NVO-TeraGrid
First Year Results

TeraGrid Utilization Annual Report
for the
National Virtual Observatory
Multi-year, Large Research Collaboration

Roy **Williams**, PI	*California Institute of Technology*
Andrew **Connolly**, coI	*University of Pittsburgh*
Jeffrey **Gardner**, coI	*Pittsburgh Supercomputing Center*

Senior Personnel

Bruce **Berriman**	*California Institute of Technology*
Leesa **Brieger**	*San Diego Supercomputer Center*
Sharon **Brunett**	*California Institute of Technology*
Tamas **Budavari**	*Johns Hopkins University*
Joerg **Colberg**	*University of Pittsburgh*
Ewa **Deelman**	*University of Southern California*
George **Djorgovski**	*California Institute of Technology*
John **Good**	*California Institute of Technology*
Jim **Gray**	*Microsoft*
Sebastian **Jester**	*Fermilab*
George **Kremenek**	*San Diego Supercomputer Center*
Simon **Krughoff**	*University of Pittsburgh*
Reagan **Moore**	*San Diego Supercomputer Center*
Robert **Nichol**	*University of Portsmouth, UK*
Ryan **Scranton**	*University of Pittsburgh*
Chris **Stoughton**	*Fermilab*
Alex **Szalay**	*Johns Hopkins University*
Dan **Vanden Berk**	*Pennsylvania State University*

Abstract

The NSF National Virtual Observatory (NVO) is a multiyear effort to build tools, services, registries, protocols, and standards that can extract the full knowledge content of massive, multi-frequency data sets. Here we detail our progress from the first year, and plans for the second year of our effort to combine the computational resources of the TeraGrid with the NVO. The work includes: creation of derived image products from multi-terabyte datasets; multiwavelength and multitemporal image federation; analysis of 3-point correlation in galaxies; fitting models to thousands of galaxy spectra; Monte-Carlo modeling of early-Universe models; processing of the ongoing 13 Tbyte Palomar-Quest synoptic survey. TeraGrid impact is broadened through science gateways and grid services, specifically for: exposing massive multi-terabyte datasets through NVO protocols; source extraction and cross-matching; mosaicking, image reprojection and coaddition, as well as authentication mechanisms for such services.

1 Introduction

The development of the National Virtual Observatory (NVO; http//www.us-vo.org), an NSF sponsored Information Technology Research project, together with related efforts ongoing world-wide (through the International Virtual Observatory Alliance) has the potential to revolutionize scientific research in astrophysics. The NVO is designed to support the federation of astronomical sky surveys; seamlessly interlacing data across the full electromagnetic spectrum. It provides tools to explore within and extract data from these massive, multi-frequency surveys. Such unprecedented access provides an astronomer with the ability to determine the detailed correlations within these data, to understand the physical processes that give rise to these correlations and to potentially identify new classes of astrophysical phenomena.

The data processing applications identified within our original proposal and in this renewal request are geared to cover a broad range of scientific questions and methodologies that are representative of the majority of future NVO grid services. They encompass algorithms that comprise large numbers of small, independent work units as well as tightly coupled parallel applications that can utilize the massive parallelism of each TeraGrid resource. Each project uses the TeraGrid to find answers to fundamental science questions as well as providing an important diagnostic capability for understanding and optimizing the interface between the Grid and the NVO.

In this renewal proposal, we request an allocation of 407,000 Service Units to complete our six original NVO Teragrid applications as well as the development of new services for the astronomical community. These new services will greatly broaden the use of TeraGrid resources, since they do not require sophisticated knowledge of TeraGrid architecture, but rather the ability to work with a web form. They constitute the first wave of “TeraGrid Science Gateways”.

All applications discussed here are full-fledged and, prior to their porting to the TeraGrid, have been developed and tested in smaller grid environments. Each is designed yield new and valuable scientific results that would not have been possible without TeraGrid resources.

1.1 Summary of our First Year on the TeraGrid

During the first year of this multi-year effort, we focused on several applications that are representative of the kinds of computing that the NVO community will need to perform on the TeraGrid. These initial projects were geared towards producing valuable scientific objectives. Over the course of our initial TeraGrid allocation we have been able to accomplish a number of scientific milestones as well as to make substantial strides in developing new software to better adapt the TeraGrid to our needs. To date, all applications discussed in the initial proposal have been implemented and run on the TeraGrid. Three of these applications have yeiled initial science results. As of the beginning of December, we have used over 130,000 SUs out of our awarded 257,000, or roughly half of our allocation.

