

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

Anterior Cerebral

Middle
Cerebral




Aneurysm

Right Interior
Carotid Artery

Basilar

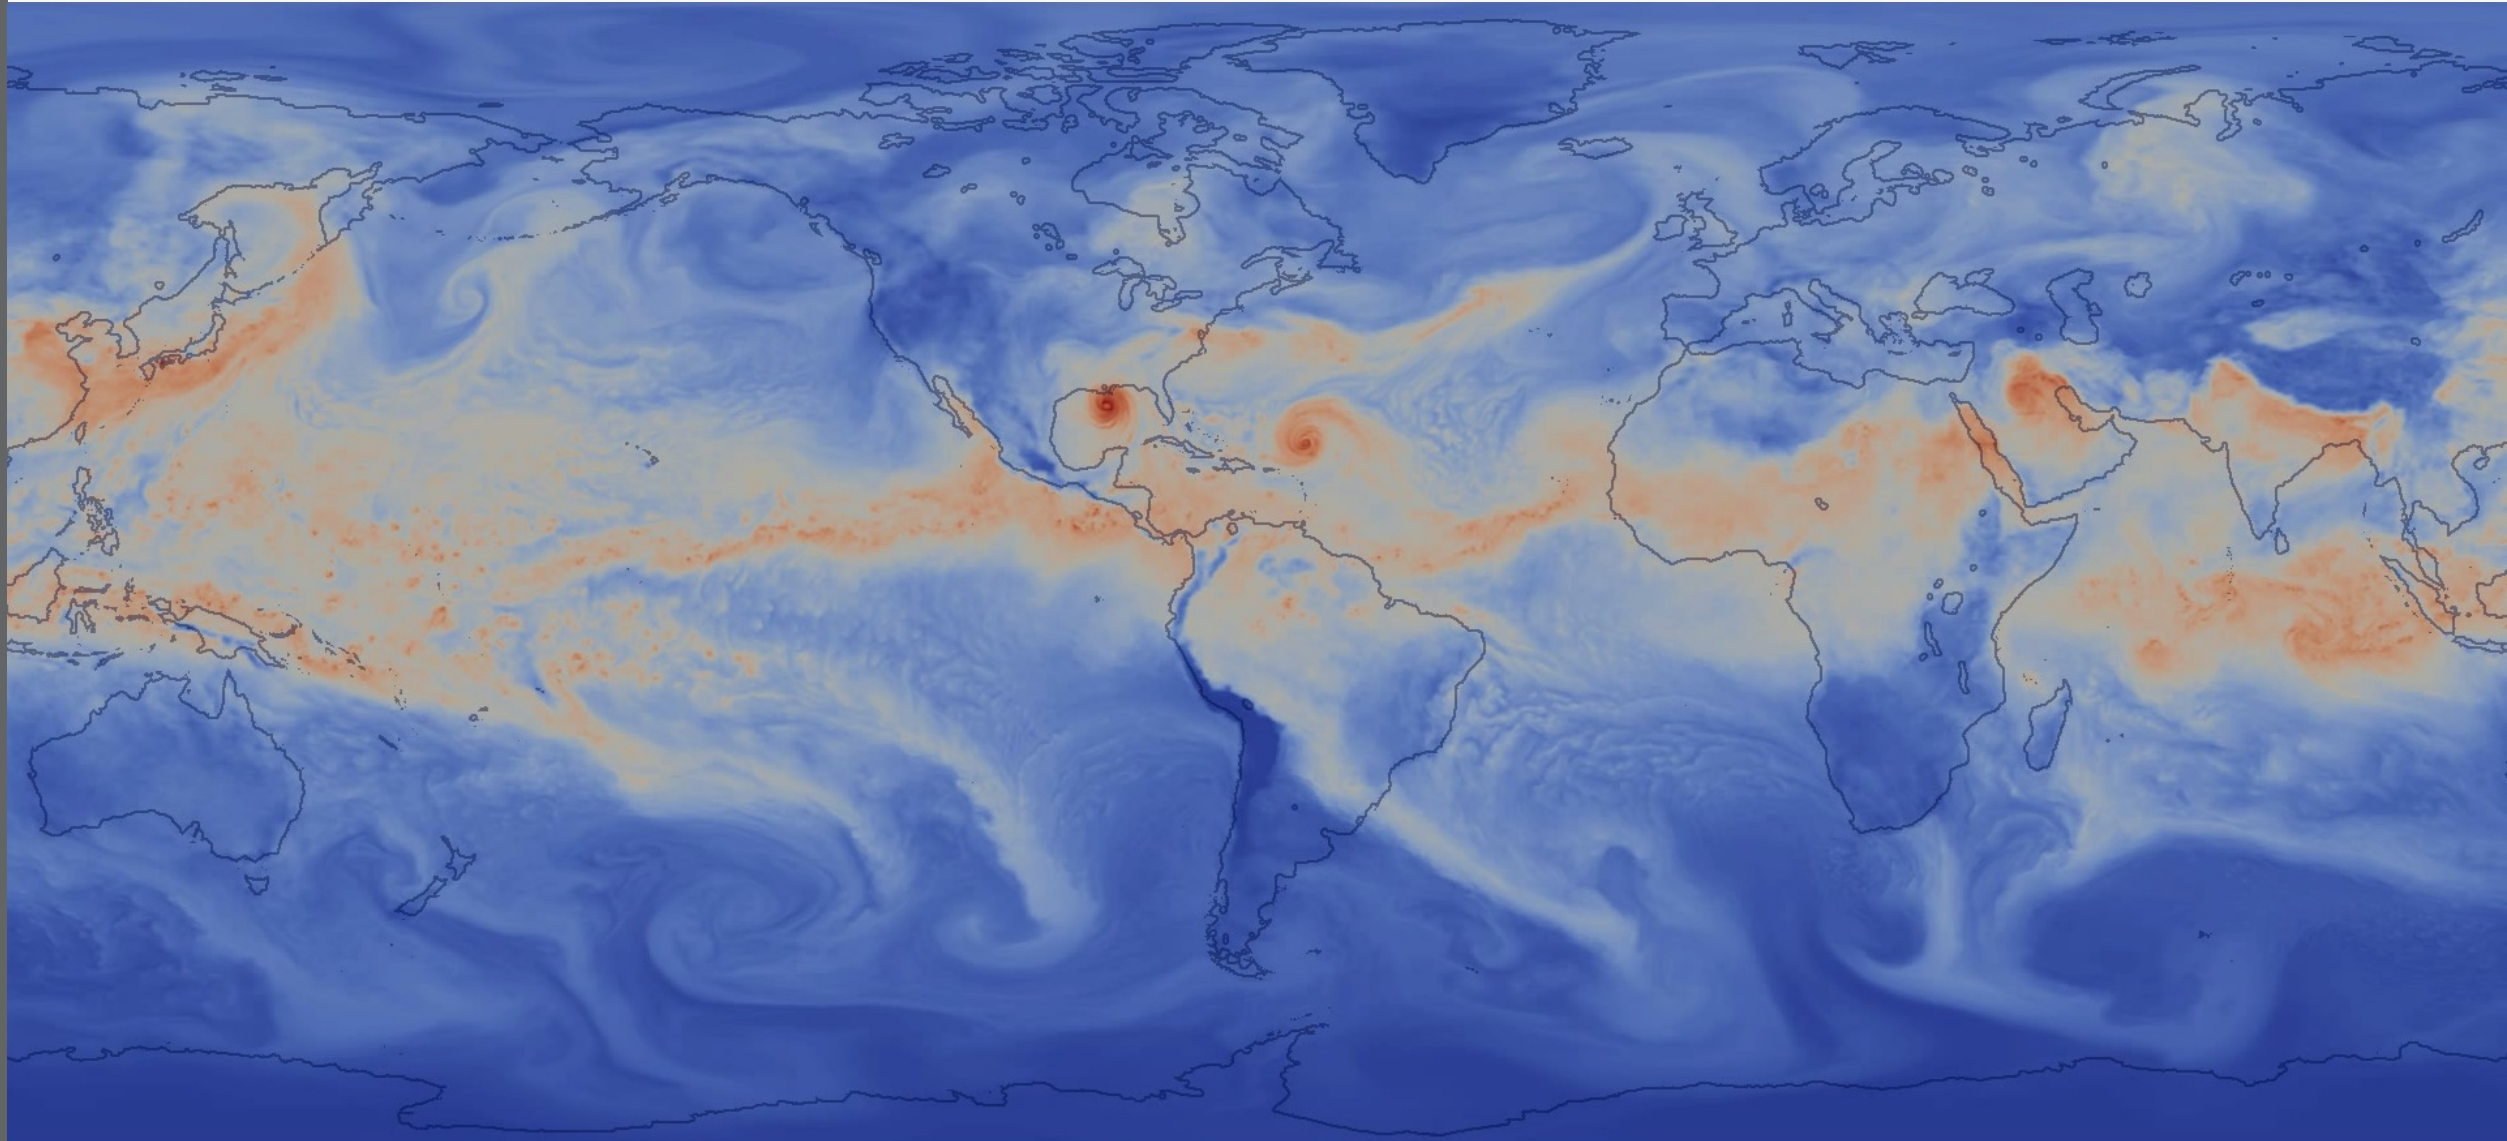
Left Interior
Carotid
Artery

Vertebral




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 Telsa 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 displays the Globus Transfer Files interface. At the top, there's a navigation bar with 'Manage Data', 'Publish', 'Groups', 'Support', and 'Account'. Below that, a secondary bar shows 'Transfer Files | Activity | Endpoints | Bookmarks | Console'. The main area is titled 'Transfer Files' and includes a prompt: 'Get Globus Connect Personal Turn your computer into an endpoint.' with three status icons. The interface is split into two panels for source and destination endpoints. The source endpoint is 'alcf#dtn_theta' with path '/projects/visualization/'. Its directory listing includes folders: DATA, fujiwara, insley, paraview, preeti, soft, srizzi, turam, and v3pkgs. The destination endpoint is 'alcf#dtn_mira' with path '/projects/visualization/'. Its directory listing includes folders: ATPESC, DATA, KOVE, SC17, cam, cooley, esias, flash_data, fmiranda, ftessier, fujiwara, kashu, marrin, mira, mlewis, preeti, sanuar, srizzi, and sshilpika. A 'Label This Transfer' field is present, with a note: 'This will be displayed in your transfer activity.' Below it, 'Transfer Settings' are listed: 'sync - only transfer new or changed files', 'delete files on destination that do not exist on source', 'preserve source file modification times', 'verify file integrity after transfer' (checked), and 'encrypt transfer'.



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

- [v13](#): shader-based volume and particle rendering

Analysis Environments

- [Matlab](#)
- [Parallel R](#)

Utilities

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

 Available on Cooley

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)

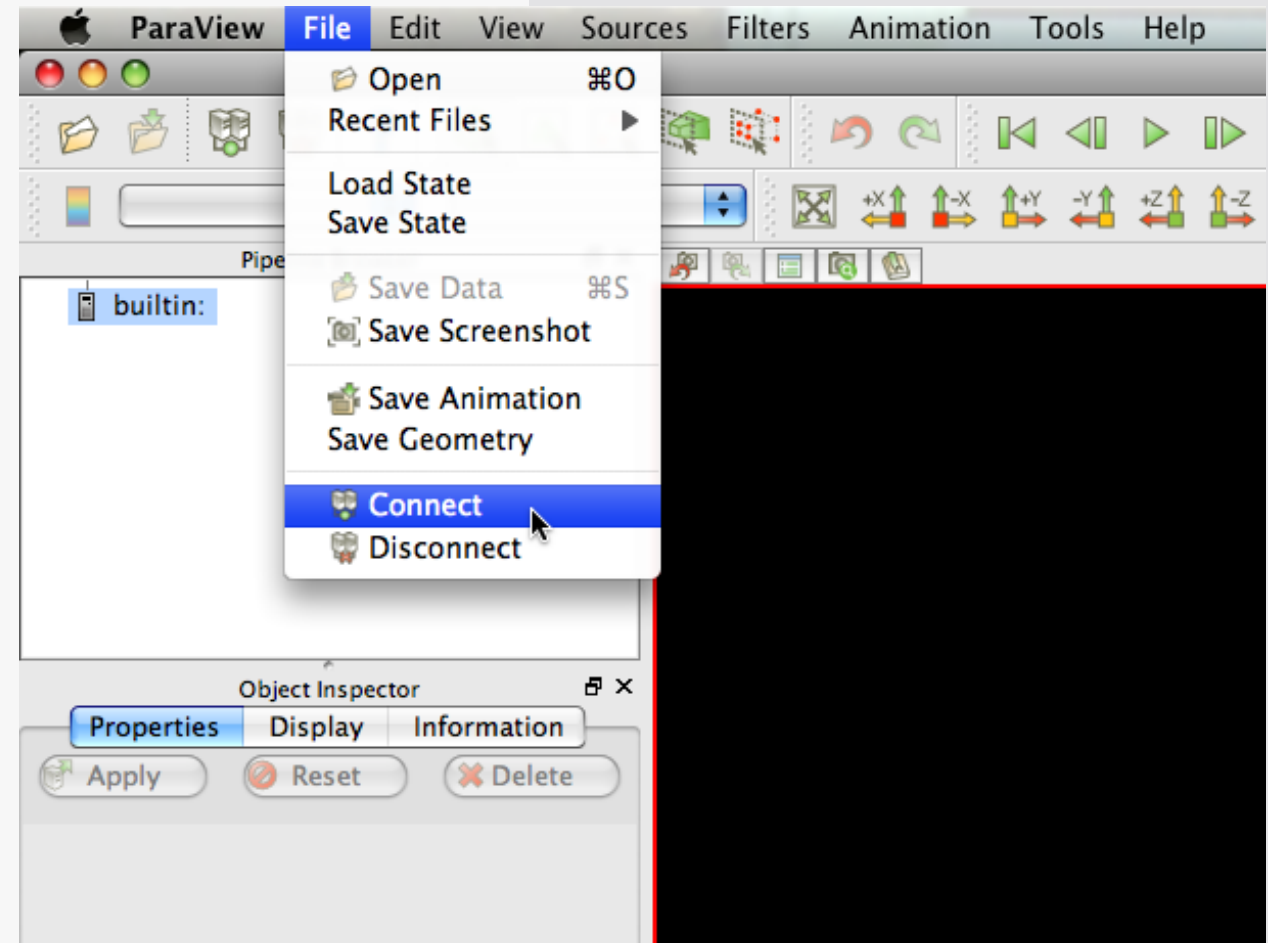
VTK	PLOT3D	Facet	Tetrad
Parallel (partitioned) VTK	SpyPlot CTH	PNG	UNIC
VTK MultiBlock (MultiGroup, Hierarchical, Hierarchical Box)	HDF5 raw image data DEM	SAF	VASP
Legacy VTK	VRML	LS-Dyna	ZeusMP
Parallel (partitioned) legacy VTK	PLY	Nek5000	ANALYZE
EnSight files	Polygonal Protein Data Bank	OVERFLOW	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**

