DOI: 10.18727/0722-6691/5206

ALMA Data Quality Assurance and the Products it Delivers — The Contribution of the European ARC

Dirk Petry[1]
Thomas Stanke[1]
Andy Biggs[1]
María Díaz Trigo[1]
Fabrizia Guglielmetti[1]
Evanthia Hatziminaoglou[1]
Eelco van Kampen[1]
Luke Maud[1]
Anna Miotello[1]
Gergö Popping[1]
Suzanna Randall[1]
Felix Stoehr[1]
Martin Zwaan[1]

[1] ESO

From its inception, the Atacama Large Millimeter/submillimeter Array (ALMA) was intended to be accessible to all astronomers, including those who are more used to carrying out their research at other wavelengths. Since the beginning of science observations in September 2011, ALMA has therefore applied a comprehensive Quality Assurance (QA) process to the observed data before delivering them to the principal investigators (PIs). This huge investment, unique for a ground-based (non-survey) observatory of this calibre, results in fully calibrated datasets as well as high-quality images that allow the PIs to assess the quality of their data upon delivery and that provide an advanced starting point for the scientific analysis. In this article we provide a summary of the purpose and status of ALMA QA, a brief description of the QA process and the resulting ALMA data products, and a discussion of how the ALMA user profits from them.

ALMA observations and data processing

The considerable effort going into ALMA QA is provided by staff at all four main locations of the ALMA project: the Joint ALMA Observatory (JAO) in Chile, and the three ALMA Regional Centres (ARCs) in East Asia, North America and Europe[a]. In Europe, the work is done by the ARC staff at ESO in Garching as well as staff in the European ARC network. For a description of the European ARC network see Hatziminaoglou et al. (2015).

In their observing programmes, ALMA PIs do not propose for observing time but for a particular sensitivity at a range of angular scales to achieve their science goal (the required strategy can also include time constraints). The observations and their scheduling are based on Scheduling Blocks (SBs) — the observing units defined within each science goal. A SB is a plan for a complete set of calibration and science target observing scans (see Figure 1). Several different SBs may be needed to define the observations necessary to achieve one science goal: for example, observations from different array configurations. The total duration of a SB execution can be up to two hours. If longer exposures are required on the target, the same SB is executed several times. The details of the SB setup and the number of required SB executions (the so-called Execution Blocks, EBs) are estimated based on the ALMA Sensitivity Calculator and the parameters provided by the PI in the proposal using the ALMA Observing Tool (OT).

Level 0 quality assurance (QA0) takes place at the telescope shortly after the completion of the execution of a SB. It aims to catch obvious problems with the observation at an early stage and ensure that the data collected during this particular execution are useful to achieving the science goal. A number of diagnostics are created in order to permit a basic check of the correct setup of the included antennas and their receivers and to quantify the overall stability of the atmosphere; in addition, QA0 verifies that the flux calibrator used in the observation has a recent flux measurement.

If the execution has achieved a significant fraction of the intended science observation, it can move on to the following QA stage and is declared “QA0 Pass”. It is stored in the ALMA Science Archive and is replicated from the JAO to archive copies at the three ARCs. A SemiPass or Fail state indicates a partially useful execution or an execution that cannot be calibrated at all, respectively, and the execution is repeated. The contribution of each EB is measured as a so-called “execution fraction”, factoring in the observing conditions and number of available antennas. The execution fraction can be larger than unity if the conditions are better than expected. Once the sum of the EB execution fractions is equal to the planned number of EBs, the SB is considered fully observed.

Level 2 quality assurance (QA2) takes place once an SB is fully observed. Note that there is also a level 1 QA which concerns the longer-term monitoring of observatory parameters, but this is not discussed in this article. The full set of executions of an SB is called a Member Observation Unit Set (MOUS). The MOUS

