

ALMA Cycle 8 2021 Technical Handbook



www.almascience.org

User Support

For further information or to comment on this document, please contact your regional Helpdesk through the ALMA User Portal at www.almascience.org. Helpdesk tickets will be directed to the appropriate ALMA Regional Center at ESO, NAOJ or NRAO.

Revision History

Cycle	Version	Date	Editors	Remarks
5	1.1	2017-03-21	R. Warmels and A. Remijan	Cycle 5
6	1.0	2018-03-20	R. Warmels	Cycle 6
7	1.0	2019-03-19	A. Remijan	Cycle 7
7	1.1	2019-06-29	A. Remijan	Cycle 7 - Fixed Typo in Table 7.2 Caption
8	1.0	2021-03-01	A. Remijan & P. C. Cortes	Cycle 8

Editor

Paulo C. Cortes - (With a special thanks to Anthony Remijan for passing the flag with a well structured and matured document!)

Authors of this Edition

Anthony Remijan - Chapter 1; Antonio Hales - Chapter 2 and Appendix A; John Carpenter - Chapter 3; Bill Dent - Chapter 4; Seiji Kameno - Chapter 5; Bill Dent and Eric Villard - Chapter 6; Brian Mason - Chapter 7; Baltasar Vila Vilaro and Paulo Cortes - Chapter 8; Andy Biggs - Chapter 9; Ed Fomalont and Paulo Cortes - Chapter 10; Bill Dent and Eric Villard - Chapter 11; Rachel Rosen - Chapter 12; Felix Stoehr - Chapter 13; Kamaljeet Saini - Appendix B

Contributors to this Edition

Tim Bastian, Crystal Brogan, John Carpenter, Chinshin Chang, Geoff Crew, Paul Fisher, Misato Fukagawa, Melissa Hoffman, Chat Hull, Hugo Messias, Todd Hunter, Ruediger Kneissl, Andy Lipnicky, Sergio Martin, Lynn Matthews, Luke Maud, Yusuke Miyamoto, Tony Mroczkowski, Hiroshi Nagai, Kouichiro Nakanishi, Masumi Shimojo, Richard Simon, Jeremy Thorley, MCarmen Toribio, Catarina Ubach, Catherine Vlahakis, Martin Zwaan

In publications, please refer to this document as

Cortes, P. C., Remijan, A., Biggs, A., Dent, B., Carpenter, J., Fomalont, E., Hales, A., Kameno, S., Mason, B., Philips, N., Saini, K., Stoehr, F., Vila Vilaro, B., Villard, E. 2020, ALMA Technical Handbook, ALMA Doc. 8.4, ver. 1.0
ISBN 978-3-923524-66-2

DOI: <https://doi.org/10.5281/zenodo.4511522>

Contents

1	Introduction	9
2	Array Components	11
2.1	The ALMA Telescope	11
2.2	The 12-m Array	11
2.3	The Atacama Compact Array	13
2.3.1	The 7-m Array	13
2.3.2	The TP Array	13
3	Principles and Concepts of Interferometry	15
3.1	Introduction	15
3.2	Single-dish Response	15
3.3	Visibilities and Aperture Synthesis	19
3.4	The Visibility or (u,v) Plane	22
3.5	Fields-of-view and Mosaics	22
3.6	Spatial Filtering	23
3.7	Multi-configuration Observations	27
3.8	Units and Conversions	28
3.9	Further Reading	28
4	Receivers	29
4.1	Local Oscillators and IF Ranges	31
4.2	The Cycle 8 Receivers	32
4.2.1	Band 3 Receiver	33
4.2.2	Band 4 Receiver	36
4.2.3	Band 5 Receiver	39
4.2.4	Band 6 Receiver	42
4.2.5	Band 7 Receiver	45
4.2.6	Band 8 Receiver	48
4.2.7	Band 9 Receiver	51
4.2.8	Band 10 Receiver	54
5	The Correlators	57
5.1	The 64-input Correlator	58
5.1.1	TDM mode	58
5.1.2	FDM mode	59
5.1.3	Correlation and Realtime Processing	60
5.2	The ACA Correlator	60
5.3	Digitizers	62
5.4	Offline WVR Correction	63
5.5	Capabilities of the Correlators	64
5.5.1	Polarization	64
5.5.2	Spectral Setup	64

5.5.3	Time Resolution	65
5.5.4	Phase Switching	68
5.6	Practical Performance	69
5.6.1	Sensitivity	69
5.6.2	Linearity	69
5.7	Final data product - the ASDM	70
6	Spectral Setups	71
6.1	Introduction	71
6.2	Frequency Definitions	73
6.3	Spectral Setups	73
6.3.1	Observing Frequencies for Continuum	74
6.3.2	Multiple Spectral Windows in the Same Baseband	76
6.3.3	Spectral Setups for Lines near the Edge of the Bands	77
6.3.4	Spectral Setup in Bands 9 and 10: DSB Considerations and 90-degree Phase Switching	77
6.4	Usable Bandwidth	78
6.5	Spectral Resolution and Channel-averaging	78
6.6	Spectral scans	79
6.7	Spurious Signals	79
6.8	Doppler Setting and Velocity Reference Frames	80
7	Imaging with ALMA	83
7.1	Introduction	83
7.2	Configurations	84
7.3	Shadowing	86
7.4	Angular Resolution and Beam Shape	90
7.4.1	Definition of angular resolution	90
7.4.2	Effects of weighting scheme	90
7.4.3	Other factors affecting beam size and shape	91
7.5	Snapshots	92
7.6	Large Spatial Scale Response	95
7.7	Mosaicing	97
7.8	Multi-array and Multi-configuration Observations	98
7.9	Multi-array and Multi-configuration Imaging	102
8	Observing Modes	107
8.1	Observing Project Structure	107
8.2	Execution Structure	109
8.3	The Observing Process	110
8.3.1	The Source Selection Algorithm	110
8.4	Single Field Interferometry	111
8.5	Pointed Mosaic Observations	112
8.6	Total Power (Single Dish) Observations	112
8.7	Polarization	113
8.7.1	Sessions	117
8.8	Multiple Region Modes	118
8.9	Observation of Ephemeris Objects	118
8.10	Solar Observations	118
8.10.1	Solar Observing Modes	120
8.10.2	Array Observations	120
8.10.3	Single-Dish Mapping	121
8.10.4	ALMA Solar Ephemeris Generator Tool	122
8.11	VLBI Observing Mode	123
8.11.1	General VLBI Considerations	123
8.11.2	ALMA Considerations for VLBI	124

8.11.3	Observing Targets Fainter than 500 mJy with Passive Phasing	128
8.11.4	Phased-Array Observing	129
8.12	High Frequency Observing	131
9	ALMA Sensitivity Calculator	133
9.1	Calculating the System Temperature	133
9.1.1	Sky Temperature	133
9.1.2	CMB Temperature	134
9.1.3	Receiver Temperatures	134
9.1.4	Ambient Temperature	135
9.1.5	DSB Receivers	135
9.1.6	System Temperature	136
9.1.7	System Temperature (web application and OT GUI)	137
9.2	The Sensitivity Calculation	137
9.2.1	12-m and 7-m Arrays	138
9.2.2	Total Power Array	139
9.3	User Interface	139
9.4	Total Time Estimates	140
10	Calibration and Calibration Strategies	141
10.1	Fundamental Synthesis Relationship	141
10.2	Calibration Fundamentals	141
10.2.1	Calibrated and Measured Visibility Data	141
10.2.2	Calibration Simplification	142
10.2.3	Reference Antenna	143
10.2.4	Calibration Measurements Using Hardware Devices	143
10.2.5	Calibration Measurements Using Astronomical Sources	143
10.2.6	Source Catalog	144
10.2.7	Flagging Data Based on Calibrations	144
10.3	Observatory Calibrations	144
10.3.1	Antenna Primary Beam Pattern	145
10.3.2	Antenna Surface Efficiency	145
10.3.3	Antenna Pointing Accuracy	145
10.3.4	Antenna Focus	145
10.3.5	Antenna Integration Calibrations	145
10.4	Execution Block Calibrations	147
10.4.1	Making the Measurement Set	147
10.4.2	Amplitude Calibration Procedure (System Temperature Measurement)	148
10.4.3	Water Vapor Measurements	148
10.4.4	Antenna Pointing Improvement	149
10.4.5	Antenna Position Improvement	150
10.4.6	Bandpass/Frequency Calibration	150
10.4.7	Flux Density Scale	150
10.4.8	Amplitude Temporal Variations	152
10.4.9	Phase Referencing Calibration	153
10.4.10	Check Source Observation	154
10.4.11	Estimate of Image Quality in QA0	155
10.4.12	Calibration for Specific Observing Modes	157
10.5	Additional Calibrations	157
10.5.1	Self-Calibration	157
10.5.2	Astrometric Observations/Calibrations	160
10.5.3	Band-to-Band Phase Calibration	162
10.5.4	Band-Width-Switching	162

11 Quality Assurance	165
11.1 Cycle 8 Quality Assurance Goals	166
11.2 QA0	166
11.2.1 QA0 criteria	167
11.2.2 Execution Fraction	168
11.3 QA0+ and imaging quality	169
11.3.1 Relationship between image quality and phase noise criteria	169
11.4 QA1: Observatory calibration and QA	170
11.5 QA2	171
11.6 QA3	173
11.7 The Quality Assurance Reports	174
11.8 QA in Solar Projects	174
11.9 QA in VLBI Projects	175
11.10 QA in Pulsar Projects	175
12 Data Flow and Logical Data Structure	177
12.1 Data and Control Flow	177
12.2 The ALMA Science Data Model	180
13 Data Archiving	181
13.1 Introduction	181
13.2 Data Flow and Archive	181
13.3 Calibrator Source Catalogue	182
13.4 PI Data and Data Delegation	182
13.5 Archive Query	184
13.6 Request Handler	184
A Antennas	187
A.1 Design and Properties	187
A.2 Antenna Foundations	189
A.3 Antenna Transportation	190
A.4 Cryostat	191
A.5 Amplitude Calibration Device	192
A.5.1 Atmospheric Calibration Procedure	195
A.6 Water Vapor Radiometers	195
B The LO and IF System	199
B.1 Functions of the LO and IF System	199
B.2 Summary of Operation	200
B.3 Frequency Generation and Distribution in ALMA	201
B.3.1 Reference and LO Signal Generation	201
B.3.2 LO1 Signal Generation and Distribution	203
B.3.3 Line Length Corrections	204
B.3.4 The First Local Oscillator (LO1)	205
B.3.5 LO2 and the IF Processor Units and IF Switch	205
B.3.6 Digitization and Signal Transmission	207
B.4 Other functions of the LO/IF	207
B.4.1 Delay Corrections	207
B.4.2 Sideband Suppression - LO Offsetting	208
B.4.3 Interference Rejection - 180 degree Phase Switching	208
B.4.4 Sideband Separation - 90 degree Walsh Switching	209
C Acronym Dictionary	211

This page was intentionally left almost blank

Chapter 1

Introduction

The Atacama Large Millimeter/Submillimeter Array (ALMA) is an aperture synthesis telescope consisting of 66 antennas that can be positioned in a number of different configurations. It operates over a broad range of observing frequencies in the millimeter and submillimeter regime of the electromagnetic spectrum.

ALMA Early Science Operations started with Cycle 0 in September 2011 and the official inauguration took place in March 2013.

Starting with Cycle 8 operations, the distinction between standard and non-standard observing modes are removed. Unlike previous Cycles, there will be no upper limit on the amount of time allocated to non-standard observing modes in Cycle 8. Like previous Cycles, however, some observing programs, such as Large Programs, will still be restricted to specified observing modes. Users should refer to Appendix A (ALMA Capabilities) in the ALMA Cycle 8 2021 Proposer's Guide, for the latest information and a description of all the available observing modes including those programs with restricted observing modes.

The Technical Handbook concentrates on the technical aspects of the Cycle 8 observing capabilities as well as on the hardware and software available for ALMA users. It should, however, not be necessary to use the Technical Handbook to prepare an ALMA proposal.

The Technical Handbook is divided into three main sections: The concepts of interferometry and the ALMA hardware components (Chapters 2–5), the observing concepts and software (Chapters 6–9), and finally the data quality and handling (Chapters 10–13). It also includes a number of appendices with expanded information on specific ALMA hardware components and calibration and an acronym list widely used by ALMA staff.

Chapter 2 describes the ALMA Array components: The 12-m Array and the Atacama Compact Array (ACA), also known as the Morita Array which is comprised of the 7-m and Total Power (TP) Array. A general description of the different array elements and hardware components is provided.

Chapter 3 gives a brief introduction to interferometry, including a description on the concepts of basic radio astronomy and the principles of aperture synthesis.

Chapter 4 describes the details of the eight receiver bands offered for Cycle 8 observations. The general technical specifications and a brief explanation on local oscillators and intermediate frequency (IF) range is presented. Plots with the atmospheric transmission and the typical system temperatures per receiver band are also included.

Chapter 5 describes the correlators and the data processing taking place in these special purpose supercomputers. A description of the 64-input Correlator (used for the 12-m Array) and the ACA Correlator (used by the 7-m and the TP Arrays) is provided. Observing modes available for continuum and spectral line observations are presented.

Chapter 6 describes how the spectral setup is done in the correlators. A description of the signal path and local oscillator (LO) chain used between the frontends and the correlators is covered and how the hardware is configured to define spectral setups for an observation.

Chapter 7 describes several aspects of imaging to consider in ALMA observations. A short description on the different configurations proposed for Cycle 8 is included. Note that in Cycle 8, the two longest baseline configurations will not be available (Configuration-9 with a maximum baseline of 13.9 km and Configuration-10 with a maximum baseline of 16.2 km). Concepts of shadowing, beam shape, and spatial scale filtering are revisited. Mosaicing and 12-m and 7-m Array data combination are also presented.

Chapter 8 describes the observing modes offered for Cycle 8 and the observing sequence of projects. Single field interferometry, mosaics, single-dish observations, polarization, multiple region modes, Solar, VLBI, astrometry and ephemeris observations are detailed in this section.

Chapter 9 gives a brief overview how sensitivities and integration times are calculated at ALMA.

Chapter 10 describes how calibration is performed at ALMA, providing a description on how to calibrate long-term and short-term effects as well as how calibrators are selected.

Chapter 11 describes the data quality assurance (QA) process. A description of the criteria used for passing quality assurance as well as the pipeline heuristics used in data reduction and calibration which are used in determining the overall quality of the data are presented.

Chapter 12 describes the data flow and structure of ALMA data. A description of all the main software subsystems involved from data acquisition to archiving is provided as well as a description of the ALMA Science Data Model (ASDM) which defines the metadata structure adopted by ALMA.

Chapter 13 describes how the data are stored, the data flow and the user interface to the ALMA Archive.

The Technical Handbook concludes with three appendices, which contain supplemental material about the antenna design and ALMA transporter (Appendix A), the Local Oscillator (LO) and Intermediate Frequency (IF) system (Appendix B) and a list of acronyms widely used by ALMA staff (Appendix C).



Figure 1.1: ALMA antennas on the Chajnantor Plateau.

Chapter 2

Array Components

This chapter describes the main characteristics of each ALMA array. Unless otherwise noted, the description is appropriate for the fully completed ALMA telescope.

2.1 The ALMA Telescope

ALMA is composed of 66 high-precision antennas. Fifty of these antennas are 12-meter antennas in the 12-m Array, used for sensitive, high-resolution imaging. These fifty 12-m antennas are complemented by the Atacama Compact Array (ACA), also known as the Morita Array¹, composed of twelve closely spaced 7-m antennas (the 7-m Array), and four 12-m antennas for single-dish (or Total Power) observations (the TP Array), to enhance wide-field imaging of extended structures. In Cycle 8, ALMA will cover most of the wavelength range from 3.6 to 0.32 mm (84–950 GHz), and during Full Operations, will cover a range from 8.5 to 0.32 mm (35–950 GHz).

The array is located on the Chajnantor plain of the Chilean Andes (lat. = -23.02917° , long. = -67.754649°), a site that normally offers the exceptionally dry and clear sky conditions required to observe at millimeter and submillimeter wavelengths². The ALMA antennas, weather stations, the two correlators and their computer interfaces, Local Oscillator generation hardware, timekeeping hardware, and the related array Real-Time Machine computer are all located at the 5000-meter altitude site referred to as the Array Operations Site (AOS). This site is connected via Gigabit fiber links to the Operation Support Facility (OSF), located at an altitude of 2900 meters, about 22 km from the AOS and 40 km from the town of San Pedro de Atacama. Science operations are conducted from the OSF and coordinated from the JAO central office in Santiago. All three ALMA arrays are controlled via control software developed on the ALMA Common Software (ACS).

There are 192 antenna foundations (stations) distributed over the Chajnantor and Pampa la Bola plateaus. The antenna foundation distribution yields baselines (distances between two antennas) ranging from 15 m to ~ 16 km, which are crucial in determining the image quality and spatial resolution of ALMA (see Chapter 7). The antenna foundations provide the stiffness required for precise antenna pointing, as well as electrical power and digital connectivity to the main AOS building (See Appendix A.2). The antennas can be re-configured into the different array configurations (Chapter 7) using the two special purpose ALMA antenna transporters (see Appendix A.3).

The number of antennas in each array component (12-m, 7-m and TP Arrays), and the specific configurations available for an observing season (e.g. Cycle 8) will be published in the Capabilities section of the document *ALMA Proposer's Guide*. Complementary background information on ALMA and its capabilities for Cycle 8 2021 can be found in the document *Observing With ALMA: A Primer*. Both documents can be found on the link <http://www.almascience.org/documents-and-tools/>.

2.2 The 12-m Array

The 12-m Array consists of fifty 12-m diameter antennas designed and built by the European and North American ALMA partners (each providing 25 units), according to the stringent ALMA Antenna Performance

¹dedicated to the honour of K.-I. Morita

²<http://www.almascience.org/about-almawweather>




Figure 2.1: The ALMA 12-m Array in its compact configuration (left hand side of the image). The ACA with all 7-m antennas (dashed orange circles) and four single-dish 12-m antennas (blue circles) are distributed in the right hand side of the image. A few unoccupied stations can be seen, to which antennas of the 12-m Array can be moved by the transporter as the array is being reconfigured. At its most extended configuration, antennas in the 12-m Array will be about 16 km apart.

specifications (see Appendix A). Each antenna contains one front end, including a cryostat (see Appendix A.4), amplitude calibration device (ACD; A.5), water vapor radiometer (WVR; A.6), and back-end electronics (analog and digital racks). The WVRs are used to correct the phase fluctuations caused by water in the atmosphere along the line of sight of each 12-m Array element. The cryostat can contain up to ten cartridges, each covering one frequency band (see Chapter 4). Only one band observes at any time, but up to three can be switched on simultaneously. Rapid switching between receiver bands is possible but currently only used for observatory calibration (see Chapter 4). Each receiver band (Chapter 4) detects two orthogonal linear polarizations and down-converts the signals to an intermediate frequency with 7.5 GHz of bandwidth per polarization³. Bands 3-8 cartridges are dual sideband (2SB) and Bands 9 and 10 are double sideband (DSB).

