Computing IPT

B.E. Glendenning (NRAO)
J. Ibsen (JAO)
G. Kosugi (NAOJ)
E. Schmid (ESO)

J. Schwarz (ESO)

Outline

- Overall Status
- Software Development Organization
- Telescope Software (CSV)
- Observatory Software (DSO)
- “Top 10” Issues - Software
- Human Computer Interface (HCI) Initiative

Overall Status

- CIPT software has enough features to support early science
  - Bands 3,6,7 & 9 CSV SB-based observing, DSO Integrated tests
- Performance, efficiency, and robustness improvements will be the focus in the coming year
  - Significant improvements in R8, especially in performance, for Archive, CASA, Control, Correlator
- Telescope software (CSV) has more schedule risk than Observatory software (DSO)
  - AIV only requires modest support; routine operations
- The Archive has been a particular concern in the past year, remedial actions have already produced marked improvement

Computing IPT Organization

- Trilateral IPT – centrally coordinated contributions from each Executive
  - Close collaboration with JAO Computing Group – will be combined in operations (Operations Software IPT)
  - New EU CIPT lead: Erich Schmid (ESO)
  - 15 development sites, 4 continents!
- Work divided into 15subsystems
  - Functional (10): Control software, Archive software, …
  - Process (5): Management, Integration & Test, Common Software, …
  - Subsystems often split between sites, although with a center of mass

CIPT Releases

- Two releases per year (December 1 & June 1)
  - R8 = December 2010, R8.1 = June 2011
    - Deployed later than the release (e.g., R8 in January 2011)
  - In principle one major and one minor, in practice they are all major at this stage of the project
  - Usually one major patch per release => updates to users ~quarterly
  - Weekly iterations – small bug fixes
  - Release procedure defined: CIPT integration & test, CIPT test period, user test period, punchlist correction, final deployment
- Release Planning
  - Participation by all internal stakeholders
  - Twice per year: at annual (!) CDR and 6 months later.

Telescope Software

- Software required for telescope operations (principally used by AIV & CSV at present)
  - Equipment monitor & control
  - Front-end archive (monitor points, observational data)
  - Operator interfaces
  - (Observing Tool Phase II)
  - (Manual data processing/ analysis (CASA))
  - Dynamic & queued/ manual SB selection/ execution
  - Online calibrations

(items which are principally being used/tested by at present by AIV & CSV although they could be considered part of Observatory software)

Telescope Software: Current Status

- SB based observing (Observing Tool Phase II produced)
  - All bands, many spectral modes (~20 tested, more available)
  - "Single Field Interferometry" – can have multiple single-pointing sources in one SB
- AIV support ~complete (holography, optical pointing, etc.)
- Initial Front-End Archive installation (18 machines) complete, repository for monitor, observational, and project data
  - Including data versioning procedures
- Full suite of online calibrations
  - WVR only applied offline at present (FP6 contribution)
- All observational data flagged/calibrated/imaged in CASA

Telescope Software: Selected R8/R8.1 developments

- Control/Correlator: Mosaicing, full correlator performance, flagging & blanking, Incremental writes, final procurement of antenna & correlator computers
  - New total power implementation in place
- Scheduling: New implementation in place; includes dynamic scheduling algorithm in online scheduler
- Online calibration: Improved atmosphere calibration (Tsys, sideband ratios)
- Improved user access to monitor & observational data, ~~install~~ Storage Area Network installed at OSF
- CASA:
  - Parallelization of key tasks in progress
  - Performance in general significantly improved
- ACA correlator: fringes, spectrum obtained (see next slide)

Integration of ACA Correlator

- ACA Correlator operated with real antennas at AOS under ALMA 7.1.1 software in November
  - Band 3
  - 3 antennas
  - Spectral Window: 31.25 MHz – 512 ch (left), 2 GHz – 512 ch (right)

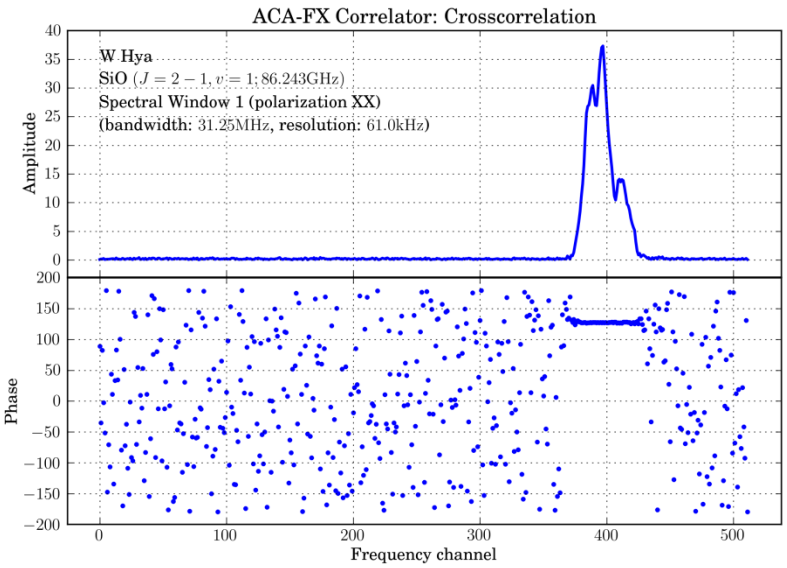

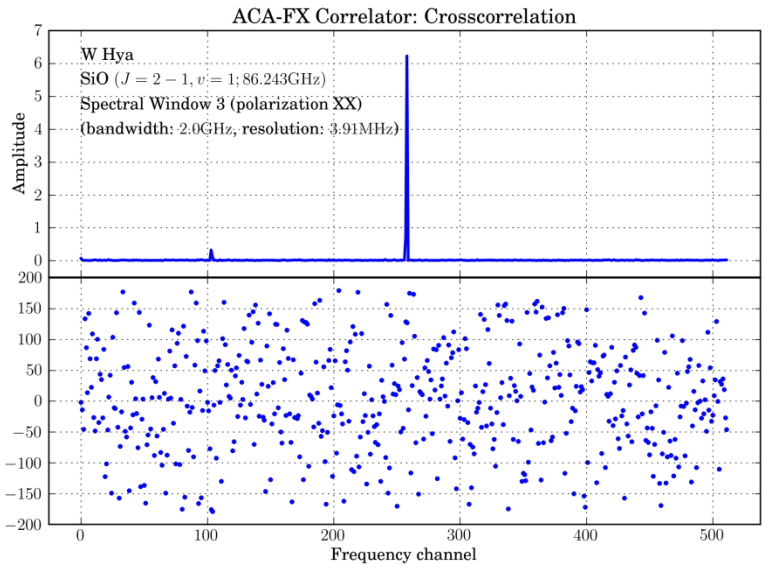

Computing IPT

Observatory Software

- Software required by ALMA to interact with the observing community, optimize the observing process, and quality check what has been observed
  - Distributed science archive
  - Science Pipeline
  - Phase 1 & phase 2 Observing Tool
  - TAC support
  - Observing project tracking
  - QA metrics extraction & tracking
  - User Portal
  - Data packaging & distribution
  - Manual data processing/ analysis (CASA)