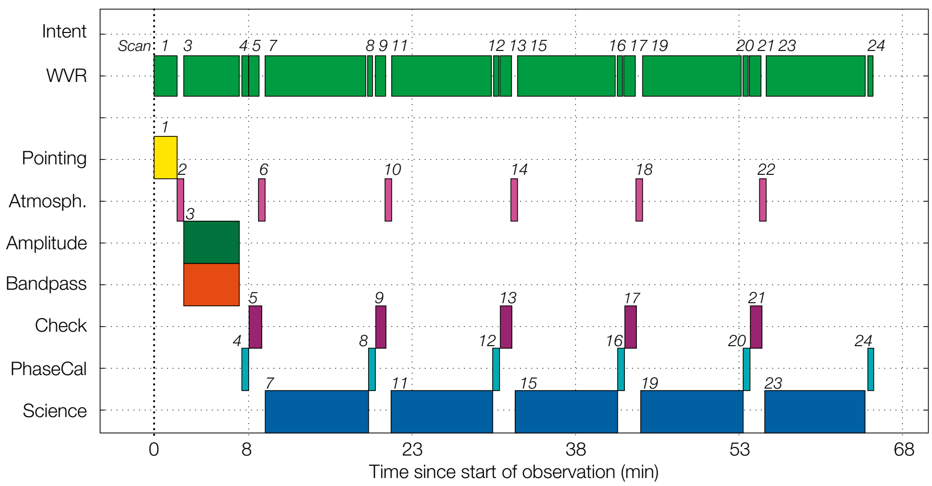

Figure 1. Typical structure of an ALMA Scheduling Block (SB) describing the interleaved observation of a single science target (dark blue), a phase calibrator (light blue), a check source (dark purple) to assess the quality of the phase transfer, and other calibrators. Also shown are the calibrator scans during which water vapour radiometer data are taken (light green) and the receiver response and the atmospheric opacity along the line of sight are measured (light purple).

is the smallest data entity that is delivered to the PI. Thus, QA2 operates on MOUSs.

While QA0 takes place at the observatory, QA2 is decentralised. Up to the end of ALMA Observing Cycle 1 (September 2014), essentially all QA2 processing took place at the three ARCs. After that, the capabilities of the JAO were extended, and QA2 processing was gradually moved there.

QA2 processing is computing intensive. In the typical case of an interferometric dataset[b], it consists of three steps: the a priori calibration, the calibration of flux, bandpass, and phase, and the imaging.
– The a priori calibration applies initial phase and intensity corrections based on the water vapour column density, the receiver response and the atmospheric opacity measured during the observations, applies small positional updates of the antennas, and runs an initial flagging of bad data.
– The flux, bandpass, and phase calibration makes use of dedicated observations of calibrator targets that are integrated into the SB, interleaved with the observations of the science target(s) (see Figure 1). The flux calibration bootstraps the flux scale to an absolute scale by comparing to recent observations of a known quasar calibrated against a Solar System object with well-known emission. The bandpass calibration corrects for the spectral response of the ALMA receivers by observing a bright quasar with a featureless, non-thermal spectrum. Finally, the phase calibration derives a correction for the atmospheric phase fluctuations from the observation of another bright quasar at a small angular distance from the target. All these steps may require iteration if bad data is found during the processing that needs to be flagged. For full-polarisation observations, another calibration step is needed which applies the information gained from the observations of a polarisation calibrator over a sufficiently large parallactic angle range.
– Finally, imaging is carried out on the calibrated data for all spectral windows and for as many of the science targets as possible with the available computing resources. While the calibration of a MOUS (i.e., all the executions of an SB) typically takes between 1 and 24 hours, the imaging of (spectrally and/or spatially) high-resolution observations can take between half a day and several weeks of computing time. For line observations, the imaging process also includes the determination and subtraction of the continuum emission before imaging the line cubes.

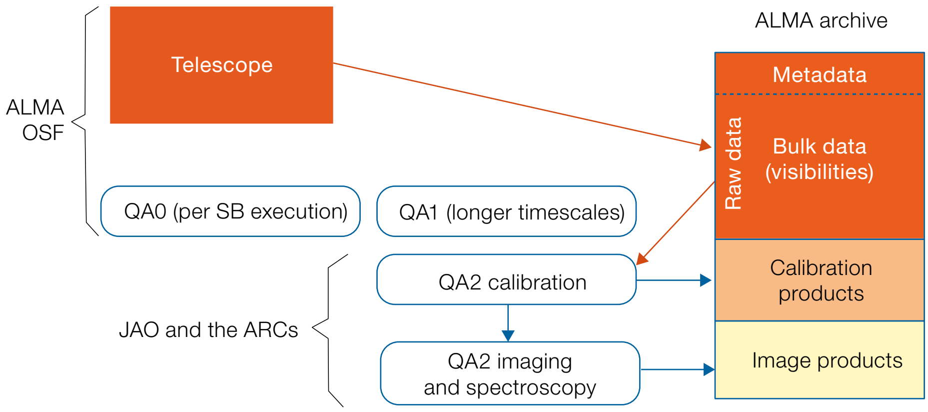

Figure 2. Simplified ALMA data flow. While the raw data go straight to the Archive after QA0, the calibration and science products are generated during QA2, which takes place at the JAO and the ARCs. Note that the calibration products do not include the calibrated visibilities in order to save Archive storage space.

