ADASS XIV in Pasadena: Dealing with Petabytes of Data

Doug Mink

Telescope Data Center
Smithsonian Astrophysical Observatory

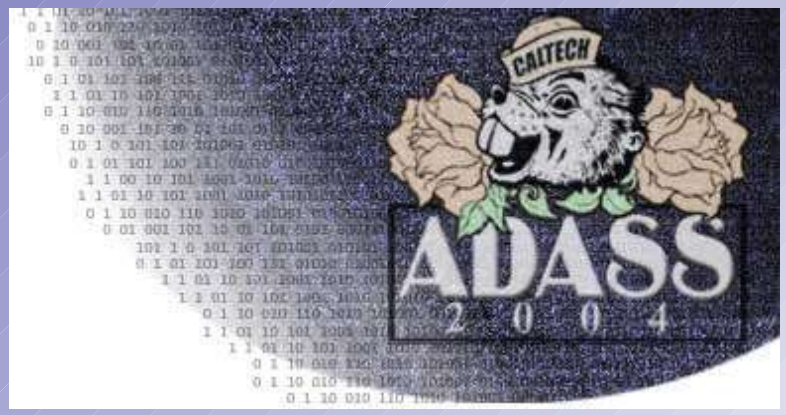

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

http://tdc-www.harvard.edu

Astronomical Data Analysis Software and Systems

- FITS standard evolves
- JPL adds planetary data
- New bigger data sets
- More computer power needed
  - -> Matt's Beowulf Cluster paper
- Virtual Observatory Update
- Pipelines
  - -> Hectospec Pipeline paper
- My To Do List

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

FITS

- Don Wells, co-inventor of FITS, retired this year
- FITS MIME types almost approved by IESG
- Spectral coordinates paper (III) almost approved
- Controversy over distortion->paper IV not quite ready
- 64-bit integers will be added to standard
- Keyword "conventions" need to be characterized
- FITS and the Virtual Observatory overlap but need to connect

JPL

- Mars Rover Data system talk
- Mars Rover Science talk at banquet

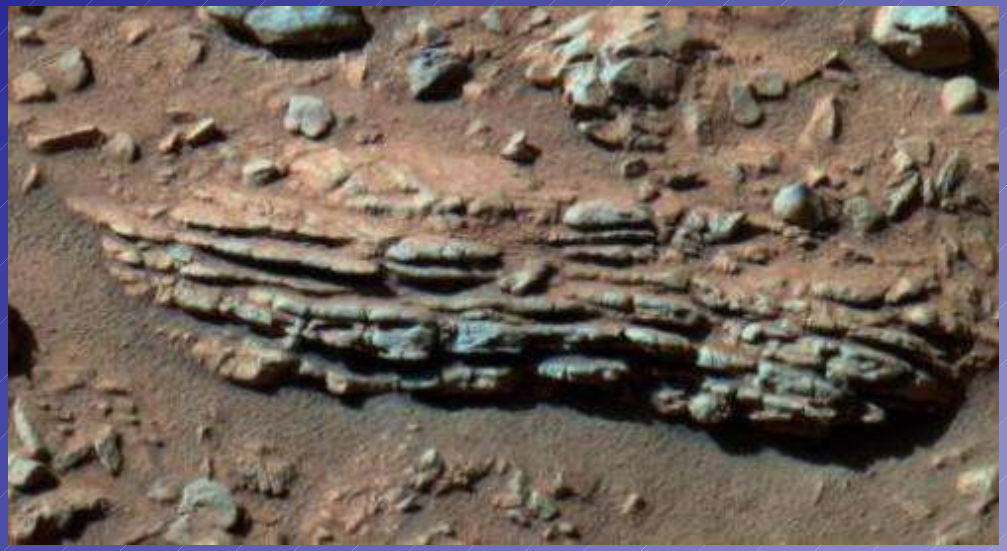

"Sedimentary" rock found by Spirit

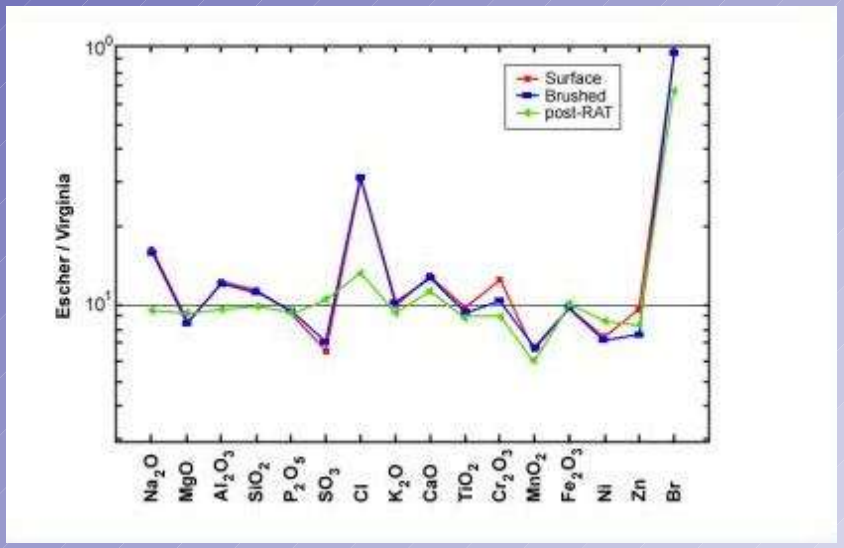

Alpha \-scattering spectrometry of rock possibly altered by water from Opportunity

Telescope Data Center
SMITHSONIAN ASTROPHYSICAL OBSERVATORY

AVO
ASTROPHYSICAL VIRTUAL OBSERVATORY

Astronomy in the XXI century

Radical changes are needed!

- Huge surveys: 100M sources at <3k spectra/night ⇒ >**100 yr**!
- Ever fainter sources: surpassed the identification limits of 8 - 10m telescopes ($R_{mag} \approx 25$)
- Huge data collections: downloading Sloan Digital Sky Survey (SDSS) DR3 (~ 1/2 of total) images (6 Tb) ⇒ ~ **2.3 months** at 1 Mb/s (ESO's speed); catalogs (2.3 Tb) ⇒ ~ **1 month**. On DVDs ⇒ ~ **1,300** of them. And analysis?? (similar size for MACHO, 2MASS etc …)
- Ever increasing amount of data (~ 1 Tb/night)

Oct. 26, 2004 P. Padovani, ADASS XIV 2

Some Big Data Sets

- 2MASS: 14 Terabytes
- PanSTARRS: 4000 Gpix images/night -> 2.5Pb/year
- VST: 150 Gbyte/night (online winter 2005)
- GRIST: 25Gbyte/night, 50/year on Palomar 48-inch

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

Number Crunching

Doug Tody (NRAO):

- Data explosion will be in archives, not necessarily near astronom
- Software is aging, 10-20 years old, with old design/usage mode
- New paradigm: distributed, multiwavelength, scalable
- Virtual Observatory provides
  - Uniform data interface
  - Data access
  - Computation facilities (GRID)
  - Applications?

CfA Cluster Requirements

- Cost to be spread over the several groups participating.
- Limited personnel resources for hardware and software installation and Management.
- Compatibility with existing environment of Linux and SPARC Solaris platforms.
- 64-bit addressing for large data sets in a single process.
- Resource allocation for each group based on size of contribution (i.e. number of machines purchased) to the cluster, however "public" access to available resources should be provided to the CfA community at large.

CfA Cluster Hardware

AMD Opteron based SunFire V20z nodes from Sun Microsystems

- Dual processor, 64-bit capable, 1U rack-mountable servers
- Runs 64-bit Linux with compatibility for 32-bit applications
- Existing Sun - CfA relationship made Educational Promotional discount possible.

Sun Compute Grid Rack System

- Pre-assembled 32 node cluster with gigabit network switch all in one cabinet.
- Just add power (a lot!) and a gigabit connection to the public network

CfA Cluster Software

ROCKS Cluster Distribution from http://rocksclusters.org

- Based on RedHat Enterprise Linux Release 3
- Supports Opteron Hardware in 64-bit mode
- "Rolls" of supporting software for HPC computing (MPI, PVFS, etc.), Globus Grid toolkit, and Sun Grid Engine batch queueing system
- Simple install of frontend node, followed by automatic Kickstart install of compute nodes via PXE boot

Sun Grid Engine Enterprise Edition

- Extemely flexible resource queueing system
- Included in ROCKS distribution
- Queues based on projects, Unix groups, or individual users
- Job suspension and checkpointing which allows jobs to be moved from node to node without losing previous work

CfA Cluster Problems

- **Weight**: Cabinet is 1800 pounds (815 kg), machine room tiles need to be reinforced.
- **Size:** Cabinet needed professional movers to go from loading dock to machine room.
- **Heat:** AC in machine room may need boosting
- **Power:** 60 A 220 V power source needed (4 big plugs [NEMA L6-20P])
- **User accounts:** ROCKS uses 411 insted of NIS for user account (and other) information which doesn't integrate with our existing environment. Seperate accounts (with matching UIDs and GIDs) are necessary.
- **Software installs:** ROCKS can install any RPM based software on all nodes, but non-RPM based software (e.g. IRAF and IDL) is problematic. An NFS mounted directory on an existing Network Appliance disk server was set up for this. It is a sub-optimal solution because of the network traffic.

