



# **Parallel Distributed-Memory Visualization with ParaView**

## **Cluster Processing Large Models**

**August 31, 2009**

**Kenneth Moreland  
Sandia National Laboratories**



Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,  
for the United States Department of Energy's National Nuclear Security Administration  
under contract DE-AC04-94AL85000.






# Golevka Asteroid vs. 10 Megaton Explosion

---

- CTH shock physics, over 1 billion cells






## Polar Vortex Breakdown

---

- SEAM Climate Modeling, 1 billion cells (500 million cells visualized).






## Objects-in-Crosswind Fire

---


- Coupled SIERRA/Fuego/Syrinx/Calore, 10 million unstructured hexahedra







# Large Scale AMR





# ParaView Architecture

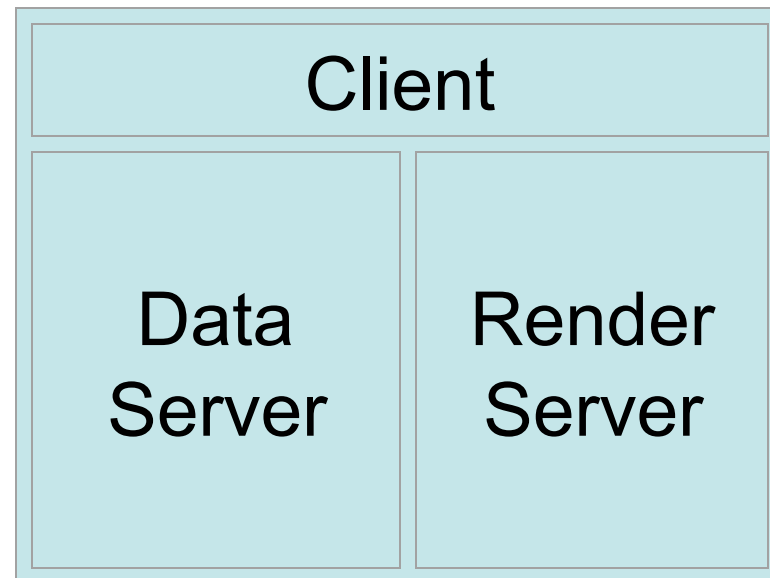
---

- **Three tier**
  - **Data Server**
  - **Render Server**
  - **Client**



# Standalone


---





# Client-Server


---





# Client-Render Server-Data Server

---







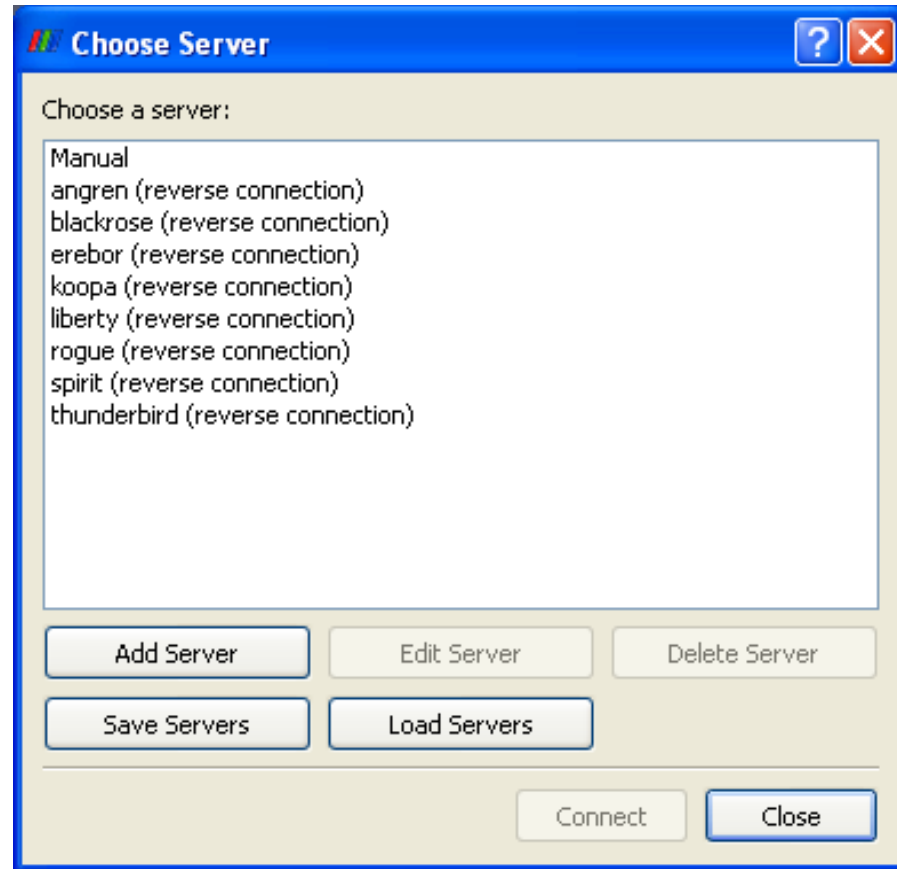
# Requirements for Installing ParaView Server

