

Montage



A Comparison of Two Methods for Building Astronomical Image Mosaics on a Grid

<http://montage.ipac.caltech.edu/>
dsk@cct.lsu.edu





Montage

- An astronomical image mosaic service for the National Virtual Observatory
 - Project web site - <http://montage.ipac.caltech.edu/>
 - Core team at JPL (NASA's Jet Propulsion Laboratory) and Caltech (IPAC - Infrared Processing and Analysis Center, CACR - Center for Advance Computing Research)
 - Grid architecture developed in collaboration with ISI - Information Sciences Institute
- | | |
|--------------------------------------|----------------------------------|
| • Attila Bergou - <i>JPL</i> | • Carl Kesselman - <i>ISI</i> |
| • Nathaniel Anagnostou - <i>IPAC</i> | • Anastasia Laity - <i>IPAC</i> |
| • Bruce Berriman - <i>IPAC</i> | • Thomas Prince - <i>Caltech</i> |
| • Ewa Deelman - <i>ISI</i> | • Gurmeet Singh - <i>ISI</i> |
| • John Good - <i>IPAC</i> | • Mei-Hui Su - <i>ISI</i> |
| • Joseph C. Jacob - <i>JPL</i> | • Roy Williams - <i>CACR</i> |
| • Daniel S. Katz - <i>JPL</i> | |



What is Montage?

- **Delivers custom, science grade image mosaics**
 - An image mosaic is a combination of many images containing individual pixel data so that they appear to be a single image from a single telescope or spacecraft
 - User specifies projection, coordinates, spatial sampling, mosaic size, image rotation
 - Preserve astrometry (to 0.1 pixels) & flux (to 0.1%)
- **Modular, portable “toolbox” design**
 - Loosely-coupled engines for image reprojection, background rectification, co-addition
 - Control testing and maintenance costs
 - Flexibility; e.g custom background algorithm; use as a reprojection and co-registration engine
 - Each engine is an executable compiled from ANSI C
- **Public service will be deployed on the *Teragrid***
 - Order mosaics through web portal




David Hockney Pearl Blossom Highway 1986



Use of Montage

- **Scientific Use Cases**
 - Structures in the sky are usually larger than individual images
 - High signal-to-noise images for studies of faint sources
 - Multiwavelength image federation
 - Images at different wavelengths have differing parameters (coordinates, projections, spatial samplings, . . .)
 - Place multiwavelength images on common set of image parameters to support faint source extraction
- **Montage supports observation planning and generation of science and E/PO products in the projects listed below.**
 - Spitzer Legacy Teams
 - SWIRE: Spitzer Wide Area Infrared Experiment
 - GLIMPSE: Galactic Legacy Infrared Mid-Plane Survey Extraordinaire
 - c2d: “From Molecular Cores to Planet-forming Disks”
 - Spitzer Space Telescope Outreach Office
 - IRSA (NASA’s InfraRed Science Archive)
 - 2 Micron All Sky Survey (2MASS)
 - COSMOS (a Hubble Treasury Program to study the distribution of galaxies in the distant Universe)
 - IPHAS: The INT/WFC Photometric H-alpha Survey of the Northern Galactic Plane
 - NSF National Virtual Observatory (NVO) Atlasmaker project
 - UK Astrogrid Virtual Observatory




Montage Use By IPHAS: The INT/WFC Photometric H-alpha Survey of the Northern Galactic Plane

Nebulosity in vicinity of HII region, IC 1396B,
in Cepheus



Crescent Nebula
NGC 6888



Study extreme
phases of stellar
evolution that
involve very large
mass loss

Supernova remnant
S147




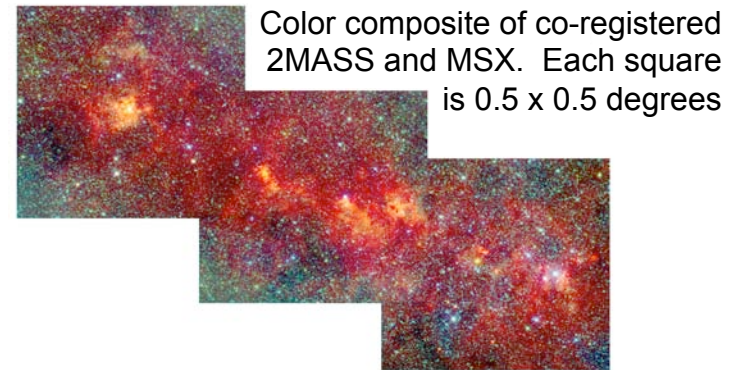
Montage Use by “Galactic Legacy Infrared Mid-Plane Survey Extraordinaire” Spitzer Legacy Team (GLIMPSE)



GLIMPSE is making the first global survey of star formation in the Galaxy

Applications of Montage:

- Federation of 2MASS J, H, K and MSX 8 μm images to act as quality assurance and validation products
- Generation of the primary data products: image mosaics at four infrared wavelengths
- Data deliveries every three months, starting in February 2005 at <http://data.spitzer.caltech.edu/popular/glimpse/>





Montage Use by “Spitzer Wide-area Infrared Extragalactic Survey” Legacy Science Program (SWIRE)



SWIRE uses Montage for:

- Building sky simulations for use in mission planning
- Generation of its primary data products:
co-registered multi-wavelength image mosaics
covering several square degrees
 - Will be used for extraction of new populations
of high-redshift galaxies
- Visit <http://data.spitzer.caltech.edu/popular/swire/>




Right: Spitzer IRAC 3 channel mosaic (3.6 μm in green, 4.5 μm in red, and i-band optical in blue); high redshift non-stellar objects are visible in the full resolution view (yellow box).



Montage Use by Hubble Cosmic Evolution Treasury Program (COSMOS)

- 72,000 x 72,000 pixel mosaic by Montage
- Comprised of 51 I-band images measured with the Hubble Space Telescope (HST) Advanced Camera for Surveys (ACS)
- Supports science goals of COSMOS
 - Study of structure of high-redshift universe in a 4 square degree area





Montage Use by Spitzer E/PO Group

Visualization of full-sky datasets


- Visualization of full-sky image surveys by end-users require novel projections that are atypical of standard astronomical schemes
- Montage supports reprojection of standard datasets into the projections needed by E/PO - e.g.; development of all-sky datasets in a format easily used for immersive viewers, backdrops for realistic 3D animations, and even maps/globes that can be distributed online

- Two examples shown from the E/PO page at:

http://coolcosmos.ipac.caltech.edu/resources/informal_education/allsky/



Panoramic view of the sky as seen by 2MASS.



100 μm sky; aggregation of COBE and IRAS maps (Schlegel, Finkbeiner and Davis, 1998). Covers 360 x 180 degrees in CAR projection.



