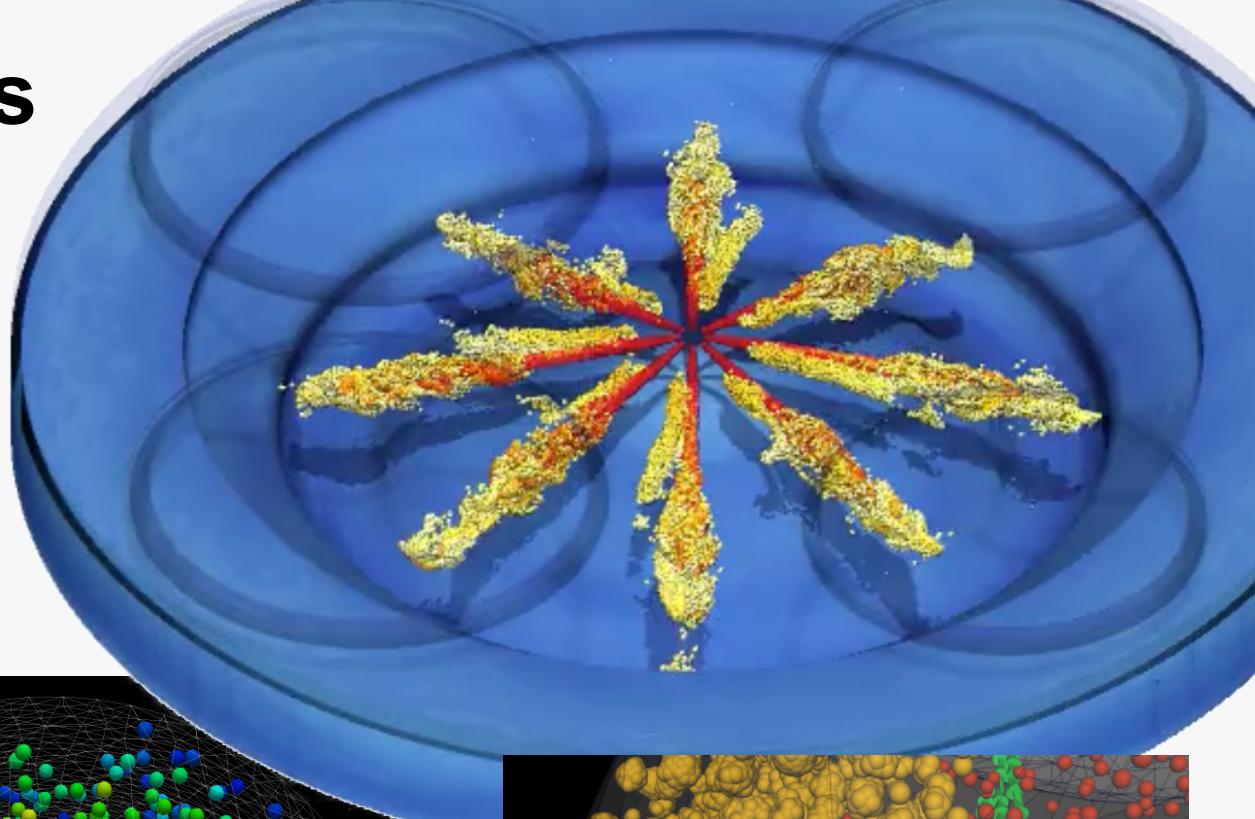
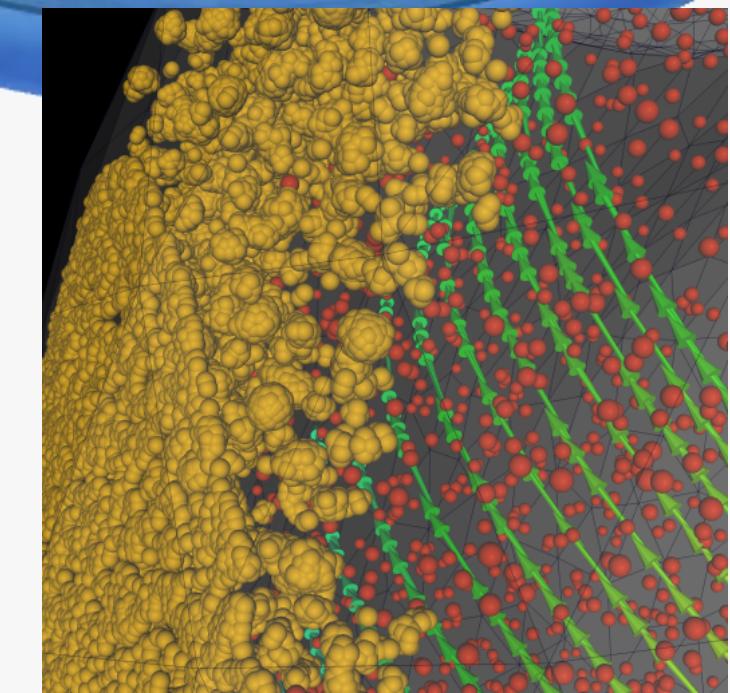
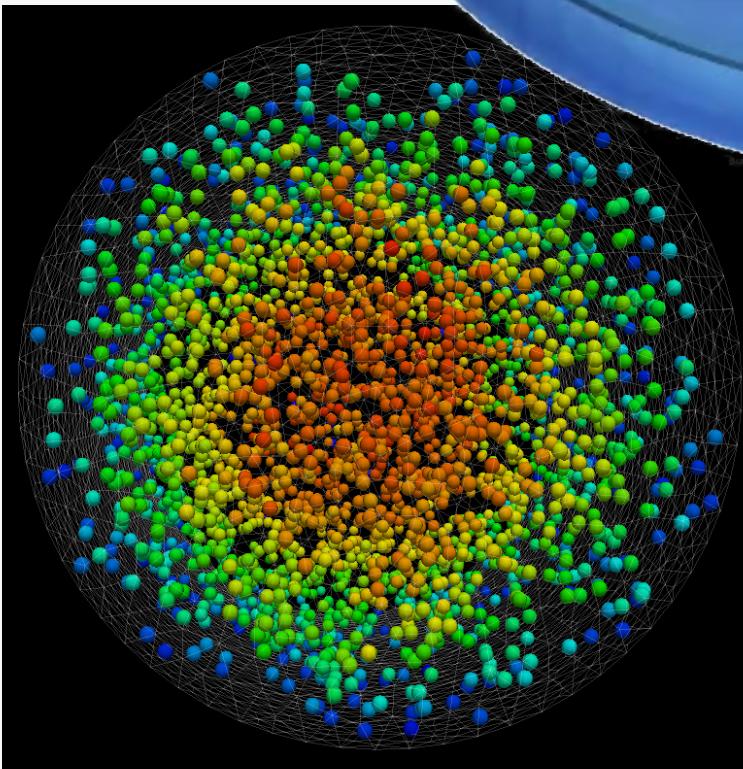
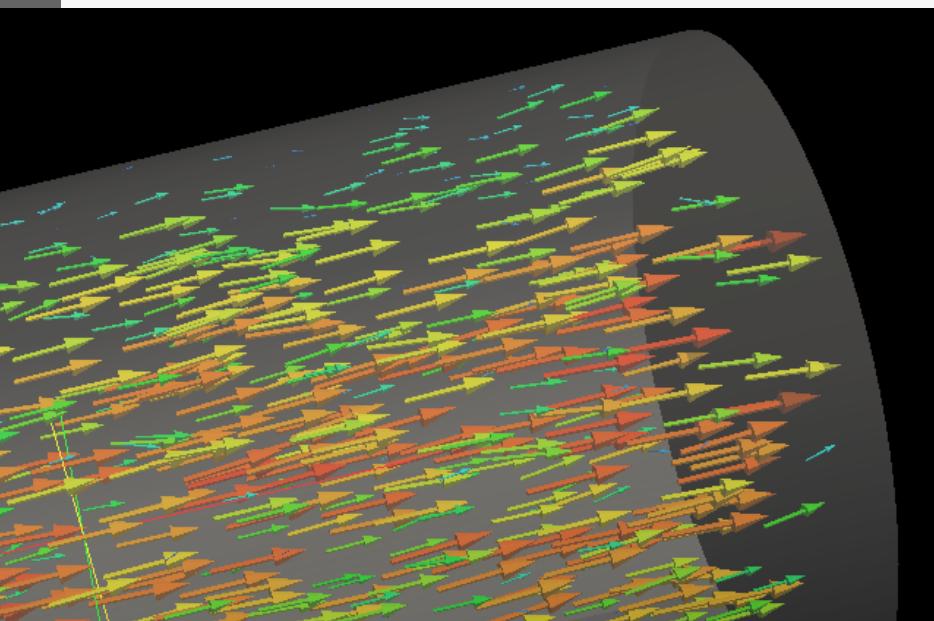
2D or 3D geometric object to represent point data

Location dictated by coordinate

- 3D location on mesh
- 2D position in table/graph

Attributes of graphical entity dictated by attributes of data

- color, size, orientation



Data Representations: Contours (Isosurfaces)

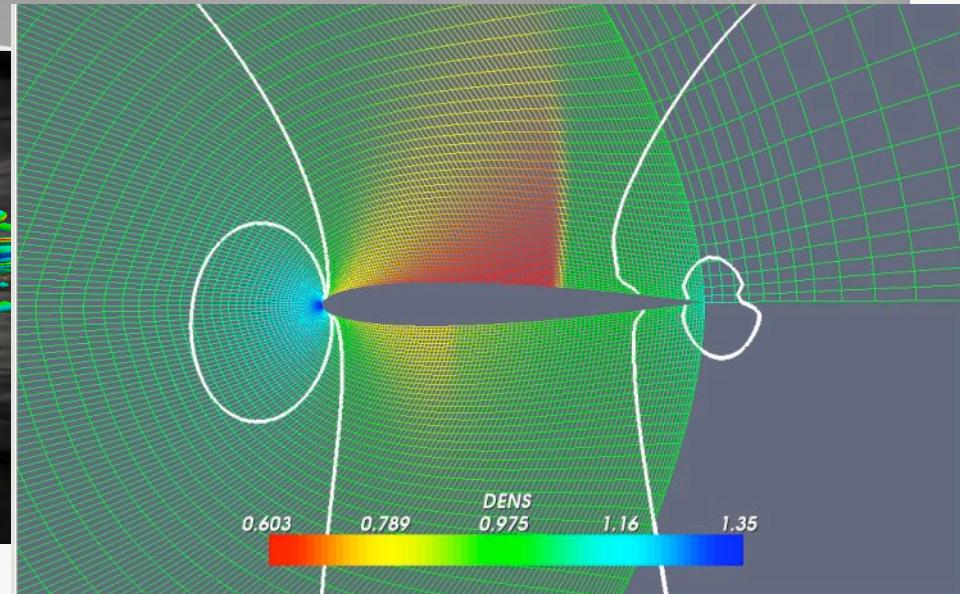
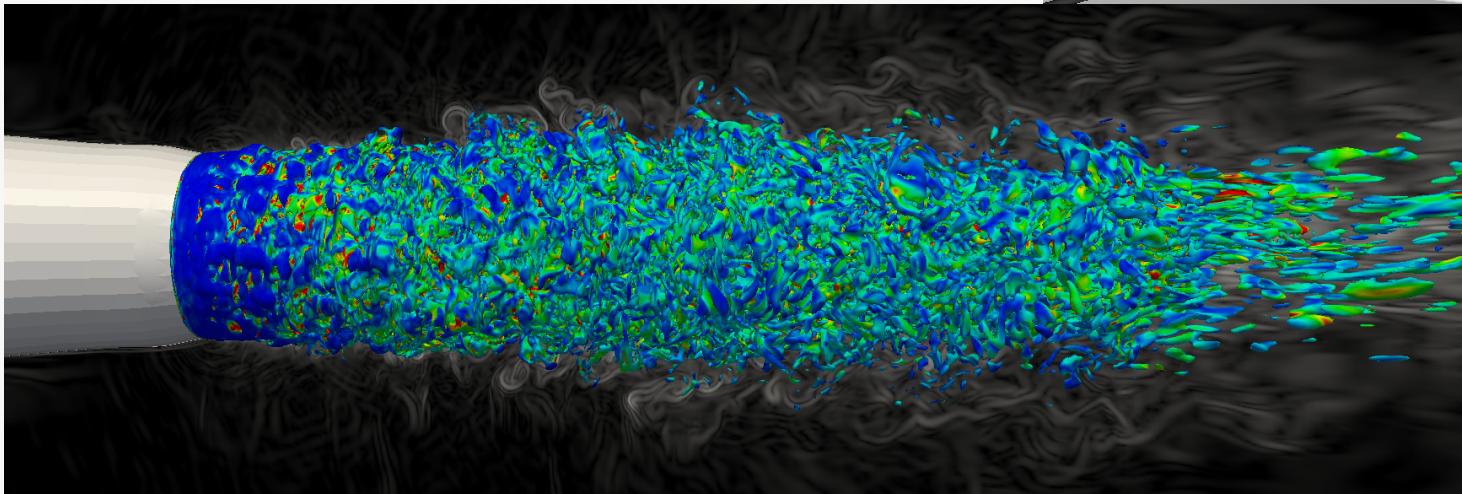
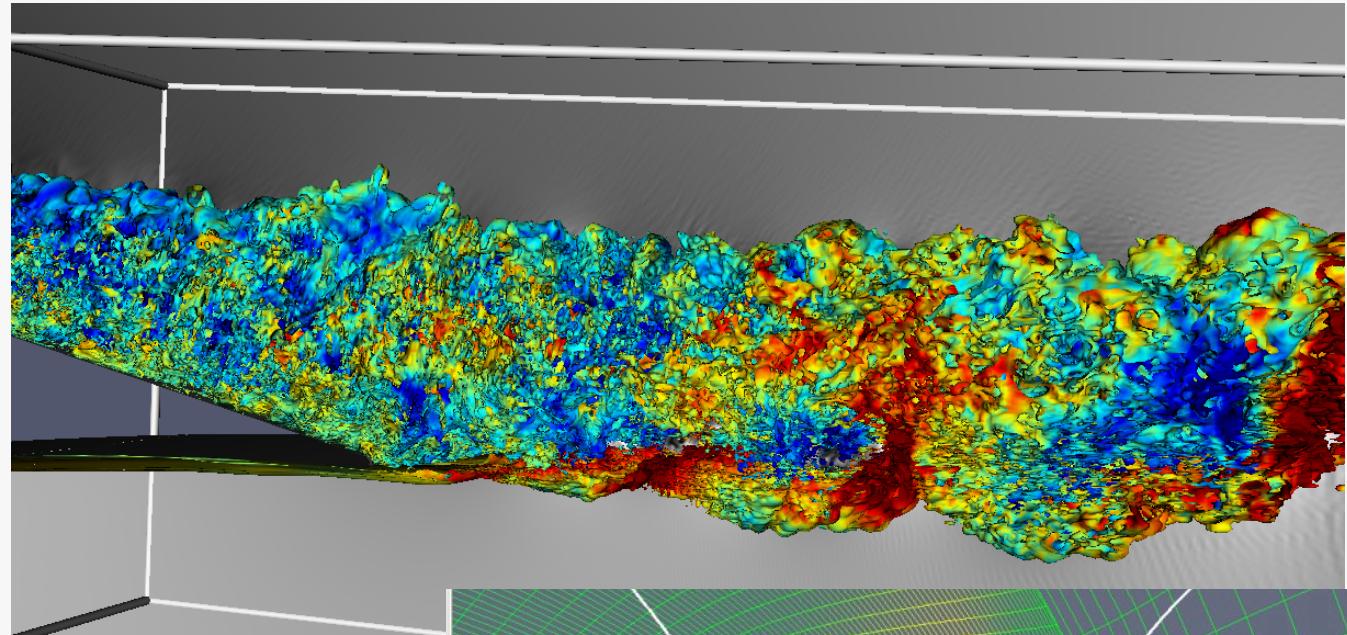
A Line (2D) or Surface (3D),
representing a constant value