---

- C++
- CMake ([www.cmake.org](http://www.cmake.org))
- MPI
- OpenGL (or Mesa3D [www.mesa3d.org](http://www.mesa3d.org))
- Qt 4.3 (optional)
- Python (optional)
- [http://www.paraview.org/Wiki/Setting\\_up\\_a\\_ParaView\\_Server#Compiling](http://www.paraview.org/Wiki/Setting_up_a_ParaView_Server#Compiling)



# Connecting to a ParaView Server




[http://www.paraview.org/Wiki/Setting\\_up\\_a\\_ParaView\\_Server#Running\\_the\\_Server](http://www.paraview.org/Wiki/Setting_up_a_ParaView_Server#Running_the_Server)



## Data Parallel Pipelines

---

- Duplicate pipelines run independently on different partitions of data.









# Data Parallel Pipelines

---

- Duplicate pipelines run independently on different partitions of data.









# Data Parallel Pipelines

---

- **Some operations will work regardless.**
  - **Example: Clipping.**











# Data Parallel Pipelines

---

- **Some operations will work regardless.**
  - **Example: Clipping.**






# Data Parallel Pipelines

---

- **Some operations will work regardless.**
  - **Example: Clipping.**









# Data Parallel Pipelines

---

- **Some operations will have problems.**
  - **Example: External Faces**









# Data Parallel Pipelines

---

- **Some operations will have problems.**
  - **Example: External Faces**






# Data Parallel Pipelines

---

- Ghost cells can solve most of these problems.








# Data Parallel Pipelines

---

- Ghost cells can solve most of these problems.






# Data Partitioning

---

- Partitions should be load balanced and spatially coherent.






# Data Partitioning

---

- **Partitions should be load balanced and spatially coherent.**







# Data Partitioning

---

- Partitions should be load balanced and spatially coherent.






## Load Balancing/Ghost Cells

---

- Automatic for Structured Meshes.
- Partitioning/ghost cells for unstructured is “manual.”
- Use the D3 filter for unstructured
  - (Filters → Alphabetical → D3)







## **Job Size Rules of Thumb**

---

- **Structured Data**
  - Try for max 20 M cell/processor.
  - Shoot for 5 – 10 M cell/processor.
- **Unstructured Data**
  - Try for max 1 M cell/processor.
  - Shoot for 250 – 500 K cell/processor.



## Rendering Modes

---


- **Still Render**
  - Full detail render.
- **Interactive Render**
  - Sacrifices detail for speed.
  - Provides quick rendering rate.
  - Used when interacting with 3D view.