We will discuss the specific accomplishments of each Grid application in Section 2. Here we note the principal challenges we have addressed in utilizing the TeraGrid. Many TeraGrid accounts were not activated until June 2004, so usage as of December 1 reflects roughly 6 months work. Thus, usage of half of our yearly allocation in the first 6 months we have had access to the TeraGrid is consistent with our expectations. A second challenge arose due to the fact that we were not able to conduct projects with large numbers of independent jobs without first developing additional software. In section 4.1 we discuss our solution, GridShell, to enable this mode of computation and its utility in a general TeraGrid framework.

2 Science Applications

The bulk of our allocation this year will focus on continued development of last year’s astronomy applications and using what we learned last year to get them into true production mode on the TeraGrid. Throughout the next section we reiterate the scientific drivers for each of these components, summarize our progress so far, and quantify our code enhancements and scientific goals for the upcoming year.

2.1 Image Reprojection and Federation

Roy Williams, John Good, California Institute of Technology
Leesa Brieger, San Diego Supercomputer Center

Many of the high-quality sky surveys concentrate on the catalog output as the primary product. However, this science product comes from the images; it is a set of large mosaics with careful reprojection to minimize astrometric and photometric distortion. All of the following surveys are being projected to the same pixel grids (the "hyperatlas" set, see Williams 2003), so that these surveys can be federated by stacking and deeply mined.

The 2MASS image dataset is a high-quality infrared sky survey of about 10 Tbyte. It has been replicated at Caltech on disk, and also at SDSC on the SRB. SDSC is using the TeraGrid Datastar machines and the NASA Montage software to reproject this dataset. Tests indicate that the whole sky will use about 50000 CPU-hours. This work will enter full production in 2005.

The Digital Palomar Observatory Sky Survey (DPOSS) is a 3 Tbyte optical sky survey, that has been rebuilt and cleaned in 2004 using TeraGrid resources. A derived version has been created and stored on the Caltech PVFS, which has registration marks cropped, with flattening and re-moding, to allow for mosaicking. Some 20 Hyperatlas pages have been built with the very slow Montage 1.7, using some 4000 node hours at Caltech to produce 30 Gbyte. The algorithm is being rebuilt, replacing the complex polynomial plate solution with a much simpler piecewise TAN projection, and we expect a 20-fold speedup. The reprojection is more complex than 2MASS, but the survey covers only the northern hemisphere, and we estimate that 50000 CPU-hours will be needed. The full DPOSS Hyperatlas will be produced in 2005.

The Palomar-Quest survey is running a catalog pipeline at NCSA and an image pipeline at Caltech, to reproject images to Hyperatlas so they can be stacked. This will allow searching for very faint quasars, as well as image subtraction to find transient sky sources.

2.2 N-Point Correlation Functions of Galaxies

Jeffrey Gardner, Pittsburgh Supercomputing Center
Andrew Connolly and Joerg Colberg, University of Pittsburgh
Jim Gray, Microsoft Corporation
Bob Nichol, University of Portsmouth

2.2.1 Science Goals and Implementation

A fundamental ambition of any scientific analysis is to quantify any clustering within the data. Cosmologists have chosen to characterize such clustering using n-point correlation functions (see Peebles 1980) which are a non-parametric measure of the higher order moments of a distribution of points. For example, the 2-point function of galaxies in the Universe quantifies the excess (or deficit) of pairs of galaxies as a function of their separations compared to a Poisson distribution. Likewise, the 3-point function is the excess (or deficit) of triplets (triangles), and so on up the hierarchy of combinations (4-point, 7-point, *etc*.). If the structure in the Universe originated from Gaussian initial conditions predicted by Inflation, then all the higher order correlation functions (3-point and above) would only depend upon the 2-point function, the lowest order correlation function. Any other set of initial conditions, or non-linear physics, would result in a non-trivial set of higher order correlation functions and therefore, it is imperative that we accurately measure these functions to test our fundamental cosmological theories and assumptions. Precision measurements of the higher-order correlations of galaxies are now possible due to the availability of high quality data from the SDSS survey.

2.2.2 Progress Report and Future Requirements

The calculation of the N-point correlation function requires parallelism that is even more tightly-coupled than many parallel simulations, and is the most difficult of our efforts in achieving scalability. So far we have calculated spatial 2- and 3-point correlation functions for the Sloan Digital Sky Survey Data Release

