ALMA quality assurance: concepts, procedures, and tools

A.M. Chavan*[a], S.L. Tanne[b], E. Akiyama[c], R. Kurowski[a], S. Randall[a], B. Vila Vilaro[b], E. Villard[b]

[a] European Southern Observatory, Karl-Schwarzschild-Str. 2, 85748 Garching, Germany;
[b] Atacama Large Millimeter/submillimeter Array (ALMA), Alonso de Córdova 3107, Vitacura, Santiago, Chile;
[c] National Astronomical Observatory of Japan, 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan

ABSTRACT

Data produced by ALMA for the community undergoes a rigorous quality assurance (QA) process, from the initial observation ("QA0") to the final science-ready data products ("QA2"), to the QA feedback given by the Principal Investigators (PIs) when they receive the data products ("QA3"). Calibration data is analyzed to measure the performance of the observatory and predict the trend of its evolution ("QA1").

The procedure develops over different steps and involves several actors across all ALMA locations; it is made possible by the support given by dedicated software tools and a complex database of science data, meta-data and operational parameters. The life-cycle of each involved entity is well-defined, and it prevents for instance that "bad" data (that is, data not meeting the minimum quality standards) is ever processed by the ALMA pipeline.

This paper describes ALMA's quality assurance concepts and procedures, including the main enabling software components.

Keywords: Radio Astronomy, Quality Assurance, Science Operations, Data-flow Software

1. INTRODUCTION

The Atacama Large Millimeter/submillimeter Array (ALMA) strives to provide the community with scientific data of the highest possible quality. An extensive Quality Assurance program was devised during the construction phase and is being constantly refined to be more effective and efficient: the product QA begins as data is being taken and follows it as data is processed, packaged and distributed to the Principal Investigators (PIs) — their feedback is also taken into account and incorporated in the process.

The plan includes four distinct tiers, named QA0 to QA3.

QA0 is performed in quasi-real-time by astronomers and operators on duty as observations proceed by examining the QuickLook (QL) display and extracting selected calibration information from the generated data. It is meant to confirm that the observation was performed within the constraints set by the PI and the data can be calibrated correctly.

The *QA1* tier monitors the overall system performance by executing special observations defined by the Observatory calibration plan, and collecting and analyzing the data. Long-term trends can be identified and used to perform preventive maintenance and problem avoidance. Not that unlike the other steps QA1 is not directly tied to a specific observing project.

All science observations are reduced (with the ALMA Pipeline or manually) and selected parameters are assessed to evaluate data quality -- that is what ALMA calls *QA2*. ALMA is committed to deliver science-ready data and this stage checks whether the data meets all scientific quality constraints; observed sensitivity and angular resolution must meet the requirements set by PI. Reduced, validated data is stored in the ALMA Science Archive where it can be

* amchavan@eso.org

accessed by the PI, and, after the proprietary period expires, the entire community. The QA0, QA1 and QA2 reports are delivered to the PI as well.

Any defects detected by the PI after data delivery (*QA3*) are collected by the ALMA Regional Centers (ARCs) and fed back to the Joint ALMA Office (JAO) in Chile for investigation and resolution.
Dedicated software tools and databases support all QA stages, from the QuickLook display to the Pipeline to the ALMA Quality Assurance tool (AQUA) to many ad-hoc scripts and displays. In the following sections we detail ALMA's quality assurance processes and the tools and infrastructure elements supporting them.

1.1 Structure of an ALMA Observing Project

ALMA Observing Projects are assembled hierarchically. The actual structure is rather complex but QA is concerned only with a handful of elements, as shown in the following, simplified example.