## Level of Detail (LOD)

---

- Geometric decimation.
- Used only with Interactive Render



Original Data



Divisions: 50x50x50

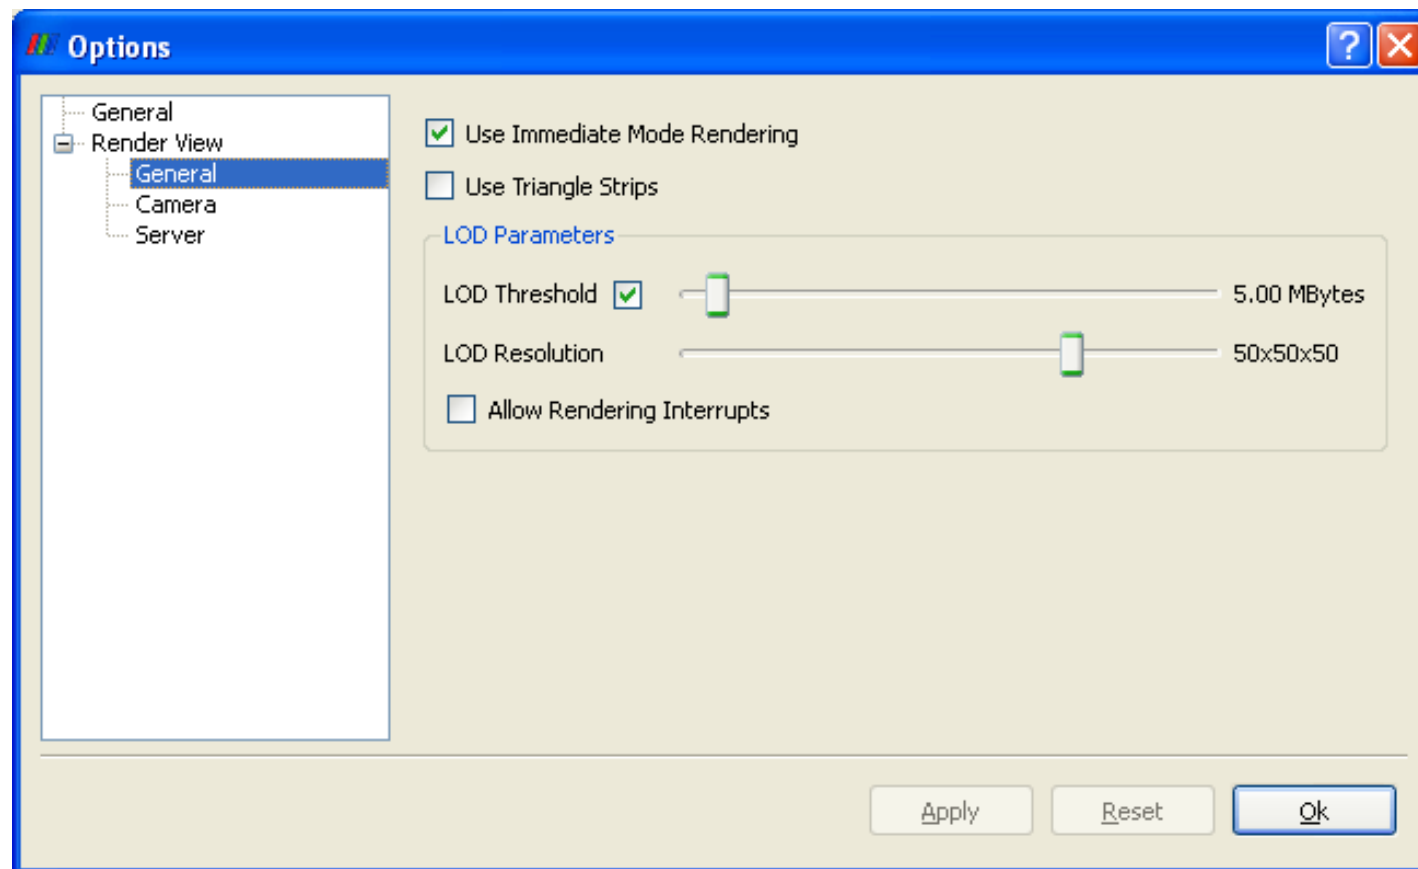


Divisions: 10x10x10



## 3D Rendering Parameters






**Edit → Settings, Render View → General**







# Parallel Rendering

---






# Parallel Rendering





# Tiled Displays

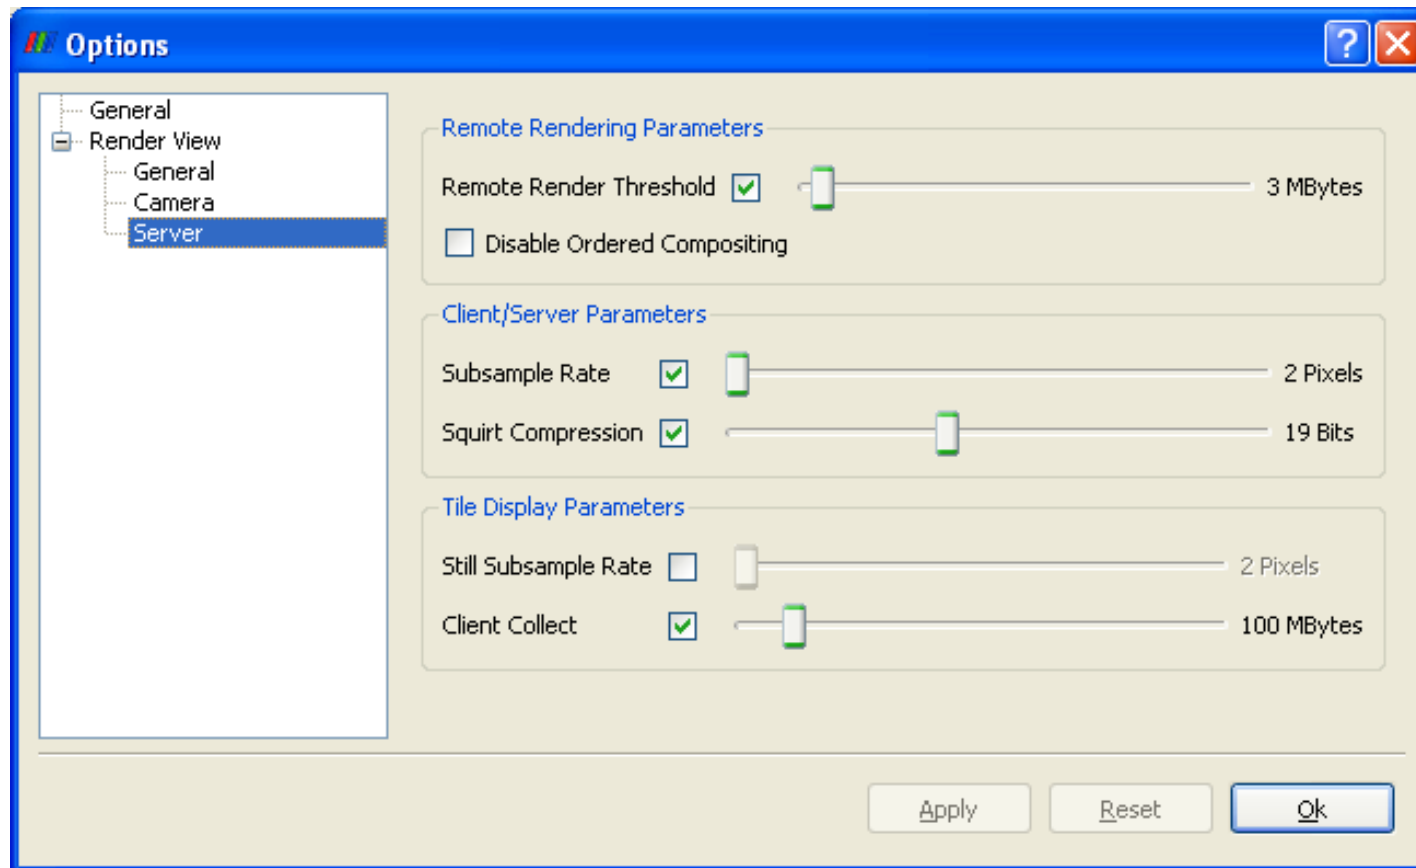
---





# Parallel Rendering Parameters

**Edit → Settings, Render View → Server**







## Parameters for Large Data

---

- Use Immediate Mode Rendering on.
- Use Triangle Strips off.
- Try LOD Threshold *off*.
  - Also try LOD Resolution 10x10x10.
- Always have remote rendering on.
- Turn on subsampling.
  - Try larger subsampling rates.
- Squirt Compression on.