First Public Release of Montage

- Version 1 emphasized accuracy in photometry and astrometry
 - Images processed serially
 - Tested and validated on 2MASS 2IDR images on Red Hat Linux 8.0 (Kernel release 2.4.18-14) on a 32-bit processor
 - Tested on 10 WCS projections with mosaics smaller than 2 x 2 degrees and coordinate transformations Equ J2000 to Galactic and Ecliptic
- Extensively tested
 - 2,595 test cases executed
 - 119 defects reported and 116 corrected
 - 3 remaining defects renamed caveats
 - Corrected in Montage v2 release




Later Public Releases of Montage

- **Second release: Montage version 2.2**
 - More efficient reprojection algorithm: up to 30x speedup
 - Improved memory efficiency: capable of building larger mosaics
 - Enabled for parallel computation with MPI
 - Enabled for processing on TeraGrid using standard grid tools (TRL 7)
- **Third release: Montage version 3.0 (currently in beta)**
 - Data access modules
 - Tiled output
 - Outreach tool to build multi-band jpeg images
 - Other improvements in processing speed and accuracy
 - Bug fixes
- **Code and User's Guide available for download at <http://montage.ipac.caltech.edu/>**




Montage v1.7 Reprojection: mProject module

Arbitrary Input Image

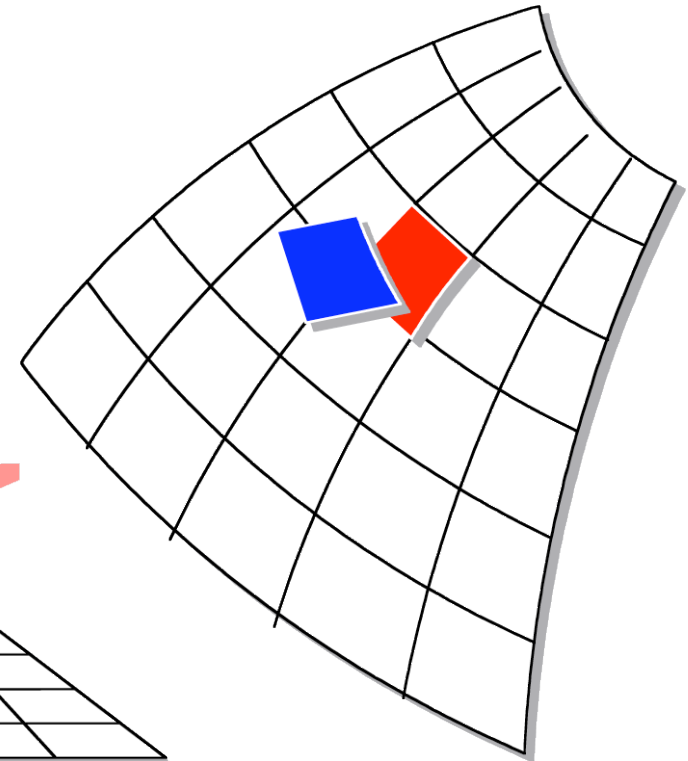


Central to the algorithm is accurate calculation of the area of spherical polygon intersection between two pixels (assumes great circle segments are adequate between pixel vertices)



Input pixels projected on celestial sphere


Output pixels projected on celestial sphere



```
SIMPLE =           T /
BITPIX=          -64 /
NAXIS  =           2 /
NAXIS1=          3000 /
NAXIS2=          3000 /
CDELTA1=         - 3.333333E-4 /
CDELTA2=         - 3.333333E-4 /
CRPIX1=           1500.5 /
CRPIX2=           1500.5 /
CTYPE1='RA---TAN'
CTYPE2='DEC--TAN'
CRVAL1=           265.91334 /
CRVAL2=           -29.35778 /
CROTA2=            0. /
END
```

