Montage: An Astronomical Image Mosaic Service for the National Virtual Observatory

http://montage.ipac.caltech.edu

Cooperative Agreement Number NCC 5-626

First Annual Report
(Period April 1 2002 – August 31 2002)

Figure 1: A 2MASS Atlas Image mosaic, covering 40′ × 40′ on the sky and binned to 2′′ × 2′′ pixels, of the ρ Ophiuchi Dark Cloud, derived with Montage and featured as the 2MASS Picture of the Week on September 23 (add in URL for 2mass POTW). (The image has been binned to.) Only about 150 pc distant, this cloud is one of the closest regions of low-mass stars in formation. As such, its stellar, dusty, and gaseous contents have been intensively studied at a number of wavelengths, from radio to X-rays. Shown in this 2MASS image are only a few components of the larger Cloud, primarily the dark nebula LDN 1688, toward the center of the image. A number of embedded protostellar and young stellar objects, as well as patchy, filamentary, and streamer-like dust clouds, are also seen. The 2MASS color-color diagram shows that the extinction within the cloud exceeds 30 visual magnitudes. *Image mosaic courtesy of R. Hurt (IPAC).*

Objective:

The Montage project will deploy a portable, compute-intensive service that will deliver science-grade custom mosaics on demand. 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. Montage leverages the image mosaic algorithms already deployed in the *yourSky* image mosaic server.

Montage will generate mosaics from input files that comply with the Flexible Image Transport System (FITS) standard and contain images whose projections comply with the World Coordinate System (WCS) standards. In operations, Montage will be deployed on the emerging Distributed Terascale Facility(hereafter, TeraGrid). where it will process requests for 2Micron All Sky Sky Survey (2MASS), Sloan Digital Sky Survey (SDSS) and Digital Palomar Observatory Sky Survey (DPOSS) image mosaics; the requests will be made through existing astronomy World Wide Web portals.

Montage's performance goal is to sustain throughput of 30 square degrees (e.g. thirty 1 degree x 1 degree mosaics, or one 5.4 degrees x 5.4 degrees mosaic, etc.) per minute on a 1024x400MHz R12K Processor Origin 3000 or machine equivalent with sustained bandwidth to disk of 160 MB/sec.

Approach:

Deep, wide area imaging surveys have assumed fundamental importance in astronomy. They are being used to address such fundamental questions as the structure and organization of galaxies in space and the dynamical history of our galaxy. One of the most powerful probes of the structure and evolution of astrophysical sources is their behavior with wavelength, but this power has yet to be fully realized in the analysis of astrophysical images because survey results are published in widely varying coordinates, map projections, sizes and spatial resolutions. Moreover, the spatial extent of many astrophysical sources is much greater than that of individual images. Astronomy therefore has need for a general image mosaic engine that will deliver image mosaics of arbitrary size in any common coordinate system, in any map projection and at any spatial sampling. Montage aims to provide this service.

The key to our technical approach is to develop a flexible framework that will support many custom user cases and processing needs. Such cases range from compute and time intensive mosaics covering wide areas of the sky, small mosaics generated on a desktop as part of a small research project or observation planning program, to mosaics generated as a standard science product as part of processing pipeline.

Scientific Accomplishments

Three Science Legacy teams from the Space Infrared Telescope Facility (SIRTF) are validating a mature prototype of Montage (see Technology Accomplishments). The SIRTF Wire Area Infrared Experiment (SWIRE) has already adopted Montage to generate its wide-area image mosaic science products, and to generate mosaics from ancillary radio and optical data that is central to source identification and interpretation. Figure 2 shows a simulation of a deep image of new extragalactic sources such as Montage will expect to produce science products.

The Galactic Legacy Infrared Mid-Plane Survey Extraordinaire (GLIMPSE) team is investigating the use of Montage as part of its pipeline, where it would generate deep, infrared mosaics of the galactic plane. Finally, “From Molecular Cores to Planet-forming Disks” will generate maps of molecular clouds and they are also considering Montage as part of their pipeline.