VisIt & ParaView:

- good at this

vtk:

- same, but again requires more effort



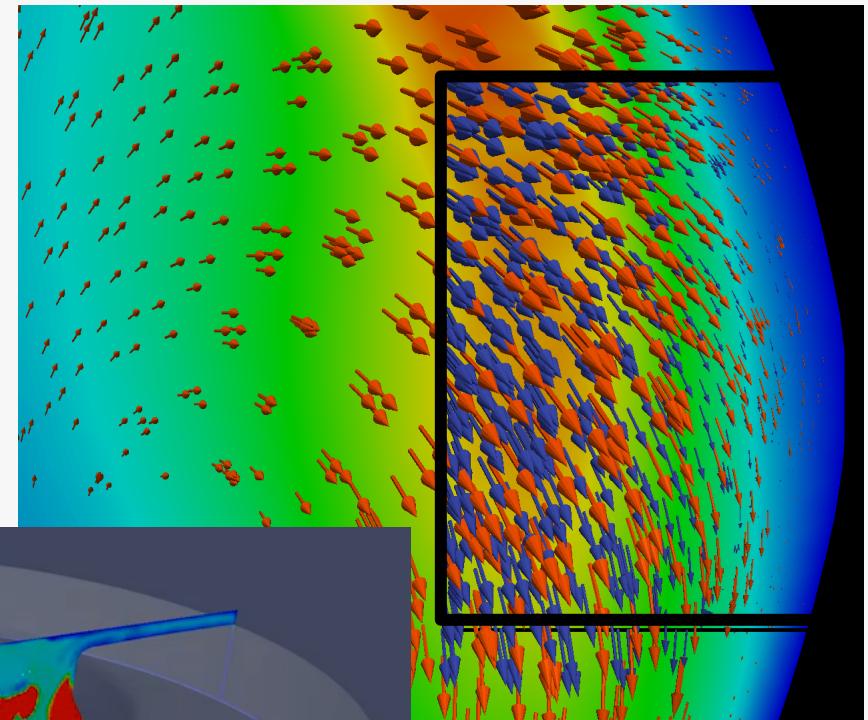
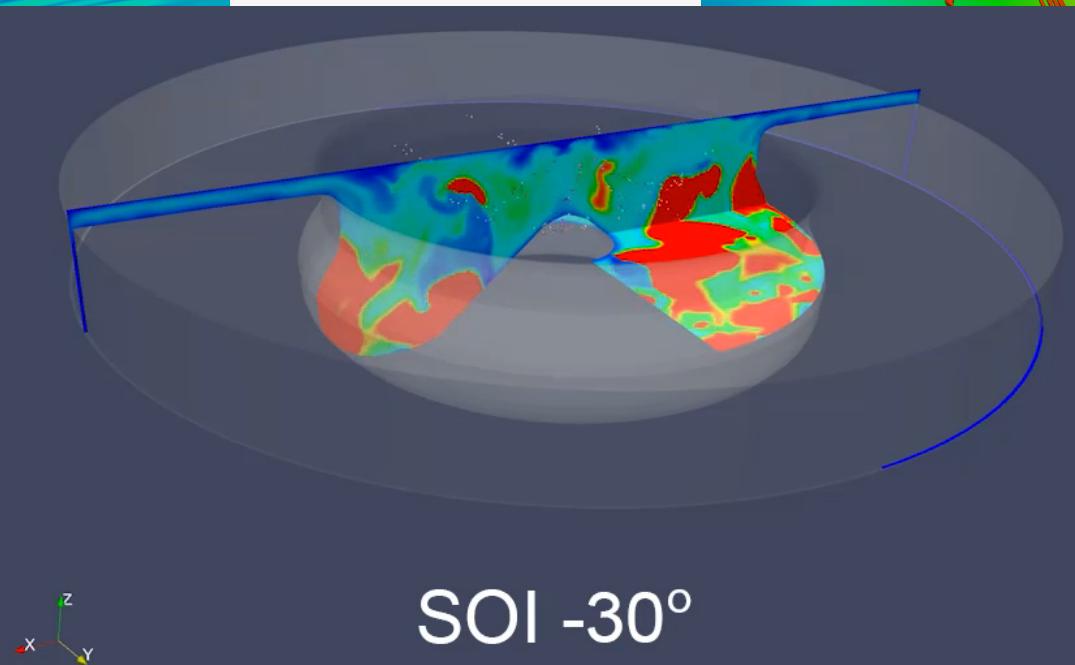
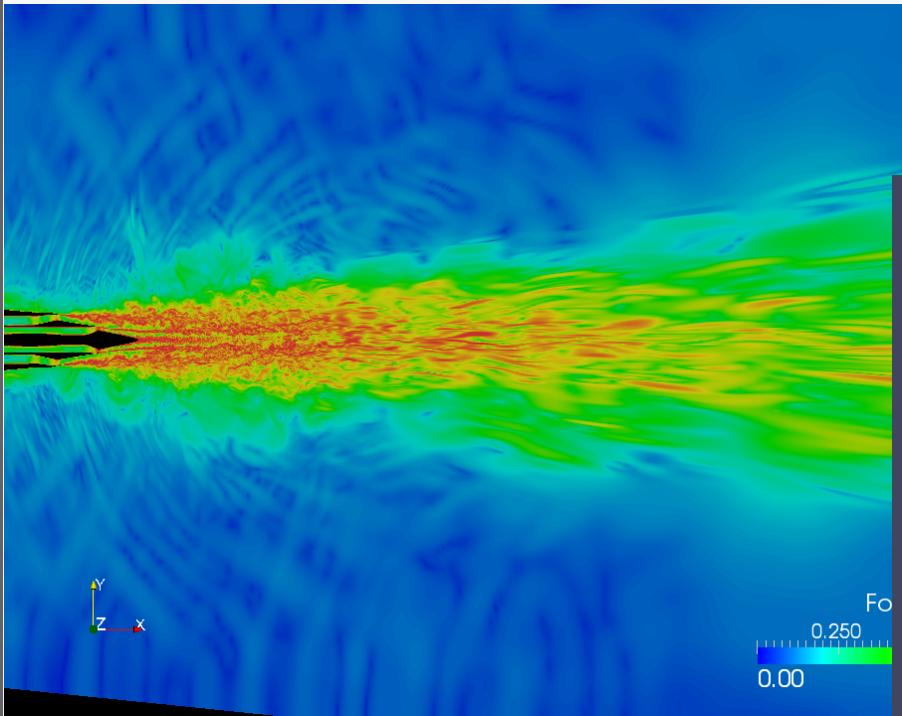
Data Representations: Cutting Planes

Slice a plane through the data

- Can apply additional visualization methods to resulting plane

VisIt & ParaView & vtk good at this

VMD has similar capabilities for some data formats

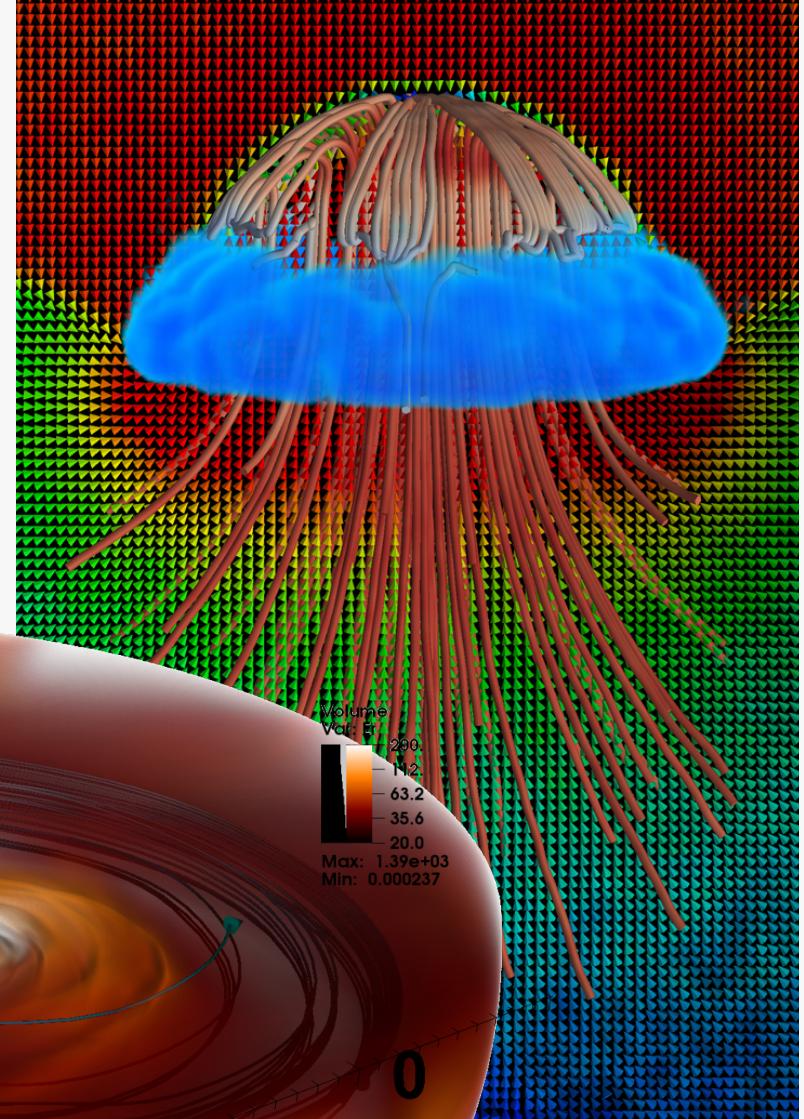
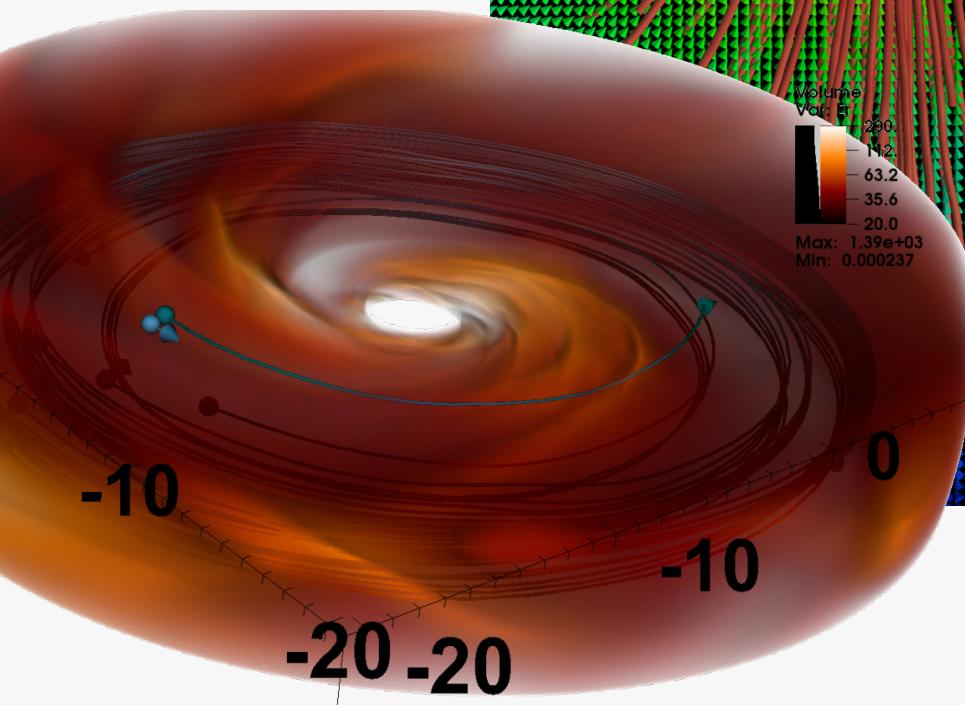
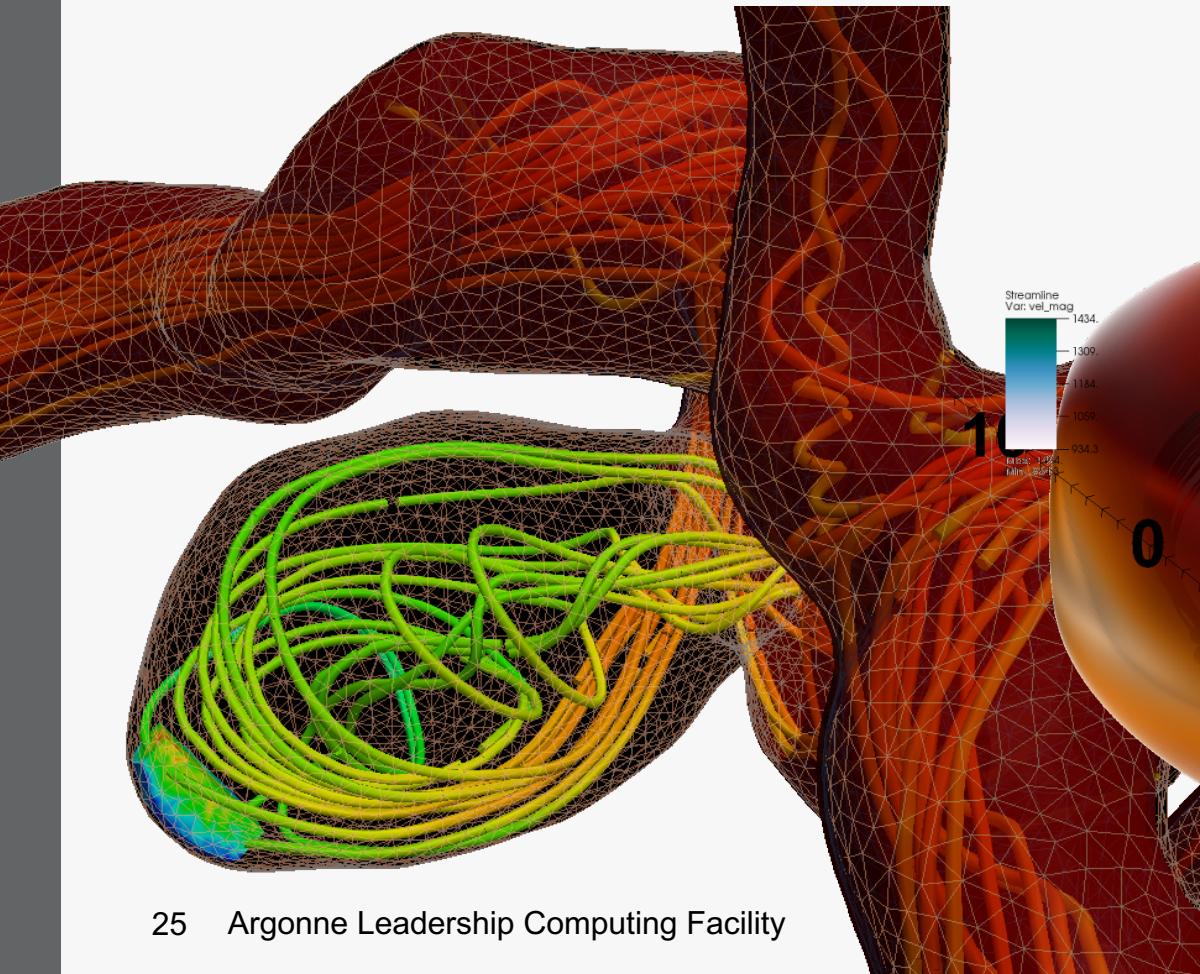


Data Representations: Streamlines

From vector field on a mesh (needs connectivity)

- Show the direction an element will travel in at any point in time.