The QA2 processing is followed by an assessment of whether the sensitivity in the representative spectral range and the achieved angular resolution match the PI's requirements as recorded in the proposal. If they do, the dataset is declared "QA2 Pass" and the calibration and image products are ingested into the Archive. The successful ingestion is followed by an email notification to the PI to advertise the availability of their data. This is called the delivery and starts the proprietary time of one year. Access to proprietary raw data is also possible upon request before the official delivery, but this request immediately starts the clock for the proprietary time and comes without any user support from the ARCs with the calibration.

If a dataset does not pass QA2 immediately (< 10% of the cases), re-observation of the SB followed by new QA2 processing is attempted until the project times out (for details see Remijan et al., 2019). This is another reason why QA speed is of the essence.

In order to save storage space, calibrated visibilities and single-dish data are not stored in the Archive and are not part of the data delivery. Instead, all products necessary for the calibration are provided. The user has to restore the calibrated data by running a script contained in the delivery package on the raw data. Since October 2019, the calibrated data can also be requested for download via a dedicated service offered in Europe[1]. Before the end of the proprietary time, the service is of course only available to PIs and data delegates.

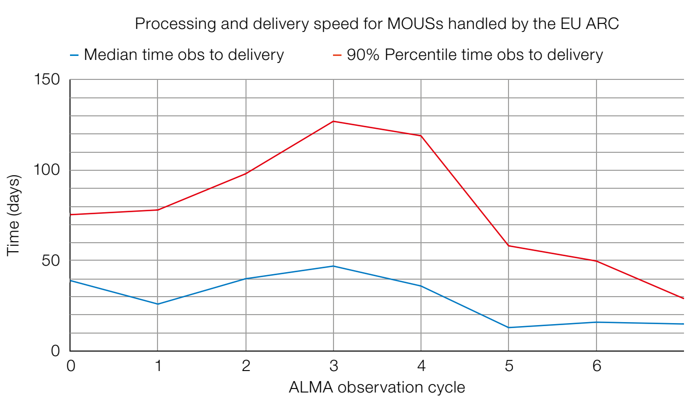

Figure 3. The processing and delivery performance of the QA work at the EU ARC. After some overload and technical problems in Cycles 3 and 4, 90% of the deliveries now take place within one month after observation.

In summary, QA2 processing of the full set of executions of an SB is a complete, high-quality, science-ready calibration of the data followed by detailed imaging with the aim of providing the PI (and later archival researchers) with a set of images that make it possible to inspect the degree to which the science goals were achieved.

A Level 3 quality assurance (QA3) process has been put in place to handle any errors that are discovered after the official data delivery. If the problem is discovered by a user, they may file an ALMA helpdesk ticket. If confirmed, a detailed investigation is started. The outcome ranges from the addition of a note to the QA2 report to a correction of the data products, followed by a re-ingestion into the Archive, or if necessary and possible, even a re-observation of the SB. In ten years of ALMA observations, QA3 cases that affected large portions of ALMA data have happened only a few times. In all cases, the observatory strove to keep the users informed about the implications of the problem for their data and about the progress of the correction. Obviously, every such campaign implies a high additional load on the QA staff and the computing facilities.

From semi-manual to pipeline processing

Full automation of QA2 processing was always planned but cannot be achieved without a period of semi-manual processing until the data are fully characterised. In ALMA observing Cycle 1 (September 2014), all QA2 processing was carried out exclusively semi-manually by analysts using the Common Astronomy Software Applications package (CASA; McMullin et al., 2007; Petry, 2012; Emonts et al., 2019) and the Calibration Script Generator, a tool that evaluates ALMA raw data and generates a draft calibration script (see Petry et al., 2014). Based on the experience gathered with this prototype pipeline, a fully automated pipeline was developed and gradually deployed cycle by cycle for more and more of the different ALMA observing modes. In particular, the heuristics for automated flagging and calibration were deployed first. The capability to automatically image the data followed in the middle of Cycle 4.