Observatory Software: Current Status

- ~~9~~ 15 highest priority science archive queries implemented (42%)
- Archive data distribution mechanisms implemented (need SCO Archive for realistic testing)
- Science pipeline heuristics for single dish and interferometric observing (including pointed mosaics) implemented and tested on data from other telescopes
- CASA at general release 3.1; supporting EVLA (open to shared-risk observing since March 2010)

DSO's Integrated Test No. 3

- Currently underway
  - Phase 1 and Phase 2 Observing Tool ~~very advanced~~ undergoing final revisions for Cycle 0 CfP
  - Dynamic scheduling algorithm implemented in simulator
  - Phase 1 manager and Project Tracker ready for Early Science
- Submission test: 500 proposals ingested in 15 minutes
- Phase 2 preparation test in progress

Observatory Software: Selected R8/ R8.1 developments

- Installed Santiago archive
- Tested bulk data OSF-Santiago
  - Metadata, ARC distribution tests imminent
- Mosaics: R8 measurement of 87% on-source time (90% theoretical)
- Complete catalog support (including offline catalog in OT)
- Preliminary parallelization of pipeline processing
- Finish ACA 7m incorporation in Observing Tool (now partial)
- Inserted realistic/measured calibration overheads (from algorithms supplied by DSO) in Observing Tool
  - Systematic measurement campaign not currently planned
- Will implement Science pipeline processing directive flowthrough from Observing Tool for Cycle 1

Observing Tool Status

- Supports all Early Science cycle 0 modes
- Ephemeris support in place
- User-selection of Calibrators in a Science Goal possible
- 1-n single field positions on one target or pointed mosaic
- Includes calibration overheads in time estimates
- Users can only view projects they are investigators on.
- Galactic to J2000 coordinate transform possible
- More image servers in place (more suitable wavelengths, *e.g.* NVSS, SUMMS, IRAS...)
- Connection to Calibrator source catalog in Archive
- Production of Phase I SBs for APRC simulation (invisible to user)
- Improved help system

“Spectral line picker”

(I/F to Splatalogue)