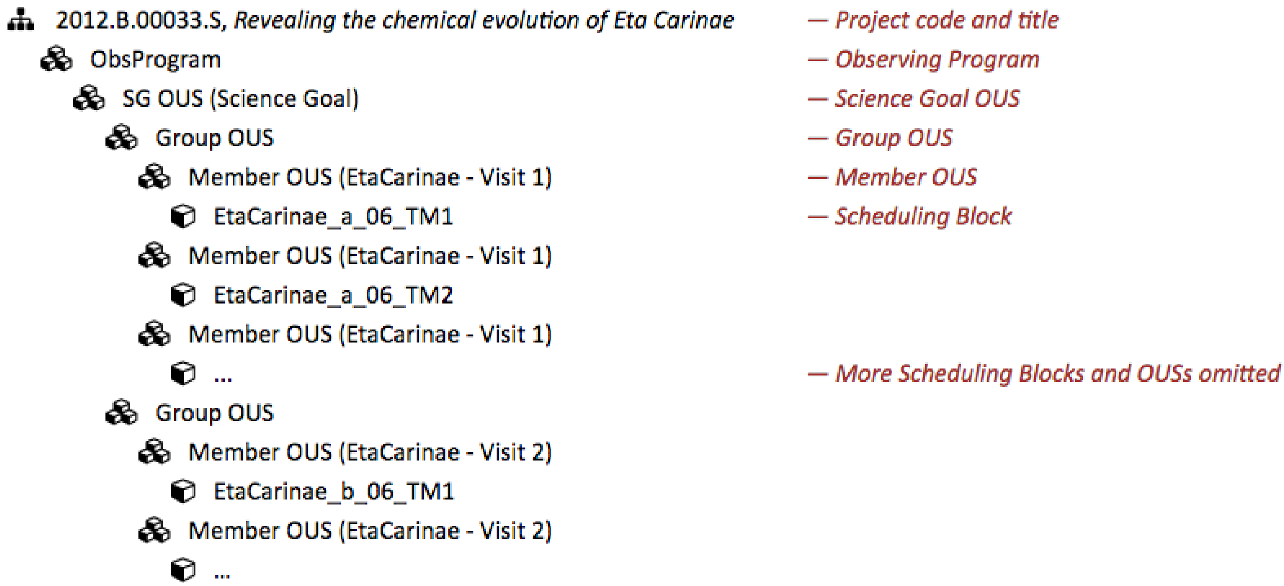

Figure 1. An example ALMA Observing Project (simplified view)

The highest level in the hierarchy is the project itself, identified by a code like *2012.B.00033.S*; the project's title is given by the original submitted observing proposal; in this case *Revealing the chemical evolution…*

The Observing Program includes one or more hierarchies of *ObsUnitSets* (OUSs), playing an important role during the Pipeline reduction and QA phases. The top-level OUSs correspond to the proposal's *Science Goals* (SOUSs); they include one or more *Group* OUSs (GOUSs), in turn containing one or more *Member* OUSs (MOUSs). Each MOUS, finally, contains one or more *Scheduling Blocks* (SBs), the smallest schedulable unit in the hierarchy.

QA0 is performed examining calibration data related to an observation; that is, to an *ExecBlock* (EB).

Finally, data reduction and QA2-QA3 are performed at the OUS level, processing and evaluating the raw data generated by all SB observations (that is, all contained EBs).

1.2 The life-cycle of ALMA entities

ALMA entities like SBs and OUSs follow a well-defined life-cycle, from creation to their end state. Life-cycles can be represented as state diagrams, every state transition being the result of an operation performed on that entity. For instance, an OUS is in the *FullyObserved* state when all its contained SBs have been observed, and becomes *ReadyForProcessing* when all generated raw data has been transferred to the ALMA Archive in Santiago. If data was taken in a mode that's currently supported by the Pipeline the OUS will then transition to the *ReadyForPLCalibration* state, otherwise to *ReadyForManualCalibration*, and so on.

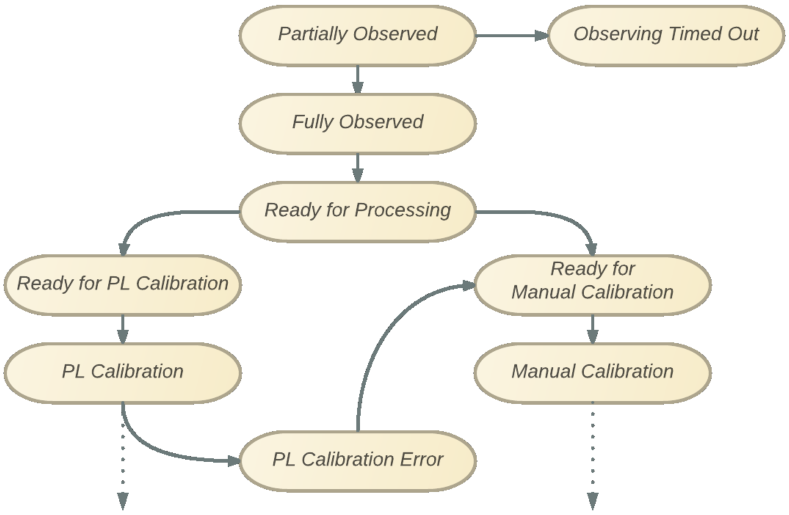

Fig. 2. Life-cycle of ObsUnitSets (partial view)