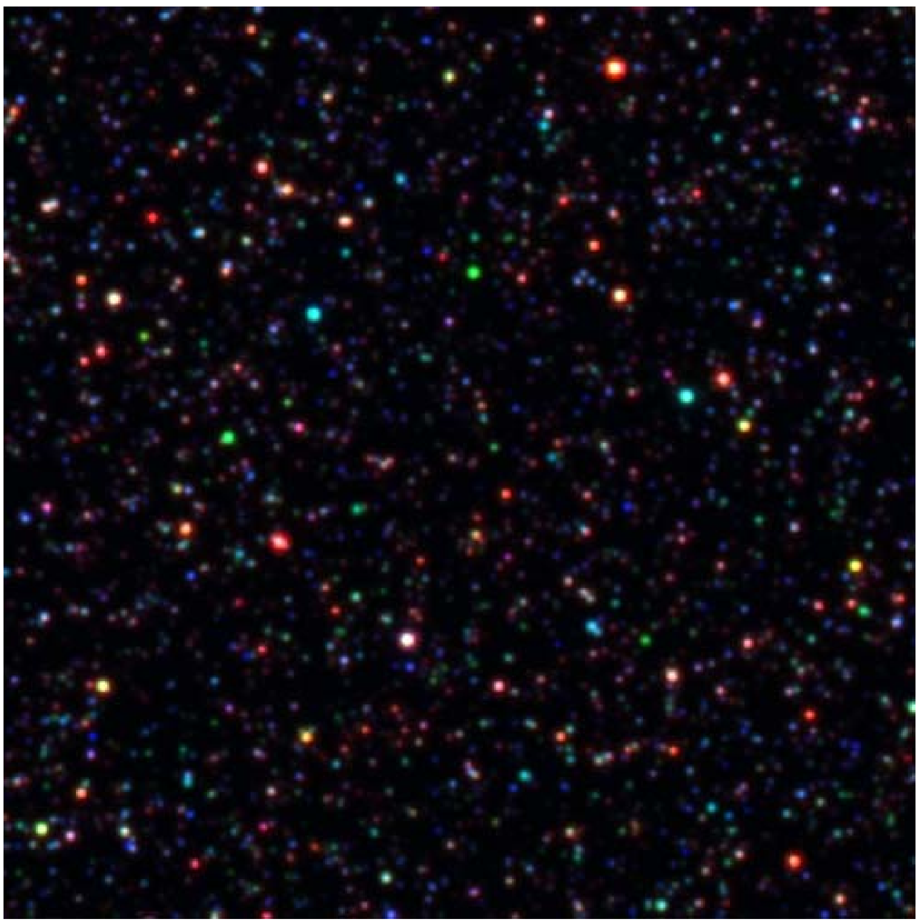

Figure 2: A Montage mosaic generated by the SWIRE team. This is a 3-color simulation of the infrared extragalactic Universe such as SIRTF is expected is reveal. The image shows a simulation of the distribution of galaxies and quasars that range in distance from a few million to over 10 billion light years away. This image is just 0.2 square degrees in size, and hints at the full dimension of science content of SIRTF data. The final delivered SWIRE mosaics will be up to 75 times larger than the sample images

displayed here. *Image and Caption courtesy of Carol Lonsdale (SWIRE Principal Investigator)*

Technology Accomplishments (including progress towards milestones)

In this, the first year of the project, we have emphasized project requirements analysis, project planning and sign and prototyping.

Progress Towards Milestones

Table 1 shows the Milestones met this year, and the products delivered. All documents are available from the Project website at http://montage.ipac.caltech.edu/documents.html

Table 1: Montage Milestones and Deliverables

Milestone	Deliverables
A) Software engineering plan completed	Software Engineering Plan Version 1.0 Software Requirements Specification 1.0 Project web page.
E) Code baseline delivered	Baseline performance for *yourSky* mosaic engine and background removal published on web page. High level design specification v1.0
H) Design Policy for Interoperability and Community Delivery Established	Design policy published on project web page Test Plan specification v1.0
B) First Annual Report delivered to project web site	This document

Technical Accomplishments

Baseline Performance

We have published on the project web page (http://montage.ipac.caltech.edu/documents.html) the baseline performance of the *yourSky* mosaic engine. Figure 3 summarizes this performance.