FITS header defines output projection

Reprojected Image







Montage v2.2 Reprojection: mProjectPP module

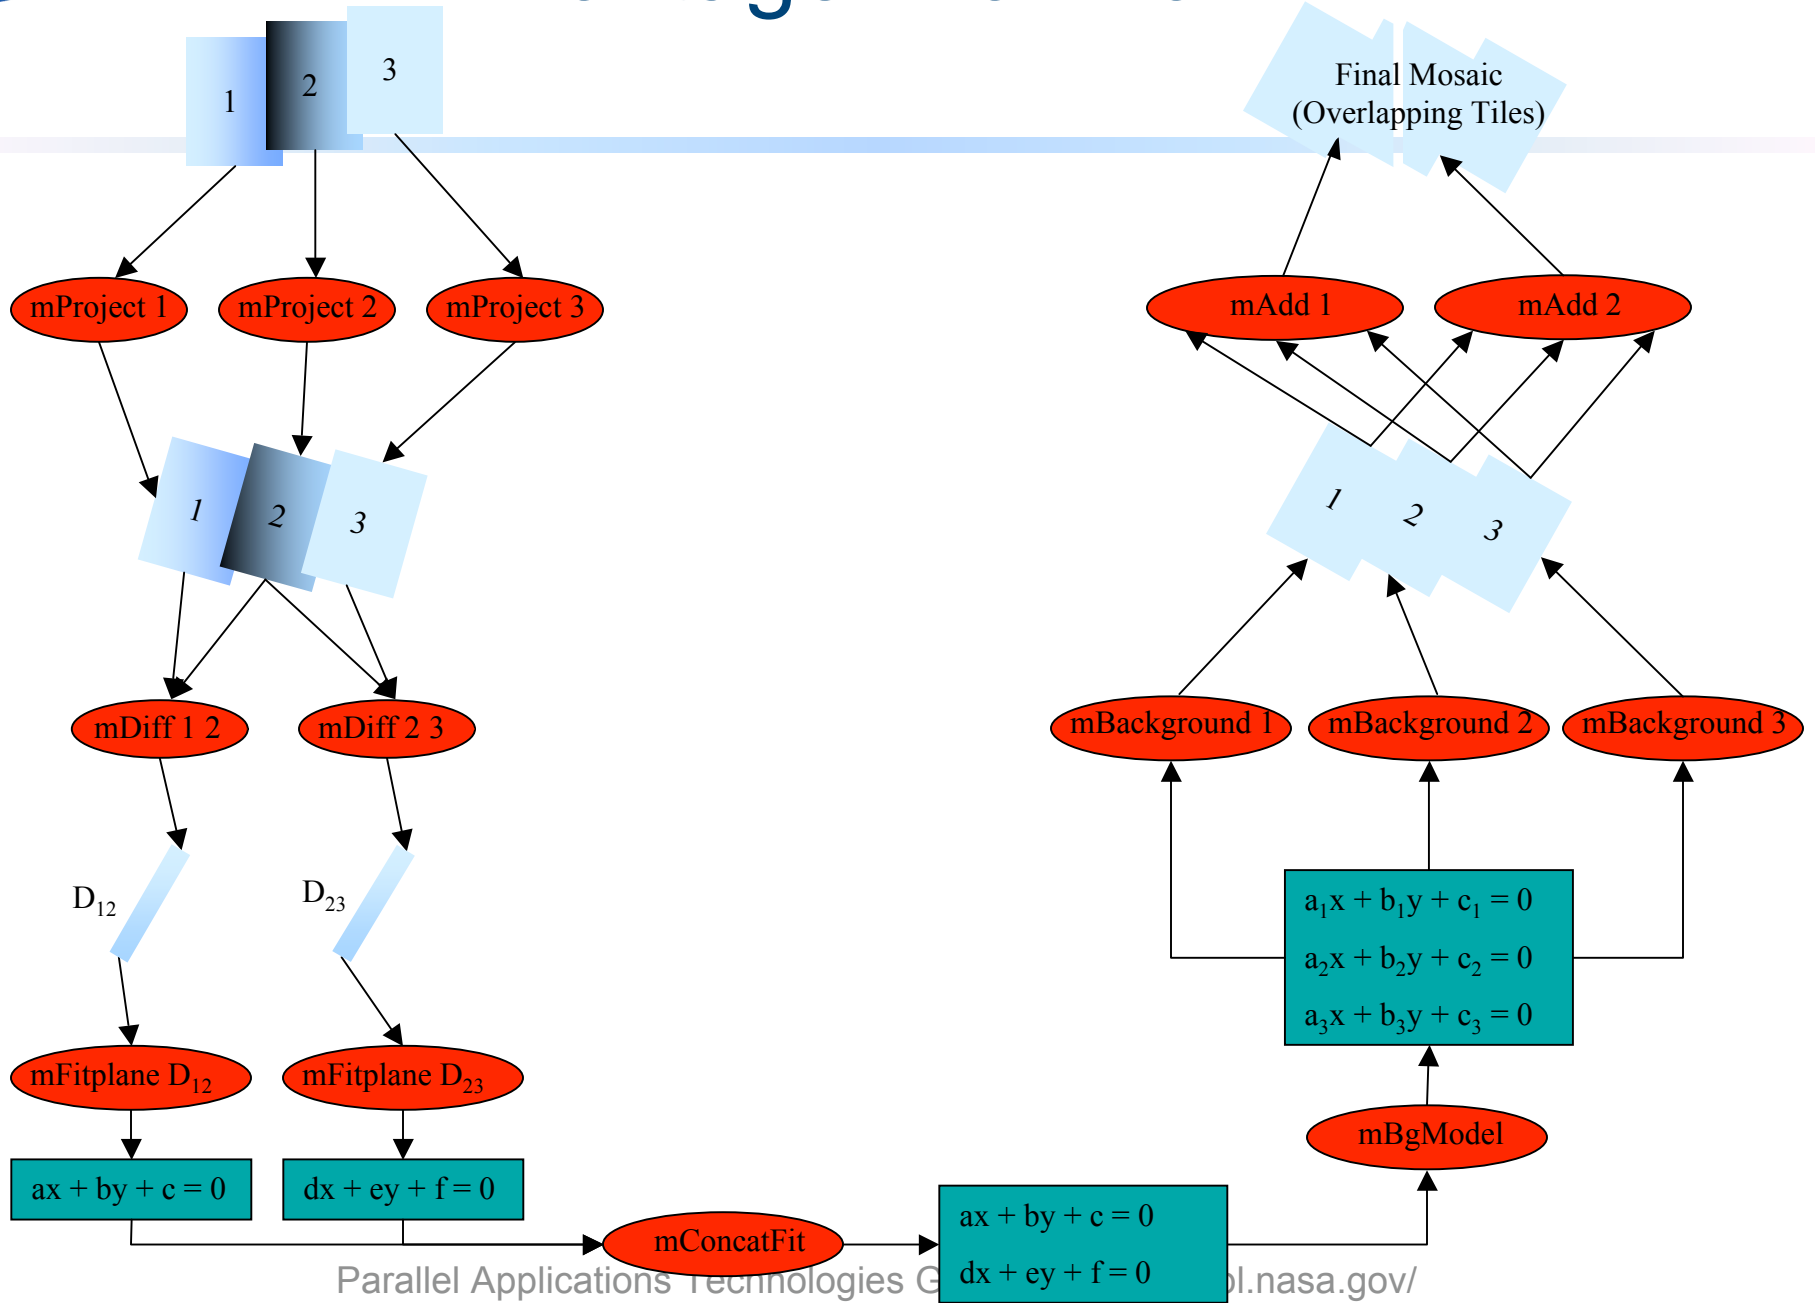


- Transform directly from input pixels to output pixels
 - Approach developed by Spitzer for tangent plane projections
 - **Performance improvement in reprojection by x 30**
- Montage version 2.2 includes a module, mTANHdr, to compute “distorted” gnomonic projections to make this approach more general
 - Allows the Spitzer algorithm to be used for other projections (in certain cases)
 - For “typical” size images, pixel locations distorted by small distance relative to image projection plane
 - Not applicable to wide area regions





Montage Workflow






Montage Parallel Performance

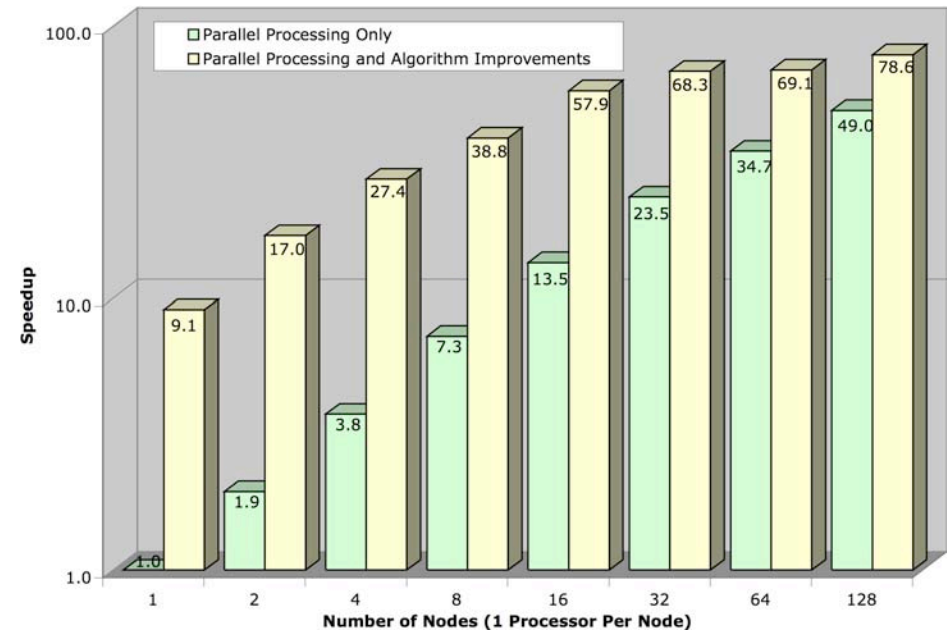
6 x 6 degrees 2MASS mosaic of M16 at 1" sampling