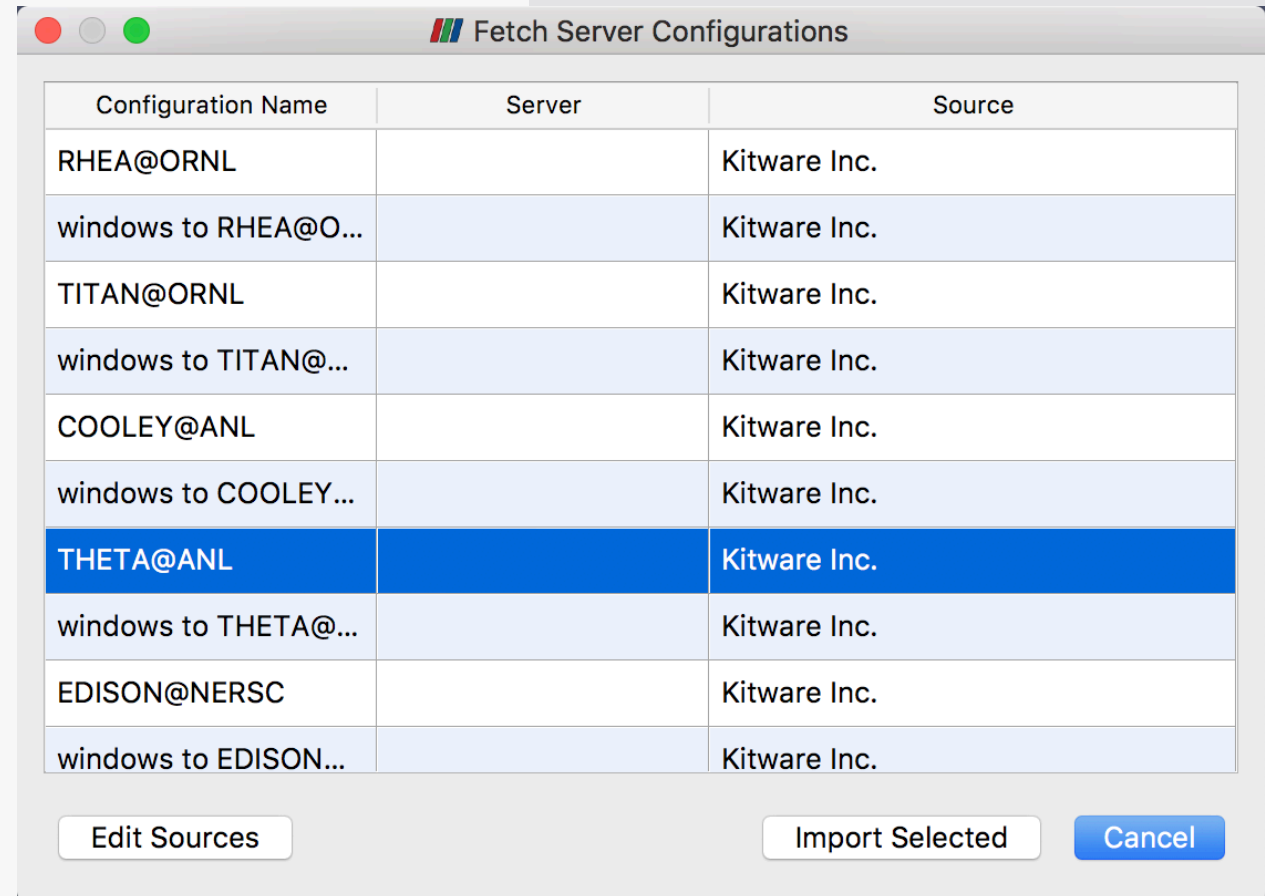


ParaView on Theta

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**



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.

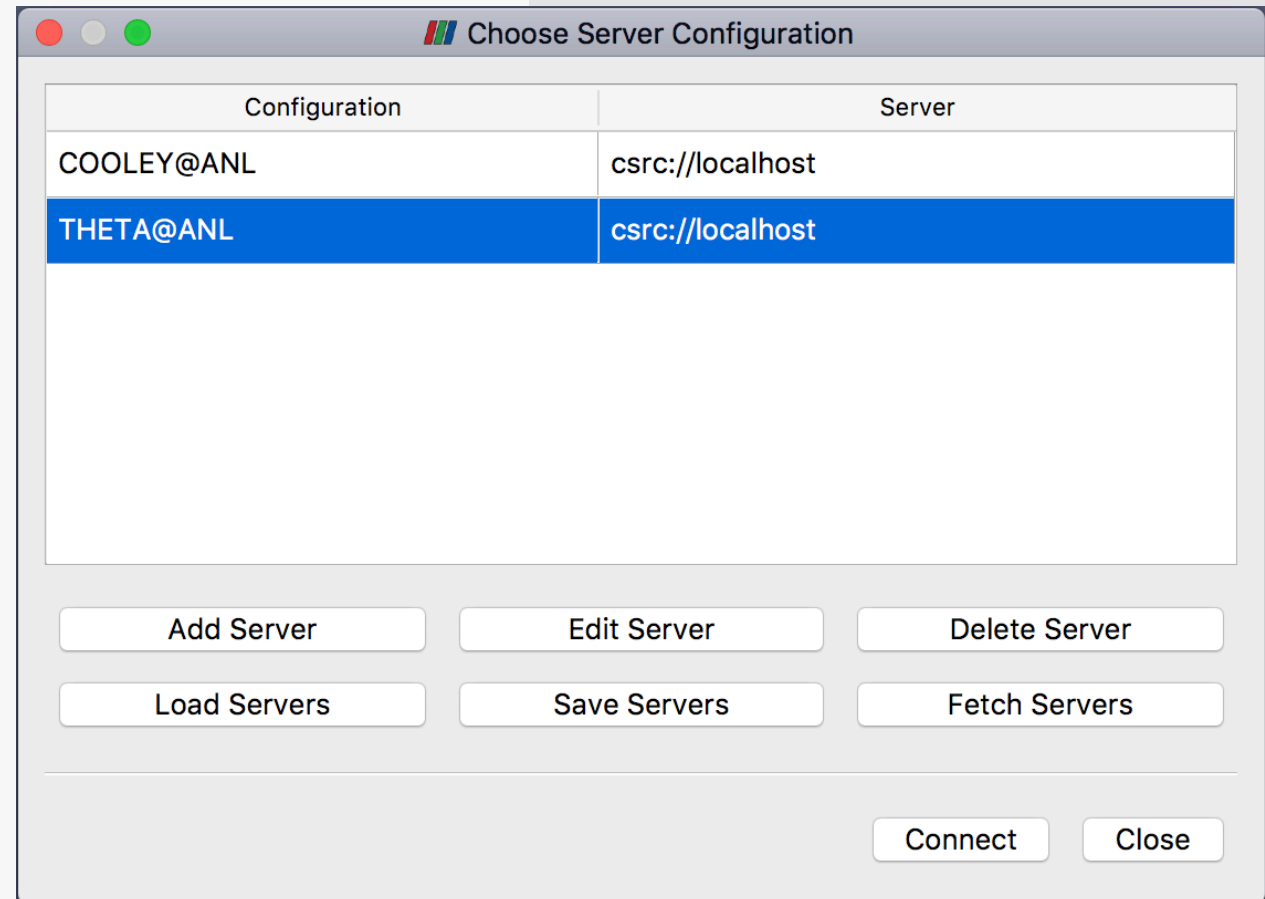
Edit Sources Import Selected Cancel

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**

Connection Options for "THETA@ANL"

Xterm executable ...