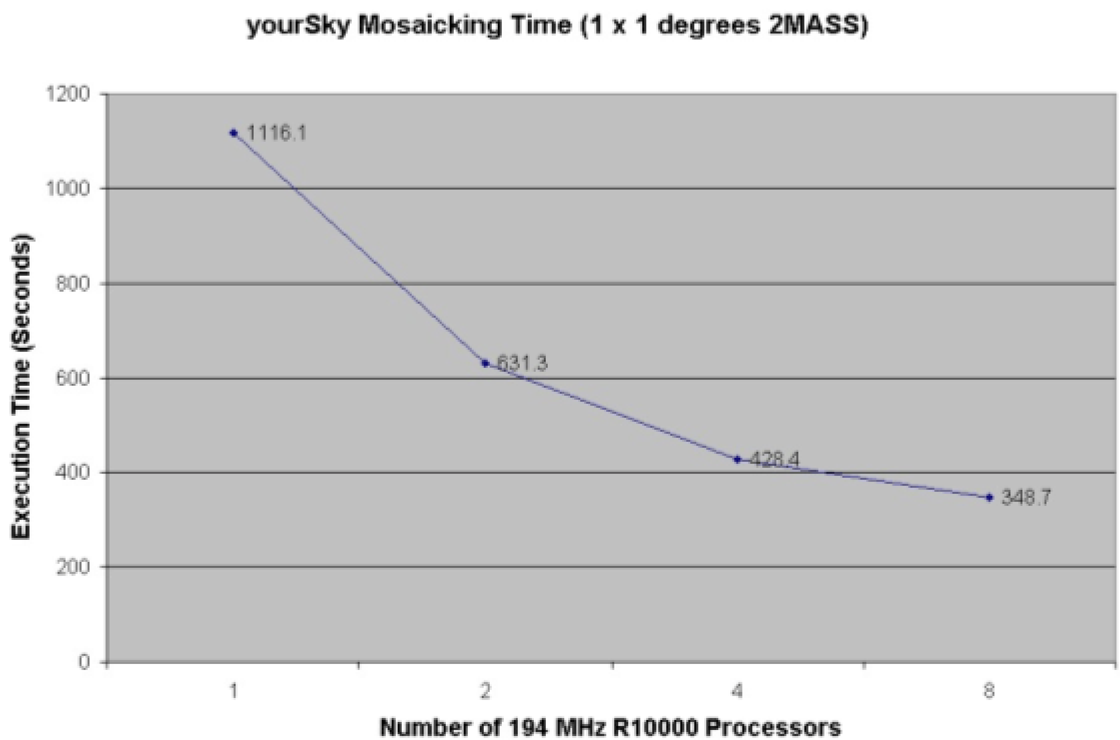

Figure 3: Baseline performance of the *yourSky* mosaic engine for processing a 1 square degree mosaic of 2MASS images.

We have also baselined the background correction calculations for 2MASS and DPOSS image collections. The 2MASS correction involved iterative fitting of surfaces to the 2MASS images. This execution took on average 13 seconds for each 2MASS image, on a Sun Ultra 5 workstation under Solaris 2.8. The full collection of 5834400 images would require 2.4 CPU years on this platform.

The DPOSS background rectification assumed a quadratic fit to the background.
The rectification was performed on a single processor of an HP Superdome (PA8600 processor at 550 MHz), and took 660 seconds for each plate; this corresponds to about 20 days for the whole 3 Tbyte DPOSS image collection.

Design Specification

The Science requirements developed in cooperation with the Customer Review Board mandated a highly modular design that would support a wide variety of use cases and custom processing needs. With our experience of yourSky as a guide, we designed an architecture that performed the tasks needed to generate image mosaics as *standalone* modules controlled by executives or scripts that are easily adapted to particular use cases. The tasks are:

- transforming the coordinates and projections of the image mosaics according to the user's specifications,

- rectifying the images for terrestrial background emission, modeled as a smoothly varying surface, and
- co-adding the images to generate the final mosaics.

Processing Algorithms

The first Code Improvement Milestone F emphasizes accuracy in the mosaics delivered by Montage. We have developed a general algorithm that preserves the photometric integrity of the input data and delivers high accuracy astrometry in the reprojected images. This algorithm is the computational heart of Montage, and is shown in Figure 4. An input pixel generally overlaps several pixels in the reprojected image. Our algorithm is based on spherical trigonometry, and weights fluxes in the reprojected pixels according to the area of overlap. The algorithm is quite general but highly parallelizable.