The Local Oscillator (LO) signals (Appendix B) are transmitted to the antennas on optical fibers with a round trip measurement to correct for changes in the fiber length. There are four independent LO reference systems so that the 7-m Array, the TP array and two subsections of the 12-m Array (e.g., two ‘sub-arrays’) can conduct simultaneous independent observations. Please note that the sub-arrays feature of the 12-m Array is not yet a capability that is offered for PI science observations. The 7.5 GHz total IF bandwidth from the selected receiver is divided into four 1.875 GHz-wide basebands which are digitized at four Gsamples/s, with three-bit resolution, and transmitted on optical fibers. Total data rates are therefore 96 Gbits/s per antenna. With formatting, the bit rate is 120 Gbits/s.

On arrival at the central building, the data are recovered and processed in one of the two correlators: the 64-input Correlator and the ACA Correlator (see Chapter 5). All antennas can feed either correlator. The 64-input Correlator (Section 5.1) is normally used for the 12-m Array, but it can also take inputs from the 7-m Array or TP antennas. It is an XF correlator (cross-correlates first, then Fourier transforms), but the correlator proper is preceded by Tunable Filter Banks, which makes it a digital hybrid XF correlator or FXF. These can select sub-bands from the 2 GHz-wide basebands in a very flexible manner. From each of these, the correlator then generates 2016 cross-correlations and 64 auto-correlations (requiring 1.7×10^{16} operations per second). Either 2-, 3- or 4-bit resolution is used, and the sampling can be Nyquist or twice Nyquist. See Table 6.1 for an example table of spectral setups. These correlated data are fed to a group of processors which do the transforms and carry out integration and data compression.

Both correlators have minimum dump rates of 16 ms for cross-correlation and 1 ms for auto-correlation, although these dump rates can only be achieved using a reduced number of channels to prevent exceeding the maximum transmission and storage rates. The systems are designed for a maximum data rate of 64 MB/s,

³See Section 4.1 for the maximum instantaneous bandwidth available for PI science observations.

although the mean data rate will be considerably less.

The 12-m Array configurations have been designed so that at the highest frequencies in the most extended configurations, the spatial angular resolution will be as small as 5 milliarseconds at 950 GHz.

2.3 The Atacama Compact Array

Using an interferometer to obtain images of extended or large-scale structures leads to the well known “zero spacing” problem. This problem arises from the constraint that, to avoid collisions, it is not possible to pack antennas closer than their diameter, leaving a hole in the distribution of baselines at short and zero baseline separations (corresponding to large angular structure). As a result, spatial information from baselines shorter than the closed-packing ratio is not recovered⁴. This problem has considerable impact on observations of extended objects, particularly those in which the emitted power is dominated by their large-scale structures.

To achieve high-fidelity imaging of sources with emission on angular scales larger than those corresponding to the minimum spacing of the 12-m Array (the “Maximum Recoverable Scale” for that array - see Section 7.6), ALMA has been designed to include the Atacama Compact Array (ACA), also known as the Morita Array.

The ACA is composed of twelve 7-m antennas for interferometry (the 7-m Array) and four 12-m antennas for single-dish observations (TP Array). The four single-dish antennas provide spatial information samples equivalent from 0 m up to 12 m spacings as auto-correlations. The 7-m Array samples baselines from 9 m to 30 m, bridging the baseline sampling gap between the 12-m Array and the TP Array. The number of array elements available is published for each observing cycle.

The ACA is controlled via control software developed on the ACS platform and is operated in a similar fashion to the 12-m Array. To achieve this unified operation, the ACA system is as compatible with the 12-m Array as possible at the level of hardware, interface, data, and observing modes. The standard observing modes for the TP Array include spectral line and continuum observations with raster or Lissajous on-the-fly (OTF) scans, or position switching. The raw time-series signals from the ACA antennas are processed in the ACA Correlator (see Section 5.2) to produce the cross-correlated and auto-correlated data.

2.3.1 The 7-m Array

The 7-m Array is composed of twelve 7-m diameter antennas designed and built by East Asia to the ALMA specifications (see Appendix A). Each antenna contains one front end, including a cryostat, amplitude calibration devices, and one back end. Similar to the 12-m antennas, each cartridge (see Chapter 4) receives two orthogonal polarizations. Unlike the antennas of the 12-m Array, the 7-m antennas do not contain Water Vapor Radiometers (WVRs). However, since the 7-m Array is quite compact, atmospheric phase fluctuations will be smaller than for the 12-m Array.

The 7-m antenna cryostats are fitted with receivers nearly identical to those on the 12-m antennas, with small differences in the warm optics. The Local Oscillator signals transmitted to the 7-m antennas are originated identically to the ones sent to the 12-m Array, i.e., from inside the AOS building.

The ACA Correlator is normally used for the ACA, and can work with two sub-arrays. It is an FX correlator (Fourier transform first, then cross-correlate) with 3-bit input and 4 bits in the correlation. The correlator generates 120 cross-correlations and 16 auto-correlations for each baseband. These are passed to a (special-purpose) data processing computer at up to ~ 0.6 Gbits/s per baseband.

Even in its most compact configuration the 12-m Array does have spacings smaller than 15 m. The array configuration of the 7-m Array is designed to fill in these shorter spacings from about 9 m to ~ 30 m (see Chapter 7).

2.3.2 The TP Array

The TP Array can fill in baseline coverage from 0 m to about 12 m, complementing the 7-m and 12-m Array’s baseline coverage. It consists of four 12 m diameter antennas built by East Asia (Appendix A). The specifications of the TP antennas are almost identical to the ones for the 12-m Array. The TP antennas are located on stations surrounding the 7-m Array.

⁴Strictly speaking, mosaicing with imaging using a joint-deconvolution algorithm allows the recovery of more spatial information than normal synthesis imaging, but the problem caused by absent short and zero spacing information still remains.

The TP Array is usually connected to the ACA Correlator, but its antennas can also be connected to the 64-input Correlator and used for cross-correlation. The ALMA Cycle 8 2021 Proposer's Guide describes the observing modes and capabilities offered for the TP Array for each cycle.

Due to the poorer point-source sensitivity of the 7-m Array, during Cycle 8 Operations, the TP Array will be routinely used in the calibration observations of the 7-m Array.



Figure 2.2: The Morita Array - In remembrance of Professor Koh-Ichiro Morita. Koh-ichiro Morita, a professor at the NAOJ Chile Observatory, was one of the world's renowned scientists in the field of aperture synthesis. He made a great contribution to designing the configuration of 16 antennas composing the Atacama Compact Array (ACA) manufactured by Japan, as well as to realizing high-resolution and high-quality imaging at millimeter/submillimeter wavelengths to further enhance the performance of ALMA. The picture above shows Professor Koh-Ichiro Morita taken at his office in the Joint ALMA Observatory in 2011 September.

Chapter 3

Principles and Concepts of Interferometry

3.1 Introduction

Interferometry is the technique ALMA uses to obtain very high angular resolution observations of astronomical phenomena. Figure 3.1 shows examples of two other observatories, the IRAM NOthern Extended Millimeter Array (NOEMA) Interferometer in France and the Submillimeter Array in Hawaii, which also use this technique. In this chapter, the principles and concepts behind interferometry are described, so that ALMA users can plan and understand their observations better.

Interferometry involves the combination of signals received from the sky by two or more physically separated antennas. The signals are interfered, allowing a sky brightness distribution to be sampled on an angular scale smaller than possible with a single antenna. The interference modifies the angular sensitivity of the antennas to include a sinusoid of constructive and destructive nodes. In this sense, the only emission measured by the interferometer is that from the scale defined by the angular extent of the sinusoidal wavelength, equivalently, the “spatial frequency”. This wavelength is inversely proportional to the projected distance between the two antennas. Each datum from the interferometer is called a visibility, and is a measure of the brightness of the emission on the angular scale sampled (i.e., related to the amplitude of the sinusoid) and the relative position of that brightness on the sky (i.e., related to the phase of the sinusoid.)

A range of discrete angular scales can be sampled by including many pairs of antennas in an array. Importantly, by tracking a source across the sky, the rotation of the Earth can be used to change the projected separations of the antenna pairs, allowing more angular scales to be sampled. An ensemble of the data (i.e., sinusoids of various amplitude and phase) can be then “summed” via the Fourier transform to produce an image of the sky brightness distribution. How well this image reflects the actual sky brightness distribution depends on how completely the relevant angular scales have been sampled. With its emphasis on delivering images at high angular resolution, interferometry works extraordinarily well for observing intrinsically compact targets.

The following sections expand upon these basic ideas. We begin by introducing the basic concepts of radio astronomy, and then move to the principles of aperture synthesis.

3.2 Single-dish Response

The *brightness*, or equivalently *specific intensity*, I_ν , is defined as the electromagnetic (EM) power δP within a range of frequencies (a bandwidth) $\delta\nu$ received from a solid angle $\delta\Omega$ and intercepted by surface area δA :

$$I_\nu = \delta P / \delta\Omega \delta A \delta\nu, \quad (3.1)$$

where I_ν has typical units of $\text{W m}^{-2} \text{ Hz}^{-1} \text{ sr}^{-1}$. In addition, the *flux density*, S_ν , is defined as the integration of brightness over the solid angle of the emitting source:

$$S_\nu = \int I_\nu d\Omega, \quad (3.2)$$




Figure 3.1: The NOthern Extended Millimeter Array (NOEMA) (top) and the Submillimeter Array (SMA) (bottom) are two of the precursors to the ALMA telescope; both are still in full operation and helped pioneer the science of millimeter wave interferometry.

where S_ν has typical units of $\text{W m}^{-2} \text{Hz}^{-1}$. In millimeter/radio astronomy, the power received is typically so weak that a convenient unit to use for S_ν is the *Jansky* (Jy), where $1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{Hz}^{-1}$. A radio telescope with effective area A_e receives power P_{rec} per unit frequency from an unpolarized source:

$$P_{rec} = \frac{1}{2} I_\nu A_e \delta\Omega. \quad (3.3)$$

The coefficient of 1/2 in Equation 3.3 comes from the fact that a receiver is generally sensitive to only one mode of polarization. (ALMA's receivers, however, have been constructed with two independent receptors so that both modes of polarization can be detected simultaneously, so the coefficient is actually 1 for ALMA.) The antennas bring incident EM power to a focus after reflecting it off a primary surface. The antenna response (i.e., its relative sensitivity) is a summation of all EM power brought to the focus.

The antenna response is actually dependent on the angle from the on-axis pointing direction of the antenna due to diffraction; i.e., self-interference. To demonstrate the angular dependence, Figure 3.2 shows in the top panel the case for EM power of wavelength λ arriving along the axis of an unobstructed antenna of diameter D . Since the source of the EM power is very distant, the EM power arrives at the primary surface essentially as plane-parallel wavefronts. Note that the antenna surface is parabolic in shape, so the path that each part of the front travels to the focus is constant. With zero path difference, the EM power arriving on-axis is coherently summed at the focus. This arrangement is only true, however, along the axis of the antenna. In the lower panel of Figure 3.2, the case for EM power arriving from an off-axis direction is shown. In this situation, the EM power does not add as constructively. In addition, the diameter of the antenna projected along the off-axis direction is less than the true diameter, decreasing the amount of power received from that direction. As a result, the antenna power response, i.e., its relative sensitivity, will be less than that found on-axis. In particular, at the off-axis angle of λ/D radians, where λ is the wavelength of observation and D is the diameter of the antenna, the path difference across the antenna diameter, the *aperture*, will equal one wavelength of the incident emission. The combination of such emission at the focus eads to destructive interference at that angle.

As an illustration, Figure 3.3 shows an example of a one-dimensional antenna power response with angle for a 12-m diameter parabolic antenna uniformly illuminated by emission of wavelength $\approx 0.85 \text{ mm}$ (350 GHz). The power response is largest on-axis but it declines to almost zero in ~ 18 arcseconds. The central Gaussian-like feature is called the *primary beam* or the *antenna beam size* and it has a Half Power Beam Width (HPBW) given by:




Figure 3.2: Schematic of an incoming plane-parallel wavefront reflecting off an antenna of diameter D and being brought to a focus. The top panel shows the case for a wavefront arriving on-axis. The bottom panel shows the case for a wavefront arriving off-axis. Note that the paths of incident EM power in the first case are all of equal length, and hence the power is summed constructively at the focus. In the second case, the path lengths differ, leading to less constructive summations at the focus.

$$\text{HPBW Primary Beam} = 1.02 \times \lambda/D. \quad (3.4)$$

For example, the FWHM of a uniformly illuminated antenna of 12 m diameter at $\lambda = 0.85$ mm is $14.9''$. HPBW is sometimes referred to as Full Width at Half Power or FWHP. In one dimension, the FWHP is equivalent to the Full Width at Half Maximum or FWHM, a quantity sometimes used to describe the width of a Gaussian approximation to the central feature.

Note that the antenna power response rises and declines repeatedly at ever larger angles. The constructive and destructive interference at larger angles leads to successive *sidelobes* (whose maxima decline with increasing angle) and *nulls* respectively. The first sidelobes have a relative response of only 1.74% that of the primary beam. Nevertheless, incident emission, if bright enough, coming in at angles well beyond those of the primary beam can make a large contribution to the received EM power. The angular distance between the first nulls is termed the Full Width Between Nulls (FWBN), and is given by:

$$\text{FWBN Primary Beam} = 2.44 \times \lambda/D. \quad (3.5)$$

Half the FWBN of the primary beam, $\sim 1.22 \lambda/D$, is considered the Rayleigh resolution of the antenna, i.e., its ability to distinguish objects on the sky separated by some angular distance. For convenience, the antenna

power response is typically normalized to 1.0 along the axis. Figure 3.3 illustrates the antenna power response in one dimension (in log units); on the actual sky, the antenna power response is two-dimensional, and is obtained by rotating the function shown in Figure 3.3 about its central axis.

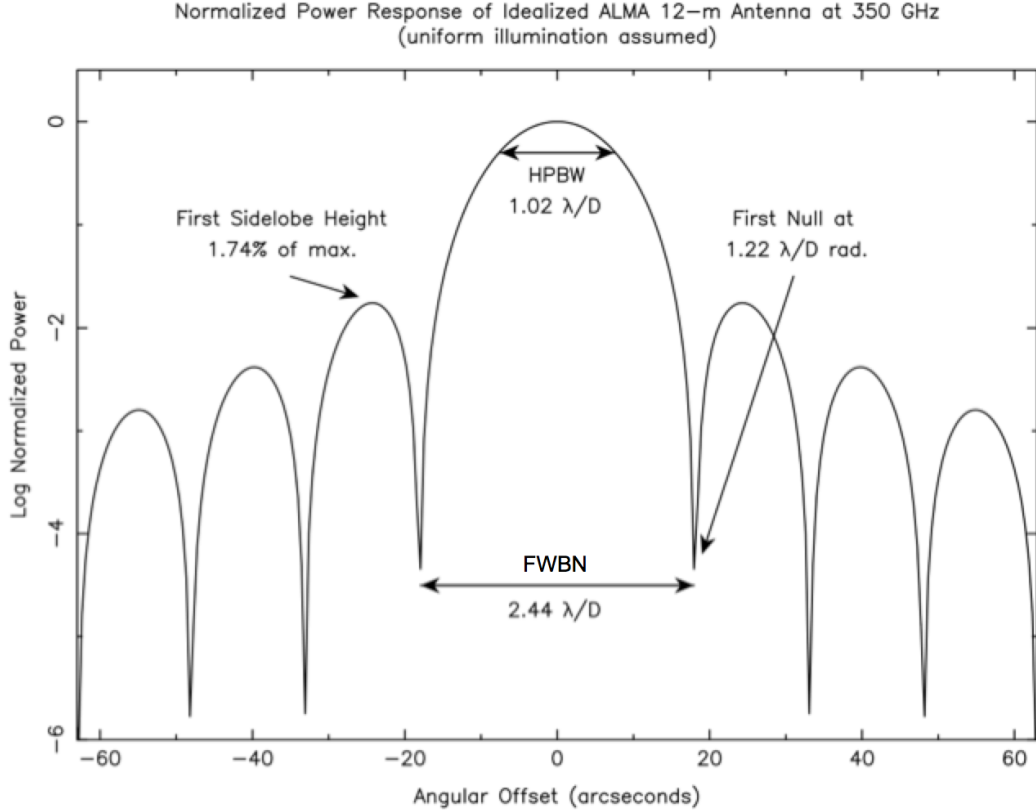


Figure 3.3: Normalized 1-D antenna power response for a 12-m antenna uniformly illuminated at 350 GHz. The power is in log units to emphasize the sidelobes. The HPBW of the primary beam is $\sim 1.02 \lambda/D$ and the FWBN is $\sim 2.44 \lambda/D$. The angle of the first null, i.e., the resolution, is $0.5 \times \text{BWFN} = \sim 1.22 \lambda/D$. Note that the HPBW measured for actual ALMA 12-m and 7-m antennas is $\sim 1.13 \lambda/D$ because the illumination is not uniform.

Up until now, an idealized antenna was described. The power response of an actual antenna, however, can be altered by various effects, including the degree by which the secondary is illuminated, diffraction by the arms supporting the secondary, and surface imperfections. For ALMA, each 12-m antenna has a secondary reflector and support arms that block an effective area of 0.75 m diameter on the primary surface. The actual ALMA feedhorns were designed to provide an antenna power response with a nearly Gaussian primary beam and low sidelobes, preserving as much resolution and sensitivity as possible. The actual ALMA 12-m antennas (and the ALMA 7-m antennas) have measured primary beam HPBW values of $\sim 1.13 \lambda/D$.

An equivalent way to consider the antenna power response is in terms of the voltage response, $V(\theta)$, where $P(\theta) \propto V^2(\theta)$. In the far-field, i.e., under the Fraunhofer approximation, a diffraction pattern at the point of observation is the Fourier transform of the field distribution at an aperture. Hence, the voltage response at the focus is the Fourier transform of the aperture shape. For an unobstructed antenna, the aperture is a uniform circle, and $V(\theta) = J_1(\theta)/\theta$, where $J_1(\theta)$ is the Bessel function of the first kind. $P(\theta)$ is correspondingly proportional to $(J_1(\theta)/\theta)^2$. The normalized version of the antenna power response, P_N , is also known as the Airy function.

Defining θ and φ as orthogonal directional variables (e.g., sky coordinates), $I_\nu(\theta, \varphi)$ and $P_N(\theta, \varphi)$ can be

defined as the directional functions of the sky brightness and the normalized antenna power, respectively. The total received power of an antenna at a given pointing is the integration over the sky of the product of the sky brightness distribution and the antenna power response:

$$P_{rec} = \frac{1}{2} A_e \int_{4\pi} I_\nu(\theta, \varphi) P_N(\theta, \varphi) \delta\Omega. \quad (3.6)$$

In addition, the solid angle of the antenna power response $P_N(\theta, \varphi)$ is defined by:

$$\Omega_A = \int_{4\pi} P_N(\theta, \varphi) \delta\Omega. \quad (3.7)$$

3.3 Visibilities and Aperture Synthesis

Observing at millimeter/radio wavelengths is essentially diffraction-limited, as the angular diameter of the Airy disk is the FWHM for a circular aperture of diameter D is $\sim 1.22 \lambda/D$. Millimeter/radio wavelength (single-dish) observations have lower resolutions than optical wavelength (single-dish) observations because λ is larger by many orders of magnitude. Though the D of millimeter/radio single-dish telescopes can also be much larger than those of optical wavelengths, the increase in D possible for single-dish telescopes is generally never enough to obtain angular resolutions comparable to ground-based optical telescopes, e.g., $1''$ or better. For example, the JCMT 15-m diameter antenna has an angular resolution at $850 \mu\text{m}$ of $\sim 14''$.

