

Montage: An On-Demand Image Mosaic Service for the NVO

G. B. Berriman, J. C. Good

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

D. Curkendall, J. Jacob, D. S. Katz

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

T. A. Prince

Division of Physics, Mathematics and Astronomy, California Institute of Technology, Pasadena, CA 91125

R. Williams

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

Abstract. Montage will deliver a generalized toolkit for generating on-demand, science-grade custom astronomical image mosaics. “Science-grade” in this context requires that terrestrial and instrumental features are removed from images in a way that can be described quantitatively. “Custom” refers to user-specified parameters of projection, coordinates, size, rotation and spatial sampling, and whether the drizzle algorithm should be invoked. The greatest value of Montage will be its ability to analyze images at multiple wavelengths, by delivering them on a common projection, coordinate system and spatial sampling and thereby allowing analysis as if they were part of the same multi-wavelength image. Montage will be deployed as a compute-intensive service through existing portals. It will be integrated into the emerging NVO architecture, and run operationally on the Teragrid, where it will process the 2MASS, DPOSS and SDSS image collections. The software will also be portable and publicly available.

1. Introduction: What Is Montage?

Montage is an astronomical image mosaic service. It will deliver science-grade custom image mosaics according to input specifications of size, WCS projection, coordinate system, image rotation, spatial sampling, and background removal. Science grade in this context has two definitions: Montage will preserve the astrometric and photometric accuracy of the input images, and it will perform

background rectification on the input images in such a fashion that its impact on the photometric quality of the data can be described quantitatively.

The principal science goal of Montage is to extend the mosaicking capabilities available to astronomers by supporting the production of multi-wavelength images. That is, images from diverse sources over many wavelengths can be combined into mosaics having common coordinates, WCS projection, spatial sampling, size and image rotation. Such mosaics are requisite for analysis tasks such as multi-wavelength source extraction and band-merging. The SIRTf Wide Area Infrared Experiment (SWIRE) has adopted Montage to support analysis of multi-wavelength images of deep SIRTf (infrared) images, and ground based optical and radio images.

2. Schedules and Platforms

The Montage software will be publicly available through the project website at <http://montage.ipac.caltech.edu>. Montage will be designed and developed as a portable software package written in ANSI C, and supported under current plans on Solaris 2.7 and 2.8, Linux 6.x and 7.x and AIX. It will be delivered incrementally at roughly six monthly intervals between February 2003 and January 2005. Early deliveries will emphasize delivery of astrometrically and photometrically accurate, science grade mosaics, while later ones will emphasize speed and throughput. On final delivery, Montage will achieve a throughput of 30 square degrees (e.g., thirty 1 degree \times 1 degree mosaics, or one 5.4 degrees \times 5.4 degrees mosaic, etc.) per minute on a 1024 \times 400MHz R12K Processor Origin 3000 or machine equivalent with sustained bandwidth to disk of 160 MB/sec.

In August 2003, Montage will be deployed as an operational service through the National Virtual Observatory (NVO; <http://www.us-vo.org>). Users will then order mosaics through existing astronomy portals, with the calculations performed on the Distributed Terascale Facility (<http://www.teragrid.org>), a high performance computational grid provided by the NSF Partnership for Advanced Computational Infrastructure. When fully deployed, the Teragrid will provide aggregate computational power on the order of 10 teraflops, aggregate disk cache on the order of 800 TB and archival storage capacity of 6 petabytes.

3. Design of Montage

Montage is designed as a portable toolkit that performs the tasks needed to compute image mosaics in stand-alone modules, controlled by simple executives or scripts. Montage performs the following steps in computing an image mosaic: re-projection of input images to a common spatial scale, coordinate system and World Coordinate System (WCS) projection; modeling of background radiation in images to achieve common flux scales and background levels; rectification of images to a common flux scale and background level (if requested); and co-addition of re-projected, background-corrected images into a final mosaic.

The most complex parts of Montage are the image reprojection and the background modeling and rectification. These are described separately below.

3.1. Image Reprojection

Image reprojection involves the redistribution of information from a set of input pixels to a set of output pixels. For astronomical data, the input pixels represent the total energy received from an area on the sky. It is critical to preserve this information when redistributed into output pixels. And, it is important to preserve the positional (astrometric) accuracy of the energy distribution, so common techniques such as adding all the energy from an input pixel to the “nearest” output pixel are inadequate.