VisIt & ParaView & vtk good at this



Molecular Dynamics Visualization

VMD:

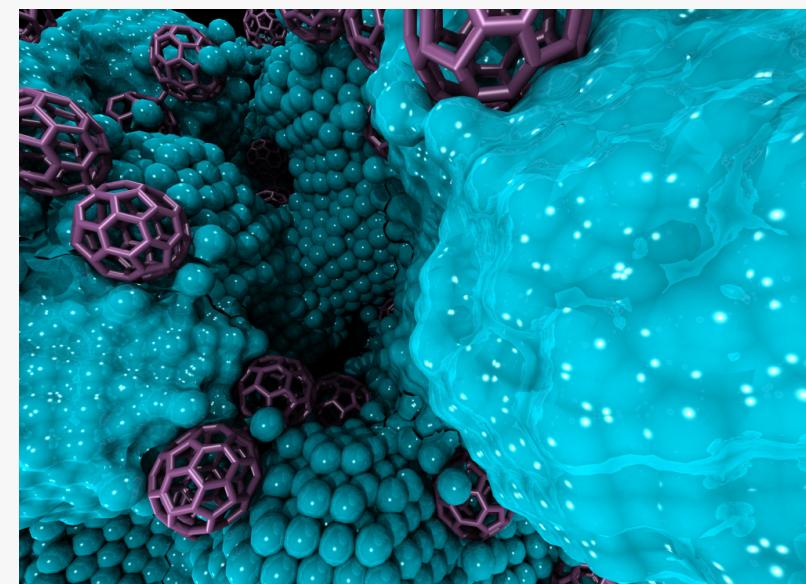
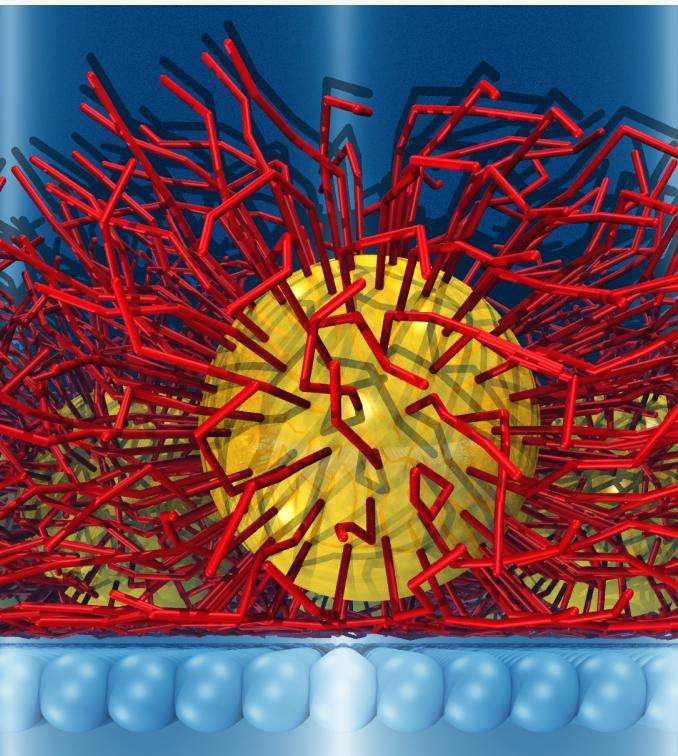
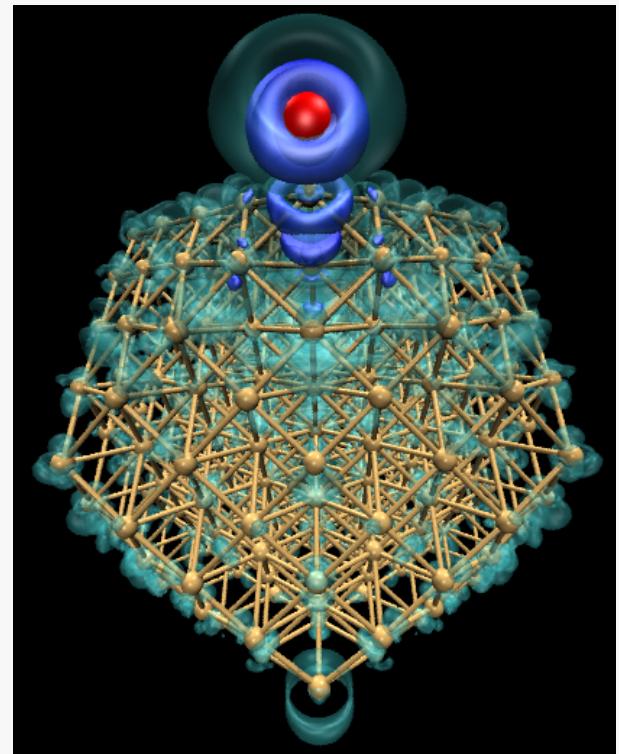
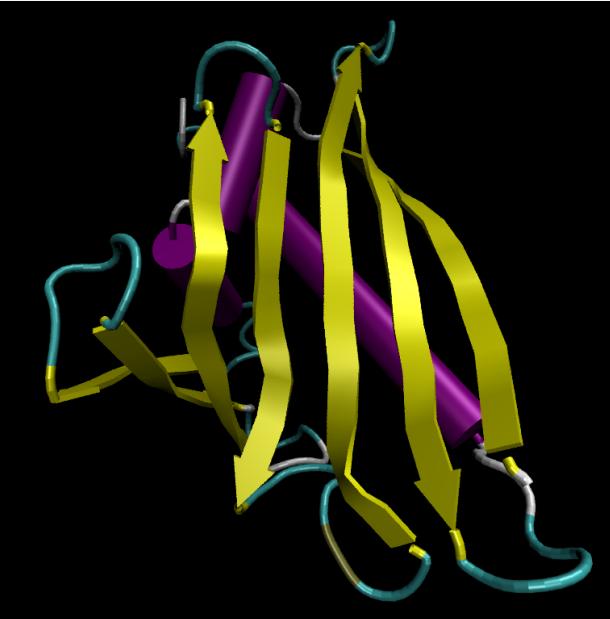
- Lots of domain-specific representations
- Many different file formats
- Animation
- Scriptable

VisIt & ParaView:

- Limited support for these types of representations, but improving

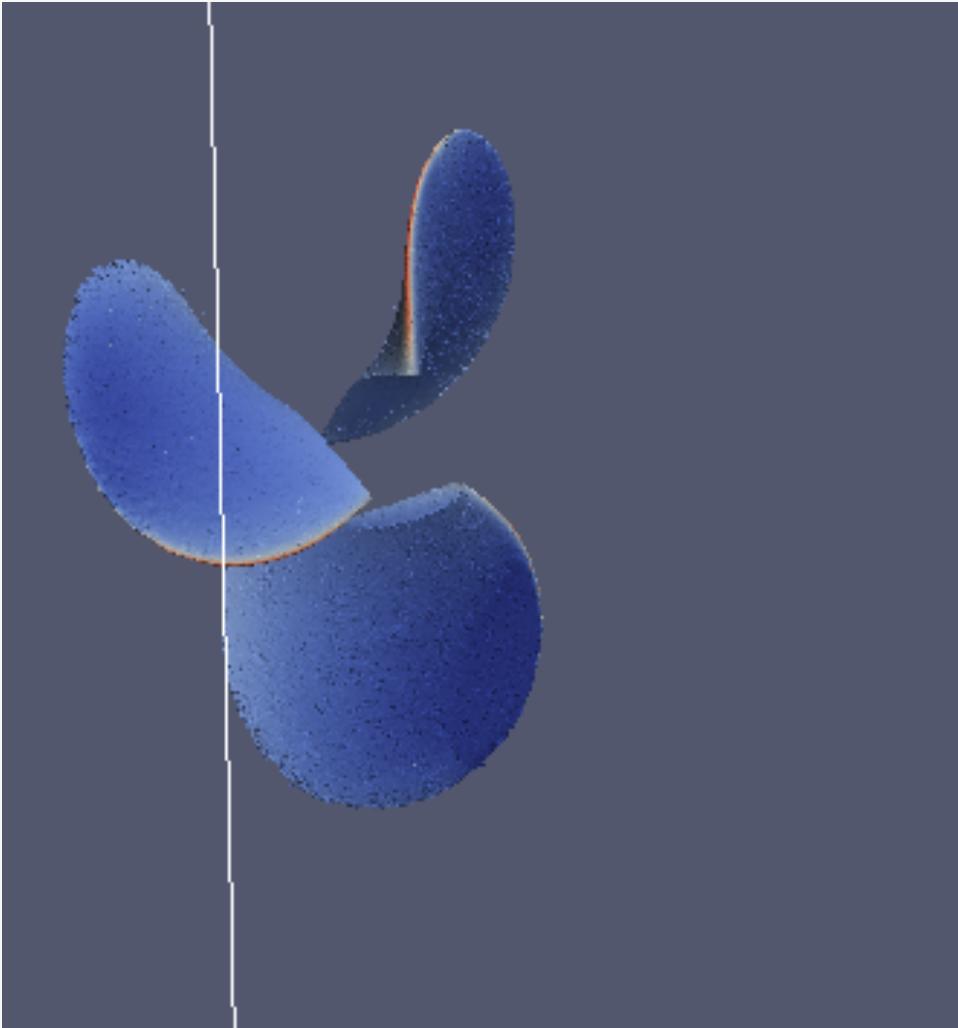
VTK:

