An Inventory of Indexes Used In Astronomical Archive Databases

Version 1.1

(March 13, 2015)

G. Greene (Space Telescope Science Institute)
G. B. Berriman and J. C. Good (NASA Infrared Processing and Analysis Center)
T. Donaldson and T. Rogers (Space Telescope Science Institute)

Abstract

One of the most complex challenges facing the future of the NASA archive data centers is providing rapid discovery of, and access to, very large data sets. Modern data collections typically contain comprehensive metadata descriptions, including spatial coverage, and are typically managed with large-scale archival systems and databases. These require intelligent and scalable queries to search and discover datasets. This report provides a high level inventory of the current indexing and query optimization practices in astronomy archives, with emphasis on the NASA archive data centers. The report will inform the technical activities of a comparative study of the performance of database indexing schemes that will be performed over the next two years to provide recommendations on optimum indexing practices.

Keywords: Databases Spherical Geometry, Spatial Index, Region, Footprints, Coverage Maps

Introduction

Indexing can be broadly defined as methodologies used to efficiently search datasets. In the context of databases, indexes are structures stored within databases that improve the speed of retrieval of records from database tables ([HNP95]). In the context of spatial searches of astronomical databases, it is appropriate to expand the definition to include related techniques, such as coverage statistics, that aid in identifying datasets and spatial coverage within larger collections. This is in line with Database Management System (DBMS) use, where column coverage histograms are often used in query planning to determine which of several potential indices to use for a particular query.

This report will inventory the current state of indexing implemented in astronomical archives and data centers, with particular emphasis on datasets at NASA-funded centers and archives. IPAC and MAST will collaborate on a comparative study of the performance of indexing schemes, with the goal of providing recommendations on the optimum use of each indexing method. The study as funded calls for a study of two schemes, described in more detail below. These schemes are a HEALPIX sky-tessellation scheme (the "Multi Order Coverage" mechanism used at CDS [MOC14]) and the other is an R-tree based scheme, the "VO Inventory" mechanism developed at IPAC as part of the VAO. The findings of this report will inform the detailed technical activities of that study. Virtual Observatory (VO) Registry records include collections of data sources with key metadata such as sky coverage, spatial scale, waveband, temporal coverage, data type, publisher, and much more, containing more than a billion records. The study will also investigate whether integrating registry coverage metadata with indexing schemes will accelerate access to data.

Overview

The comparative study described above will evaluate the performance, applicability and scalability of spatial indexing schemes implemented primarily at the NASA data centers. We will document comparative performance of the indexing schemes, describe installation trade-offs, and deliver recommendations for optimal implementation strategies, such as whether the indices should be installed as part of the registry standard coverage metadata description or locally near the datasets. Individual NASA archives will be responsible for implementing such updates to their indexing mechanisms, as they deem necessary to achieve optimum performance. We will make the results of this investigation available to the larger astronomical community. We will also work with the International Virtual Observatory Alliance (IVOA) Technical Coordination Group on the possible impact on IVOA standards.

While our primary focus is on performance of position based indices, our effort will consider other key data attributes which require indices for best performance. What might be called "smart indexing" mechanisms supported by modern commercial and open-source databases are widely used in the NASA archives. The Hierarchical Triangular Mesh (HTM) and the Multi-Order Coverage (MOC) scheme, a HEALPix-based ([Gor84]), sky-mesh indexing mechanism developed at CDS, are both in use at MAST ([MOC14]). There has been substantial progress in developing a MOC IVOA standard. Spherical index libraries integrated into databases are in use at both HEASARC (PostgreSQL PGSphere) and MAST (SQL Server Spatial Lib) which offer high performance for positional queries. IPAC uses HTM [Sza05, Bud10]) for point-source catalogs but does not integrate it into the DBMS. Both HTM and HEALPix are sky-partitioning quadrature schemes that store indices in binary trees (B-trees). A scheme using R-trees ([Gut84]) that stores rectangle-based indices in memory-mapped files was developed for the NVO and the VAO and is used at IPAC to support image metadata searches for large datasets such as Spitzer. This scheme, known as VOInventory, is adaptive to the density of records and will index spatially extended sources (*e.g.* images). Quadrature schemes involve less set-up time and storage and respond better to dynamic datasets, but are best used for point source tables.

The various technologies for implementation of classical indexing used by astronomical data centers are described in more detail below. These fall primarily into two categories: tessellation schemes (spatial geometric partitioning) and database indices. There is a third category of spatial information, which defines sky coverage and region statistics. While not strictly indexing in the search sense, it is closely related and valuable technology and is included in this study as a key way of characterizing datasets as a whole. In many instances, more than one of these methods may be used in combination to create the full data location methodology.