Instead, we must redistribute input pixel energy to the output based on the exact overlap of these pixels, possibly even with a weighting function across the pixels based on the point spread function for the original instrument. The goal is to create an output image which is as close as possible to that which would have been created if the sky had been observed using an instrument with the output image’s pixel pattern. We are also committed to building a system which handles all astronomical projections and coordinate systems equally well.

The most common approach to determining pixel overlap is to project the input pixel into the output pixel Cartesian space, but Montage will instead project both input and output pixels onto the celestial sphere. Since all such “forward” projections are well defined, the rest of the problem reduces to calculating the area of overlap of two convex polygons on a sphere (with no further consideration of the projections involved). The issue of handling reprojections therefore becomes a problem of classical spherical trigonometry. General algorithms exist for determining the overlap of polygons in Cartesian space (O’Rourke 1998). We have modified this approach for use in spherical coordinates to determine the intersection polygon on the sphere (a convex hull) and applied Girard’s Theorem to calculate the polygon’s area.

The result is that for any two overlapping pixels, we can determine the area of the sky from the input pixel that contributes energy to the output pixel. This provides not only a mechanism for accurately distributing input energy to output pixels but, as we shall see, a natural weighting mechanism when combining overlapping images. Our approach implicitly assumes that the polygon defining a single pixel can be approximated by the set of great circle segments connecting the pixel’s corners. Since even the largest pixels in any realistic image are on the order of a degree across, the non-linearities along a pixel edge are insignificant. Furthermore, the only effect of this assumption would be on the astrometric accuracy of the energy location information and would amount to a very small fraction (typically less than 0.01) of the size of a pixel. Total energy is still conserved. Finally, this design easily supports parallel processing because each image can be processed independently of the others.

3.2. Background Modeling and Rectification

If several images are to be combined into a mosaic, they must all be projected onto a common coordinate system (see above) and then any discrepancies in brightness or background must be removed. Our assumption is that the input images are all calibrated to an absolute energy scale (i.e., brightnesses are absolute and should not be modified) and that any discrepancies between the images are due to variations in their background levels that are terrestrial or instrumental in origin.

The Montage background matching algorithm is based on the assumption that terrestrial and instrumental backgrounds can be described by simple functions or surfaces (e.g., slopes and offsets). Stated more generally, we assume that the “non-sky” background has very little energy in any but the lowest spatial frequencies. If this not the case, it is unlikely that any generalized background matching algorithm will be able distinguish between “sky” and rapidly varying “background”; background removal will then require an approach that depends on detailed knowledge of an individual data set.

Given a set of overlapping images, characterization of the overlap differences is key to determining how each image should be adjusted before combining them. We take the approach of considering each image individually with respect to its neighbors. Specifically, we determine the areas of overlap between each image and its neighbors and use the complete set of overlap pixels in a least-squares fit to determine how each image should be adjusted (e.g., what gradient and offset should be added) to bring it “best” in line with its neighbors. In practice, we only adjust the image by half this amount, since all the neighbors are also being analyzed and adjusted and we want to avoid ringing in the algorithm. After doing this for all the images, we iterate (currently for a fixed number of times though we may later introduce convergence criteria). The final effect is to have subtracted a low-frequency (currently a gradient/offset) background from each image in such a way that the cumulative image-to-image differences are minimized. To speed the computation (and minimize memory usage), we approximate the gradient and offset values by a planar surface fit to the overlap area difference images rather than perform a least squares fit.

4. Conclusions

The Montage image mosaic service will deliver science grade astronomical image mosaics. These science grade images will preserve the photometric and astrometric accuracy of the input images, and rectify the images for terrestrial background emission. Montage will be deployed on the Distributed Terascale Facility, but the code will be portable and available for download. We invite comments and participation from the community, especially for testing and validation. Interested parties should contact the project at montage@ipac.caltech.edu.

Acknowledgments. We wish to thank Dr. Ewa Deelman, Dr. Reagan Moore and Ms. Leesa Brieger for technical guidance and assistance in our efforts to run Montage on parallel processors. Montage is supported by the NASA Earth Sciences Technology Office Computing Technologies program, under Cooperative Agreement Notice NCC 5-6261.

References

O’ Rourke, J. 1998, in *Computational Geometry*, (Cambridge University Press), 220.