Montage_v2.1 Execution Times on NCSA TeraGrid Cluster



Montage_v2.1 Speedup on NCSA TeraGrid Cluster



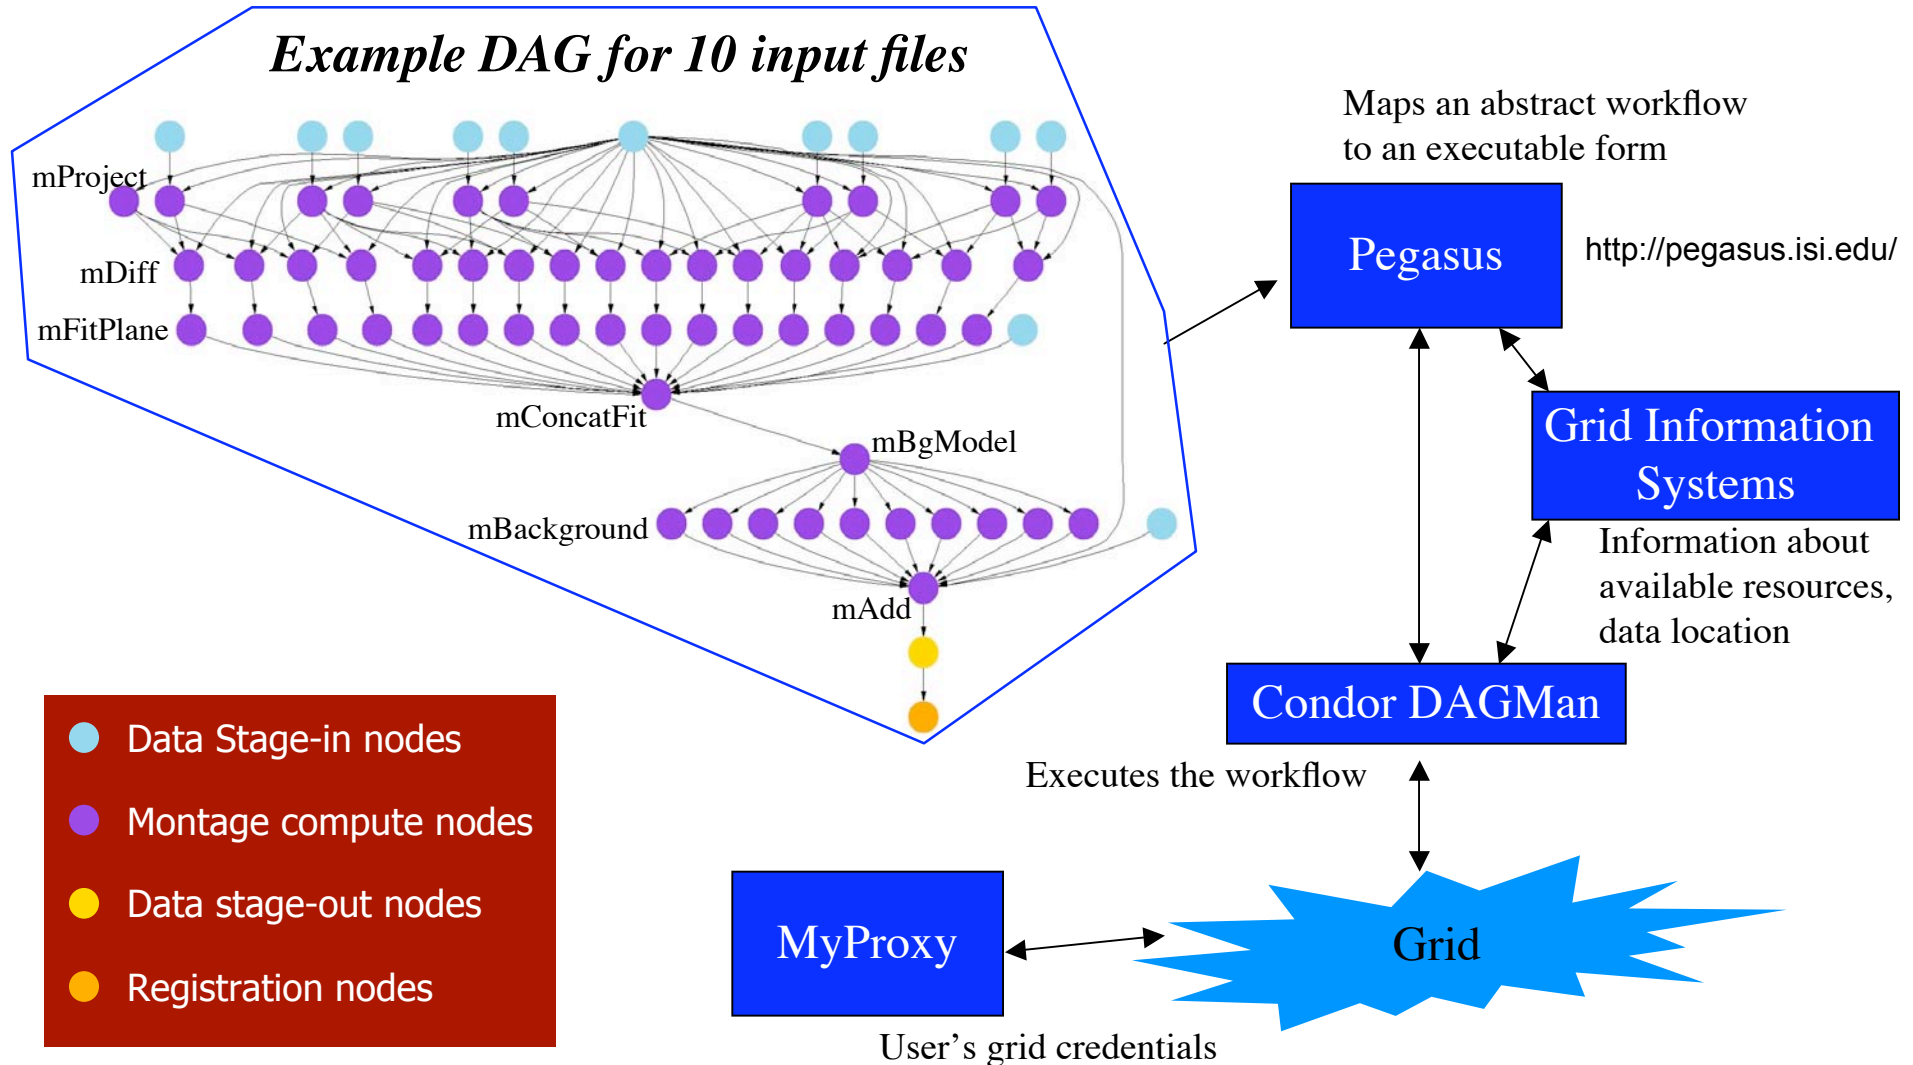


Montage on the Grid

- “Grid” is an abstraction
 - Array of processors, grid of clusters, ...
- Use a methodology for running on any “grid environment”
 - Exploit Montage’s modular design in an approach applicable to any grid environment
 - Describe flow of data and processing (in a Directed Acyclic Graph - DAG), including:
 - Which data are needed by which part of the job
 - What is to be run and when
 - Use standard grid tools to exploit the parallelization inherent in the Montage design
- Build an architecture for ordering a mosaic through a web portal
 - Request can be processed on a grid
 - Our prototype uses the Distributed Terascale Facility (TeraGrid)
- This is just one example of how Montage could run on a grid