SSH executable ...

Remote machine

Username

ParaView version

Client port ▾ ▸

Server port ▾ ▸

Number of nodes to reserve ▾ ▸

Number of minutes to reserve ▾ ▸

Account

Queue


Job name

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

Exploring the data interactively, the scientist is able to identify structures previously unseen in the data.



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



Can We....

Enable use of any in situ framework?

Develop analysis routines that are portable between codes?

Make it easy to use?

OUR APPROACH

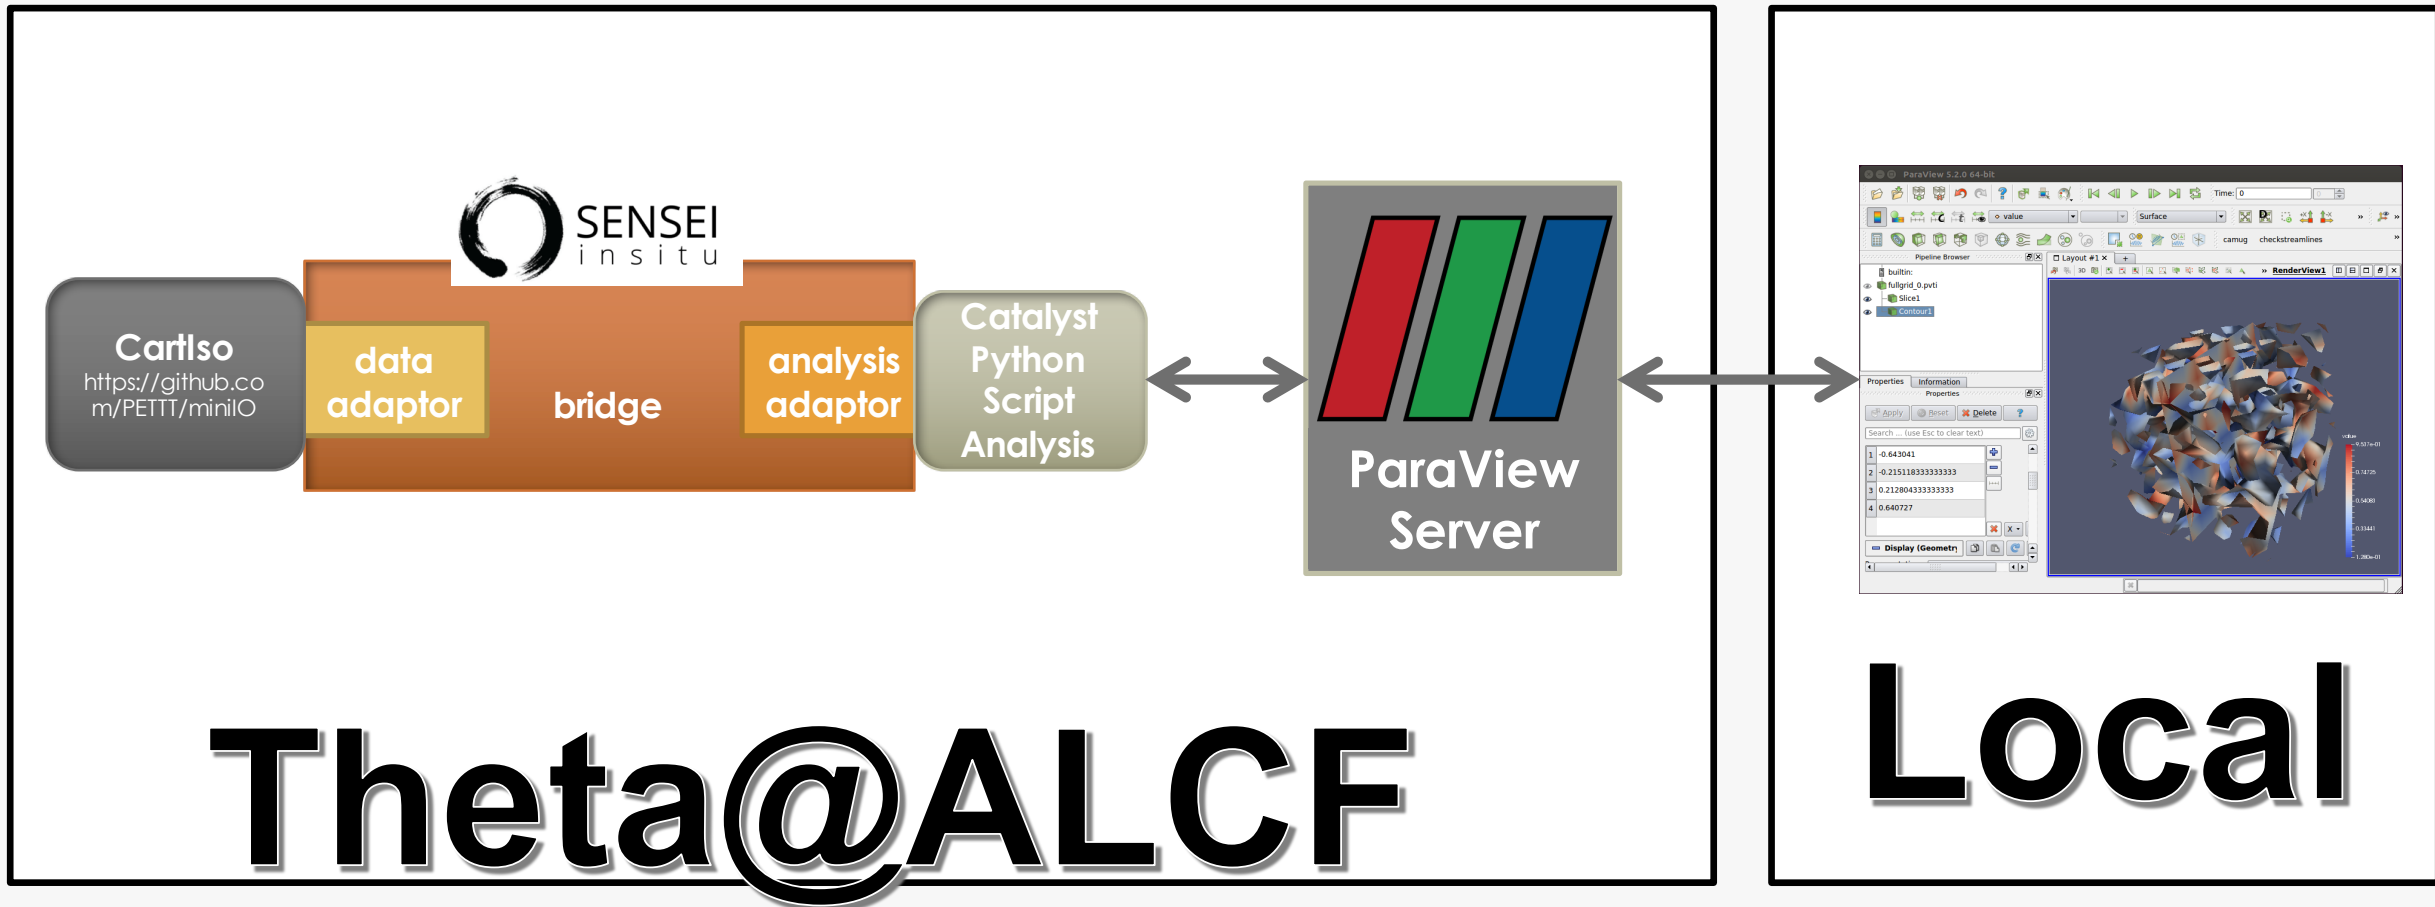
Data model – to pass data between
Simulation & Analysis