2 (SDSS DR2) which consists of 260,000 galaxies with measured redshifts. By the time our current allocation expires, we will have analyzed SDSS DR4 which was just released to the collaboration in October and contains 478,000 galaxies with measured redshifts. In addition to spatial N-point (in which galaxy positions are specified in traditional 3-D Cartesian coordinates) we have also performed "projected" N-point in which distances are calculated in coordinates projected on the sky. Using the projected coordinate system, while substantially more time-consuming, minimizes errors due to uncertainties in measuring a galaxy's distance. We have tested the accuracy of our algorithm on mock datasets (in a cubical volume) of up to 10 million particles, achieving roughly 90% speedup when running on 256 processors. In analyzing the actual SDSS datasets, however, we discovered that the complex geometry of the survey volume introduces extreme load-balancing requirements. Our existing code attempted to load balance by predicting the computation requirements of each galaxy before the start of the calculation, then distributing galaxies in such as way as to equalize the workload amongst processors. Unfortunately, this approach could not deal with the complex survey boundaries of SDSS. We were able to work around this limitation for spatial N-point, but projected N-point for the SDSS required a substantial retooling of our software to add dynamic load balancing.

In our new parallel N-point code "*N*Tropy" we implement two new features: dynamic load balancing and permutation tracking. With dynamic load-balancing, *N*Tropy threads will be able to negotiate their workload on the fly, allowing all of the processors to be kept busy throughout the calculation. Effective load-balancing will allow us to achieve the scalability required for projected 3-point calculations of the SDSS DR4 and beyond. An early implementation of *N*Tropy was presented at ADASS 2004.

Our second enhancement in the new *N*Tropy code aims at increasing the flexibility of our analysis and improving the value of scientific results. Previously, we were limited to N-point configurations that either had all sides the same, or all sides different (e.g. for 3-point we had to use an equilateral triangle, or a triangle with all 3 sides different lengths). Allowing arbitrary N-point configurations, yet ensuring that each set of points is only counted once no matter how many permutations of the N-point configuration it satisfies introduces substantial bookkeeping overhead. *N*Tropy maintains a permutation list to make sure that each N-tuple is counted only once, even if it satisfies more than one N-point configuration. This added capability introduces roughly a factor of 4 expense in run-time.

*N*Tropy incorporates all of the scalability enhancements of our previous generation code detailed in our original proposal, plus it introduces dynamic load balancing and permutation tracking. To calculate the projected 3-point function for SDSS DR4 (478,000 galaxies) for 1000 triangle configurations will require approximately 100,000 node-hours on a 1.3 GHz Itanium2 machine. Spatial N-point for SDSS DR4 will require an additional 20,000 node-hours (this is less than last year's spatial N-point request, because based on last year's finding we will use fewer simulated Universes for our error estimates).

2.3 Resolving Star Formation in Galaxies

Andrew Connolly, University of Pittsburgh
Alex Szalay and Tamas Budavari, Johns Hopkins University

2.3.1 Science Goals and Implementation

The science goals for this program were to study the physical processes that drive the formation and evolution of galaxies within our universe by analyzing the distribution of star formation within individual galaxies. Earlier work has shown that the environment in which a galaxy resides can significantly affect star formation activity within the galaxy itself (Gomez et al 2003, Lewis et al 2002). Mechanisms proposed to explain this relation range from galaxy harassment through to ram pressure stripping of gas as a galaxy enters an overdense environment. Current studies have, however, been limited due to the relative small sample of galaxies. We proposed to gain an order of magnitude increase in sample size by utilizing the data available through the NVO. From multicolor photometric data, estimates of the age, metallicity, dust and star formation rate in the individual pixels within galaxies are derived by fitting spectral energy distributions to their fluxes (Conti et al 2003).

We proposed applying the fitting techniques described above to a representative sample of galaxies selected from the SDSS (approximately 500,000 resolved galaxy images). The galaxies in this sample have a median redshift of z=0.1 and a typical size of 50x50 pixels (corresponding to a spatial resolution of 0.5 kpc per pixel). Our implementation of the SED fitting procedure is built upon a multi-resolutional grid (standard hill climbing search techniques are not well suited to this task due to the presence of multiple minima within the likelihood space and the need to characterize these local minima not just find a simple global minimum for the fit). Local minima within the four dimensional space are initially identified in the lowest resolution mode of the grid. The grid is then progressively refined to increase the resolution of the likelihood space around each of these local minima. We marginalize the likelihoods for each parameter (star formation, age, metallicity and dust) in order to characterize likelihood functions. All software is written in ANSI C++ and has been compiled on the NCSA TeraGrid cluster.

