

PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://spiedigitallibrary.org/conference-proceedings-of-spie)

yourSky: rapid desktop access to custom astronomical image mosaics

Joseph C. Jacob, Robert J. Brunner, David W. Curkendall,
Stan George Djorgovski, John C. Good, et al.

Joseph C. Jacob, Robert J. Brunner, David W. Curkendall, Stan George Djorgovski, John C. Good, Laura Husman, George Kremenek, Ashish Mahabal, "yourSky: rapid desktop access to custom astronomical image mosaics," Proc. SPIE 4846, Virtual Observatories, (16 December 2002); doi: 10.1117/12.461514

SPIE.

Event: Astronomical Telescopes and Instrumentation, 2002, Waikoloa, Hawai'i, United States

yourSky^{*} : Rapid Desktop Access to Custom Astronomical Image Mosaics

Joseph C. Jacob^{†a}, Robert Brunner^b, David W. Curkendall^a, S. George Djorgovski^b, John C. Good^c,
Laura Husman^a, George Kremenek^d, and Ashish Mahabal^b

^aJet Propulsion Laboratory, California Institute of Technology

^bAstronomy Department, California Institute of Technology

^cInfrared Processing and Analysis Center, California Institute of Technology

^dSan Diego Supercomputer Center

ABSTRACT

The yourSky custom astronomical image mosaicking software has a Web portal architecture that allows access via ordinary desktop computers with low bandwidth network connections to high performance and highly customizable mosaicking software deployed in a high performance computing and communications environment. The emphasis is on custom access to image mosaics constructed from terabytes of raw image data stored in remote archives. In this context, custom access refers to new technology that enables on the fly mosaicking to meet user-specified criteria for region of the sky to be mosaicked, datasets to be used, resolution, coordinate system, projection, data type and image format. The yourSky server is a fully automated end-to-end system that handles all aspects of the mosaic construction. This includes management of mosaic requests, determining which input images are required to fulfill each request, management of a data cache for both input image plates and output mosaics, retrieval of input image plates from massive remote archives, image mosaic construction on a multiprocessor system, and making the result accessible to the user on the desktop. The URL for yourSky is <http://yourSky.jpl.nasa.gov>.

Keywords: Astronomical image mosaics, Web portal, Image reprojection, Data access.

1 INTRODUCTION

In recent years the Astronomy community has witnessed an explosion in the size and complexity of datasets that are available for astronomical research due to rapid advances in remote sensing technology. The massive datasets that now exist collectively contain tens of terabytes of imagery and catalog data in wavelengths spanning the entire electromagnetic spectrum. Although this rich data store represents a significant opportunity for new scientific discoveries, it also represents a serious challenge to the community: *How does one effectively and efficiently extract information from such a large and complex collection of data?* To address this question, the Astronomy and Information Technology communities have joined forces to launch the National Virtual Observatory (NVO) [1,2,3] in the United States and related efforts elsewhere in the world [4,5,6,7].

As a community effort, the NVO will necessarily exhibit a loosely coupled, distributed architecture, with an emphasis on interoperability between components developed and deployed by domain experts in various areas. Since many of these components will require an enormous amount of computation and/or data movement, the NVO will need to be deployed in a distributed, high performance, scalable computing environment. However, a significant fraction of astronomical research is conducted by scientists and students with limited resources, ordinary desktop computers and low bandwidth network connections. Therefore, to be effective the NVO will also need to provide portals to its high performance infrastructure that will make it usable by researchers anywhere. This paper summarizes some of our early efforts to design, develop and deploy a useful service called yourSky, for custom astronomical image mosaicking, with an

* <http://yourSky.jpl.nasa.gov/>

† Joseph.Jacob@jpl.nasa.gov; phone 1 818 354-0673; fax 1 818 393-6141; <http://pat.jpl.nasa.gov/public/jjacob>; Jet Propulsion Laboratory, 4800 Oak Grove Drive, Mail Stop 126-104, Pasadena, CA 91109-8099

architecture that supports both high performance computations and lightweight portals, and matches the loosely coupled, distributed nature of the NVO infrastructure.