To obtain higher angular resolution images at millimeter/radio wavelengths, signals from physically separated antennas can be combined through interferometry. With this technique, sometimes called *aperture synthesis*, it is possible achieve a resolution that emulates the effect of having a radio telescope with a larger diameter. Observers, however, must contend with the reality that only certain angular scales (i.e., those determined by the projected separations of each pair of antennas) will be sampled. This section builds on the concepts introduced previously to discuss aperture synthesis in more detail. This powerful technique of aperture synthesis was pioneered by Sir Martin Ryle, who shared the Nobel Prize in Physics in 1974 in part for its development.

As there are no path differences in the case of a plane-parallel wavefront arriving on-axis an antenna, EM power from across the antenna is brought together in phase at the focus. Now imagine that the parabolic surface is divided into N smaller contiguous areas; i.e., *elements*. In this situation, the received voltage $V(t)$ is the sum of contributions $\Delta V_i(t)$ from each of element i :

$$V(t) = \sum_i \Delta V_i(t) \quad (3.8)$$

The power received by the antenna is proportional to the running time average of the square of the contributions from each element. Assuming illumination is the same for each element, the expression for received power in terms of the sum of time averages of the products of voltages from element pairs can be rewritten:

$$\langle P \rangle \propto \left\langle \left[\sum \Delta V_i \right]^2 \right\rangle = \sum \sum \langle \Delta V_i \Delta V_k \rangle. \quad (3.9)$$

Next, this expression can further be rewritten in terms of the sums of element pairs which are the same and those which are not as:

$$\langle P \rangle \propto \sum \langle \Delta V_i^2 \rangle + \sum_{i \neq k} \sum \langle \Delta V_i \Delta V_k \rangle. \quad (3.10)$$

The first and second sets of terms in Equation 3.10 are called *auto-correlation* and *cross-correlation* terms, respectively, since the voltages multiplied in each term are from either the same or different elements, respectively.

From Equation 3.10, any measurement with a large filled-aperture telescope can be understood as being a sum in which each term depends on contributions from only two of the N elements. As long as the contributions from each element arrive at the focus in phase, *there is no need for the elements to be physically contiguous*. Generalizing, each cross-correlation term $\langle \Delta V_i \Delta V_k \rangle$ in Equation 3.10 can be measured with two smaller, physically separated antennas (at locations i and k) by measuring the average product of their output voltages with a correlating, or multiplying, receiver. Moreover, if the source properties do not change, there is

no need to measure all pairs at the same time. A given parabolic surface with N elements has $N(N - 1)/2$ pairs of elements, and these could be observed sequentially to “synthesize” a measurement by a large filled-aperture telescope. Alternatively, numerous pairs of antennas, with each antenna considered an element, can be distributed to positions at distances much larger than it is possible to build a single filled-aperture telescope, and the signals received by these antennas can be combined in phase to approximate the resolving power of a single filled-aperture telescope.

The above situation only describes the emission received on-axis from antenna pairs. Of course, as noted above, emission also arrives at the antennas from other directions, leading to phase differences. To understand the power response expected from a pair of antennas, let’s look at the ideal 1-D situation of a two-antenna interferometer.

Figure 3.4 shows a schematic picture of a two-antenna interferometer separated by distance b , known as a *baseline*. This distance can be measured in units of the observing wavelength, λ . In terms of familiar units of length, $b = L/\lambda$, where L is the distance between antennas and λ is the wavelength in the same unit; e.g., meters. Both antennas observe a common position s_o located at an angle θ from the meridian. The projected separation of the two antennas towards s_o from the perspective of the source is $u = b \cos \theta$. In this example, an on-axis wavefront incident to both telescopes reaches antenna 2 first and the wavefront reaches antenna 1 a little later, having traversed an extra path length of $b \cdot s_o = b \sin \theta$. In other words, emission received by antenna 1 experiences a *geometrical delay* relative to that received by antenna 2, where the time equals $\tau_g = b \cdot s_o / c$. To compensate for the geometrical delay, an artificial delay can be inserted into the signal path of antenna 2 (e.g., electronically) so that the signals from both antennas arrive at the correlator with the same phase.




Figure 3.4: An ideal 1-D interferometer consisting of two antennas, 1 and 2, separated by physical distance (i.e., a baseline) b . The antennas are both pointed towards a sky location given by s_o , which is at an angle θ from the meridian. The projected distance between the two antennas in that direction is thus $u = b \cos \theta$. The two antennas are connected to a correlator where the voltages detected from each are combined.

Moving slightly off-axis, a small angle from the axis can be described as α , and its 1-D sky position as

$l = \sin \alpha$, i.e., the direction cosine. At angle α , an off-axis signal reaching antenna 1 will have to travel a slightly longer path than an off-axis signal reaching antenna 2, even with the geometrical delay introduced to compensate for an on-axis signal. This extra path length is $x = u \sin \alpha = ul$. All distances can be considered in this situation in units of the wavelength of the emission, λ , so that x is the number of wavelengths within a given distance. The extra path lengths result in phase differences with α that can be characterized where the voltage response of antenna 2, V_2 , can be written in terms of the product of the voltage response of antenna 1, V_1 , and a phase delay factor sinusoidally varying as a function of angle:

$$V_2 = V_1 e^{2\pi i(ul)}. \quad (3.11)$$

Expanding to two dimensions introduces β , a direction on the sky orthogonal to α . Also, $m = \sin \beta$ is defined as the small angle analog to l in this new direction. Note that the baseline here is actually a 2-D vector with components in both dimensions, i.e., b_1 and b_2 . Hence, we define $u = b_1 \cos \theta$ and $v = b_2 \cos \varphi$ where φ is the angle of the position s_o on the sky from the reference position orthogonal to θ . Finally, y is defined as the extra path length introduced in this new direction, in units of the wavelength of emission, i.e., $y = v \sin \beta = vm$. With these changes, the two-dimensional voltage response of antenna 2 is:

$$V_2 = V_1 e^{2\pi i(ul+vm)}, \quad (3.12)$$

where u and v are identified as specific *spatial frequency* components of the sinusoid in the E-W and N-S directions respectively, and these are the projected lengths of the antenna separations measured in units of the wavelength at the time of observation. Also, l and m are identified as direction cosines relative to a reference position in the E-W and N-S directions, respectively. Typically, the on-axis position s_o has $l = 0$ and $m = 0$ and is called the *phase center*.

The correlator acts as a multiplying and time-averaging device for the incoming signals from antennas 1 and 2. Hence, its output is:

$$\langle V_1 V_2 \rangle = \langle \iint V_1(l, m) dldm \iint V_2(l, m) dldm \rangle. \quad (3.13)$$

Under the assumption that signals emanating from different parts of the sky are incoherent (i.e., they have no similarities in phase), the time averages of the correlation of those signals will be zero. Thus, the product of the integrals in Equation 3.13 can be simplified to:

$$\langle V_1 V_2 \rangle = \langle \iint V_1(l, m) V_2(l, m) dldm \rangle \quad (3.14)$$

$$\langle V_1 V_2 \rangle = \iint \langle V_1(l, m) V_2(l, m) \rangle dldm \quad (3.15)$$

$$\langle V_1 V_2 \rangle = \iint \langle V_1(l, m)^2 \rangle e^{2\pi i(ul+vm)} dldm. \quad (3.16)$$

As $V^2 \propto P$ (see Equation 3.9) and $P \propto I_v$ (see Equation 3.3),

$$\langle V_1 V_2 \rangle \propto \iint I(l, m) e^{2\pi i(ul+vm)} dldm \quad (3.17)$$

where $I(l, m)$ is the intensity distribution on the sky. The correlator therefore measures a quantity known as the *complex visibility*, \mathcal{V} , which is formally the Fourier transform of the intensity distribution on the sky:

$$\mathcal{V}(u, v) = \iint I(l, m) e^{2\pi i(ul+vm)} dldm = A e^{i\phi}. \quad (3.18)$$

Note that \mathcal{V} is a complex number, and can be described by an amplitude, A , and a phase, ϕ . The amplitude and phase contain information about the source brightness and its location relative to the phase center, respectively, at spatial frequencies u and v .

3.4 The Visibility or (u, v) Plane

The relationship between the sky brightness distribution and the complex visibility distribution is governed by the van Cittert-Zernike theorem and it is the basis of aperture synthesis. Given that the complex visibility is the Fourier transform of the sky brightness distribution in the image plane, it follows that the sky brightness distribution is in turn the inverse Fourier transform of the complex visibility distribution in the visibility plane:

$$\mathcal{V}(u, v) = \iint I(l, m) e^{2\pi i(ul+vm)} dl dm \quad (3.19)$$

$$I(l, m) = \iint \mathcal{V}(u, v) e^{-2\pi i(ul+vm)} du dv. \quad (3.20)$$

Here, $\mathcal{A}(l, m)$ and $1/\sqrt{(1-l^2-m^2)}$ are omitted, as they are usually close to unity. By measuring the distribution of complex visibilities (in the visibility or (u, v) plane), in principle the sky brightness distribution can be recovered. In essence, an image is a “sum” (i.e., the Fourier transform) of the visibilities where each visibility has an amplitude and phase representing the brightness and relative position of emission on a specific angular scale. The image and its Fourier transform are conjugates of each other, and each contains the same amount of information.

Two antennas at a given physical distance b can have signals interfered to sample the sky brightness distribution on a scale inversely proportional to the projection of that distance on the sky. As shown above, the response of the interferometer is sinusoidal, and is sometimes referred to as a *fringe*, with spacing on the sky in the 1-D case of:

$$\text{Fringe Spacing} = 1/u \text{ (radians)} = 1/(b \cos \theta) = \lambda/(L \cos \theta). \quad (3.21)$$

In effect, the interference of the signals modifies the angular response of the antennas and the antennas can “see” the true sky brightness distribution only on the scale defined by the wavelength of the sinusoid. As the fringe spacing depends inversely on the projected distance, antennas closer together measure emission on larger scales. Conversely, those antennas spaced further apart measure emission on smaller scales. Since fringe spacing also depends on the wavelength of emission, as b is measured in numbers of wavelengths, observing shorter or longer wavelengths also can sample smaller or larger scales, respectively. These ideas can be easily generalized to two dimensions, with the fringe spacing and on-sky orientation depending on the relative magnitudes of u and v .

A given pair of antennas will only instantaneously sample a single scale of the sky brightness distribution. Given the E-W and N-S separations of the pair, a visibility in the (u, v) plane is measured. Since visibilities are samples of a complex-valued function with Hermitian symmetry, a single sampling gives two visibilities, one at (u, v) and its complex conjugate at $(-u, -v)$. To recover the true sky brightness distribution, however, knowledge of the distribution of visibilities across the (u, v) plane would be needed. Improving coverage of visibilities over the (u, v) plane can be done in several ways. First, multiple antennas can be incorporated into an array, with each at a different distance from the others to prevent redundancy. An array of N antennas will have $N(N-1)/2$ independent baselines, with each pair providing a single pair of samples in the (u, v) plane. Second, a target can be observed repeatedly by the array as it appears to move across the sky due to the Earth’s rotation. Though the physical distances between the antennas do not change, their projected distances do change depending on the altitude and azimuth of the target. Hence, repeated observations by all the pairs in an array can sample many visibilities across the (u, v) plane. Finally, antennas in the array may be arrangeable in several configurations so that pairs of antennas have different distances and can sample different parts of the (u, v) plane. Assuming the source emission is not variable, the combination of these schemes can reasonably sample the (u, v) plane, yielding an image that can resemble the true sky brightness distribution.

3.5 Fields-of-view and Mosaics

Each antenna of an actual interferometer has finite diameter. As noted before, such antennas have their own power responses on the sky $P_N = \mathcal{A}(l, m)$; e.g., the Airy function for an unobstructed, uniformly illuminated aperture. Indeed, the individual antenna response fundamentally limits the extent of an interferometric image

made with a single pointing. In practice, the HPBW of the primary beam serves as the “field-of-view” of the single-pointing interferometric image. Moreover, $\mathcal{A}(l, m)$, the Airy function, is actually included formally in the correlator output:

$$\mathcal{V}(u, v) = \iint \mathcal{A}(l, m) I(l, m) e^{-2\pi i(ul+vm)} dl dm. \quad (3.22)$$

Hence, an interferometer actually measures the Fourier transform of the sky brightness distribution multiplied by the antenna power response. To recover $I(l, m)$, the image resulting from the Fourier transform of the complex visibilities must be divided by $\mathcal{A}(l, m)$ as the last step of image processing. This so-called *primary beam correction* should be performed if the image contains scientifically relevant emission off the phase center. Otherwise, intensities and flux densities measured at those locations will not be accurate.

To counteract the angular fall-off of sensitivity due to the primary beam response, or even to sample emission over areas on the sky larger than the primary beam, an interferometer can observe adjacent positions, producing a *mosaic*. Sensitivity across a mosaic depends on the spacing of the individual positions observed. A mosaic can have close to uniform sensitivity with a minimum number of pointings if the positions observed are arranged in a grid of equilateral triangles spaced by $\lambda/(\sqrt{3}D)$, where D is the diameter of the antenna. With this spacing, the fall-off of the primary beam response at one pointing is made up by the responses of the primary beams at adjacent pointings, except of course at the edge of the mosaic (see Section 7.7 for a more complete description). Mosaics can be made with adjacent positions that are spaced either closer or more distant, with non-uniform sensitivities. Mosaics provide increased areal coverage but at a cost of more observing time. For example, uniform sensitivity can be obtained across the primary beam of a single pointing but requires a minimum of six other pointings around the single pointing. These six pointings are arranged in a hexagonal pattern with each vertex spaced $\lambda/(\sqrt{3}D)$ from each other (and the single pointing). A single image is produced by combining the visibilities obtained at all pointings into a single ensemble that is simultaneously Fourier transformed.

3.6 Spatial Filtering

Through the principles described in Sections 3.3-3.5, the true sky brightness distribution can be recovered. It is, however, impossible in practice to sample completely the (u, v) plane and obtain all visibilities. The incomplete (u, v) plane sampling effectively provides a fundamental limit to the level of detail discernible in the sky brightness distribution; i.e., down to a minimum scale defined as the resolution. In addition, incomplete sampling results in *spatial filtering* of the true sky brightness distribution; i.e., the resulting images do not contain information on angular scales unobserved by the interferometer. In particular, the lack of coverage at the shortest baselines (i.e., lower than those sampled by the smallest baselines) results in an intrinsic lack of sensitivity to large-scale emission. It is crucial for ALMA users to understand these limitations. Below, we describe the ideas behind spatial filtering and provide a few illustrative examples. Further examples can be found in Sections 7.6 and 7.7.

First let’s discuss resolution. The resolution of any interferometric image depends on the distribution of visibilities sampled. Assuming a finite number of M visibilities has been obtained, the (u, v) plane has been sampled at $2M$ discrete points. The sampling distribution can then be characterized as an ensemble of $2M$ (Dirac) delta functions:

$$B(u, v) = \sum_{k=1}^{2M} \delta(u - u_k, v - v_k). \quad (3.23)$$

Given the sampling function $B(u, v)$, we can calculate the function I_V^D , which is the *dirty image*:

$$I_V^D(l, m) = \iint \mathcal{V}(u, v) B(u, v) e^{-2\pi i(ul+vm)} du dv. \quad (3.24)$$

The inverse Fourier transform of this ensemble of visibilities can then be written as:

$$I^D(l, m) = FT^{-1}\{B(u, v)\mathcal{V}(u, v)\}. \quad (3.25)$$

Following the convolution theorem, the Fourier Transform of a convolution of two functions is the product of the Fourier Transforms of those functions. Hence, Equation 3.24 can be rewritten as

$$I^D(l, m) = b(l, m) * I(l, m)\mathcal{A}(l, m). \quad (3.26)$$

In effect, the image obtained is the convolution of the true sky brightness distribution (modified by the antenna power response $\mathcal{A}(l, m)$) with the point spread function, $b(l, m) = FT^{-1}\{B(u, v)\}$, which is the Fourier transform of the (u, v) plane sampling distribution. The point spread function is sometimes called the *synthesized beam* or the *dirty beam*. It is important to distinguish this beam from the single-dish response function $\mathcal{A}(l, m)$, which in the interferometry context is called the primary beam. The image resulting from the Fourier transform of a finite number of visibilities, $I^D(l, m)$, is sometimes referred to as the *dirty image*.

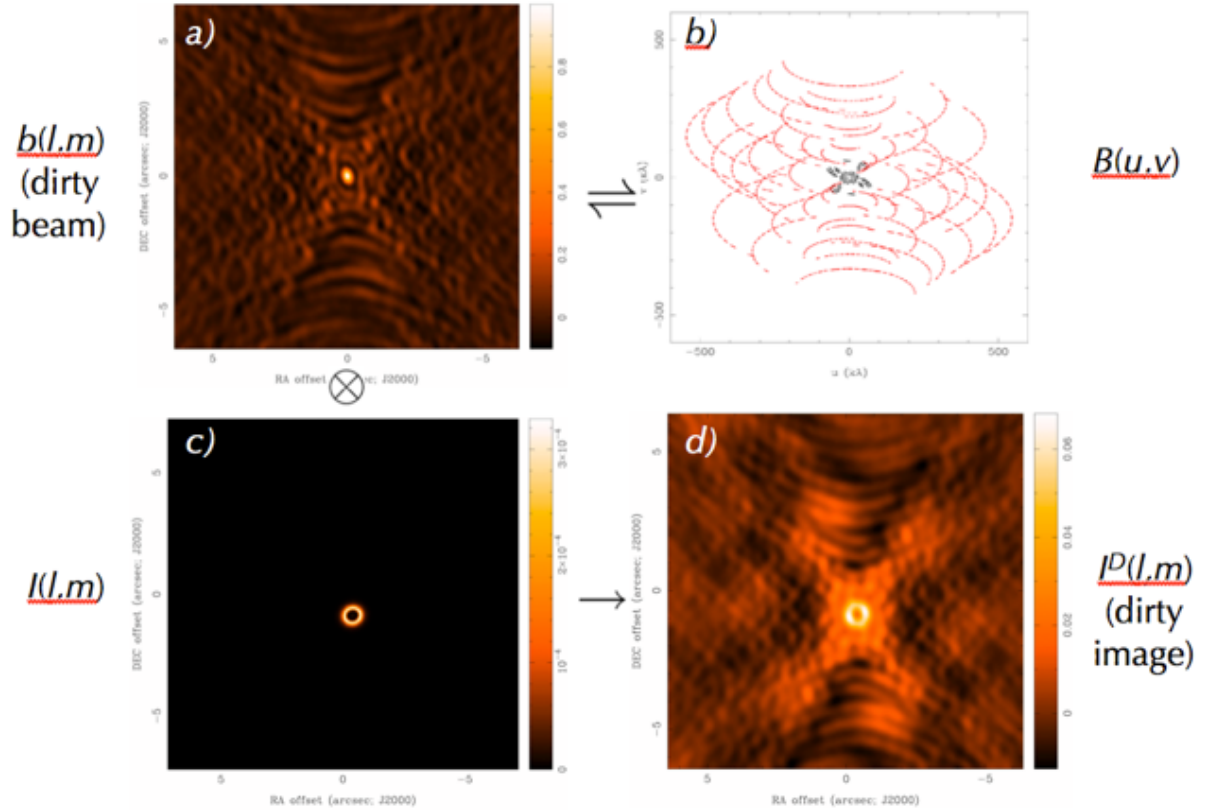


Figure 3.5: Imaging concepts. *Panel a (upper left)*: Example of a dirty beam, $b(l, m)$. *Panel b (upper right)*: The related ensemble of discrete points sampled in the (u, v) plane, $B(u, v)$. The black points were obtained from a compact configuration while the red ones were obtained from an extended configuration. *Panel c (lower left)*: Example of a true sky distribution, $I(l, m)$. *Panel d (lower right)*: The dirty image $I^D(l, m)$ resulting from observing $I(l, m)$ over the baselines of $B(u, v)$, or equivalently the convolution of $I(l, m)$ by $b(l, m)$. The antenna power response, $\mathcal{A}(l, m)$, has been ignored in this illustration since it is much wider than the true sky brightness distribution. (Figure courtesy of D. Wilner.)