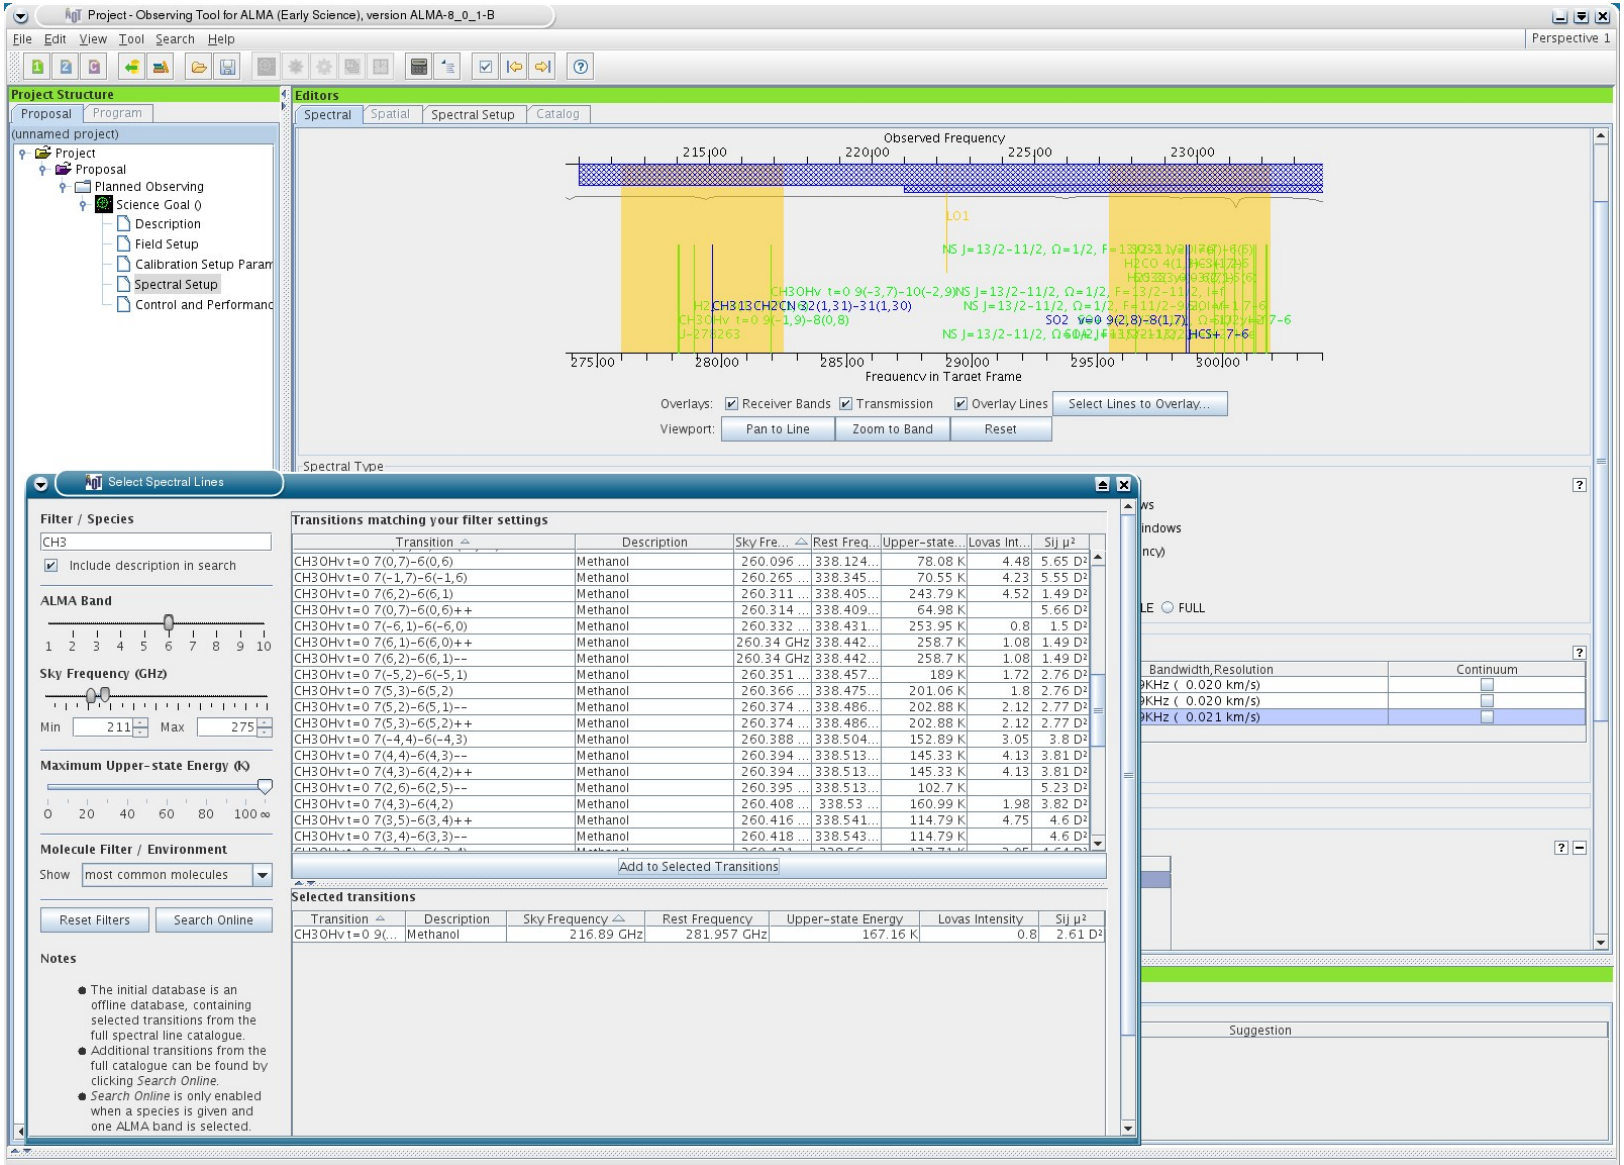

“Top 10” Issues - Software

- System robustness: reliable startup (Priority=2)
  - Two aspects: restart sometimes has to be repeated, restart takes 30m
  - First aspect hard to reproduce
  - Startup reliable, restarting the computers takes ~10m, arranging GUIs and getting ready for observations takes 20+ minutes
    - Latter aspect receiving attention (e.g., automatically reopen GUIs)
- System robustness: tolerance to hardware failures (1)
  - Change from “stop on unexpected problem” to “flag data on unexpected problem”
    - Now flag rather than fail when antenna or subreflector not on source
  - Antenna flagging in R8 (worst problem), will add other sources in priority order (probably ACD & lock problems)
  - Special mission in early 2011

“Top 10” Issues – Software (2)

- Overheads are far too high (4)
  - Concentrate on delays internal to a SB initially
  - Inter-subscan delay: was 6s, then 3s, became 2.0s in R8
  - Pretune frequencies (25s -> 1.5s), R8.1
  - CIPT/CSV team setup to analyze delays in observing mode
    - First investigation astronomical holography: 70% on source time
    - Mosaics (87% on source time R8 (90% theoretical))
- Telescope Monitor and Control Database (2)
  - CIPT plan presented last Friday to SE, CSV, DSO
    - Software impact variable
    - Plan generally accepted, to be iterated/finalized by CDR9
  - Zenoss trending/display infrastructure developed by ADC
    - Go/no-go decision by CDR9 at the latest

“Top 10” Issues – Software (3)

- ~~Shift log tool is not fit for purpose (4)~~
  - “this version is GREAT by the way, [some] improvements will be helpful but this version is a HUGE improvement from previously and is looking like a good tool that we are going to be able to use, so congratulations.”
  - The additional improvements requested are in progress
- Control of focus and pointing offsets not complete (2)
  - Band-dependent focus and pointing offsets in R8, follow-on tip-tilt changes can probably be considered normal work
- Data rate limitations (2)
  - ~~(Microwave link – temporary issue)~~ Fibre link in place (10 Gb/s)
    - Most non-real-time computers moved successfully from AOS to OSF
    - Data rates of ~ 18 MB/s routinely achieved (3 x specified average)
      - Often > 60 MB/s but still some sharp drops (under investigation)
  - Correlator Data Processor (16 computer cluster) performance
    - Was implemented as an end-to-end realtime system
      - Downstream hiccups stopped processing, made it hard to use all cores
    - First part of redesign/reimplementation completed in R8, More powerful computers (currently under test)

“Top 10” Issues - Archive

- Archive reliability; efficiency and ease of use of the Archive (2)
  - Reliability
    - Front-end Archive hardware / databases have excellent formal availability
    - In practice the Oracle part encountered occasional internal congestion which resulted in data dropouts and general sluggish performance
      - We very recently discovered a third-party library/Oracle version incompatibility that we believe is the cause of this problem
      - ~~In addition we have purchased and will shortly install a~~ Storage Area Network (SAN) installed and has ~~which will~~ greatly increased the Oracle throughput
    - The bulk data transmission software ~~can cause~~ crashes/hangups largely resolved

