

Streaming Multi-Resolution Distance Visualization

Jon Woodring, Los Alamos National Laboratory

**James P. Ahrens¹, Jonathan Woodring¹, David E. DeMarle²,
John Patchett¹, and Mathew Maltrud¹**

¹Los Alamos National Laboratory

²Kitware, Inc.



UNCLASSIFIED

Operated by Los Alamos National Security, LLC for the U.S. Department of Energy's NNSA

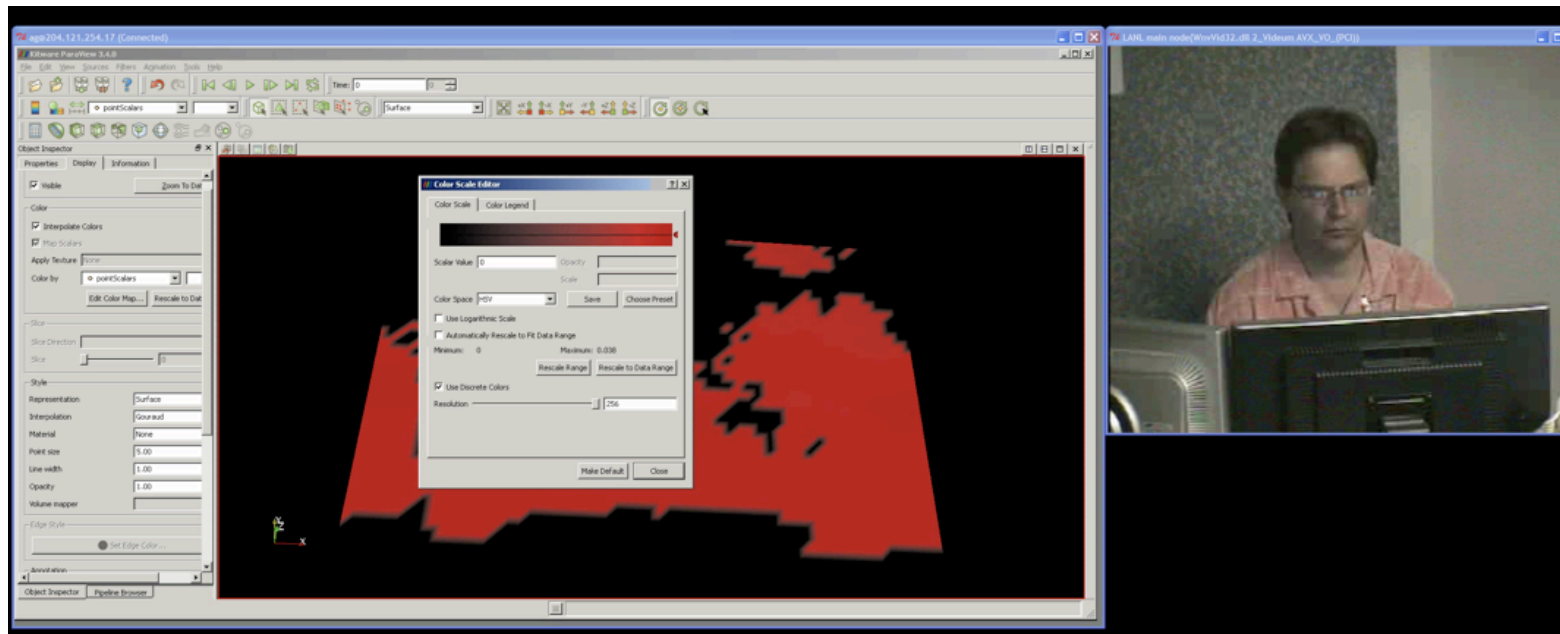


Executive Summary

- **Multi-resolution streaming visualization system for large scale data distance visualization**
 - Representation-based distance visualization (process data, send data, render client-side)
 - Alternative approach to image-based (process data, render data, send images)
 - Send low resolution data initially
 - Incrementally send (stream) increasing resolution data pieces over time and progressively render on the client side
 - Sends pieces in a prioritized manner, culling pieces that do not contribute
 - Implemented in ParaView/VTK and is publically available in the ParaView developer CVS archive
 - Works with most filters – the structural system changes only take place in the reader, renderer, and new pipeline meta-data messages

Remote Data

- **Mat Maltrud works at LANL on the Climate team and runs climate simulations at ORNL on Jaguar**
 - Mat is responsible for generating and analyzing the simulations
- **Other scientists we collaborate with use off-site resources as well**




Remote LARGE Data

- **Using 100 TeraFLOPs of Jaguar (ORNL)**
 - 6 fields at 1.4GB each 20x a day
 - 3600 x 2400 x 42 floats
- **Transfer to LANL would take > 74 hours (~3 days)**
 - ~5 Mbps between LANL and ORNL (this was measured last year, it might have improved slightly since then; many software firewalls)
- **Transferring all of the data from ORNL to LANL will not work!**
 - 250 TeraFLOPs
 - 12 fields
 - 1 PetaFLOP
 - 24 fields and 40x a day = 740 hours (~1 month)