- Anything's possible if you try hard enough

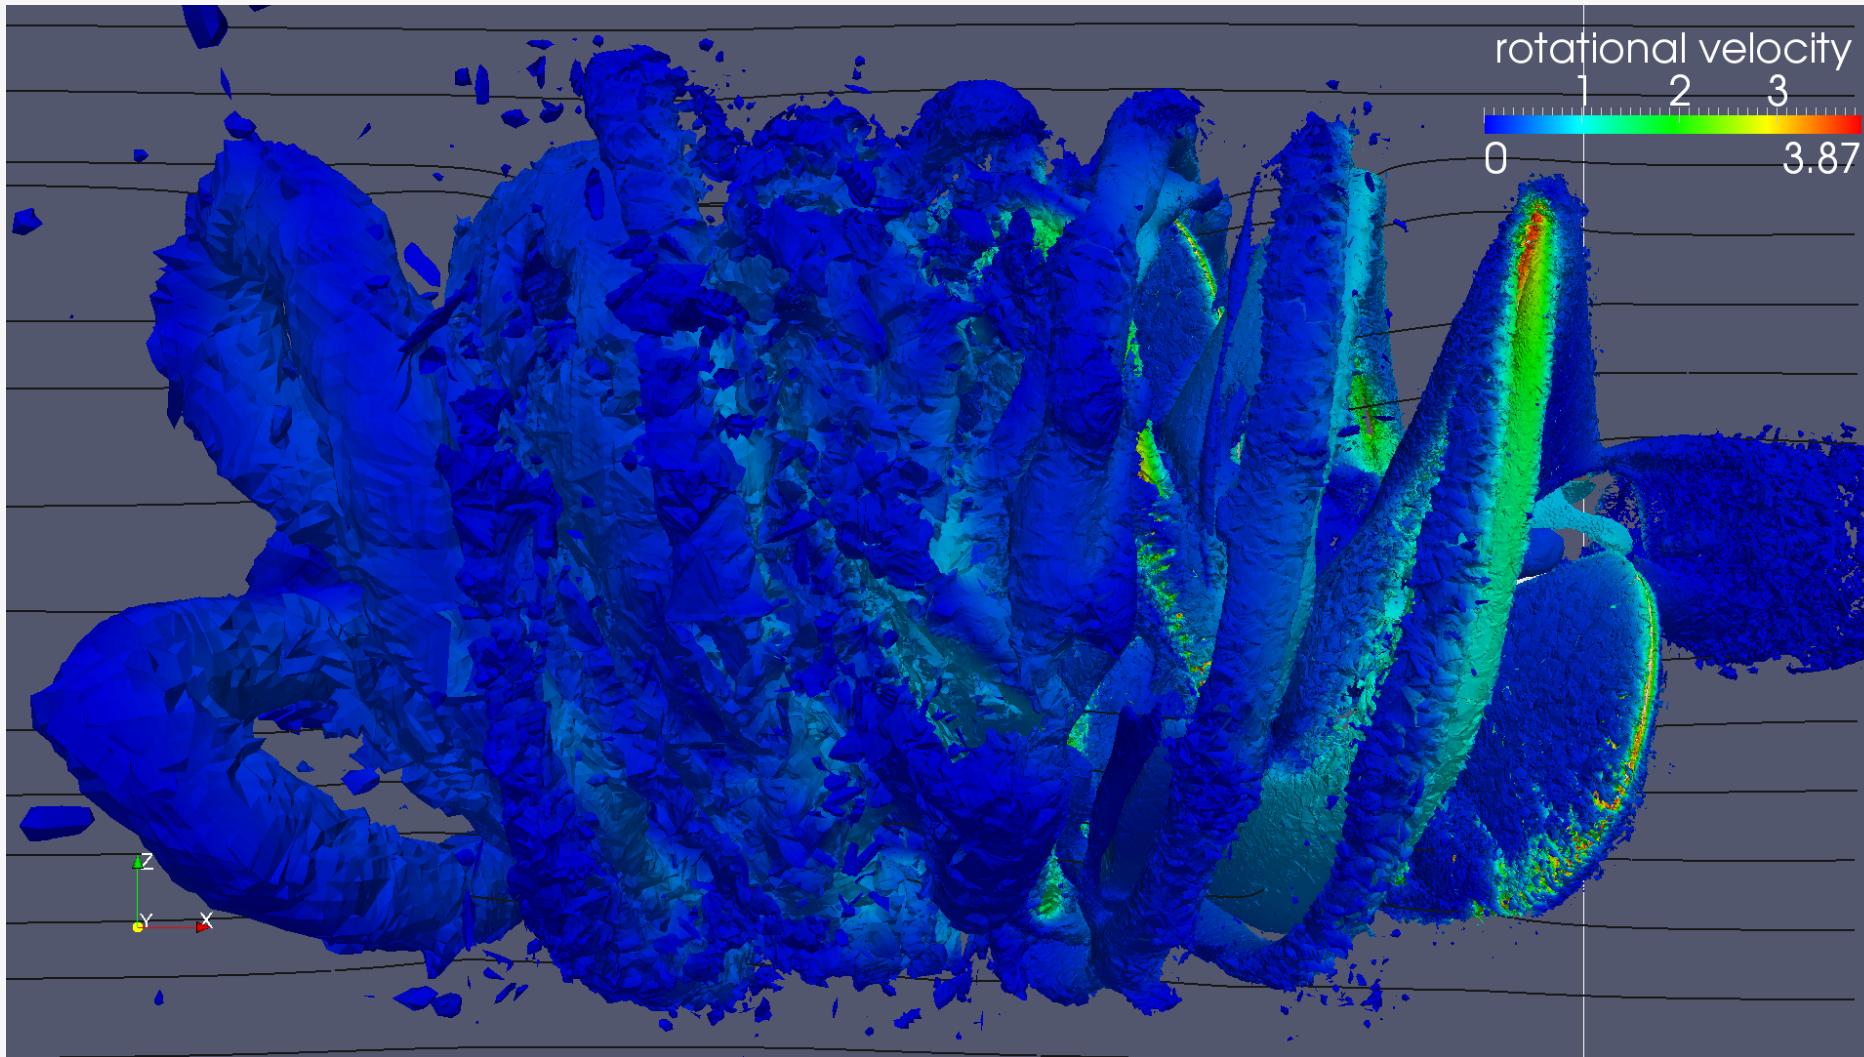


Visualization for Debugging

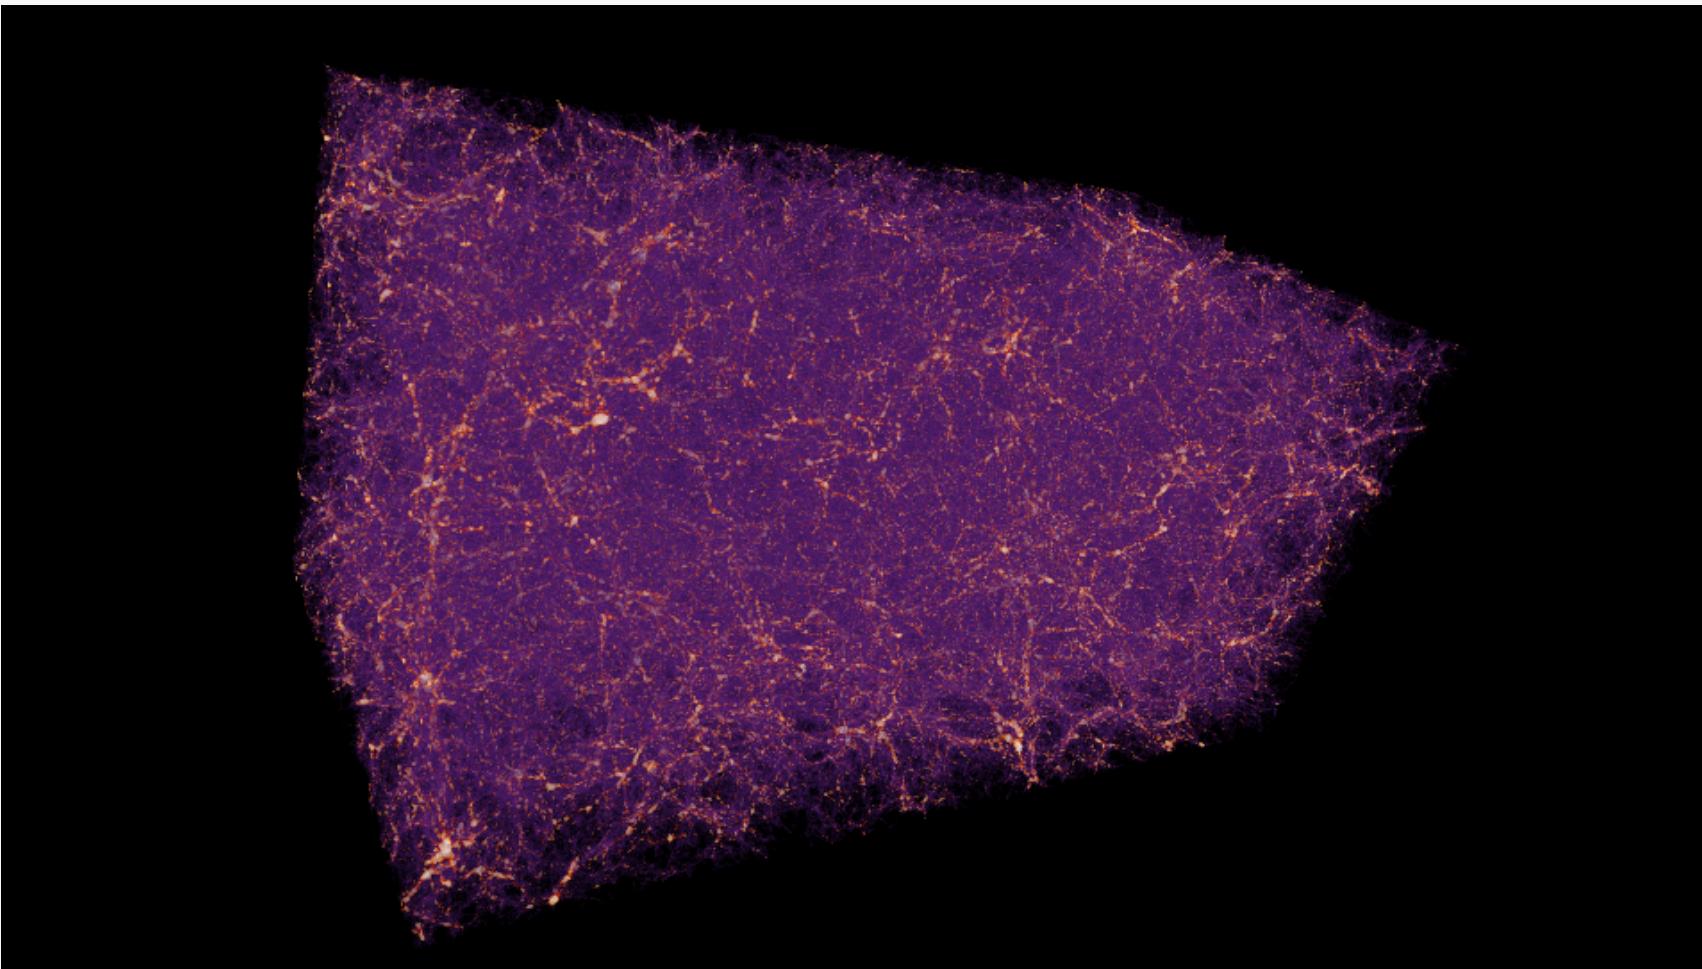
Visualization for Debugging



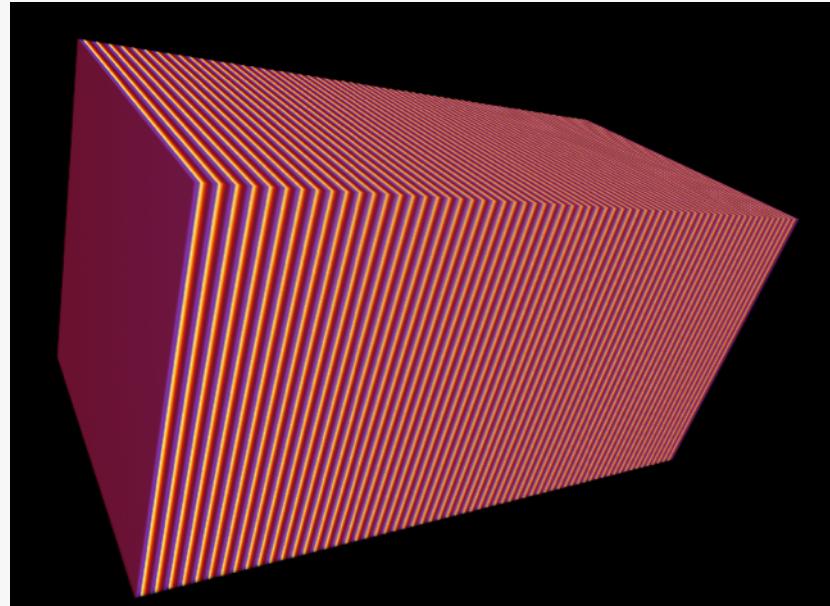
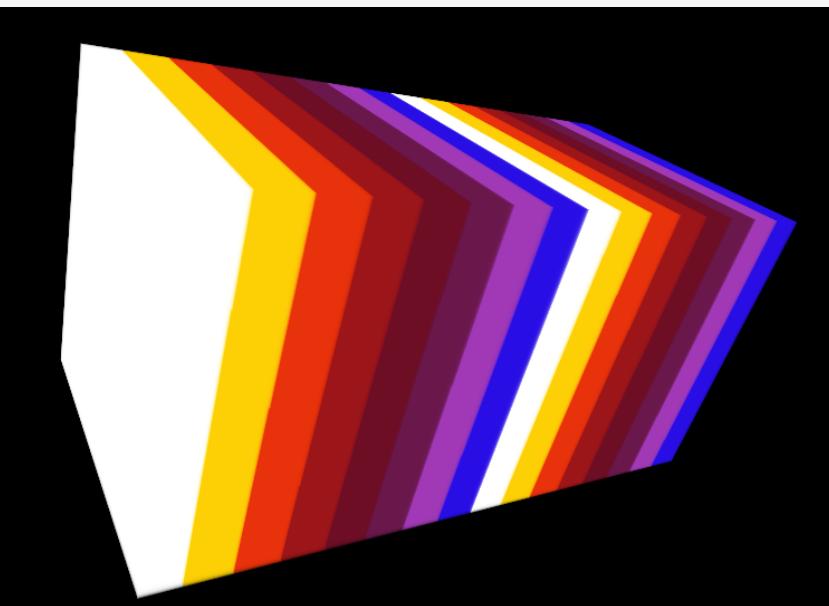
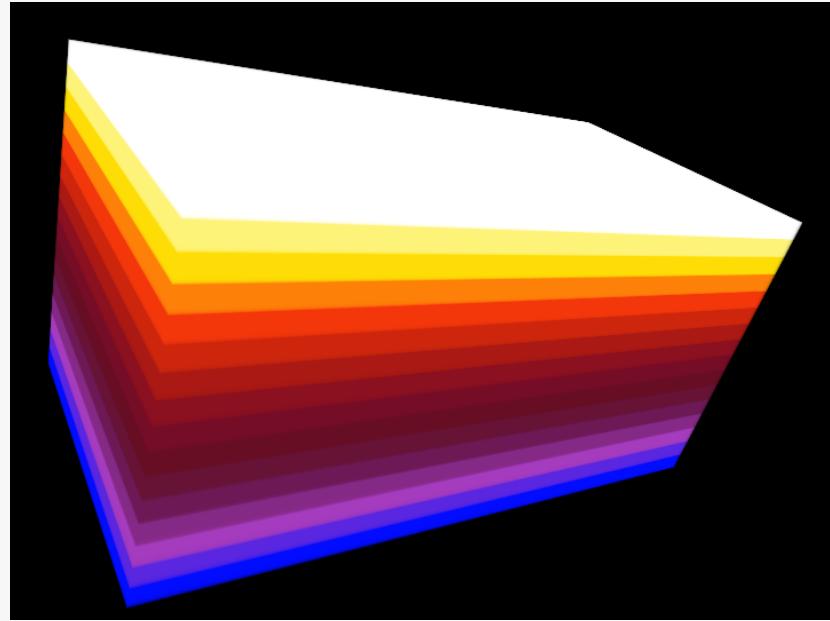
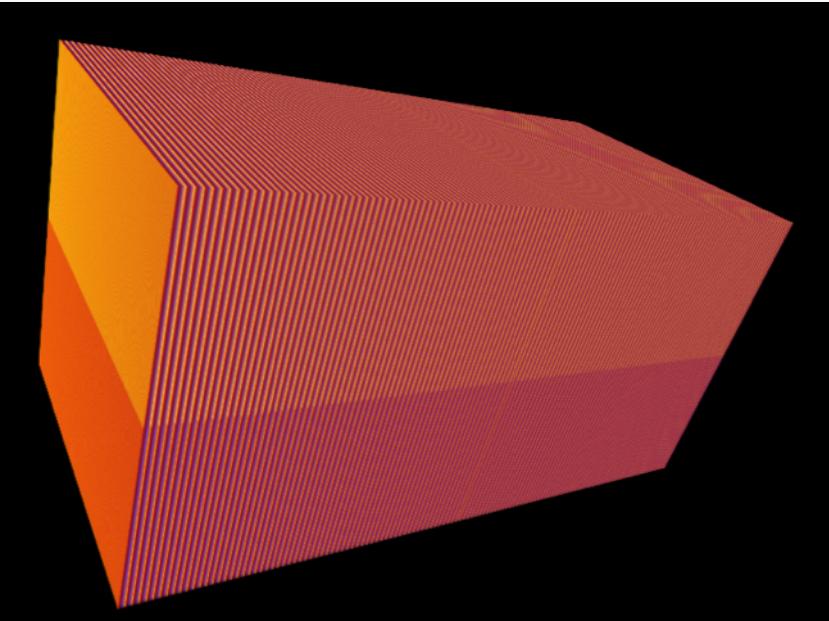
Visualization for Debugging



Visualization as Diagnostics: Color by Thread ID



Visualization as Diagnostics: Color by Thread ID



Defining Workflows with Science Teams

Investigating Tooth Enamel Fractures



Large-Scale Computing and Visualization on the Connectomes of the Brain

Objectives:

- development of imaging and analytical pipelines for full mammalian brains at the level of individual cells, axons and blood vessels
- integration on large-scale computing systems

Imaging technique:

- X-Ray extended tomography (or Mosaic Tomography) with 1micron resolution done at the beamline 32-ID-C on the Advanced Photon Source

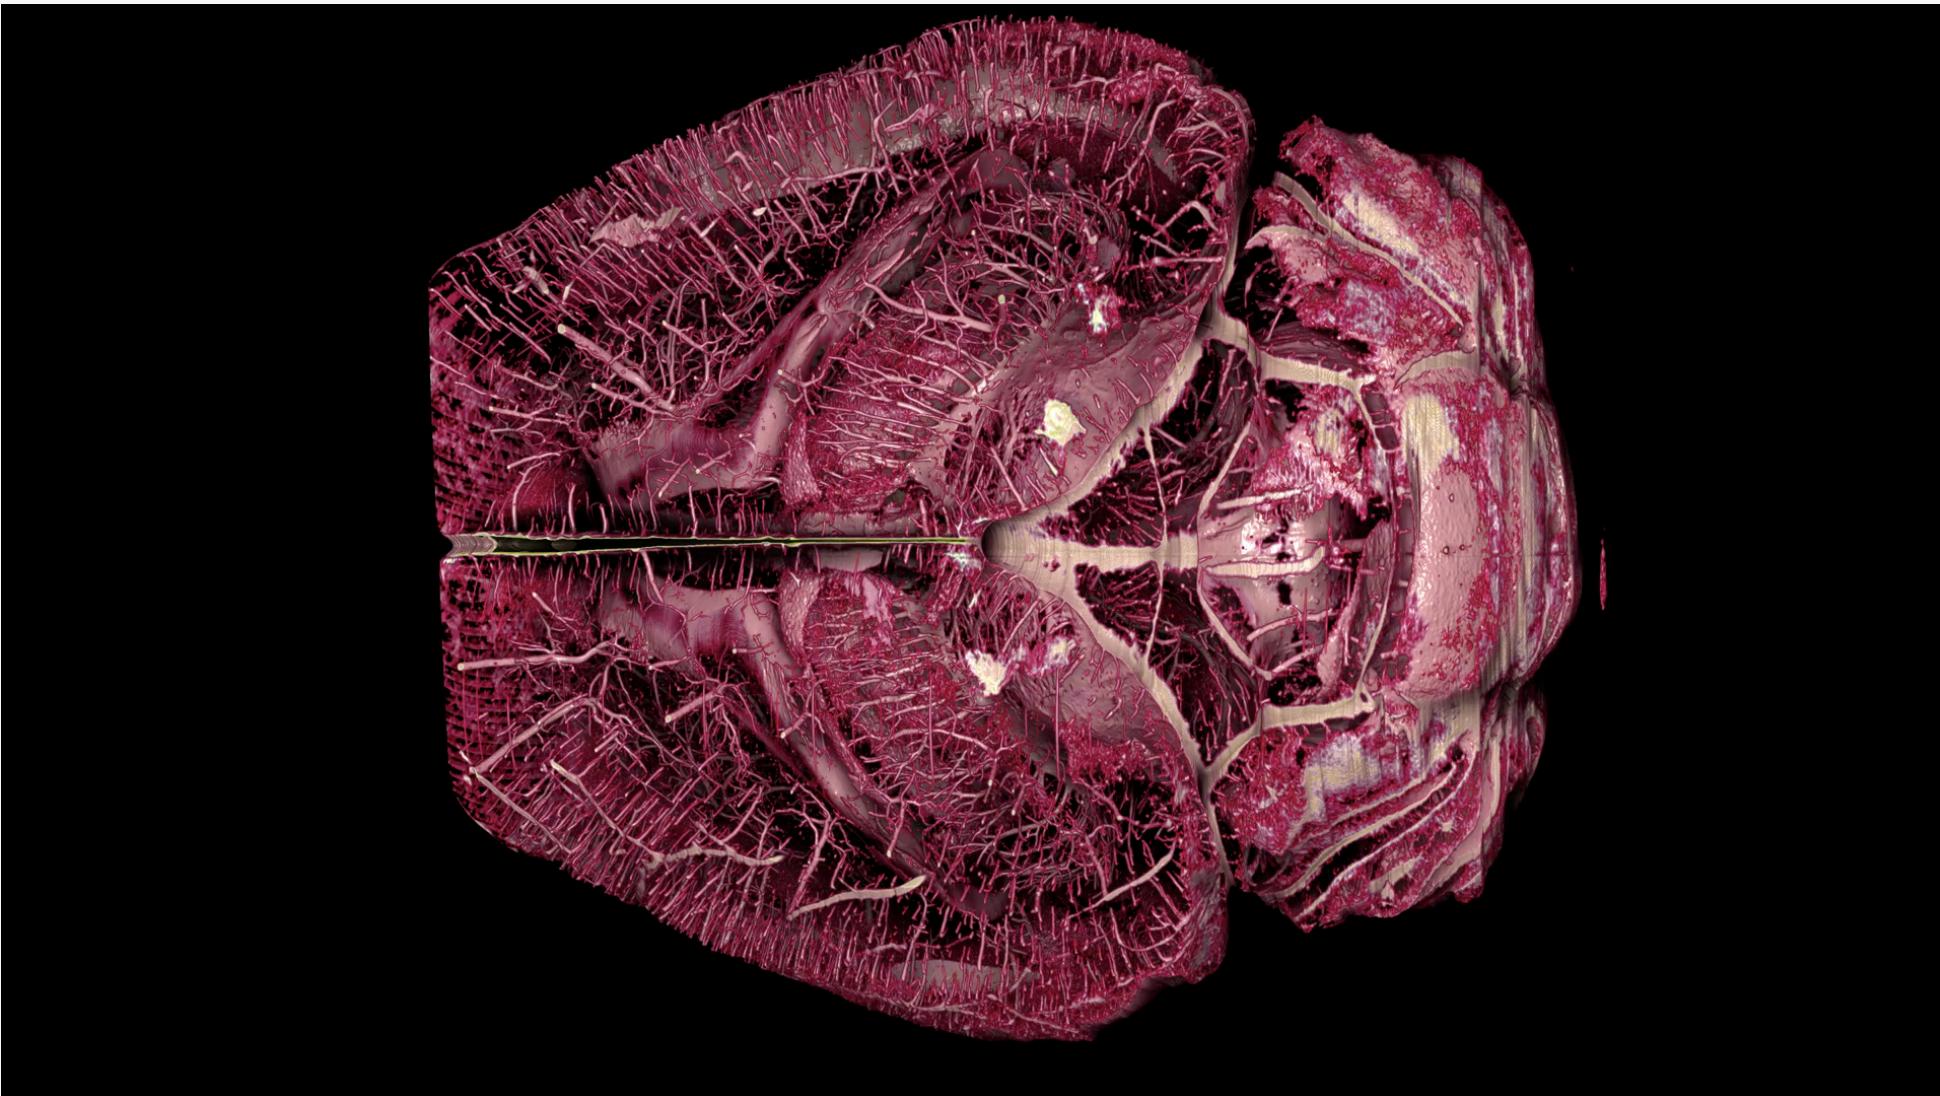
Segmentation :