The measure of how similar an image is to the true sky distribution is sometimes referred to as *image fidelity*. Image fidelity depends on the specifics of coverage of the (u, v) plane sampled by the interferometer. Since the numbers of samples are necessarily finite and discrete, there are invariably gaps in any practical sampling of the (u, v) plane. These gaps mean that no information is obtained about the true sky brightness distribution on those specific angular scales. Note that visibilities corresponding to those unobserved scales can have any value. With no information, however, it is typically assumed that $\mathcal{V}(u, v) = 0$ at unsampled locations in the (u, v) plane. Including these visibility domain gaps through the Fourier transform produces aliased features in the resulting image, the magnitude of which depends on the extents and locations of gaps in the (u, v) plane and the brightness of emission on sampled scales. If the (u, v) plane has been reasonably well sampled, the synthesized beam will consist of a compact positive feature surrounded by positive and negative features of lower relative

amplitude. These latter features, also called *sidelobes*, can complicate the image since brightness is distributed via the point spread function throughout the image. The resulting image can have significant artifacts depending on the sky brightness distribution and the sampling of the (u, v) plane. A dirty image, however, can be improved through deconvolution techniques to minimize the effect of incomplete spatial frequency sampling (e.g., CLEAN and its variants; see Chapter 7).

Though speaking generally about true sky brightness distributions so far, a special note should be made for the case of point sources. Obviously, a point source is a distribution of emission that is not extended relative to the resolution of the observation. In this case, the morphology of the dirty image will equal that of the dirty beam (e.g., see Equation 3.25). Moreover, the complex visibilities of the point source have the same amplitudes on all observed angular scales. Of course, sources may appear point-like at low resolutions but may appear extended in higher-resolution observations.

Figure 3.5 illustrates the concepts of dirty beam and dirty image and their impact on the recovered image. Panels *a* and *b* (upper pair) show respectively the dirty beam and the related ensemble of locations observed in the (u, v) plane; i.e., the (u, v) coverage. Note in panel *a* the positive feature in the center of the dirty beam distribution and the surrounding positive and negative features of lower amplitude. These latter features arise from the incomplete sampling of the (u, v) plane seen in panel *b*. As an aside, note that two sets of (u, v) plane samples are identified in panel *b*; these result from observations by the same antennas in two different configurations, a compact one (black) and an extended one (red). Panels *c* and *d* (lower pair) show respectively an example of a true sky distribution (here, a model of a ring of emission) and the dirty image. The dirty image is the convolution of the true sky distribution by the dirty beam, and one can easily see how incomplete sampling of the (u, v) plane leads to the appearance of significant artifacts in the resulting dirty image.

The resolution of the dirty image is defined effectively by the compactness of the central feature of the dirty beam; e.g., half its FWBN. Since the structure of the dirty beam is generally more complicated than that of a single-dish antenna (e.g., the beam from a uniformly illuminated antenna shown in Figure 3.3), it is not so easy to measure FWBN. Instead, the resolution is typically approximated to first order by the FWHM of a Gaussian fit to the central feature of the dirty beam. The resolution of the dirty image depends ultimately on how the interferometer antennas are arranged. In general, distributions connected via a Fourier transform scale inversely to each other. For example, narrow distributions in one domain have wide ones in the other, and vice versa. By analogy, an ensemble of discrete points, $B(u, v)$, clustered around the (u, v) plane origin provided by a compact configuration yields a low-resolution image since the central beam feature $b(l, m)$ is wide. Conversely, an ensemble of discrete points distributed more widely from the (u, v) plane origin yields a high-resolution image since the central beam feature is narrow. Indeed, resolution is fundamentally limited by the extent of the longest baselines in a given configuration. The minimum scale discernible in the image is limited by these maximum baselines. A handy formula for the approximate resolution provided by an interferometer is:

$$\text{Interferometer Resolution} = \theta_{res} = k \lambda / L_{max}, \quad (3.27)$$

where k is a factor that depends on how the visibilities are weighted during inversion (typically ~ 1 ; see Figure 7.6) and L_{max} is the longest baseline in the array.

Another important limitation of interferometric array observations is insensitivity to large angular scales. This insensitivity arises because interferometric arrays alone cannot sample spatial frequencies lower than those sampled by a baseline equal to an antenna diameter. In effect, visibilities at locations on the (u, v) plane at or near its origin are not sampled, leading to the so-called *zero-spacing problem*. The lack of sensitivity to larger scale emission due to the zero-spacing problem biases the resulting image to the compact, small-scale emission of the true sky brightness distribution. As a guideline, the interferometer image has a *maximum recoverable scale* given roughly by:

$$\text{Maximum Recoverable Scale} = \theta_{MRS} \approx 0.6\lambda / L_{min}, \quad (3.28)$$

where L_{min} is the minimum baseline in the array configuration. Strictly speaking, for an input Gaussian visibility distribution of FWHM θ_{MRS} , the ratio of the brightness at source center of an image made by an array with L_{min} to the same made with no central hole in its visibility sampling (i.e., no zero spacing problem) is $1/e^1$. The smallest baseline possible in an array occurs when two antennas are adjacent to each other. Of

¹see Wilner, D.J., Welch, W.J. 1994, ApJ, 427, 898




Figure 3.6: Examples of spatial filtering using the CASA task *simobserve* on *representative* ALMA configurations for Cycle 8. *Panel a* (upper left): An optical image of the galaxy M51 used as a template for a true sky brightness distribution for the simulations. The frequency of the emission has been changed to 100 GHz, the image size has been scaled to $\sim 3' \times 3'$, and its declination has been changed to -40° to allow ALMA observations to be simulated. For the simulations, the galaxy was “observed” over a mosaic of 33 pointings, for ~ 10 hours in total. The resulting dirty images were CLEANed. *Panel b* (upper right): The high-resolution image of the galaxy obtained when observed in the ALMA C43-6 configuration with maximum baseline of 2516.9 m, respectively. The resulting synthesized beam is $\sim 0.47''$ and the maximum recoverable scale is $\sim 4''$. *Panel c* (lower left): Medium-resolution image of the galaxy when observed in the ALMA C43-4 configuration with maximum baseline of 783.5 m, respectively. The resulting synthesized beam is $\sim 1.3''$ and the maximum recoverable scale is $\sim 11''$. *Panel d* (lower right): Low-resolution image of the galaxy when observed in ALMA configuration C43-1 with maximum baselines of 160.7 m, respectively. The resulting synthesized beam is $\sim 4.3''$ and the maximum recoverable scale is $\sim 28''$.

course, the antennas cannot be moved physically closer together than their diameters. Note that in projection, antennas can appear closer than their diameters. In those cases, found typically when observing low elevation sources in compact configurations, the antenna in front partially blocks the reception by the antenna in the rear. The resulting visibilities are distorted; e.g., the antenna in the rear receives less power than the one in front. This situation is called *shadowing* and affected data are typically removed from the ensemble of observed visibilities.

Figure 3.6 illustrates the idea of spatial filtering using simulations of *notional* ALMA configurations. In Cycle 8, as in the three previous ALMA Cycles, ALMA aims to observe with at least 43 antennas on the 12-m Array. However, there are often times during the year where more than 43 antennas are available for PI science observations - even up to 50! As such, starting with Cycle 8, the “C43-X” convention for configurations has

been deprecated in favor of a "C-X" notation to account for the fact that more than 43 antennas may be used. See Chapter 7 and Table 7.1 for the full range of configurations offered for Cycle 8 and their corresponding resolutions. Simulations should still be done using the *notional* ALMA "C43-X" configuration files as that represents the *minimum* expected results from an observation. In practice, data will often be of higher quality than what the simulations predict.

The simulations were performed using the Common Astronomy Software Applications (CASA²) package, the suite of data reduction and analysis software used by ALMA. Here, panel *a* shows an optical image of the galaxy M51 used as an example of a true sky brightness distribution, after changing the frequency of the image to 100 GHz and placing the image center at $\delta = -40^\circ$, a declination easily observable with ALMA. Panels *b*, *c*, and *d* show the recovered images obtained by simulating 43-antenna observations of the galaxy using the CASA task *simobserve* in extended (C43-6), somewhat extended (C43-4), and compact (C43-1) configurations, respectively, and CLEANing. Angular resolutions of $\sim 0.47''$, $1.3''$, and $4.3''$ are obtained, respectively, and the corresponding maximum recoverable scales are $\sim 4''$, $\sim 11''$, and $\sim 28''$, respectively. Larger-scale emission from the galaxy has been filtered out in the extended configuration observations (panel *b*), leaving only the compact structures of its arms. On the other hand, these compact structures are not very discernible in the compact configuration observations (panel *d*). A reasonable compromise is found in the somewhat extended configuration observations (panel *c*), yet some small-scale detail and larger-scale emission remain missing. Note that though these images are missing angular scales, good science can still be obtained with them, as long as their limitations are properly understood. Combining data obtained from multiple configurations, or having more antennas, would increase the fidelity of recovered images.

3.7 Multi-configuration Observations

As previously discussed, a given configuration with baselines ranging from L_{min} to L_{max} is sensitive to angular scales from $\sim \theta_{MRS}$ to θ_{res} . Sensitivity to a broader range of angular scales is possible by combining data obtained in multiple configurations, where more extensive coverage of the (u, v) plane is attained. For example, the same source can be observed in different configurations of the 12-m Array, with more extended configurations providing higher angular resolutions. For sensitivity to extended structures, the more compact 12-m Array configurations or the more tightly clustered 7-m Array can be used. Finally, the individual Total Power (TP) Array antennas of the ACA can be used to map the largest angular scales and address the zero-spacing problem. In this section, we briefly describe so-called *multi-configuration observations*. Discussion of how ALMA implements such observations in practice can be found in Sections 7.8 and 7.9.

During proposal preparation, ALMA users should take note of the maximum recoverable scale needed to ensure that the proposed observations will be able to recover the scales needed to address the science in question. The ALMA Observing Tool will determine which combination of configurations will yield the desired angular resolution and maximum recoverable scale.

Data combination appears to work best when the signal-to-noise ratios (SNR) of the datasets are similar on overlapping spatial scales. Otherwise the noise in the lower-SNR dataset will corrupt the overlapping spatial scales in the map and degrade the image quality. In addition, the accuracies of the astrometry and calibration of the different datasets are important. Assuming the SNR of the individual datasets is high, combination is best done in the visibility domain rather than the image domain. By combining data in the visibility plane, the (non-linear) deconvolution has access to all available information, which tends to produce more robust, higher fidelity images.

Single-dish and interferometric data are typically combined by a technique called "feathering", which is described in greater detail in Chapter 7. To feather the single dish and interferometric data together, they are Fourier transformed to the spatial frequency domain; weighted appropriately, down-weighting high spatial frequency information in the single dish map and down-weighting low spatial frequency information in the interferometric map; added together; and inverse Fourier transformed back into a single image or cube. This procedure works best if the single-dish and interferometric data sample some common (overlapping) spatial scales, for example, if the single-dish antenna has a diameter which is twice the minimum baseline sampled by the interferometer. Multiple interferometer pointings (mosaics) can also in some circumstances partially recover missing low spatial frequency information in the absence of single dish data.

²For more information on CASA, see <http://casa.nrao.edu>

3.8 Units and Conversions

This chapter ends with a discussion of various units used in millimeter and radio astronomy and describe some useful conversions. Returning to the concept of specific intensity, this quantity can be described alternatively in terms of a temperature:

$$I_\nu(\theta, \varphi) = \frac{2k\nu^2}{c^2} T_B(\theta, \varphi). \quad (3.29)$$

In this equation, T_B is the *brightness temperature*, the temperature of a blackbody with the same specific intensity at a given frequency in the Rayleigh-Jeans limit, i.e., $h\nu/kT \ll 1$. Brightness temperature serves as an equivalent way of expressing the specific intensity of an astronomical source. The unit of brightness temperature is Kelvin (K).

In turn, brightness temperature can be included into the definition of flux density, S_ν (Equation 3.2), where

$$S_\nu = \frac{2k\nu^2}{c^2} \int T_B d\Omega. \quad (3.30)$$

Assuming the beam is Gaussian, one can then connect brightness temperature to flux density following:

$$\left(\frac{T}{1 \text{ K}} \right) = \left(\frac{S_\nu}{1 \text{ Jy}} \right) \left[13.6 \left(\frac{300 \text{ GHz}}{\nu} \right)^2 \left(\frac{1''}{\theta_{max}} \right) \left(\frac{1''}{\theta_{min}} \right) \right], \quad (3.31)$$

where ν is the observing frequency, and θ_{max} and θ_{min} are the major and minor axis of the synthesized beam, respectively. Note again that flux densities observed by ALMA are typically in units of Janskys, where $1 \text{ Jy} = 10^{-26} \text{ W m}^{-2} \text{ Hz}^{-1} = 10^{-23} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}$. The ALMA Observing Tool converts between temperatures and flux densities using these formulae.

An important point ALMA users must consider when proposing their projects is the dependence of brightness temperature sensitivity on synthesized beam size. Using Equation 3.30, one sees that an rms value in flux density (ΔS) can translate to an rms value in brightness temperature (ΔT), assuming a given synthesized beam size. Larger beam sizes correspond to lower ΔT , i.e., the surface brightness sensitivity increases. In turn, extended low surface brightness objects may be harder to detect at higher angular resolutions as the corresponding sensitivities may be too low. Typically, a compromise must be obtained between angular resolution and brightness sensitivity when planning interferometric observations.

3.9 Further Reading

If more information about interferometry is desired, the topic is covered in more detail in the following seminal texts:

- *Interferometry and Synthesis in Radio Astronomy - Third Edition*, by Thompson, A. R., Moran, J. M., & Swenson, G. W. (Springer), now available online at <https://link.springer.com/content/pdf/10.1007%2F978-3-319-44431-4.pdf>
- *Tools of Radio Astronomy - Sixth Edition*, by Wilson, T. L., Rohlfs, K., & Hüttemeister, S. (Springer), Fifth Edition now available for free online at <https://link.springer.com/content/pdf/10.1007%2F978-3-540-85122-6.pdf>
- *Synthesis Imaging in Radio Astronomy II*, PASP Conference Series, Vol. 180, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco - ASP)
- *Millimeter Interferometry: Proceedings of the IRAM Millimeter Interferometry School*, ed. A. Dutrey, available online at <http://iram.fr/IRAMFR/IS/IS2002/archive.html>.

Chapter 4

Receivers

The ALMA front end can accommodate up to 10 receiver bands covering most of the wavelength range from 8.5 to 0.32 mm (35–950 GHz). Each band is designed to cover a tuning range which is approximately tailored to the atmospheric transmission windows. These windows and the tuning ranges are outlined in Figure 4.1. This illustrates the broad, deep absorption features, mostly due to H₂O in the lower few km of the atmosphere, as well as some O₂ transitions. The many narrow features seen in this plot are mostly from stratospheric O₃, along with some transitions of CO and other trace species. In Cycle 8, Bands 3, 4, 5, 6, 7, 8, 9, and 10 are available, and the basic characteristics of the bands are outlined in Table 4.1. Each of the ALMA receiver bands is described in more detail in the following sections as well as in the references listed in Table 4.2.

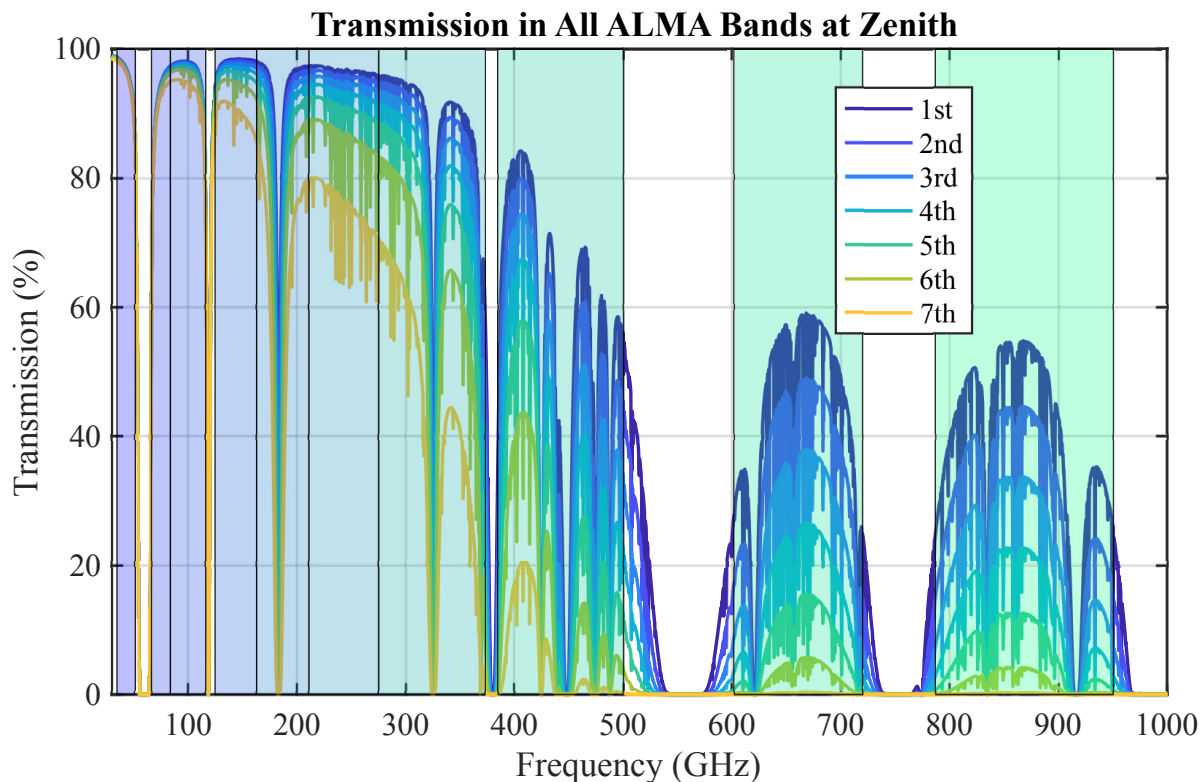


Figure 4.1: The ten ALMA receiver bands along with atmospheric transmission. The receiver coverage is shown shaded, superimposed on a zenith atmospheric transparency plot at the Array Operation Site (AOS) for the first seven octiles of observing conditions at ALMA (see Table 9.1).