Two Remote Visualization Approaches


■ Server side rendering

- Run data server and render server on the supercomputer – send images



■ Client side rendering

- Run data server on the supercomputer – send geometry
- Render client side



Evaluating Client Side Rendering

- **Image-based distance vis: it works, but...**
 - Completely server side based (dumb client)
 - Frame rate is network latency and bandwidth limited
- **Client side rendering?**
 - Higher potential frame
 - Can render without needing the server (caching)
 - Can still use a render server to deliver imagery if image-based distance visualization is still required
- **Though, this is large data – too big for the client, network, and display...
Is it even practical?**


Subset and Downscale the Data to Fit Displays, Networks, and I/O

Prefix	Mega	Giga	Tera	Peta	Exa
10^n	10^6	10^9	10^{12}	10^{15}	10^{18}
Technology	Displays, networks, I/O		Data sizes and super- computing		

Downscaling
Sampling
Feature Extraction

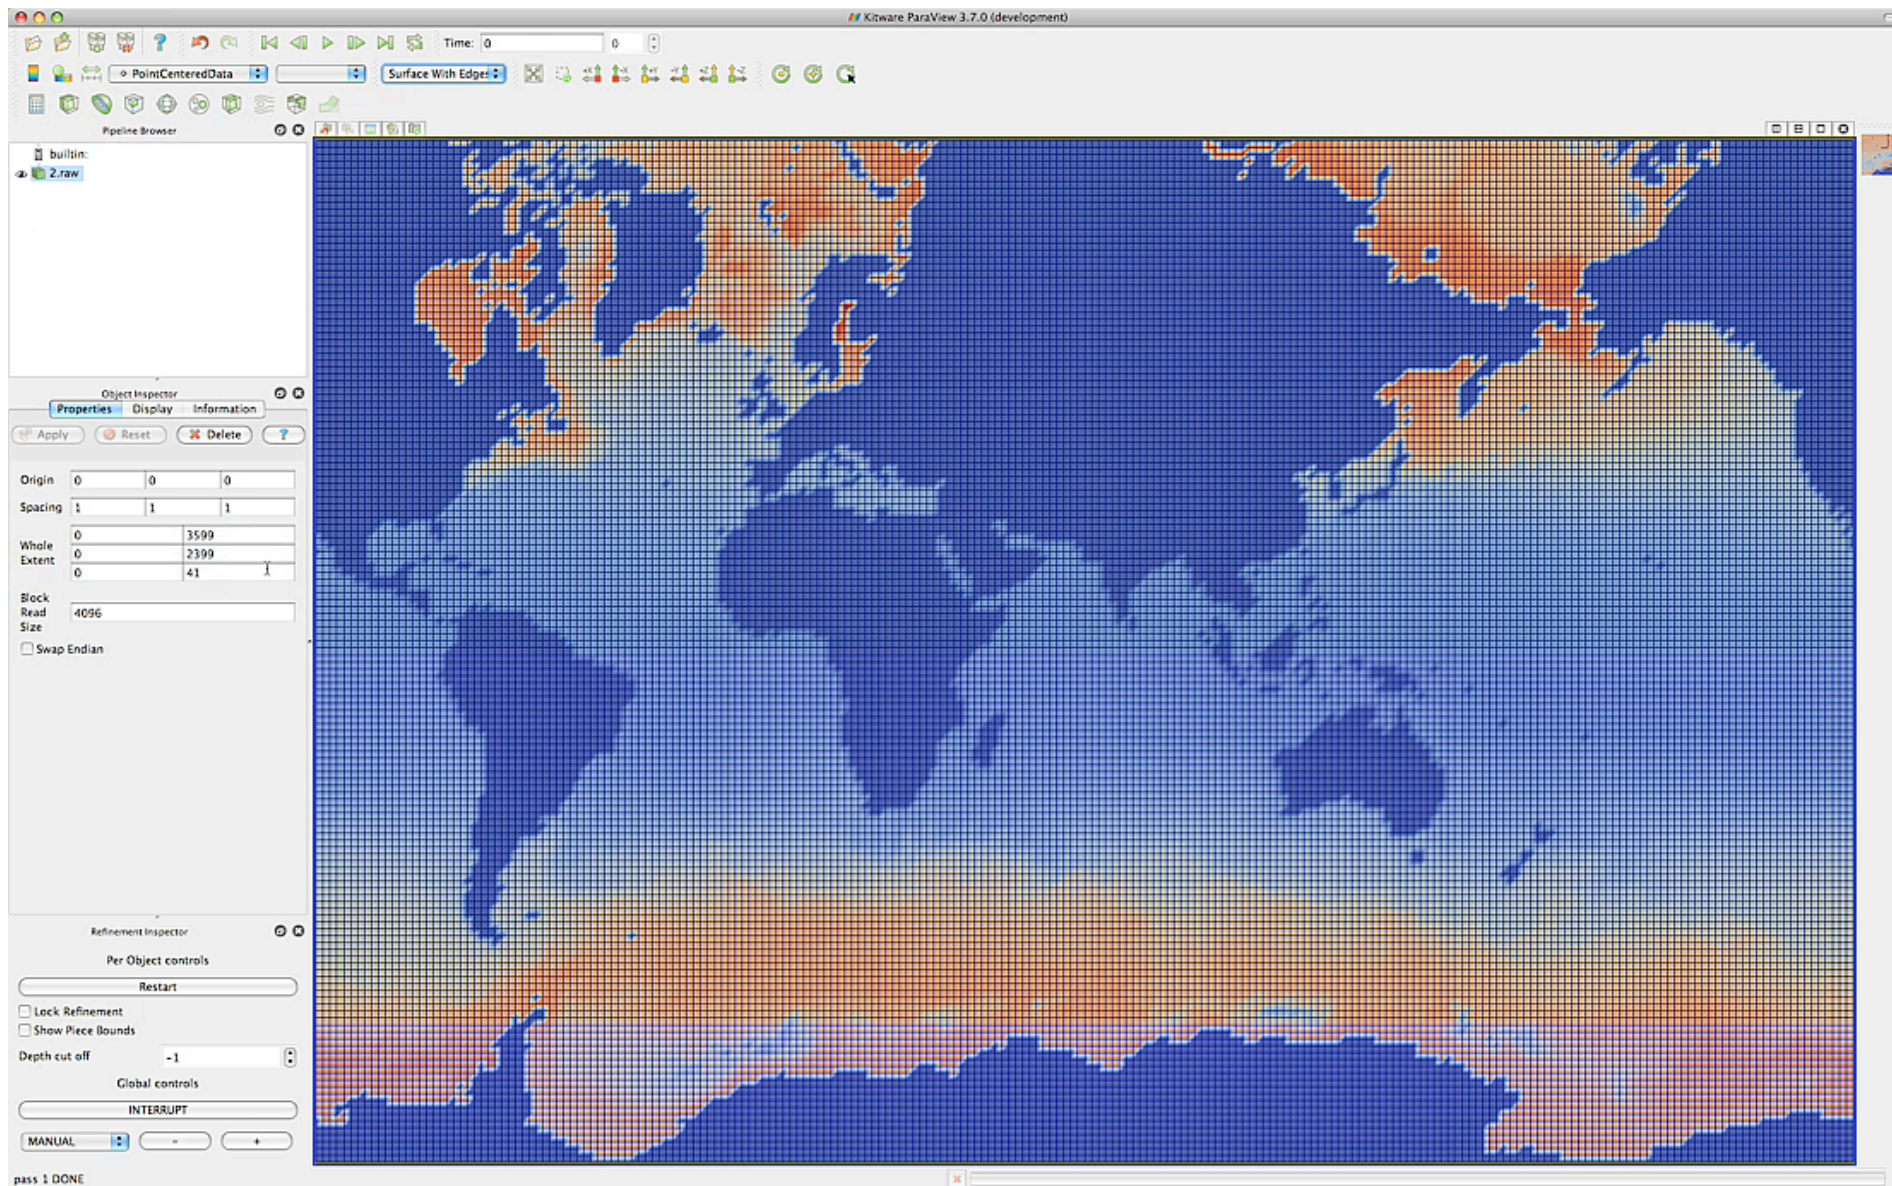
The data has more points than
available display pixels...
Post-sim I/O is expensive...
We need to reduce the data, anyways