- Tensor flow based segmentation to extract features like cell bodies, myelinated axons and blood vessels



Science: Narayanan (Bobby) Kasthuri and team
Slide courtesy Rafael Vescovi, Hanyu Li

Large-Scale Computing and Visualization on the Connectomes of the Brain



In Situ Visualization and Analysis

The Need of *In Situ* Analysis and Visualization

Research challenges for enabling scientific knowledge discovery at extreme-scale concurrency

Widening gap between FLOPs and I/O capacity

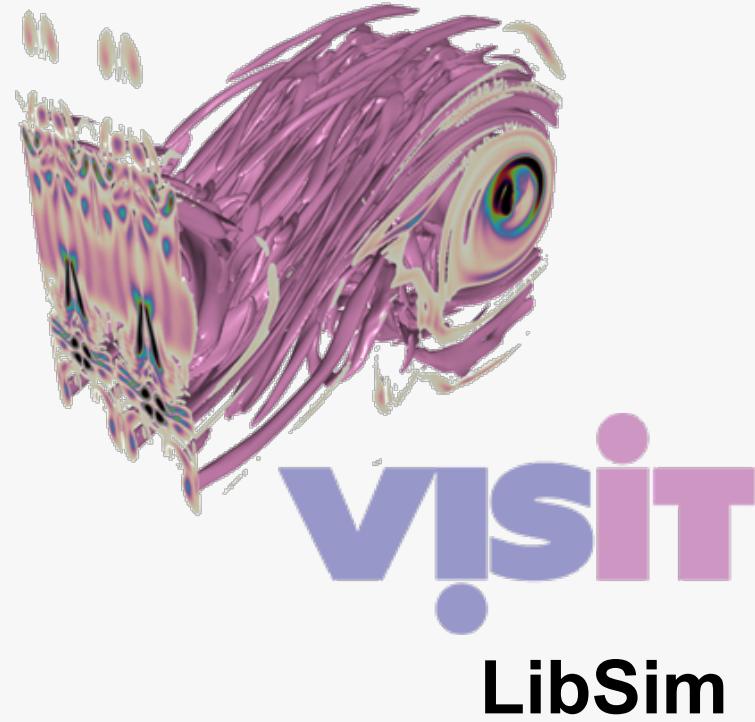
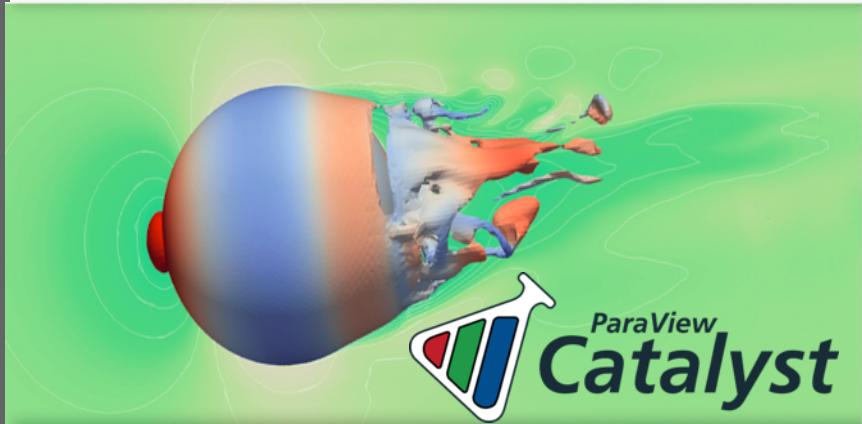
- will make full-resolution, I/O-intensive post hoc analysis prohibitively expensive, if not impossible.

Slides courtesy SENSEI in situ project:

www.sensei-insitu.org



Multiple in-situ infrastructures



ADIOS

Can We....

Enable use of any in situ framework?

Develop analysis routines that are portable between codes?

Make it easy to use?