API – for instrumenting simulation and
analysis codes



SENSEI
i n s i t u

Miniapp instrumentation with SENSEI



Time: 0

Surface

camug checkstreamlines imageblanking makeverts

Pipeline Browser

- ThetalsoCartSENSEILive (csrc://localh)
 - Extract: ExtractSurface1
 - Extract: Slice3
 - catalyst:
 - input
 - PassArrays1
 - Contour1

Properties Information

Properties

Apply Reset Delete ?

Search ... (use Esc to clear text)

Properties (Ext)

Display (Geom)

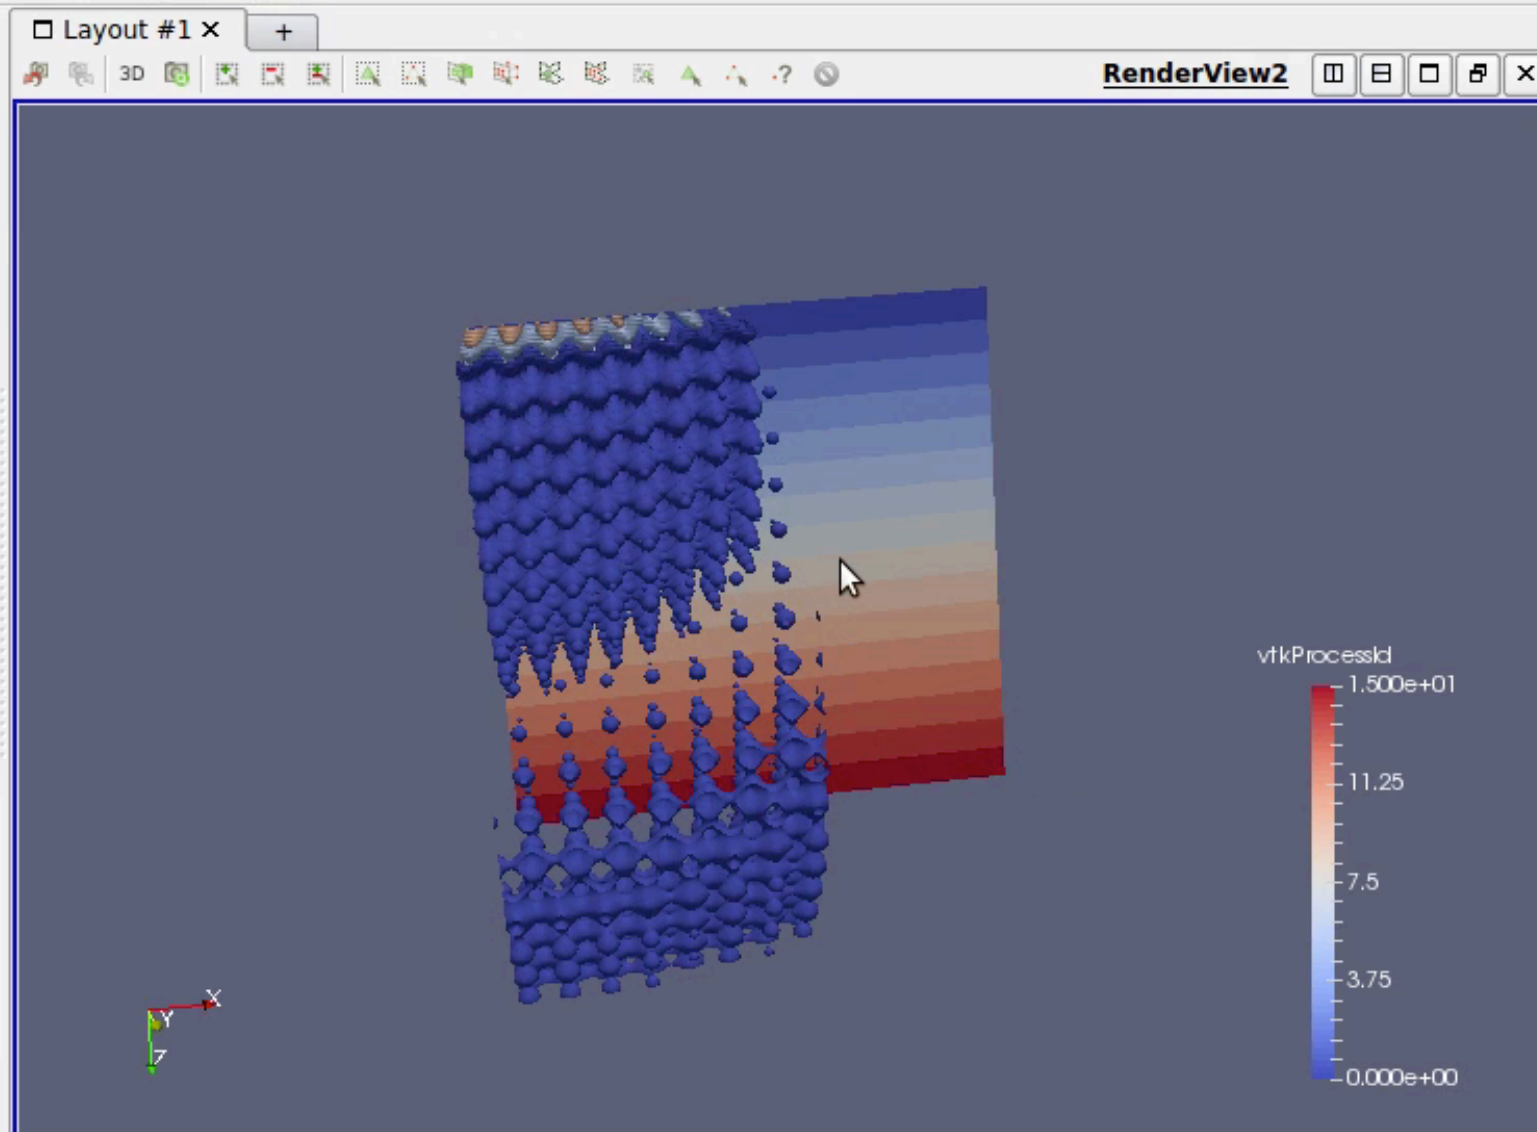
Representation Surface

Coloring

vtkProcessId

Edit

Styling



Ospray for Interactive *In Situ* Visualization




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

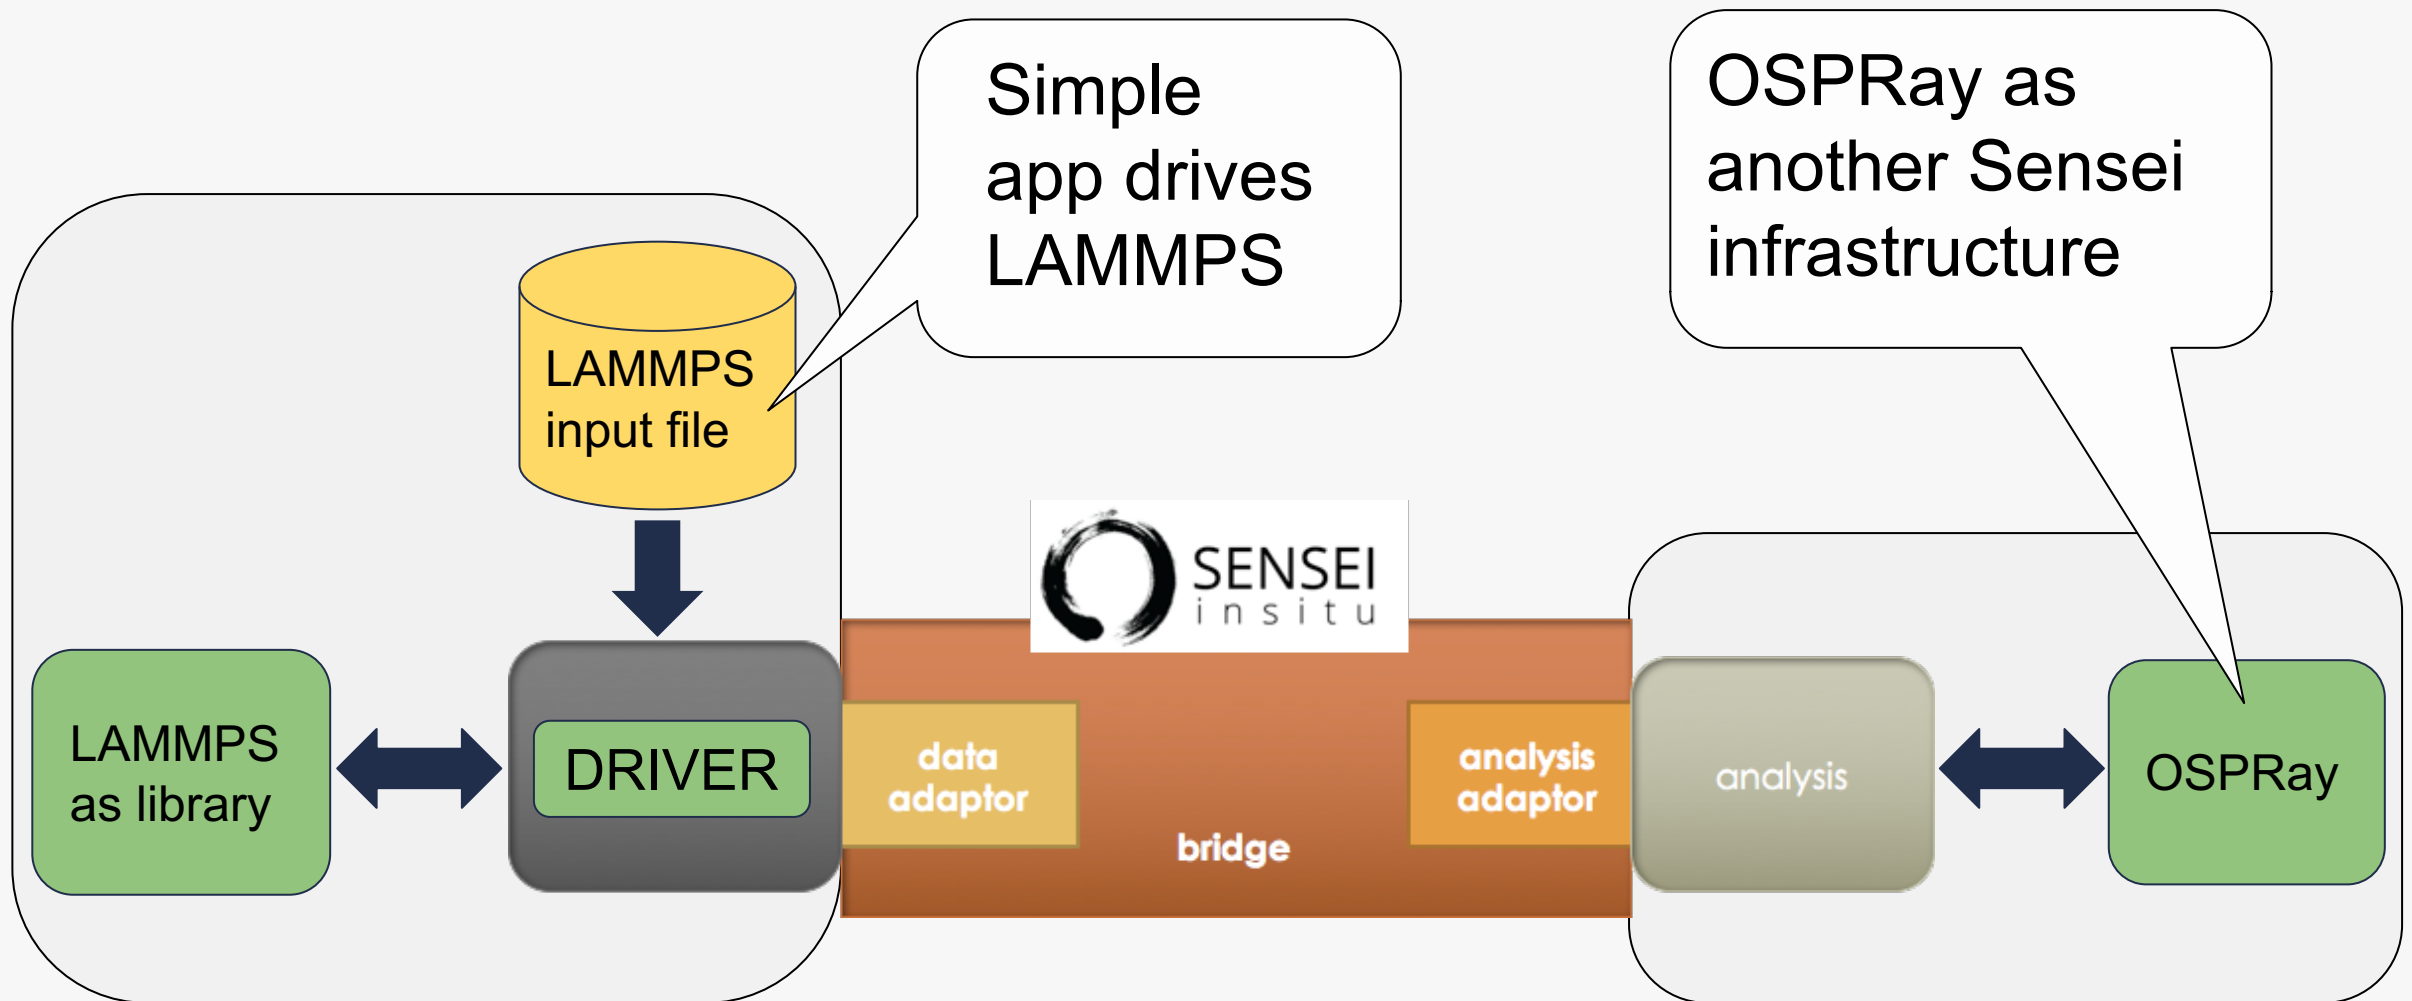
- Volumes, triangle meshes, non-polygonal geometry (spheres, cylinders,...)
- Ray traced shading effects for shadows, ambient occlusion

