Integration of Scientific Analyses and Storage

Reagan W. Moore
University of North Carolina at Chapel Hill

A major challenge for "Big Data and Extreme-scale Computing" is the massive amount of data that is being generated by modern observational and simulation technologies. The indexing of the data to support discovery, the management of the migration of the data across file systems and archives, the verification of data integrity, and the arrangement of the data in collections becomes onerous as collection sizes increase. To minimize the amount of labor needed to manage the data, data grids are used to enforce management policies, automate administrative functions, verify assessment criteria, and analyze data products. The knowledge required to implement the management functions can be captured as procedures that are enforced through policies managed by a data grid.

The traditional approach for data manipulation relies on the movement of the data to a compute engine. An emerging capability is the movement of computation to storage systems to enable processing within storage controllers. This provides a scalable approach to the management and processing of exabyte data sets. The mantra for this style of computing is that "a disk spin that is not utilized to support analysis is a revolution lost forever."

The technology that enables processing of data within storage controllers is a policy-based data management system such as the integrated Rule Oriented Data System. Policy-based data grids are middleware that is installed at each location where data are stored. The middleware translates the desired action to the local storage protocol. The middleware also includes a rule engine and a local rule base at each storage location. The data grid middleware is instrumented with policy enforcement points to control actions that are performed upon the data. Policy enforcement points are typically associated with the ingestion of a file into a shared collection, or with enforcement of access controls on file access, or with file deletion, or with creation of descriptive metadata. When a policy enforcment point is tripped, the rule engine selects the appropriate policy to apply.

Within a storage controller, a virtual machine environment can run the data grid server, including the rule engine that automates the application of workflows to data streams. A simple example is the automation of feature detection within a file. On deposition into a storage system, a procedure is invoked to parse the file for the desired feature. If the feature is found, metadata is linked to the file and stored within a metadata catalog. The catalog can then be queried to identify the files that contain the desired feature.

This approach translates management and analysis policies into computer actionable rules. Each policy manages the application of procedures that encode the knowledge needed to either manage the data, analyze the data, or manipulate the data. The policies that are applied to the file can also be captured in workflows that are in turn registered as objects within the data grid. The workflows can be shared, modified, and re-executed. The input and output files associated with the workflow can be shared and modified. This enables reproducibility of data-driven research. Collaborators can modify a workflow, compare with prior execution results, and iterate analyses until the research question is answered.

An implication of the ability to automate the application of procedures to files at the storage location is that now knowledge can be managed along with information and data. This is a third step in the evolution of data management systems:

1. Data management – originally supported by file systems with a standard interface such as Posix I/O operations.
2. Information management – supported by digital libraries which define a context to associate with each file. The context is queried through SQL.
3. Knowledge management – supported by policy-based data grids which implement knowledge through computer executable procedures under the control of computer actionable rules.

Knowledge management requires the application of procedures that evaluate relationships between information attributes. The relationships can be quite sophisticated, ranging from semantic/logical, to temporal/procedural, to spatial/structural, to functional/algorithmic, to systemic/epistemological. Thus the transformation of a data format can be viewed as the application of structural relationships to transform the data set to the desired data type. The relationships are captured in procedures that can be executed under the control of the data grid. Within the iRODS data grid, the procedures are composed by chaining together basic functions, called micro-services that exchange information either through structures in memory, or through persistent state information stored in a metadata catalog, or through files, or across networks to remote locations through serialization of the in-memory structures.

For extreme computing on massive data collections, the integration of computing with storage can lead to a scalable solution. As the size of the data increases, the associated computing capability also increases through the addition of more storage controllers. There are restrictions on this approach. The complexity (number of operations per byte) of the knowledge procedures that are applied must be less than the ratio of the storage controller processing speed to the storage controller bandwidth; and the granularity of features that are manipulated should be less than the available memory buffer size. For application of procedures that require more processing, the data grid will need to move the data to a computing resource. Examples of these technologies have been demonstrated in the iRODS software development project (NSF grant OCI-1032732), NSF DataNet Federation Consortium (NSF grant OCI-0940841), and the NSF EarthCube Layered Architecture concept award (NSF grant EAR-1239678).