

Montage: A Grid Enabled Image Mosaic Service for the National Virtual Observatory

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

Infrared Processing and Analysis Center, California Institute Of Technology, Pasadena, CA 91125

A. Bergou, J. Jacob, D. S. Katz

Jet Propulsion Laboratory, California Institute Of Technology, Pasadena, CA 91109

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

USC Information Sciences Institute, Marina del Rey, CA 90292

R. Williams

Center for Advanced Computing Research, California Institute Of Technology, Pasadena, CA 91125

Abstract.

The architecture of Montage, which delivers custom science-grade astronomical images, was presented at ADASS XII. That architecture has been tested on 2MASS images computed on single processor Linux machines that hold all image data in memory. This year, we describe the design of a grid-enabled version of Montage, suitable for large scale processing of the sky. It exploits to the maximum the parallelization inherent in the Montage architecture, whereby image reprojections are performed in parallel. All the re-projection jobs can be added to a pool of tasks and performed by as many processors as are available. We show how we can describe the Montage application in terms of an abstract workflow so that a planning tool such as Pegasus can derive an executable workflow that can be run in the Grid environment. The execution of the workflow is performed by the workflow manager DAGMan and the associated Condor-G. The grid processing will support tiling of images to a manageable size when the input images can no longer be held in memory. When fully tested, Montage will ultimately run operationally on the Teragrid. We will present processing metrics and describe how Montage is being used, including its application to science product generation by SIRTf Legacy Program teams and large-scale image processing projects such as Atlasmaker (this conference).

1. What Is Montage?

Montage has the broad goals of providing astronomers with software for the computation of custom, science grade image mosaics in FITS format. Custom refers to user specification of the parameters describing the mosaic: WCS projection, coordinate system, mosaic size, image rotation and spatial sampling. Science grade mosaics preserve the calibration and astrometric fidelity of the input images (Berriman et al. 2003).

Production of an image mosaic consists of four steps:

1. Re-projection of input images to a common spatial scale, coordinate system and World Coordinate System (WCS) projection
2. Modeling of background radiation in images to achieve common flux scales and background levels (by minimizing the inter-image differences)
3. Rectification of images to a common flux scale and background level and
4. Co-addition of re-projected, background-corrected images into a final mosaic.

Montage accomplishes these steps in independent modules, written in ANSI C for portability. This "toolkit" approach controls testing and maintenance costs and provides considerable flexibility to end users. They can, for example, use Montage simply to reproject sets of images and coregister them on the sky, or implement a custom background removal algorithm without impact on the other steps, or define a specific processing flow through custom scripts.

2. Distribution of the Montage Code

Version 1.7.1 of Montage is available for download at the project web site at <http://montage.ipac.caltech.edu>. The distribution consists of 20 modules containing 7560 lines of code, on which have been run 2595 test cases that exposed 119 defects. It includes a complete User Guide. This release emphasizes accuracy in photometry and astrometry over performance, and supports serial processing only. Montage has been built and used on many platforms (<http://montage.ipac.caltech.edu/docs/platforms.html>), but has been formally tested on Linux Red Hat 8.0 (Kernel release 2.4.18-14) running on 32-bit Intel processors. The tests involved building mosaics of 2MASS 2nd Incremental Release Full Resolution input images in 10 WCS projections; the mosaics were up to 2 x 2 degrees on a side, and output in Equatorial J2000, Galactic and Ecliptic coordinates.

3. Applications of Montage

Another paper in this volume (Williams et al. 2003) describes how Montage underpins Atlasmaker, which aims to deliver a multi-wavelength, science grade image atlas of the sky. Two of the SIRTf Legacy Program teams, GLIMPSE and SWIRE, are actively using Montage. Apart from its obvious role in generating science image products, Montage is also supporting data simulation and quality assurance. GLIMPSE will generate an infrared atlas of the galactic plane, and to support quality assurance of this atlas, GLIMPSE has used Montage as an engine to co-register 2MASS images in the J, H and K bands and MSX images at 8 μm .

SWIRE will perform a wide area imaging survey of high galactic latitude fields to study the evolution of stellar populations up to $z=3$, and has co-registered SIRTf Infrared Array Camera synthetic images and 2MASS images to support observation planning and pipeline testing. In support of their Cool Cosmos project (<http://coolcosmos.ipac.caltech.edu>), the IPAC Education and Public Outreach team is using Photoshop to combine single color mosaics made by Montage into spectacular multi-color images of regions of the sky covering spatial extents of several square degrees and more. A striking example is a three-color 2MASS mosaic of Rho Ophiuchus and the Galactic Center.

