

Next Generation I/O and Visualization Requirements for In Situ Processing

SC 2011 Birds of a Feather

November 15, 2011

Kenneth Moreland
Sandia National Laboratories

Release Marking (e.g. Not Approved for Release, SAND XXXX, etc.)

Who I am, Why I Care

- Spent the last decade in HPC vis
- ParaView as develop/deploy platform
- Before 2007, refused in situ projects




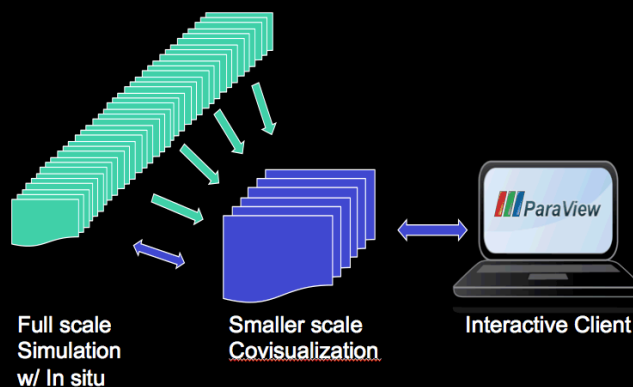
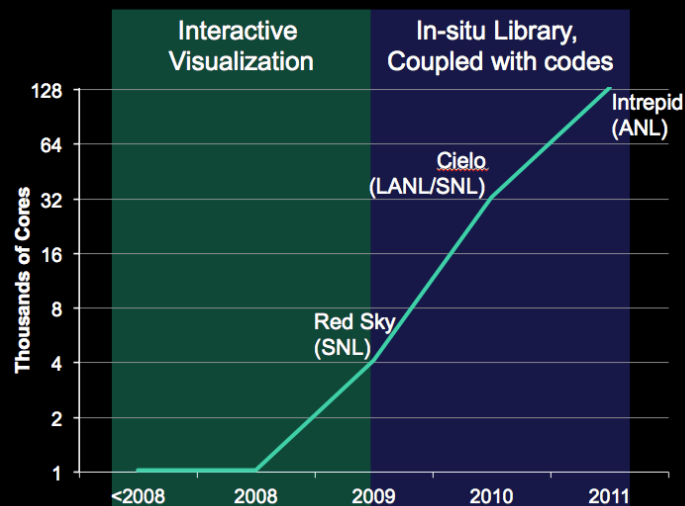
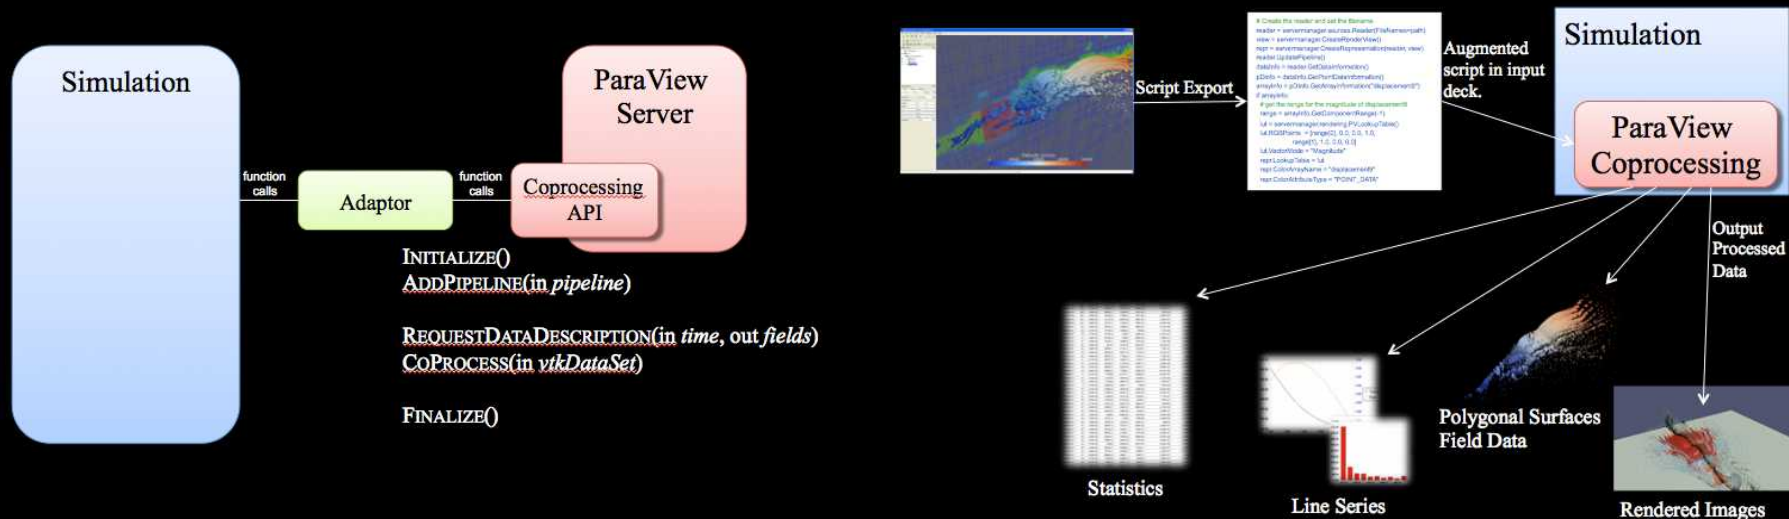
1 mesh per 100 steps
~1% write time