CfA Cluster Resources

- Sun Fire V20z
  http://www.sun.com/servers/entry/v20z/

  - Sun Compute Grid Rack System
    http://www.sun.com/servers/computegrid/

  - ROCKS Open Source High Performance Linux Cluster Solution
    http://rocksclusters.org

  - Sun Grid Engine
    http://gridengine.sunsource.net/

Registry Model

Full Searchable Registry

VO Projects

Local Publishing Registry

Full Searchable Registry

Data Centers

Local Publishing Registry

Local Searchable Registry

Specialized Portals & Services

Registry Model

harvest
(pull)

Full Searchable Registry

VO Projects

Full Searchable Registry

Local Publishing Registry

Data Centers

Local Publishing Registry

Local Searchable Registry

Specialized Portals & Services

Registry Model

Full Searchable Registry

VO Projects

harvest (pull)

replicate

Local Publishing Registry

Full Searchable Registry

Data Centers

Local Publishing Registry

Local Searchable Registry

Specialized Portals & Services

Registry Model

Full Searchable Registry

VO Projects

harvest (pull)

replicate

Local Publishing Registry

Full Searchable Registry

Data Centers

selective harvesting

Local Publishing Registry

Local Searchable Registry

Specialized Portals & Services

Registry Model

Full Searchable Registry

VO Projects

Full Searchable Registry

Local Publishing Registry

Data Centers

Local Publishing Registry

search queries

Client Applications

Local Searchable Registry

Specialized Portals & Services

Registry Model

Full Searchable Registry

VO Projects

Local Publishing Registry

Full Searchable Registry

Data Centers

Local Publishing Registry

search queries

Local Searchable Registry

Client Applications

Specialized Portals & Services

Registries in Use: DataScope

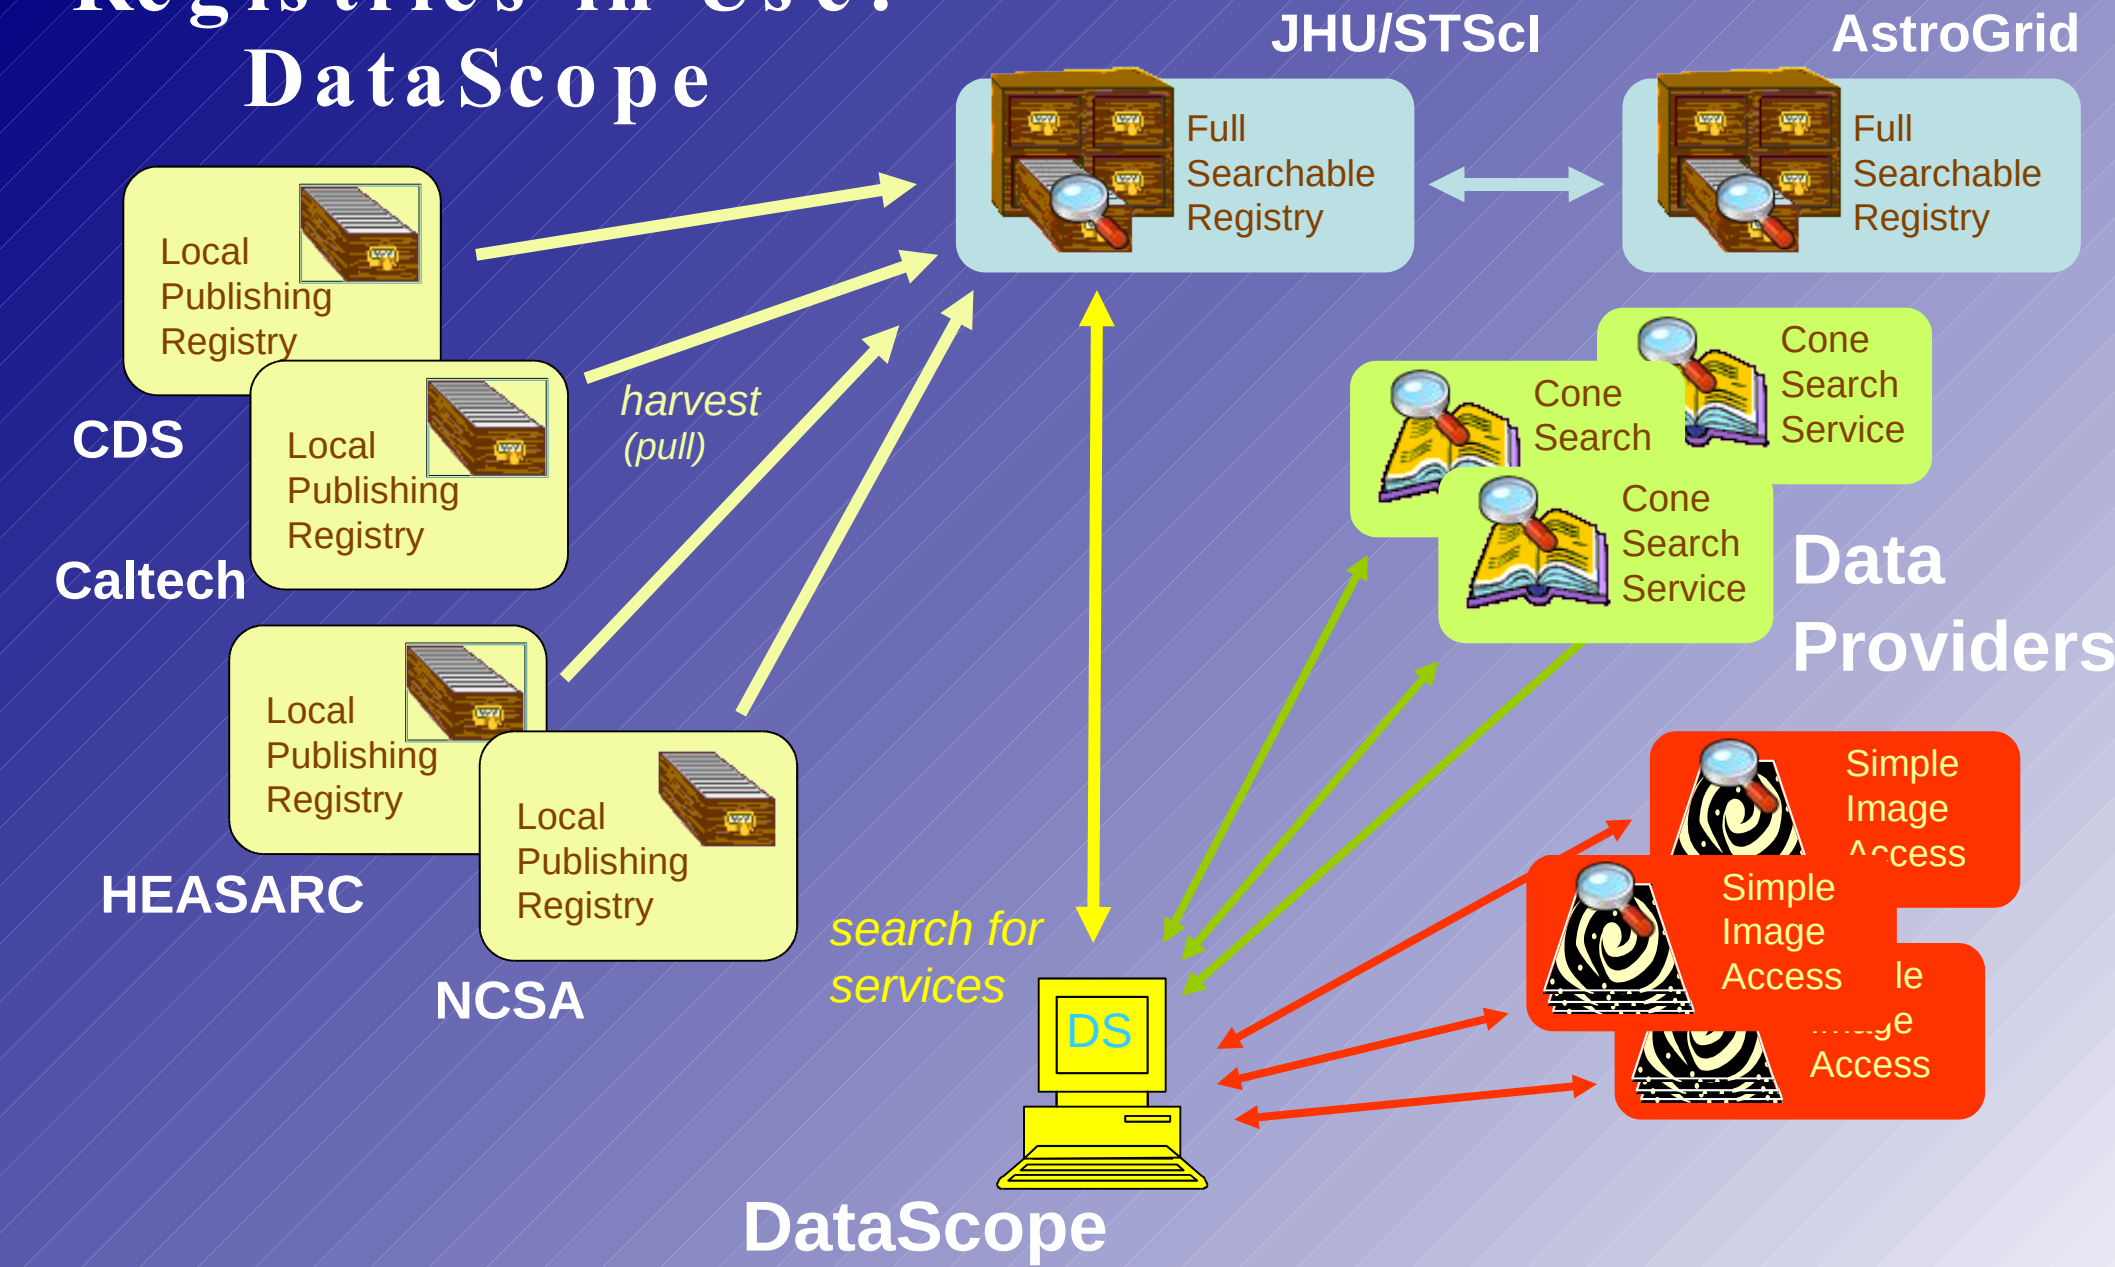

Resource Metadata Model

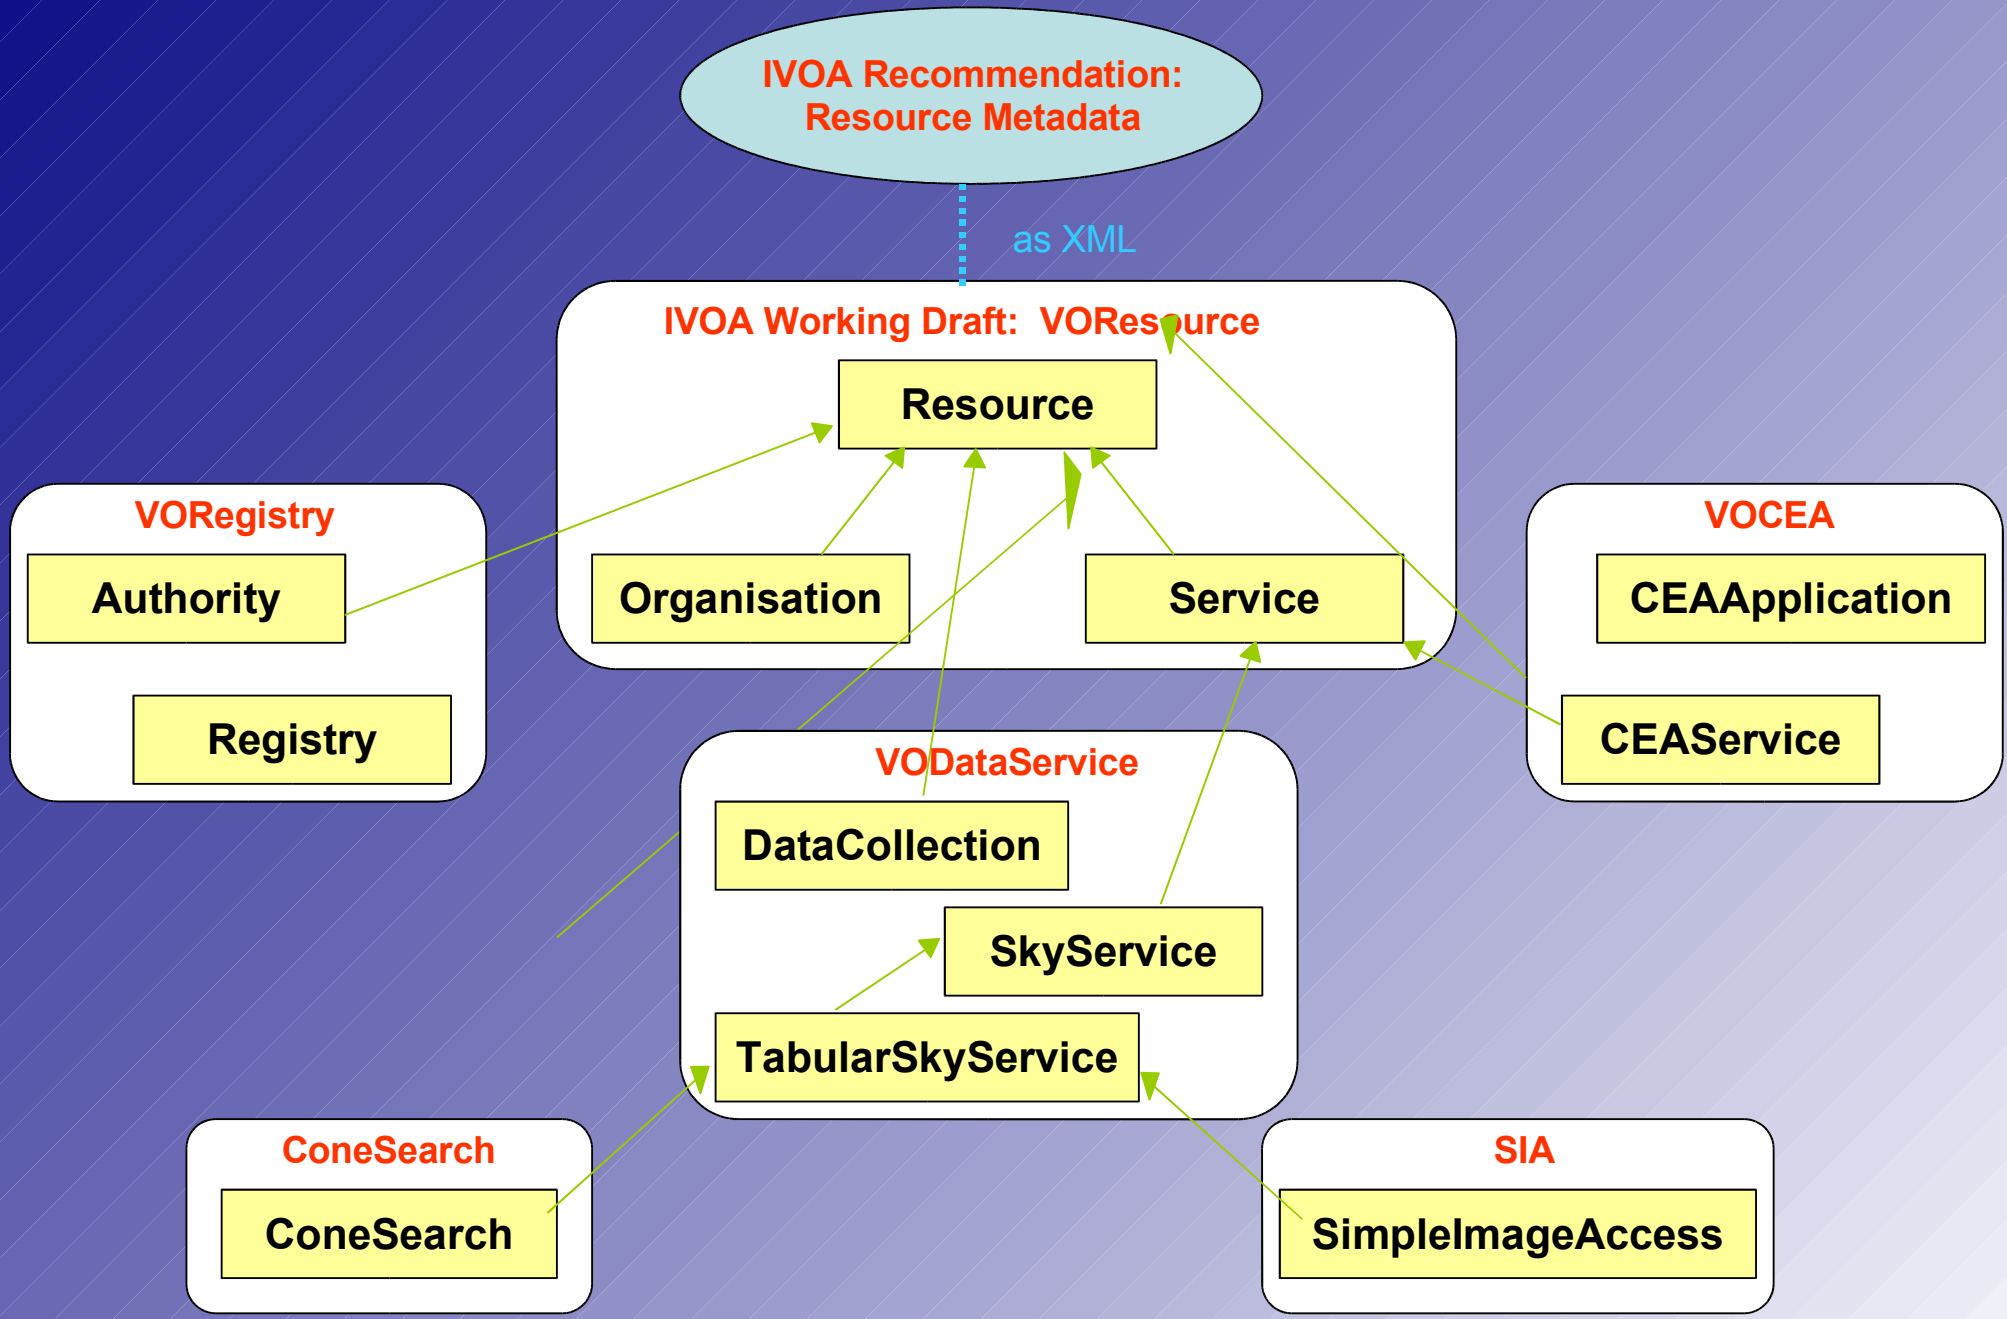

Virtual Observatory in Action

AVO
ASTROPHYSICAL VIRTUAL OBSERVATORY

VO First Science!

