SDAV's Software Infrastructure and Tools

Hank Childs, Berk Geveci, and Scott Klasky (SDAV's Software Infrastructure Committee)

Overview

SDAV provides comprehensive expertise in scientific data management, analysis, and visualization aimed at transferring state of the art techniques into operational use by application scientists on leadership-class computing facilities over its five year period. A primary goal of SDAV is to provide technical solutions that are broadly used by the computational science community. These technical solutions will come in the form of software. This poster describes the SDAV Software Plan and also gives an overview of the tools supported by SDAV.

SDAV Software Plan

Our strategy for customer impact is to work directly with application scientists and to assist them by applying the best tools and technologies at our disposal. Further, we will learn from the scientists where our tools fall short. Technical solutions to any shortcomings will be implemented to ensure that our tools overcome mission-critical challenges in the scientific discovery process. These tools will be further developed and improved upon as computing platforms change over the life of SDAV.

This strategy has several implications for software engineering:

- *(i) we will deploy existing tools*
- *(ii) we will extend existing tools to meet customer needs*
- *(iii) we may need to develop new tools if customer have needs where there is no existing software*

SDAV distinguishes the "customer readiness" of our software: production software, emerging software, and prototype software. Our preference is to deploy production software whenever possible.

SDAV recognizes software engineering as an important aspect to deploying quality software. As we develop new capabilities, we endeavor to follow invest in the software engineering activities to keep our software strong.

SDAV's Data Management Tools

ADIOS: ADIOS is a data management end-user tool that deals with parallel I/O for large-scale scientific simulations. The purpose of this tool is to provide fast, adaptable and scalable I/O interfaces so that scientific codes can run highly efficiently across all computing platforms. ADIOS provides a simple, flexible way for scientists to describe the data (either files or streams) in their code that may need to be written, read, or processed outside of the running simulation. ADIOS received a 2013 R&D 100 Award.

ALACRITY Compressed Data Indexing: This tool offers storage-lightweight indexing over compressed scientific data, for the purpose of enabling accelerated response times on exploratory query-driven analysis and visualization. It is designed to function either during post processing or in situ. The ultimate goal is to facilitate scientific discovery by reducing time to results.

Darshan: Darshan is an I/O instrumentation library that collects I/O access pattern information from large-scale applications with minimal overhead. Data collected with Darshan can be leveraged by application users, administrators, and I/O researchers to better understand and tune I/O behavior. Darshan is deployed in production on a variety of systems including major platforms at NERSC and ANL.

DataSpaces: DataSpaces is a data management framework that enables dynamic and asynchronous interactions between applications. It provides support for coupled applications workflows, including loose and tight coupling, data analysis workflows with in-situ and in-transit processing..

FastBit: FastBit is a tool for querying large multidimensional data sets with a large number of numerical variables. It supports subsetting and filtering with predicate conditions over the variable, as well as histograming and similar aggregation operations. Its primary strength is the extensive set of efficient bitmap indexing methods.

FastQuery: FastQuery intends to provide querying capability on top of data files in scientific data file formats including ADIOS, HDF5, netCDF and pNetCDF. This tool gives scientists advanced searching capability without requiring them to load their data into a data management system. FastQuery was optimized to run on a large number of cores.

FlexPath: FlexPath is a component of Adios designed to provide memory-to-memory data transfers and transport-level data reduction through 'data conditioning plugins'. It is designed to aid such things as data management, analysis and visualization.

GLEAN: GLEAN is an I/O acceleration and simulation-time analysis framework.

HDF5: Production quality, high performance parallel I/O library with self-describing data format, and accompanying tool suite.

ISOBAR. *ISOBAR compression provides effective in situ lossless compression to improve end-to-end I/O throughput when writing (or reading) scientific data, including other types of "hard-to-compress" data. ISOBAR compression accomplishes this by identifying highly entropic content within a provided byte buffer ("hard-to-compress" data) and then preconditions, compresses, and writes out the resultant.*

MLOC. MLOC offers multi-level storage layout optimization for scientific datasets to speed up queries with heterogeneous access patterns. It reorganizes datasets to a hybrid multi-level layout scheme to improve the locality for space-constrained, value-constrained scenarios. It supports a byte-level precision-based multi-resolution data storage and access.

Parallel-NetCDF. Parallel netCDF (PnetCDF) is a library providing high-performance parallel I/O while still maintaining file-format compatibility with Unidata's NetCDF. The library provides a data model based on multi-dimensional arrays of typed data, which allows for annotation of datasets and variables with additional information (e.g. timestamps, provenance, or units).