1.3 Databases

A complex network of databases support ALMA operations and QA in particular. A complete description of the storage facilities is beyond the scope of this paper; we'll only mention here that the range of persistence technologies include relational databases, XML storage, file repositories (*NGAS*, see [2]) and so-called "NoSQL" databases (*Cassandra*, see `www.planetcassandra.org/what-is-apache-cassandra`).

2. QUALITY ASSURANCE 0

QA0 targets *atmospheric conditions* like weather parameters, WVR measurements, system temperature, sky opacity, phase fluctuations and total power levels; *Antenna issues* like position, pointing, focus, target tracking; and finally *Receiver and Correlator parameters* like bandpass, image rejection ratio (sideband ratio), receiver temperature, ALMA Calibration Device output, delay and output signal level.

The QA0 procedure is carried out by the Astronomer on Duty (AoD). Each execution of as SB produces an EB, which is then stored in the Archive: both binary data and metadata in XML format. The AoD checks a number of parameters using the AQUA Web, validating the EB or flagging it as "QA0 Fail". The process can be listed as follows:

- Check of the integrity of the EB: metadata, all calibration intent (Bandpass, Flux, Phase)
- Check of the pointing and focus of the antennas
- Check of the science target: position, mosaicing, time integration
- Check of the Atmosphere quality: PWV, system and receiver temperature, according to fixed criteria
- Check of the Phase quality (See Fig. XX): one important criterion is the phase RMS (after correction of WVR)

- A summary page is produced, displaying information about EB content, estimated vs. expected EB execution fraction and any warnings about the execution
- Finally, a QA0 Flag is set to either *Pass*, *Fail*, or *SemiPass*

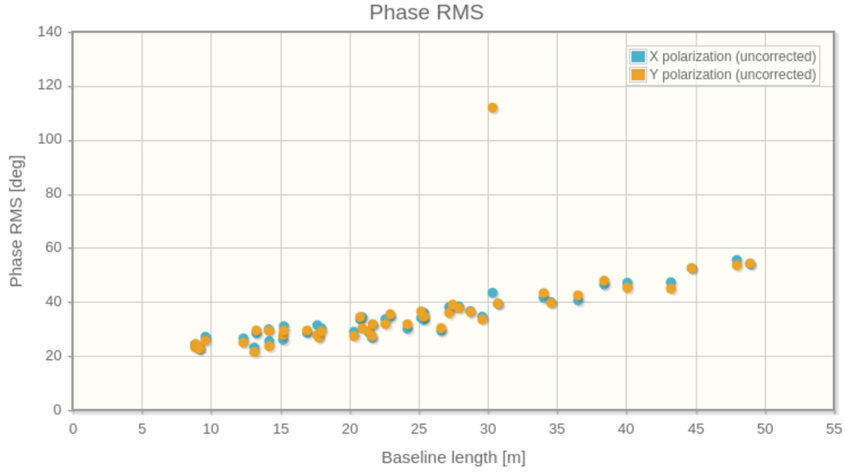

Fig. 3. AQUA plot of the phase RMS for a Phase Calibrator in one of the basebands; an outlier is clearly visible in YY polarization

Some of the parameters for the QA0 are in fact checked in real-time during the observation, using the Quicklook tool (see below). AQUA uses as well external tools (e.g. *aos-check.py*) called via a Web service infrastructure.

2.1 QuickLook

QuickLook is an application running on the operator's console and dedicated to show calibration information in near real-time, that is, as an observation proceeds. It processes data generated by the *Telescope Calibration* module (TelCal) and presents it in easy-to-read graphical format.