The ALMA receivers in each antenna are situated in a single front-end assembly (see Appendix A, Section

Band	Frequency/ Wavelength range (GHz) ¹ /(mm)	LO range (GHz)	Sideband mode ²	IF range (GHz)	Inst. IF bandw. (GHz) ⁴	T_{rx} over 80% of band (K) ⁶	T_{rx} at any frq. (K) ⁶
3	84.0 - 116.0/ 2.59 - 3.57	92 - 108	2SB	4-8	7.5	<39 ⁷	<43 ⁷
4	125.0 - 163.0/ 1.84 - 2.40	133 - 155	2SB	4-8	7.5	<51	<82
5	158.0 - 211.0/ 1.42 - 1.90	166 - 203	2SB	4-8	7.5	<55	<75
6	211.0 - 275.0/ 1.09 - 1.42	221 - 265	2SB	4.5-10 ³	7.5	<83	<136
7	275.0 - 373.0/ 0.80 - 1.09	283 - 365	2SB	4-8	7.5	<147	<219
8	385.0 - 500.0/ 0.60 - 0.78	393 - 492	2SB	4-8	7.5	<196	<292
9	602.0 - 720.0/ 0.42 - 0.50	610 - 712	DSB	4-12	7.5(15) ⁵	<175 (DSB)	<261 (DSB)
10	787.0 - 950.0/ 0.32 - 0.38	795 - 942	DSB	4-12	7.5(15) ⁵	<230 ⁸ (DSB)	<344 (DSB)

Table 4.1: ALMA Receiver Specifications. *Notes to Table:* **1.** Frequency ranges are the maximum available, at the extreme upper and lower limits of the IF passband. In reality, because of filter roll-off, the coverage is ~ 60 MHz less (See Section 6.4) **2.** Sideband modes: 2SB means dual sideband receiver where the two sidebands are available simultaneously, DSB means double sideband receiver. See text for details. **3.** Usable IF range is extended to allow simultaneous observations of multiple lines specifically of Carbon Monoxide (CO). However, the auto-correlation noise performance is degraded by a factor of up to about 1.2 below 5.5 GHz (Section 4.2.4) **4.** Maximum instantaneous IF bandwidth, limited by the back-end filters and spectrometers. As both upper and lower sidebands both pass through the same IF bandwidth but are subsequently separated, the effective signal bandwidth given in this column for 2SB receivers is twice the actual IF filter bandwidth. In addition, this is per polarization, so the total effective bandwidth for each receiver is then another factor of 2 higher. Note that the effects of the anti-aliasing filters have been included (see Section 6.4). **5.** In Cycle 8, the maximum bandwidth is approximately double in cross-correlation mode, because both sidebands can be separated and correlated using 90-degree phase switching (see Section 6.3.4 and Section B.4.4). **6.** The *maximum* specified SSB receiver temperatures (T_{rx}) are given, unless otherwise noted. These values are the average over the whole IF band. In many cases, the average realised results are better than specifications; the sections on individual receiver bands describe the typical values measured. The numbers adopted in the observing tool (OT) and ALMA sensitivity calculator (ASC) are conservative values of the average value over the band. **7.** The specification for Band 3 receivers are described slightly differently to other bands: at LO1=104GHz, $T_{rx}<39$ K (including warm optics), and $T_{rx}<43$ K for any other valid LO setting. **8.** The specification for Band 10 receivers is $T_{rx}<230$ K within a selected 80 % portion of that band (787-950 GHz).

Topic	Author/Year	Technical papers or Meeting proceedings	ADS identifier
B3	Claude et al. 2008	SPIE 7020	2008SPIE.7020E..33C
B4	Asayama et al. 2014	PASJ, 66 (3), 57(1-13)	2014PASJ...66...57A
B5	Belitsky et al. 2017	A&A (doi=10.1051/0004-6361/201731883; in press)	
B6	G. A. Ediss, et. al 2004	15th Intl Symp Space Terahertz Tech	2004stt..conf..181E
B7	Mahieu et al. 2012	Trans. THz Sci. and Tech., 2(1) 29-39	2012ITTST...2...29M
B8	Sekimoto et al. 2008	19th Intl Symp Space Terahertz Tech	2008stt..conf..253S
B9	Baryshev et al. 2007	19th Intl Symp Space Terahertz Tech	2008stt..conf..258B
B10	Uzawa et al. 2009	20th Intl Symp Space Terahertz Tech	2009stt..conf...12U
Optics	Rudolf et al. 2007	IEEE Trans. on Antennas & Propagation	2007ITAP...55.2966R
WVR	Emrich et al. 2009	20th Intl Symp Space Terahertz Tech	2009stt..conf..174E

Table 4.2: Technical papers describing the receiver bands, optics and the water vapor radiometer.

A.4). The front-end assembly consists of a large cryostat containing the receiver cold cartridge assembly (CCA) for each band (including the Superconductor-Insulator-Superconductor (SIS) mixers and Local Oscillator (LO) injection) and the Intermediate Frequency (IF) and LO room-temperature electronics of each band (the warm cartridge assembly, WCA). The cryostat is kept at a temperature of 4 K through a closed cycle cooling system. The Amplitude Calibration Device (ACD) used for ambient/hot load calibration is mounted above the front end (see Appendix A.5). Each receiver cartridge contains two complete receiving systems sensitive to orthogonal linear polarizations. The designs of the mixers, optics, LO injection scheme, and polarization splitting vary from band to band, depending on the optimum technology available at the different frequencies; each receiver is described in more detail in the sections below.

Up to three bands can be switched on at a time (more would risk overloading the cryostat cooler). From a hardware point of view, it takes only about 1-2 seconds to switch between two bands which are powered up (limited by the LO locking and antenna movement). For bands that are not switched on, it takes around a minute to power-up the band, after which it is necessary to set a nominal LO power and allow around 5 minutes for thermal stabilization before science scans commence¹.

4.1 Local Oscillators and IF Ranges

The observed sky frequencies need to be mixed down to frequencies in the range 2 to 4 GHz in order to send the digitized signals to the correlator (see Section 6.1). This frequency downconversion involves a set of LOs, and the LO and IF systems are described in detail in Appendix B.

The front-end mixer uses LO1 to down-convert the sky frequencies into an IF band with a range of 4–12 GHz. This covers the needs of all the ALMA bands, since the mixers for Bands 3, 4, 5, 7, and 8 have an output range of 4–8 GHz, Band 6 a range of 4.5–10 GHz and Bands 9 and 10 a range of 4–12 GHz (Table 4.1). The possible sky frequency ranges covered by each receiver with the first LO, LO1, set to a frequency F_{LO1} are:

- For the lower sideband (LSB): $(F_{LO1} - IF_{lo})$ to $(F_{LO1} - IF_{hi})$
- For the upper sideband (USB): $(F_{LO1} + IF_{lo})$ to $(F_{LO1} + IF_{hi})$

where IF_{lo} and IF_{hi} are the lower and upper IF ranges in the “IF Range” column of Table 4.1, and the IF bandwidth (per sideband) is $IF_{hi} - IF_{lo}$. The ranges of LO1 are given in column 3 of Table 4.1. This is illustrated in Figure 4.2. Note that the maximum IF bandwidth in Table 4.1 may be a few percent less than the IF range in Table 4.1 (see Section 6.4).




Figure 4.2: IF ranges for the two sidebands in a heterodyne receiver.

¹The original ALMA specifications dictate that after 15 minutes of thermal stabilization time the receivers should meet all the ALMA technical specifications, although in practice this has been found to be excessively large, as gain drifts calibrate out in real observations.

4.2 The Cycle 8 Receivers

The receivers for Bands 3 to 8 are two-sideband (2SB) receivers, where both the upper and lower sidebands are provided separately and simultaneously. There are 4 outputs from each of the receivers, comprising the upper and lower sidebands in each of the two polarizations. Each output has a bandwidth of 4 GHz (slightly more for Band 6) which is reduced to an effective total bandwidth of 3.75 GHz due to the anti-aliasing filters, etc., see Appendix B.3.5. The mixers give 10 dB or more unwanted sideband rejection, which is adequate for reducing the degradation of S/N from noise in the unwanted sideband, but not adequate for suppressing astronomical signals in the unwanted sideband. Further suppression is performed by offsetting LO1 and LO2 (and eventually the tunable filter bank LO, TFB LO) by small and opposite amounts, as well as 180-degree phase-switching, the offsets of which depend on the antenna such that the signals from two antennas in the image sideband do not correlate (see Chapter 6).

Bands 9 and 10 use double sideband (DSB) receivers, where the IF contains noise and signals from both sidebands. They only have two outputs, one per polarization. The IF effective bandwidth is 7.5 GHz per polarization (after passing through the IF processing units), so the total instantaneous receiver bandwidth is the same as Bands 3–8. However, both sidebands can be simultaneously correlated using 90-degree phase switching (see Chapter 6), although this does not remove the noise from the opposite sideband. The effective system bandwidth available in Bands 9 and 10 is therefore double that of the 2SB receivers (15 GHz).

Each of the ALMA receiver bands is different in several aspects, and the following sections describe the individual bands in more detail.

Figure 4.3 shows the layout of the ALMA receiver bands in the ALMA cryostat.

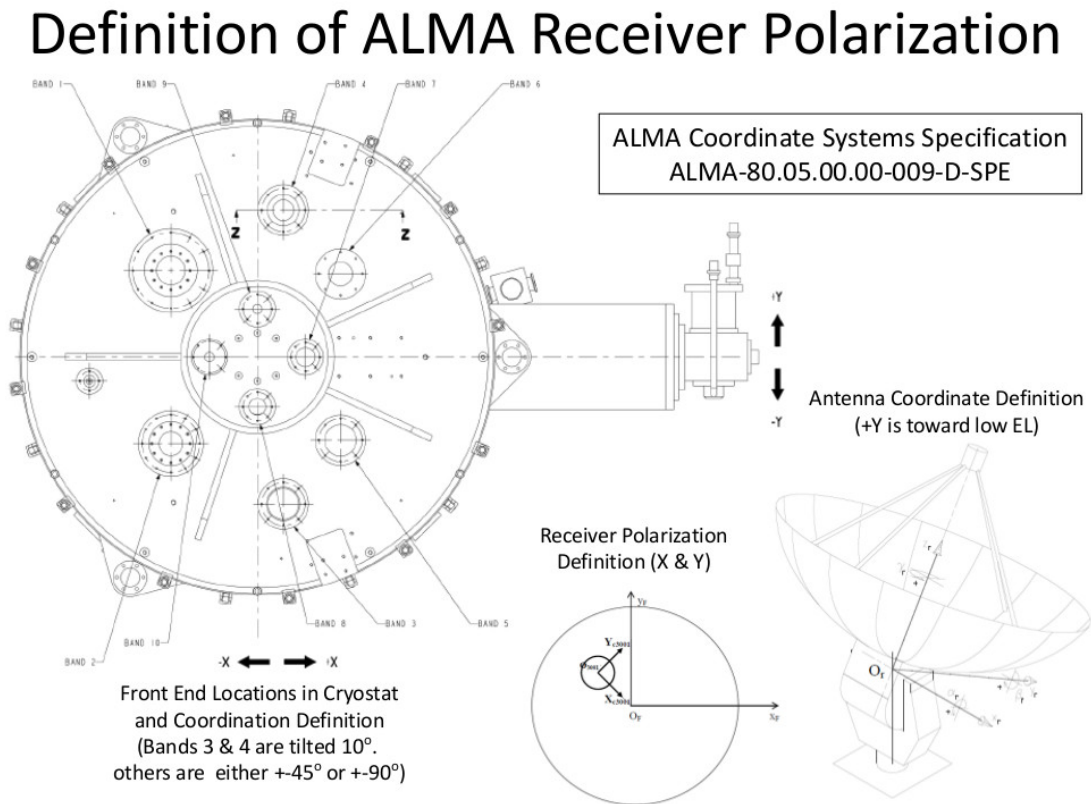


Figure 4.3: The layout of the ALMA receiver bands within the ALMA cryostat. Also shown is the orientation of the receiver polarization position angles.

4.2.1 Band 3 Receiver

Band 3 is the lowest frequency band available in Cycle 8, covering 84 to 116 GHz of spectral range (the 3 mm atmospheric window). The cartridge is fed by a “periscope” pair of ellipsoidal pickoff mirrors located outside the cryostat, which refocus the beam through the cryostat window, allowing for a smaller window diameter (Figure 4.4). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the incoming signal into two linear orthogonal polarizations and feeds the SIS mixers.




Figure 4.4: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the solid line, and the Cassegrain focus is shown by the small circle to the upper right.

A block diagram of the Band 3 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.5. The CCA contains the corrugated feedhorn, OMT, SIS mixers and the low-noise High Electron Mobility Transistor (HEMT) first IF amplifiers. At room temperature, the WCA includes further IF amplification and the LO covering 92–108 GHz.

The acceptance criteria for the Band 3 receiver noise performance (T_{rx}) is <39 K at LO1=104 GHz, and <43 K for any other valid LO setting (SSB T_{rx}). This includes a 2 K component due to the losses in the external warm optics. The achieved median performance is typically somewhat better than this, with $T_{rx} \sim 37 \pm 6$ K. The OT assumes $T_{rx}=40$ K. The atmospheric transmission over most of Band 3 is very high, even with a large PWV (Figure 4.6) which means observations in Band 3 can (and do) take place with up to 10 mm or more of PWV. The system temperature (T_{sys}) shows the expected rise at the higher end due to the edge of an atmospheric oxygen line (Figure 4.7), so this increase is mostly independent of the PWV.

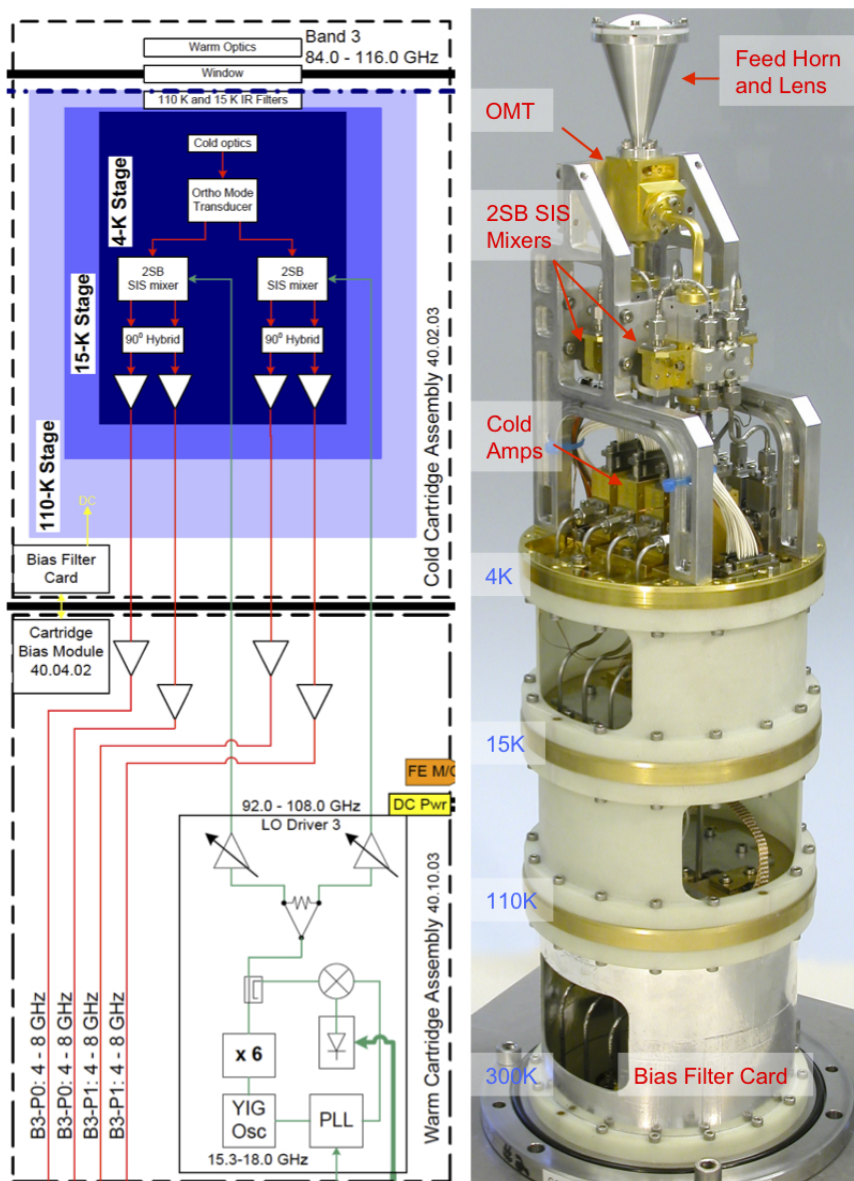


Figure 4.5: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges were constructed in Canada at the National Research Council – Herzberg Institute of Astrophysics (NRC-HIA) in Victoria, British Columbia.

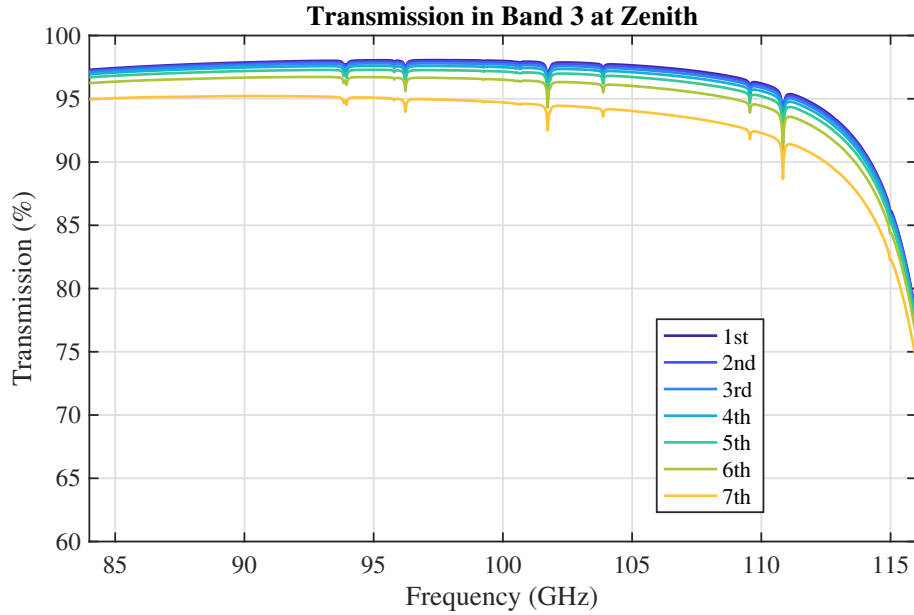


Figure 4.6: Band 3 zenith transmission for the first seven octiles of observing conditions at ALMA. The decrease at the high frequency end is due to the edge of an O_2 absorption line at 118.75 GHz.