Like the semi-manual analysis, the ALMA pipeline (ALMA pipeline Team, 2019) is based on the CASA package. It is distributed together with CASA and thus also published to ALMA users. For QA2, it runs at the JAO and the ARCs, controlled by additional infrastructure software. For each run, the pipeline creates a set of diagnostic plots and tables (wrapped in a system of html pages), called the weblog. This weblog is then reviewed manually in order to judge whether the pipeline run was successful, and the observing parameters were met.

Today, the ALMA pipeline is capable of processing most ALMA data without much human intervention other than reviewing the weblog. Only about 10% of the datasets still require semi-manual processing by analysts. For projects supported by the European ARC, these analysts are based at the European ARC network including the ESO ARC department. Similar efforts are ongoing in the other ALMA regions, and of course at the JAO.

During the first observing cycles, the delays between data taking and data delivery were significant. Today, thanks to an enormous effort at the JAO and the three ARCs, a complete redesign of the data flow system, and the increased usability of the ALMA pipeline, QA2 processing has been accelerated to the point where 90% of the deliveries take place within one month after the observation (the median is 2 weeks).

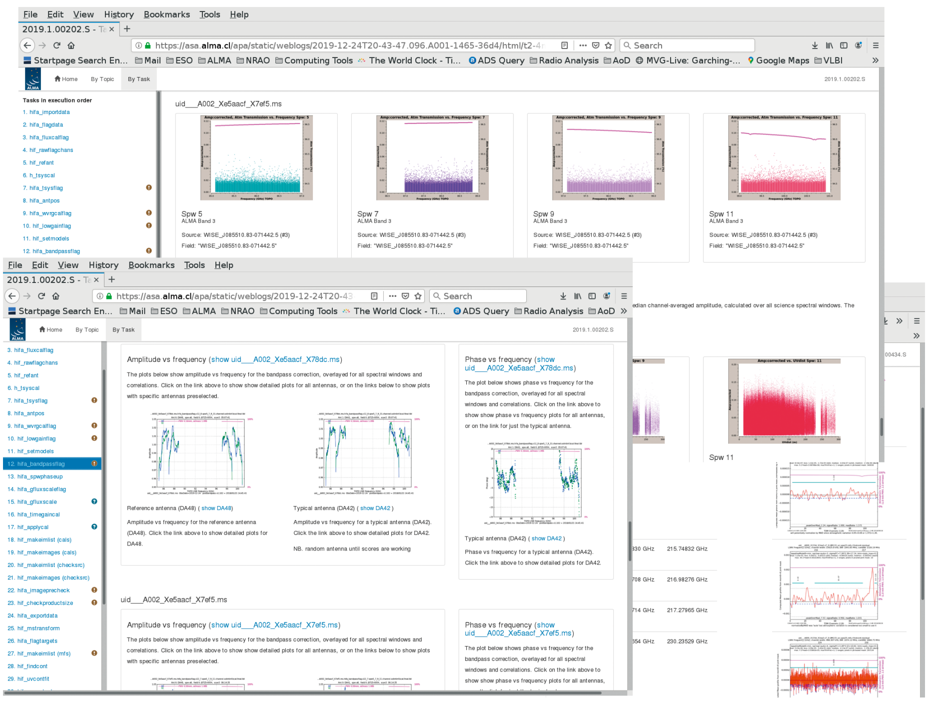

Why is QA2 necessary and what does it provide to the user?

ALMA is a large, complex project that needs to perform detailed bookkeeping and monitoring to make sure that the observatory performs reliably and to specification. The QA effort is part of this process, providing a vital link in the chain from the proposal of an observation to the publication of its scientific results. Experience has shown that subtle problems are often only noticed when trying to extract scientific information from the data. Furthermore, the only precise method of determining the achieved sensitivity is to fully calibrate the data and create an image of the spectral range that is of interest to the PI. Finally, to optimise the extraction of scientific results from its archive, ALMA would like to provide

Figure 4. A sample of ALMA pipeline weblog pages showing different diagnostic plots.