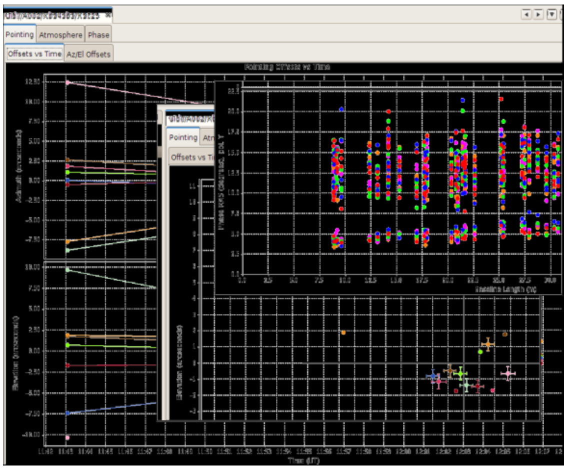

Figure 4. Some of the information presented by QuickLook

A *ScanView* pane is refreshed whenever a *ScanReducedEvent* is received, indicating that a reduction of the latest scan completed. The pane presents scans in the context of their EB, and allows the user to display plots related to a single scan or a complete EB (that is, all scans of that EB). EB plots are updated automatically when a new event is received.

A filter allows to select restrict the plots to data specific to individual antennas or baselines.

2.2 Scripts

A number of ad-hoc scripts have been written in the course of time to extract and analyze Archive information, or check the consistency and correctness of the data itself. They are typically written by staff scientists (as opposed to software engineers) and subject to change as circumstances require it. To isolate ALMA's "official" applications from rapidly changing external software, the scripts run in as Web services in a dedicated environment.

Scripts like *AOS Check* and the *metadata checker* typically process the metadata generated for an EB.

For instance, *AOS Check* is a Python script analyzing the quality of the data produced by each antenna. It checks the phase quality after correction of the WVR and the system temperature, and produces a list of flags and an estimate of the final QA0 status; for instance

```
=== QA0 summary for id___A002_Xb21386_X195 ===
Usable antennas: 18
Phase rms (Antenna,phaseCal): 0.0 deg (=0.1um)
PWV: 0.81 mm   No online WVR-corrected data available: assuming correction factor of 1.41378396845
Phase cal: 0522-364
Flux: 1353.06  Jy
Number of cycles of science/phaseCal: 99
Band observed: 3
Highest recommended: 8-8
QA0 PASS (no significant problems)
```

2.3 AQUA, QA0 tab

The *ALMA Quality Assurance* tool (AQUA) is a Web-based application supporting QA activities. It became operational in Cycle 3 (Oct 2015) allowing a centralized entry point for QA0 information. Support for QA2 is being added and will integrate with the imaging Pipeline when it becomes available in Cycle 4 (Oct 2016). Features related to the different QA tiers are grouped in separate *tabs*.

The left side of the QA0 tab lists EBs (that is, SB observations); by default, only those with no QA0 status are shown, most recent first, producing a kind of "to-do" list:

Project	SB Name	Duration	Start Time	QA0 State	EB UID
2015.1.00631.S	sn1987a_c_06_TE	1.9h	20151102-03:00	Unset	uid://A002/Xac5575/Xa01b
2015.1.01384.S	Ceres_b_06_TE	1.7h	20151102-00:39	Unset	uid://A002/Xac5575/X9bb6
2015.1.01384.S	Ceres_b_06_TE	1.7h	20151101-22:45	Unset	uid://A002/Xac5575/X9911
2015.1.00502.S	WISE_J18_a_06_TE	1.4h	20151101-20:40	Unset	uid://A002/Xac5575/X944c

Figure 5. The list of EBs in the QA0 tab of AQUA

Selecting an EB in the list leads to displaying a large information set about that EB being shown in the right-hand pane. A summary execution pane gives all key information, including whether the observation was successful or had to be interrupted, what the weather was like at the time (water vapor, wind speed, humidity, atm. pressure), the observed frequency and whether any PI-defined timing constraints were violated (for instance, "*Time critical project* *observed outside the correct window, 1.07 days off*"). The time spent on all science and calibrator sources is also indicated.

Calibration information is shown in graphical form in right-hand pane tabs dedicated to Atmosphere, Phase and Pointing calibration; Source (for mosaics) and Baseline coverage are available as well.