In this paper, mosaicking refers to reprojecting input image plates to a common coordinate system, projection, equinox and epoch, and combining the resulting plates to produce a single output image. There are strong science drivers for mosaicking. The most obvious is that large image mosaics enable analysis of celestial objects that either do not fit on a single image plate in the native image partitioning scheme used by a survey, or fall at the boundary between two or more neighboring plates. Also, mosaics enable analysis of the large-scale structure of the universe. In addition, mosaicking datasets in different wavelengths or from different surveys to the same grid enables multi-spectral analysis, which could be essential for identifying new, previously unknown, types of objects, or for identifying new objects that are so faint in a single wavelength that they are overlooked until combined with the signals from other wavelengths.

The only client software required to use the yourSky custom astronomical image mosaic server is the ubiquitous web browser. By filling out and submitting the form at <http://yourSky.jpl.nasa.gov>, users have custom access on their desktop to all of the publicly released data from the member surveys. In this context, “custom access” refers to new technology that enables on-the-fly astronomical image mosaicking to meet user-specified criteria for region of the sky to be mosaicked, dataset to be used, resolution, coordinate system, projection, data type, and image format. All mosaic requests are fulfilled from the original archive data so that the domain experts maintain control and responsibility for their data and data corruption due to resampling is minimized because only one reprojection is done from the raw input data. Currently the data archives that are accessible with yourSky are the Digitized Palomar Observatory Sky Survey (DPOSS) [8,9] and the Two Micron All Sky Survey (2MASS) [10]. DPOSS has captured the entire northern Sky at 1 arc second resolution in three visible wavelengths. 2MASS has captured the entire sky at 1 arc second resolution in three infrared wavelengths. The yourSky architecture supports expansion to include other surveys, without regard to the native image partitioning scheme used by a particular survey.

The overall architecture for yourSky is described in Section 2, followed by more detail about the components of yourSky in later sections. These components include the custom astronomical image mosaicking software described in Section 3, the request management system described in Section 4, the plate coverage database described in Section 5, and the management of local data caches described in Section 6. The member archives that are accessible with yourSky are described in Section 7. Finally, a summary and description of our future plans for yourSky are provided in Section 8.

2 ARCHITECTURE

The overall architecture for yourSky is illustrated in Figure 1. In the figure, the numbered descriptions on some of the arrows give the steps taken to fulfill a typical mosaic request. The procedure is as follows. The clients at the top left of the illustration are the Web browsers that may be used to submit requests to yourSky. A simple HTML form interface, shown in Figure 2, is used to specify the parameters that are to be passed to the custom astronomical image mosaicking software. The mosaicking software and the mosaic parameters are described in detail in Section 3. The yourSky Mosaic Request Manager running on the yourSky server checks for mosaic requests and hands them off to the Mosaic Request Handler, using the user priority scheme described in Section 4. The Mosaic Request Handler queries the Plate Coverage Database, described in Section 5, to determine which input image plates from DPOSS or 2MASS are required to fulfill the mosaic request. A fixed size data cache is maintained on the yourSky server to store the input image plates required to build recent mosaic requests. If all of the required input image plates are already present locally in the data cache, the mosaic is constructed immediately using the custom astronomical image mosaicking software, described in Section 3. Currently the mosaic processing is performed on an SGI Poweronyx with eight R10000 processors running at 194 MHz, but this machine is quite outdated and will need to be updated in the near future. If some of the required input image plates are not already cached locally, they need to be retrieved from their respective archives. Therefore a “tape request” is issued. The yourSky Tape Request Manager checks for tape requests and hands them off to the Tape Request Handler, which retrieves the required input image plates from the appropriate remote archive. Once all of the input image plates for a request have been cached on a local disk, the custom astronomical image mosaicking software is launched to construct the mosaic. When the mosaic, built precisely to match the user’s request parameters, is ready an email is sent back to the user with the URL where the image mosaic can be downloaded.




Figure 1. The architecture of yourSky supports fully automated mosaicking, including retrieval of the input image plates from the original survey archives.