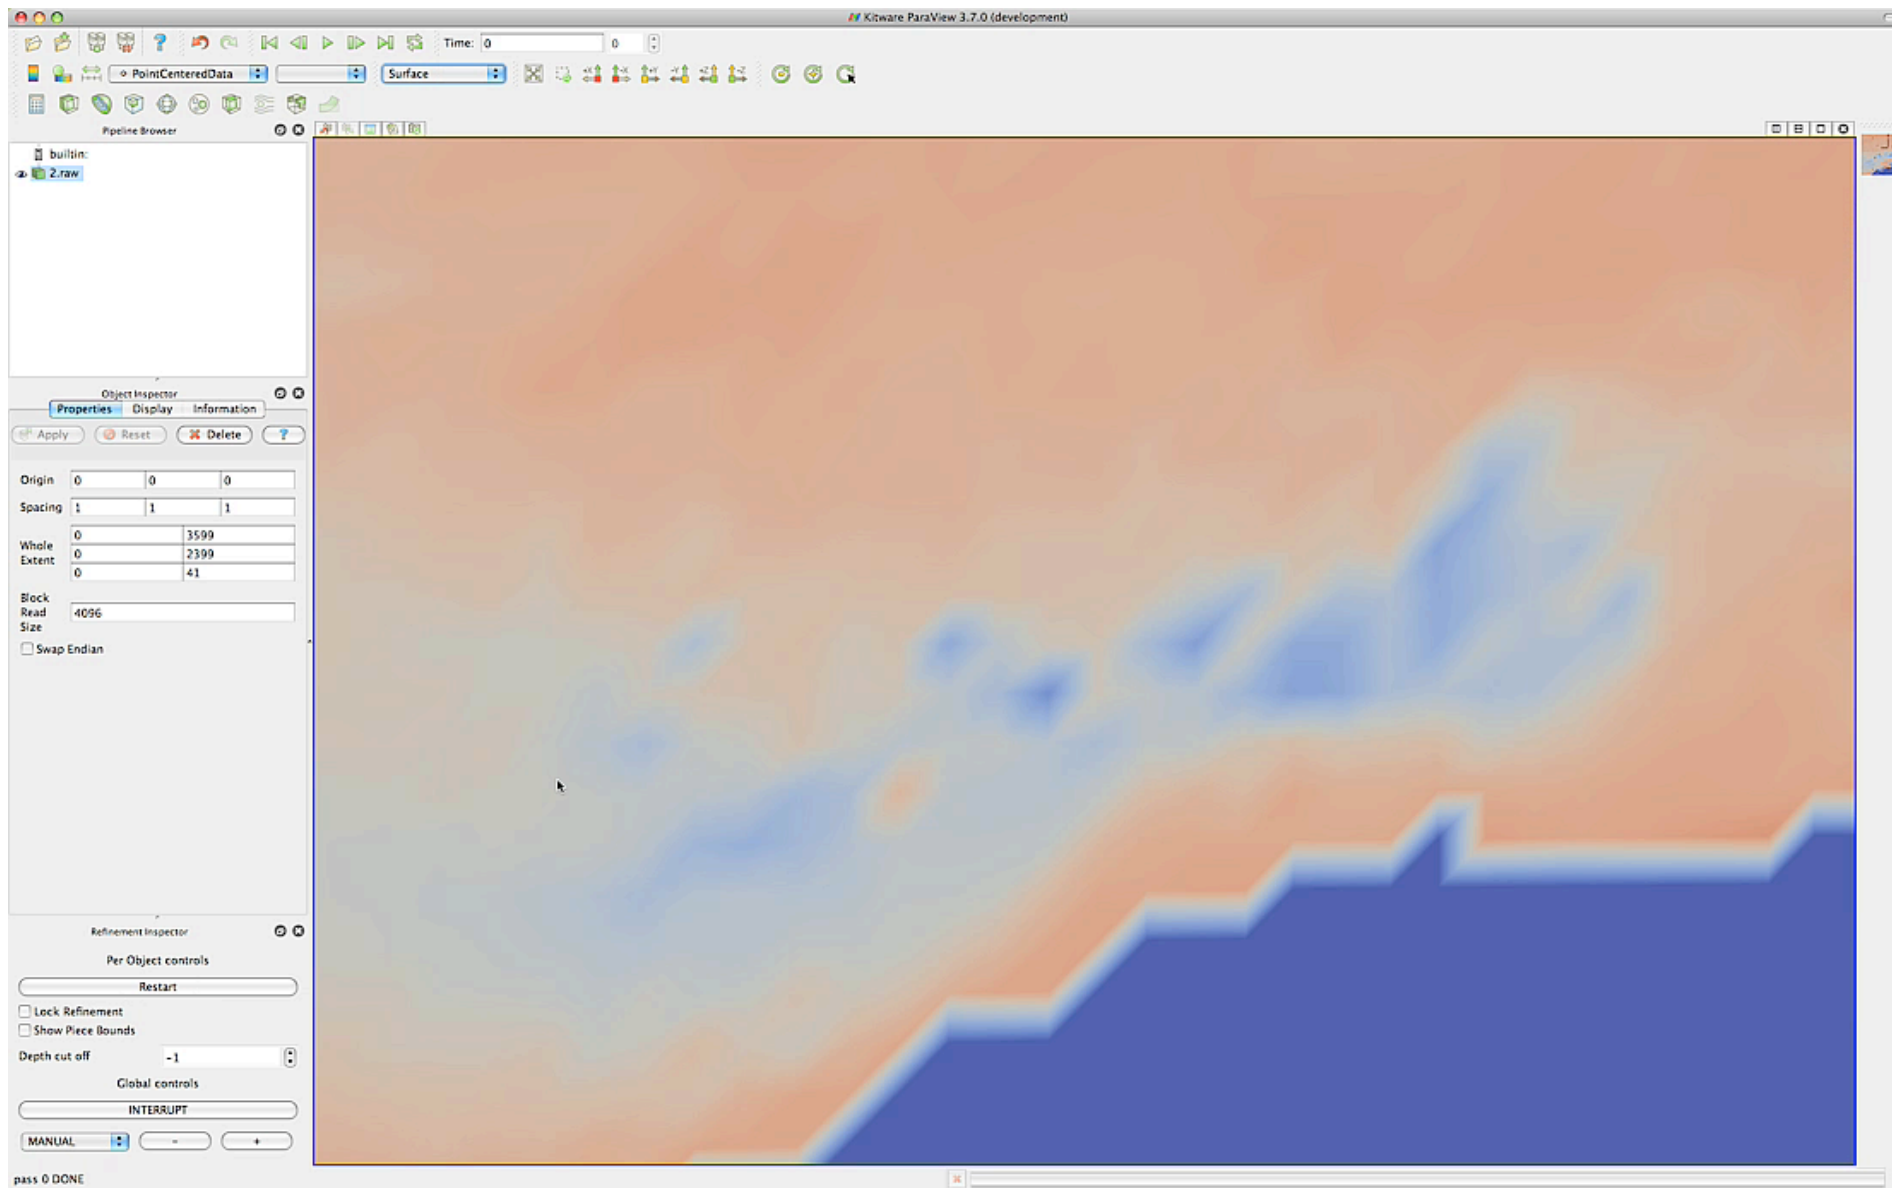
Standard, Streaming, and Adaptive Streaming Pipelines