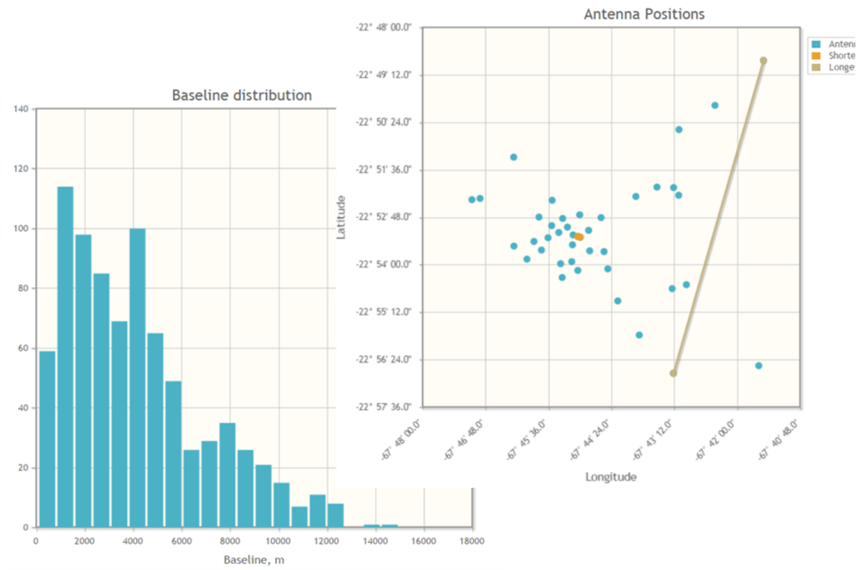

Figure 6. Antenna and Baseline coverage plots in the QA0 tab of AQUA

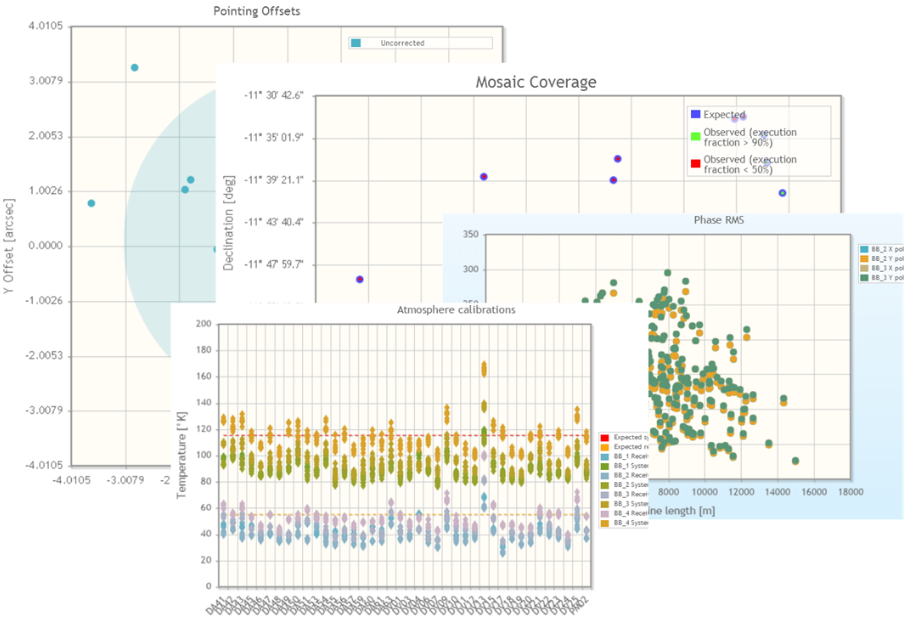

Figure 7. More AQUA QA0 calibration plots: Pointing Offsets, Source Coverage, Phase and Atmosphere calibration (clockwise from top)

That extensive information set includes all the QA astronomer needs to know to assign one of the foreseen QA0 scores to an observation: *Pass*, *Fail* and *SemiPass*. The latter applies to EBs whose data can be used for science analysis but do not reach the minimum requested to be processed for QA2; for instance, observing time is 20% or less of what was required, less than 50% of the required antennas were available, etc**.**

A simple pop-up window allows the operator to enter the final score and comment; the latter is pre-initialized with the output of some of the scripts described above.