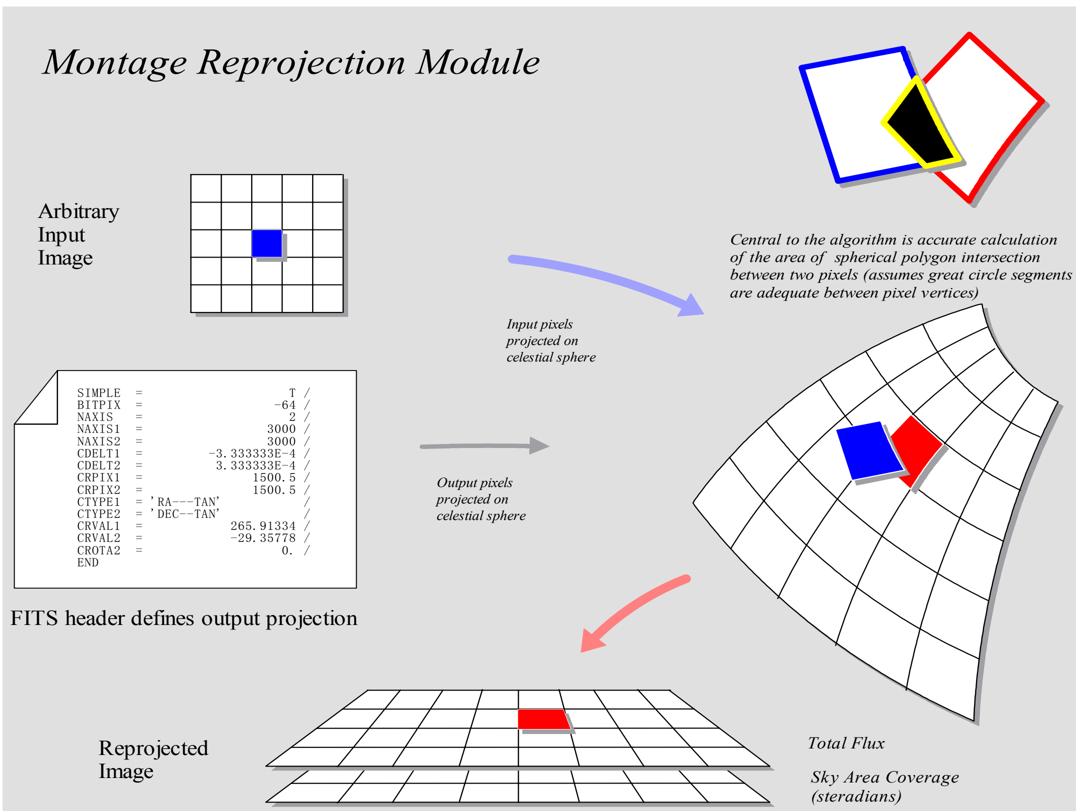

Figure 4: The Montage reprojection algorithm. This algorithm preserves the scientific fidelity and astrometric accuracy of the input images by employing spherical trigonometry to calculate the area of overlap between input and output pixels, and weighting the distribution of flux between input and output image by the areas of overlap.

Code Development

We have developed a mature end-to-end prototype of Montage that implements the design and algorithm described above. This prototype runs on Solaris and Linux 6.3, and is the version currently under evaluation by prospective users. Figures 5(a) and 5(b) show the improvement of the reprojection algorithm over the most common method used to date in astronomy, in which flux from an input pixel is assigned to the closest neighbor pixel in the reprojected image.