2.3.2 Progress Report and Future Requirements

The lack of a Condor queue in the initial version of the TeraGrid held up development of this application as the initial implementation had been written to utilize Condor. From September onward we have run the pixel-z application on the TeraGrid using the GridShell framework for both images and catalog data. The implementation assumes access to a common set of FITS format files for each of the images. The SED fitting procedure built upon a multi-resolutional grid has been demonstrated to work effectively and returns the estimated star formation rate, age, extinction, metallicity and associated errors in a FITS format file. Implementation has focused on the IA64 cluster at NCSA and has been run on up to 128 processor elements at a single time. A total of approximately 3000 service units have been utilized by this part of the program. Over the coming months we expect to enter a production run to analyze all images for the SDSS spectroscopic data set. Tests on the NCSA cluster show that at full resolution we require approximately 320s to analyze a 50x50 pixel image (the typical image resolution of the SDSS at these bright magnitudes). To characterize the 500,000 images within the SDSS, we will require approximately 44,000 CPU hours. Storing just the minima and their associated errors (two-sided as the errors are not symmetric) will require approximately 100-200 GB of disk space. By the start of the second year of this proposal we estimate that we will have completed runs for approximately 30% of the current data and, therefore, request a further 30,000 service units for this application in this proposal.

2.4 Fitting QSO Spectra

Dan Vanden Berk, Pennsylvania State University
Sebastian Jester and Chris Stoughton, Fermilab

2.4.1 Science Goals and Implementation

The goal of this project is to measure the physical parameters of quasars and other active galactic nuclei by fitting models to their UV/optical spectra. For example, the mass of the central black hole can be estimated from the width of the broad emission lines and the luminosity of the continuum emission. The difficulty lies in the complexity of the quasar spectra – there are several different continuum components, contributions from the host galaxy, and numerous emission lines that are always blended together. The dimensions of the model parameter space usually exceed 100 per spectrum, and all of the components must be fit simultaneously. Our dataset consists of over 50,000 spectra from the SDSS component of the NVO – we needed a robust technique that does not require manual intervention or monitoring of individual spectra. We developed a genetic algorithm-based (GA) fitting routine that satisfies these requirements. The routine employs many of the standard tools of GA’s, such as crossover and mutation, and we have added a number of novel techniques that improve the performance of the algorithm. The use of GA’s for spectral fitting is not new (e.g. McIntosh et al. 1998) but they are not yet widely employed in astronomy. Our routine is quite general and it has been easily adapted to other parameter search problems.

All software is written in ANSI C++ and has been compiled on the NCSA TeraGrid cluster. Jobs are run in trivially parallel mode, one spectrum per processor. Individual ASCII spectrum files are copied to a local disk (each file occupies about 200k of memory, so storage issues are not a concern). A perl script is used to generate RSL (resource specific language) files that specify the executable, command line arguments, stdout/stdin, and maximum running time for each job. There is one RSL file for each input file. Then globusrun is called to do the job submission. A log file is created that stores the globusrun job identifier string. The job progress is monitored by looking at the grid monitor web page. Output ASCII files of the fitting progress and the parameter solution and are written to the local disk. Each job takes about 1 hour to complete.

2.4.2 Progress Report and Future Requirements

We have so far run the fitting routine on about 17,000 real spectra, and we have run tests on about 10,000 synthetic quasar spectra to check the accuracy and robustness of the algorithm. The algorithm has been shown by these tests to work quite well in most cases. There appear to be some classes of quasars that are difficult to fit with a unique solution, particularly for the continuum parameters, although the emission lines are still fit consistently. We have made a number of small changes to the algorithm to handle these cases. We are currently writing a paper describing the algorithm and the initial results of running it on the TeraGrid. We plan to run the fitting routine on up to 50,000 quasar spectra from the latest data release of the SDSS (this includes multiple spectra of some objects, and does not include the 17,000 spectra already processed). Since it takes on average about 1 SU per spectrum, we require 50,000 SUs.

In the future, we will utilize the GridShell framework to more easily submit the large number of jobs we require. We also plan to run the routine using a web service. The user will specify a spectrum or list of spectra, the spectra will be retrieved from a remote server such as the NVO spectrum server, the job will be submitted, and the results made available to the user. Ideally, the solutions will be stored after they are completed to avoid redundancy, so that the solution set of a previously run spectrum can be retrieved immediately. A number of issues will certainly have to be resolved to make this work, but it will mark a significant step in allowing the “typical” astronomer to use our routine.

2.5 The Cosmic Microwave Background Grid

Ryan Scranton and Andrew Connolly, University of Pittsburgh

2.5.1 Science Goals and Implementation

The science goal for this program was the measurement of the Integrated Sachs-Wolfe (ISW) by cross-correlating maps of the Cosmic Microwave Background Radiation (CMBR) and foreground galaxy density. In a universe with dark energy, the accelerated expansion of the universe at late times exceeds the growth rate of large scale structure. This leads to a decay of the gravitational potentials containing these galaxies. During this process, photons from the surface of last scattering that we will detect as the CMBR pass through these potentials. In the absence of this decay, the gravitational blueshift these photons experience entering the potentials is balanced by an equal redshift upon leaving the potential. With the acceleration driving their decay, the potentials are flatter when the photons leave than when they entered, leading to a net increase in energy. This means that the CMBR temperature will be slightly higher in the direction of the highest galaxy densities and vice versa.