- First refereed astronomical paper enabled via end-to-end use of VO tools and systems: "*Discovery of optically faint obscured quasars with Virtual Observatory tools*", Padovani, Allen, Rosati, & Walton, 2004, A&A, 424, 545

ESA/ESO Press release May 28

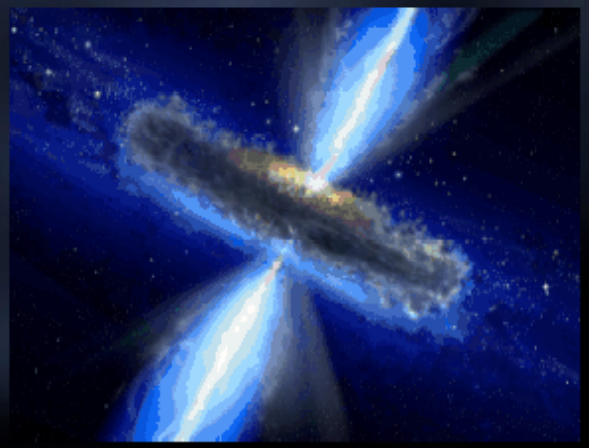

Oct. 26, 2004 P. Padovani, ADASS XIV 7

Virtual Observatory in Action

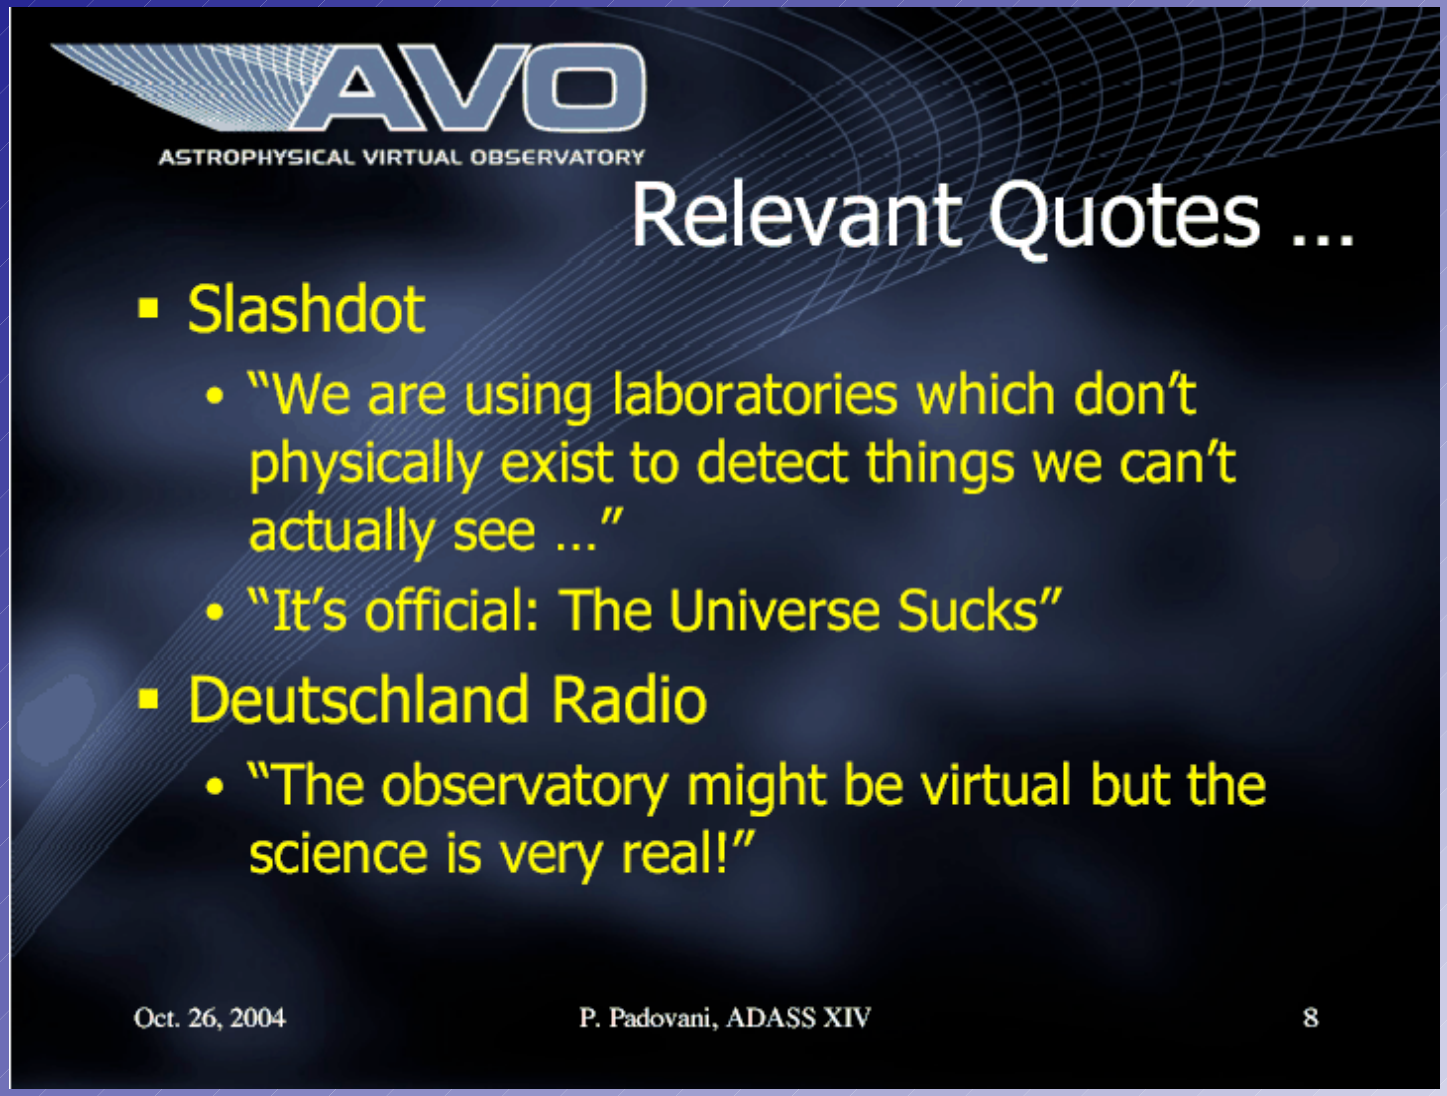

Pipelines

Data reduction and analysis pipelines are scripted, but
there is as yet no agreeement on a scripting language:

- Python (always up and coming, but not yet there)
- Perl (it can do anything)
- Java (everyone else uses it)
- IRAF (it's portable and tuned to astronomy)
- IDL (astronomers know how to use it)

Creating Data that Never Die: Building a Spectrograph Data Pipeline in the Virtual Observatory Era

by D.J. Mink, W.F. Wyatt, J.B. Roll, S.P. Tokarz, M.A. Conroy, N. Caldwell, M.J. Kurtz, M.J. Geller

Telescope Data Center
Smithsonian Astrophysical Observatory

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

http://tdc-www.harvard.edu

Abstract

Data pipelines for modern complex astronomical instruments do not begin when the data is taken and end when it is delivered to the user. Information must flow between the observatory and the observer from the time a project is conceived and between the observatory and the world well past the time when the original observers have extracted all the information they want from the data. For the 300-fiber Hectospec low dispersion spectrograph on the MMT, the SAO Telescope Data Center is constructing a data pipeline which provides assistance from preparing and submitting observing proposals through observation, reduction, and analysis to publication and an afterlife in the Virtual Observatory. We will describe our semi-automatic pipeline and how it has evolved over the first nine months of operation.