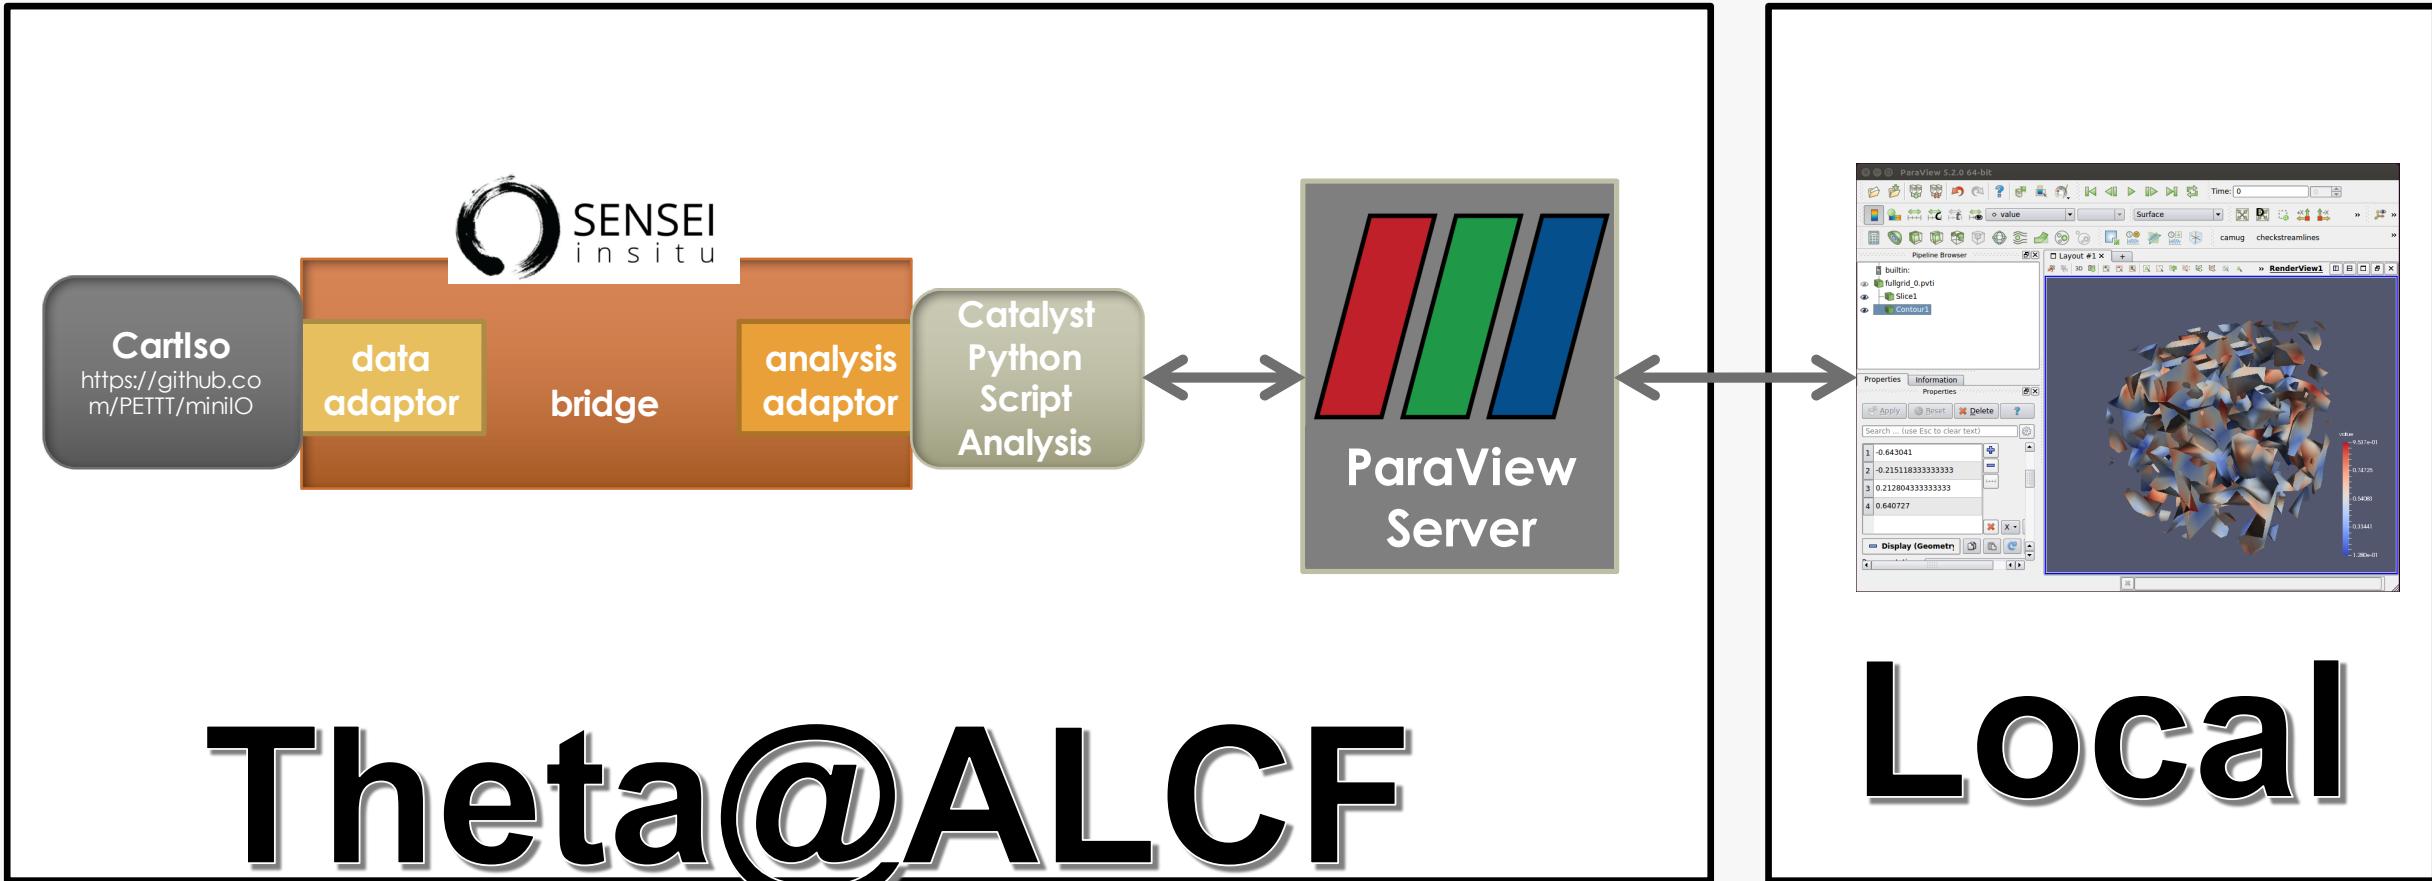
OUR APPROACH

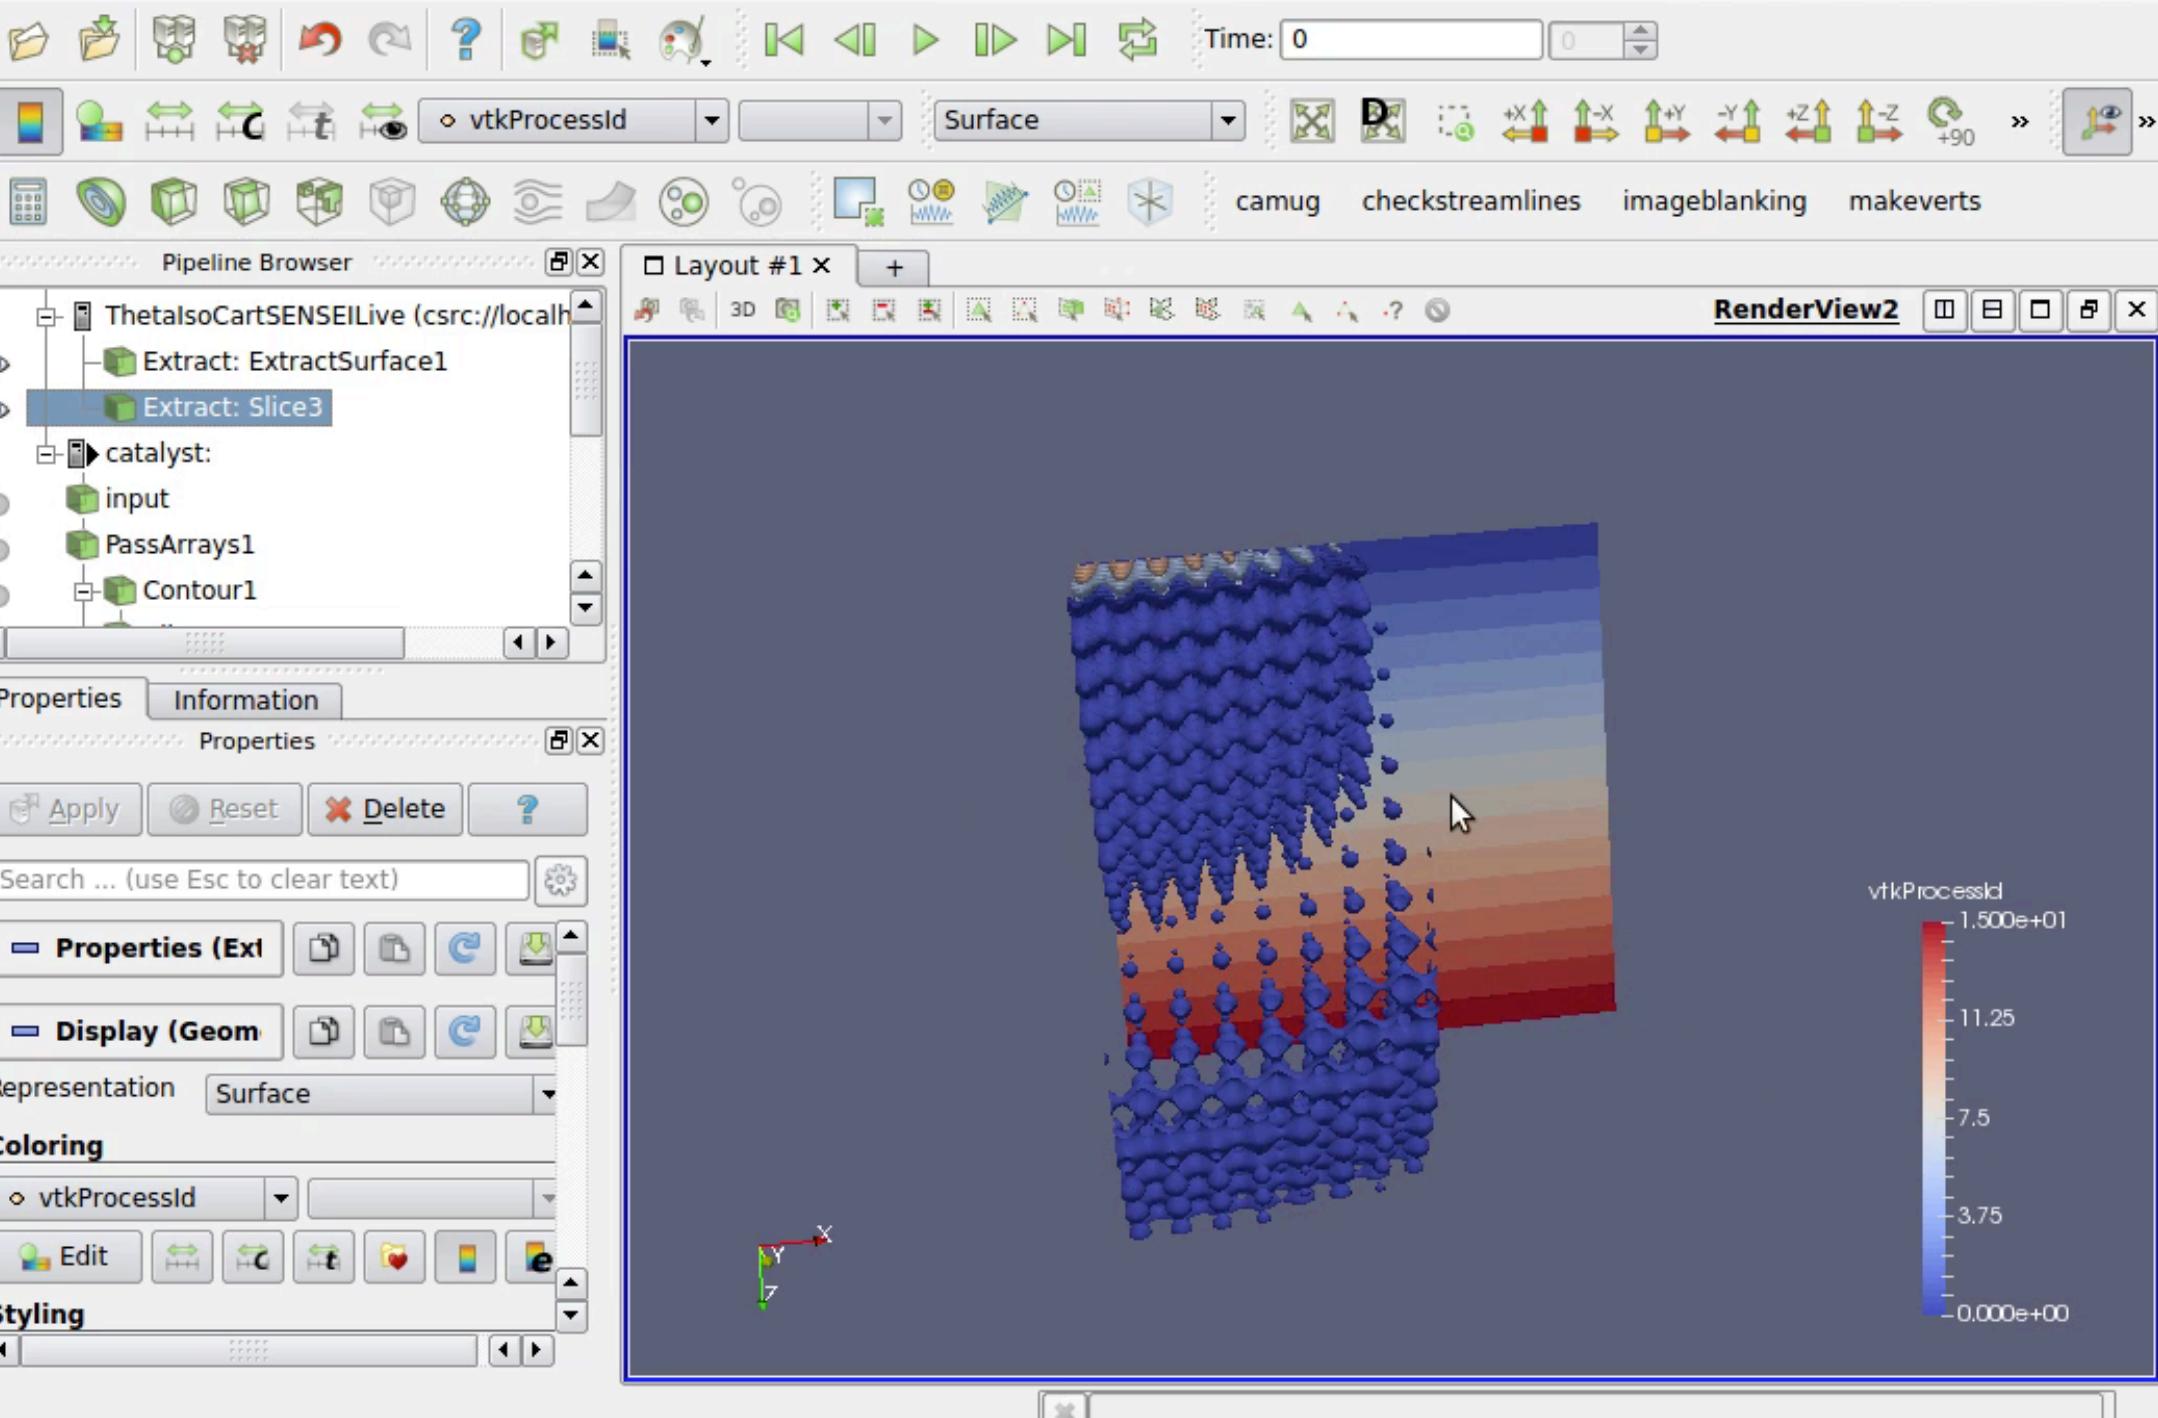
Data model – to pass data between
Simulation & Analysis

API – for instrumenting simulation and
analysis codes



Miniapp instrumentation with SENSEI





Interactive *In Situ* Visualization for MD simulations

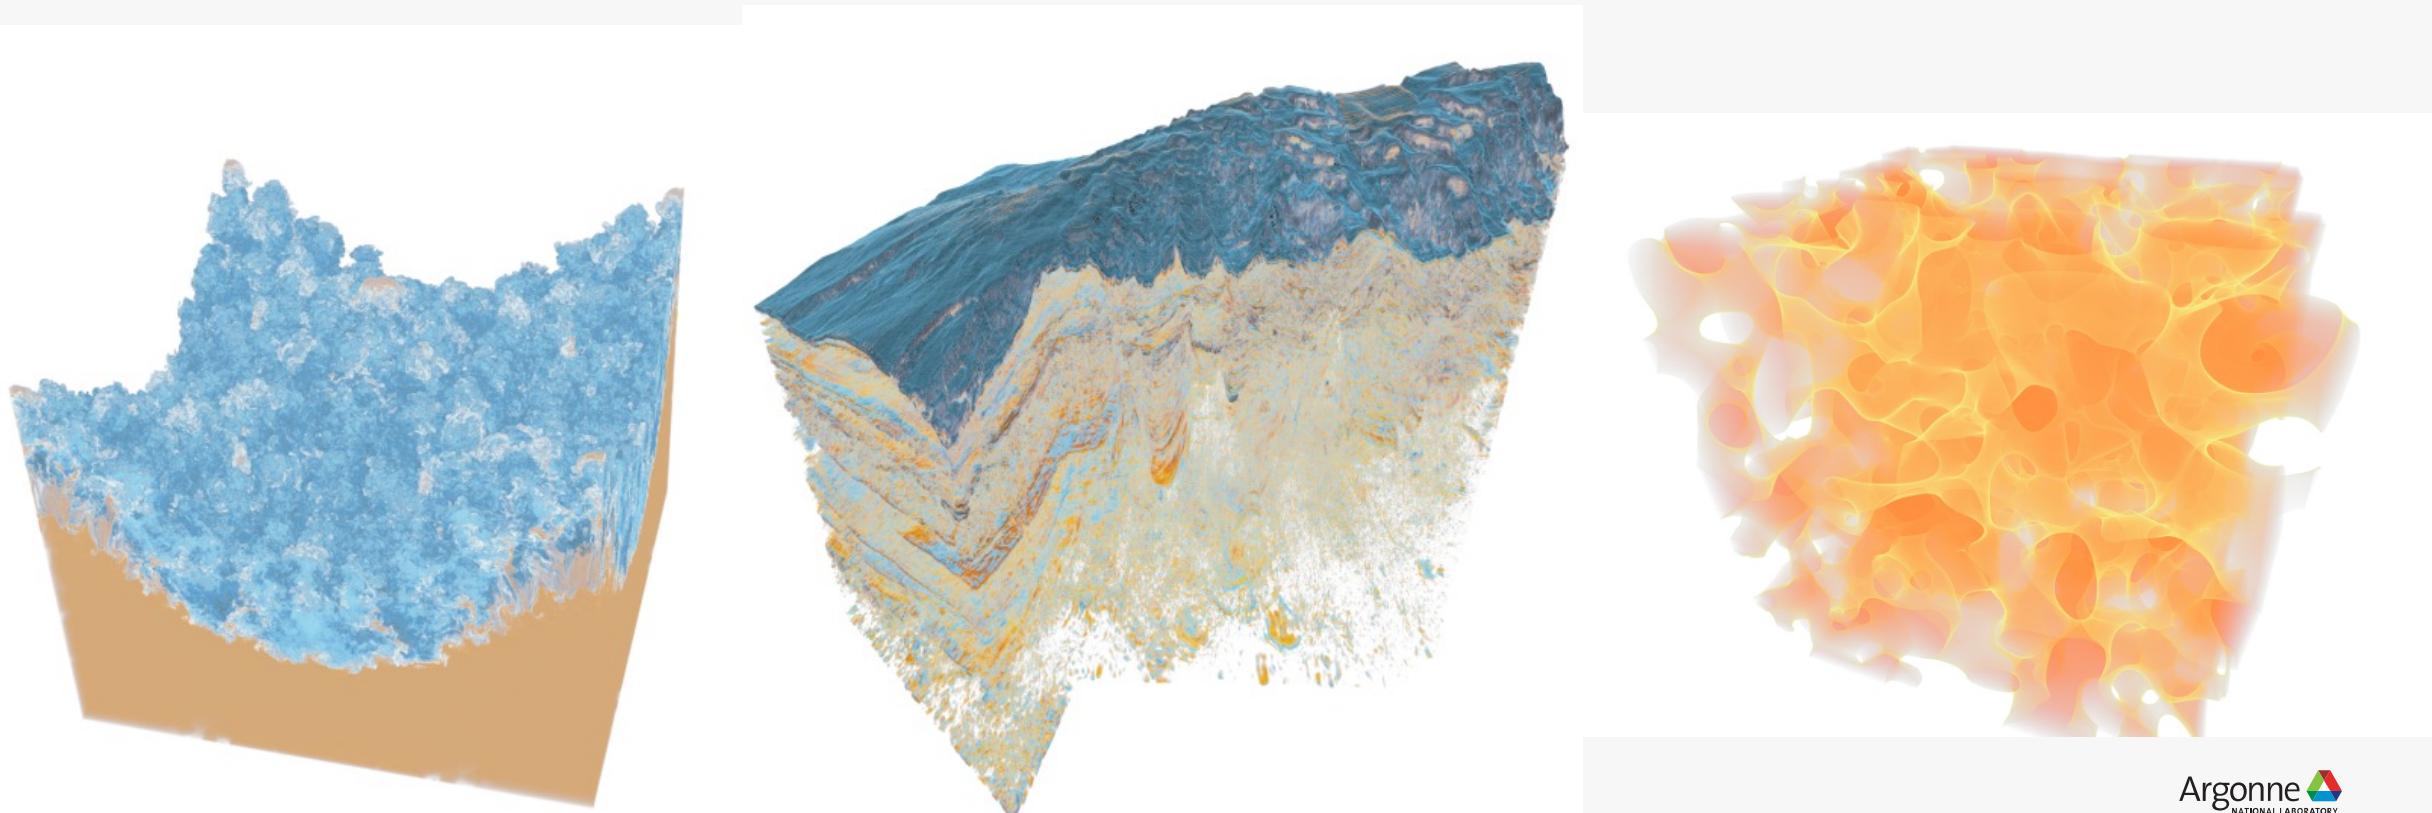
OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion



OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion



OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion



[Wald et al. '15]