Multi-resolution and Streaming Related Work

- Norton and Rockwood
- Clyne and Rast
- Pascucci and Frank
- Wang, Gao, Li, and Shen
- LaMar, Hamann, and Joy
- Prohaska, Hutanu, Kahler, and Hege
- Rusinkiewicz and Levoy
- Childs, Duchaineau, and Ma
- Ahrens, Desai, McCormick, Martin, and Woodring







Pipeline Approaches in ParaView



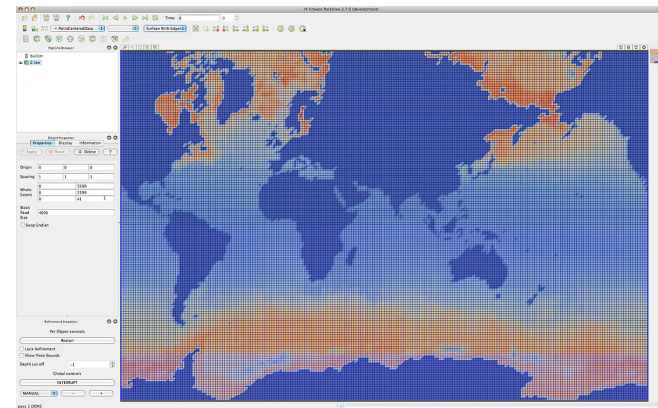
standard



streaming


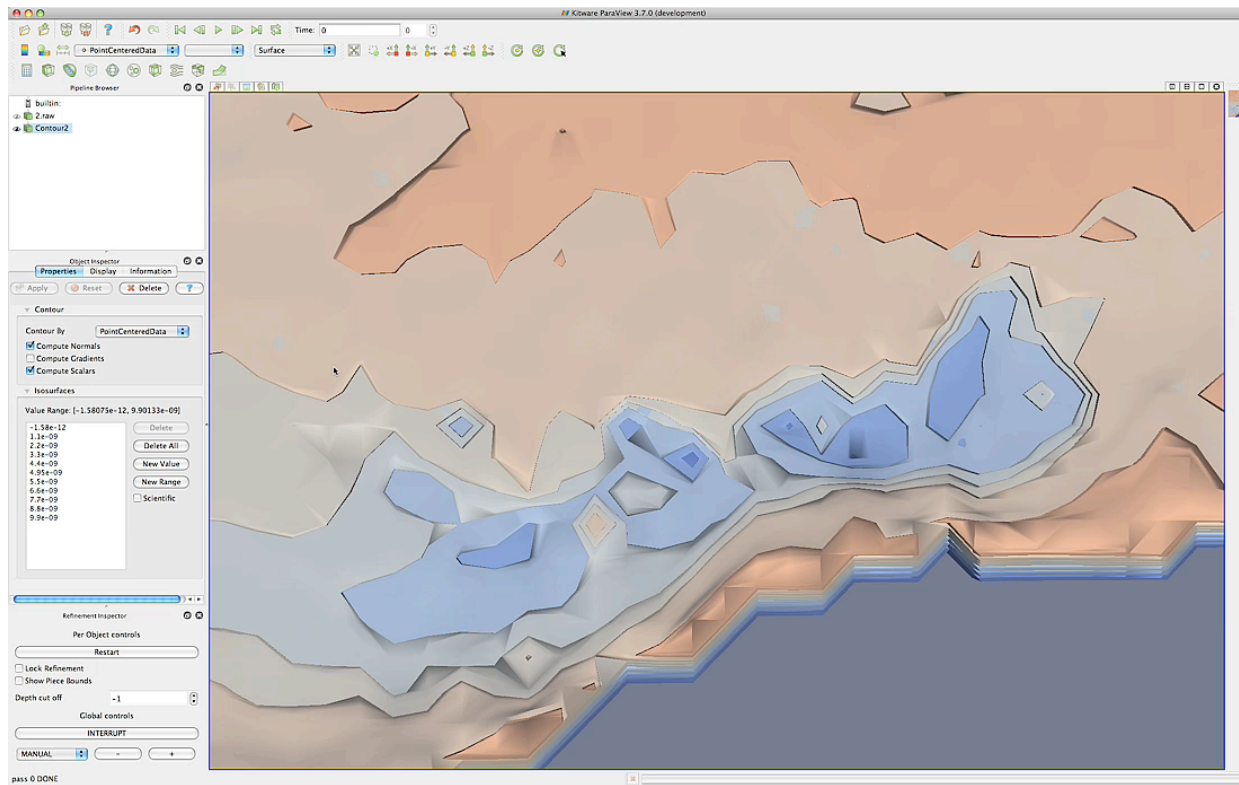


prioritized streaming

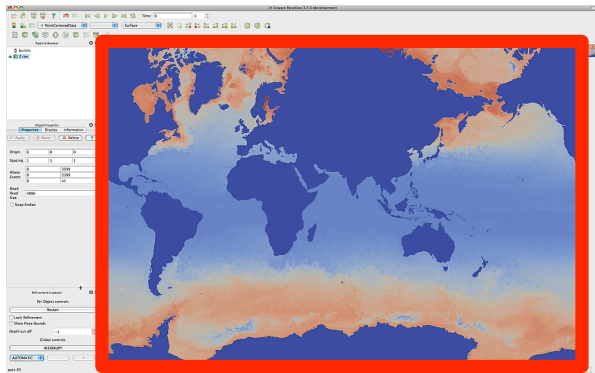


multi-resolution prioritized streaming

Multi-resolution Visualization System





Multi-resolution Prioritized Streaming




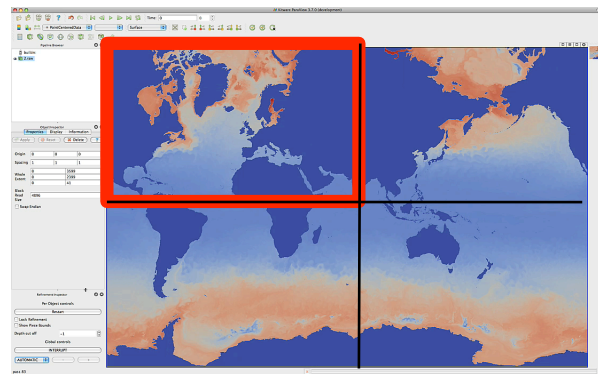
1) Send and render
lowest resolution data

Multi-resolution Prioritized Streaming






- 1) Send and render lowest resolution data
- 2) Virtually split into spatial pieces and prioritize pieces