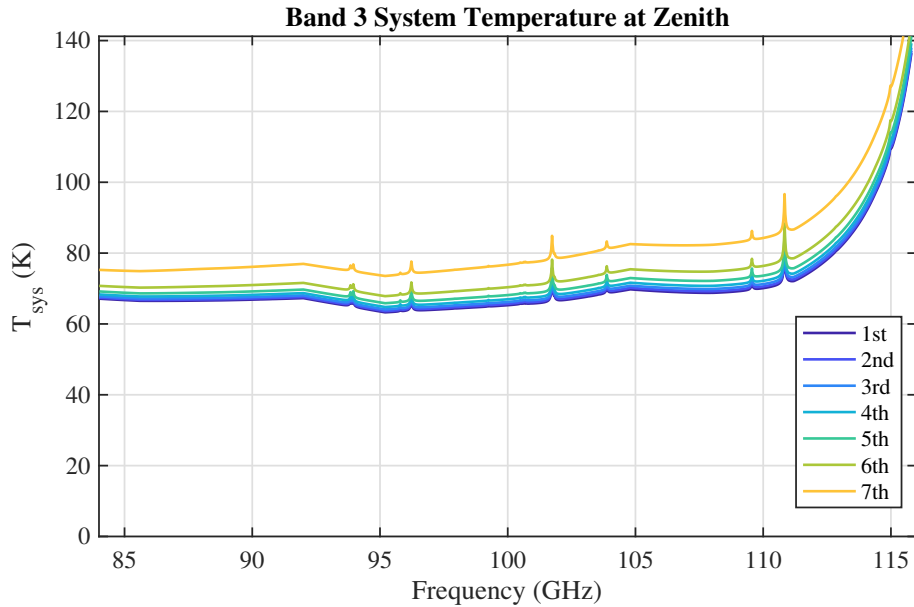


Figure 4.7: Typical system temperature (T_{sys}) at zenith for Band 3 for the first seven octiles of observing conditions at ALMA. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

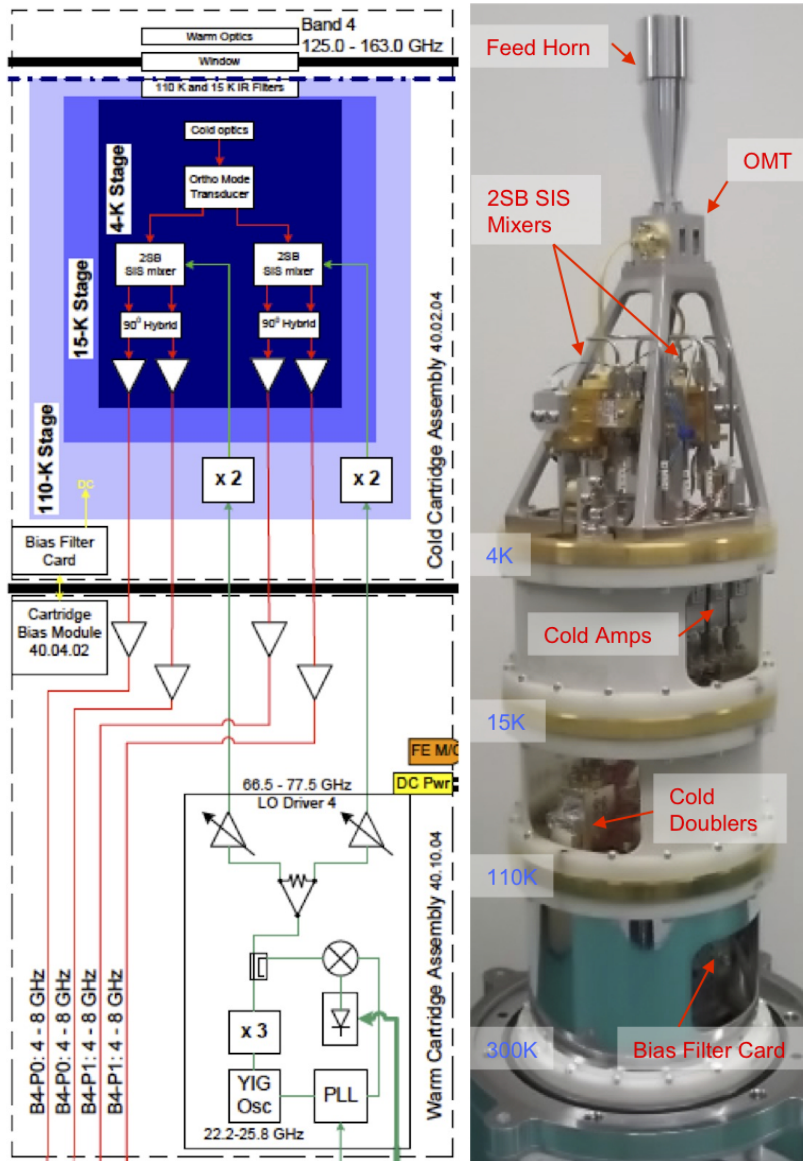


Figure 4.9: Block diagram of the Band 4 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 4 CCA. A single feedhorn feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The LO is generated in the WCA, with the final x2 multiplier on the 110 K stage. The Band 4 cartridges were constructed in Japan at the NAOJ Advanced Technology Center (ATC) in Mitaka, Japan.




Figure 4.10: Band 4 zenith transmission for the first seven octiles of observing conditions at ALMA.

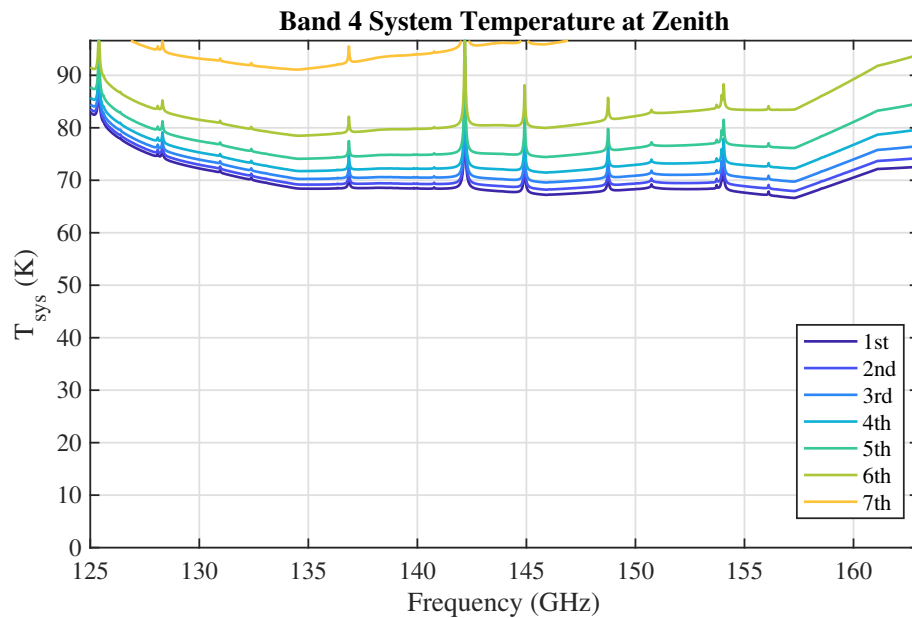


Figure 4.11: System temperature (T_{sys}) at zenith for Band 4 for the first seven octiles of observing conditions at ALMA. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.3 Band 5 Receiver

The Band 5 receiver covers the 158 to 211 GHz spectral range (the 1.7 mm atmospheric window), slightly extended from the original 163–211 GHz coverage. The signal collected by the telescope is focused to the Band 5 cartridge using a set of cold mirrors (Figure 4.12). A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the 2SB SIS mixers.




Figure 4.12: Optical layout of the Band 5. Placing the mixer assembly along the side of mirror M1 and the streamlined layout of the IF system are the key to fulfill the requirements on $5w_o$ optical beam clearance at the entrance to the cryostat.

A block diagram of the Band 5 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.13. The Band 5 CCA contains a feed horn, an OMT as a polarization splitter, 2SB SIS mixer assemblies, cold IF amplifiers, isolators, and LO frequency doubler. The RF signal is down converted to 4–8 GHz using a 2SB mixer unit.

The atmospheric transmission in Band 5 is shown in Figure 4.14 for different PWV values. The specification for Band 5 receiver noise performance (T_{rx}) is <55 K over 80% of the 163–211 GHz range, and <75 K over the whole 158 - 211 GHz extended range (SSB T_{rx}). However, the realized performance of the receivers are often better than 50 K over the band. The OT assumes 50 K. The resulting system temperatures (T_{sys}) are shown in Figure 4.15.

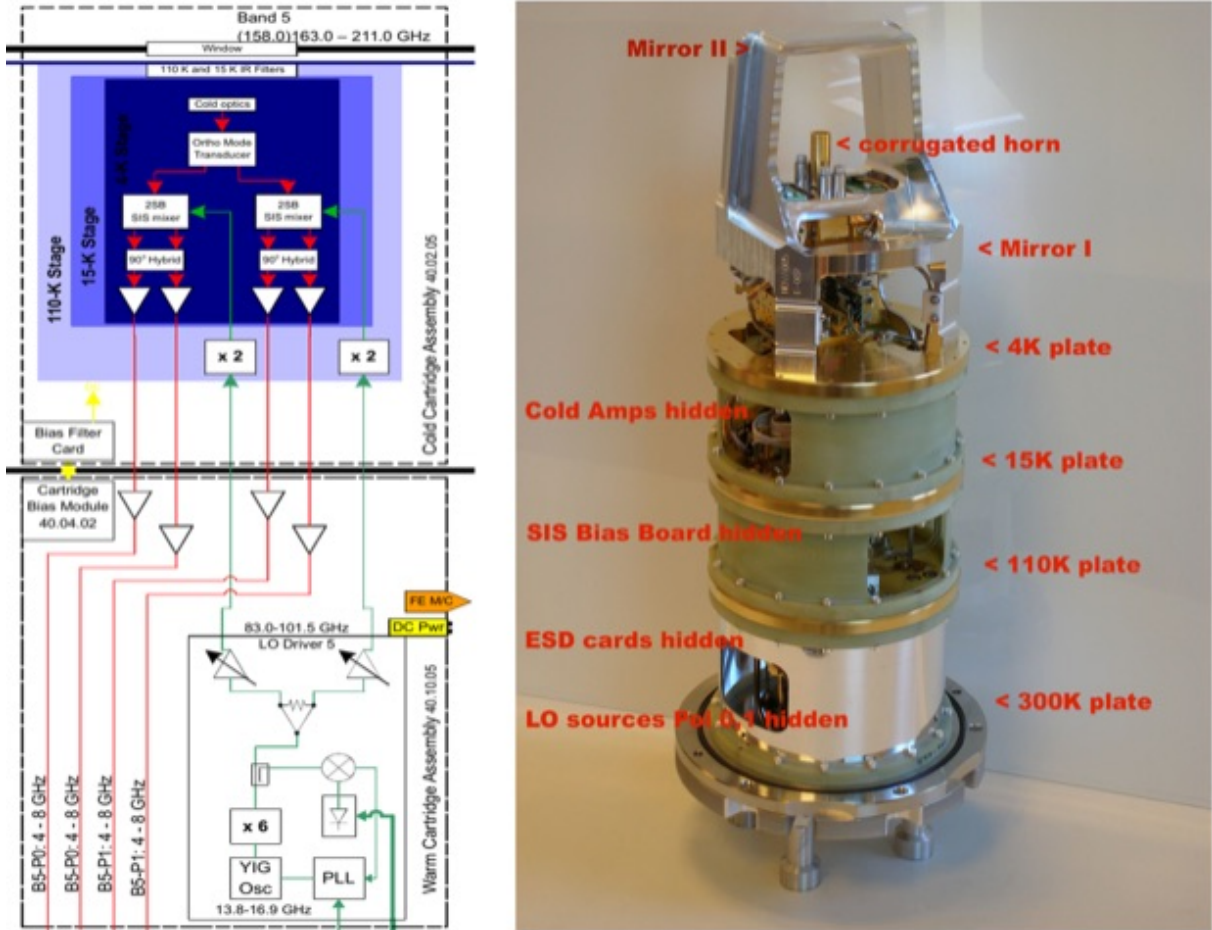


Figure 4.13: Block diagram of the Band 5 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 5 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The production design has a room-temperature $\times 6$ multiplier in the WCA, then a $\times 2$ cold multiplier in the CCA (at 110K). The Band 4 cartridges were constructed in Sweden by the Group for Advanced Receiver Development (GARD) at Onsala Space Observatory, Chalmers University of Technology, Sweden.




Figure 4.14: Band 5 zenith transmission for the first seven octiles of observing conditions at ALMA. The strong absorption due to H_2O at 183.31 GHz allows for observations at this frequency only when the PWV is $< 1\text{mm}$. Continuum observations in Band 5 should avoid this part of the band if possible (See Chapter 6 – Spectral Setups)




Figure 4.15: Typical system temperature (T_{sys}) at zenith for Band 5 for the first seven octiles of observing conditions at ALMA. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.4 Band 6 Receiver

The Band 6 receiver covers the 211 to 275 GHz spectral range (the 1.3 mm atmospheric window). This receiver has a window with a pair of off-axis ellipsoidal mirrors inside the cryostat (Figure 4.16). A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the SIS mixers. A block diagram of the Band 6 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.17.




Figure 4.16: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocking filter. Note the slightly inclined inner window, designed to minimize standing waves.

The Band 6 allowed IF frequency range was 5–10 GHz, but has been extended further since Cycle 6 to allow for multiple simultaneous line observations with more bandwidth per spw²; it now covers the range 4.5–10.0 GHz. Although there is up to ~10-20% excess receiver noise below 5.5 GHz due to LO1 (but the increase of T_{sys} is less), the CO multi-transition setup is still considerably more efficient than observing each line separately. It is recommended that for continuum observations, the IF range 6–10 GHz is still used. Also, it should be noted that the full range 4.5–10 GHz cannot be completely sampled because of the limited 4 GHz width of the two basebands per polarization.

The atmospheric transmission in Band 6 is shown in Figure 4.18 for different PWV values. Most of the narrow absorption lines are from ozone. The specification for Band 6 receiver noise performance (T_{rx}) is <83 K over 80% of the band, and <136 K over the whole band (SSB T_{rx}). The achieved on-array results are considerably better, typically ~40–50 K over most of the band. The OT assumes 50 K. The resulting system temperatures (T_{sys}) for different PWV values are shown in Figure 4.19.

²Specifically, the $^{12}\text{CO}/^{13}\text{CO}/\text{C}^{18}\text{O}$ $J=2-1$ combination at 230.538/220.398/219.560 GHz, which has a minimum separation of 10.14 GHz. Since Cycle 6, the lower end of the allowed IF range has been extended to 4.5 GHz in order to cover all three lines with broader spws.

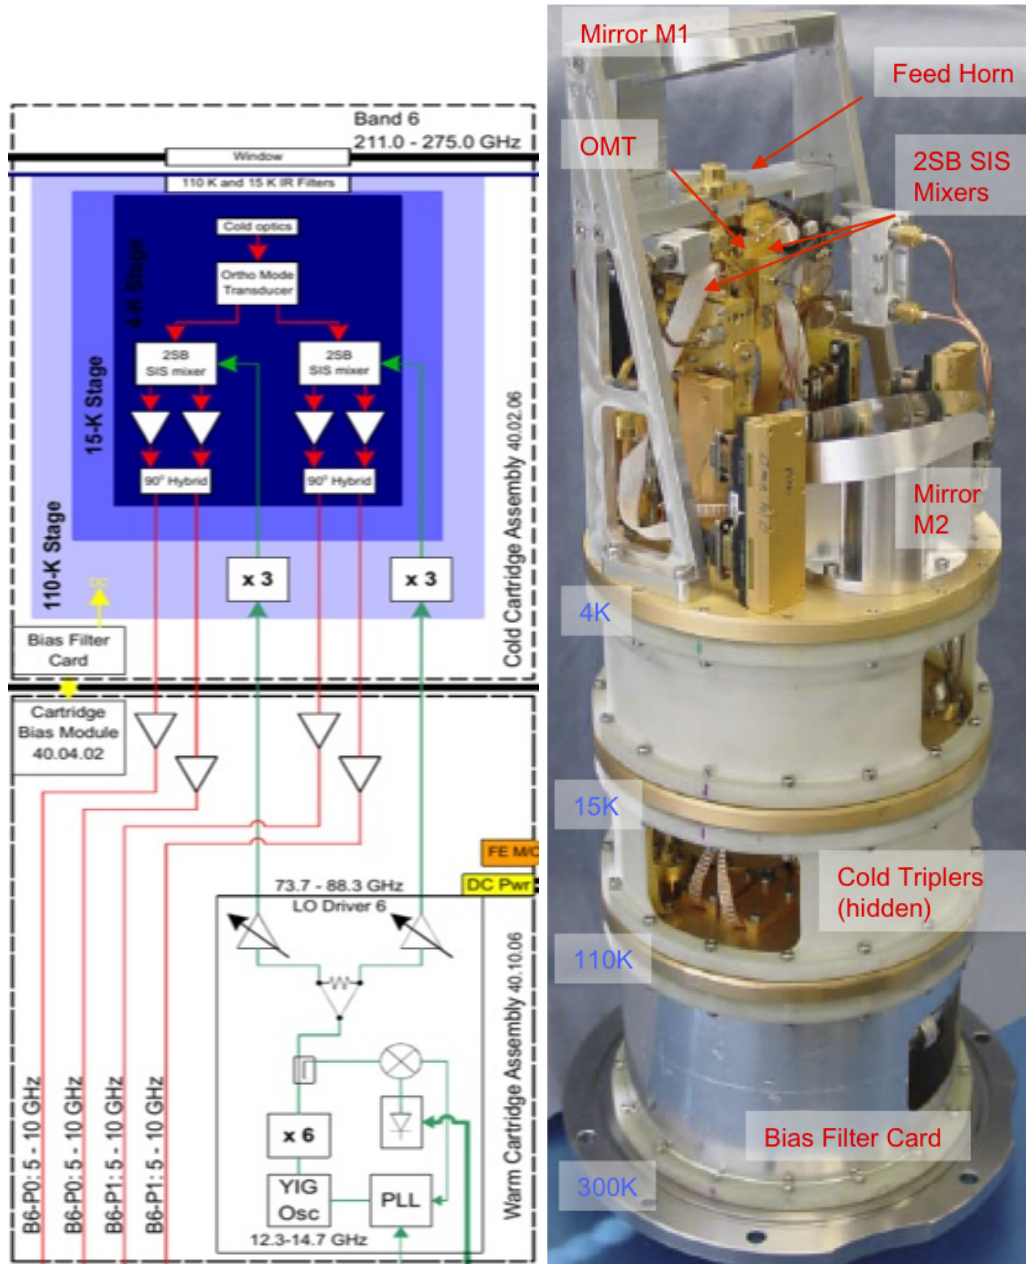


Figure 4.17: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra $\times 3$ multiplier inside the cryostat. Note that the IF output range has been extended to cover 4.5–10 GHz; the range shown is the one recommended for continuum observations (see text). The Band 6 cartridges were built in the United States at the NRAO Technology Center in Charlottesville, Virginia.




Figure 4.18: Band 6 zenith transmission for the first seven octiles of observing conditions at ALMA. Most of the narrow absorption lines are from ozone.




Figure 4.19: Typical system temperature (T_{sys}) at zenith for Band 6 for the first seven octiles of observing conditions at ALMA, based on measured values of the receiver temperatures. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.5 Band 7 Receiver

The Band 7 receiver covers the 275 to 373 GHz spectral range (the 0.85 mm atmospheric window). It has a similar cold optics design as Band 6, but uses a wire-grid polarization splitter instead of an OMT (Figure 4.20). A block diagram of the Band 7 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.21.




Figure 4.20: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid.

The atmospheric transmission in Band 7 is shown in Figure 4.22 for different PWV values. The original specification of the Band 7 receiver noise temperature was $T_{rx} < 147$ K over 80% of the range and < 219 K over the whole tuning range, except at the upper end of the band (370-373 GHz), where the specifications were < 300 K (SSB T_{rx}). However, the performance of the receiver as measured in the lab and on the array is considerably better than this, with typically 65 K achieved, in mid-band. The OT assumes 72 K over the whole band. The resulting system temperatures (T_{sys}) for different PWV values are shown in Figure 4.23. Note that the atmospheric transmission (and hence T_{sys}) at frequencies below 300 GHz is considerably better than that of the top half of Band 7; in that respect the performance is closer to that of Band 6. Also to maximise continuum sensitivity, the deep water absorption lines around 325 GHz should be avoided if possible.