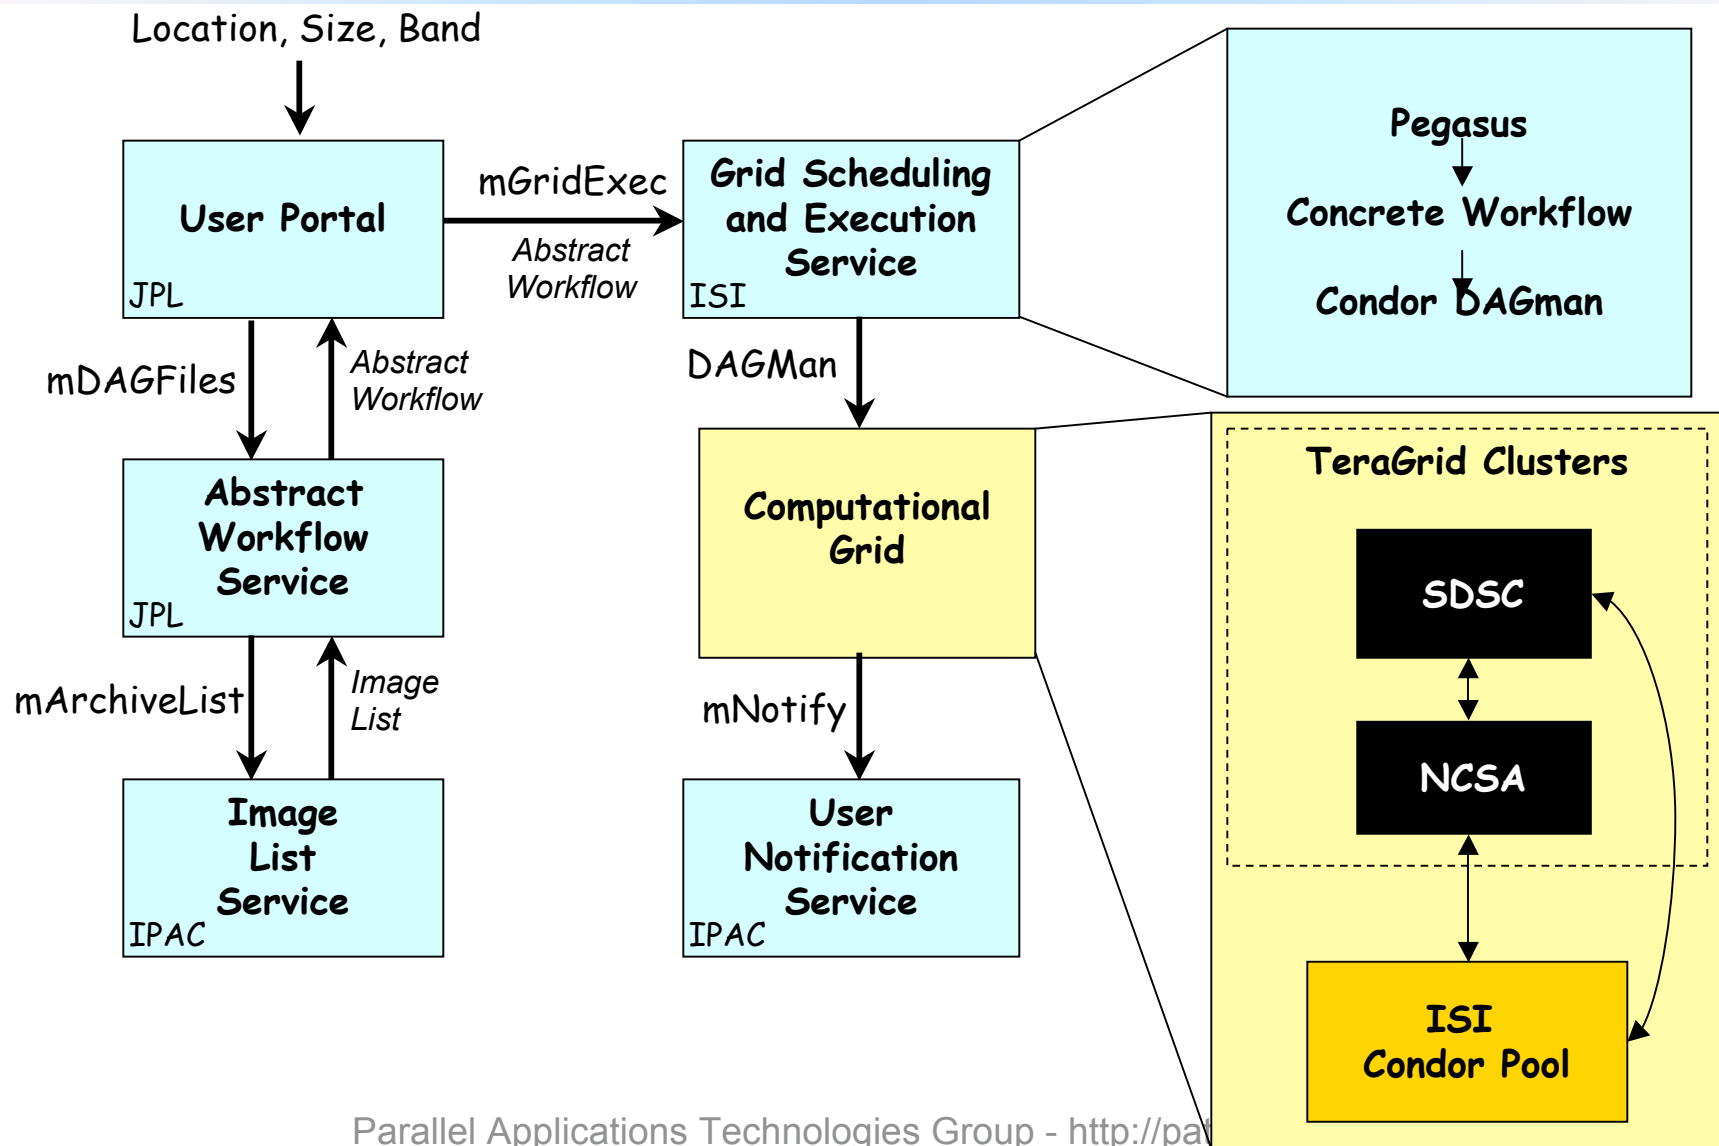


Montage on the Grid Using Pegasus (Planning for Execution on Grids)