PIDX. PIDX is a high performance parallel I/O library that writes data in hierarchical multi-resolution Z order format suitable for analysis and visualization of large data-sets. The library integrated with the visualization elements also aids in real-time monitoring of large scale simulations.

TALASS. The TALASS toolkit (which is part of the Topologika suite) incorporates both a highly effective streaming computation of feature-based segmentations as well as an interactive environment to explore the temporal evolution of features and their properties.

SDAV's Visualization & Analysis Tools

DAX: Dax supports the fine-grained concurrency for data analysis and visualization algorithms required to drive exascale computing and simplifies the development of parallel visualization algorithms. Dax achieves these goals by encapsulating algorithms in "worklets," which enable an unlimited number of threads without the complications of memory clashes or other race conditions.

DBSCAN: DBSCAN is a density-based clustering algorithm, which is capable of discovering arbitrary shaped clusters and eliminating noise. The algorithm has been used in various scientific areas, such as astrophysics, satellite image segmentation, noise filtering and outlier detection, unsupervised document clustering, and bioinformatics.

DIY: DIY is a library to assist with writing large-scale parallel data analysis programs, with functions for decomposing one or more domains, communicating information within and among domains, and managing I/O. Its lightweight and scalable performance makes it particularly attractive for directly embedded in situ analysis.

EAVL: EAVL is a visualization library that targets the approaching hardware/software ecosystem by (1) updating the traditional data model to handle modern simulation codes and a wider range of data, (2) leveraging an updated data and execution model to achieve necessary computational, I/O, and memory efficiency, and (3) enabling developers to realize these efficiency gains through high productivity programming constructs.

IceT: The Image Composition Engine for Tiles (IceT) is a high-performance sort-last parallel rendering library. IceT is principally used in large-scale scientific visualization tools to provide scalable and interactive rendering.

ImageVis3D: ImageVis3D (and it's companion 'Tuvok' library) is a visualization tool meant to transparently solve the problem of large data visualization. The tool can easily visualize terascale data sets on commodity workstation hardware.

MSCEER: MSCEER (Morse-Smale Complex Extraction Exploration and Reasoning) is a set of libraries and a visualization tool for data analysis and visualization using Morse-Smale complexes, which is part of the Topologika suite. The software provides command-line tools for pre-computing and storing an MS complex, simplifying and extracting features, and interactive exploration and visualization.

OSUFlow: OSUflow is a parallel flow analysis and particle tracing library that supports analysis of large scale flow data sets. The library allows for scalable computation of streamlines, pathlines, stream surfaces, and streak surfaces, and can be deployed on DOE's leadership computing facility.

ParaView: ParaView is a general-purpose scientific data analysis and visualization tool. It was designed from grounds up to work with large data. It supports an MPI-enabled data-parallel back-end (server) and client-server computing.

ParaView-Catalyst: *Catalyst is an open-source, scalable data analysis and visualization library designed to be tightly coupled with simulation codes. It can be directly embedded into parallel simulation codes to perform in situ analysis at run time. Furthermore, it can be coupled with the ParaView In Situ Analysis Framework to perform run-time visualization of data extracts and steering of the data analysis pipeline.*

PISTON: PISTON is a framework for creating such operators that employs the data-parallel programming model, which enables portable performance over many architectures. PISTON consists of a user-extensible set of operators (such as isosurface, threshold, and cut surfaces) implemented using this model.

Ultravis-P: Ultravis-P supports the visualization and statistical analysis of large-scale particle data through a suite of techniques that allow users to create, cluster, classify, and visualize particle trajectories, in both phase and physical space. The tool uses GPU clusters to maintain interactivity.

VisIt: VisIt is a reachly featured visualization and analysis tool for massive data sets. It has a client-server design and has been scaled up to tens of thousands of cores and to operate on data sets with trillions of cells. VisIt reads over one hundred file formats and is used heavily in the DOE.

ViSUS: The ViSUS Client ad Apache server is used to provide fast access to large scale data on a variety of architectures including Windows PC, Mac, Linux and IPhone/IPad. The viewer has limited visualization capabilities but is able to deal with massive volumes using very modest computing resources.

VTK: The Visualization Toolkit aims to provide general-purpose data analysis and visualization to users and developers. It is the foundation for highly popular production tools including VisIt and ParaView. It is also used by many scientists as part of their daily workflow either through its scripting interface or as part of a bigger C++ framework.

To learn more, visit
http://www.sdav-scidac.org

SDAV
Scalable Data Management, Analysis, and Visualization

U.S. DEPARTMENT OF ENERGY | Office of Science