3 CUSTOM ASTRONOMICAL IMAGE MOSAICKING SOFTWARE

The heart of the yourSky server is the custom astronomical image mosaicking software that is used to construct an image mosaic precisely matching user-specified parameters. The inputs to the mosaicking software are a list of input images to be mosaicked and the custom parameters that determine the properties of the mosaic to be constructed. The only requirements on the input images are the following. First, they must comply with the standard dictated by the Flexible Image Transport System (FITS), a data format that is well understood by the astronomy community and has long been used as the de facto method for sharing data within the community [11]. FITS format images encapsulate the image data with keyword-value pairs that give additional information about how the data values in the image map to locations on the sky. The second requirement for input images to the mosaicking software is that the FITS header must contain valid World Coordinate System (WCS) information. The WCS defines pixel-to-sky and sky-to-pixel coordinate transformations for a variety of coordinate systems and projections commonly used by the astronomy community [12].

3.1 Custom Access

With yourSky, the emphasis is on custom access to astronomical image mosaics. The following parameters may be used to specify the mosaic to be constructed:

1. **Center right ascension and declination:** Required parameters, analogous to the CRVAL1 and CRVAL2 FITS keywords, which specify the location on the celestial sphere of the tangent point for the image projection plane. By default, this center of projection is placed at the center pixel in the mosaic, analogous to the CRPIX1 and CRPIX2 FITS keywords.




Figure 2. yourSky custom mosaic form interface.

2. **Resolution:** Required parameter(s), analogous to the CDELTA1 and CDELTA2 FITS keywords, which specify the pixel size in degrees in each of the two image dimensions *at the mosaic center of projection*.
3. **Radius in degrees:** Optional parameter that limits the mosaic size using degrees from the mosaic center. If not specified, the radius is determined automatically from the region of coverage of the input image plates.
4. **Width and height in pixels:** Optional parameters, analogous to the NAXIS1 and NAXIS2 FITS keywords, that limit the mosaic size using a specific number of pixels. These parameters supersede the radius in degrees if that is given as well. If not specified, the radius is determined automatically from the region of coverage of the input image plates.
5. **Coordinate system:** Required parameter, analogous to the first half of the CTYPE1 and CTYPE2 FITS keyword values, that specifies the alignment of the mosaic axes in 3-D space. Four coordinate systems are supported:
 - Galactic.
 - Ecliptic.
 - J2000 Equatorial.
 - B1950 Equatorial.
6. **Projection:** Required parameter, analogous to the second half of the CTYPE1 and CTYPE2 FITS keyword values, that specifies how locations on the celestial sphere are mapped to the image projection plane. All of the projections specified by WCS are supported:

- Linear (LIN).
 - Gnomonic or Tangent Plane (TAN).
 - Orthographic/Synthesis (SIN).
 - Stereographic (STG).
 - Zenithal/Azimuthal Perspective (AZP).
 - Zenithal/Azimuthal Equidistant (ARC).
 - Zenithal/Azimuthal Polynomial (ZPN).
 - Zenithal/Azimuthal Equal Area (ZEA).
 - Airy (AIR).
 - Cylindrical Perspective (CYP).
 - Cartesian (CAR).
 - Mercator (MER).
 - Cylindrical Equal Area (CEA).
 - Conic Perspective (COP).
 - Conic Equidistant (COD).
 - Conic Equal Area (COE).
 - Conic Orthomorphic (COO).
 - Bonne (BON).
 - Polyconic (PCO).
 - Sanson-Flamsteed Sinusoidal (SFL).
 - Parabolic (PAR).
 - Hammer-Aitoff (AIT).
 - Mollweide (MOL).
 - COBE Quadrilateralized Spherical Cube (CSC).
 - Quadrilateralized Spherical Cube (QSC).
 - Tangential Spherical Cube (TSC).
 - Digitized Sky Survey Plate Solution (DSS).
 - Plate fit polynomials (PLT).
7. **Image Format:** Required parameter that specifies the output mosaic image format (FITS is recommended but not required). The following image formats are currently supported:
- FITS.
 - JPEG.
 - PGM.
 - PNG.
 - TIFF.
 - Raw Data.
8. **Data Type:** Required parameter that specifies the data type of the mosaic pixels. This is analogous to the BITPIX FITS keyword, but for the Raw Data image format some additional data types are provided that are not supported by FITS. The following data types are supported:
- 8-bit unsigned integer.
 - 8-bit signed integer (not supported by FITS).
 - 16-bit unsigned integer (not supported by FITS).
 - 16-bit signed integer.
 - 32-bit unsigned integer (not supported by FITS).
 - 32-bit signed integer.
 - Single precision floating point.
 - Double precision floating point.
9. **Quantization Extrema:** Optional parameters that specify the minimum and maximum over which to stretch the input pixel values for those data types that require quantization to a limited number of output bits per pixel (especially, 8-bit and 16-bit integers). The user can specify these values to control how many gray levels in the output mosaic are assigned to low or high intensity regions of the sky.