3. QUALITY ASSURANCE 1

QA1 targets *Array correction* (baseline and delay); *Antenna calibration* (pointing, beam pattern, main beam efficiency, surface roughness or deformation); and *monitoring of Calibrator flux* (solar system objects and quasars).

QA1 tackles array performance parameters that change slowly, over timescales longer than a week. They are measured by the AoD by executing "standard" SBs created as per the Calibration Plan. QA1 parameters are usually measured as "Observatory Tasks" at predefined periods during the month, or when detecting significant performance deterioration during operations — measuring those parameters is currently performed by running ad-hoc software packages (Python scripts, Tpoint, etc). In the near future most of those packages will integrated within Telescope Calibration and/or CASA thanks to a joint effort of the AoDs and the Data Management Group astronomers; their output will include all the information relevant to the specific measurement and a summary section than can be read by AQUA for further processing.

The main QA1 parameters are stored in the Telescope Configuration and Monitoring Database (TMCDB), broken down by array (12m and 7m). The database can be modified directly through a dedicated GUI (*TMCDB Explorer*) or programmatically.
The most critical parameters will also be displayed on the main operator's console.

QA1 checks are performed following a predefined schedule as part of science operations, and analyzed by the ALMA staff. The results of the analysis are reported to the PI as a summary of the array behavior. Note that all QA0 and QA1 data is processed by Science Operations to perform trending analysis on the different areas of the ALMA system (pointing, focus, antenna position, delay, weather condition, antenna efficiency, observation efficiency, imaging quality, calibrators, etc)

3.1 AOS Monitor

AOS Monitor is a software tool that characterises the performance of all ALMA antennas every day, utilising all observations made at the ALMA Operations Site (AOS), where the array is located. This facilitates assessing the status of antennas, and helps in the detection and solving of the problems. The goal is enable improving the performance of the ALMA array every day.

At the moment, the monitor is limited to the analysis of the pointing and focus calibrations in all ALMA bands, checking the models in the TMCDB and proposing new models. In the future, this tool will include new capabilities helping in the issues associated with tiltmeters, delays, antenna motions, primary surfaces and aperture efficiency.

Param	TMCDB	New Fit	Expected	Difference	Error	Diff/Error
IA	+607.790	+616.118	--	+8.328	+0.638	+13.056
CA	+125.770	+120.471	130.0+/-5	-5.299	+0.409	+12.965
IE	+81.010	+79.464	--	-1.546	+0.409	+3.784
AN	+12.930	+13.218	--	+0.288	+0.472	+0.609
AW	+1.200	+1.262	--	+0.062	+0.444	+0.139

