



Volume rendering of angular momentum and streamlines of shock wave vector field in supernova dataset, courtesy of John Blondin.



Time lag between first snowfall and green-up in MODIS dataset, courtesy of NASA.



Applications


From astrophysics to climate modeling, we are working one-on-one with scientists and their data to meet their large-scale visualization requirements.



The Argonne Leadership Computing Facility's IBM Blue Gene/P (left) and the Oak Ridge Leadership Computing Facility's Cray XT4 (right). Images courtesy ALCF and NCCS.



Architectural diagram of the Blue Gene/P I/O system



Architectures


At scales of tens of thousands of cores, visualization algorithms need to be tuned to specific architectures, so we study systems in detail, I/O for example.

Parallel Visualization on Leadership Computing Resources


One approach to meeting the increasing demands for analysis and visualization is to perform more of these tasks on supercomputers traditionally reserved for simulations. This can lead to increased performance, reduced cost, and tighter integration of analysis and visualization in computational science. Our team is developing and optimizing visualization algorithms to scale to architectures at over 10,000 cores. Performance is analyzed for ultrascale applications.

Algorithms

Parallelizing algorithms such as volume rendering into pipelines, consisting each of many cores, maximizes performance on these architectures.




The cost of I/O in rendering a time series can be masked by visualizing multiple time steps in parallel pipelines. Each of the pipelines below is further parallelized among multiple nodes. The forwarder daemon runs on the login node and serializes final results.




Analyses


Aggregate and component results are analyzed to determine bottlenecks. Because ultrascale visualization is dominated by I/O, our team devotes considerable effort to its study, both from systems and application perspectives.



Scalability over a variety of data, image, and system sizes.



The relative percentage of time in the stages of volume rendering as a function of system size.



Different numbers of processes performing I/O can affect overall writing performance.