10. **Pixel Masks:** Optional masks may be specified to discard pixels around the outer perimeter or from particular rectangular regions in each input image.
11. **Background Matching:** Logical parameter that specifies whether or not yourSky should attempt to match the background intensities among the input images that comprise a mosaic in an attempt to produce a mosaic that is as seamless as possible.

3.2 Parallel Mosaicking Algorithm

The yourSky mosaicking algorithm is designed to be able to handle arbitrarily large mosaic requests from typical small requests covering a single celestial object to all-sky mosaics at full resolution. Also, the algorithm is efficient in the face of arbitrarily sized input image plates, so that yourSky can be extended to support other image archives without consideration of the native image partitioning scheme used by the archive. For example, the two surveys currently accessible by yourSky have drastically different native image partitioning schemes, from millions of small 2 MB image plates in the case of 2MASS to thousands of much larger 1 GB plates in the case of DPOSS. In addition, the mosaicking algorithm is designed to support arbitrary mappings from input image pixels to output mosaic pixels.

The mosaicking proceeds in two phases, *Analysis* and *Build*. During *Analysis*, the following is accomplished. First, the mosaic width and height are determined if they are not provided explicitly as part of the user-specified parameter set. Second, the pixel coordinates that intersect the mosaic are determined for each input image along with the corresponding intersection coordinates from the mosaic. These coordinates are used to set loop bounds and buffer sizes during the *Build* phase. Third, in cases where the data type requires quantization to a limited number of output bits per pixel in the output mosaic (e.g. 8-bit and 16-bit integers), the minimum and maximum over which the pixel values should be quantized are determined if these extrema are not specified explicitly as part of the user-specified parameter set. Fourth, if background matching is to be done, the intensity correction for each input image plate is determined.

During the *Build* phase the information gathered during *Analysis* is used to construct the custom mosaic. If the mosaic is to be lower resolution than the input image plates, the outer loop is over the input image pixels and the mosaic pixel values are calculated as the average of the input pixels for which the pixel center falls within the mosaic pixel region of coverage. If the mosaic is to be roughly the same or higher resolution than the input image plates, the outer loop is over the mosaic pixels and the mosaic pixel values are computed to be the result of sampling from the input images using bilinear interpolation. In either case, mapping from input pixel coordinates to output pixel coordinates is done by first mapping from input pixels to a location on the sky, then mapping from the sky coordinates to the output pixel coordinates, as illustrated in Figure 3.




Figure 3. Mapping from input pixel coordinates to output pixel coordinates is done in two steps. First the input coordinates are mapped to a position on the sky, then that position on the sky is mapped to the output mosaic coordinates.

The mosaicking proceeds in parallel during both Analysis and Build, with each processor being assigned a subset of the input image pixels. By default the input images are assigned to processors in a round robin fashion, with one processor per image, but the user can reconfigure this at run-time by specifying the number of processors to be assigned to each input image. Assigning multiple processors to each input image plate dramatically improves efficiency for archives,




Figure 4. IRAS all-sky mosaic in the Cartesian (CAR) projection constructed at 90 arc second resolution from 430 IRAS image plates in each of four wavelengths. The full resolution mosaic is 14,400 x 7,200 eight-bit unsigned integer pixels.