Montage TeraGrid Portal

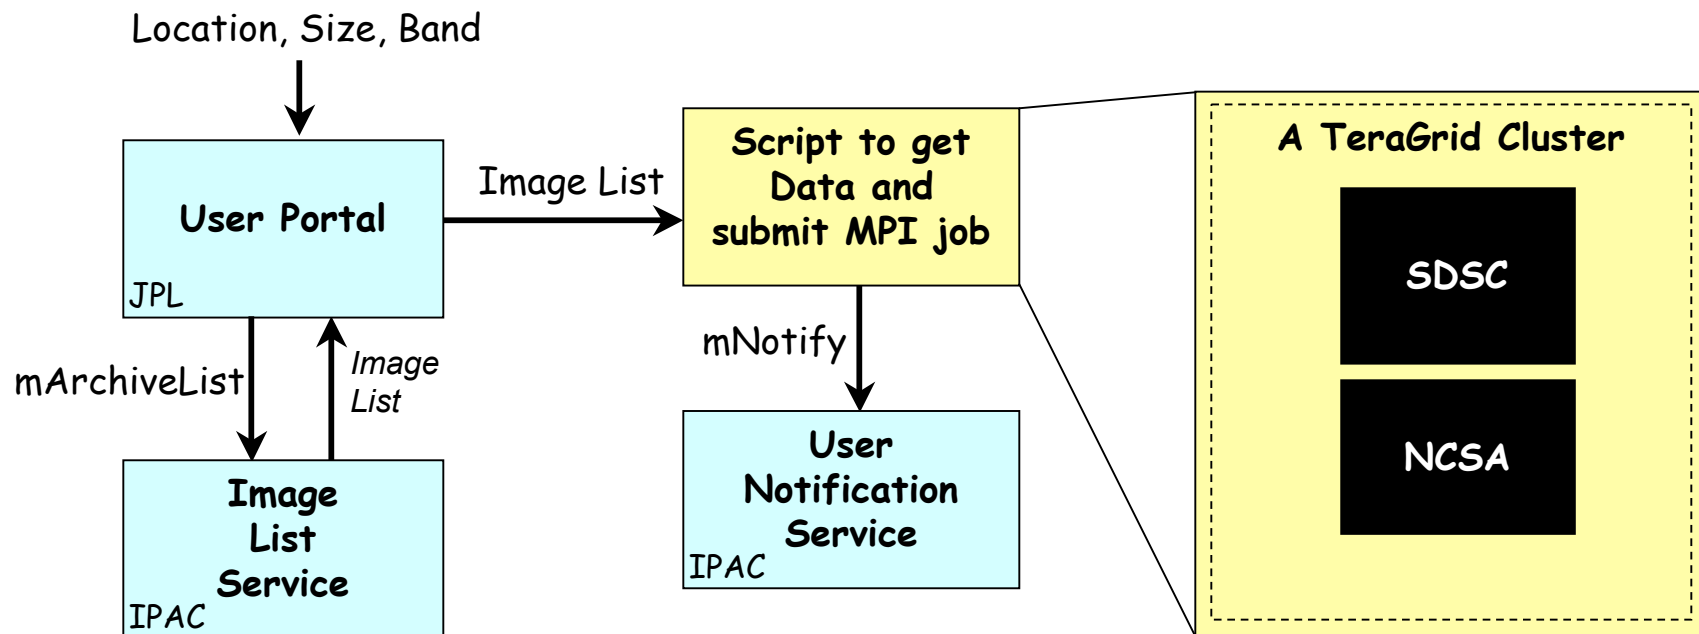




Alternative Montage

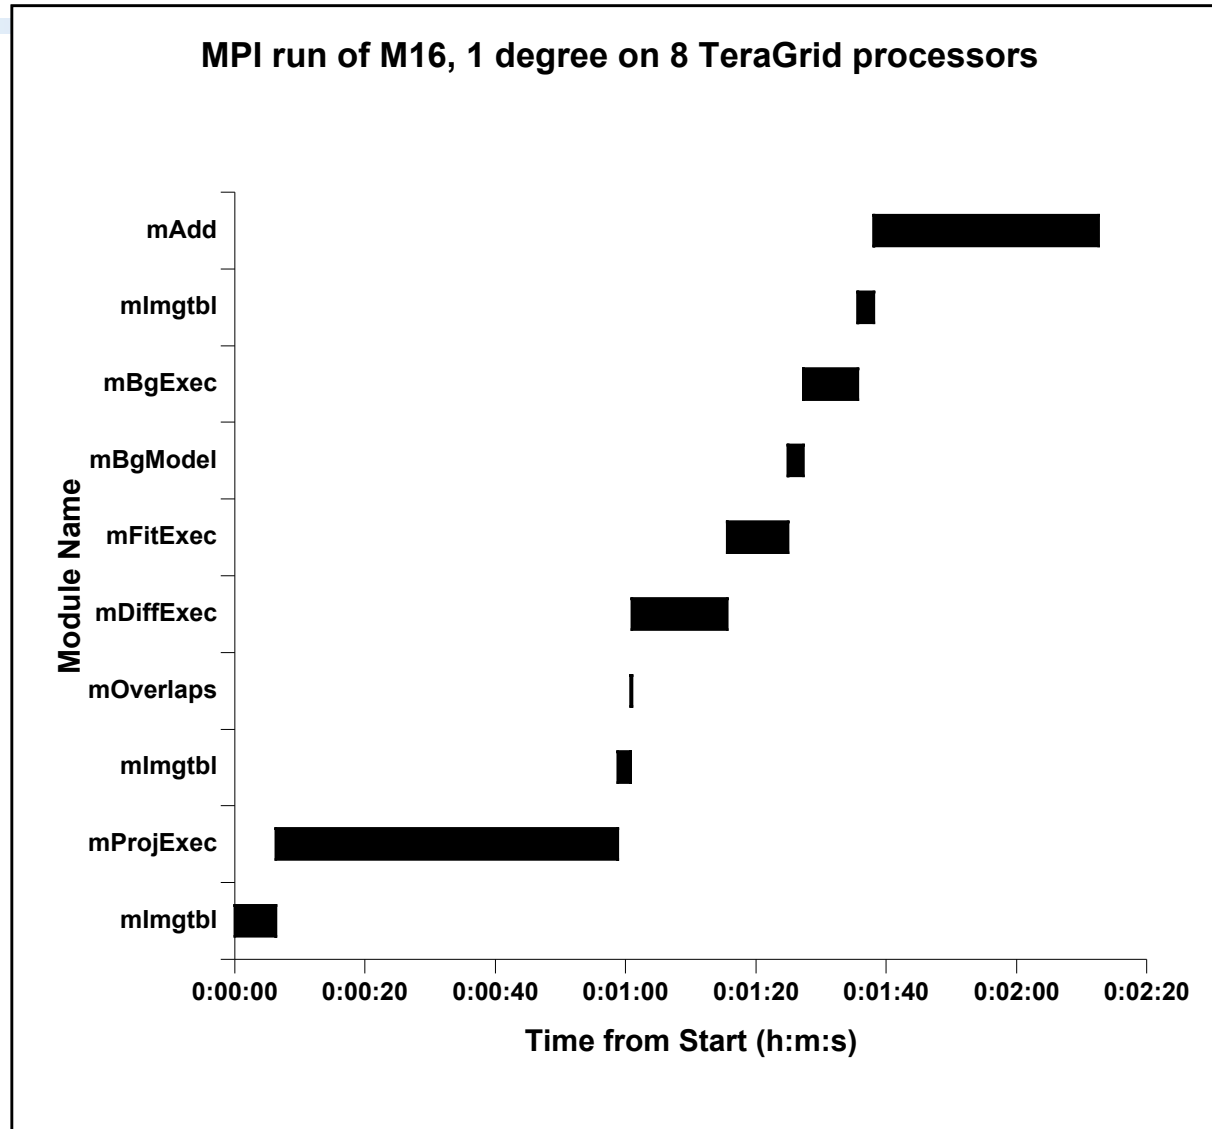


TeraGrid Portal (vaporPortal)