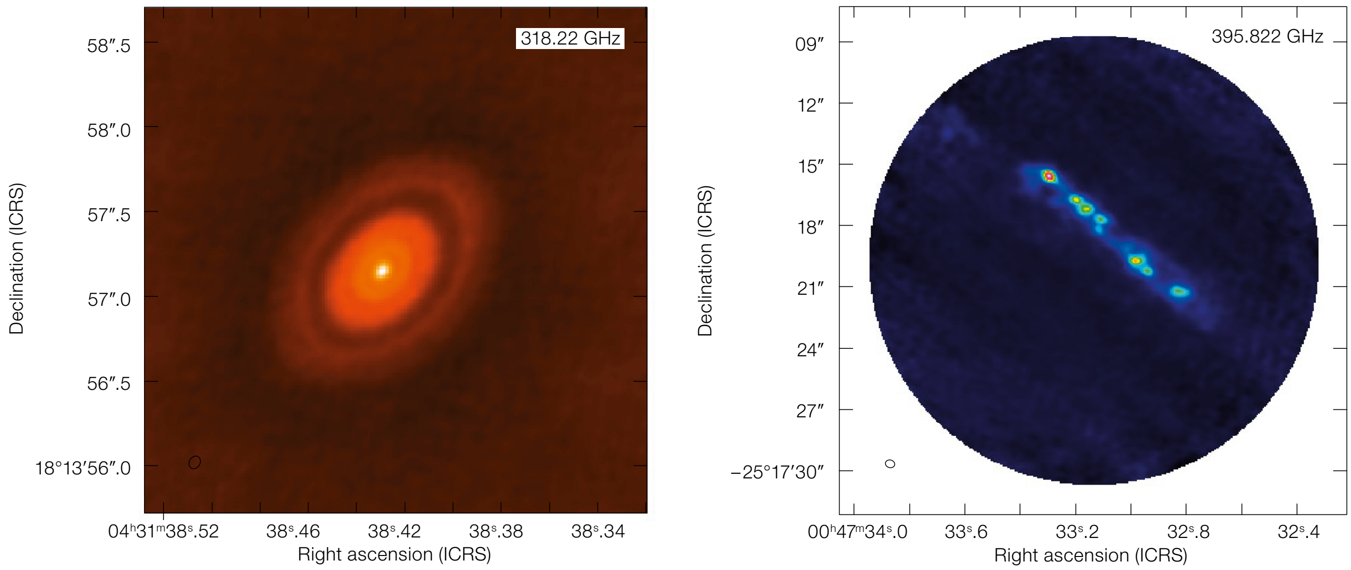

Figure 5. Two examples of science target images created during ALMA QA2 and delivered to the corresponding PIs.

valid data products to facilitate archival research. The QA2 process combines all three of these processes and presents ALMA users with two major advantages:
1. An advanced starting point for their own data analysis which helps PIs to publish sooner and better.
2. A valuable, comprehensive, and most of all, homogeneous set of high-quality data products for more efficient archival research.

As a result of the QA2 effort and in addition to the raw data, PIs receive for each individual dataset (MOUS) a standard package with a wealth of important information:
– The science-grade calibration information. The user can take their raw data and calibrate it in a reproducible way by running the CASA script “scriptForPI.py” included in the package. Each data set is processed with a specific CASA version. The user is required to run the scriptForPI under that CASA version. Alternatively, European users can request the generation of the calibrated data from the European ARC via the helpdesk.
– A detailed summary of the QA stages in the form of QA0 and QA2 reports.
– All calibration and imaging diagnostic plots (for pipeline-processed data in the form of the weblog, see previous section). This permits the user to assess nearly all details of the data properties without having to touch the raw data or even starting CASA. The weblog is simply opened in a web browser (see Figure 4).
– The QA2 imaging products as FITS files (following FITS standard 3.0). This package contains aggregate bandwidth, continuum, and line-cube images depending on the science goal. Calibrator and check source images are also typically included. The completeness of the set is typically better for pipeline-processed datasets since the semi-manual imaging process is slower and needs to save time for quick delivery.

For a detailed description of the ALMA QA2 data products see Petry et al. (2018).

Are ALMA image products science-ready?