Figure 5. 2MASS H band center of the galaxy mosaic constructed from 16 2MASS image plates at 1 arc second resolution without (left) and with (right) background matching.

such as DPOSS, that have such large image plates that only a single or a few input image plates are required for a typical mosaic request. If multiple processors are assigned to each input image, a group synchronization among the processors assigned to the same image is required for each image so that Analysis results can be accumulated and shared. Also, in all cases, a global synchronization is required between the Analysis and Build phases so that Analysis results that relate to the entire mosaic, such as pixel value distributions required to calculate the appropriate quantization extrema, can be accumulated and shared. The software should be portable because it is written in ANSI C and all inter-processor communication and synchronization is done using Message Passing Interface (MPI), which has been implemented on many platforms [13].

3.3 Sample Mosaics

Some sample image mosaics[‡], constructed with the yourSky custom image mosaicking software, are shown in Figures 4, 5 and 6. Figure 4 shows a full resolution, all-sky, 90 arc second resolution, 14,400 x 7,200 pixel mosaic constructed from 430 Infrared Astronomical Satellite (IRAS) [14] image plates in each of 4 wavelengths, 12, 25, 60, and 100 microns. High performance exploration of this and other large datasets is possible using visualization software developed previously at JPL [15,16]. Figure 5 shows a center of the galaxy mosaic from the 2MASS H band before and after background matching is performed. The striped appearance without background matching is primarily due to atmospheric effects that become more pronounced as the path length through the atmosphere gets longer at different look angles. The background matching algorithm used by yourSky results in a more seamless mosaic, but edge effects are still visible. This background matching algorithm will be improved under a NASA Grand Challenge effort to develop




Figure 6. DPOSS F band mosaic of M31 at 1 arc second resolution. The image shown here is the center of a much larger 34,816 x 36,352 single precision floating point mosaic constructed from 9 DPOSS plates.

[‡] The input data used to construct these mosaics was provided courtesy of the respective archives, but the images shown here are derived products, constructed by yourSky, and as such have not been endorsed by the respective archives, should not be expected to be suitable for all scientific purposes, and should not be taken to be indicative of the image quality of the raw data provided by the respective archives.

Mosaicking Time on Origin 2000




Figure 7. Mosaicking software performance on an SGI Origin 2000 for two different mosaic parameters, a 2x2 degree 1 arc second resolution mosaic of the galactic center in 2MASS with one processor assigned to each input image plate and a 2x2 degree 10 arc second resolution mosaic of M31 in DPOSS with all processors assigned to each input image plate.

high performance cornerstone technologies for the NVO. This effort, called Montage, has staged code improvement deliverables with the final deliverable scheduled for January, 2005 [17]. Figure 6 shows a DPOSS F band mosaic of M31 in a Galactic Tangent Plane projection. The mosaic shown in the figure is the center part of a larger 34,816 x 36,352 single precision floating point mosaic constructed from 9 DPOSS plates at full 1 arc second resolution.

3.4 Performance Results

In this section we show timing results for the yourSky mosaicking software running on a SGI Origin 2000 with 300 MHz R12000 processors and 512 MB of RAM per processor. The first test mosaic is a 2 x 2 degree 2MASS mosaic of the galactic center at full one arc second resolution. The resulting mosaic is 7,201 x 7,201 single precision floating point pixels (207 MB) in size in a Galactic Tangent Plane projection, constructed from 174 2MASS plates totaling 365 MB in size. The second test mosaic is a 2 x 2 degree DPOSS mosaic of M31 at 10 arc seconds resolution, resulting in a 721 x 721 single precision floating point mosaic in the Galactic Tangent Plane projection. This 2.1 MB mosaic was constructed from 2 DPOSS plates (total size 2.2 GB). No background matching was performed for this test. Figure 7 shows the wall clock time required to construct the mosaics on different numbers of processors on the Origin 2000. For the 2MASS mosaic, one processor was assigned to each input image plate, but all the processors were assigned to each input image plate for DPOSS. The plot shows the scaling curves for up to 64 processors.