The choice of indexing scheme and details of implementation will be explored in our study to understand how to optimize their performance. This includes choices such as how deeply the indexing is integrated into the DBMS, which specific tessellation scheme or R-Tree splitting / combining algorithms are used, the issues associated with extended objects, and how to use region statistics to augment spatial indexing.

This study will draw heavily on the real-world experience of NASA data archive centers but at the same time we recognize that the GIS community heavily uses spatial indexing for terrestrial spherically based datasets and will take this into consideration where relevant.

Science Applications and Use Cases For Indexing Methodologies

The approach for development of an indexed astronomical data repository is driven by the data content (volume, heterogeneity, *etc.*), application dependencies, and the organizational goals for both internal systems optimization and end use.

1. Search Use Cases and Data Mining strategy

We outline here starting points in defining astronomical science use cases of data discovery and/or retrieval that are dependent on the implementation of spatial indexing. Throughout the study of the NASA archive data center indexing methods, these use cases will drive our development strategy for optimizing search and access of data using indices. .

- **Basic Searchable Index Units include:**

  A. Catalogs of point sources.
     - Individual records.
     - Spatial regions with statistics.

  B. Extended Regions.
     - Individual images (observation regions), fields of view, user-defined areas, extended astronomical objects.
     - Collected subsets of images.
     - General coverage information for a whole collection.
     - Regions with statistics (*e.g.* counts, average brightness, *etc.*)

Our study will focus on the basic indexing units outlined above. As resources permit, we will consider more complex uses involving joins of spatial regions and the use of standards for spatial coordinate representation which are derivative of the basic unit implementation.

- **Advanced Spatial Geometry**

A. Regions Combined
    - Unions
    - Intersections
B. Derived Quantities
    - Region Area
C. Spatial Reference Systems
    - Space Time Coordinate (STC) – the IVOA standard for region representation

- **Discovery vs. Retrieval**

There is a distinction between the discovery of data and retrieval of data from an indexed search. Starting with the assumption that the user is performing data searches over wide areas, initial searches generally explore coverage, and usually exploit coverage maps in tandem with tessellation schemes. Finer grained queries for specific datasets with parametric constraints or sub-filtered lists will benefit from different types of indices, usually embedded in relational databases.

Users usually respond to coverage results with a *request* to retrieve data. Retrieval speeds depends on the details of the implementation of the indexing scheme as well as the type of service and how it packages and returns data. For point sources, requests return a catalog of database records, usually one record per source. For extended regions, requests return a set of image URLs or a list of observations that map to a set of images or spectra. With this investigation we will be characterizing the current state of supported APIs for utilizing indexes.

2. Performance and Optimization

Performance information of the NASA indexed data will be assessed for identified indexed science data collection(s). The measurement of performance and process for generating these metrics will be explained in detail. We also will describe limitations where documented.

3. Cross-correlation

Catalog cross-comparison can be optimized with the use of indices. With this study we will not be developing correlation schemes. However we will be able to document the underlying index and describe how this potentially benefits this science application use case.

Indexing Methodologies

1. Sky Tessellation

Hierarchical tessellation schemes are one of the primary methods for spatially indexing spherical geometric data and therefore a natural fit for astronomical data measurements on the celestial sphere. These are often used as a primary index for pruning big datasets to smaller subsets of regions.

Tessellation is achieved by partitioning of spatial cells. In the case of astronomical data, the celestial sphere is broken down into smaller and smaller sub-cells and each record is assigned a cell value (or identifier). Occasionally multiple cell values are assigned for extended objects. These cell values are B+Tree-indexed. Searching involves making a list of cells from a geometrical sky constraint and asking the DBMS for any records with matching cell IDs. Detailed filtering follows but the indexing governs performance; filtering speed (whether done inside or outside the DBMS) is secondary.

Examples of tessellation schemes currently in practice at the NASA archive data centers include the use of the HTM, HEALPix, declination zoning, and custom rectangular or area binning. The use of HEALPix tessellation has been documented in an IVOA standard, the HealPix Multi-Order Coverage (MOC) map [MOC14]. Applications which typically take advantage of the use of tessellation schemes include visualization and graphical representations of regions. For example, the MAST Data Discovery portal Astroview takes advantage of the HEALPix tessellation which provides the 4 vertice polygon region representation to graphically display areas of coverage for observational data. Pre-generated images can be quickly displayed with HEALPix identifier tagging by fast lookup for spatial regions in the tessellated regions. The HTM tessellation scheme is also used in a similar convention.