Although measuring this effect (using data from the Sloan Digital Sky Survey and Wilkinson Microwave Anisotropy Probe in our case) is easily accomplished on a single CPU, determining the significance of the detection is a demanding task. The calculation of the covariance matrix for the measurement is accomplished by generating large numbers of statistically equivalent CMBR realizations and cross-correlating them with the set of galaxy maps. In our case, we have four galaxy maps, each selected to span a narrow range is redshift. To fully exploit the statistical power of this data set, we need to generate a covrariance matrix that combines all four galaxy data sets cross-correlated with a given CMBR map, requiring an even larger number of realizations.

2.5.2 Progress Report and Future Requirements

Beginning in early October, we used the Condor queue in the Gridshell environment on the IA64 cluster at NCSA to generate the measurements for our covariance matrix. For each CMBR map, we generated 5000 realizations and cross-correlated them against each of the four galaxy maps. Each of these 5000 iterations took roughly 10 minutes and we ran them in groups of 50 spread across 120 processor elements using GridShell. Over the course of testing the routines and performing the final measurements, we have used approximately 12000 SUs. This analysis is currently being written up (Scranton et al 2005). Over the coming months, we plan to expand this measurement to incorporate all of the existing large galaxy surveys. This will require approximately 20000 SUs. The size of the final data set for this measurement is small (< 10 GB), but the intermediate stages of the pipeline require on order 300 GB of temporary storage at any given time.

3 Grid Services, Gateways, and Infrastructure

Part of this multi-year proposal involves enabling new grid services for use by the NVO as part of our to provide community-tailored science gateways to the TeraGrid. In this renewal we discuss the next set of services that are in prototype, that we propose to bring into production over the coming twelve months. For the development of these procedures we request a total of 87,000 service units (to ensure that the initial science goals can be achieved).

3.1 Data Replication and the NVO

George Kremenek, Reagan Moore, San Diego Supercomputer Center
Sharon Brunett, Roy Williams, California Institute of Technology

Many current sky surveys are available through NVO protocols from their curators, but the bandwidth may be quite low. For example, the Sloan Digital Sky Survey (~15 TB when complete) could be downloaded from Fermilab in 30 days. Many of the NVO grid applications described in this proposal require faster data access rates to remain compute and not I/O bound. To provide an order of magnitude increase in data access, the surveys and their NVO access services must be replicated on higher performance storage systems. Based upon the access rates that have already been demonstrated, it will be possible to analyze an entire 10 TB survey in 3 hours once it is replicated onto the TeraGrid storage. This will enable analyses that process an entire sky survey to a derivative product, to mine the entire archive for patterns, or to federate multiple such surveys.

Tests of the performance of the TeraGrid infrastructure have been tracked, with the following verified results:

- Sustained data transfer rates between SDSC and NCSA have been demonstrated at 250 MB/sec over the TeraGrid Network. A 10 TB collection can, therefore, be moved between TeraGrid sites in 12 hours.
- Data has moved at rates up to 5 GB/sec between the SAN and the TeraGrid Linux cluster at SDSC. With this rate a 10 TB collection can be accessed in 34 minutes.
- Replication of the 2-Micron All Sky Survey (10 TB) and the Digital Palomar Observatory Sky Survey (3 TB) onto the SAN at SDSC through use of the Storage Resource Broker. Final image products have been successfully installed on the TeraGrid.

Datasets that have been replicated to SDSC include the Sloan Digital Sky Survey (~15 TB), the United States Naval Observatory B collection (~8 TB), and the MACHO survey (~8 TB).

Datasets replicated at Caltech include the 2MASS atlas images (~10 TB), DPOSS (~3TB), and Palomar-Quest (~13 TB). Derivative products available include a tiled version with simplified plate solutions, and a cropped, flattened, re-moded version that is suitable for dynamic mosaicking services.

The 13 Tbyte Palomar-Quest database is available from Caltech on spinning disk, and is archived at NCSA on the tape-based Unitree system. Each clear night at the Palomar observatory adds another ~ 40 Gbyte to this archive. Derived products are also available from Caltech, only as a prototype in Jan 04, but the quantity and diversity will be greatly expanded in 2005. The speed of the TeraGrid trunk allows terabyte datasets to be effectively shared between widely separated sites (Caltech, SDSC, NCSA).