4 REQUEST MANAGEMENT

Simultaneous mosaic requests are accepted from a simple HTML form submitted from the yourSky mosaic request web page at <http://yourSky.jpl.nasa.gov>, and queued on the yourSky server by a Common Gateway Interface (CGI) program interfacing with the Apache web server. The mosaic parameters for each request are stored on the yourSky server along

with the identity of the user that submitted the request. The yourSky Mosaic Request Manager, shown in the architecture diagram in Figure 1, needs to locate these mosaic requests and assign them one at a time to the Mosaic Request Handler. A user priority scheme is in place that starts off all users with equal priority. As requests are processed the user priorities change based on the number of mosaic pixels produced by each user in the past period referred to as the Priority Window, currently set to 24 hours. Users with the least number of mosaic pixels produced in the Priority Window period have highest priority for future mosaic requests. Furthermore, a mosaic request in progress that has had to wait for input image plate retrieval from a tape archive gets the highest priority to run next once all of the required input images have been retrieved. This ensures that all users get a chance to have their mosaic constructed and no single user will dominate all the available resources.

5 PLATE COVERAGE DATABASE

In order to be accessible by yourSky, all member surveys have to be included in the Plate Coverage Database that contains the minima and maxima of the longitudes (right ascensions) and latitudes (declinations) in each of the supported coordinate systems for all of the input image plates. The yourSky Mosaic Request Handler queries this database to determine which input image plates are needed to fulfill each mosaic request. The open source database, MySQL, is used to store this plate coverage information [18,19]. The result of the query to the Plate Coverage Database is a list of the input image plates that are required to fulfill the mosaic request. These input image plates are retrieved from the appropriate tape archives, staged in a local data cache, and provided as an input to the image mosaicking software described in Section 3. The plate coverage database is also accessible as a standalone service, called the yourSky Archive Database Query, at <http://yourSky.jpl.nasa.gov/query/index.html>.

6 DATA MANAGEMENT

A data management scheme is implemented on the yourSky server to manage both a data cache for the input image plates used to fulfill recent mosaic requests and a work area used to store recently constructed mosaics until they are downloaded.

The input data cache is maintained at a fixed size with image plates discarded on a least recently used basis. This enables mosaics to be recomputed with some changes to the custom request parameters without having to repeat the input image plate retrieval from the tape archives if the new request is resubmitted shortly after the original request. Also, mosaics of popular regions of the sky are likely to have their input image plates already cached on the yourSky server from previous requests, so they can be constructed more quickly.

Mosaics that have been completed are stored in a work area from where the appropriate users may download them. Currently mosaics are purged after a 24 hour period expires.

7 DATA ARCHIVE ACCESS

All of the publicly released data from two archives are currently accessible by yourSky, the Digitized Palomar Observatory Sky Survey (DPOSS) and the Two Micron All Sky Survey (2MASS).

DPOSS has captured nearly the entire northern sky at 1 arc second resolution in three wavelengths, 480 nm (J Band - blue), 650 nm (F Band - red), and 850 nm (N Band - near-infrared). The survey data were captured on photographic plates by the 48" Oschin Telescope at the Palomar Observatory in California [8,9]. The total size of the DPOSS data accessible by yourSky is roughly 3 TB, stored in over 2,600 overlapping image plates on the High Performance Storage System (HPSS) [20] at the Center for Advanced Computing Research (CACR) at the California Institute of Technology. The DPOSS plates are each about 1 GB in size and contain 23,552 x 23,552 pixels covering a roughly 6.5 x 6.5 degree region of the sky. The yourSky server uses a client program called the Hierarchical Storage Interface (HSI) to retrieve selected DPOSS plates in batch mode from the HPSS [21].

