Rendering on the Supercomputing Platform

Jon Woodring

Los Alamos National Laboratory

Visualization process

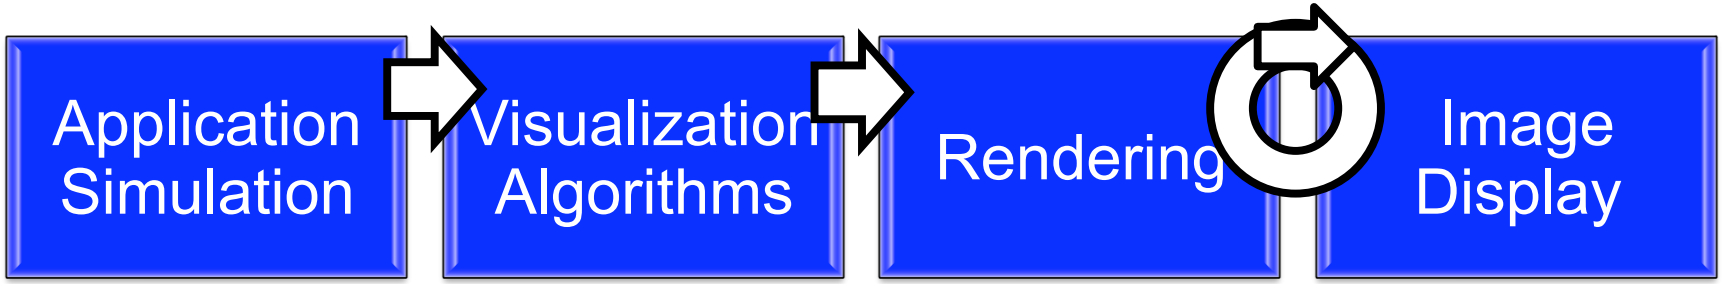

Visualization process

Simulation
Results

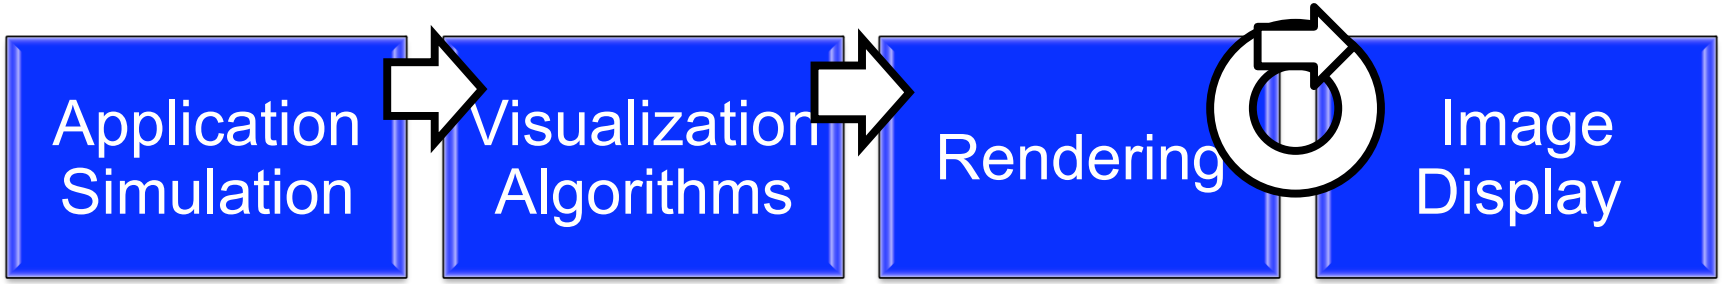

Visualization process

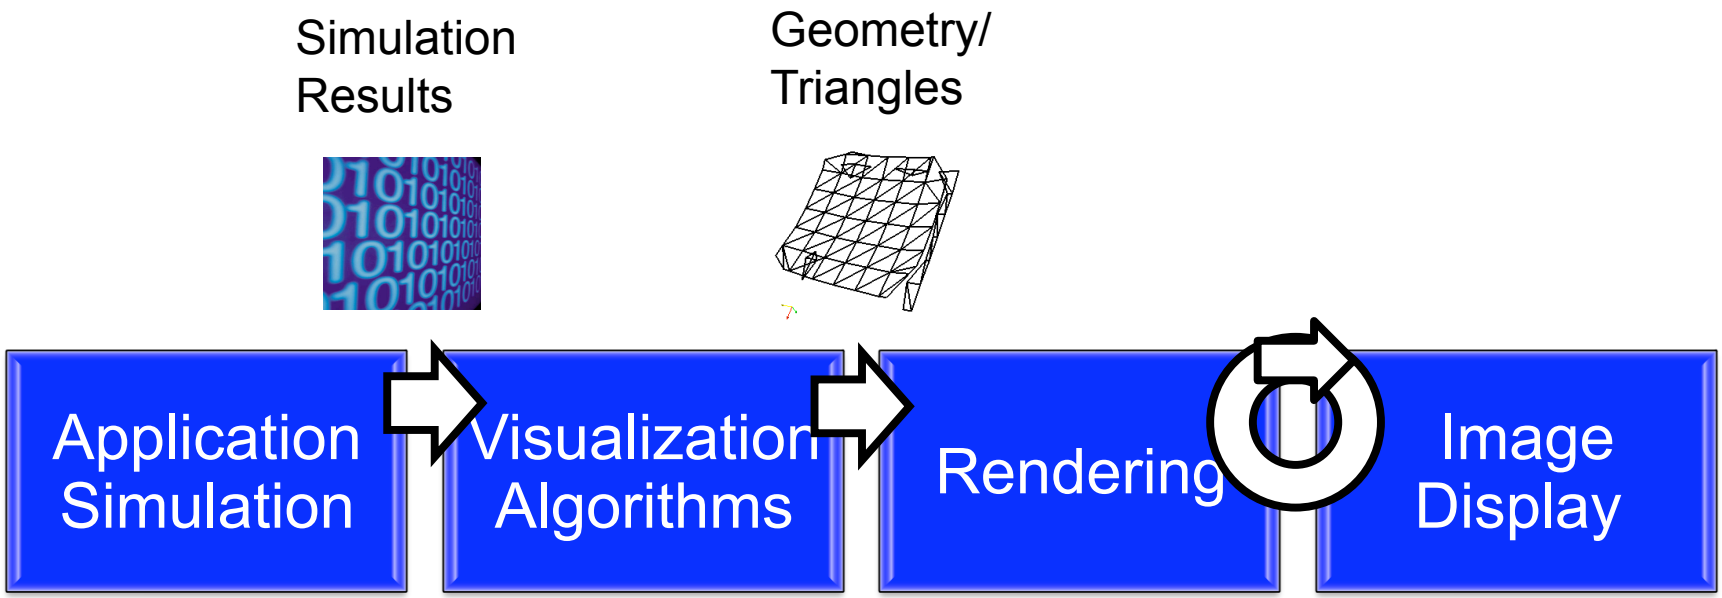

Visualization process

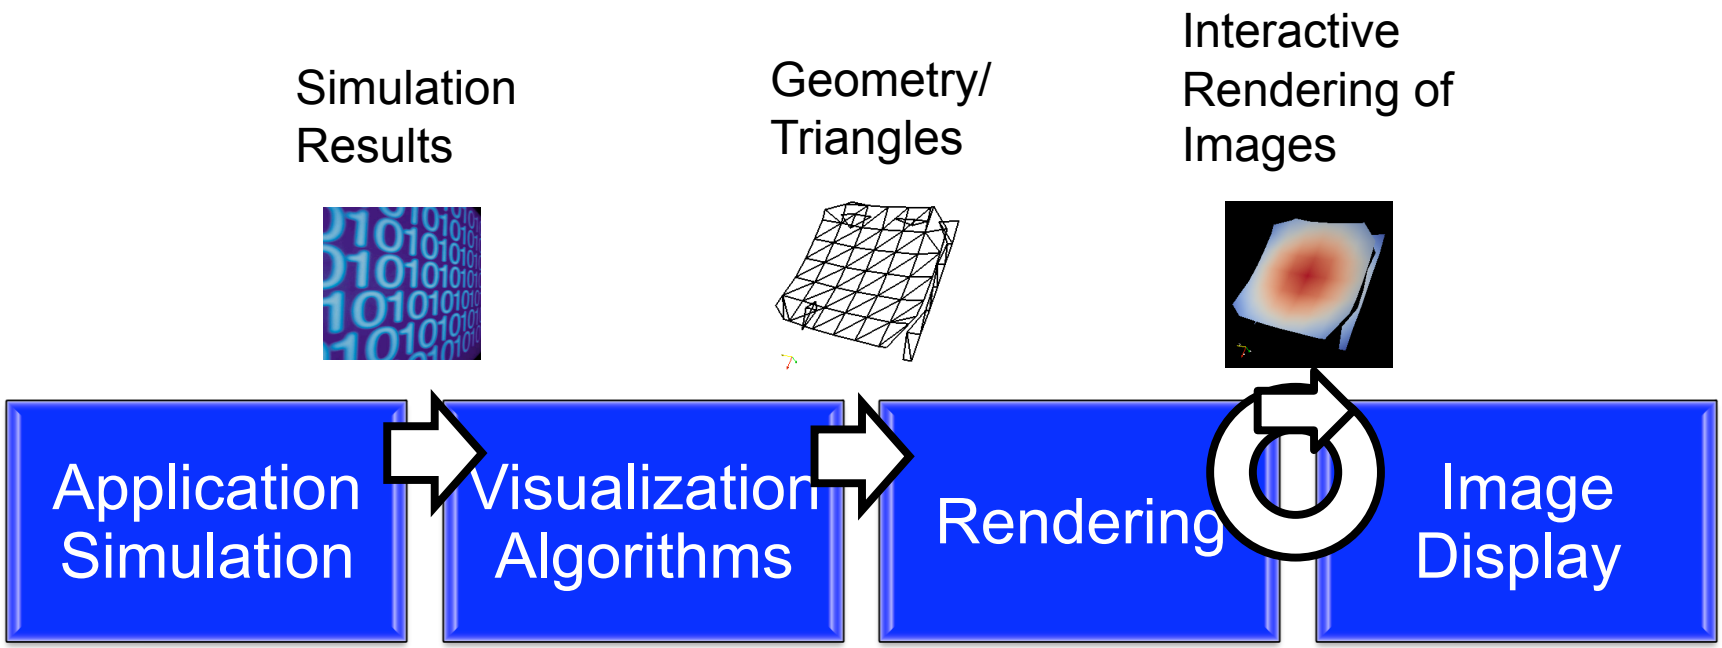

Current LANL approach

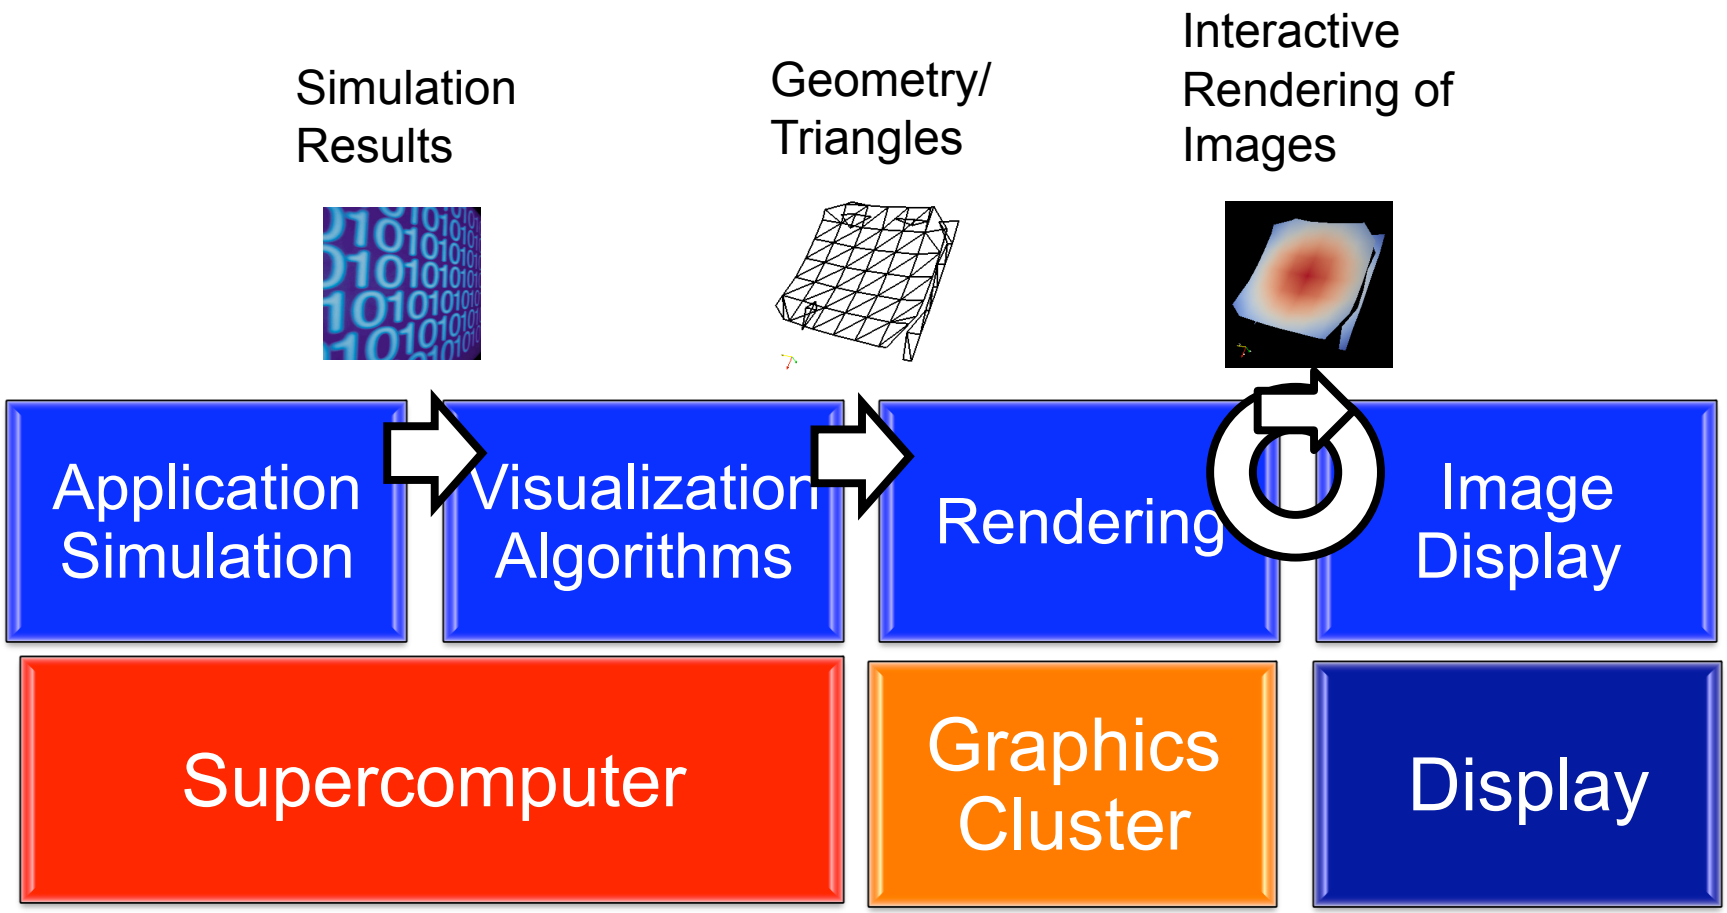

Visualization Process

- **The visualization and analysis process is composed of a number of activities**
  - Analysis and statistics
  - Representation
    - Map simulation data to a visual representation (i.e., geometry)
  - Rendering
    - Map geometry to imagery on the screen
- **Already runs on the supercomputer**
  - Analysis, statistics, and representation
- **Rendering is performed by separate hardware**

Can we interactively render on the supercomputer, as well?

- **Interactivity is critically important for insight**
  - 5-10 fps minimum
- **High-performance requirements**
  - Provided by GPU in graphics cluster
- **Maximum parallel compositing**
  - 20-30 frames per second (and lower)
  - Network limited
- **To meet this target frame rate**
  - GPU rendering
    - 300-350 frames per second – overkill
  - Maybe CPU rendering? Mesa 3D is too slow…

