IVOA Support & Protocols for File-Oriented Catalog Data

Gregory Dubois-Felsmann
Caltech/IPAC
(Rubin Observatory/LSST & NASA-SPHEREx)

IVOA Virtual InterOp
19 November 2020

File-Oriented Catalog Data – conversation starter

- Speaking directly from concerns/interests of Rubin Observatory/LSST and of IPAC
  - Issues seem broadly applicable, though

- Bulk processing of wide-area catalog data is becoming increasingly important even as the sizes of the catalogs grow (billions to tens of billions of objects are becoming the baseline, e.g., Gaia and ZTF now, soon LSST and SPHEREx)
  - Machine learning
  - Computation of derived characteristics (e.g., photo-Z) for an entire catalog
  - Matched-template searches of time series data or spectral data (e.g., SPHEREx, ~1-5B spectra)
  - Cross-matching of multiple catalogs and joint analysis

File-oriented catalog data processing

- Contemporary data-science tooling uses shared-nothing, often file-oriented approaches
  - Parallel processing systems like Spark, Dask
  - Column-store data formats for optimized I/O for "table scan" processing, e.g., Apache Parquet
  - Viewed as a peer to providing data in more traditional database query systems parquet.apache.org
  - Data are often spatially sharded into files in order to support the parallelism
- Rich ecosystem of analysis packages that can be used once the environment is in place
- Many projects are planning server-side analysis environments – "Science Platforms" – in part to support these analysis models
  - Their resources will always be limited and people will want to bring the data to their own systems
- Joint analysis of data from multiple projects / missions is key to the 2020s science
  - Interoperability will help to facilitate data exchange

Bulk download becomes more user-facing and more widely needed

- Data publishers have always provided means for “bulk download” of catalog data
  - For example: https://irsa.ipac.caltech.edu/data/download/
  - Relatively small number of wide-area surveys with datasets large enough that they couldn’t be exported as a single file
  - Often done in a slightly different way for each project

  - Often these systems were used either by other data centers who wished to mirror the datasets, or by highly motivated individuals who had an overriding scientific need
  - The datasets were rarely used “in place” via live access to the (e.g., FTP) addresses of the files

- All this is changing
  - More users
  - More users of multiple datasets simultaneously
  - Datasets held in the cloud and much more suitable for in-place access

The picture is shifting…

- from one in which big piles of CSV-type files are dropped in a directory on an FTP server or Web site and then used by experts as an intermediate toward re-creating databases on a user's system, limited to a small number of high-value datasets for which users are willing to do unique work in reading documentation,

- to one in which end users interact directly with the (e.g., Parquet) files, and wish to do so with standardized data science tooling in Python or R, or with GUI tools of their choice such as TOPCAT, Firefly, and Aladin.

Coming soon…

- The Rubin Observatory plans to make the major catalogs from the LSST survey available through TAP, to be sure, but also as sets of Parquet files, available for analysis in a Rubin-hosted Science Platform, but also for external access.

- IRSA is developing both hosted and externally available environments for bulk analysis of its catalog holdings in Parquet format.

- The time scales of planned work in both institutions are such that discussions initiated this month and carried out over the months to come would be directly valuable to us.

Five challenges for us to pursue

All of these trends create a strong sense that improving interoperability for file-oriented catalog holdings would be a valuable area for conversation and possible standardization within the IVOA.

- Discovery of the existence of large file-oriented catalog datasets;
  - Registration of the equivalence of a file-oriented view of a dataset and a corresponding TAP view;
- Metadata service for the content and spatio-temporal coverage of the individual files within such datasets, especially when non-compact spatial sharding is involved, such as in AXS, creating tension with the ObsCore STC-region-based coverage model;
- Query capabilities within such metadata services, allowing users to identify the subset of files which cover a region of interest to them (perhaps along the lines of MOC-based queries);
- Means for applying the valuable infrastructure of VOTable and DataLink in documenting the schemas of columnar data, while enabling the use of Parquet or other comparable future formats for the efficient access to the actual data; and
- Openness of the resulting conventions / standards to the implementation of interoperable APIs in commonly used languages in the community such as Python and R.