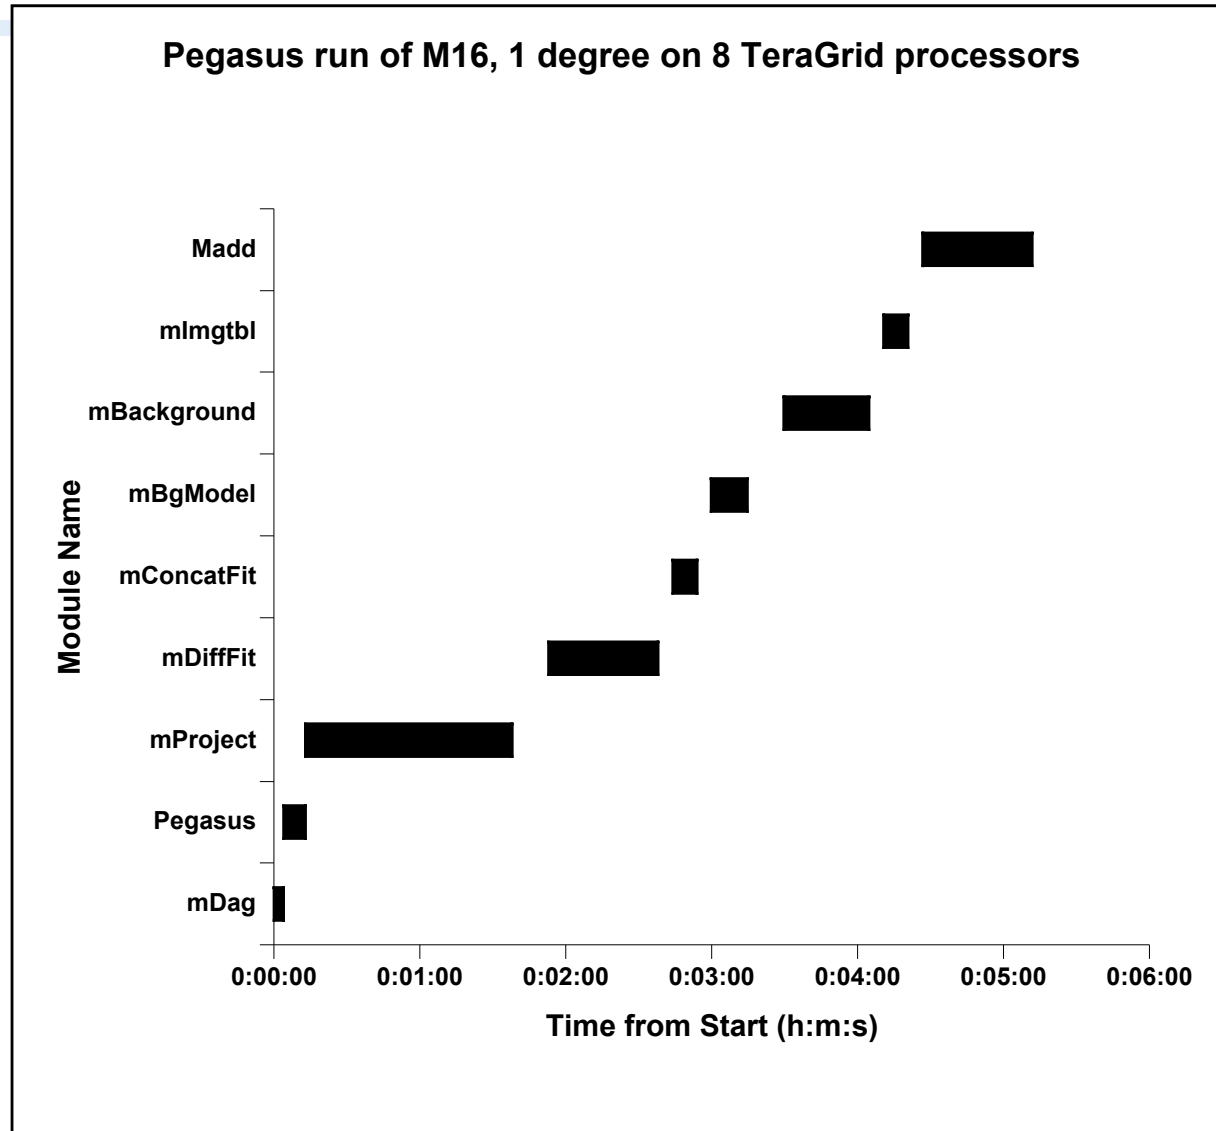


Montage Performance on Small Problem





Montage Performance on Small Problem





Timing Discussion

- Both MPI and Pegasus timings ignore time to start job (queuing delay)
 - MPI - script is placed in queue that calls both serial and parallel tasks, in sequence, on the nodes that are obtained from the queue
 - Pegasus - Condor Glide-in is used to allow single processor jobs to work in parallel in a the pool that is obtained from the queue
 - For efficiency, jobs are clustered and each cluster is submitted to the pool
 - Condor overhead for each item submitted is between 1 and 5 seconds
- Tasks are different
 - MPI - mImgtbl, mProjExec, mImgtbl, mOverlaps, mDiffExec, mFitExec, mBgModel, mBgExec, mImgtbl, mAdd
 - *Exec tasks are parallel tasks, others are sequential
 - Flow is dynamic, based on resulting files from previous stages
 - Pegasus - mDag, Pegasus, mProject(s), mDiffFit(s), mConcatFit, mBgModel, mBackground(s), mImgtbl, mAdd
 - *(s) tasks are multiple, clustered by Pegasus/Condor
 - Flow is fixed, based on output of mDag