As described above, ALMA QA2 is a standardised and semi-automated process that is not meant to cover all of the specific scientific needs of the PIs. Feedback from users indicates that the delivered standard images and cubes are often close to optimal and can be used as a basis for scientific analysis. However, it should be clear to all ALMA users, PIs and archive researchers that the delivered images cannot, in general, be classified as science ready. The user may well have to go back to the calibrated data to optimise the parameters of continuum subtraction and/or imaging for their scientific goal, in order to obtain the image or cube for publication. Extensive help in assessing the standard products and improving on them is provided to all European users by the European ARC network and the helpdesk.

Although ALMA strives to provide high-quality and homogeneous informative imaging products for all datasets, the imaging products from the earlier Cycles 2 to 4 were produced semi-manually and are often only based on a fraction (at least 200 channels of most spectral windows) of the total spectral coverage. Additional Representative Images for Legacy (ARI-L) is an ongoing ALMA development project that aims at increasing the legacy value of the ALMA Science Archive by bringing the reduction level of ALMA data from Cycles 2 to 4 close to that of the more recent cycles, for which the imaging pipeline was used. These re-processed images and cubes are being included in the Archive as value-added products. Future Messenger articles will present the ARI-L project in more detail and also the ALMA Science Archive.

Other value-added data products come from the ALMA Large Programmes. The award of a Large Programme carries with

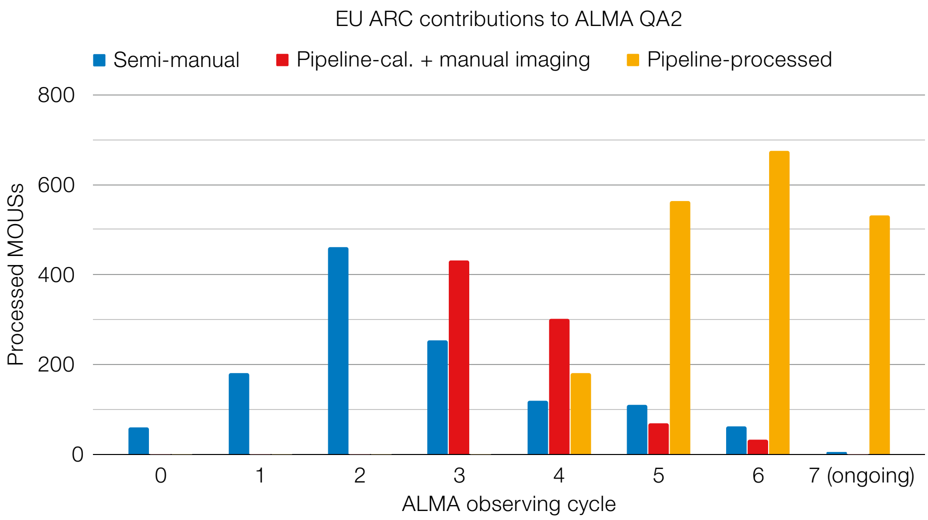

Figure 6. The MOUSs (individual ALMA datasets) QA2-processed and delivered by the European ARC and its nodes since the beginning of ALMA observations. The numbers are presented separately for the three main processing workflows: script-generator-assisted semi-manual processing (blue), pipeline calibration followed by manual imaging (red), and full pipeline processing (yellow). Note that pipeline processing also requires a human review of the results (called the weblog review).

it the responsibility to deliver back to ALMA a set of enhanced data products which supplement the standard ALMA products generated by the ALMA observatory during QA2.

Conclusions and forward look

Since the start of science observations, the ALMA project has put a huge effort into providing science-grade calibrated data and informative image cubes, a first for a large ground-based astronomical observatory. This effort is populating the ALMA Science Archive with homogeneous, high-quality data while making ALMA more accessible to all astronomers, regardless of their scientific background.

ALMA aspires to extend the user support even further: one of the project's longer-term goals is to produce higher level data products such as catalogues or images combining data at different angular resolutions from different SBs. Such products will further enable the use of the facility and its archive by non-expert users, increasing at the same time the scientific impact of the observatory. Making science-ready data products available will shift the focus of the users from the (hardware) limitations and the technicalities related to interferometric data reduction to the scientific exploitation of the data.