4. Montage: The Grid Years

Montage's capability of supporting any WCS projection and preserving the calibration and astrometric fidelity of the input images comes with a substantial computational burden. For example, reprojecting and resampling a single 2MASS image (1024 x 512 pixels; 2 MB in size) requires 100 seconds of processing time on a 1.4 GHz Linux box. This computational burden is a consequence of the accuracy and generality in the algorithm used to perform the reprojection. This algorithm uses classical spherical trigonometry to compute the overlap between the input and output (reprojected) images. The redistribution of flux from the input pixels to the output pixels consumes the bulk of the processing time; as a very good rule of thumb, the redistribution of flux in the reprojection of 2MASS images took up over 90% of the processing time in our test program.

We will investigate ways to speed-up the processing, but a simple way of providing speed-up is to run Montage on parallel processing environments, now relatively inexpensive. Many parts of the Montage processing can be run in parallel. Reprojection of input images can obviously be run on as many processors as are available, but many of the background removal can be run in parallel too. Fitting planes to the overlap between pairs of images can be performed in parallel as soon as the reprojection has been completed. Calculation of the parameters of the best-fit background model requires that all the fitting has been completed, but subsequent application of the model to the individual images can be done in parallel.

We have been developing a parallel processing architecture for Montage that takes advantage of the parallelization just described and uses the processing modules already tested and delivered. There are two broad aims to this work:

1. Take advantage of the modular design to develop grid-enabled software that will run on any parallel processing environment, such as an array of processors, or a grid of clusters; such software provides maximal flexibility to end users wishing to build large scale mosaics on local machines
2. Deliver a distributed image mosaic service that runs on the Teragrid (see www.teragrid.org); users order mosaics through a web form and receive e-mail notification when the mosaic is ready for pick-up. When complete, this service will take advantage of emerging protocols proposed by the NVO for discovering image data and making them accessible to Teragrid processors.

The first goal is met through a script that is built by Montage at run time. For a particular job, this script describes the flow of data and processing, specifies

which data are needed where, and which processes are to be run and when. This script, a directed acyclical graph (DAG), must be built at run time because the overlaps between the input images, requisite to the background removal, depend on their footprints on the sky, and therefore cannot be defined in advance of the processing. This DAG is what is submitted to standard tools for execution on available processors.

The second goal is met through a "DAG building script". It is part of a prototype architecture that has successfully run Montage on the 64-bit Linux Teragrid clusters at SDSC, NCSA, PSC, ANL, and CACR, and on the 32-bit processors of the USC-ISI/Wisconsin Grid.

A web service located at JPL creates an abstract workflow description of a mosaic request made to Montage by calling the DAG-building script. This workflow description, written in XML, is submitted to Pegasus (Planning for Execution in Grids), developed at ISI as part of the GriPhyN project. It interprets the "abstract workflow" to produce a "concrete workflow", which specifies the location of the data and the execution platforms. The concrete workflow is optimized by Pegasus to take advantage of previous processing runs, through the concept of Virtual Data. If Pegasus finds that data products described within the abstract workflow are already available (via queries to the Globus Replica Location Service), it reuses them and thus reduces the complexity of the concrete workflow. Pegasus also adds transfer nodes in the concrete DAG for staging the input image files and transferring out the generated mosaic. When the concrete workflow has been prepared, Pegasus submits it to Condor's DAGMan for execution.

Thus far, we have performed test runs to build 2MASS mosaics of the M42 region (Orion Nebula) that are 1.5 degrees on a side. One run on a single pool at SDSC consisted of 951 jobs, including 117 data transfer jobs, to process 113 input image files, and ran to completion in 94 minutes.

Acknowledgments. Montage is funded by NASA's Earth Science Technology Office, Computational Technologies Project, under Cooperative Agreement Number NCC5-626 between NASA and the California Institute of Technology.

References

- Berriman, G. B. et al., 2003, in APS Conf. Ser., Vol. 295, Astronomical Data Analysis Software and Systems XII, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook (San Francisco: ASP), [O6-3]
Williams, R. et al., 2004, this volume, [O4-3]