

Advancing the Power & Utility of Server-Side Aggregation

Panel summary from the ESIP Summer Meeting, 16-July-2015, by Dave Fulker

This paper reflects a panel discussion held on day three of the Earth System Information Partners (ESIP) Federation's 2015 Summer Meeting in Pacific Grove, California. Organized and convened by **Dave Fulker** and **James Gallagher** of OPeNDAP, Inc., the session¹ was supported by NASA/GSFC under Raytheon Co. contract number NNG10HP02.

The invited panelists were **Jon Currey** of NASA/Langley, **Mike Folk** of The HDF Group, **Mohan Ramamurthy** of UCAR/Unidata and **Bob Simons** of NOAA's Environmental Research Division. Each provided a slide set that may be found on the ESIP Web site (<http://commons.esipfed.org/node/7993>), but Currey was unable to attend, so his slides (co-authored by Aron Bartle of Mechdyne) were delivered by Gallagher. Fulker introduced the panelists and set the context in four bullet points:

- Data systems often contain files or images (i.e., granules) that may be accessed only independently, even when kept in collections of highly similar entities.
- Such granularity typically reflects how data are *collected*, unrelated or even contrary to data *utility* for many use cases.
- Panel members are experts on diverse approaches to and needs for actual or virtual aggregation (in conjunction with a variety of data-access methodologies).
- EOSDIS recently invested in enhancements to OpenDAP's aggregation features. (This included means by which users may access potentially-huge lists of granules, applying the same query to each and receiving the results as a single, aggregated response.)

Below are this author's impressions from the presentations, meant to illustrate the range of viewpoints expressed and the richness of aggregation capabilities that may accompany the provision of data as a service.

Mike Folk, citing examples and metaphors spanning the spectrum from nature and art to Legos and satellite imagery, spoke about the ways aggregation—and disaggregation—can add value. His title emphasized an HDF (Hierarchical Data Format) perspective, but his observations were broadly applicable.

- Aggregation often is prerequisite to seeing the big picture. "If a tool only can work with a single object, [actual or virtual] aggregation can combine together in a single object all the information we want the tool to use."
- *The whole may well be more than the sum of its parts*, as with artistic mosaics, e.g.
- Aggregation can yield efficiencies of storage and transport, as with railroads, e.g.

¹ The session abstract read: *Tying this session to the overall ESIP-Summer-Meeting theme of "data-driven community resilience," we enlist speakers with expertise on: the diverse ways that large collections of data granules can now or potentially be used for scientific, economic or societal advancement. We specifically will aim to enlist speakers from communities where important needs for data (such as very long time series, e.g.) may be orthogonal to how the data are collected and stored. To set the stage for discussion, we describe new OPeNDAP aggregation features (developed for NASA/EOSDIS) that we think may open doors to unprecedented forms of data-mining and statistical supports for resiliency.*

- Effective aggregation often is preceded by wise or clever forms of *disaggregation*, as with Legos, e.g., and this can lead to greater flexibility of use.
- The Hierarchical Data Format (HDF) is based on (computer science) notions of aggregation, including hierarchies, arrays, groups, chunks and even strings, some nested.
- The HDF toolset includes a number of aggregation-related capabilities.
 - Virtual datasets may be accessed as single objects even as their contents comprise multiple HDF files.
 - Swath data from polar-orbiting satellites reflect an especially rich set of HDF-based aggregation forms, due to the complexities of the associated geolocation and profiling information.
 - "Nagg is a tool for rearranging Suomi-NPP data granules from existing files to create new files with [differing aggregation or packaging arrangements]."
 - Work on a "scalable HDF server" is underway, built around the notion that a suitable cloud-based object store (in contrast to a more traditional HDF file store) might enable more flexible and efficient forms of disaggregation and aggregation.

Folk suggested the possibility of an "aggregation reference model," and asked: "Could there be common vocabulary and model that spans the [various] types of aggregation?"

Bob Simons discussed aggregation as realized in the Environmental Research Division's Data Access Program (ERDDAP). Widely used, especially for in-situ observations, ERDDAP is an extended implementation of DAP (the OPeNDAP protocol). ERDDAP treats a given dataset either as a set of multidimensional gridded variables (like many DAP servers) or as a database-like table (which DAP calls a "sequence"). Because ERDDAP has aggregation capabilities for gridded data that are similar to those in other DAP realizations, Simons focused on ERDDAP's support for tabular data, a less common type in DAP servers.

- For tabular datasets in ERDDAP:
 - Users may employ DAP's predicate-style query constraints (e.g., latitude>30 AND latitude<35) to request subsets, so selecting rows from a table works very much like SQL does for relational databases.
 - The tabular data model is very general and may be applied to a range of discrete sampling geometry (DSG) data (points, time-series, profiles, trajectories, time-series profiles, trajectory profiles) and non-geospatial data (e.g., from lab experiments). ERDDAP retains and exploits the higher-level structures of DSG data.
 - ERDDAP can, on-the-fly, *output* data in various user-specified formats (including .csv, .json, .mat, .nc, .kml, .png, .pdf) some of which may be quite highly structured, such as netCDF files employing the CF DSG contiguous ragged array conventions.
 - ERDDAP, even more than other DAP servers, accepts a variety of *input* data types. This, combined with the preceding point, implies that ERDDAP serves as middleware, mapping source data types onto those required by end-user tools and programs.
- Under ERDDAP, a tabular dataset is always represented as a single, database-like table, without nesting or references to other tables (external keys). Normalized databases thus

are denormalized (joined) before being served via ERDDAP. The implications of this design choice are:

- For tabular data, aggregation is merely the concatenation of tables, i.e, the appending of rows from tables with matching columns.
- Compliance with the CF (Climate and Forecast) conventions implies consistency of spatial-temporal query constraints, with or without aggregation.
- The absence of external keys simplifies both aggregation and querying, as neither ever requires "join" operations.