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

Hectospec Pipeline

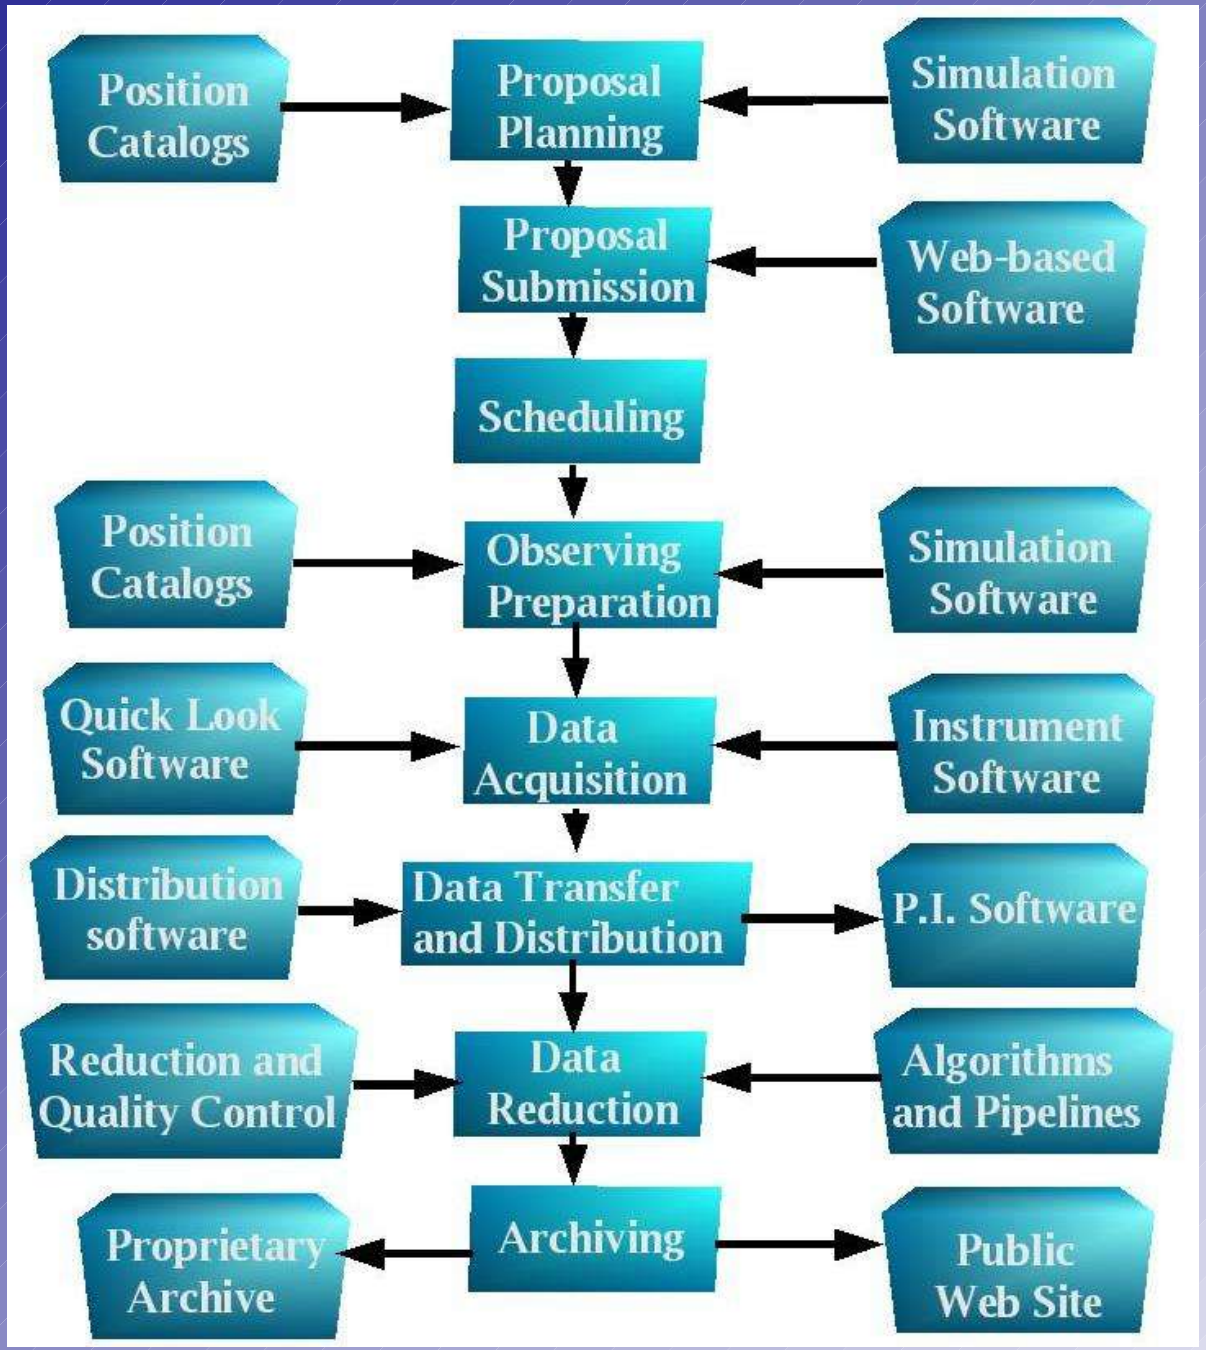

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

Program Conception

- Instrument information from the Web
- Previous observations from ADS and VO
- Catalogs from VO
- Positions from images reduced using VO catalogs and WCSTools

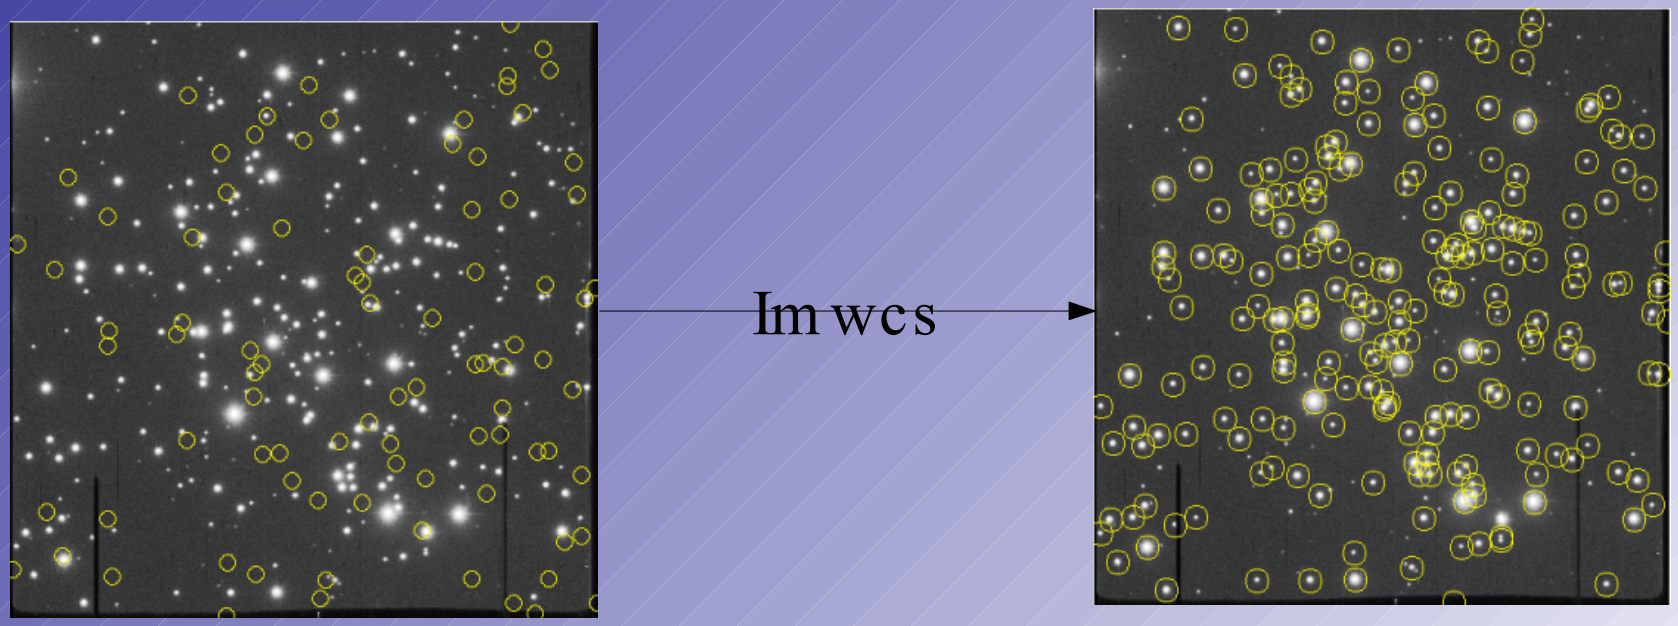

Proposal Submission and Evaluation

- Web-based Proposal-> online archive and database of proposals
- Web-based TAC evaluation software-> database of accepted proposals

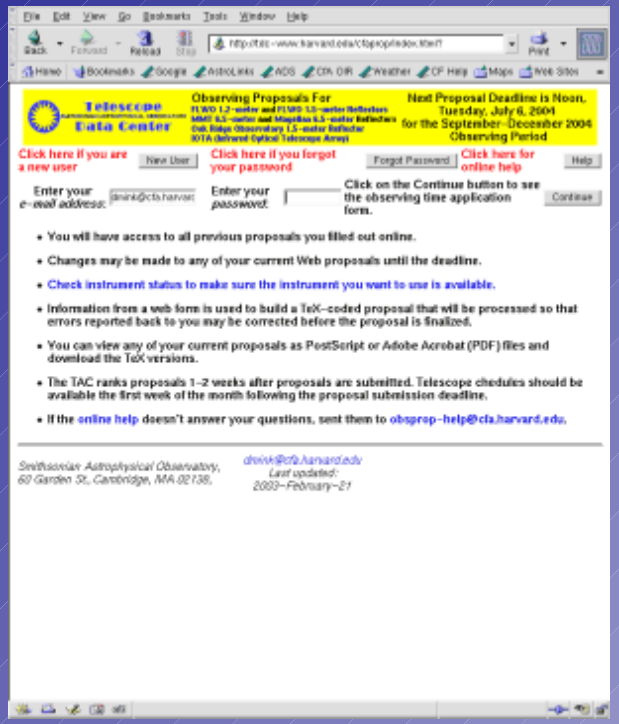

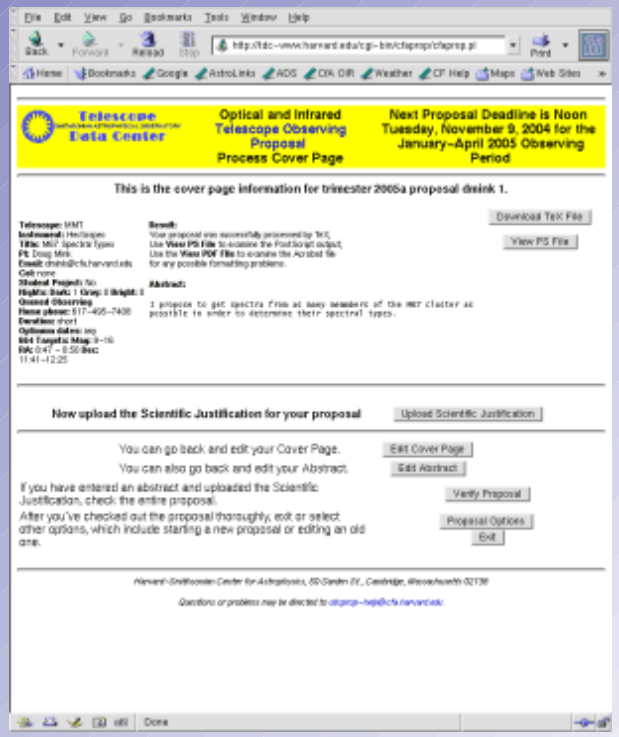

Telescope SMITHSONIAN ASTROPHYSICAL OBSERVATORY Data Center

Scheduling

- Web-accessible telescope schedules
- TAC-assigned program numbers follow data throughout its life (<year><A|B|C>-<TAC>-<number>, i.e. 2004C-SAO-11)