Fig. 8.1. Example of pointing analysis with AOS Monitor

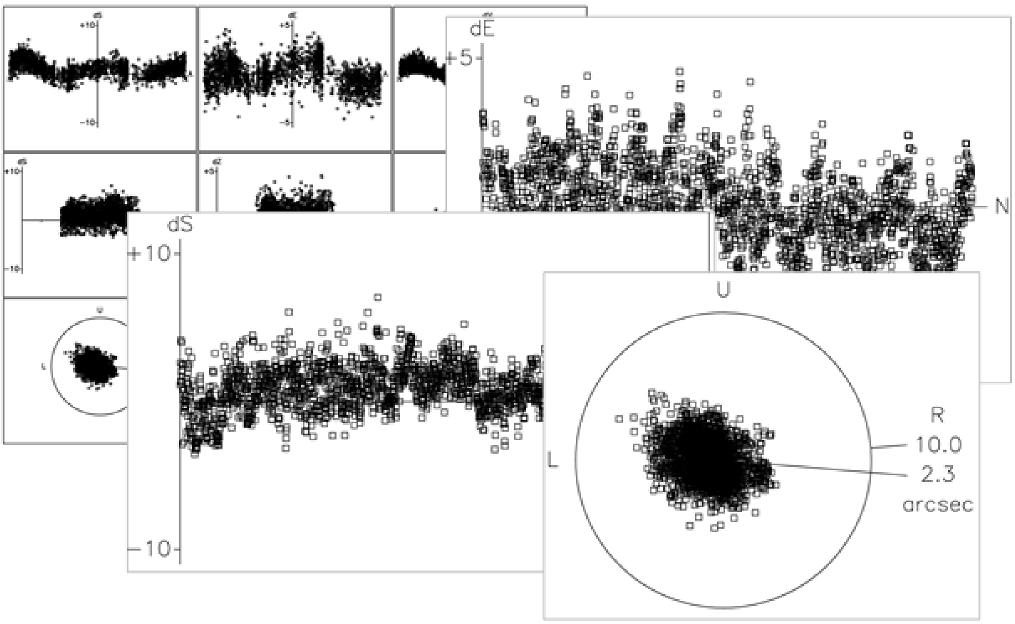

Fig. 8.2. Example of pointing analysis with AOS Monitor

4. QUALITY ASSURANCE 2

QA2 constitutes the final quality assurance step before the data are released to the PI. It is performed on the fully calibrated and imaged data products available either from the ALMA pipeline or manual processing. In Cycle 3, some 75% of ALMA data are calibrated by the pipeline at JAO and subsequently imaged manually at the ARCs, while the remaining 25% are processed completely manually. In the future, the pipeline is expected to fully reduce at least 80% of the ALMA QA0 Pass data.

Many aspects of QA2 were covered in [1].

While QA0 pertains to the EB level and is designed to check that the data are useful and calibratable, QA2 is performed at the OUS level and assures the data products meet the specifications requested by the PI in their proposal. The checks currently focus on the sensitivity achieved for a given frequency and bandwidth as well as the angular resolution (AR) of the image. The recoverable largest angular scale (LAS) requested by the PI is not specifically checked during QA2, but accommodated from the outset by adding MOUSes containing SBs scheduled for more compact configurations in the same GOUS if necessary. In Cycle 3, the QA2 sensitivity and AR checks are applied only to the MOUS observed with the most extended configuration necessary, while the SBs coming from complimentary more compact arrays (including the Atacama Compact Array, also known as the Morita Array) receive only basic checks such as that the expected time on source has been reached. Once the procedure for making and assessing combined images has been finalised it is expected that a final QA2 evaluation will be carried out at the GOUS as well as the MOUS level.

QA2 targets *calibration measures* like WVR, Tsys and antenna position, data flagging (atmospheric absorption, bad visibility), bandpass and flux calibration, phase stability; imaging part is reconstructing the image from the calibrated visibilities and measures like deconvolution by CLEAN method, equalization and concatenation among EBs, setting channel ranges for science targets, and image and grid size; and finally *products* (images and scripts).

QA2 is performed ALMA's Regional support Centers by ARCs in EA, EU, and NA, as well as in Chile (pipeline reduction).

4.1 WebLog

The WebLog is a set of HTML pages giving a summary of how the calibration proceeded, and is included in all ALMA deliveries. It can be viewed using most standard Web browsers, providing a portable familiar interface to a rich hierarchical repository. The WebLog aims to provide a quick overview of the datasets as well as give methods for exploring the pipeline processing in detail. Most pages give an initial single "representative" view of a particular step (e.g. calibration or flagging); users can then drill down to a more detailed view of all the plots associated with that step. These plots can easily be filtered by a combination of outlier, antenna and spectral window criteria. In many places histograms of appropriate metrics are displayed, in modern web browsers it is possible to draw boxes on multiple histograms to select the plots associated with those data points. All steps also undergo an automated scoring to give an "at a glance" indication of any trouble points. Throughout the WebLog, links and enlargeable thumbnails are used to provide an interactive environment allowing intuitive data and processing exploration.