3.2 Image coaddition grid services for the NVO

The next decade will see the implementation of many imaging survey programs for astrophysics (e.g., Palomar-Quest, PanSTARRS, LSST, JDEM). These programs will focus on the time domain exploration of the local and distant universe. One product of these surveys will be multiple observations of the same part of the sky, both to identify variable sources and to detect faint objects. Coadding these time domain data (as a function of time and wavelength) will produce a deeper image of the sky and probe the formation of structure at high redshift. The first of the new NVO grid service will focus on developing a general TeraGrid facility to coadd series of images each taken under different observing conditions.

These applications will build upon the development work of Atlasmaker and Montage (see Section 2.1). Atlasmaker and Montage provide a series of images that are repixelized to a common scale and projection. These geometrically matched images will be coadded across the time and frequency domain using the techniques of Szalay, Connolly and Szokoly (1999) for stacking multi-frequency data and will build on the work of Kaiser (2002) for accounting for variable sky conditions.

3.3 Palomar-Quest Stacking

George Djorgovski, Ashish Mahabal, Roy Williams, Caltech

The Palomar-Quest sky survey creates approximately 40 Gbytes of new imaging data for each clear night at Palomar. The archive is currently ~13 Tbytes, stored on tape at NCSA, and replicated at Caltech on spinning disk, from which it is exposed through NVO services.

The Palomar-Quest image pipeline then creates “stackable” images from these, cleaning, flattening, and reprojecting to the Hyperatlas standard set of map projections. These derived products are already served in small quantity from the Caltech TeraGrid machine. In 2005, we will (a) create a much larger cache of this derivative data, and expose it through NVO services, and (b) implement the image pipeline as a service to run on demand. At this stage, we will have images from different observation times that are co-registered to exactly the same map projection, and can therefore be stacked, either additively for deep detection of faint quasars, or subtractively to detect transients on the sky.

3.4 SDSS Stacking

Andrew Connolly, University of Pittsburgh

In the next year of NVO-TeraGrid, we will also analyze a 100 square degree region along the celestial equator that has been observed multiple times by the SDSS. This region has been observed 15 times in each of 5 passbands with a total of 800 GB of imaging data divided into 130,000 images. The stacking of the 15 time steps (for all five colors) requires approximately 7 minutes of CPU time on a single 1.3 GHz Itanium2 processor (after the geometric distortions have been removed). Completion of a single stack of all of the 130,000 images will, therefore, require approximately 15,000 service units. This data set in of itself will provide one of the largest deep imaging surveys available to the astronomical community. The TeraGrid application for the stacking will provide a general framework for other time domain surveys such as QUEST and PanSTARRS.

3.5 Montage

Bruce Berriman, IPAC, California Institute of Technology,
Ewa Deelman, ISI, University of Southern California, and many others

The NASA Earth Science Technology Office Compute Technologies (ESTO CT) program has funded Montage, an image mosaic general image reprojection and mosaicing engine. Montage can be used both for simple individual science projects on a small machine or in parallel across large-scale computational Grids. Montage is being used by several NVO-related projects, in particular a collaboration between IRSA and USC Information Sciences Institute to use their Pegasus software system to schedule and manage complex data processing workflows on Grid (specifically the TeraGrid) resources.

This effort has been ongoing for some time but has started to ramp up toward full operational status, supplying Grid-based image mosaicing services to the astronomical community on demand. An NVO SIAP-compliant prototype to illustrate this capability (using the ten Terabyte 2MASS image set) has been developed at IPAC, with an anonymous web-form input. Users need to store their credentials in a MyProxy server in order to use the portal. The username and password used to store the credential can be used to login to the portal. Once logged in, the user can specify the grid resources that can be used and the location of Montage executables on those resources. The portal makes it easy for the users to use Pegasus to submit Montage computation on the grid. The main purpose of the portal is to shield the user from the complexity of various tools required to execute the Montage workflow over the TeraGrid. The portal takes care to authenticating with the grid resources and executing the Montage computations on the resources using the credentials retrieved from the MyProxy Server. The portal has been build using java server pages and java servlet technology and java COG Kit. It uses HTTPS for secure communication between the web browser and the server when the MyProxy username and password is being sent from the browser to the server.

3.6 Automated source extraction and cross matching

Andrew Connolly, Simon Krughoff, University of Pittsburgh

Upon coaddition of images source identification is a necessary next step in order that these data be scientifically useful. A web enabled source identification and cross-matching service (WESIX; http://nvo.phyast.pitt.edu/) has been developed to analyze imaging data and to cross-match these catalogs with existing multi-frequency data sets. A major component of the new services described here is to grid-enable this application to undertake source detection on the coadded data described above. WESIX is currently deployed as a webservice running on a single processor. One coadded SDSS image takes approximately 60 seconds of CPU time (on a 1.3 GHz Itanium2 processor) for source identification to a signal-to-noise level of 5 and a further 120 seconds of time for the cross-matching with existing catalogs. To source extract and cross-match a single realization of the 100 square degree coadded SDSS data will require approximately 6500 service units. A total of 10,000 service units is, therefore, requested to enable the development of the catalogs for the coadded SDSS data.