1 image per 1 step
~0.1% write time

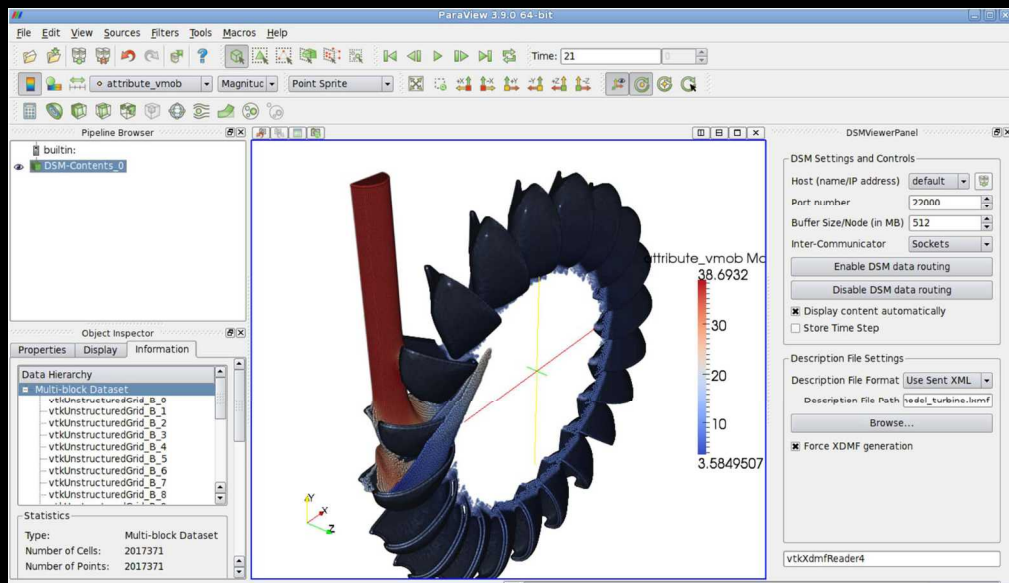
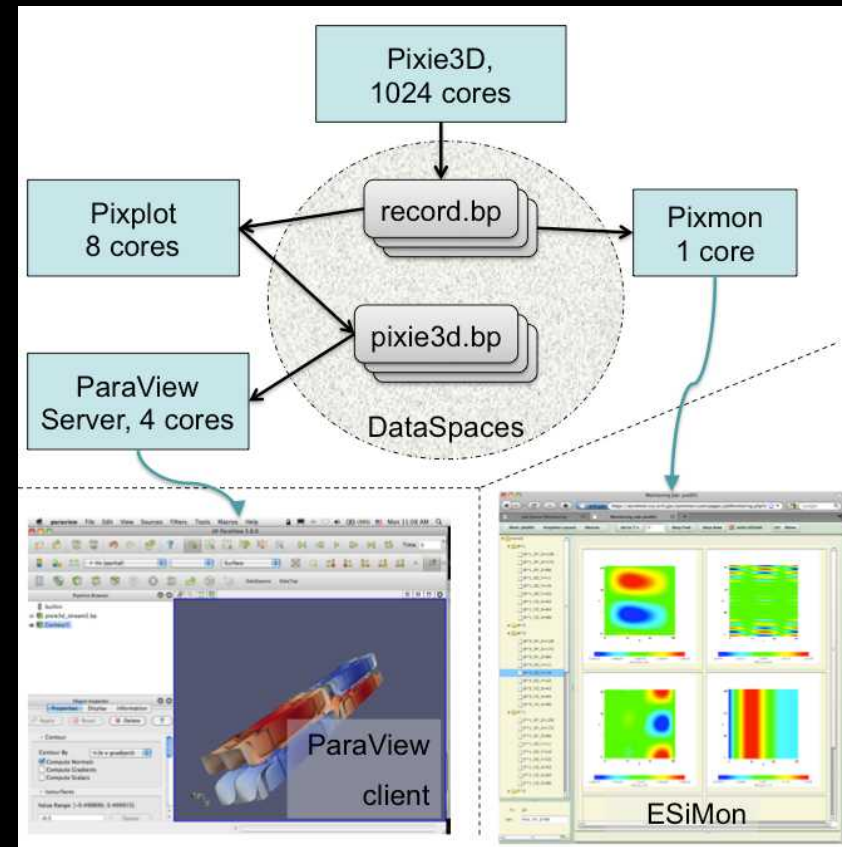
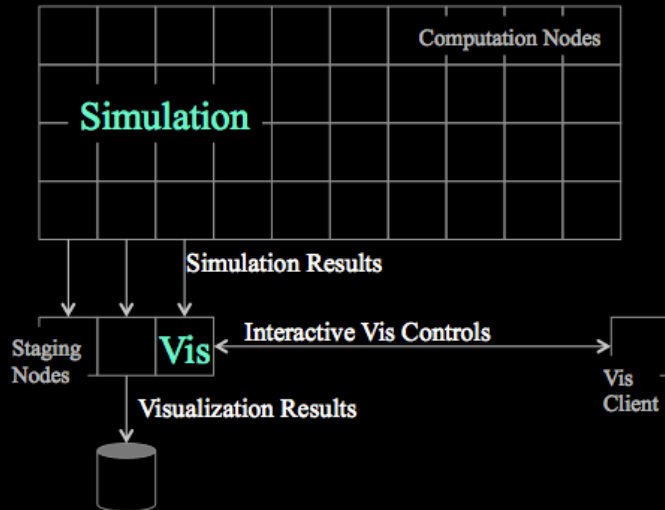
- Trend: storage rates will be insufficient to do science
 - Disk rate ASCI purple: 140 GB/s
 - Disk rate JaguarPF: 200 GB/s

Current Research 1: ParaView Coprocessing Library

<http://paraview.org>



Current Research 2: In Transit Visualization



Gaps

- Memory-sharing in situ
 - Memory footprint.
 - Integrating 15 years vis R&D, > 2 million LOC.
 - Hour long compile can be intimidating.
 - Data layout.
 - Fault propagation.
- In transit in situ
 - Ask Scott.
 - Not production ready.
 - A single interface for data from any simulation.
 - General purpose schema?

The Great Kendini Predicts the Future



- A collection of federated runtime components.
 - (Yeah, I drank that Cool-Aid.)
 - Components can dynamically change seamlessly through a virtual file system.
- I/O libraries move up a level.
 - “Low” level (POSIX, MPI/IO, sockets, pipes, portals) are unmanageable.
 - Move up to higher level (HDF5, ADIOS).
- More forethought in simulation data collection.
 - Cannot collect it all.
 - Becomes more like experimentation.
- JIT simulation results.
 - Rather than store simulation results, run it on demand.
 - Simulations run repetitively as analysis progresses.