Archive Ingest of Monitor Data

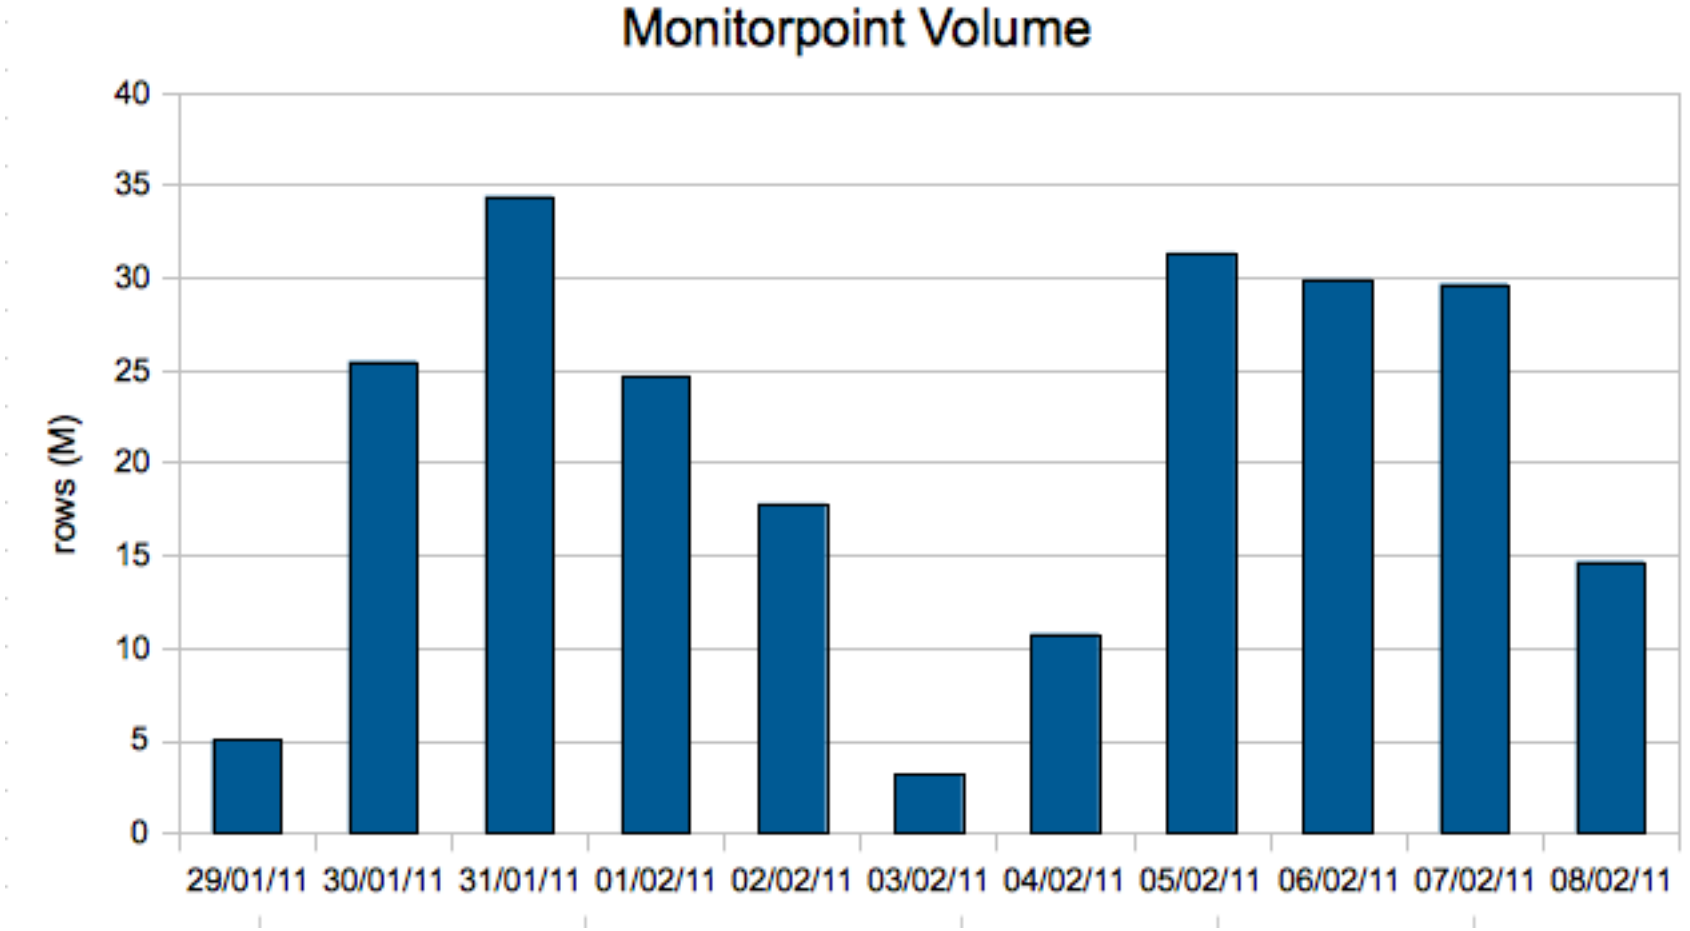

Hourly Breakdown

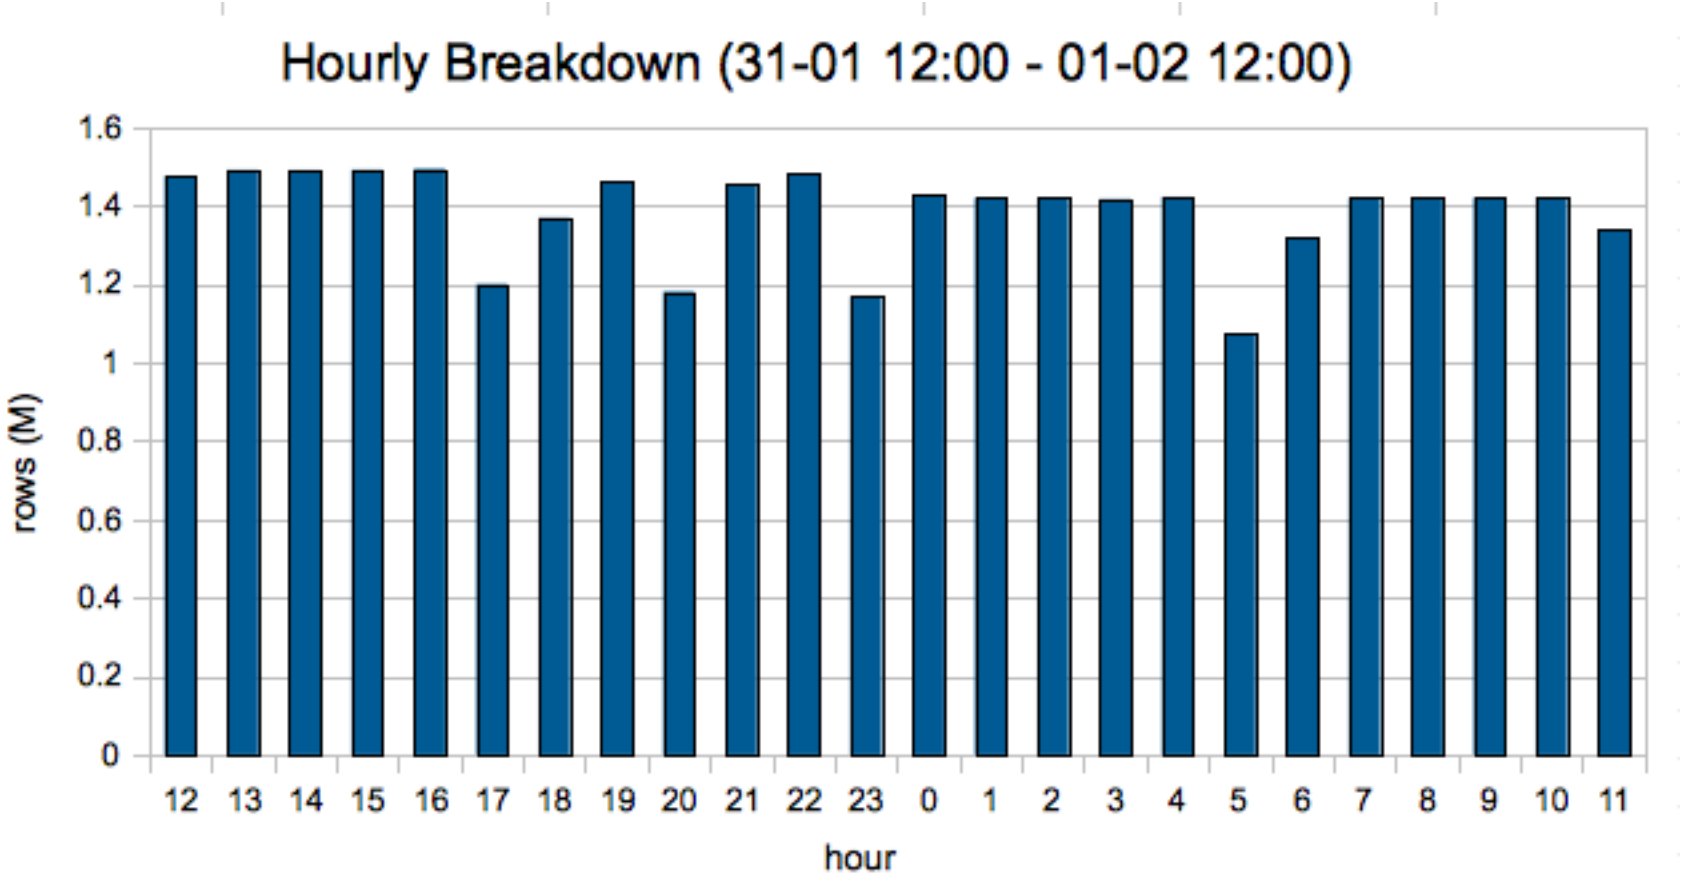

Display Monitor Points w/Zenoss

- Retrieve information from text files (prototype done by ADC)