Figure 4.21: Band 7 front-end receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built in France at the Institut de Radio Astronomie Millimétrique (IRAM) in Grenoble, France.

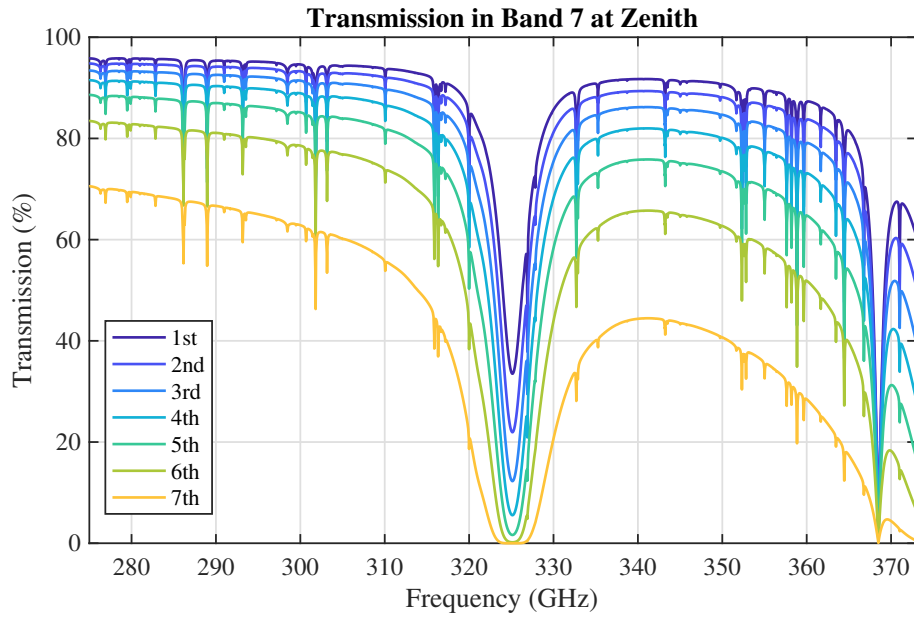


Figure 4.22: Band 7 atmospheric zenith transmission for the first seven octiles of observing conditions at ALMA. The deep atmospheric absorption at 325 GHz is due to water, and the narrower feature at 369 GHz is due to oxygen.




Figure 4.23: Typical T_{sys} at zenith for Band 7 for the first seven octiles of observing conditions at ALMA. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.6 Band 8 Receiver

The Band 8 receiver covers the 385 to 500 GHz spectral range (the 650 μm atmospheric window). The cryogenic optics of this receiver adopts a single mirror to couple a feed horn in front of an SIS mixer block to the sub-reflector. A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the 2SB SIS mixers (Figure 4.24).




Figure 4.24: Optical layout of the Band 8 receiver.

A block diagram of the Band 8 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.25. The Band 8 CCA consists of a cold optics, a feed horn, an OMT, 2SB SIS mixers assemblies, cold IF amplifiers, isolators, and LO frequency sextuplers.

The atmospheric transmission in Band 8 is shown in Figure 4.26 for different PWV values. The specification of the Band 8 receiver noise temperature is $T_{\text{rx}} < 196$ K over 80% of the range and < 292 K over the whole tuning range (SSB T_{rx}). However, the performance of the receiver as measured in the lab and on-array is considerably better than this, with typical values of 70-120 K. The OT assumes 135 K. The resulting system temperatures (T_{sys}) for different PWV values are shown in Figure 4.27. Note that transmission in the lower part of the band is almost as high as Band 7, whereas at the high frequency end of the band (e.g. around the CI line at 492 GHz) the transmission is as low as the center of Band 9.

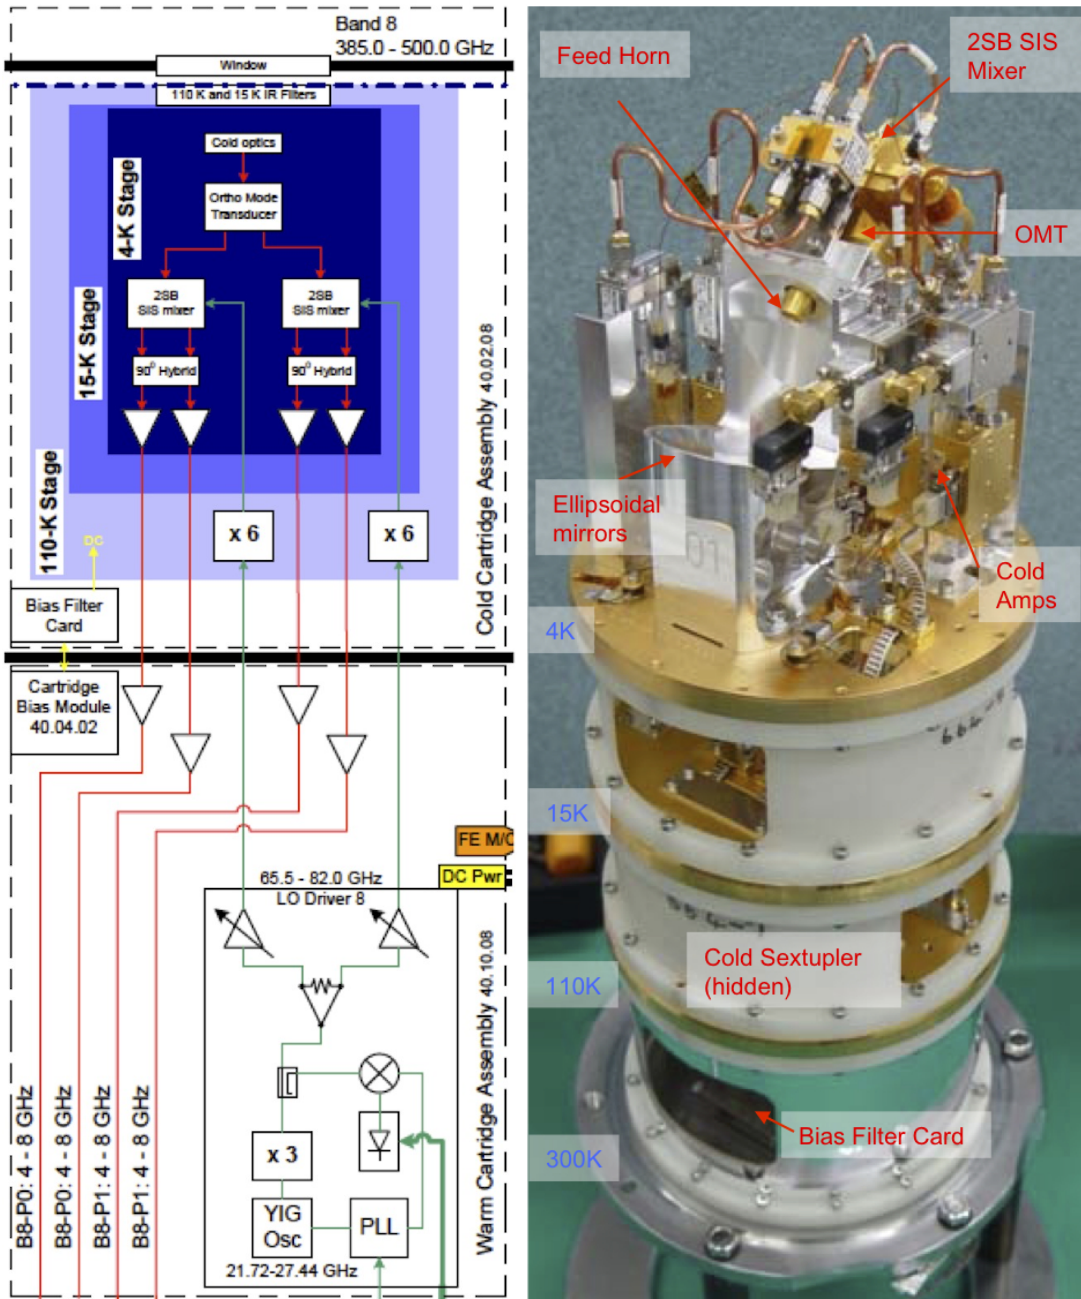


Figure 4.25: Block diagram of the Band 8 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 8 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 8 cartridges were constructed in Japan at the NAOJ Advanced Technology Center (ATC) in Mitaka, Japan.




Figure 4.26: Band 8 atmospheric zenith transmission for the first seven octiles of observing conditions at ALMA. The atmosphere in the Band 8 frequency range has some deep absorption features due to water and oxygen.




Figure 4.27: Typical T_{sys} at zenith for Band 8 for the first seven octiles of observing conditions at ALMA. Note that some of the values for T_{sys} in the worst octiles, which should never be used for high-frequency observations, are off scale in this plot. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.7 Band 9 Receiver

The Band 9 receiver covers the 602 to 720 GHz spectral range (the 450 μm atmospheric window). It uses a wire grid in order to separate the two orthogonal polarizations, as well as to provide the LO injection scheme (Figure 4.28).




Figure 4.28: Basic Band 9 optics layout. The signal path is symmetrical for the vertical polarization (P0) and the horizontal polarization (P1). For P1 the path is as follows: the telescope focal point (FP) is followed by mirror M3, wire grid polarization splitter, mirror M4, a beam splitter for LO insertion, and finally the mixer horn U1. For P0 the signal follows from FP to the same mirror M3 and then, reflected by the grid, comes to mirror M4, a beam splitter, and mixer horn U1'.

The mixers are double sideband (DSB), and therefore additional techniques must be employed during the observations to either separate the sidebands or reject the unwanted sideband. LO offsetting can be used to reject one of the two sidebands, which can be chosen independently for each spectral window (spw). Note that LO offsetting does not reject the noise from the unwanted sideband, it simply moves any correlated signal to a high fringe rate so that the signal is smeared over a larger bandwidth increasing noise incoherently. Since Cycle 6, 90-degree phase switching is used to correlate both USB and LSB simultaneously (see Chapter 6). The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 6.4), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.29.

The Band 9 atmospheric transmission is strongly dependent on the PWV, as illustrated in Figure 4.30, where only the lower values of PWV have significant transmission. The specifications for the receiver are $T_{\text{rx}} < 175$ K over 80% of the band and < 261 K over all the band. However, the performance is considerably better than this, with engineering and typical on-array values for T_{rx} of 65–120 K. The OT assumes 105 K. Figure 4.31 shows the expected T_{sys} for different PWV values using the given OT default expected receiver noise values. However, the achieved T_{sys} is extremely dependent on the line-of-sight PWV. Phase stability also limits when observations can be made, therefore, most observations in Band 9 are done at night and early morning (and more commonly during austral winter). As well as having a lower atmospheric transmission and a less stable

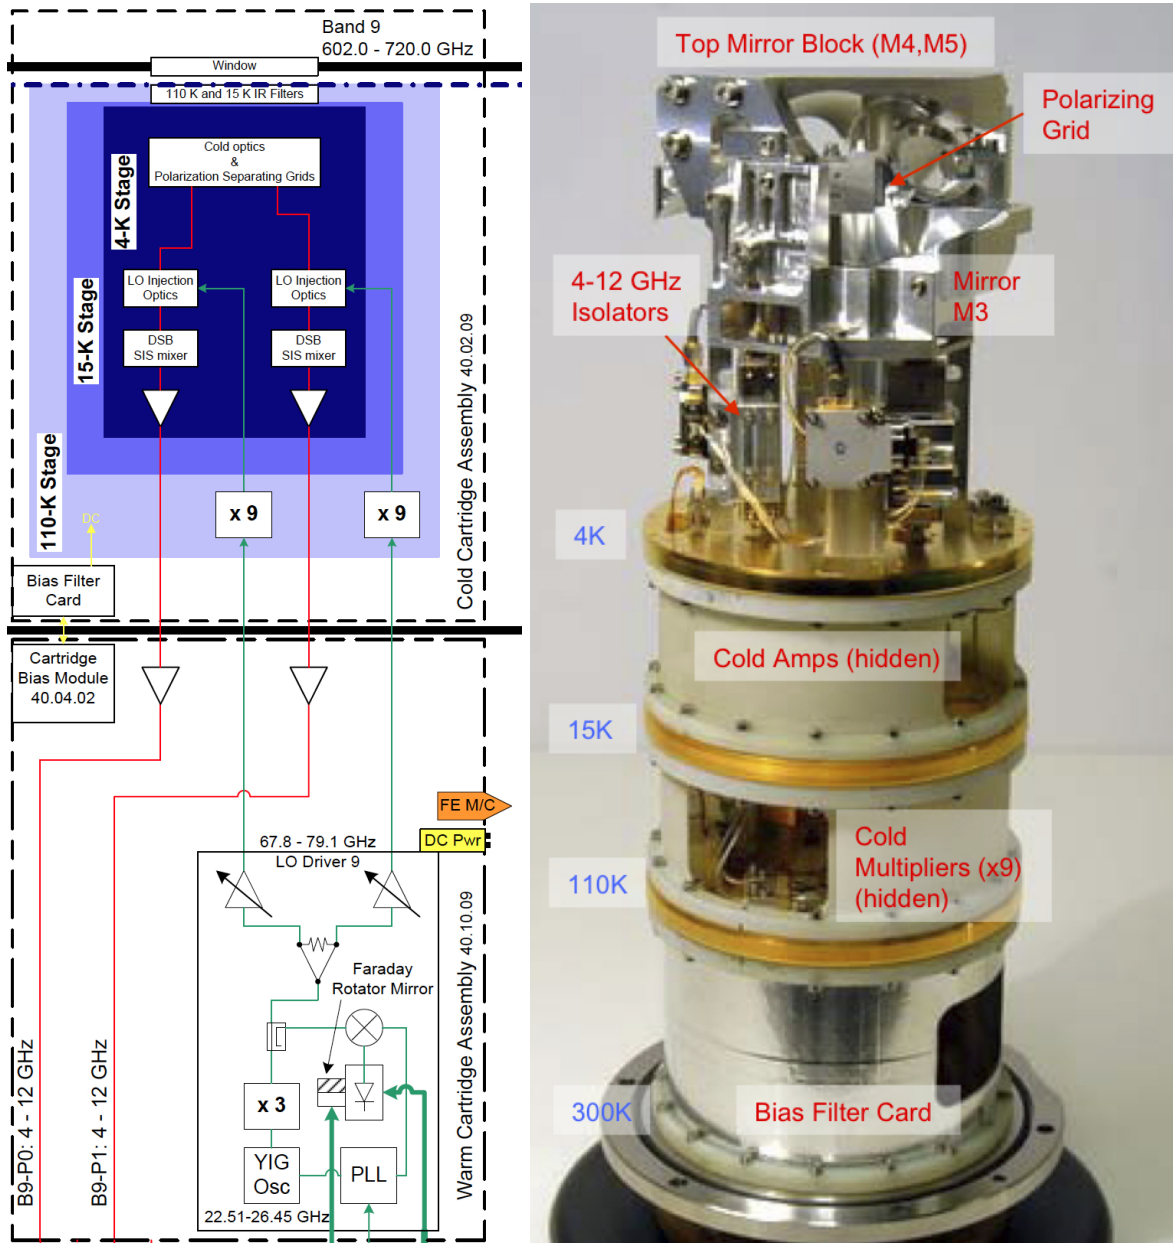


Figure 4.29: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The extra Faraday rotation mirror in the LO system is part of the ALMA fiber optic Line Length Corrector (LLC) system (see Appendix B.3.3), and means that the Band 9 receiver must be available on all the antennas for this to work. The Band 9 receiver was built in the Netherlands at the Nederlandse Onderzoekschool voor Astronomie (NOVA) in Leiden, Netherlands.

atmosphere, Band 9 observing provides additional challenges, including finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing for the relatively small primary beam (although the actual pointing is done in Band 6), and the need for the highest level of stability in the rest of the system.




Figure 4.30: Band 9 zenith transmission for the first seven octiles of observing conditions at ALMA. Most of the band is affected by the wings of pressure-broadened H₂O at 620 GHz.




Figure 4.31: Typical T_{sys} at zenith for Band 9 for the first seven octiles of observing conditions at ALMA. Note that the values for T_{sys} in the worst octiles, which should never be used for high-frequency observations, are off scale in this plot. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.8 Band 10 Receiver

The Band 10 receiver covers the 787 to 950 GHz spectral range (the 350 μm atmospheric window). Band 10 is the highest frequency receiver of the ten bands envisioned for the ALMA front-end system. The development of the Band 10 receiver was extremely difficult and faced many technical challenges from its material selection. Niobium (Nb) superconducting tuning circuits, which are used in other ALMA receiver bands, cannot be used for Band 10 SIS mixers due to large losses from pair-breaking above a superconducting gap frequency of about 700 GHz. Therefore, niobium-titanium-nitride (NbTiN) with a critical temperature of about 15 K, has been utilized in the tuning circuit of Band 10 mixers. The Band 10 Nb/AlO_x/Nb tunnel junctions with NbTiN-based tuning circuitry achieved ALMA requirements and the best DSB receiver noise temperature was 125 K, corresponding to about 3 times the quantum limit for 4 K operation.

It uses a wire grid in order to separate the two orthogonal polarizations, as well as to provide the LO injection scheme (Figure 4.32). The mixers are double sideband (DSB), and therefore LO offsetting can be used to reject one of the two sidebands. As in the case of Band 9, 90-degree phase switching is used to correlate both USB and LSB simultaneously (see Chapter 6). The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 6.4), covering 4–12 GHz. A block diagram of the Band 10 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.33.




Figure 4.32: Schematic of ALMA Band 10 optics. ALMA Band 10 optics are composed of two elliptical mirrors, M1 and M2, a wire grid and two corrugated horns. The wire grid is used to separate the two linear polarizations, P0 and P1, and it is located after the two elliptical mirrors to minimize the number of optical components required.

The Band 10 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.34 for different values of PWV. The specifications for this receiver are $T_{\text{rx}} < 230$ K over 80% of the band and < 344 K over all the band. Typical on-array T_{rx} values are 175–275 K. The OT assumes 230 K. Figure 4.35 shows the expected T_{sys} for 0.472 mm of PWV, over most of the band given the OT-assumed receiver noise. Phase stability also limits when observations can be made, therefore, most observations in Band 10 will generally only be performed at night and early morning, when the sky is most stable. Also the antenna surface accuracy has a significant effect on the aperture efficiency at this high-frequency band; improvements to the surface accuracy of the 12-m antennas during 2017 has doubled the aperture efficiencies in Band 10. As well as having a lower atmospheric transmission and a less stable atmosphere, Band 10 observing provides the most challenges for observing, including finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing and focus for the relatively small primary beam (although pointing and focus is normally done in lower Bands), and the need for the highest level of stability in the rest of the system.