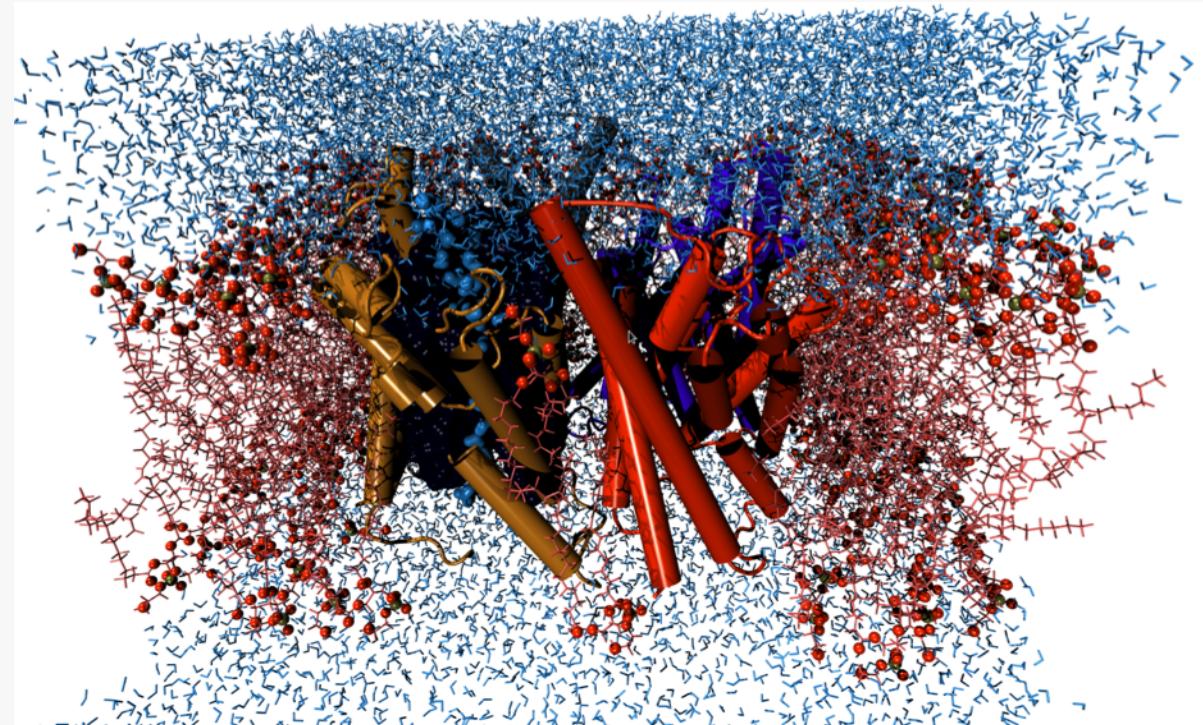
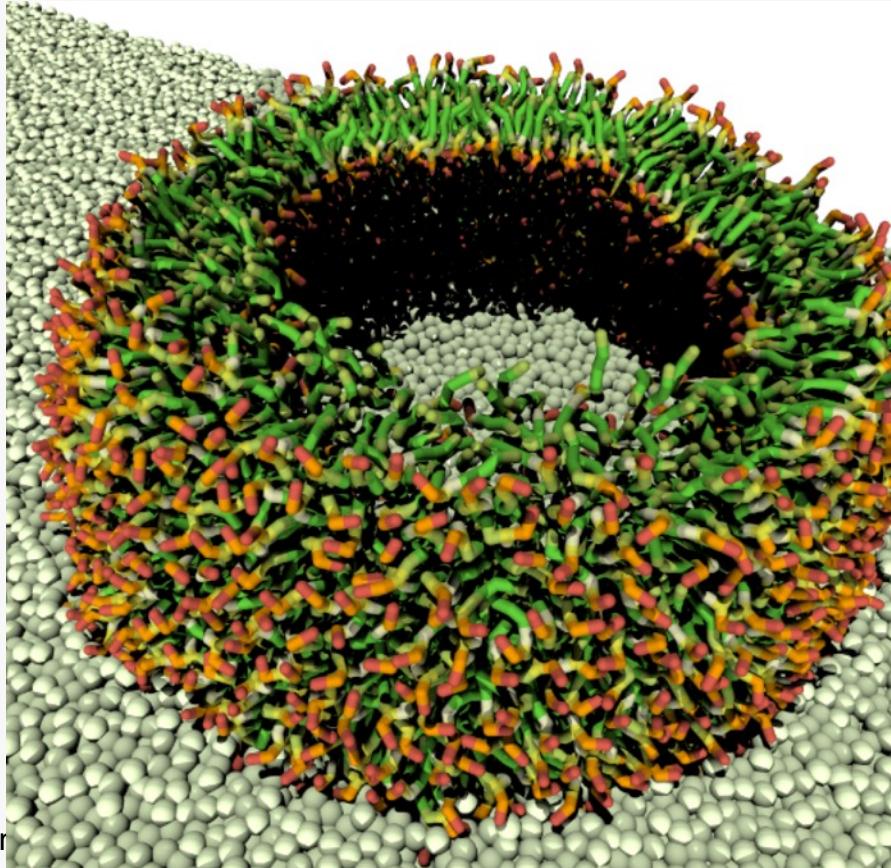
OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)



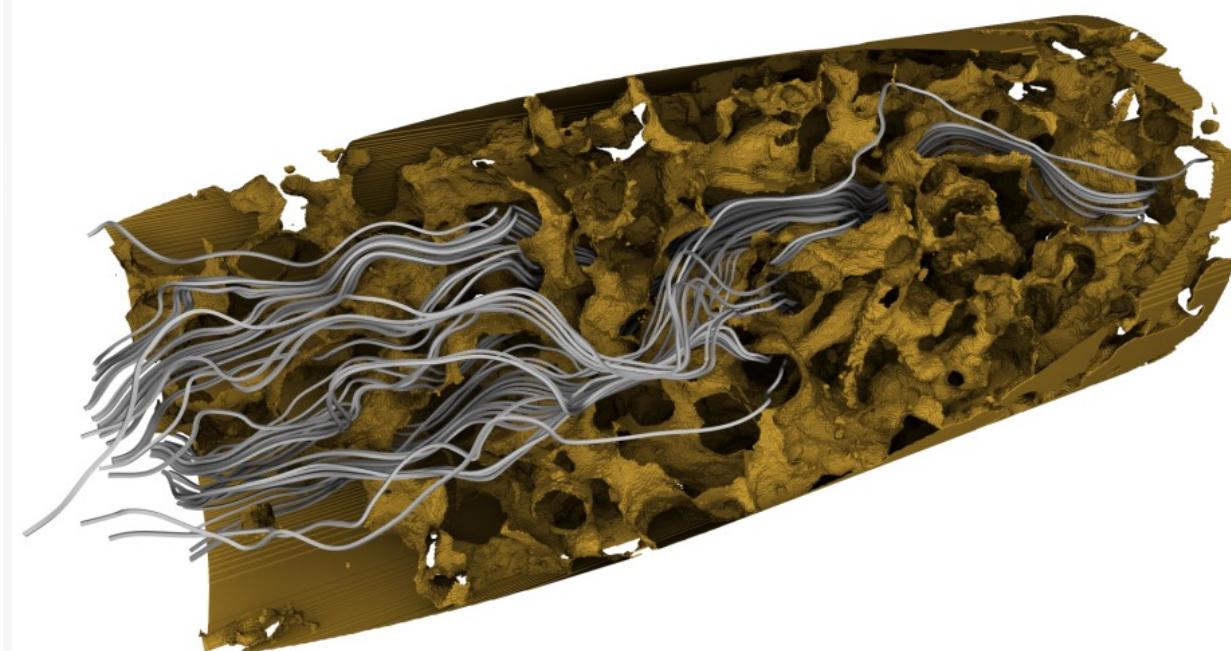
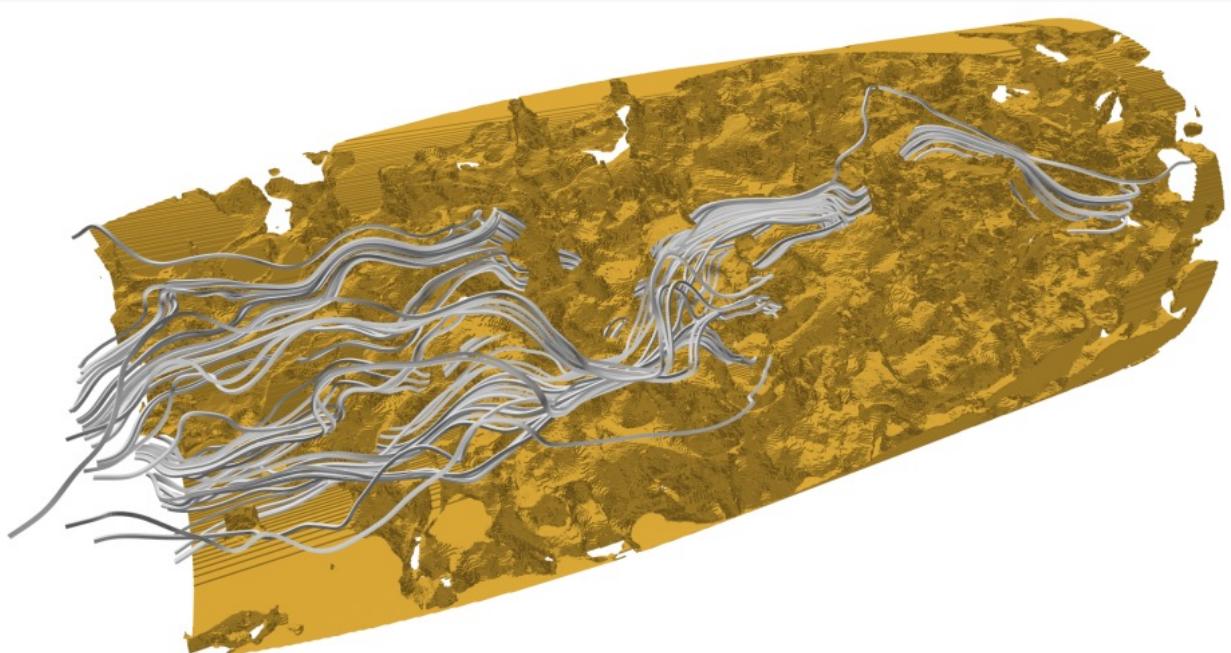
OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion



OSPRay

Slide courtesy OSPRay team @ Intel

Wald, Ingo, Gregory P. Johnson, J. Amstutz, Carson Brownlee, Aaron Knoll, J. Jeffers, J. Günther, and P. Navratil. "OSPRay-A CPU Ray Tracing Framework for Scientific Visualization." IEEE transactions on visualization and computer graphics 23, no. 1 (2017): 931-940.

Ray tracer for interactive scientific visualization-style rendering

- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion

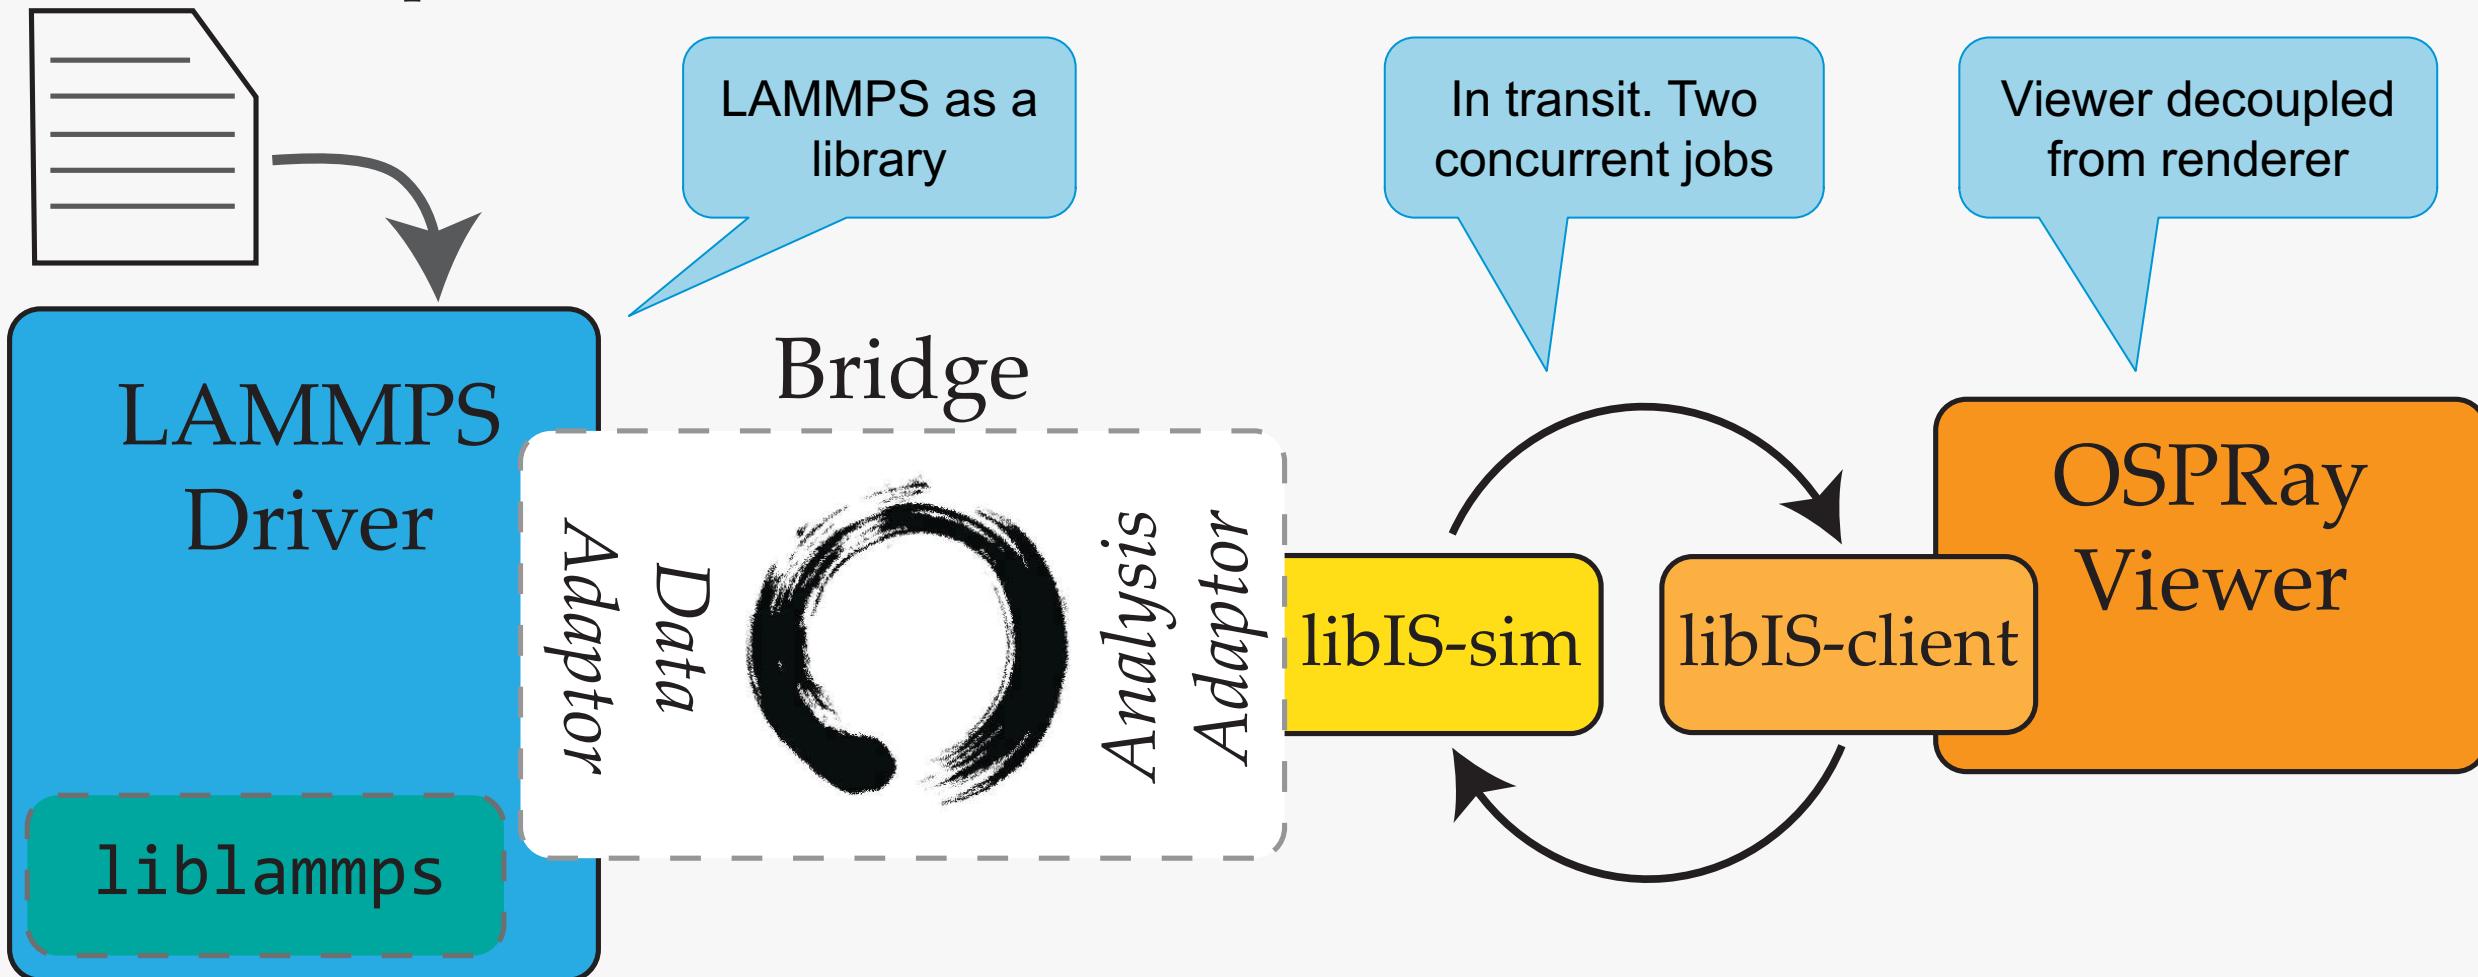
Free & open source: Apache 2.0 License

- <http://ospray.org/>

Built on top of Embree, extensive use of ISPC for vectorization

LAMMPS instrumentation with SENSEI and OSPRay

LAMMPS Input File



Callback function from LAMMPS (every timestep)

```
void LAMMPS_CALLBACK(void *ptr, bigint ntimestep,
                     int nlocal, int *id, double **x, double **f)
{
    Info *info = (Info *) ptr;

    // extents
    double boxxlo = *((double *) lammps_extract_global(info->lmp, "boxxlo"));
    double boxxhi = *((double *) lammps_extract_global(info->lmp, "boxxhi"));
    double boxylo = *((double *) lammps_extract_global(info->lmp, "boxylo"));
    double boxyhi = *((double *) lammps_extract_global(info->lmp, "boxyhi"));
    double boxzlo = *((double *) lammps_extract_global(info->lmp, "boxzlo"));
    double boxzhi = *((double *) lammps_extract_global(info->lmp, "boxzhi"));

    // get pointer to atom types
    int *type = (int *) lammps_extract_atom(info->lmp, "type");

    // update SENSEI bridge
    bridge::Set_data(nlocal, id, type, x, boxxlo, boxylo, boxzlo, boxxhi, boxyhi, boxzhi);

    // visualize
    bridge::Execute();
}
```