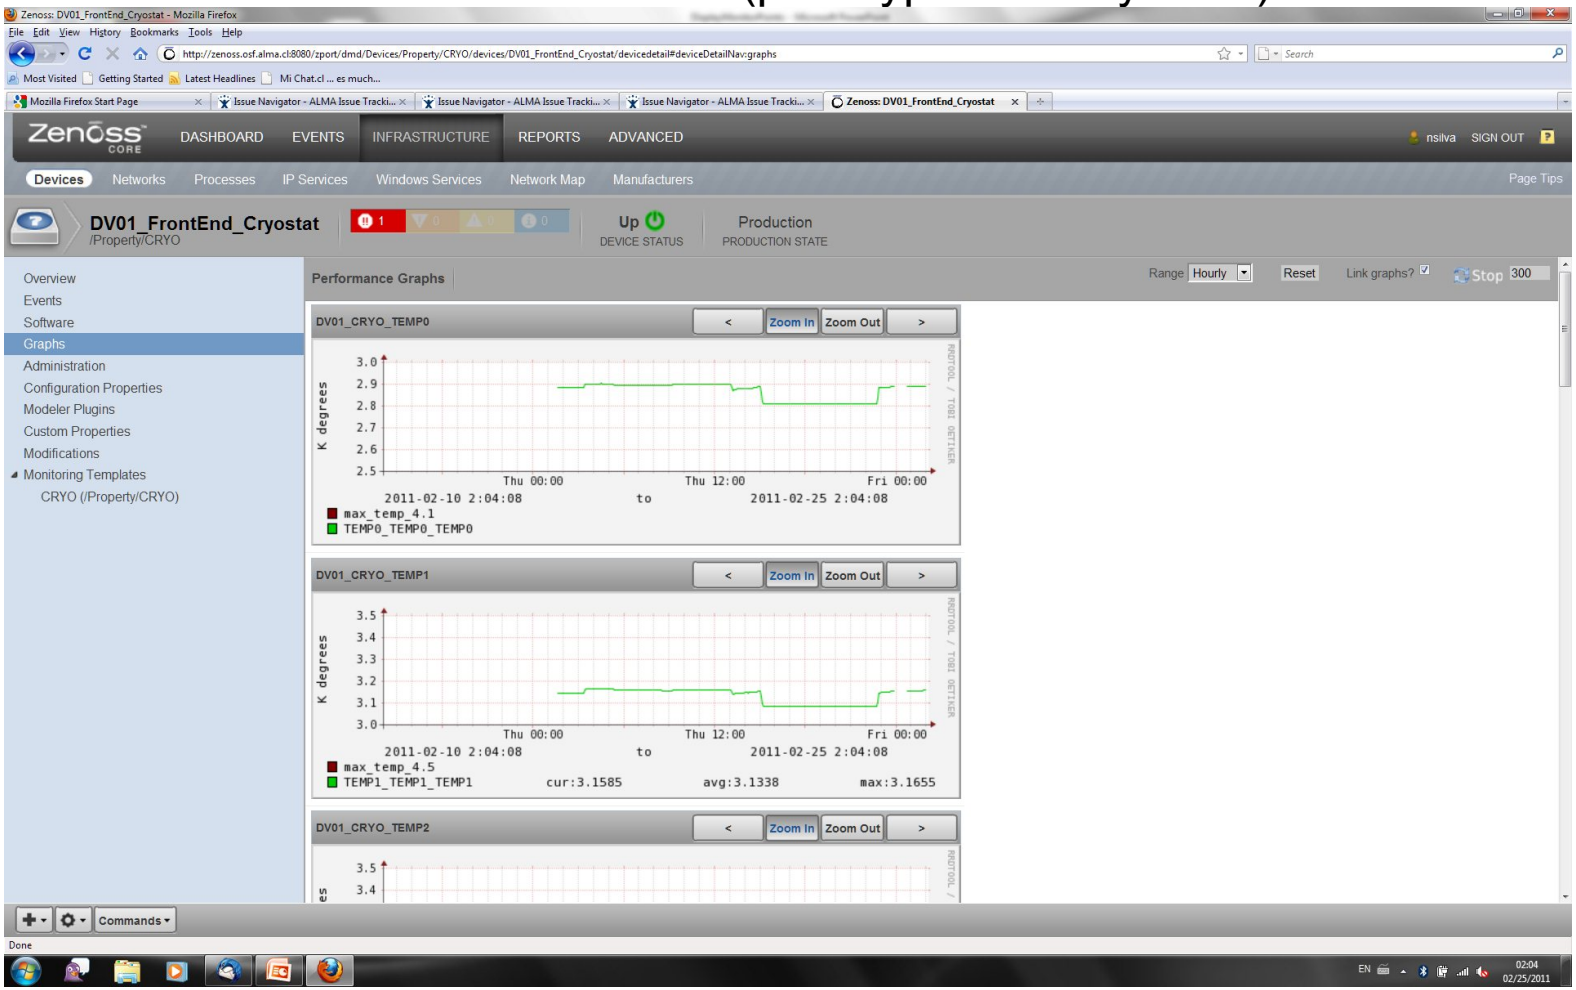

“Top 10” Issues – Archive (2)

- Efficiency
  - We believe Oracle performance issues are now resolved
    - More indexing need to be enabled
    - Harvesting of **selected** XML attributes to relational fields underway
  - Bulk data archiving (NGAS part) at > peak data rate demonstrated in a standalone test environment
    - ~~Cannot~~ Have tested at ALMA with new fibre link ~~yet~~ (peak data rate ~500 Mb/s~~, microwave link is ~100 Mb/s (single duplex)~~)
- Ease of Use
  - CSV uses command line tools, shared directories, and periodic dumps of monitor points into text files
    - We will install the science archive portal at the OSF
  - Common (“canned”) queries not available
    - Multi-parameter queries to be delivered to CSV In May
    - CIPT staff totally consumed by front-end archive issues, “tiger team” (Ops/ARC staff) formed to implement queries (15 (of 36) highest priority queries implemented so far)
    - Low-level (SQL) access provided to nominated CSV members

CASA General Description

- CASA is the post-processing package for ALMA (and EVLA) both interferometric and single dish
- The ALMA pipeline is is being built from CASA *toolkit*
- Toolkit packaged into most commonly used *tasks* for users
- Designed with parallelization in mind
- Bypasses many of the "shortcuts" inherent in older packages (that were created to process much lower dynamic range data)

- Other details:
  - The algorithms are in C++; interface in python/ipython/Qt
  - Fully scriptable, with in-line help and scientist-written documentation (notably the user manual/cookbook)
  - Interferometric calibration and imaging capabilities implemented via the Hamaker, Bregman, Sault formalism (Measurement Equation)
  - Telescope data (visibility and single-dish) are stored in a Measurement Set (MS); a filler converts ALMA and EVLA raw data (SDM) to the MS

CASA Status

- MS Filler is in daily use by both ALMA and EVLA.
- CASA has addressed approximately 73% of the specified interferometric requirements.
  - 53% are deemed satisfactory
  - 20% although implemented, require enhancement.
- The core functionality needed to calibrate and reduce data from ALMA exists and is in use.
  - Improvements will simplify and improve the data reduction path.