Free & open source: Apache 2.0 License

- <http://ospray.org/>

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

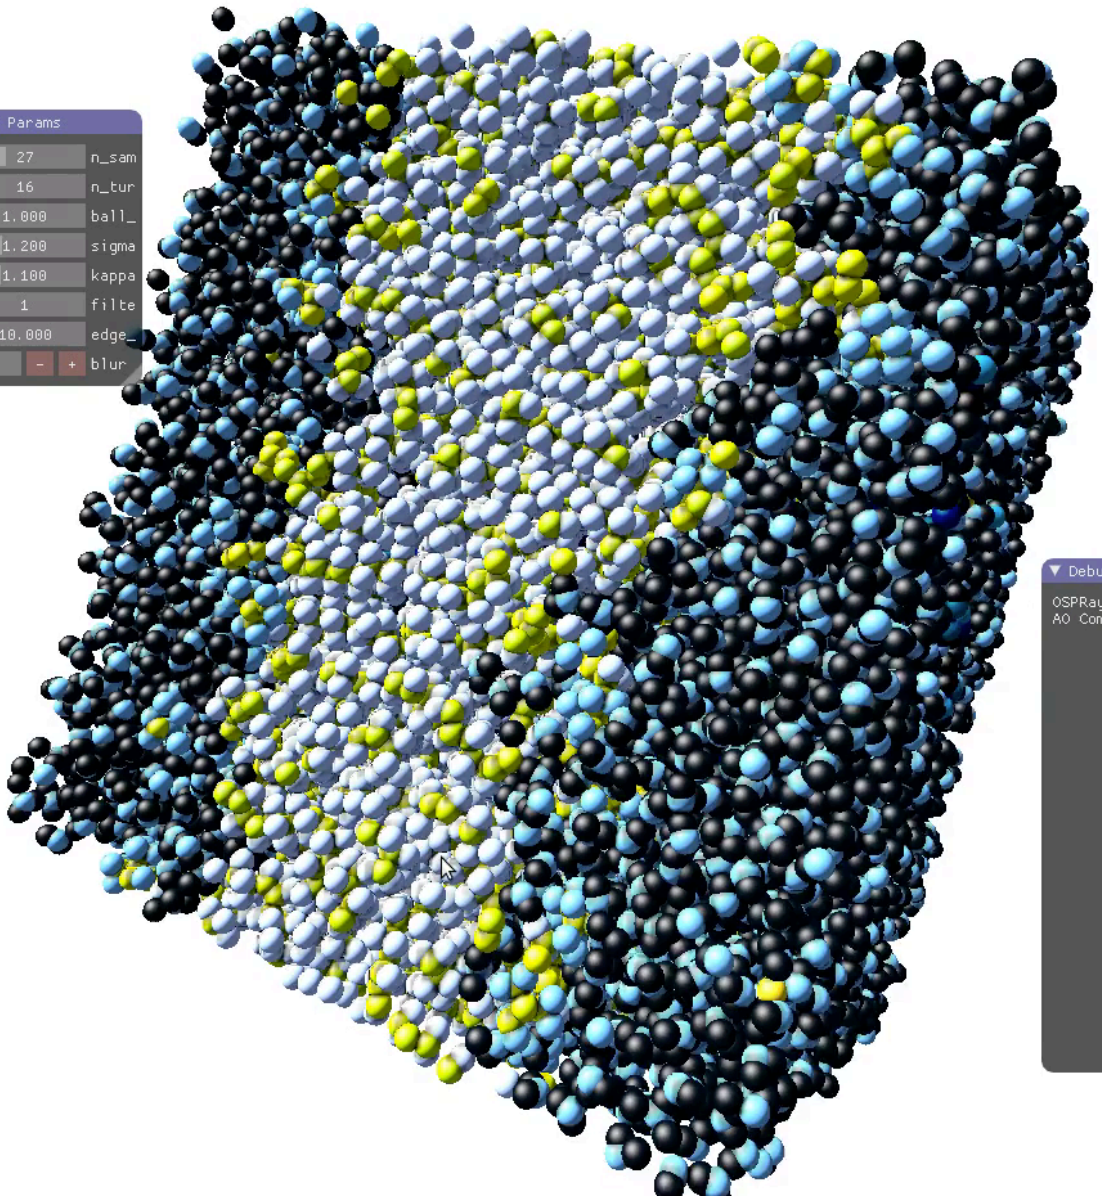
LAMMPS instrumentation with SENSEI and ospray



```
srizzi@porota: /d0/SC17livedemo/lammps_sensei_ospray/viewer/build
srizzi@porota: /d0/SC17livedemo/lammps_sens... x srizzi@porota: /d0/SC17livedemo/lammps_sens... x srizzi@porota: /d0/SC17livedemo/lammps_sens... x
Got atom type range {0, 67}
Queried region with 15897 particles
bounds = [(Got atom type range {0, 67}
-27.5, -38.5, -36.3761), (27.5, 38.5, 36.373)]
Queried region with 16103 particles
bounds = [(-27.5, -38.5, -36.3761), (27.5, 38.5, 36.373)]
Got atom type range {0, 67}
Queried region with 16105 particles
bounds = [(Got atom type range {0, 67}
-27.5, -38.5, -36.3764), (27.5, 38.5, 36.3733)]
Queried region with 15895 particles
bounds = [(-27.5, -38.5, -36.3764), (27.5, 38.5, 36.3733)]
Got atom type range {0, 67}
Queried region with 16106 particles
bounds = [(-27.5, -38.5, Got atom type range {0-36.3764), (27.5, 38.5, 36.3733)]
, 67}
Queried region with 15894 particles
bounds = [(-27.5, -38.5, -36.3764), (27.5, 38.5, 36.3733)]
Got atom type range {0, 67}
Got atom type range {0, 67}
Queried region with 15902 particles
bounds = [(Queried region with 16098 particles
bounds = [(-27.5, -27.5, -38.5, -36.3767), (27.5-38.5, -36.3767), (27.5, , 38.5, 36.3736)]
38.5, 36.3736)]
Got atom type range {0, 67}
Got atom type range {0, 67}
Queried region with 15898 particles
bounds = [(-27.5, -38.5, -36.3767), (27.5, 38.5, 36.3736)]
Queried region with 16102 particles
bounds = [(-27.5, -38.5, -36.3767), (27.5, 38.5, 36.3736)]
Got atom type range {0, 67}
Got atom type range {0, 67}
Queried region with 15900 particles
bounds = [(-27.5, -38.5, -36.377), (27.5, 38.5, 36.3739)]
Queried region with 16100 particles
bounds = [(-27.5, -38.5, -36.377), (27.5, 38.5, 36.3739)]
Got atom type range {0, 67}
Got atom type range {0, 67}
Queried region with 16100 particles
bounds = [(Queried region with 15900 particles
bounds = [(-27.5, -38.5, -36.377), (27.5, 38.5, 36.3739)]
-27.5, -38.5, -36.377), (27.5, 38.5, 36.3739)]