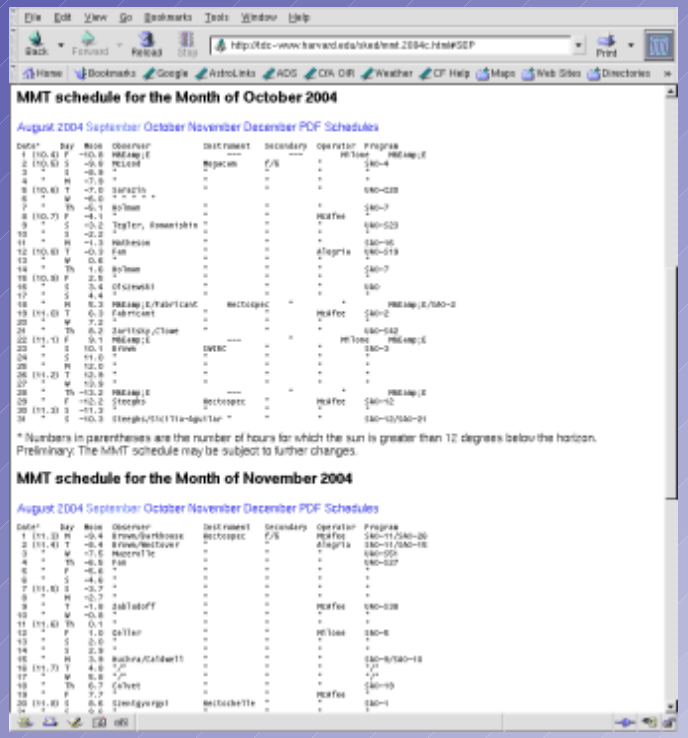

MMT schedule for the Month of October 2004

August 2004 September October November December PDF Schedules

* Numbers in parentheses are the number of hours for which the sun is greater than 12 degrees below the horizon.
Preliminary. The MMT schedule may be subject to further changes.

MMT schedule for the Month of November 2004

August 2004 September October November December PDF Schedules

Telescope Data Center
SMITHSONIAN ASTROPHYSICAL OBSERVATORY

Observing Preparation

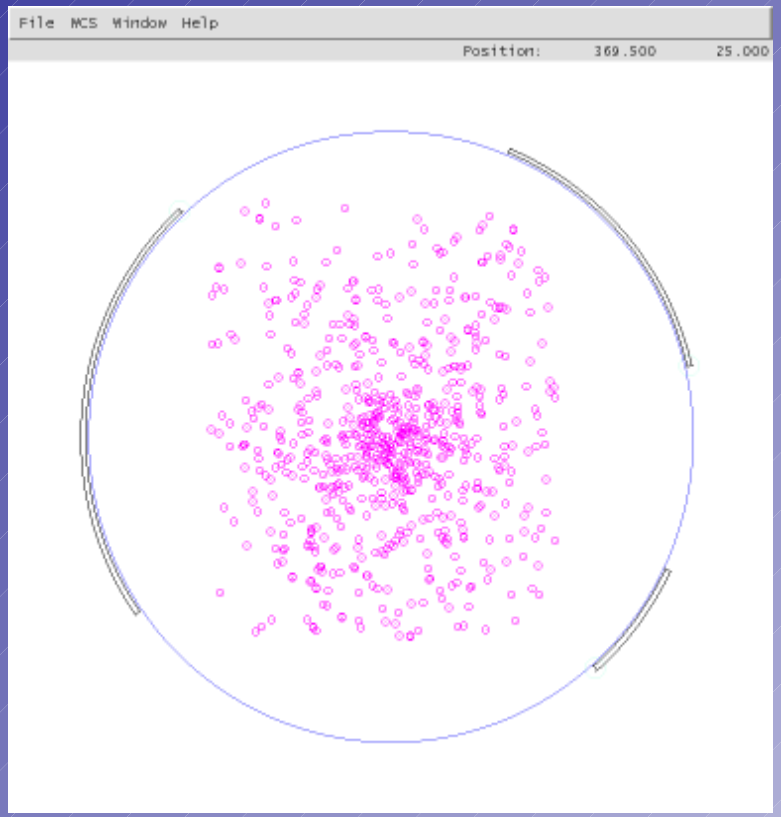

- Fiber positions from PI object catalogs
- Sky positions found using deep all-sky catalogs
- Fibers interactively assigned to program objects

Data Acquisition

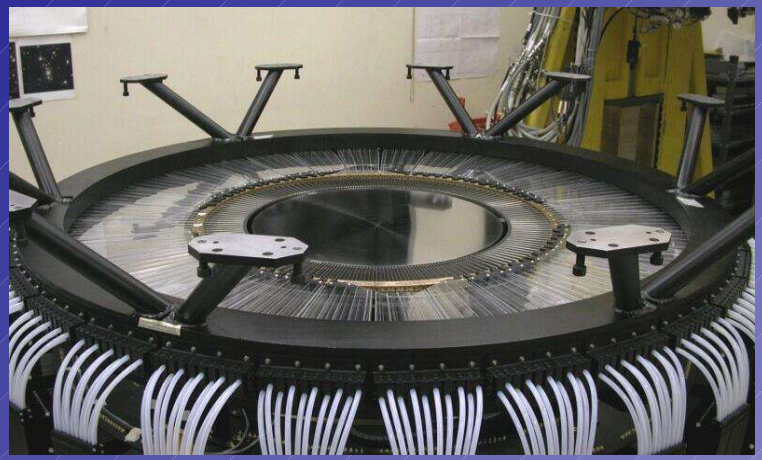

- 3 exposures per pointing are taken so
  cosmic rays can be removed
- Control software uses IRAF ICE
- Data is immediately ftp'ed to Cambridge