Although there are some features which need development and improvement, currently the primary concern for CASA is performance.

CASA Performance

- Bugs (Fixed)
  - Improper configuration of storage manager causing each channel to be accessed separately (50x slowdown in Split)
  - Erroneous access pattern in split causing 10x data access
    - Consolidation of data access patterns will minimize the number of bugs.
- Optimizations
  - Development of Asynchronous-I/O methods
  - Removal of “Scratch” Columns to decrease I/O and disk consumption.
  - “Pipelined Processing”
- Parallelization
  - Our approach is the same for both clusters and workstations.
  - CASA 3.2 will have a filler to imager parallel pipeline (uses multiple cores)
    - *Restrictions: Continuum Only, Single SDM.*

CASA Timing and Parallelization

- We are currently working with the CSV staff to generate a representative (~100 GB) Early Science ALMA data set.
- Until that is available we are using an EVLA data set of comparable size (100 GB; 16x128 channels, dual polarization)
- Still a work in progress, we are still in the process of analyzing the timings, identifying bottlenecks and beginning to design solutions.

Reduction Time vs. Number of Engines

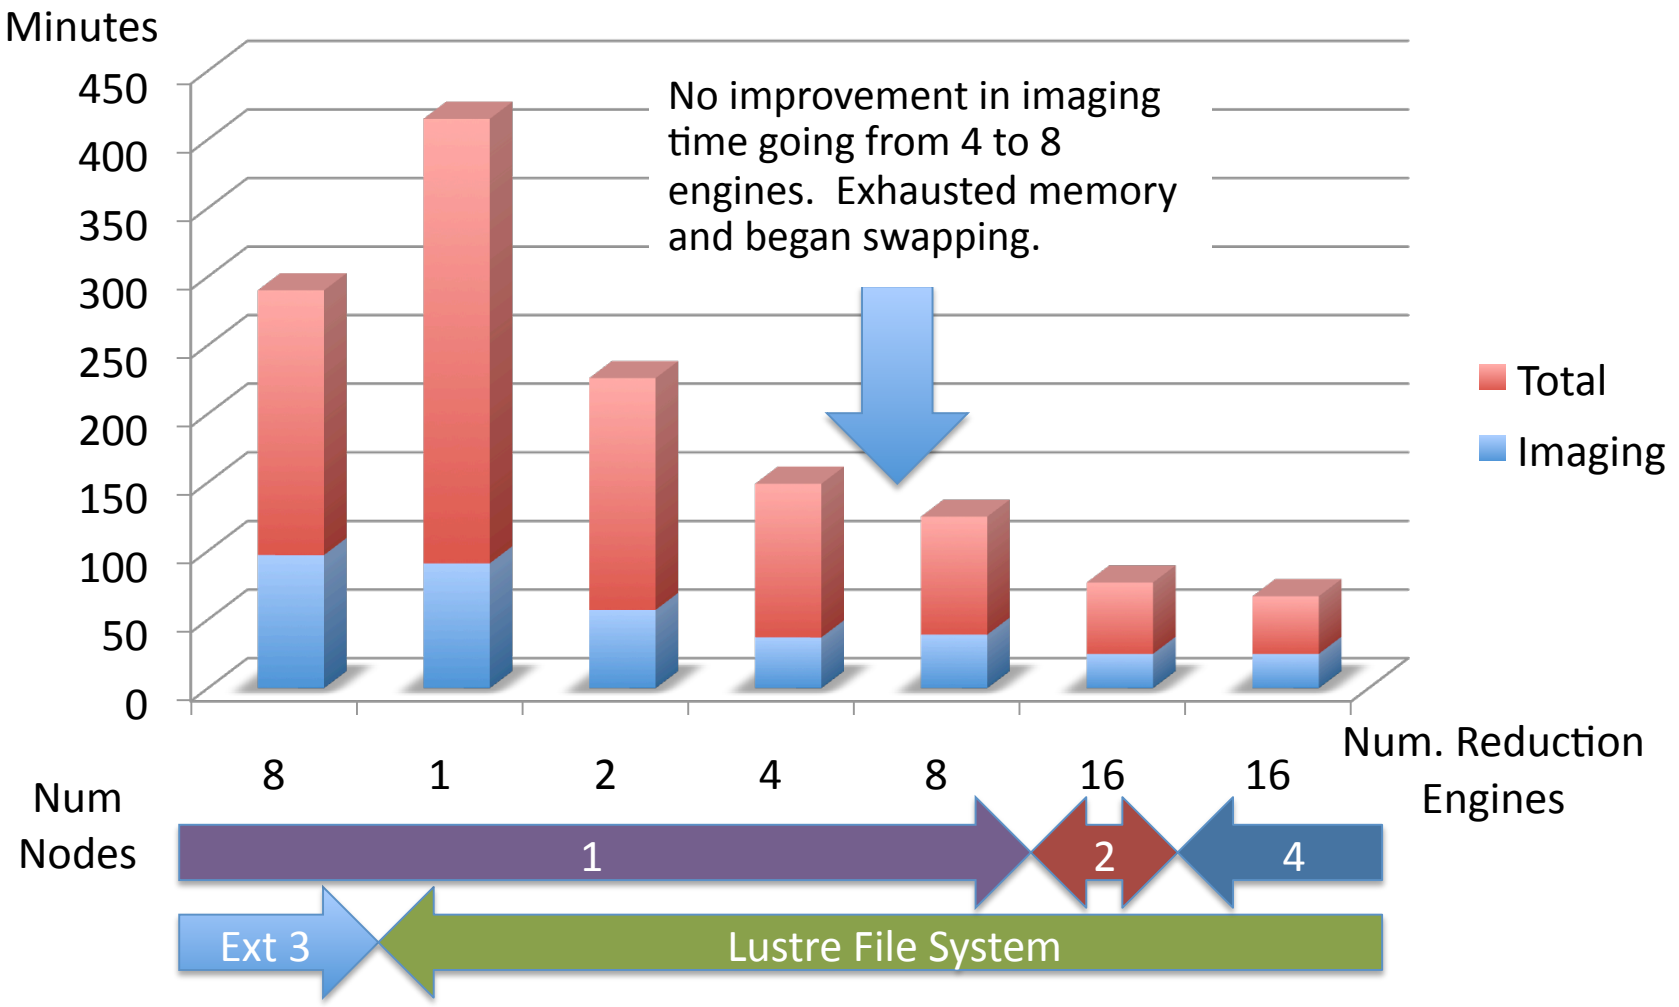

Task Timings

Serial

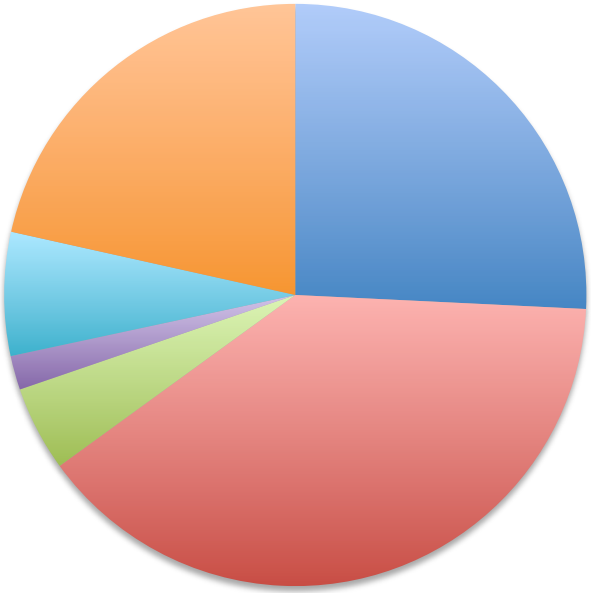

8 Engines

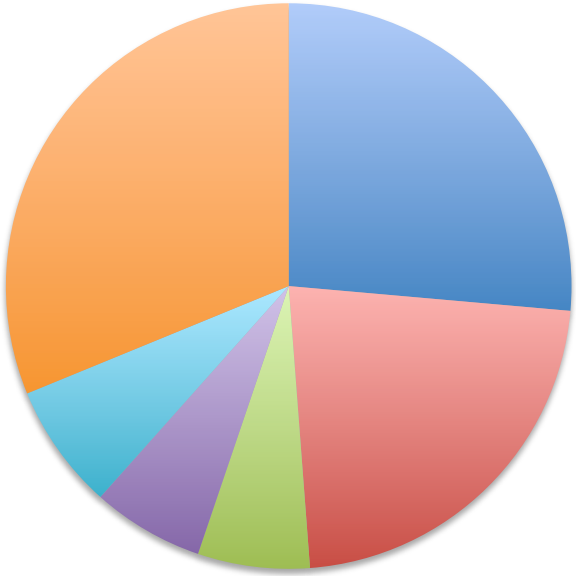

16 Engines

Split Flag ClearCal Cal Solve Apply Cal Imaging

CASA Single Dish
Sunspots & Solar prominences

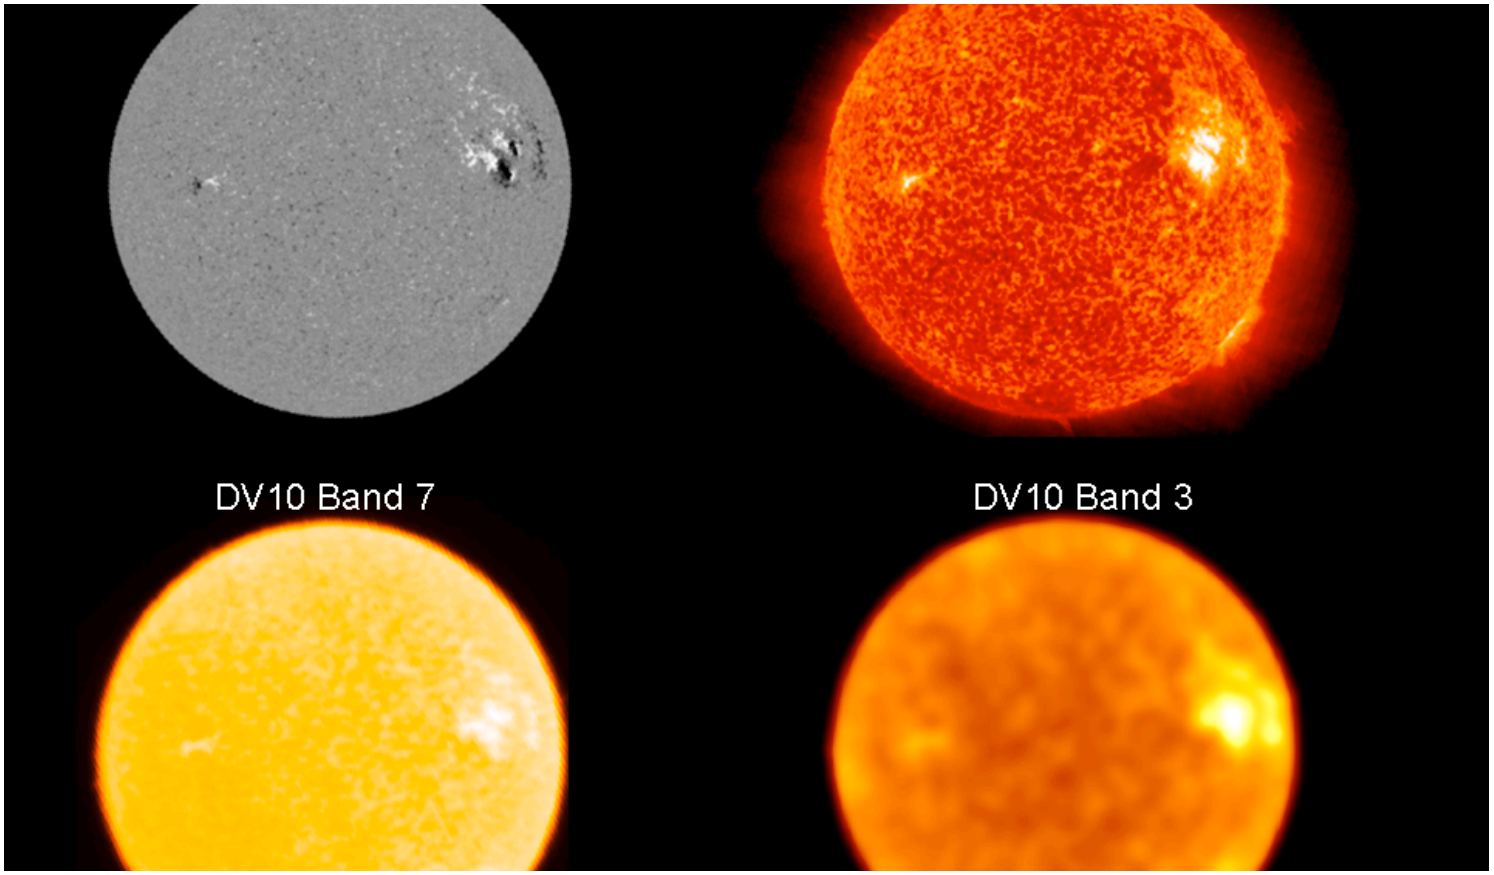

From R. Hills, http://wikis.alma.cl/twiki/pub/AIV/MaterialsForTalksAndLectures/Sun_Test_Scans_Feb2011.ppt

CASA Single Dish

- Single Dish Capabilities are well advanced.
  - NAOJ deliverable to ALMA project
  - Single dish functionality is based on ATNF Spectral Analysis Software (ASAP)

Performance issues are being identified and corrected.