Simons concluded by pointing out simplifications that result from aggregation, linking these to the conference theme of data-driven community resilience.

Mohan Ramamurthy opened with background information on Unidata, the UCAR program he directs and from which he drew aggregation examples. His talk covered distinct scientific motivations for aggregation, including:

- Ensemble forecasts are sets of outputs from multiple model runs, each producing the same set of 4-dimensional variables.
- Climate and forecast model runs often are distinct (multi-variable) datasets for each of many (i.e., hundreds) time steps.
- Even a single variable may be organized as datasets that cover distinct time intervals or geographic regions.
- For each of the above, the end-user ideal may be to view the entire set of entities as a single (aggregated) entity, where, e.g., each variable is as a single array containing values drawn from all constituents. Such views facilitate time-series analysis and visualization via meteograms and mosaics that encompass many tiles.

For some 15 years Unidata has offered *virtual* aggregations that users access via OPeNDAP services. The NetCDF Markup Language (NcML) was extended to include means for *configuring* these. The supported forms have evolved to now include:

- *JoinExisting*: variables of the same name in different files are concatenated along their outer dimension (i.e., an existing coordinate), thereby lengthening that dimension.
- *JoinNew*: resembling lamination, variables of the same name in different files are concatenated along a *new* outer dimension and a new coordinate variable is created.
- *Union*: all the dimensions, attributes, and variables in multiple NetCDF files are joined into a single accessible entity.

Ramamurthy discussed recent advances for Point Feature Collections and Forecast Model Run Collections (FMRCs). A diagram from John Caron illustrated the ways that output files from a set of model runs (with differing initialization times) may be aggregated to assemble a) all files from a single model run (i.e., from one initialization), b) all files that share the same "valid time" (i.e., that estimate a given point in real time), c) all files for same forecast-offset (6 hours into each forecast, e.g.) and d) the set of best-estimates (i.e., minimum forecast offsets) for all hours in the collection. The complexities of such

aggregations have led to indexing schemes and other efficiency enhancements in the THREDDS Data Server (TDS).

The use of semantics and the hiding of coordinate-system irregularities (such as may be encountered with radar images, e.g.) rounded out Ramamurthy's presentation.

James Gallagher, as proxy for Jon Currey and Aron Bartle, discussed aggregation in conjunction with a set of (OPeNDAP) server functions developed at NASA Langley as part of the Multi-Instrument Inter-Calibration (MIIC) project. In addition to inter-calibration, per se, Currey's talk touched on validating level-2 retrievals and surface-site data mining (to discover deep convection clouds, e.g.).

In each of these contexts, the challenges are exacerbated by needs to aggregate data *from distinct spacecraft* (unlike the more traditional pattern—under OPeNDAP—of aggregating data with identical structures and schemas). As indicated below, this has implications for how the needed data are selected and processed as well as how they are aggregated.

- For inter-calibration and validation studies, comparison data must be drawn from sensors viewing (roughly) the same portion of Earth at the same time with the same viewing angle, i.e., the ray-paths of the observations are nearly collinear, as are those of the sun upon the observational footprints.
- Hence (in key inter-calibration and validation contexts) the data-retrieval work-flows necessarily entail the prediction of "events," i.e., occurrences of suitable alignments among multiple combinations of spacecraft orbits.
- The relative infrequency of such events greatly limits the amount of relevant data, but the associated data-transfers are reduced only if the data-retrieval services support non-trivial subset-selection, sufficient to exploit footprints from the alignment events.
- Such footprints often cross granule boundaries and do not fall along natural (row/column) delineations within the images, so efficient data-retrieval requires server processing as well as aggregation.
- When inter-calibration, validation or data-mining studies involve convolutions or statistical functions, such as histograms, the benefits from (pre-retrieval) server processing can reduce data transfers by *several orders* of magnitude.

Curry's talk concluded by reiterating the importance of efficient data access *across multi-agency data centers*. He also summarized how MIIC's tiered architecture meshed with the mechanisms that Hyrax (OPeNDAP's data server) offers for adding tailored server functions, thus enabling the aggregation and matching of data from multiple platforms and dissimilar instruments.

Audience discussion and conclusions

This panel² drew one of the larger audiences of the ESIP conference, but little time remained for questions or answers following the presentations. One audience member

² This session and some of the motivating advances were supported by NASA/GSFC under Raytheon Co. contract number NNG10HP02C.

observed that metadata too may be aggregated, with benefits similar to those of aggregating the underlying data.

Conclusions from the session are elusive, in part because the presentations were so beautifully distinct. Perhaps a common theme was Folk's point about disaggregation often preceding aggregation, as all other talks mentioned similar concepts. However this author is unable to see an aggregation reference model (or anything similarly definitive) arising from the excitingly diverse work presently underway.