

An All-Sky 2MASS Mosaic Constructed on the TeraGrid

Anastasia Laity, G. B. Berriman, J. C. Good

Infrared Processing and Analysis Center, California Institute of Technology

D. S. Katz, J. C. Jacob

Jet Propulsion Laboratory, California Institute of Technology

L. Brieger, R. Moore

San Diego Supercomputer Center, University of California, San Diego

R. Williams

Center for Advanced Computing Research, California Institute of Technology

E. Deelman, G. Singh, M.-H. Su

Information Sciences Institute, University of Southern California

Abstract. The Montage mosaic engine supplies on-request image mosaic services for the NVO astronomical community. A companion paper describes scientific applications of Montage. This paper describes one application in detail: the generation at SDSC of a mosaic of the 2MASS All-sky Image Atlas on the NSF TeraGrid. The goals of the project are: to provide a value-added 2MASS product that combines overlapping images to improve sensitivity; to demonstrate applicability of computing at-scale to astronomical missions and surveys, especially projects such as LSST; and to demonstrate the utility of the NVO Hyperatlas format. The numerical processing of an 8-TB 32-bit survey to produce a 64-bit 20-TB output atlas presented multiple scalability and operational challenges. An MPI Python module, MYMPI, was used to manage the alternately sequential and parallel steps of the Montage process. This allowed us to parallelize all steps of the mosaic process: that of many, sequential steps executing simultaneously for independent mosaics and that of a single MPI parallel job executing on many CPUs for a single mosaic. The Storage Resource Broker (SRB) developed at SDSC has been used to archive the output results in the Hyperatlas. The 2MASS mosaics are now being assessed for scientific quality. The input images consist of 4,121,440 files, each 2MB in size. The input files that fall on mosaic boundaries are opened, read, and used multiple times in the processing of adjacent mosaics, so that a total of 14 TB in 6,275,494 files are actually opened and read in the creation of mosaics across the entire survey. Around 130,000 CPU-hours were used to complete the mosaics. The output consists of

1734 6-degree plates for each of 3 bands. Each of the 5202 mosaics is roughly 4 GB in size, and each has been tiled into a 12x12 array of 26-MB files for ease of handling. The total size is about 20 TB in 750,000 tiles.

1. Project Goals

The goals of the project are:

- To provide a "value-added" 2MASS¹ product that combines overlapping images to improve sensitivity;
- To demonstrate applicability of computing at-scale to astronomical missions and surveys, especially projects such as LSST; and
- To demonstrate the utility of the NVO Hyperatlas² format.

The scientific and technical applications of wide-area mosaics involve visualization of the content and observational coverage of large regions. Examples of such questions, and the role of the mosaic in addressing them, are:

- *Science question: I am planning an observing proposal of a cluster of galaxies. What observations have been made in this area?*
 - Superimpose footprints of observations on a mosaic of the cluster
- *Science Question: I am studying a large molecular cloud. Are there candidate sources not recorded in individual source catalogs?*
 - Superimpose locations of sources from existing catalogs on mosaics from surveys at different wavelengths
 - Examine multi-wavelengths mosaics for new candidate sources seen in at least two wavelengths
- *Science Question: Have I found a high proper motion object or an artifact in an image? Is it a T-dwarf candidate or is it in a region of higher than average extinction?*
 - Examine mosaics from 2MASS, SDSS and DPOSS, all at different epochs
 - Examine mosaics at different wavelengths to study distribution of color across images.

Another application of a mosaic is in quality assurance. Wide area coverage simplifies identification of artifacts present over spatial scales of several degrees, which may be caused by diffraction spikes, ghost images, and filter and dichroic glints. Figure 1 shows such artifacts image defects in a 3-color cutout, 5 deg on a side, from the all-sky mosaic centered on Wolf 359.

¹<http://www.ipac.caltech.edu/2mass>

²<http://www.us-vo.org/pubs/files/hyperatlas.pdf>




Figure 1. 3-color cutout from the all sky 2MASS mosaic, centered on Wolf 359 and 5 deg on a side

2. Calculation of Mosaic

The mosaic was calculated with the Montage image mosaic toolkit on the TeraGrid. Montage:

- Is highly parallelizable
- Processes FITS images with any WCS projection
- Reprojects images while preserving flux and astrometry
- Rectifies variable sky emission to a common level

The mosaic was computed from the 4,121,440 files (12.4 TB; 32-bit) of the 2MASS All Sky Survey Atlas Image Data at J, H and Ks; each 2 MBs in size (32-bit). The All-sky mosaic was performed on the Distributed Terascale Facility³ node at San Diego Supercomputer Center (SDSC)⁴, with parallelization through the Message Passing Interface (MPI). The images distributed to multiple nodes of a cluster for parallel steps. The input files that fell on mosaic boundaries were opened, read, and used multiple times in the processing of adjacent mosaics, so that a total of 14 TB in 6,275,494 files were actually opened and read. The processing time to complete the mosaics was more than 100,000 CPU-hours on a TeraGrid Linux cluster (1.5 Ghz Intel Itanium 2 processors); this included time for shakedown of the hardware. Further details of processing and data management are as follows:

- The entire 2MASS image archive was replicated onto TeraGrid disks

³<http://www.teragrid.org>

⁴<http://www.sdsc.edu>

- An archive of the images was created on the IBM HPSS High Performance Storage System after aggregating the images into 147,000 containers to minimize the impact on the HPSS name space.
- Checksums were used to validate the integrity of the images. Continued validation was essential to recognize and correct problems with data transport, RAID controllers, and media failure
- The computation was decomposed into 50 20-hour runs on 128 processors
- The mosaic process generated 64-bit image files, resulting in 20 TBs of output data
- The output images were truncated back to 32-bit image files, registered into the Storage Resource Broker, and replicated into the IBM HPSS archive.

3. Processing Summary

The total size of the output mosaic over all three bands is 10 TB (64-bit). It is organized in 1734 6-degree plates per band, 5205 plates total, 2 GB in size, each tiled into a 12x12 array of 13-MB files for ease of handling. There are 750,000 files total. The output is organized in Hyperatlas structure, implemented using the SDSC Storage Resource Broker⁵ data grid. The Hyperatlas organizes images in an Atlas consisting of multiple pages that are standard WCS sphere-to-plane projection mappings defined by a pointing center, projection type rotation, and scale. The All Sky Mosaic uses the TM-5-TAN-20 Atlas of 1734 pages covering the celestial sphere, with every point of the sphere at most 3.5 degrees ($= 5/\sqrt{2}$) from a pointing center, a scale of 1 arcsecond per pixel, and the WCS projection TAN.

4. Acknowledgements

This work made use of Montage, funded by the National Aeronautics and Space Administration's Earth Science Technology Office, Computational Technologies Project, under Cooperative Agreement Number NCC5-626 between NASA and the California Institute of Technology. The 2MASS project is a collaboration between The University of Massachusetts and the Infrared Processing and Analysis Center (JPL/ Caltech). Funding is provided primarily by NASA and the NSF. 2MASS image data products generated by Montage are not endorsed by the 2MASS project. Part of the work was done under funding support from the National Science Foundation National Virtual Observatory.

⁵<http://www.sdsc.edu/srb/>