Acknowledgements

Over the past seven and a half ALMA observing cycles, the European ARC network has made major contributions to the ALMA QA effort. Scientists from ESO and the ARC network have worked more than 150 seven-day Astronomer-on-Duty shifts at the observatory and performed QA0 assessments as part of them. They have so far served as QA2 analysts for more than 1250 semi-manually calibrated, more than 830 manually imaged, and more than 1940 pipeline-processed datasets from European projects. Furthermore, they have contributed significantly to the development of the CASA package, the ALMA calibration and imaging script generators, and the ALMA pipeline.

We would like to acknowledge all 87 scientists from the EU ARC network who made these contributions: Abhijeet Borkar, Adam Avison, Alison Shan Man, Alvaro Sanchez Monge, Ana Lopez Sepulcre, Andy Biggs, Anita Richards, Anna Miotello, Arek Berlicki, Arturo Mignano, Baobab Liu, Benjamin Magnelli, Carmen Toribio, Christos Karoumpis, Ciriaco Goddi, Ciro Pappalardo, Claudia Lagos, Daniel Harsono, Daniel Tafoya, Devaky Kunneriath, Dirk Petry, Eamon O'Gorman, Edwige Chapillon, Eelco van Kampen, Elisabetta Liuzzo, Eskil Varenius, Fabrizia Guglielmetti, Federica Palla, Federico Lelli, Felipe Alvez, Fernando Morales, Gaelle Dumas, George Bendo, Gergö Popping, Hau-Yu Lu, Hsi-Wei Yen, Hugo Messias, Ilsang Yoon, Israel Matute, Ivan Marti Vidal, Ivica Skokic, Jan Brand, Javier Adrián Rodón, Ka Tat Wong, Kazi Rygl, Ke Wang, Kevin Harrington, Liz Guzman, Liz Humphreys, Lizette Guzman, Luke Maud, Lydia Moser, Marcella Massardi, Maria Diaz Trigo, Markus Schmalzl, Martin Zwaan, Matthias Maercker, Matthieu Bethermin, Miroslav Barta, Nataliia Muratova, Nuria Marcelino, Pamela Klaassen, Paola Andreani, Pavel Jachym, Rebeca Aladro, Rein Warmels, Rosita Paladino, Ruud Visser, Sabine Koenig, Sandra Burkutean, Sebastian Muller, Sergio Martin, Silvio Lorenzoni, Stefanie Muehle, Stefano Facchini, Stergios Amarantidis, Suzanna Randall, Thomas Stanke, Timothy Davis, Tobia Carozzi, Toma Badescu, Vinod Arumugam, Wenjuan Liu, Wouter Vlemmings, Yanett Contreras, Yujin Yang, and Yurii Pidopryhora.

We would also like to acknowledge the work of our colleagues at the Joint ALMA Observatory in Chile who have been working with us on the QA of the European part of the ALMA data since Cycle 0, our colleagues from the North American and the East Asian ARCs for their contribution to ALMA QA and its procedures in general, and our colleagues at the ESO Archive Operations Group who provided indispensable support for data transfers and deliveries. We would like to thank all of them for the great teamwork.

References

ALMA pipeline Team 2019, ALMA Science pipeline User's Guide, Doc. 7.13, ver. 1
Hatziminaoglou, E. et al. 2015, The Messenger, 162, 24
Emonts, B. et al. 2019, ADASS XXIX, ASP Conf. Ser., in prep., arXiv:1912.09437
McMullin, J. P. et al. 2007, in Proc. ADASS XVI, ASP Conf. Ser., 376, 127
Petry, D. 2012, in proc. ADASS XXI, ASP Conf. Ser., 461, 849
Petry, D. et al. 2014, Proc. SPIE, 9152, 91520J
Petry, D. et al. 2018, ALMA QA2 Data Products for Cycle 5, ALMA doc 5.12, ver. 2.0
Remijan, A. et al. 2019, ALMA Cycle 7 Technical Handbook, ALMA doc 7.3, ver 1.1

Links

[1] EU ARC webpage on requesting calibrated data in Europe: https://almascience.eso.org/local-news/requesting-calibrated-measurement-sets-in-europe

Notes

[a] Further details of the ALMA partners and the organisation of this large international project: http://almascience.org.
[b] There are other classes of observations that are not described here but which also undergo QA. These include single-dish SBs and special modes like solar or VLBI observations.