2MASS has captured nearly the entire sky at 1 arc second resolution in three near-infrared wavelengths, 1.25 μm (J Band), 1.65 μm (H Band), and 2.17 μm (K_s Band). The survey data were captured using two 1.3 meter telescopes, one

at Mt. Hopkins, AZ and one at the Cerro Tololo Inter-American Observatory (CTIO) in Chile [10]. The 2MASS archives contain roughly 10 TB of images and the subset that has been publicly released to date, nearly 4 TB, is fully accessible by yourSky. This 4 TB of data is stored in about 1.8 million overlapping plates managed by the Storage Resource Broker (SRB) at the San Diego Supercomputer Center (SDSC). Each 2MASS plate is about 2 MB in size and contains 512 x 1,024 pixels covering a roughly 0.15 x 0.30 degree region of the sky. The SRB is a scalable client-server system that provides a uniform interface for connecting to heterogeneous data resources, transparently manages replicas of data collections, and organizes data into “containers” for efficient access [22]. The yourSky server uses a set of client programs called SRB Tools to access selected 2MASS plates in batch mode from the SRB.

8 CONCLUSION

The yourSky architecture and software for desktop access to custom astronomical image mosaics has been presented. The architecture of yourSky allows it to exploit high performance computing resources, supercomputers and high bandwidth networks, on the server side. However, at the same time it is widely usable from virtually anywhere because the architecture also supports very lightweight computing resources on the client side, e.g., ordinary desktop computers with low bandwidth network connections. Since the user interface to yourSky is a simple HTML form submitted at <http://yourSky.jpl.nasa.gov>, the only client software required to use yourSky is the ubiquitous web browser, which most of the potential users probably already has and knows how to use. This combination of being deployed in a high performance computing and communications environment while allowing access through simple portals running on the desktop makes yourSky a good match for the loosely coupled, distributed architecture of the National Virtual Observatory (NVO).

The yourSky software includes subsystems for: (i) construction of the image mosaics on multiprocessor systems, (ii) managing simultaneous user requests, (iii) determining which image plates from member surveys are required to fulfill a given request, (iv) caching input image plates and the output mosaics between requests, and (v) retrieving input image plates from remote archives. The parallel image mosaicking software emphasizes custom access to mosaics, allowing the user to specify parameters that describe the mosaic to be built, including datasets to be used, location on the sky, size of the mosaic, resolution, coordinate system, projection, data type, and image format. Performance results indicate that the yourSky mosaicking software scales well on up to 64 processors, the largest size partition available on the SGI Origin 2000 at the Jet Propulsion Laboratory.

A number of current and future efforts are planned to extend the capabilities of yourSky. First, a graphical front end to yourSky is under construction that will allow web-based pan and zoom over a medium resolution, all sky representation of each dataset, along with the capability of plotting symbol overlays from catalogs on top of the imagery. This browse mode for yourSky will be interoperable with the existing yourSky server so that it can be used to help formulate mosaic requests to yourSky. Another effort underway is to modify yourSky so that the mosaicking code is run on computational grids, such as NASA’s Information Power Grid (IPG) [23] or NSF’s TeraGrid [24], rather than on a local server. The collective computational power of the existing Grid infrastructure far exceeds that of any local server so this would enable handling of much larger or more numerous mosaic requests. Finally, yourSky serves as a pathfinder for the Montage project [17], which is sponsored by the NASA Earth Science Technology Office (ESTO) – Computational Technologies (CT) Project as a Grand Challenge effort. Montage will extend yourSky capabilities in the following ways: (i) science quality will be improved by ensuring flux per area is preserved in the output mosaics; (ii) performance will be improved to meet aggressive throughput milestones; (iii) Montage will be made fully interoperable with the NVO infrastructure, including interoperability with other archives, analysis and visualization tools, and use of any data format standards adopted by the NVO; and (iv) Montage will be fully interoperable with the grid infrastructure being established by the TeraGrid. Montage has staged code improvement deliverables with the final deliverable scheduled for January, 2005.