Data Acquisition

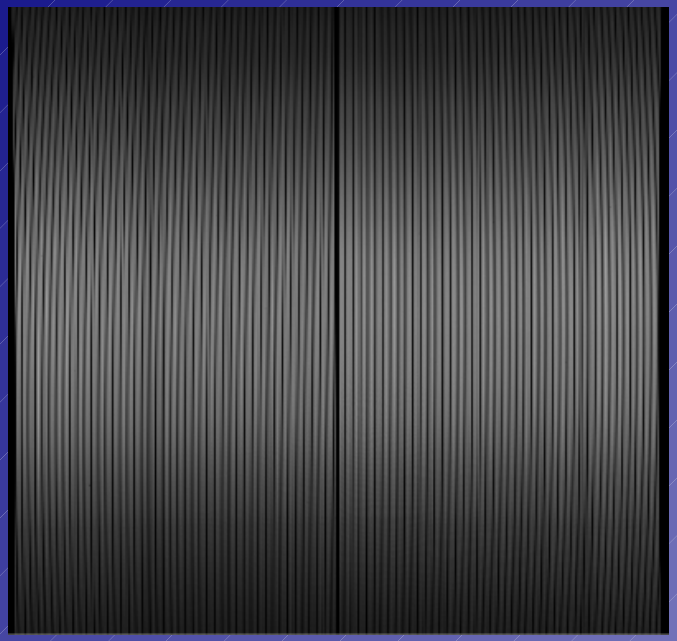

Dome flat spectra

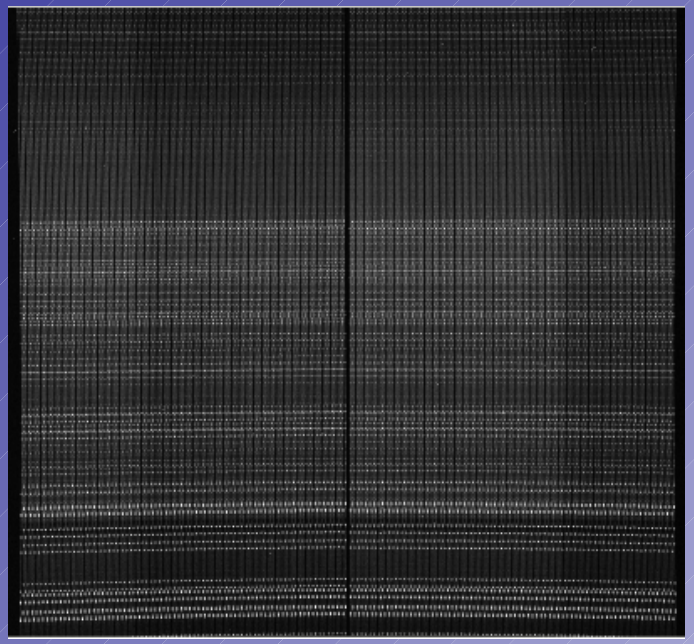

Calibration lamp spectra

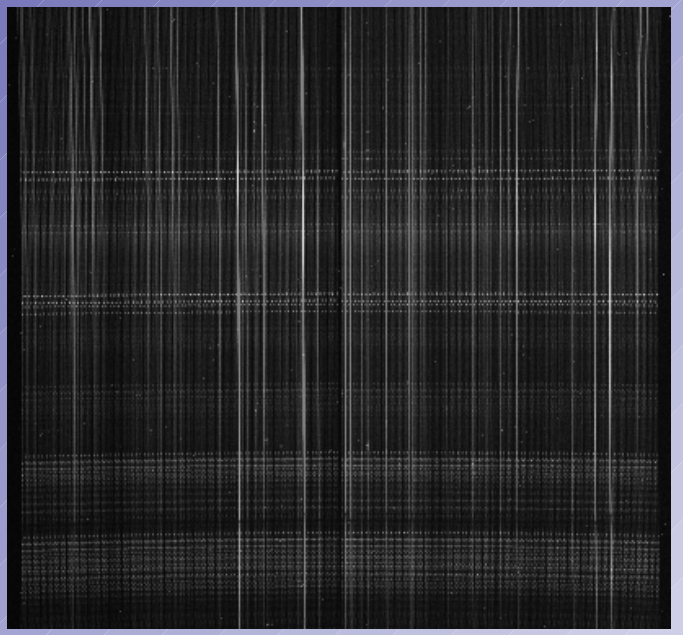

Object spectra with night sky

Telescope Data Center
SMITHSONIAN ASTROPHYSICAL OBSERVATORY

Data Reduction

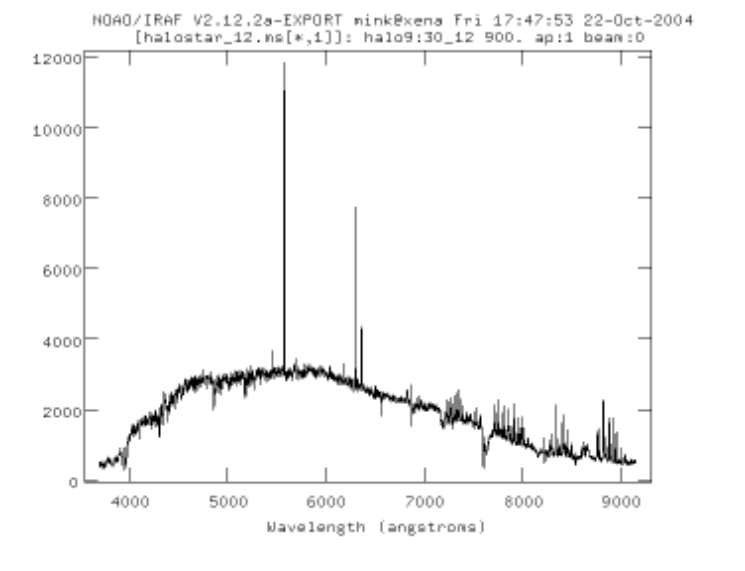

Extracted spectrum is mostly sky

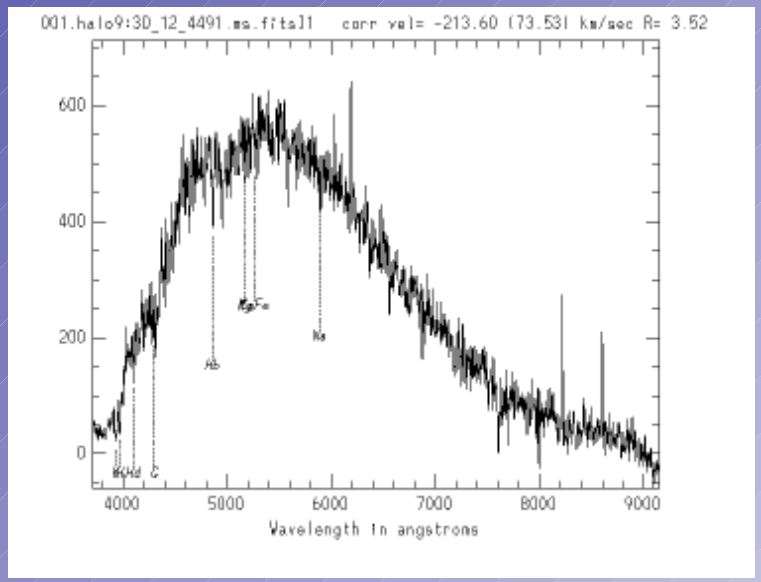

Same spectrum without sky

- Local IRAF CL and KSH scripts
- IRAF noao.imred.specred and noao.twodspec

Data Archiving

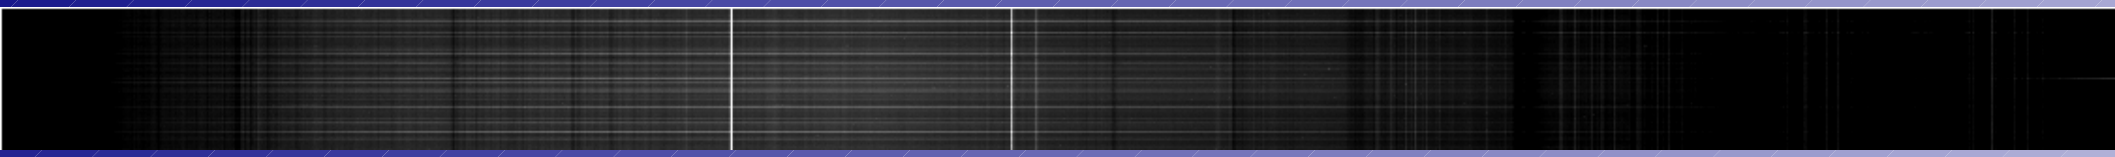

300 stacked spectra as seen by SAOimage

- Raw data
- Reduced data in single file per pointing
- Reduced data in one file per object
- PI's have access to their data

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

Virtual Observatory

Telescope Data Center — SMITHSONIAN ASTROPHYSICAL OBSERVATORY

Hectospec 300-Fiber Spectrograph

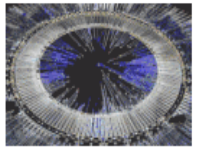

TDC HOME | TDC SEARCH | CONTACT TDC | OIR HOME | OIR SEARCH | CONTACT OIR

Hardware	Hectospec is a moderate-resolution, multiobject optical spectrograph fed by 300 optical fibers. It operates at the f/5 Cassegrain focus of the 6.7-meter MMT on Mount Hopkins in Arizona. The instrument is composed of a robotic positioner and a large, bench-mounted spectrograph, joined by a bundle of 25m long optical fibers. Dual robots, dubbed Fred and Ginger, reconfigure all 300 optical fibers in just 300 seconds.
Processing	[Preparation] [Observation] [Reduction] [Distribution] [Archiving] Getting Good Coordinates for Hectospec Hectospec Header Keyword Dictionary
Sample Image *(Click for larger image)*	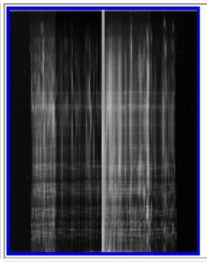
Sample Extracted Spectra	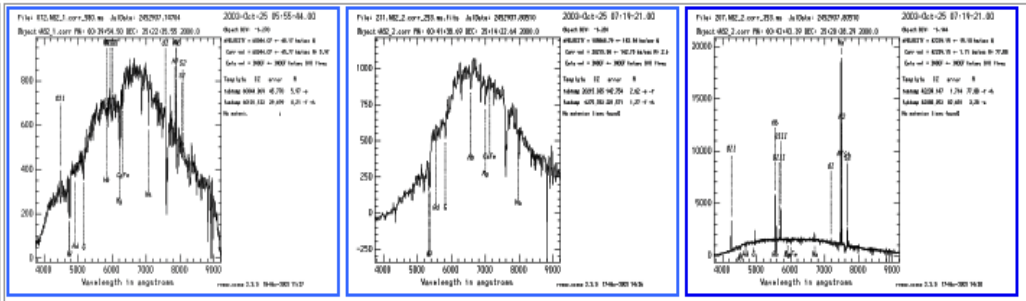 Good signal-to-noise data from a very preliminary reduction of engineering data taken in October 2003 (absorption and emission lines are marked based on cross-correlation redshifts) *(Click for larger graphs)* Here is some information on sky subtraction testing.