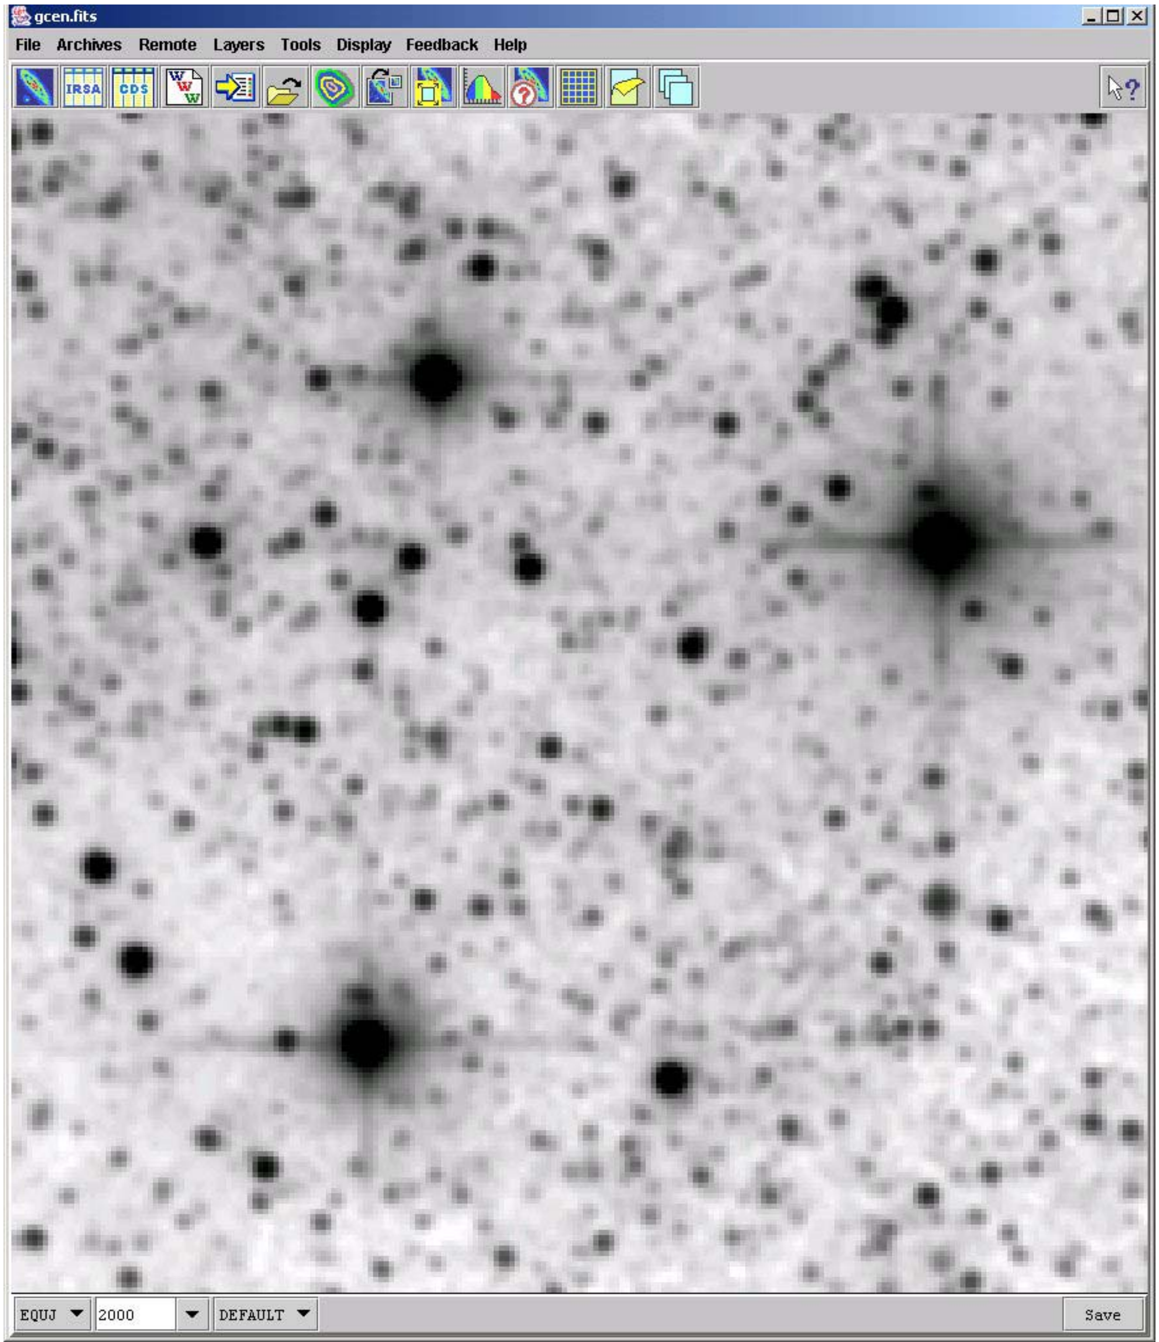

Figure 5(a) : Mosaic of the Galactic Center at 2.2 μm generated with the reprojection algorithm used in Montage. Compare with Figure 5(b).