Cuts across Apps, DAL, DM – we look forward to organizing joint efforts through the WGs

Discovery of the existence of large file-oriented catalog datasets

- Registration of the equivalence of a file-oriented view of a dataset and a corresponding TAP view;

Metadata service

- … for the content and spatio-temporal coverage of the individual files within such datasets, especially when non-compact spatial sharding is involved, such as in AXS, creating tension with the ObsCore STC-region-based coverage model

- ObsCore is a good basic model for metadata on a set of observations and on how to access them remotely, *but*
  - Time, wavelength, instrument/facility
- Spatial coverage model is limited by an STC-S-ish set of primitives available for the s_region attribute
- Important current tools like AXS use non-compact spatial sharding – no way to represent

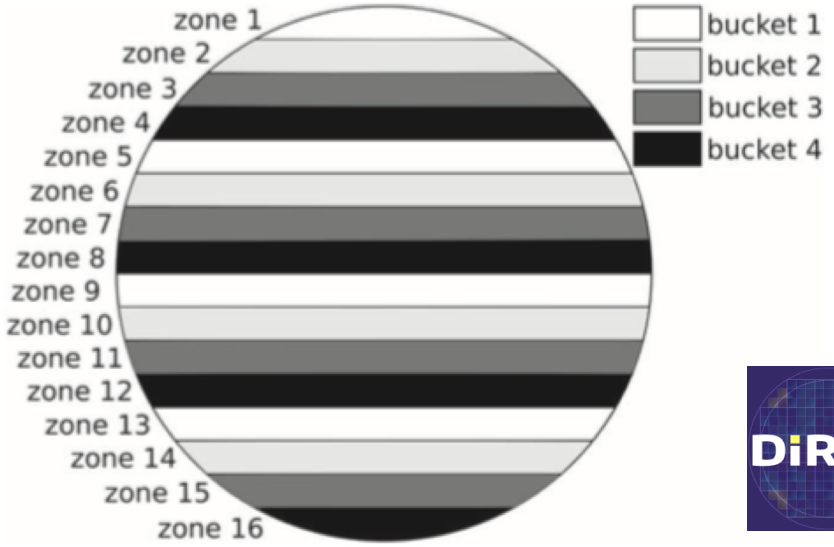

Figure 3. Partitioning the sky into zones and placing zones into buckets. The example shows the sky partitioned into 16 horizontal zones. Objects from each zone get placed into buckets sequentially. In reality, zones are much narrower and are counted in thousands.

axs.readthedocs.io
arXiv:1905.09034

Metadata service – spatial coverage definition

- Some datasets will be sharded in overlapping tiles (conventional)
- Some will be organized by HEALPix, HTM, or other all-sky binnings
- Some will be highly non-compact

- Should we consider MOC for a general (conservative) solution enabling query? Can we add to ObsCore a way to represent coverage in MOC?

- Augmented by an interoperable way to report the precise shard definitions?

Query capabilities

- Within such metadata services, allow users to identify the subset of files which cover a region of interest to them (perhaps along the lines of MOC-based queries)
  - MOCserver relevant here?
  - Application of SIAv2-style query capabilities?

Schema definition

- Means for applying the valuable infrastructure of VOTable and DataLink in documenting the schemas of columnar data, while enabling the use of Parquet or other comparable future formats for the efficient access to the actual data

  - Parquet as VOTable data payload? (Two-file solution)
  - Standardized convention for injecting a VOTable-style schema (as an XML blob?) into a Parquet file?

Think about language APIs from the beginning

- Openness of the resulting conventions / standards to the implementation of inter-operable APIs in commonly used languages in the community such as Python and R.

Next steps

- Our projects have a sense of urgency about this
- This is a highly cross-WG topic, affecting Apps, DAL, and DM
- Can we kick off coordinated activity in the WGs over the next 6 months, with the aim of making a concrete proposal or set of proposals for the next InterOp to consider?
- Can the TCG assist us in organizing such an effort?