```

LAMMPS + SENSEI Viewer



AO Params


27	n_sam
16	n_tur
1.000	ball_
1.200	sigma
1.100	kappa
1	filte
10.000	edge_
2	blur

Debug

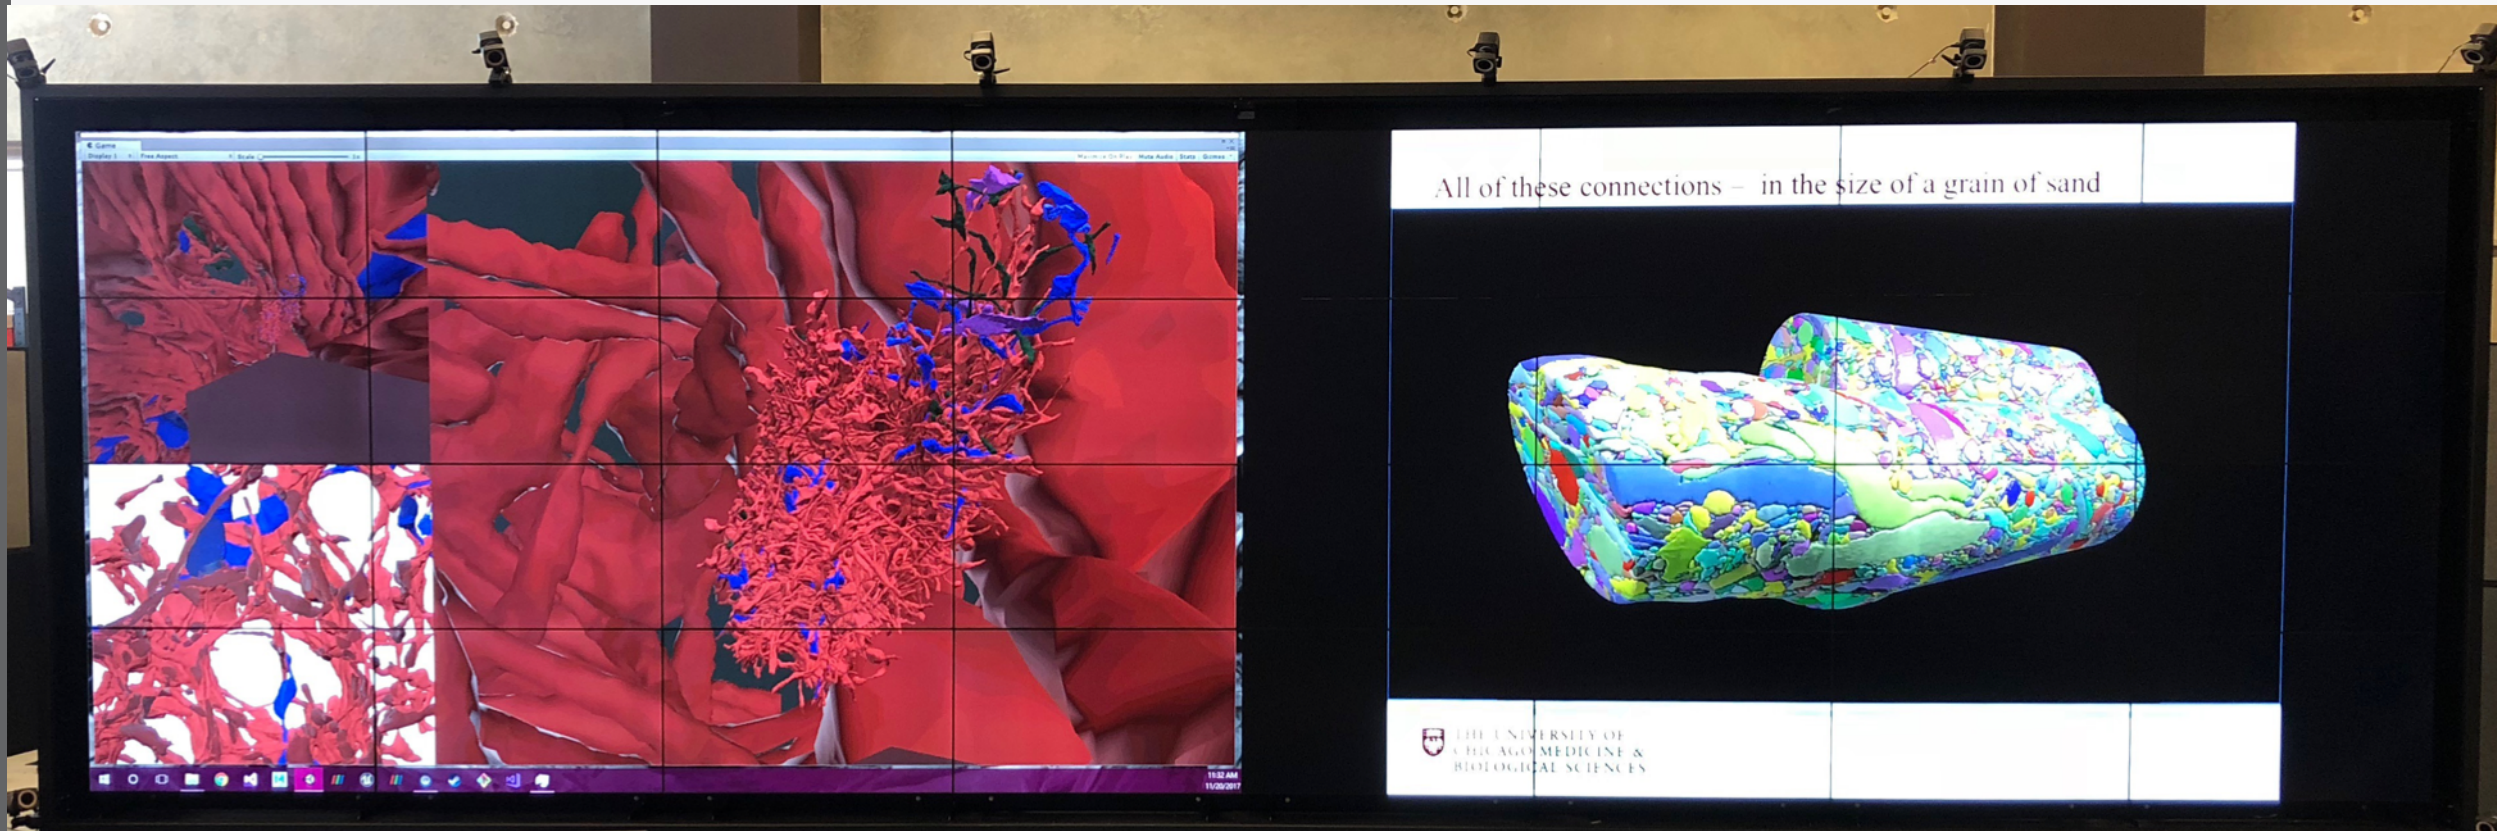
OSPray Render Time: 5
AO Computation Time:




Immersive Visualization

Immersive visualization



Immersive visualization





Slide courtesy Edouard Brooks, ALCF

QUESTIONS?

Joe Insley
insley@anl.gov

Silvio Rizzi
srizzi@anl.gov