The coaddition, source detection and cross matching will be interconnected through the data grid services to provide an application that enables an arbitrary stack of images to be constructed (selected based on frequency or quality of data), sources detected and cross-matched. This, we believe, will be an example of one of the core products of the NVO (the ability for a user to select, analyze and return arbitrary sets of data). The interplay between data selected by a user, the resulting object catalogs and existing databases of catalogs will provide a detailed test of the grid operations and services.

3.7 Authenticated Science Gateways

Roy Williams, Conrad Steenberg, and Julian Bunn, California Institute of Technology

A growing thrust within the TeraGrid project is to broaden access through the so-called “HotGrid” mechanism, which gives a small amount of immediate TeraGrid resources to anyone who fills in a form. The HotGrid allocation would be used through a Science Gateway, of which several are being built for astronomers. Science Gateways are being built with the “Clarens” software that has been developed in the High Energy Physics community for analysis of CERN/LHC data. Clarens is a secure web server that is deployed on TeraGrid that allows graduated security access: an anonymous user can run small jobs, in stages of authentication to the power user who can log in and run arbitrary code. One of the intermediate stages is the “HotGrid” concept, whereby the user fills in a form, and gets in exchange a very restricted certificate credential, so she can use a specific Science Gateway in a specific time interval, with a restricted number of CPU-hours.

4 Infrastructure

4.1 GridShell enables Condor on TeraGrid

Many of the types of grid services required by the NVO operate as large numbers of serial jobs (e.g. sweeps through cosmological parameters for simulated CMB skies or the analysis of large numbers of individual galaxy images or spectra). In its initial implementation the TeraGrid did not offer a solution for running large numbers of serial jobs. PBS/LSF installations on TeraGrid compute resources are not optimized for large numbers of single-processor (or few processor) jobs. Some sites actively discourage such practices. Condor is an example of a scheduler that is designed to handle such a task, and although at the time of our initial proposal Condor was listed as an available scheduler on the TeraGrid website, it was not implemented in any workable way on TeraGrid machines. Most sites do not have Condor installed. On those that do, for example the NCSA TeraGrid cluster, Condor will only schedule jobs on nodes not currently running PBS tasks meaning that Condor jobs can sit in the queue for days or even weeks. Using Globus or Condor-G does not strictly solve our problem in that they are merely a front-end for the PBS or LSF scheduler at each TeraGrid site.

The need to schedule parameter sweeps caused us to work closely with Edward Walker at TACC in the development of GridShell, a grid-enabled shell scripting environment. GridShell uses Globus to spawn large multi-processor PBS/LSF jobs across the TeraGrid. These PBS/LSF jobs in turn start up Condor services (using the Condor GlideIn feature) on each processor within them, converting each PBS/LSF job into a Condor pool. The Condor pools then advertise themselves to a central Condor metascheduler to which the user can submit their 10,000 or 100,000 job parameter sweep. The result is that the user can submit a *single* Condor script that describes upwards of 100,000 independent jobs, and GridShell will automatically distribute these jobs across the TeraGrid. Users can use the built-in Condor automatic file transfer mechanism to stage data in and out of compute resources, or use simple GridShell scripting commands to transfer their files via GridFTP. Jeff Gardner demonstrated using GridShell to perform calculations for our *Cosmic Microwave Background Grid* project at Supercomputing 2004 (see http://www.tacc.utexas.edu/gridshell and http://www.psc.edu/~gardnerj/talks/SC04-Gridshell.ppt for more information). With the addition of GridShell to our arsenal, we have been able to make significant headway on our projects that require large numbers of independent computations: *Resolving Star Formation in Galaxies, Fitting Quasar Spectra,* and *The Cosmic Microwave Background Grid.* Thus, GridShell became an unforeseen yet vitally important milestone towards NVO's productivity on the TeraGrid.

5 Summary of Resources Requested

The bulk of the Service Units requested will be used by the 6 astronomy data processing applications we have described, and a further 87,000 SUs will be needed to begin integrating NVO services with the TeraGrid as part of our GIG Gateway partnership. All allocations are requested to be “R” type.