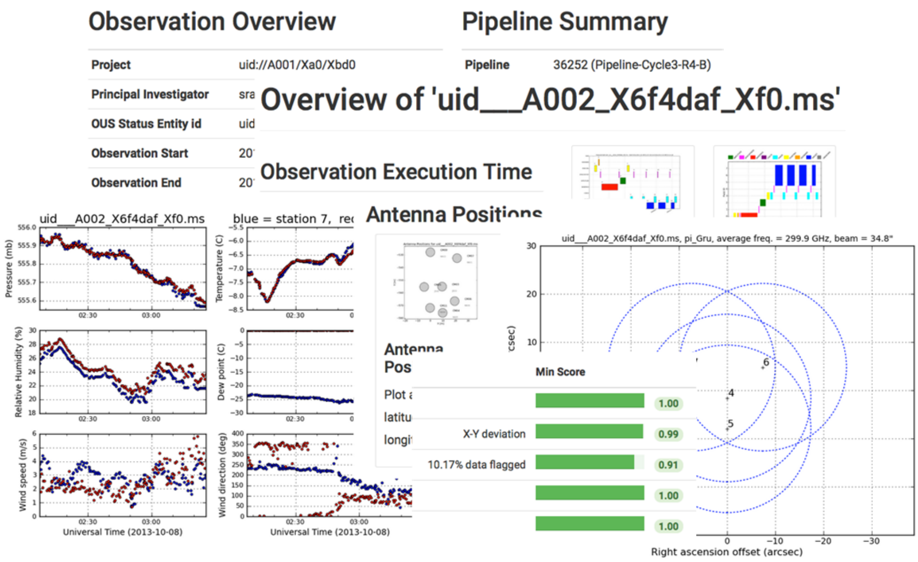

Figure 9. Some of the information presented by the WebLog

4.2 AQUA, QA2 tab

Support for QA2 is currently being built into AQUA, with the goal of integrating that activity in the course of Cycle 4. The paradigm will be the same as QA0: QA staff will be presented with a "to-do" list of OUSs awaiting QA2, and selecting one will display all relevant detailed information. As indicated above, QA2 information is generated by the ALMA pipeline, including a set of QA scores related to the Pipeline calibration stages: data flagging, flux, TSys and bandpass calibrations, and so on. For detailed info, direct access to the Weblog is also available.

Once the imaging pipeline is fully commissioned, the most relevant images and imaging parameters such as the achieved sensitivity, AR and LAS will also be accessible from within the AQUA interface together with the PI requested values. The plan is for the QA2 assessment to become increasingly automated as operations become more and more routine.

5. QUALITY ASSURANCE 3

ALMA recommends that PIs perform a post-reduction evaluation of the data products delivered to them; the process is called QA3. Any problems that they find to is reported to Contact Scientist via the ALMA Helpdesk. (The QA3 process may be also triggered by contact scientists themselves or other ARC personnel.) Problems with the data products may reflect an underlying problem with the original data, observing procedure or calibration.

Problem resolution begins with retrieval of the data from the archive to evaluate the nature of the problem, including an assessment of whether the problem applies to a specific dataset or data taken under similar set-ups and conditions is also affected. Depending on the nature and gravity of the problem the ARC may request involvement of the Observatory; ultimately, it may require repeating the observation(s).

QA3 is currently a very ad-hoc activity and no dedicated software tools were developed to support it.

6. CONCLUSIONS

ALMA's success with the community is due in large part to the quality of the data it delivers. Ensuring that raw and reduced data maintain ALMA's rigorous quality standards requires the application of focused policies and instruments at all stages of data generation, processing and distribution. Over the years, ALMA Science Operations have worked in close collaboration with the Engineering and Computing teams to develop and refine those policies and the procedures to implement them, coupled with the software tools and computing infrastructure required to support the entire process.

QA at ALMA is a constantly evolving and still manpower-intensive process; the goal of supporting a wider community by providing science-ready data imposes an additional burden to the staff. The efficiency and productivity of the QA team needs to increase and it's important that processes get automated and tools take over the most repetitive tasks. While quality assessment is ultimately a human responsibility, the process of forming an opinion about a data set or product can be facilitated greatly by providing scientists with a well-defined set of summary information as well as software tools to distribute work and collect the results.

The challenge is to maintain ALMA's level of excellence while providing for a sustainable workload and manpower budget.

7. ACKNOWLEDGEMENTS

The authors would like to thank Christian Lopez for his contributions to this paper.

REFERENCES

1. Petry, D, Vila-Vilaro, B., Villard, E., Komugi, S., and Schnee, S., "ALMA service data analysis and level 2 quality assurance with CASA", SPIE 9152, Software and Cyberinfrastructure for Astronomy III, 91520J (July 18, 2014); doi:10.1117/12.2054715.
2. J. Knudstrup, A. Wicenec, and S. Johnston, 2002, NG/AMS: The Next Generation Archive Management System. in ASP Conf. Ser. 281, ADASS XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP)[O-17]