The Microsoft World-Wide Telescope (WWT) uses a custom tessellation scheme called TOAST (Tessellated Octahedral Adaptive Subdivision Transform), which is essentially an adaptation of HTM.

These tessellation schemes do not allow for exact region descriptions. Locations are specific only down to the level of the size of the smallest cell defined. Thus they are valuable for identifying potential candidates, but a final determination must be done with secondary filtering.

2. Database Indices

The use of spatial indexing in relation to RDBMS may have varying implementations depending on whether the index scheme is associated with the database rather than integrated into the database. There are a large variety of these indices available, ranging from open source RDBMS to commercially available solutions. The choice depends on the trade-offs and decisions for locking into a feature set available for a given application domain.

The NASA/IPAC Infrared Science Archive (IRSA) and the NASA Exoplanet Archive have implemented R-Trees in conjunction with database indices to optimize searches for images. These methods are inherently N-dimensional and can rapidly find potential matches (or more accurately exclude non-matching branches) through tree pruning. Because many of the tables in IRSA are large and dynamic, it has developed a custom tree- and file-based index called "ChunkIndex" that is applicable to point-sources and that updates records for large tables without re-building the indices from scratch. IRSA has also employed a file-based implementation of HTM called "TinyHTM," used to index and serve the Planck Time Ordered Information data products. The NASA Exoplanet Archive has incorporated HTM into their SQLite database, used to serve tables of the properties and characteristics of exoplanets and their host stars.

The NASA/IPAC Extragalactic Database (NED) is currently using the PostgreSQL object-relational database system, and Bayer-tree (B-Tree) indexes and extensions. NED has improved performance over time by optimizing the use of these indexes in various ways, including partitioning and constraining, at the cost of additional maintenance and storage. Most commercial DBMSs use R-Trees, though often as an add-on or separate product:

- Oracle
- MySQL
- Informix
- Ingres
- Sybase
- SQLite
    - SpatiaLite is spatial extension of SQLite with advanced spatial querying (SQLite already has R-Tree support).

STScI MAST is currently implementing a data model, Common Archive Observational Model (CAOM), and populating observation spatial region footprints using SQL Server Spatial database indices. This is a combined internal DBMS B-Tree and geometric tessellation scheme.

HEASARC uses PostgreSQL PGSphere, which employs. GiST (Generalized Search Tree; http://www.sai.msu.su/~megera/postgres/gist/), an indexing structure introduced by Hellerstein et al. [HNP95] The best way to look at GiST is as a generalization of R-Tree that can be applied to other indexing schemes (including B+Tree) [Gut84]. Postgres claims the GiST implementation is as fast as native R-Tree but for our purposes we may consider it as another implementation of the same general algorithm.

3. Coverage Maps

Coverage "maps" and coverage statistics can be approached in several ways and different techniques are appropriate to different use cases. The study will examine several of these as resources permit:

- Simple all-sky (or region) density diagrams. An example is an Aitoff projection FITS image where the pixel values are a count of the number of sources in the pixel. Such maps are good as a visual reference for the user and uniform sets of these can be used to identify regions with high densities of data in multiple datasets. A binary version of this (data/no data without count) shows general coverage.

- Hierarchical schemes can work the same way to define coverage/counts but have the additional advantage that level in the tree can contain cumulative information. This speeds up large region processing since traversal of the tree to gather information can be short-cut for any intermediate node that is fully included in the region of interest. This approach works for any of the tessellation schemes mentioned and for R-Trees as well.
- Both of the above are limited to treating the lowest level "cells" as atomic units. Though one can approximate fractional statistics based on cell/region-of-interest overlap geometries, this is only a guess and often wrong. An alternate approach involves exact geometries (outlines of images or bounding polygons around sets of points). Such regions can be merged upward into exact bounding polygons for sets. The down side for this is that computation is slow, especially for heavily fragmented sets (*e.g.* a large set of non-overlapping images) but this can be ameliorated through combination with an associated hierarchical scheme.

In general coverage information serves two purposes. It can be used to quickly determine the appropriateness of a given dataset (*e.g.* whether to bother with a spatial search at all) and it can be used as an alternative to spatial searching when only basic information (and especially approximate information) is all that is needed.

4. Very Large Scale Indexing Strategies

This study may consider future effort in exploring technologies that extend beyond the classical indexing schemes defined above. The use of Hadoop clusters and HIVE for management of very large scale datasets is of interest, but beyond the scope of the study.