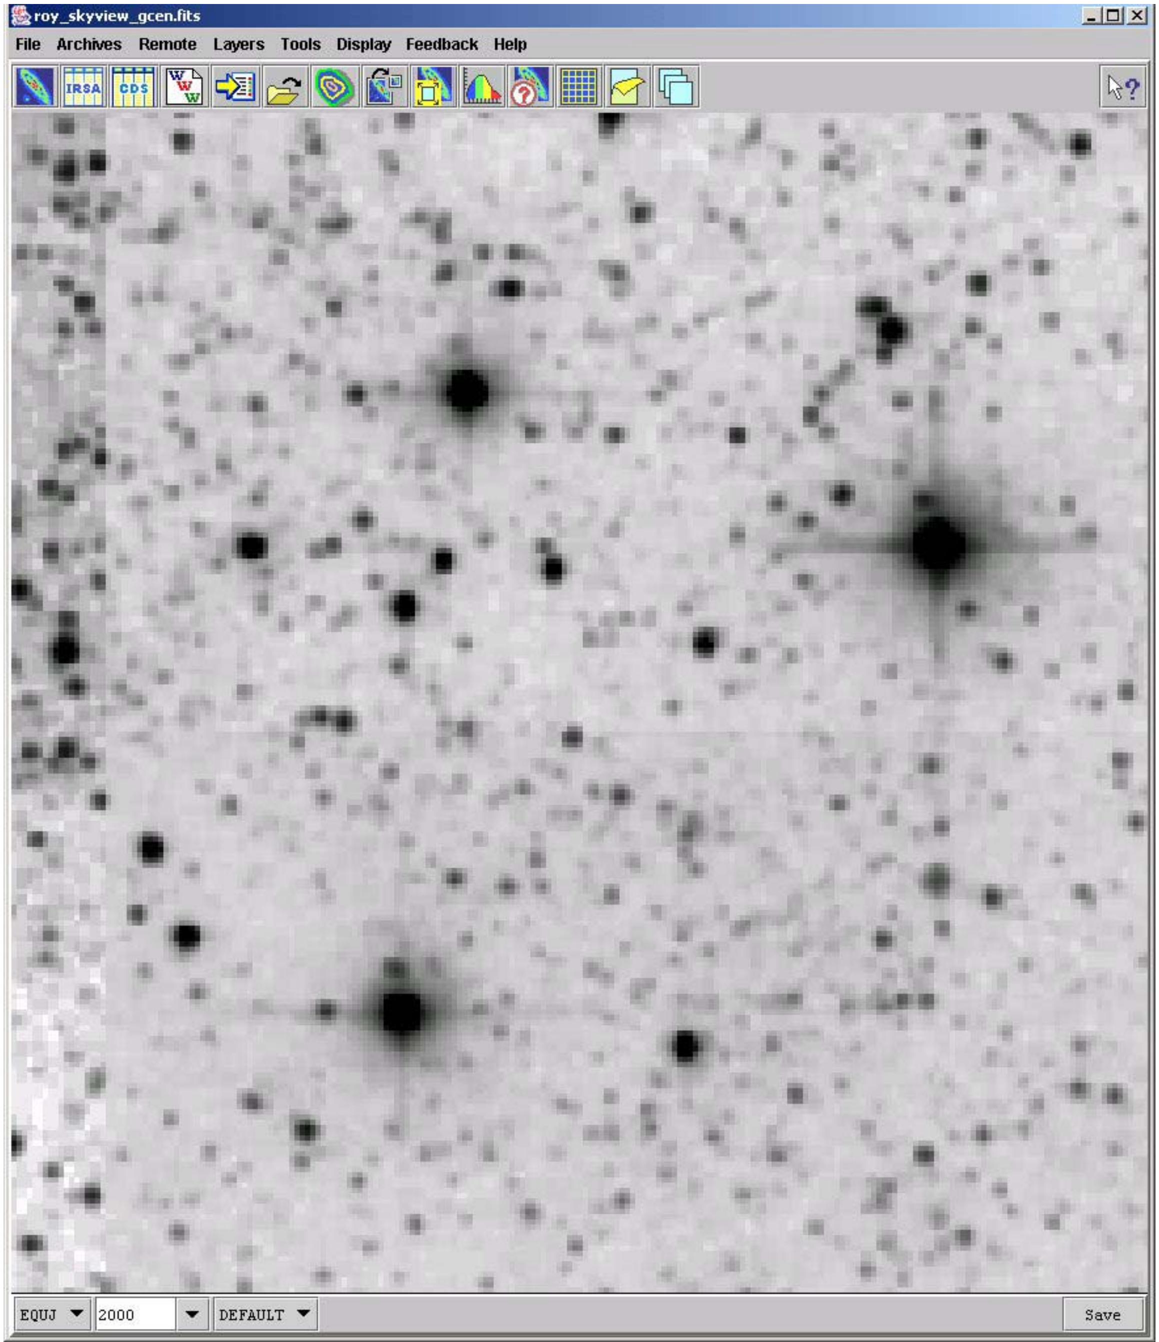

Figure 5(b): As Fig 5(a) except that it uses the nearest neighbor algorithm in reprojection. Note the squared-off source profiles and the choppy appearance of the image. The algorithm conserves neither flux nor source astrometry at the fractional pixel level.

Status/Plans

Milestones/Deliveries

The thrust of our effort in the next year is to meet the next two milestones:

F) First Code Improvement (Due Feb 28 2003)

We are currently completing the documentation of the detailed design of version 1.0 of Montage, which will be followed by delivery of the code to the test bed in late October. We are on schedule for delivery in February 2003. This code will run on a single processor.

I)Interoperability Prototype (Due July 30 2003)

Users will place an order for an image mosaic through existing astronomy portals, and the request will be processed by Montage will run on the TeraGrid. For this milestone, we will optimize the code delivered in Milestone F) to run on multiple processors on the TeraGrid. Members of the NVO consortium have agreed to take part in efforts to prototype a version of Montage that will run on the TeraGrid, and to take part in the formal testing of the delivered code. The 2MASS and DPOSS image collections will be available through the HPSS; a substantial fraction of these image collections will be copied to spinning disk on the TeraGrid, and made available to the processors running Montage through high-speed fiber channel connections.

Meetings

We will present an oral paper on Montage at ADASS 2002, Baltimore (October 12 through 16): "Montage: An On-Demand Image Mosaic Service for the NVO," G. B. Berriman , D. Curkendall , J. Good, J. Jacob , D. S. Katz, T. Prince, R. Williams.