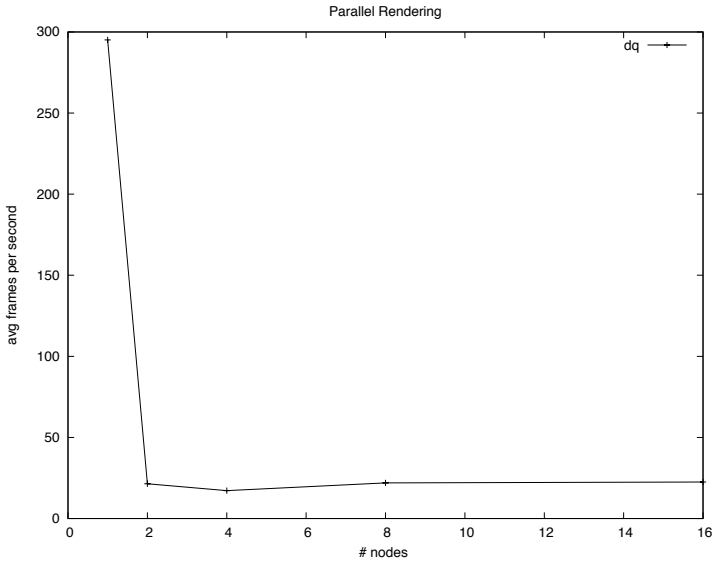

Manta ray tracer 1 million polygons

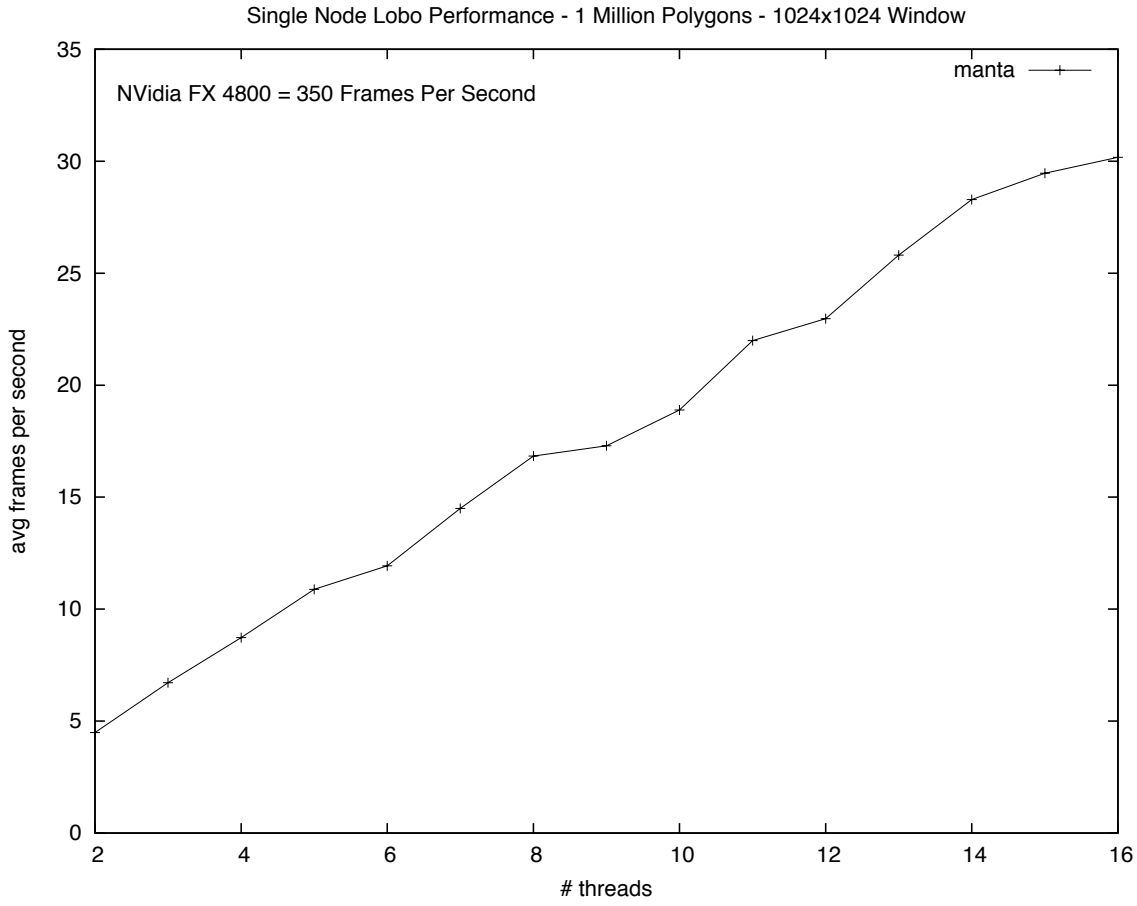

Alternative approach we are evaluating

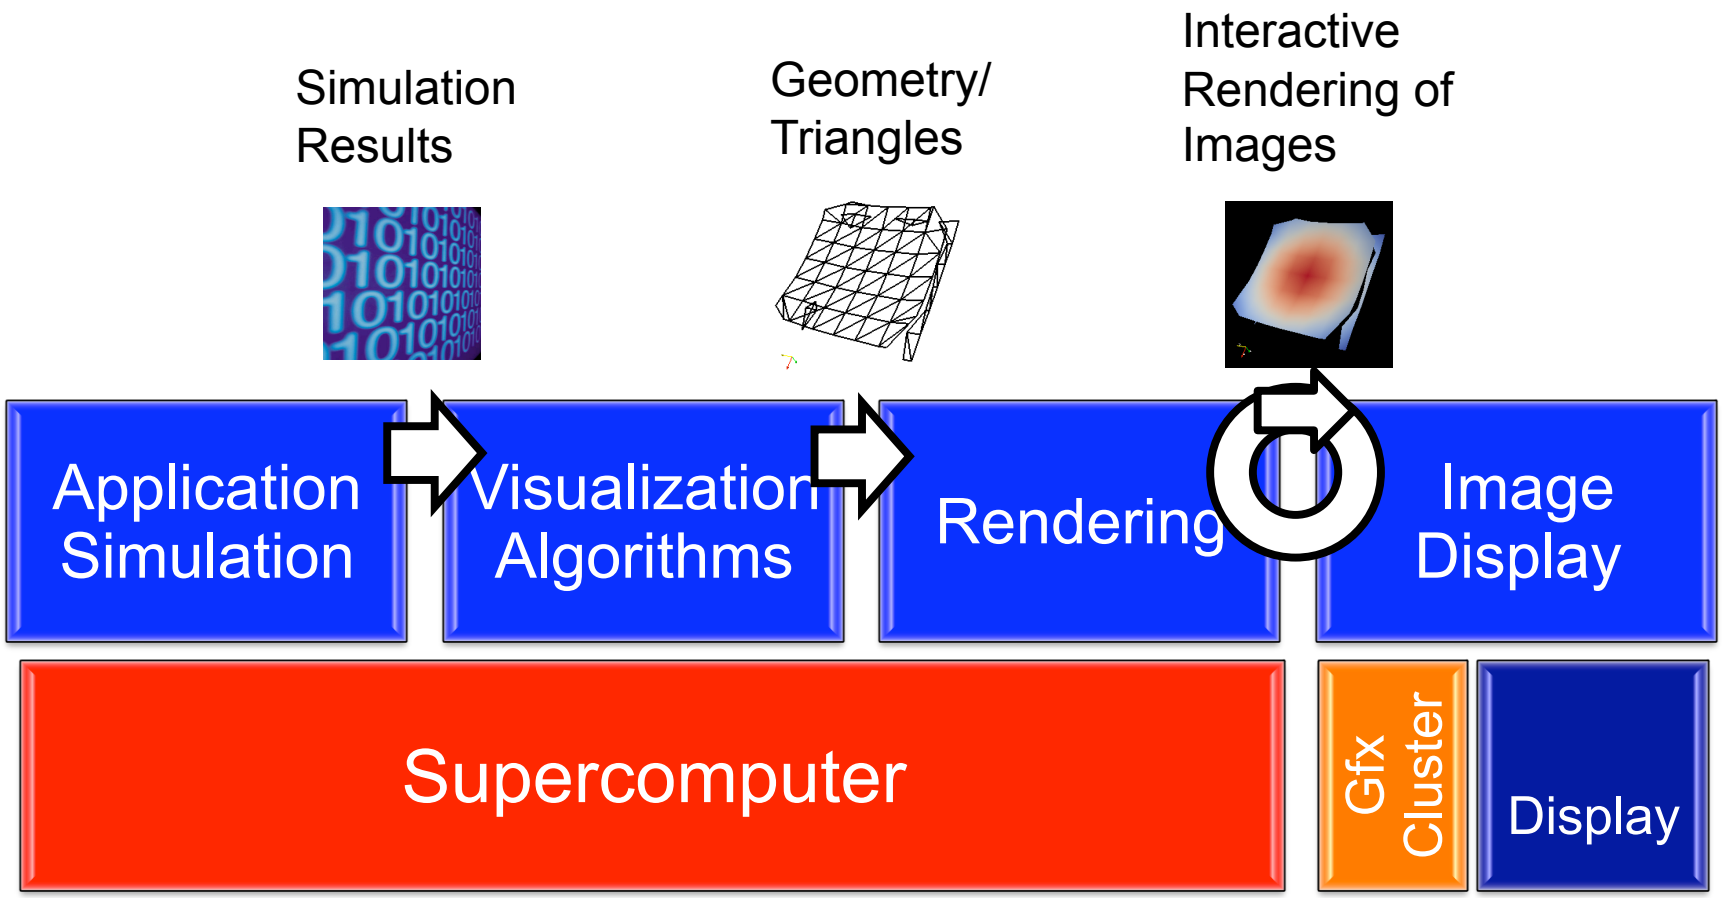

Manta Demo in ParaView

- **Show Demo**

- **Manta View is a ParaView plugin that implements a 3D view**
  - Wraps Manta, provides data, makes render requests
- **Use standard VTK depth compositing classes**
  - Ice-T or binary swap
- **Z (depth) channel added to Manta**
- **vtkMantaPolyDataMapper**
  - Hands triangles to Manta (like a OpenGL display list)
  - Generates tubes and spheres for lines and points