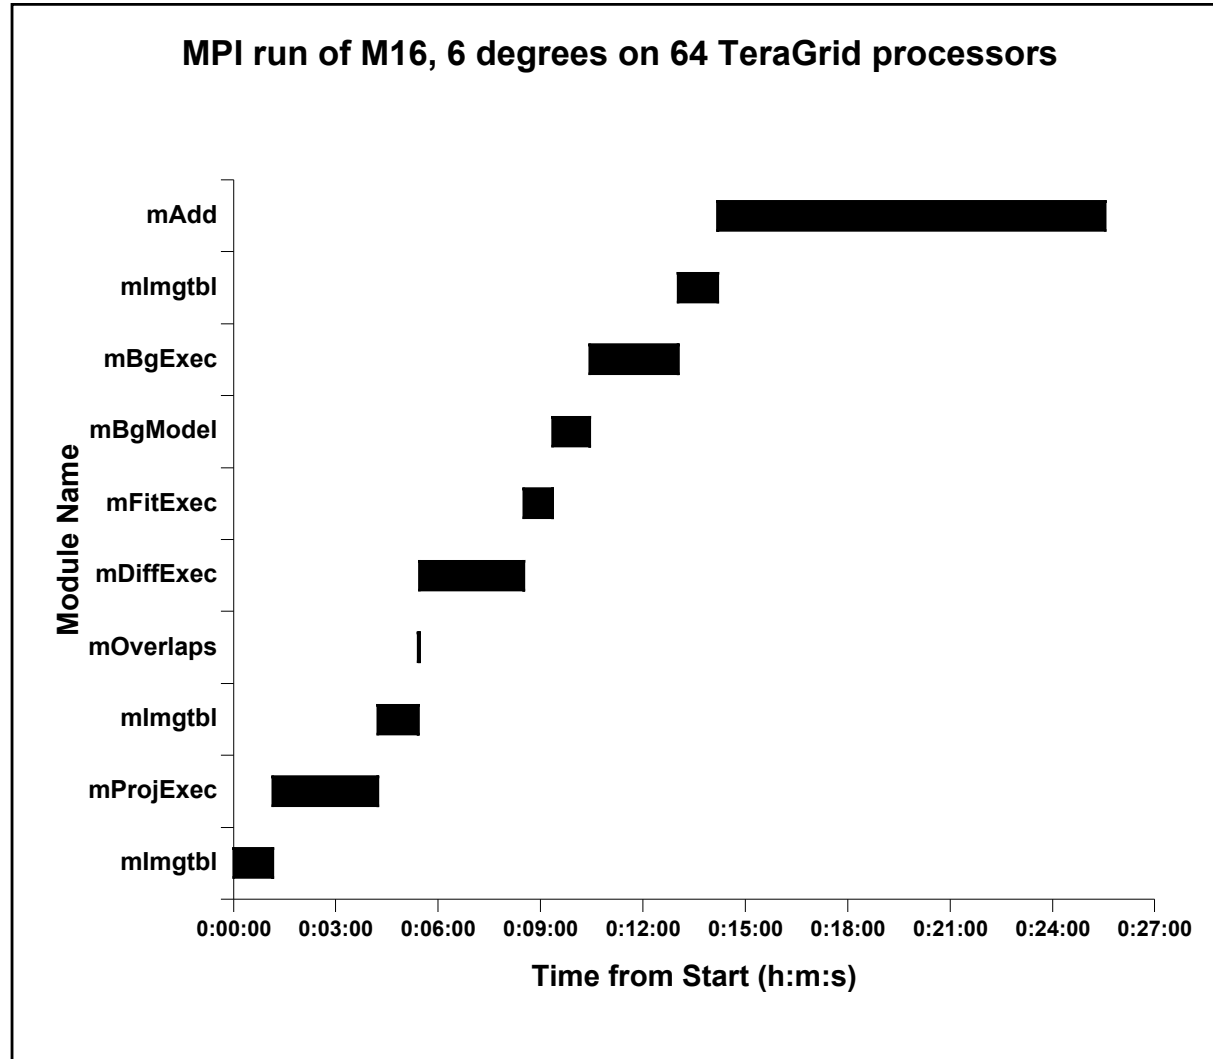


More Timing Discussion

- Gaps between tasks
 - MPI - no gaps, other than MPI job shutdown and startup
 - Pegasus - gaps of up to 5 seconds from Condor/DAGMan
- Accuracy
 - I/O dominates many of the computational tasks
 - On the TeraGrid in a multi-user environment, none of this is very precise
- Overall
 - MPI - job finishes in 00:02:12
 - Pegasus - job finishes in 00:05:12

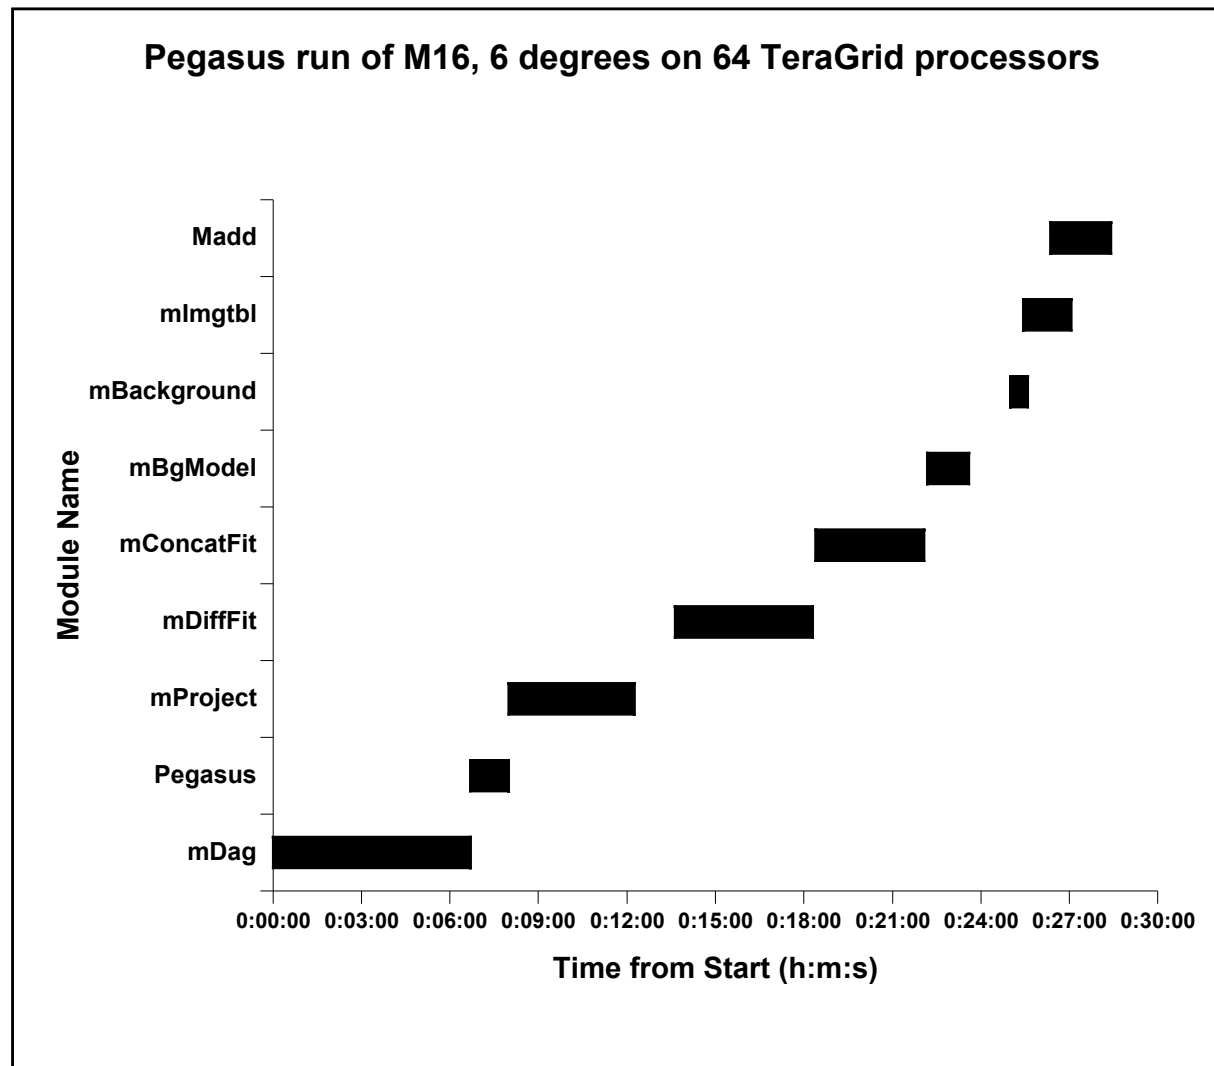


Montage Performance on Large Problem





Montage Performance on Large Problem





Timing Discussion

- Most things are the same as for the small job
- Gaps between tasks are less important, as tasks are longer
- Accuracy is more of a question, as the parallel file system is being hit harder
- Overall
 - MPI - job finishes in 00:25:33
 - Pegasus - job finishes in 00:28:25



Summary

- Montage is a custom astronomical image mosaicking service that emphasizes astrometric and photometric accuracy
- Final public release, Montage version 3b, available for download at the Montage website: <http://montage.ipac.caltech.edu/>
- A prototype Montage service has been deployed on the TeraGrid
 - It ties together distributed services at JPL, Caltech IPAC, and ISI
- MPI version of Montage:
 - Best performance
 - Requires a set of processors with a shared file system
- Pegasus (<http://pegasus.isi.edu/>) / DAGman version of Montage:
 - Almost equivalent performance for large problems
 - Built-in fault tolerance
 - Can use multiple sets of processors
- An open Montage service will be deployed on the TeraGrid by 9/05
 - Will allow MPI/Pegasus/Serial processing