SDM Center Technologies for Accelerating Scientific Discoveries

With the increasing volume and complexity of data produced by ultra-scale simulations and high-throughput experiments, understanding the science is largely hampered by the lack of comprehensive, end-to-end data management solutions ranging from initial data acquisition to final analysis and visualization. The SciDAC-1 Scientific Data Management (SDM) Center succeeded in bringing an initial set of advanced data management technologies to DOE application scientists in astrophysics, climate, fusion, and biology. Equally important, it established collaborations with these scientists to better understand their science as well as their forthcoming data management and data analytics challenges.

Our future focus is on improving the SDM framework to address the needs of ultra-scale science during SciDAC-2. Specifically, we are enhancing and extending our existing tools to allow for more interactivity and fault tolerance when managing scientists' workflows, for better parallelism and feature extraction capabilities in their data analytics operations, and for greater efficiency and functionality in users' interactions with local parallel file systems and remote storage. These improvements are necessary for the scalability and complexity challenges presented by hardware and applications at ultra scale, and are complemented by continued efforts to work with application scientists in various domains.

http://sdmcenter.lbl.gov

PI: A. Shoshani (LBNL)

Participating institutions

- DOE Laboratories:
  ANL, LBNL, LLNL, ORNL, PNNL
- Universities:
  NCSU, NWU, SDSC, UCD, UUtah

Framework and Technologies

The 3-layer framework provides integrated functionality

- **scientific workflow automation**
- **data searching and analysis tools**
- **parallel file I/O and storage management**

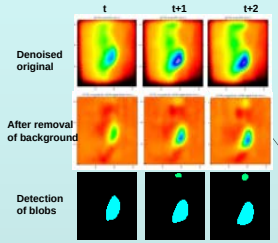

The slides show bright blobs extracted from experimental images from the National Spherical Torus Experiment (NSTX). The blobs are high energy regions. If they hit the torus wall that confines the plasma, it can vaporize. The figure show movement of the blobs over time. Top row is the original image after removing camera noise. Second row is after removal of ambient or background intensity. Third row is the blobs identified by our software. The goal is validate and refine the theory of plasma turbulence.

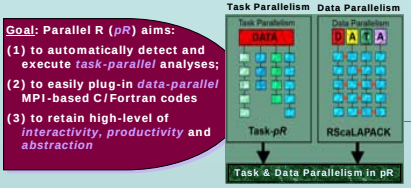

Parallel R lets scientists employ a wide range of statistical analysis routines on high performance architectures without having to deal with the intricacies of parallelizing these routines. Through Parallel R the user can distribute data and carry out the required parallel computation but maintain the same look-and-feel interface of the R system. Two major levels of parallelism are supported: data parallelism (k-means clustering, Principal Component Analysis, Hierarchical Clustering, Distance matrix, Histogram) and task parallelism (Likelihood Maximization, Bootstrap and Jackknife Re-sampling, Markov Chain Monte Carlo, Animations).

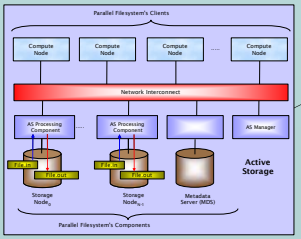

In a traditional cluster systems, I/O-intensive tasks must be performed in the compute nodes. This produces a high volume of network traffic. The Active Storage approach allows moving computations involving data stored in a parallel filesystem from the compute nodes to the storage nodes. Benefits of Active Storage include: low network traffic, local I/O operations, and better overall performance. The SDM center has implemented Active Storage on Lustre and PVFS parallel file systems.

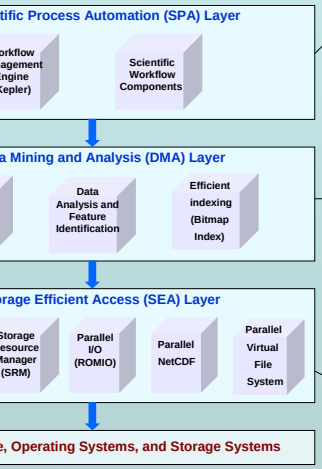

In this example, the Kepler Scientific Workflow Tool is used to automate coupling XGC-0 and M3D codes for the Center for Plasma Edge Simulation (CPES) fusion project. The processing loop within the workflow transfers data regularly from the machine that runs XGC-0 to another machine for equilibrium and linear stability computations. If the linear stability test fails, a job is prepared and submitted to perform nonlinear parallel M3D-MPP computation. Kepler supports task- and pipeline-parallel execution required for this workflow.

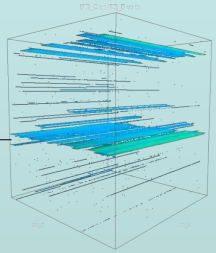

In this example, an efficient bitmap indexing tool, called FastBit, was used to identify malicious network traffic. The query-driven visualization reveals consecutive regions that represent coordinated attacks. The figure shows a 3D histogram over the IP address space and time. FastBit is fast enough to dynamically index the streaming traffic, thus permitting analysis in real-time. FastBit has been applied to various applications, including combustion, astrophysics, and Fusion.

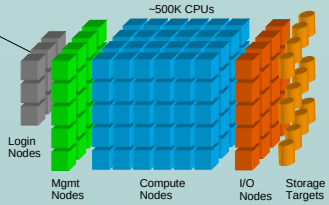

Upcoming systems will incorporate hundreds of thousands of compute processors along with a collection of support nodes. Using POSIX and MPI-IO interfaces, I/O operations are forwarded through a set of I/O nodes to storage targets. On the Argonne Leadership Computing Facility, the PVFS parallel file system, supported under the SDM Center, will provide a system-wide storage space for applications storing hundreds of terabytes of data.

SciDAC
Scientific Discovery through Advanced Computing