Multi-resolution Prioritized Streaming



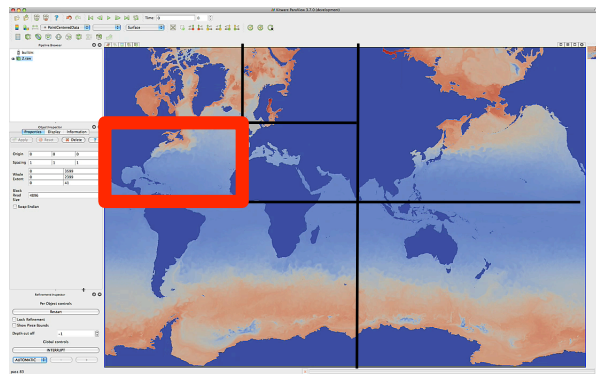
- 1) Send and render lowest resolution data
- 2) Virtually split into spatial pieces and prioritize pieces
- 3) Send and render highest priority piece at higher resolution

Multi-resolution Prioritized Streaming




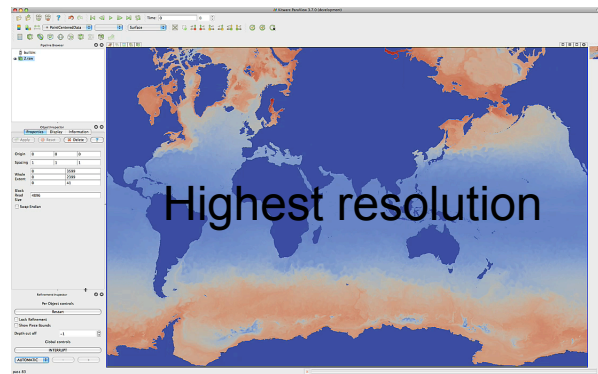
- 1) Send and render lowest resolution data
- 2) Virtually split into spatial pieces and prioritize pieces
- 3) Send and render highest priority piece at higher resolution
- 4) Goto step 2 until the data is at the highest resolution

Multi-resolution Prioritized Streaming



- 1) Send and render lowest resolution data
- 2) Virtually split into spatial pieces and prioritize pieces
- 3) Send and render highest priority piece at higher resolution
- 4) Goto step 2 until the data is at the highest resolution