Human Computer Interface (HCI) Initiative

- Remaining action from the November 2008 JAO review of the CIPT: Get some expert HCI advice
- First workshop at OSF with HCI researchers/consultants from INRIA & CNAM, France
  - Included Operators, AIV/CSV/DSO Scientists, Nov 29 – Dec 7
  - Produced first phase of new interface (next slide):
    - Geographical antenna map, semantic zoom: to be part of R8.1
    - Adjacency matrix to show status of ~2000 baselines in single window: projected for R9.0
- OT Progress
  - Observing Tool: Modest changes, to be incrementally accomplished with existing effort:
    - Steadily working on them, many in place, some not

New Antenna Map, Baseline I/F

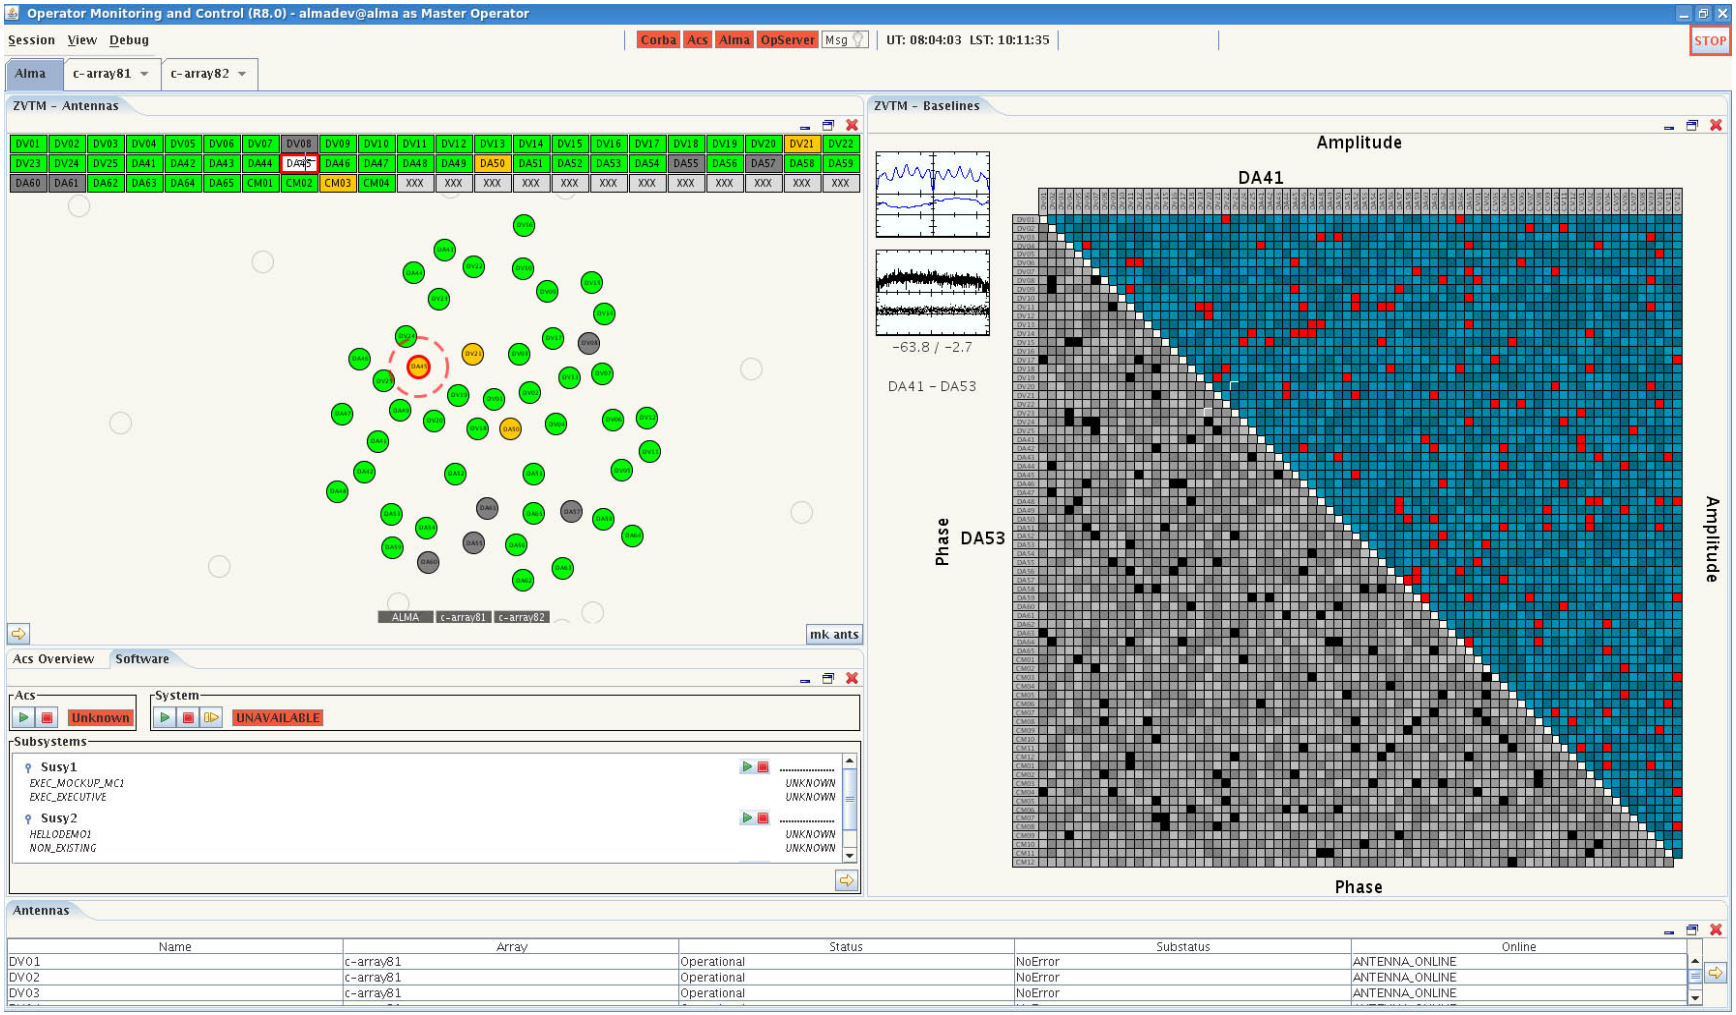

Semantic Zoom