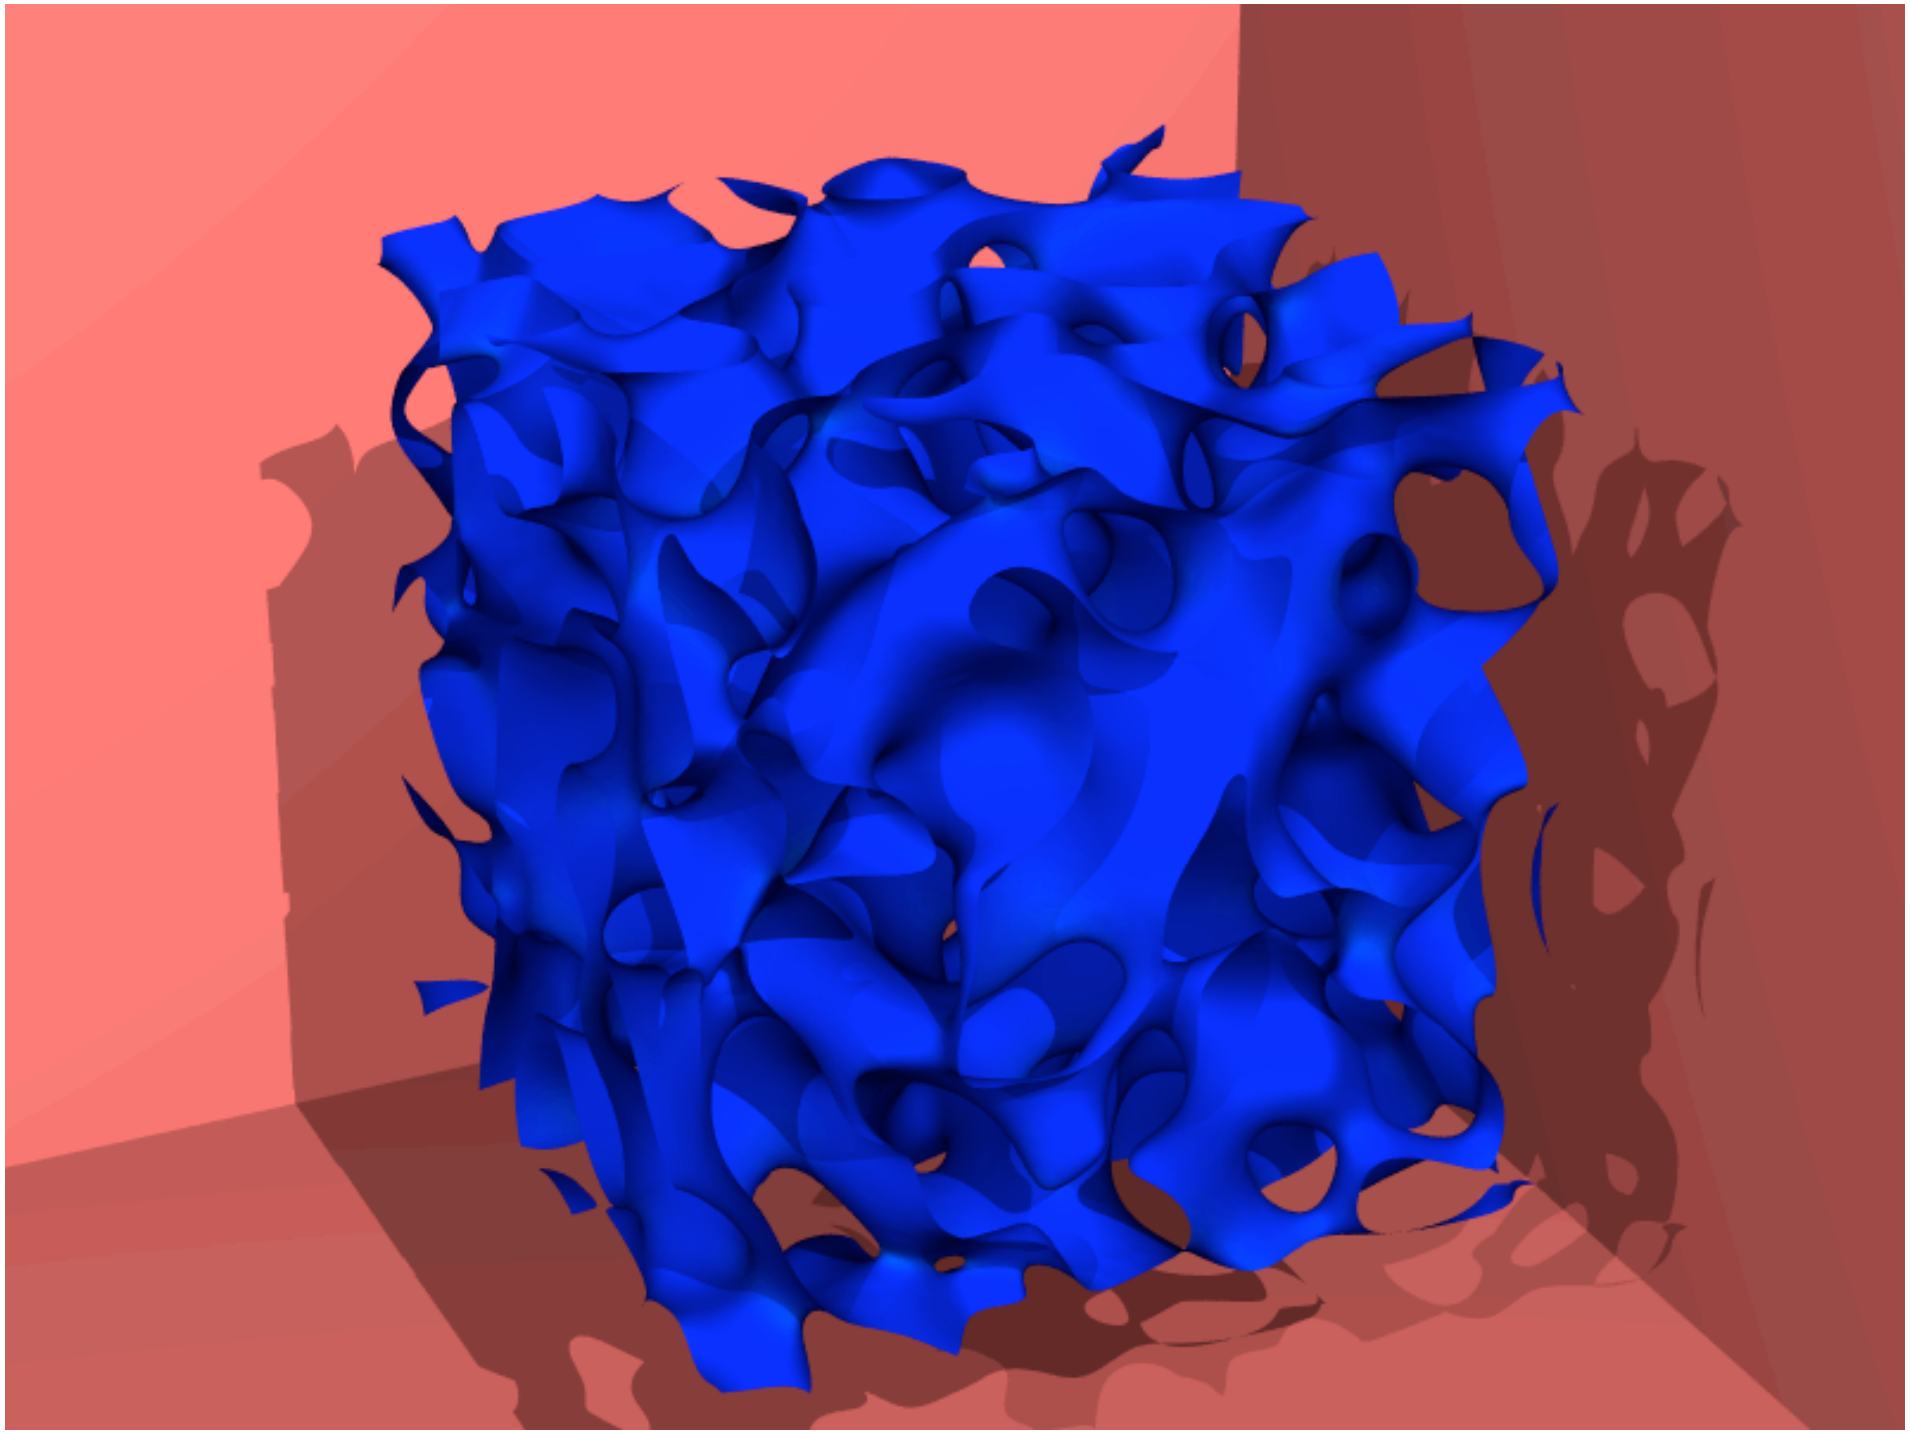

Parallel Sort-Last Rendering with Manta in ParaView

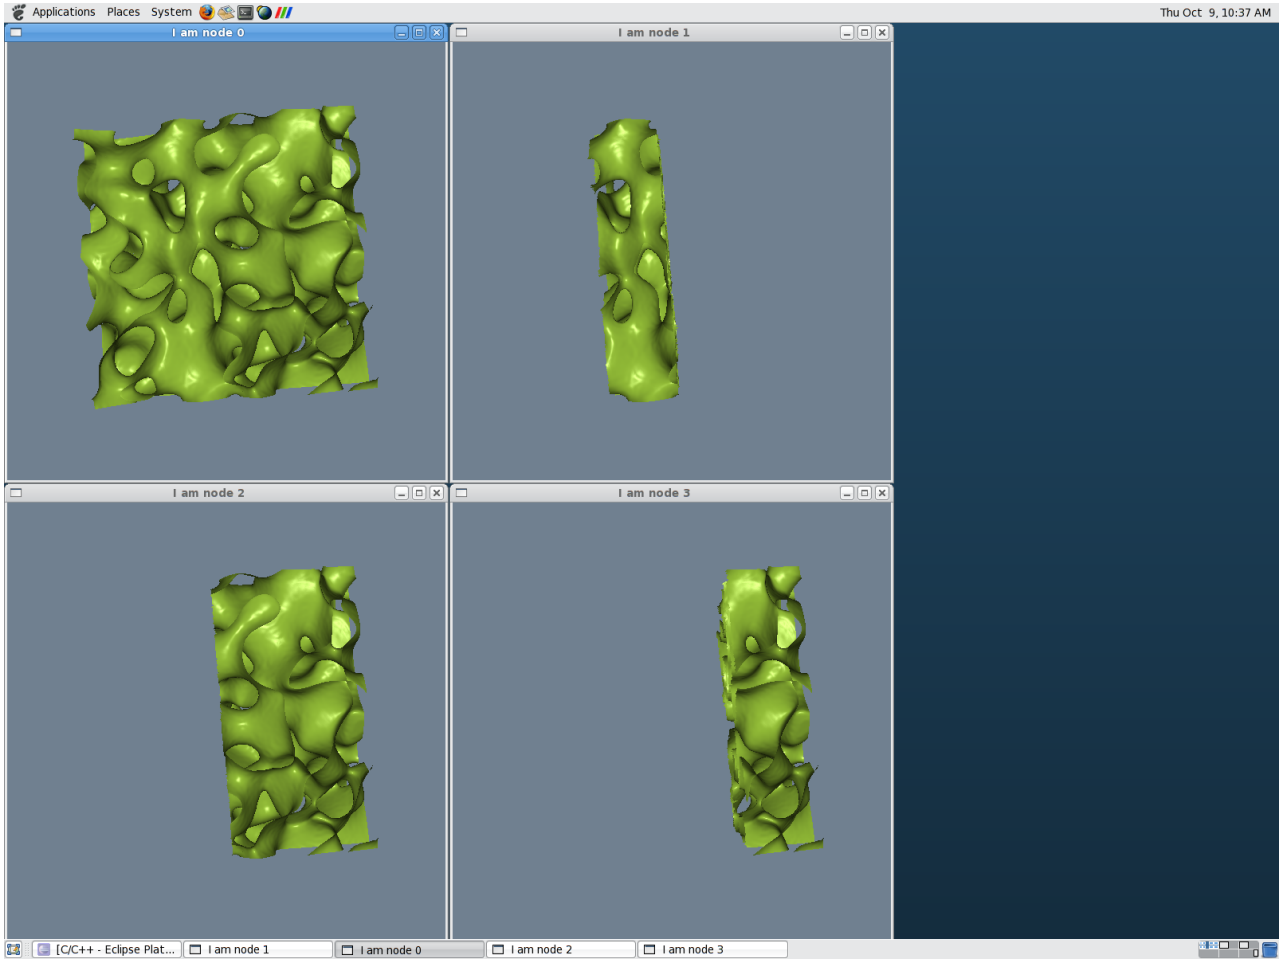

Benefits

- **Fewer specialized visualization requirements**
  - Visualization becomes a supercomputing application
  - One HPC resource to manage
- **Data is already there – no need to move it**
- **Scalable to the supercomputer size**
- **Software rendering**
  - Potentially high quality images and flexibility – shadows, multi-sampling, reflection, refraction, etc.
  - Manta gets faster the more cores you throw at it

Drawbacks

- **Interactive Queue?**
  - Supercomputing queues are batch
  - An uphill battle to get good interactive queues for vis
- **No specialized hardware for visualization task**
- **GPU is dominant for graphics and rendering**
  - Up front cost of going back/developing for CPU rendering
- **Frame rate is not high enough for stereo/RAVE/CAVE**

VPIC – A case study of interactive visualization on the RoadRunner platform

- **Running simulation on 4096 RR processors**
  - Computing a 8096 x 8096 x 448 grid
- **The VPIC team ran their visualization on 128 RR processors**
  - Striding and sub-setting to explore and understand their data
- **The VPIC team considers interactive visualization critical to the success of their project**
  - Bill Daughton, Brian Albright
- **The following movie was done on the platform with Mesa 3D rendering**
  - Bill interacts with the data on the supercomputer with ParaView
  - He queues up batch movies to be generated after interaction
  - Currently evaluating rendering the VPIC data with Manta