ACKNOWLEDGEMENTS

This research was sponsored by the Space Science Applications of Information Technology (SAIT) Program, Office of Space Science, National Aeronautics and Space Administration (NASA). S.G.D., R.B., and A.M. acknowledge a partial support from the NASA Applied Information Systems Research Program (AISRP). Processing of DPOSS was

supported in part by a generous grant from the Norris Foundation and other private donors. Atlas images obtained as part of the Two Micron All Sky Survey (2MASS), a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by NASA and the National Science Foundation.

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

1. The National Virtual Observatory (USA), <http://us-vo.org>.
2. National Virtual Observatory Science Definition Team Report: *Towards the National Virtual Observatory*, April, 2002, <http://nvosdt.org>.
3. NVO White Paper: *Toward a National Virtual Observatory: Science Goals, Technical Challenges, and Implementation Plan*, in *Virtual Observatories of the Future*, R.J. Brunner, S.G. Djorgovski, and A.S. Szalay, eds., Astronomical Society of the Pacific Conference Series, Vol. 225, pp. 353-372, 2001. Also available online at <http://www.arxiv.org/abs/astro-ph/0108115> and at <http://nvosdt.org>.
4. The AstroGrid (UK), <http://www.astrogrid.org>.
5. The Astrophysical Virtual Observatory (Europe), <http://www.eso.org/avo>.
6. The Australian Virtual Observatory (Australia), <http://www.atnf.csiro.au/people/rnorris/AVO>.
7. The Virtual Observatory – India, <http://vo.iucaa.ernet.in/~voi/html/homepage.html>.
8. S.G. Djorgovski, R.R. Gal, S.C. Odewahn, R.R. de Carvalho, R. Brunner, G. Longo, and R. Scaramella, *The Palomar Digital Sky Survey (DPOSS)*, <http://www.arXiv.org/abs/astro-ph/9809187>.
9. The Digitized Palomar Observatory Sky Survey (DPOSS), <http://www.astro.caltech.edu/~george/dposs>.
10. The Two Micron All Sky Survey (2MASS), <http://www.ipac.caltech.edu/2mass>.
11. The Flexible Image Transport System (FITS), <http://fits.gsfc.nasa.gov>, <http://www.cv.nrao.edu/fits>.
12. E.W. Greisen and M. Calabretta, *Representation of Celestial Coordinates In FITS*, <http://www.atnf.csiro.au/people/mcalabre/WCS.htm>.
13. MPI, The Message Passing Interface Standard, <http://www-unix.mcs.anl.gov/mpi>.
14. The Infrared Astronomical Satellite (IRAS), <http://www.ipac.caltech.edu/ipac/iras/iras.html>.
15. J.C. Jacob and L. Plesea, *Fusion, Visualization and Analysis Framework for Large, Distributed Datasets*, IEEE Aerospace Conference, 2001, ISBN 0-7803-6600-X.
16. J.C. Jacob and L.E. Husman, *Large Scale Visualization of Digital Sky Surveys*, *Virtual Observatories of the Future*, R.J. Brunner, S.G. Djorgovski, and A.S. Szalay, eds., Astronomical Society of the Pacific Conference Series, Vol. 225, pp. 291-296, 2001.
17. Montage: High Performance Cornerstone Technologies for the National Virtual Observatory, <http://montage.ipac.caltech.edu>.
18. M. Widenius, MySQL AB, D. Axmark, *MySQL Reference Manual: Documentation from the Source*, O'Reilly, June, 2002, ISBN 0-59600-265-3.
19. MySQL, <http://www.mysql.com>.
20. The High Performance Storage System (HPSS) at Caltech's Center for Advanced Computing Research, <http://www.cacr.caltech.edu/resources/HPSS/index.html>.
21. Hierarchical Storage Interface (HSI) for High Performance Storage Systems (HPSS), <http://www.sdsc.edu/Storage/hsi>.
22. The Storage Resource Broker (SRB) at San Diego Supercomputer Center (SDSC), <http://www.npaci.edu/SRB>.
23. The Information Power Grid (IPG), <http://www.ipg.nasa.gov>.
24. The TeraGrid, <http://www.teragrid.org>.