Application	Service Units
DPOSS hyperatlas	50000
2MASS hyperatlas	50000
N-Point Correlation Functions of Galaxies	120000
Resolving Star Formation in Galaxies	30000
Fitting Quasar Spectra	50000
Cosmic Microwave Background Grid	20000
Palomar-Quest processing	50000
Image coaddition grid services for the NVO	15000
Mosaicking gateway (Montage)	10000
Source extraction and cross matching gateway	10000
Authenticated Science Gateways	2000
TOTAL:	407000

6 References

- Brunner, R., Djorgovski, S.G., Szalay, A., *Virtual Observatories of the Future*, Astronomical Society of the Pacific Conference Series, Vol. 225, pp. 257-264, June 2000.
- Budavari, T., Szalay, A., Connolly, A., Csabai, I., 2003, AJ, in preparation
- Conti, A., Connolly, A., Hopkins, A., Budavari, T., Csabai,I., Szalay, A., Schmidt, S., Petrovich, N., 2003, AJ, 126, 2330.
- Deelman, E., Blythe, J., Gil, Y., Kesselman, C., *Workflow Management in GriPhyN*, in Grid Resource Management, J. Nabrzyski, J. Schopf, and J. Weglarz editors, Kluwer, 2003.
- Deelman, E., *et al*., *Mapping Abstract Complex Workflows onto Grid Environments*, Journal of Grid Computing, Vol.1, no. 1, 2003, pp. 25-39.
- Deelman, E., *et al*., *Grid-Based Galaxy Morphology Analysis for the National Virtual Observatory*, Proceedings of SC 2003.
- Gomez, P., et al, 2003, ApJ, 548, 210
- Jing, Y.P., & Boener, G., 2003, Astroph. J., submitted, astro-ph/0311585
- Lewis, I. et al. 2002, MNRAS, 334, 673
- Miller, Nichol, Genovese, & Wasserman, 2002, Astroph. J., 565, L67
- Miller, C.J., Nichol, R.C., Genovese, C., & Wasserman, L., 2002, Astroph. J., 565, L67
- Montage 2002, Science Use Cases, `http://montage.ipac.caltech.edu/docs/usecases.html`
- Peebles, P.J.E., 1980, *The Large-Scale Structure of the Universe*, Princeton, NJ: Princeton UP.
- Rajasekar, A., Wan, M., Moore, R., Kremenek, G., Guptil, T. *Data Grids, Collections, and Grid Bricks*,

Proc. 20th IEEE Symposium on Mass Storage Systems and Eleventh Goddard Conference on Mass Storage Systems and Technologies, San Diego, 2003.

- Rajasekar, A., Wan, M., Moore, R., Jagatheesan, A., Kremenek, G., *Real Experiences with Data Grids - Case studies in using the SRB*", Int. Symp.High-Performance Computer Architecture, Kyushu, Japan, 2002.
- Rajasekar, A., Moore, R., Ludäscher, B., Zaslavsky, I., *The Grid Adventures: SDSC's Storage Resource Broker and Web Services in Digital Library Applications*:, 4th Russian Conference on Digital Libraries, Dubna, Russia, 2002.
- Rose, J.A., Gaba, A.E., Caldwell, N., & Chaboyer, B. 2001, AJ, 121, 793
- Seljak, U., & Zaldarriaga, M., 1996, ApJ, 469, 437
- Smoot, G. F. et al., 1992, Astroph. J., 396, L1
- Scranton et al., 2005, in preparation
- Szalay, A. S., Connolly, A. J., Szokoly, G. P., *Simultaneous Multicolor Detection of Faint Galaxies in the Hubble Deep* Field, AJ, 117, 68, also astroph/9811086.
- Takada, M., & Jain, B., 2003, Mon. Not. Royal Ast. Soc., 344, 857
- Williams, R. D., Djorgovski, S. G., Feldmann, M. T., and Jacob, J. C., 2003a, *Hyperatlas, A New Framework for Image Federation*, http://arxiv.org/abs/astro-ph/0312195
- Williams, R. D., Djorgovski, S. G., Feldmann, M. T., and Jacob, J. C., 2003b, *Atlasmaker, A Grid-based Implementation of the Hyperatlas*, http://arxiv.org/abs/astro-ph/0312196
- Williams, R. D., Feldmann, M. T., 2003c, *Atlasmaker, Software to build Atlases of Federated Imagery*, 2003, http://www.cacr.caltech.edu/projects/nvo/atlasmaker/
- Williams, R. D., et. al., *Multi-Wavelength Image Space: Another Grid-Enabled Science*, Concurrency & Computation, vol 15 (2003), pp 539-549.
- Williams R. D., *Grids and the Virtual Observatory*, pp 837-858 in *Grid Computing: Making The Global Infrastructure*, eds. F. Berman, A. J.G. Hey, and G. C. Fox, Wiley, 2003.