We also anticipate giving demonstrations at AAS meetings in January and June 2003; the demos will most likely be accompanied by a brochure.

Outreach

The E/PO group at IPAC is evaluating Montage as a tool to generate large scale mosaics of the sky at multiple wavelengths, with the intent of displaying these images in schools and planetaria. Figure 1 was generated as part of this evaluation effort.

Demonstration at SC 2002

In cooperation with the NVO project and SDSC, we are planning to have a live demo of Montage at SC 2002. The mosaics will be generated on the 8-processor IBM Blue Horizon at SDSC. We are collaborating with the SDSC staff to develop a parallelized prototype of Montage to support this demo.

Publications:

"An Architecture for Access to A Compute Intensive Image Mosaic Service in the NVO"
G. Bruce Berriman, David Curkendall, John Good , Joseph Jacob, Daniel S. Katz, Mihseh Kong, Serge Monkewitz , Reagan Moore, Thomas Prince, Roy Williams

Paper presented at SPIE Conference 4686 "Virtual Observatories", Hawaii, August 2003. Available at http://montage.ipac.caltech.edu/documents.html

Presentations:

"An Architecture for Access to A Compute Intensive Image Mosaic Service in the NVO"
G. Bruce Berriman, David Curkendall, John Good , Joseph Jacob, Daniel S. Katz, Mihseh Kong, Serge Monkewitz, Reagan Moore, Thomas Prince, Roy Williams

Presentation at SPIE Conference 4686 "Virtual Observatories", Hawaii, August 2003.
Available at http://montage.ipac.caltech.edu/documents.html

Point of Contact

Dr. G. Bruce Berriman, 100-22 Caltech, Pasadena, CA 91125
Phone (626) 395-1817; FAX (626) 397-7354
Email: gbb@ipac.caltech.edu

Media References

Not applicable.

Patents

Not applicable

Graduate Students/Post-docs

Not applicable