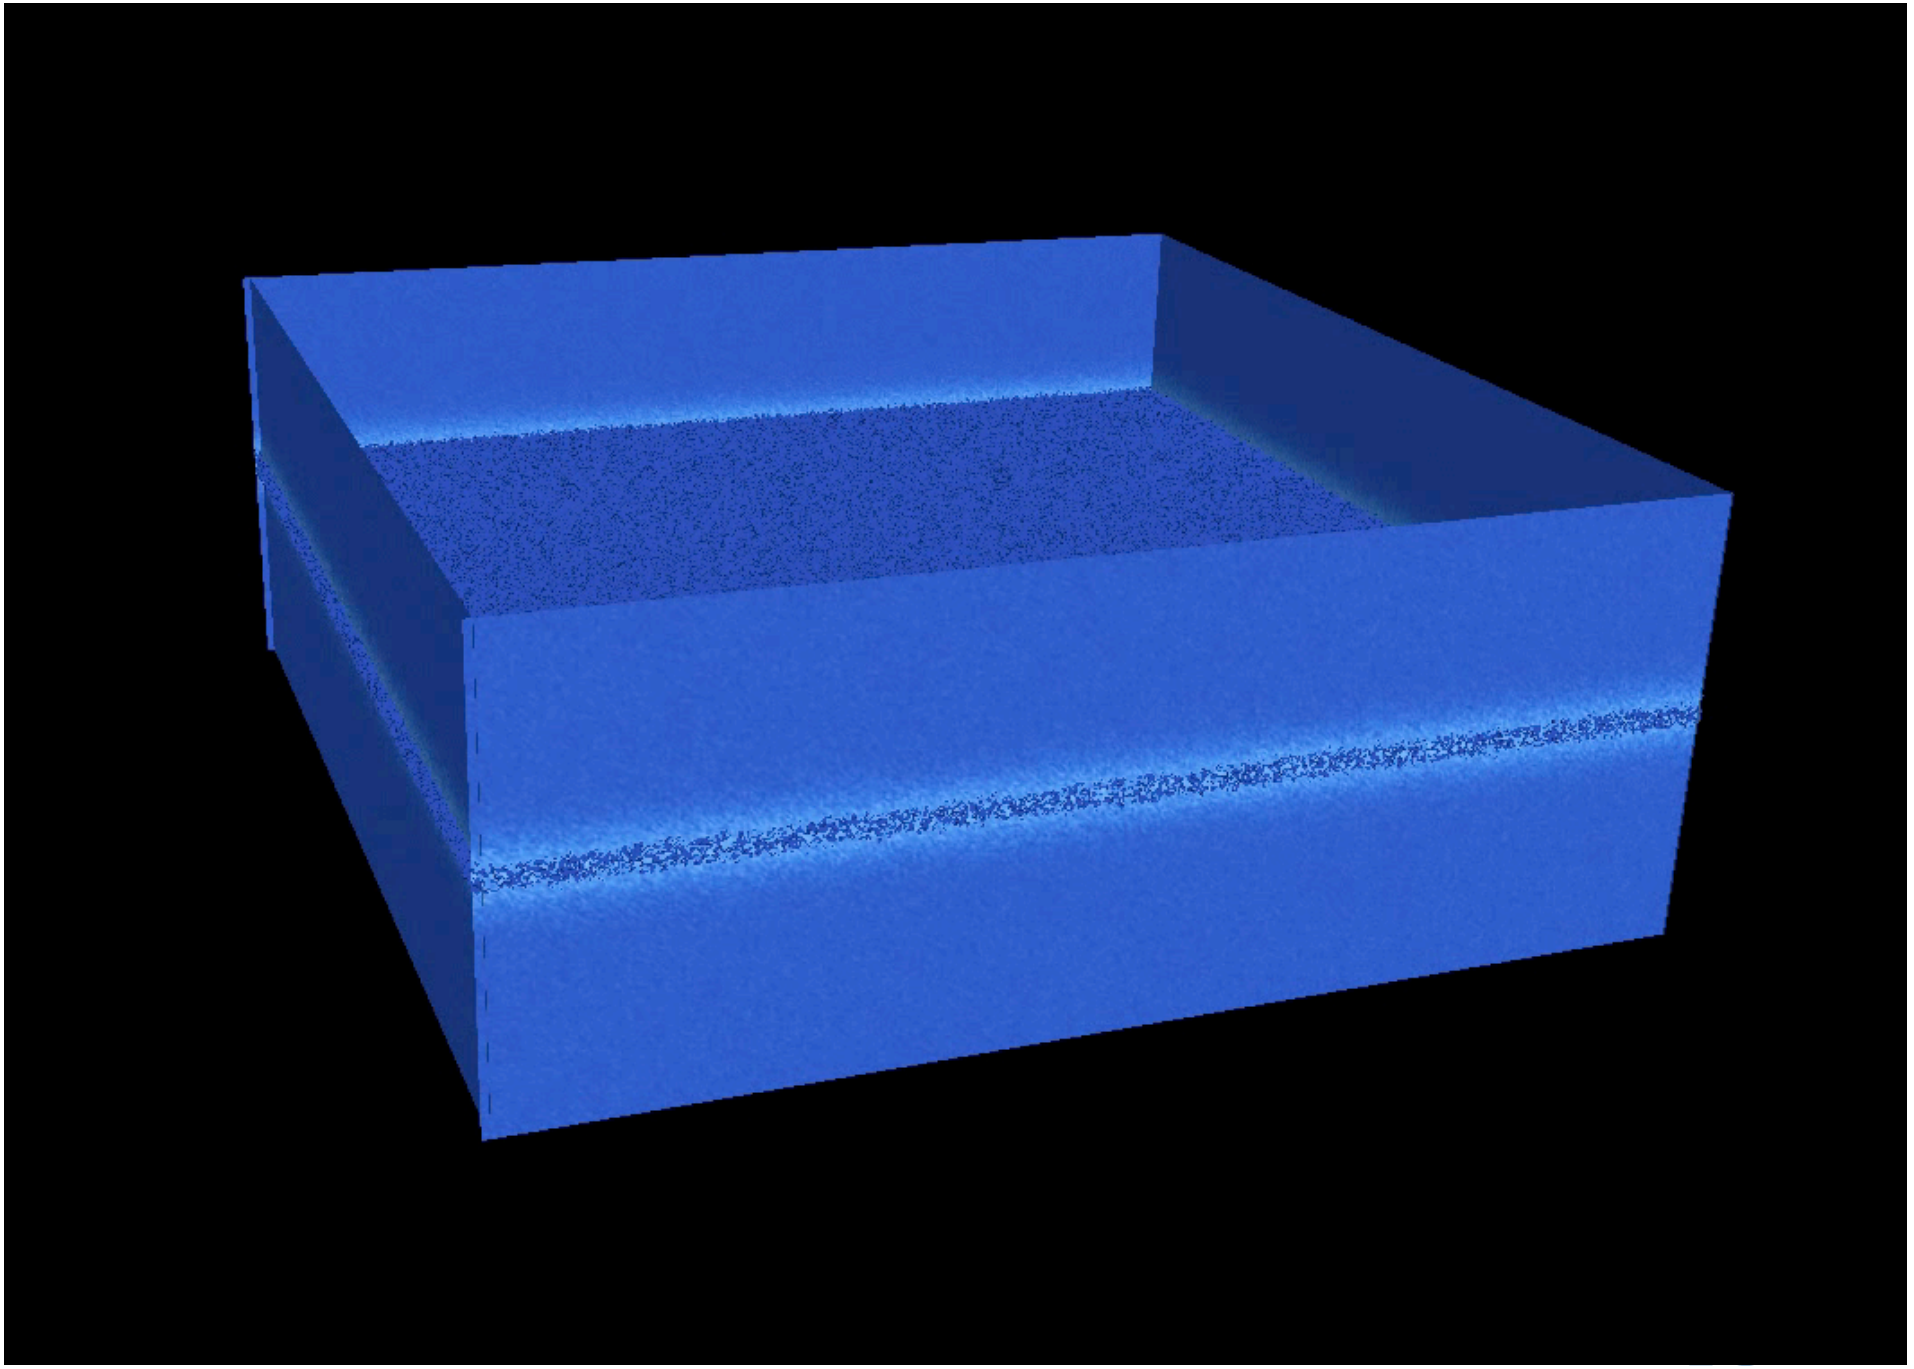

Contact

- **Jim Ahrens: ahrens@lanl.gov**
- **Li-Ta Lo: ollie@lanl.gov**
- **Dave DeMarle: dave.demarle@kitware.com**
- **John Patchett: patchett@lanl.gov**
- **Jon Woodring: woodring@lanl.gov**

- **Research is funded by ASC**