Initial Inventory of Indexed Data Collections

Table 1: NASA Archive Data Center Indexed Data Collections

Archive/Data Center	Index Method	Example Data Sets
IRSA	HTM tessellation and Oracle B-Tree database indices for point source catalogs and selected image metadata. R-Tree for e.g. Spitzer,. External to DBMS. ChunkIndex for dynamic databases, e.g. WISE. TinyHTM for Planck.	Mission Datasets: • WISE • 2MASS • IRAS • Planck • Herschel • Spitzer • COSMOS • *etc.* Catalogs • 2MASS • WISE • USNO
NASA Exoplanet Archive	HTM tessellation and Oracle B-Tree database indices.	Kepler Exoplanet Light Curves.
STScI MAST	(1) CAOM data model with HTM based MS SServer Spatial Lib (2) HTM Tesselation Level 20, MS SServer (3) HEALPix (GALEX only), TOAST/HTM Tesselated Tiles (4) Declination Zone partitioning + HTM	(1) Mission Datasets: • HST • Hubble Legacy Archive • GALEX • FUSE • Kepler • SwiftUVOT • IUE • EUVE • HUT • WUPPE • TUES • BEFS (2) Catalogs: GSC-2 (3) AstroView Image Surveys • GALEX • SDSS • DSS (4) PanSTARRS Database

GSFC HEASARC	(1)Standard Postgres BTREE indexing of tables sorted on declination (2) Postgres PGSphere-based GIS index on position	(1)All HEASARC tables with positional information (~800 tables) (2) HEASARC master tables of position (i.e., a table that contains the table_name, ra, dec for every positional table in the HEASARC) Primary Archive Domain Missions: • Fermi • NuSTAR • Swift • Compton • Chandar • Suzaku • RXTE • INTEGRAL • WMAP • XMM-Newton • ROSAT • ASCA • Einstein
NED	B-Tree and B+Tree	(1) Many database tables that join data from >93,000 journal articles and catalogs
Chandra	(1)SybaseB-Tree (2)SQLite HEALPix	(1) Chandra Archive (2) Chandra Source Catalog

Table 2: Partner Astronomical Data Center Indexed Data Collections

Archive Data Center	Index Method	Example Data Sets
CDS	HEALPix Tessellation	Vizier Data Collection
LSST	HTM Tessellation	Raw data
CADC	PGSphere/PostgreSQL employing GiST	Multiple Telescope Data Products and Advanced Data Products
SDSS	HTM Tessellation	SDSS Sky Survey

External Community Indexings

In the GIS (Geographical Information Systems) community, custom databases used to be the norm, though this is giving way to a layered approach where the GIS analysis software (ArcGIS from ESRI, MapInfo, GRASS originally from the Army Corps of Engineers, and even AutoCAD) are used in conjunction with spatially-extended DBMSs like PostGIS and SpatiaLite. [ArcGIS]

Conclusions

The NASA archives are exploiting a variety of indexing schemes in their databases. They range from vendor or technology specific solutions such as PGSphere, to custom solutions such as IRSA's ChunkIndex, to freely available indexing solutions. Generally, these schemes are instantiations of a few common indexing schemes, primarily B-tree, R-tree and sky-tessalation (HTM, HEALPix) indexing schemes. Thus this inventory of indexing methods supports the view that proposed a performance comparison of MOC (HEALPix) and VOInventory (R-tree), is appropriate and necessary, but with one extension. Given the widespread use of HTM, the study should attempt include it, as far as resources and funding permit.

References

[BUD10] Budavàri T. et al., Searchable sky coverage of astronomical observations: footprints and exposures - Astronomical Society of the Pacific, 122, 1375 – 2010 2010PASP..122.1375B

[Gor84] Gorski K. et al ., HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere, The Astrophysical Journal, Volume 622, Issue 2, pp. 759-771. 2005, 2005ApJ...622..759G

[Gut84] Antonin Guttman. R-trees: a dynamic index structure for spatial searching. In ACM SIGMOD International Conference on Management of Data, pages 47-54, 1984.

[HNP95] J. M. Hellerstein, J. F. Naughton, and Avi Pfeffer.
Generalized search trees for database systems. In Proceedings of the 21st International Conference on Very Large Data Bases, Zurich, Switzerland, 1995.

[MOC14] HealPix Multi-Order Coverage Map –
http://www.ivoa.net/documents/MOC/20140602/index.html

[Sza05] A. Szalay, J. Gray, Gy. Fekete, P. Kunszt, P. Kukol, and A. Thakar (2005). Indexing the Sphere with the Hierarchical Triangular Mesh. Microsoft Research Technical Report, MSR-TR-2005-123.

[ArcGIS]
http://resources.arcgis.com/en/help/main/10.1/index.html#//003n0000001q000000