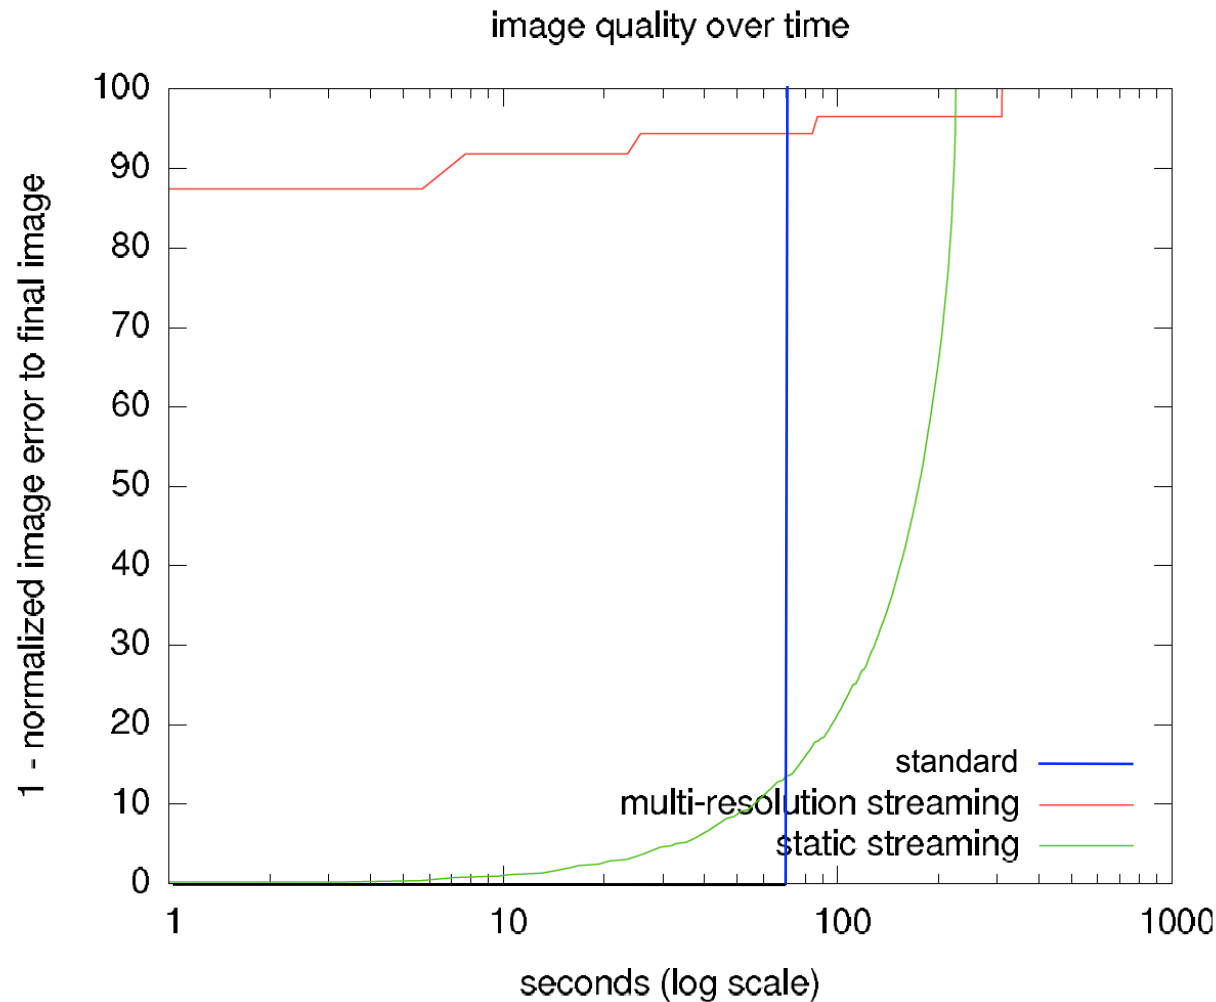
Multi-resolution Prioritized Streaming



ParaView Implementation


- **Progressive multi-resolution renderer (sink)**
 - Implements the high level algorithm on the previous slides
- **Multi-resolution preprocessor (generating multi-resolution tree)**
 - Our implementation for structured data uses strided sampling – fast
 - Doesn't modify the original data – left as-is (highest resolution); worst case uses x1 additional space, realistic cases use 1/3rd or 1/7th additional space
- **Multi-resolution reader (source)**
 - The reader provides data pieces based on resolution request and piece request
 - Uses the preprocessed multi-resolution data for fast linear reads
- **Meta-information keys**
 - RESOLUTION request (what resolution is needed)
 - UPDATE_EXTENT request (what is the spatial extent of the piece needed)
 - PRIORITY keys (for prioritization sorting and culling)

Image Quality over Time for Whole Extent (Single Node) (POP 3600 x 2400 x 42 floats, 10 MBps, 100 ms latency)




Whole Extent (POP data, 100 ms latency) (Single Node)