Virtual Observatory

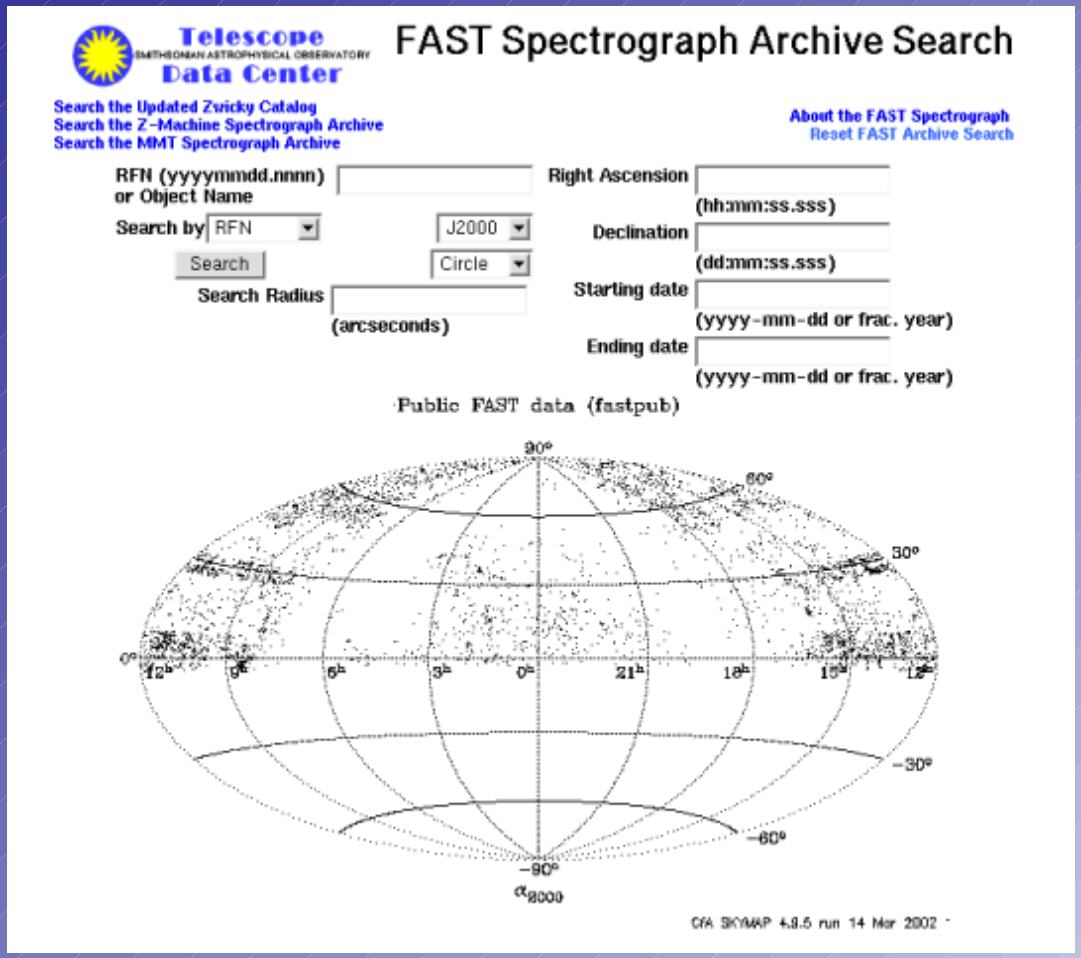

After a proprietary period all spectra will be made available through the VO as we already do for some of the spectra from our FAST spectrograph.

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center

Software

- An IRAF package containing all of the scripts used to reduce Hectospec software is available on demand.

- Complete documentation, including this presentation is available at:

http://tdc-www.harvard.edu/instruments/hectospec

My To Do List

- WCSTools: Set up user mailing list
  Upgrade to WCSLIB 3.6
  Add XML/VOTable parser to catalog search
- Archives: Add last two years of FAST calibration data
  Start setting up access to data by program
- VO: Set up registry in a box for TDC archives

Annular Solar Eclipse

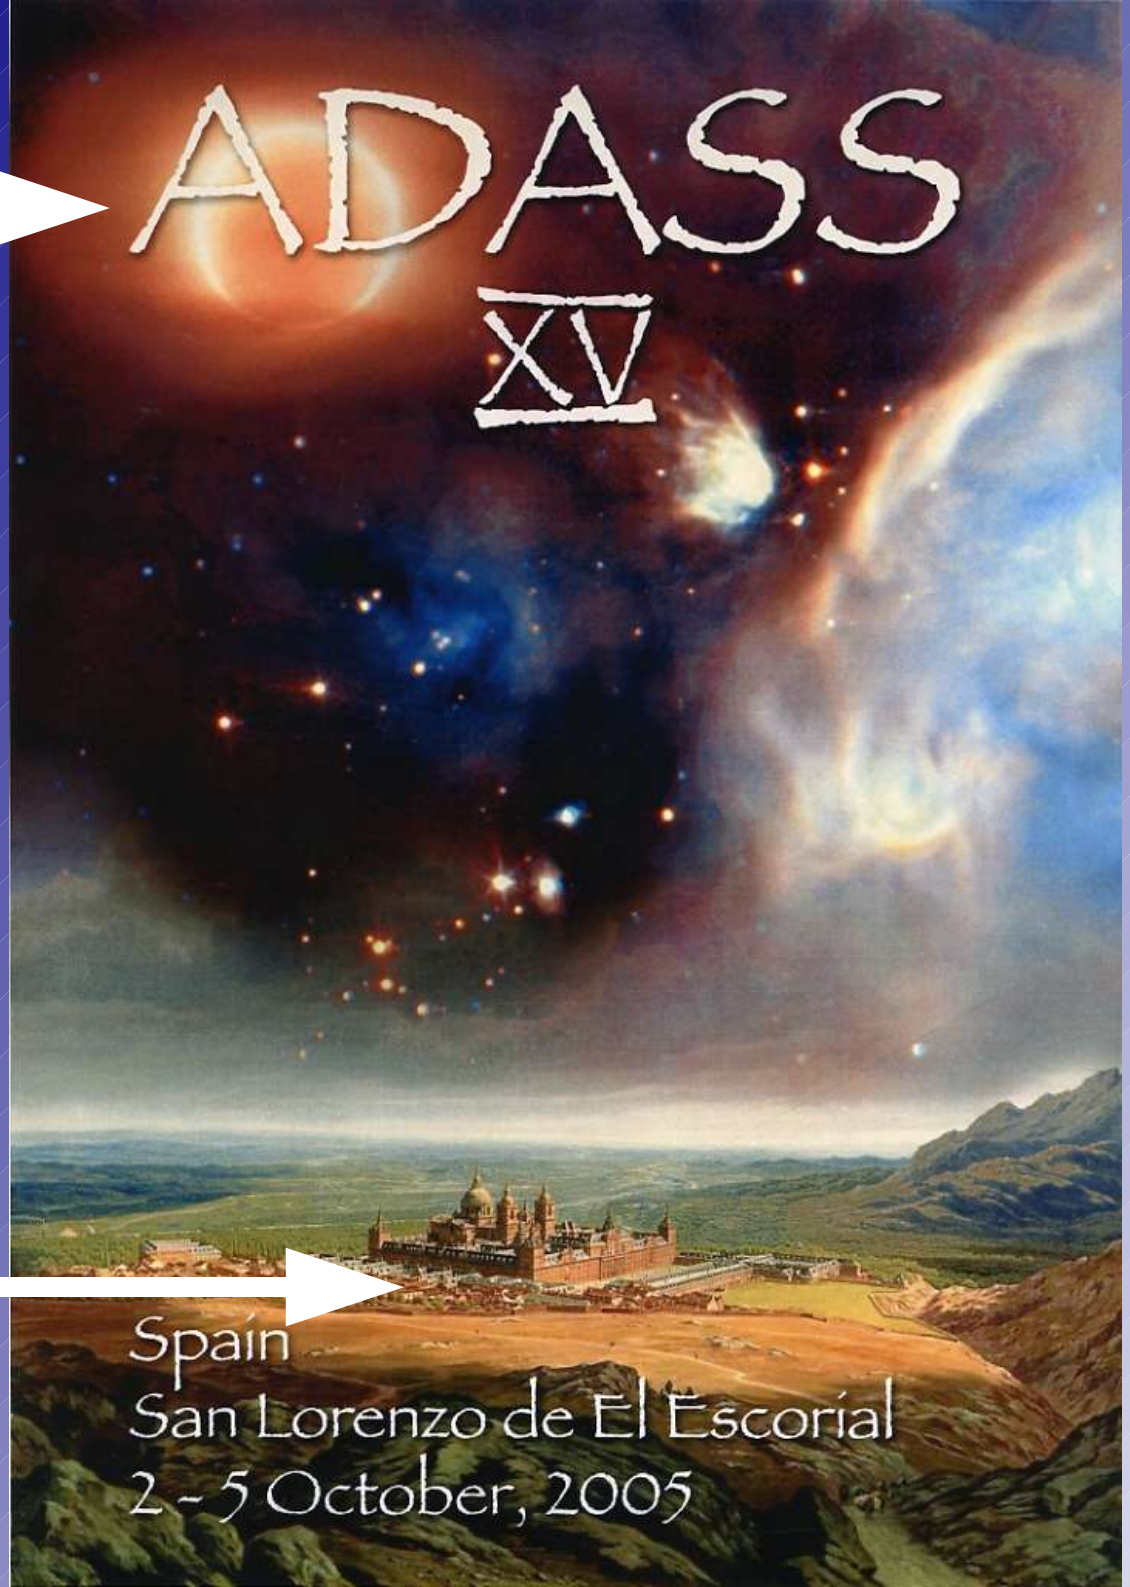

UNESCO World Heritage Site

Telescope
SMITHSONIAN ASTROPHYSICAL OBSERVATORY
Data Center