XYZ atom coords from LAMMPS

get atom types from LAMMPS

Visualize

Update SENSEI bridge

File Edit View Search Terminal Help

Pair	0.14076	0.1448	0.14897	0.7	86.63
Neigh	0	0	0	0.0	0.00
Comm	0.00045204	0.0046191	0.0085671	4.1	2.76
Output	0.000664	0.00082179	0.0014629	0.0	0.49
Modify	0.01554	0.01573	0.015936	0.1	9.41
Other				0.71	
					0.001185

Nlocal: 28506.7 ave 28612 max 28386 min

Histogram: 1 0 1 1 0 1 0 0 0 2

Nghost: 16954.2 ave 17068 max 16811 min

Histogram: 1 0 1 0 1 0 0 1 1 1

Neighs: 556196 ave 568688 max 532730 min

Histogram: 1 0 0 0 0 1 1 1 0 2

FullNghs: 18474.7 ave 19644 max 17332 min

Histogram: 1 0 0 0 0 1 0 1 1 2

Total # of neighbors = 3337173

Ave neighs/atom = 19.5111

Neighbor list builds = 0

Dangerous builds = 0

WARNING: One or more dynamic groups may not be updated at correct point in timestep (./fix_

WARNING: One or more atoms are time integrated more than once (./modify.cpp:275)

Setting up Verlet run ...

Unit style : metal

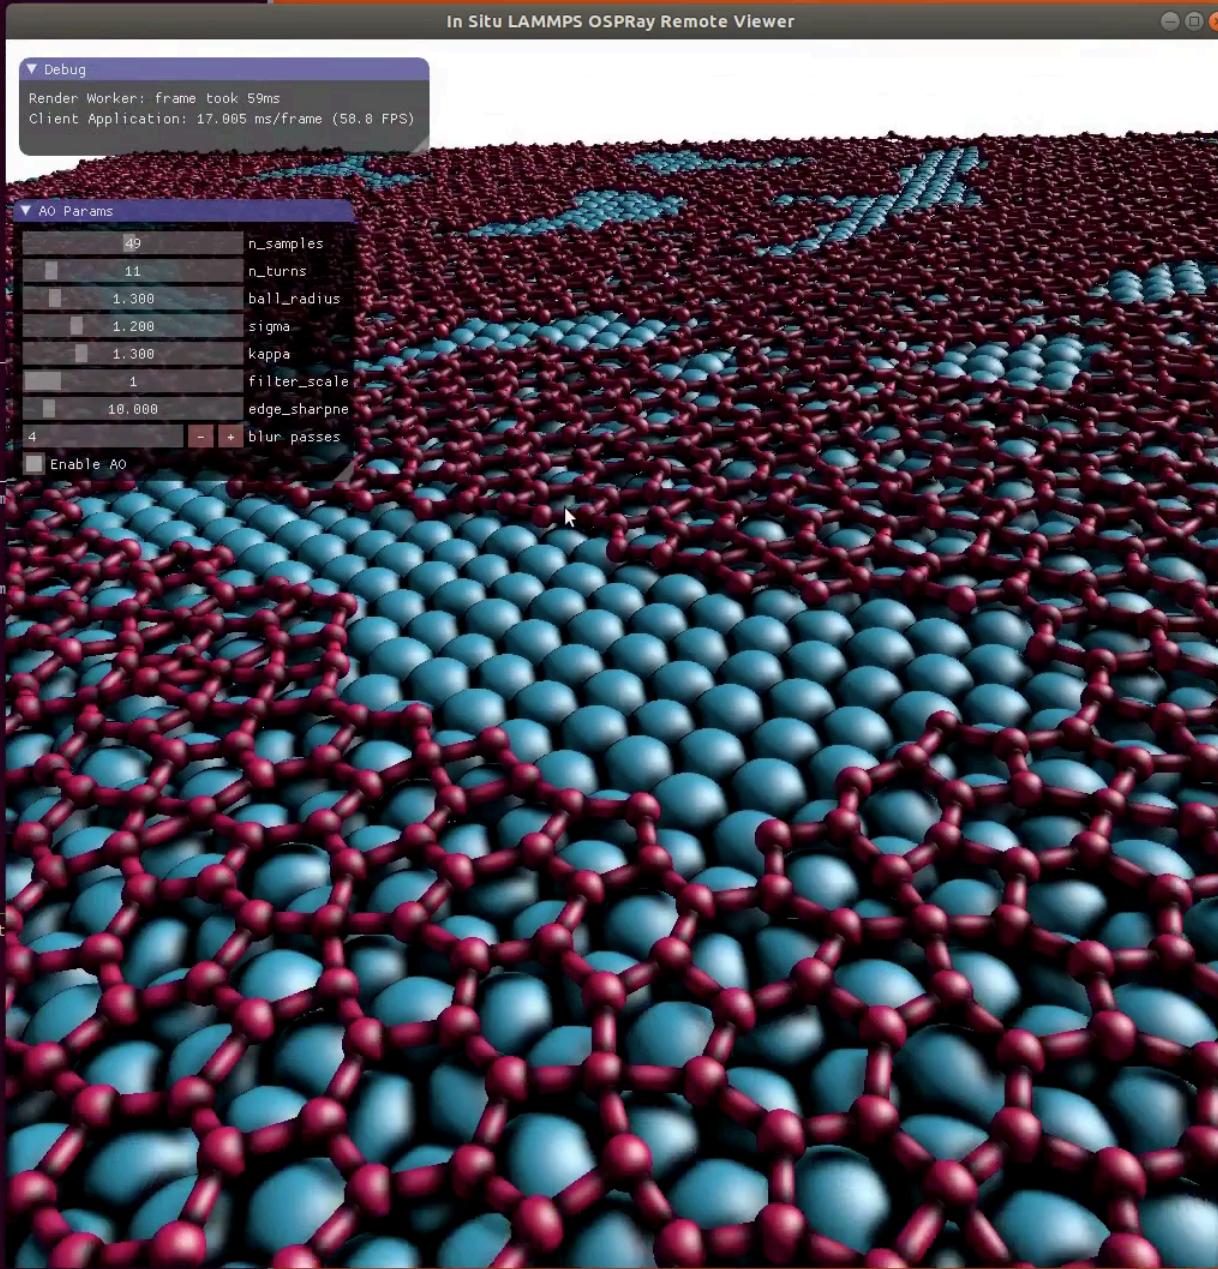
Current step : 54

Time step : 0.0005

```
will@sci1952:~/repos/lammps_sensei_ospray/viewer/build$ I_MPI_FABRIC=shm I_MPI_SHM_LMT=shm m
is_render_worker -sim-host localhost -sim-port 29374 -port 6910 -bond 1 2.7
OSPRay with rank 0, world size: 1
Connecting over the network to the simulation
[0] DAPL startup: RLIMIT_MEMLOCK too small
Sending connect cmd, got MPI port name from open 'tag#0$description#sci1952$port#38231$ifnam
rank 0 running on sci1952
Now listening for client on sci1952:6910
```

```
will@sci1952:~/repos/lammps_sensei_ospray/viewer/build$ ./lammps_is_viewer -server localhost
Got world bounds = [(-0.0750253, -0.0592656, -5):(282.256, 244.435, 41.3664)]
```

█

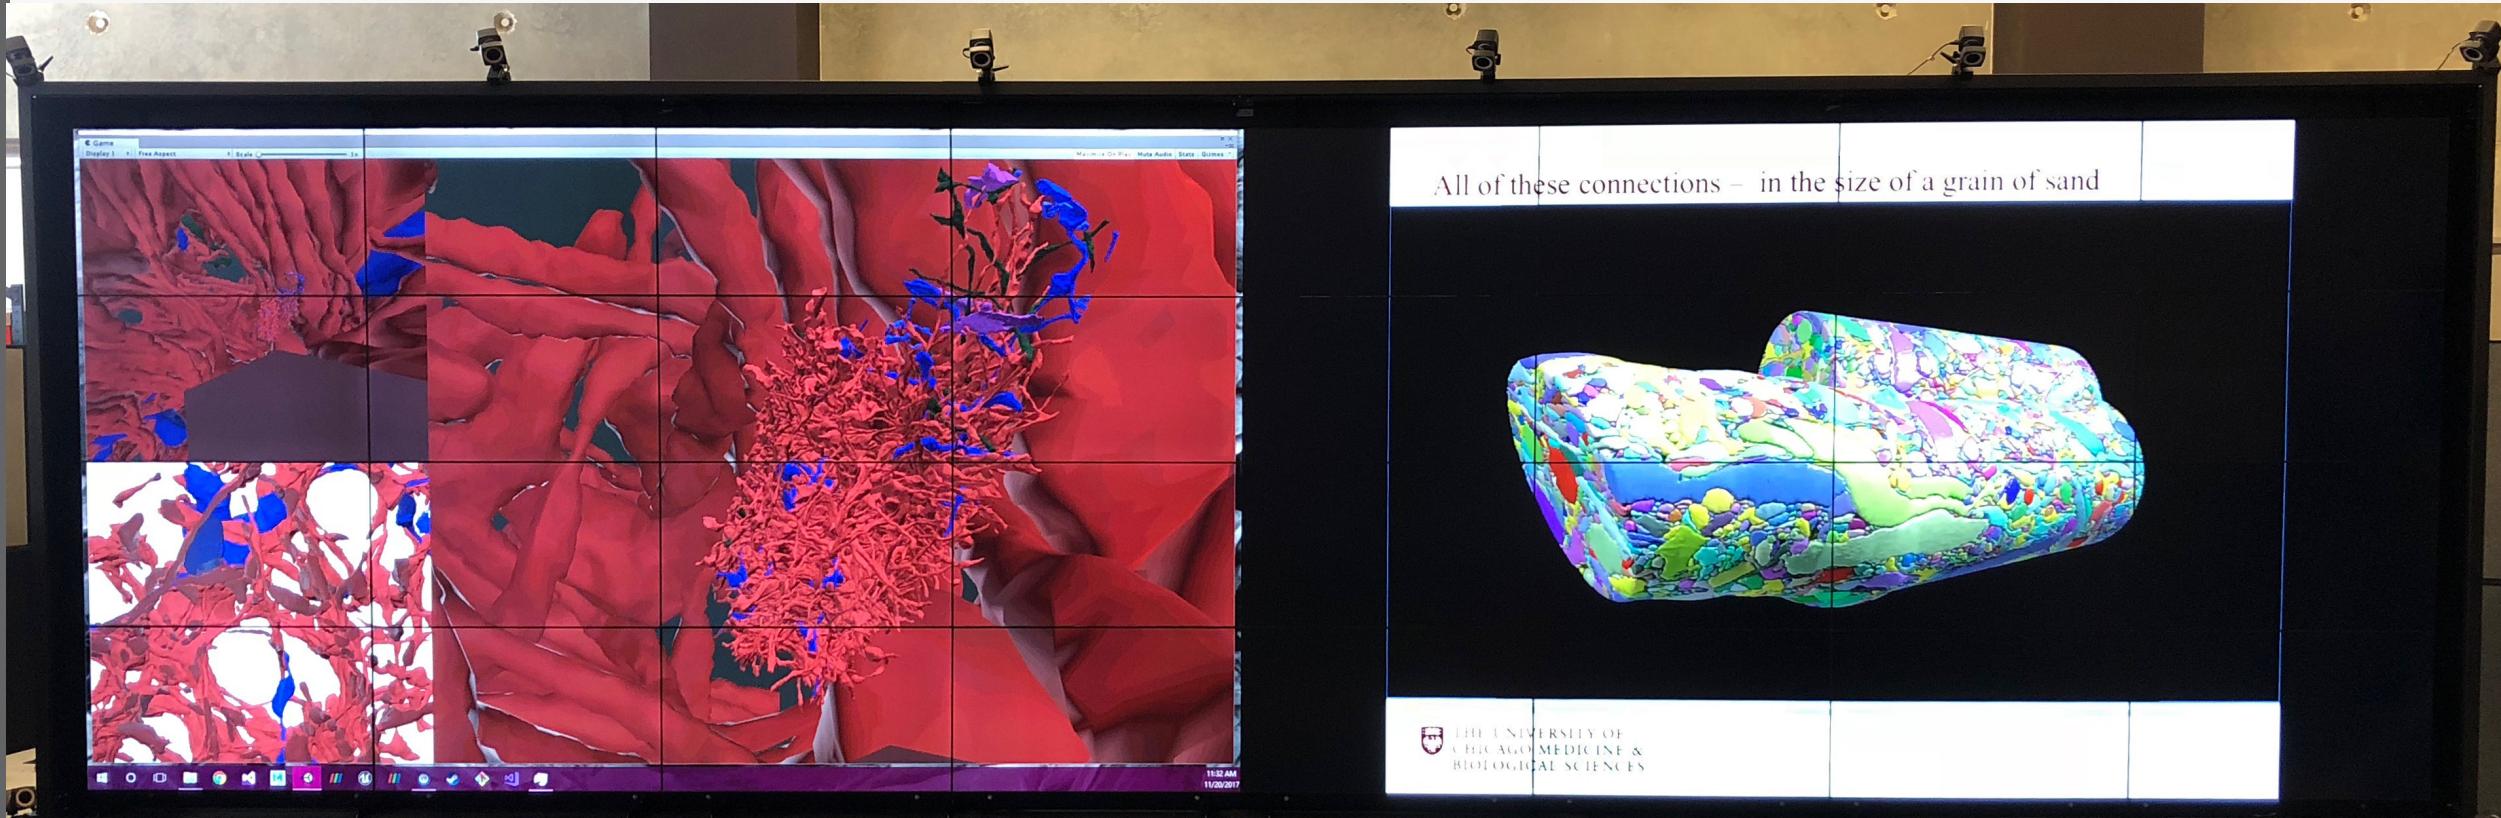


Immersive Visualization

Immersive visualization



Immersive visualization





Slide courtesy Edouard Brooks, ALCF

QUESTIONS?

Joe Insley
insley@anl.gov

Silvio Rizzi
srizzi@anl.gov

www.anl.gov

Argonne 
NATIONAL LABORATORY