Total Rendering, Client Rendering, and Send Time

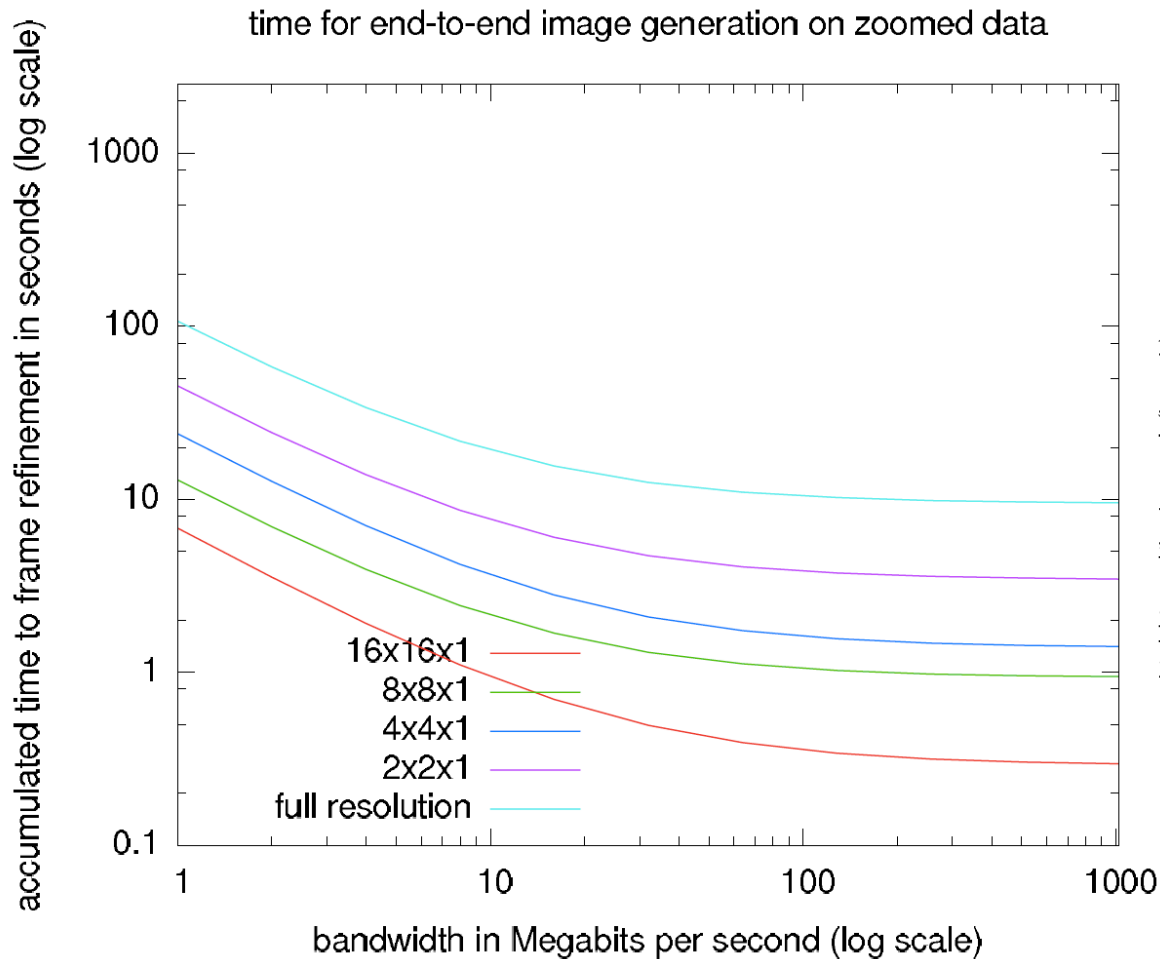


Full Extent	16x16x1	8x8x1	4x4x1	2x2x1	Full
Render	0.03 s	0.10 s	0.38 s	1.4 s	5.6 s




Zoomed In (Culling and Prioritization) (same params)

Total Rendering, Client Rendering, and Send Time



Zoomed	16x16x1	8x8x1	4x4x1	2x2x1	Full
Render	0.03 s	0.03 s	0.05 s	0.09 s	0.27 s



Single Node

Cold Start Read and Write Timings (POP data)

Time to Read Whole File	Time to Read and Create 4 levels	Time to Read and Create 20 levels
30.0 s	46.5 s	150.6 s

Full Extent	16x16x1	8x8x1	4x4x1	2x2x1	Full
Read	0.18 s	0.92 s	5.0 s	11.8 s	35.6 s
Accum. Read	0.18 s	1.1 s	6.1 s	17.9 s	53.5 s

Zoomed	16x16x1	8x8x1	4x4x1	2x2x1	Full
Read	0.18 s	0.47 s	1.4 s	2.8 s	6.1 s
Accum. Read	0.18 s	0.64 s	2.0 s	4.8 s	11.0 s

Contact

- Jim Ahrens: ahrens@lanl.gov
- Jon Woodring: woodring@lanl.gov
- Dave DeMarle, Kitware: dave.demarle@kitware.com
- John Patchett: patchett@lanl.gov
- Mat Maltrud: maltrud@lanl.gov

- Research is funded by Office of Science ASCR