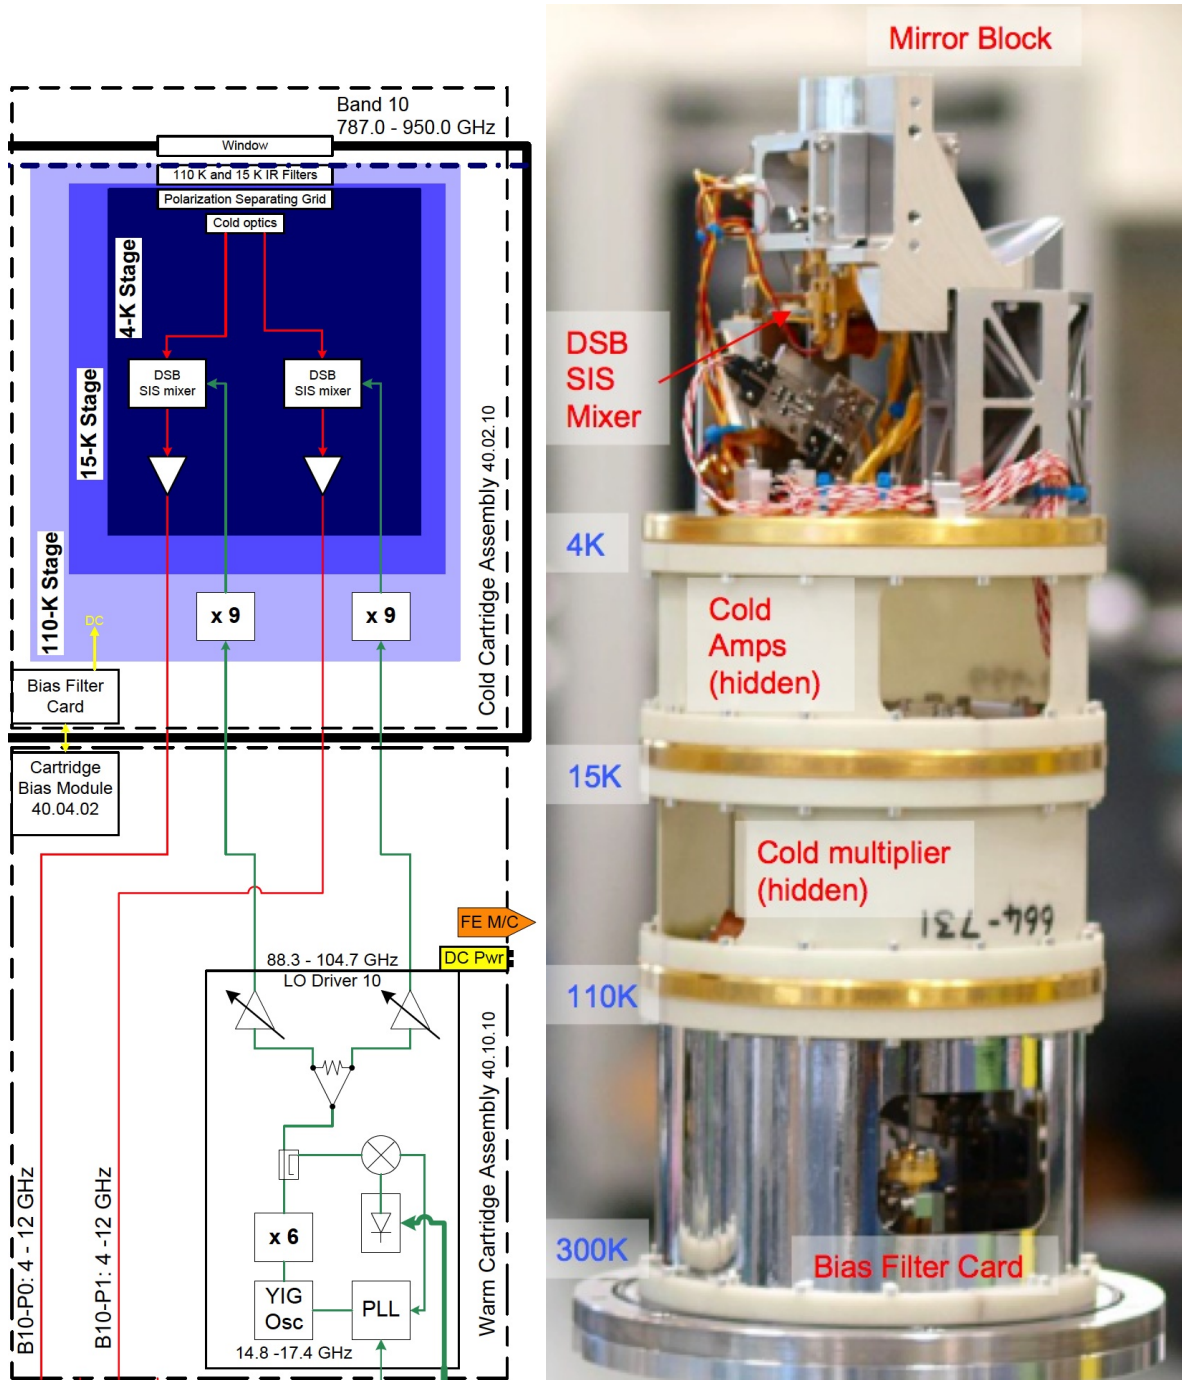


Figure 4.33: Block diagram of Band 10 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 10 cartridges were constructed in Japan at the NAOJ Advanced Technology Center (ATC) in Mitaka, Japan.




Figure 4.34: Band 10 zenith transmission for the first seven octiles of observing conditions at ALMA. Transmission of most of the band is affected by the wings of pressure-broadened lines of H_2O at 841 and 916 GHz.




Figure 4.35: Typical T_{sys} at zenith with Band 10 for the first seven octiles of observing conditions at ALMA. Note that the values for T_{sys} in the worst octiles, which should never be used for high-frequency observations, are off scale in this plot. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

Chapter 5

The Correlators

A correlator can be viewed as the virtual focal plane of an interferometer array where all voltage-based signals from all individual antennas are processed to derive both the cross-correlation products from all independent antenna pairs and the auto-correlation function of each antenna. It delivers complex fringe visibilities that are used to synthesize astronomical images after calibrations in amplitude and phase. Another function is spectroscopy by applying Fourier transform to the correlation products or the voltage-based signals. A correlator tracks delay and phase to adjust the wavefronts of received signals in order to maintain the coherence of the complex visibilities. Walsh-switching modulations for sideband separation (90° phase switching) and spurious-signal suppression (180° phase switching) are demodulated in the correlator.¹

All signals² received by ALMA antennas are processed in one of two correlators: the 64-input Correlator (also known as the Baseline Correlator or BLC) and the ACA Correlator. The 64-input Correlator is used primarily for the main 12-m Array, while the ACA Correlator is used for the Morita Array comprised of the ACA 7-m Array and the Total Power Array (i.e. single-dish), respectively³. Both correlators run simultaneously and independently. Thus, while the 12-m Array observes an object using the 64-input Correlator, the ACA Correlator can work with the 7-m Array and/or the Total Power Array observing either the same or a different object via a sub-array capability.

Celestial signals received by the antennas are down converted to lower frequency bands using a set of Local Oscillators (LOs) and mixers as described in Appendix B. The outputs from the IF system form four Base Bands (BBs), each covering a bandwidth of 2 GHz in two orthogonal linear polarizations. These analog BB signals are sampled at the sampling frequency of 4 GHz and quantized with eight quantization levels (3 bits per sample) in digitizers, and then transferred via fiberoptic cable to one of the two correlators.

Both correlators generate auto-correlation and cross-correlation products at the same time. The auto-correlation is used not only for spectroscopic TP-Array observations but also for normalization of cross power spectra and measurements of system noise temperatures (see Section 10.4.2). The cross-correlations are used for interferometry with the 12-m Array and the ACA 7-m Array, and also for pointing and focus calibrations for all Arrays.

This chapter addresses capabilities of the correlators to be offered for Cycle 8 observations. The specifications of the correlators determine a lot of observational performance such as bandwidth, spectral resolution, time resolution, and polarimetry. The phase tracking performance provides coherence and improves phase stability in correlated data delivered to users. Imperfect correction for the non-linear response in the correlators invokes systematic errors in complex visibilities. The response of digital signal processing is also presented in this chapter.

¹Some of these features can be employed outside the correlator. In the case of the ALMA correlators, the phase tracking is taken in the LO1 and LO2. The 180° phase switching is modulated in the LO1 and demodulated in the DTS transmitter after digitization. See also Sections 5.5.4, B.4.4, and also Emerson 2005, ALMA memo No. 537.

²Except total power (TP) continuum observations for solar observations (see Section 8.10.3) that employ square-law detectors (SQLDs).

³Crossbar switching allows for some flexibility in this arrangement.



Figure 5.1: The ALMA 64-input Correlator. This view shows lights glowing on some of the racks of the correlator in the ALMA Array Operations Site Technical Building and shows one of four quadrants of the correlator. Credit: ALMA (ESO/NAOJ/NRAO), S. Argandoña

5.1 The 64-input Correlator

The 64-input Correlator (Figure 5.1) employs a hybrid design, also known as the FXF⁴ system (Escoffier et al. 2007, A&A 462, 801), that increases by a factor of 32 the spectral resolution of its traditional lag (XF) part. It operates in two basic modes, Time Division Mode (TDM) – equivalent to an XF correlator with a wide bandwidth and a coarse spectral resolution for mainly continuum observations, and Frequency Division Mode (FDM) with fine spectral resolutions for spectral-line observations. A simplified overview diagram of the 64-input Correlator is shown in Figure 5.2. It consists of four quadrants, all of which are available for Cycle 8. Each quadrant can handle a 2 GHz dual-polarization BB for up to 64 antennas⁵. The full set of four quadrants is capable of accepting four BBs to cover a total 8 GHz bandwidth per polarization, that is 16 GHz of instantaneous bandwidth.

5.1.1 TDM mode

TDM is mostly used for continuum observations. Its simplicity, compared with FDM, offers advantages of a lower data rate and better linearity. Therefore, it is used for standard setups such as pointing, focus, delay calibration, system temperature measurements and sideband ratio measurements. TDM provides also a higher time resolution capability as described in Section 5.5.3. The full 2 GHz BB is directly sent to the correlator bypassing the Tunable Filter Banks (TFBs). The correlator cuts off the least significant bit to reduce quantization levels from 3- to 2-bits per sample (Section 5.3).

The TDM mode provides a spectral window (spw)⁶ that consists of up to $256/N_{\text{pol}}$ channels per BB, where N_{pol} is the number of polarization products per BB⁷. As the full 2000 MHz BB is covered, this requires some truncation of the band-edge channels in offline data processing yielding 1875 MHz of usable bandwidth – see Section 6.4.

⁴F, X and F stand for filtering, correlation and Fourier transform, respectively.

⁵ $\frac{N_{\text{ant}}(N_{\text{ant}} - 1)}{2} = 2016$ baselines and 64 auto-correlations for $N_{\text{ant}} = 64$.

⁶An spw is a continuous spectrum composed of uniformly spaced frequency channels. See also Chapter 6 about the relation between BB and spw.

⁷ $N_{\text{pol}} = 2$ for standard Stokes-I observations that employ XX and YY products, and $N_{\text{pol}} = 4$ for full-Stokes observations.

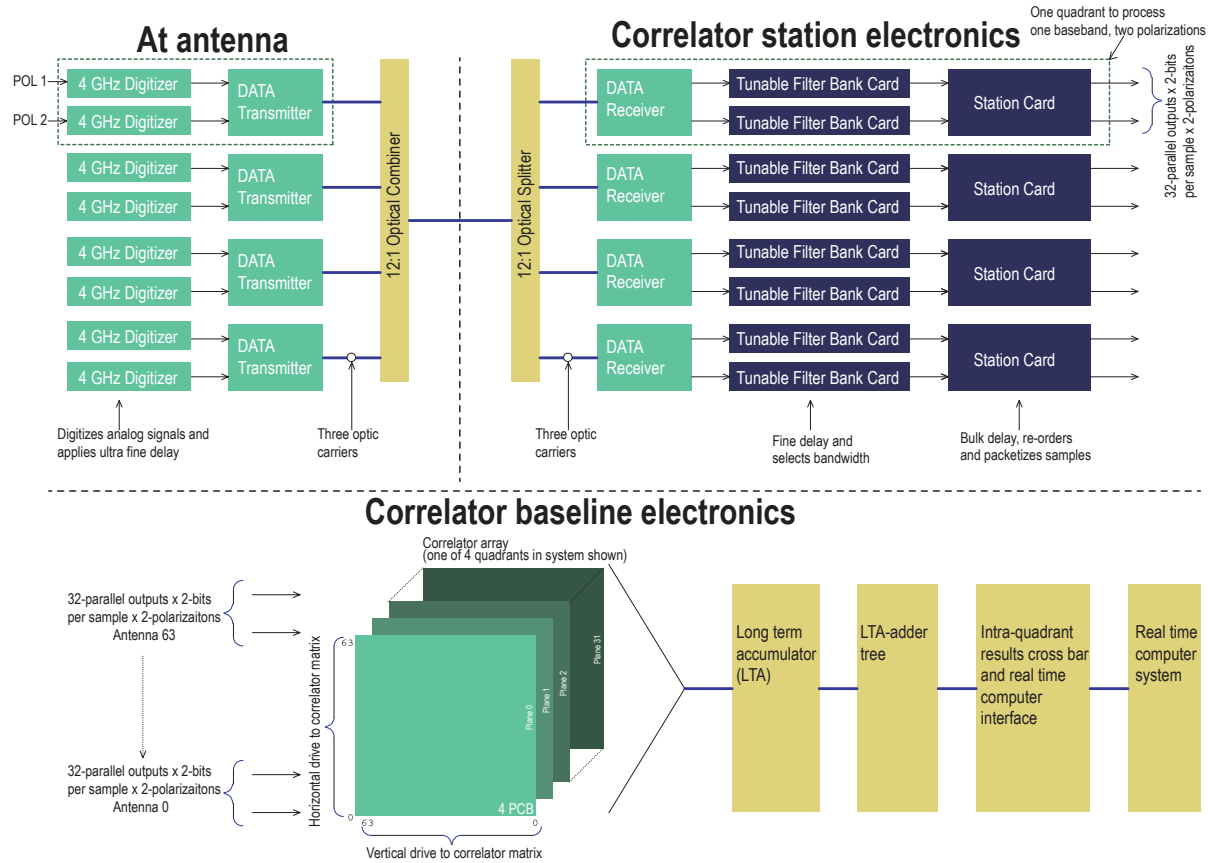


Figure 5.2: Overview diagram of the ALMA digitizers, data transmission system and 64-input Correlator (Escoffier et al. 2007, A&A 462, 801). The digitized data from the individual BBs are transferred to the Tunable Filter Banks (TFBs) and station cards (top right) for the parallelization and processing required on each antenna data stream. The correlator array (lower half) produces auto-correlations on the main diagonal and cross-correlations elsewhere in the correlator matrix. The correlations are then integrated in the Long-Term Accumulators (LTAs, lower right).

5.1.2 FDM mode

FDM is primarily used for spectral-line observations that require a spectral resolution higher than that of TDM. Each 2-GHz BB in an antenna is processed in a TFB card, where digital filtering, a digital mixer and LO (LO4) are implemented in a Field Programmable Gate Array (FPGA). A TFB consists of 32 independent 62.5-MHz bandpass filters (sub-bands). To avoid edge artifacts, sub-band center frequencies are set such that their skirts overlap and the effective bandwidth is thus reduced by a factor of 15/16 that yields 58.59375 MHz per sub-band. The correlator software then combines together contiguous sub-bands to output a seamless cross power spectrum to form an spw with as many as $\frac{15}{16} \times 8192/N_{\text{pol}}$ channels.

The number of channels in each spw can be reduced by averaging 2, 4, 8, or 16 channels into one to accommodate the maximum data rate. Table 5.1 lists the spectral setups for Stokes I ($N_{\text{pol}} = 2$) observations. The number of valid sub-bands to form an spw is selectable from 32, 16, 8, 4, 2, or 1. That is, spw bandwidths of 1875, 937.5, 468.75, 234.375, 117.1875, and 58.59375 MHz, respectively. It is possible to have multiple (up to 4) spws in the same BB. While different spectral setups can be set for spws in different BBs, all of spws in the same BB must have the same channel spacing.

The center frequency of each spw can be tuned over the 2 GHz-wide BB using the digital LO (LO4) in the TFB card. However, the edges of the full bandwidth of the sub-bands cannot fall outside the 2 GHz BB range, and the frequency tuning is made in steps of 30.5 kHz set by the frequency resolution of the TFB digital mixer.

See also Section 5.5.2 for details about spectral setting.

Correlator / Mode	Bandwidth (MHz)	Num. of ch. per pol. (in dual-pol.)	Ch. spacing (MHz)
BLC/TDM	2000*	128	15.625
BLC/FDM	1875	3840	0.488
BLC/FDM	938	3840	0.244
BLC/FDM	469	3840	0.122
BLC/FDM	234	3840	0.061
BLC/FDM	117	3840	0.0305
BLC/FDM	58.6	3840	0.0153
ACA	2000*	4096	0.488
ACA	1000	4096	0.244
ACA	500	4096	0.122
ACA	250	4096	0.061
ACA	125	4096	0.0305
ACA	62.5	4096	0.0153

Table 5.1: Correlator modes and spectral performances per BB for dual parallel polarization (XX and YY).

*Usable bandwidth excluding band edges is ~ 1875 MHz.

5.1.3 Correlation and Realtime Processing

The correlation cards perform the multiply-and-add operations to produce the correlation functions at a clock rate of 125 MHz (4 GHz samples demultiplexed by 32). Four quadrants handle four BBs. A quadrant of the correlator consists of 32 planes of 64×64 256-lag correlator circuits, and it yields auto-correlations and cross-correlations for 64 antennas. In FDM there are always 8192 spectral points per BB for both X and Y polarizations, and higher spectral resolutions can be achieved by reducing the number of sub-bands (i.e. bandwidth) used to form an spw. It is possible to set different modes in different quadrants (i.e. BBs); for example, while one BB is set in TDM, other BBs can be set in FDM. For spectral setup details, see Chapter 6. The Long Term Accumulator (LTA; see Figure 5.2) takes short 1 ms or 16 ms integrations from the correlator circuits and provides longer term integration. Further time averaging through multiple dump intervals is performed in the Correlator Data Processor (CDP) computers. See also Section 5.5.3 about time resolution. Section 5.5 describes in detail the correlator data processing.

5.2 The ACA Correlator

The ACA Correlator (Figure 5.4) is dedicated to observations with the Atacama Compact Array that consists of the 7-m Array with twelve 7-m antennas and the Total Power Array (TP Array) with four 12-m antennas. The ACA Correlator is based on the FX⁸ design in which the incoming time-domain data stream is converted into frequency-domain spectrum via an FFT (Fast Fourier Transform) module before cross multiplication to form the power spectrum. The FFT part always accepts a full 2 GHz bandwidth and outputs 524288 channel spectra with a frequency resolution of 3.815 kHz (Figure 5.6). This design enables flexible spectral handling as shown in Figure 5.3 producing multiple spws with different channel spacings by averaging multiple channels (spectral binning). Since the spectral resolution function is different from that of the 64-input Correlator, a frequency profile synthesis (FPS) is performed in the Correlator Data Processor computer so that the outputs of both correlators are matched (Kamazaki et al. 2008, ALMA Memo 580).

A detailed diagram of the signal processing in the ACA Correlator itself is shown in Figure 5.3. Similar to the 64-input Correlator, each quadrant of the ACA Correlator processes a BB pair; so the complete system can process up to sixteen antennas independently from the 64-input Correlator.

⁸F and X stand for Fourier transform and multiplication, respectively.

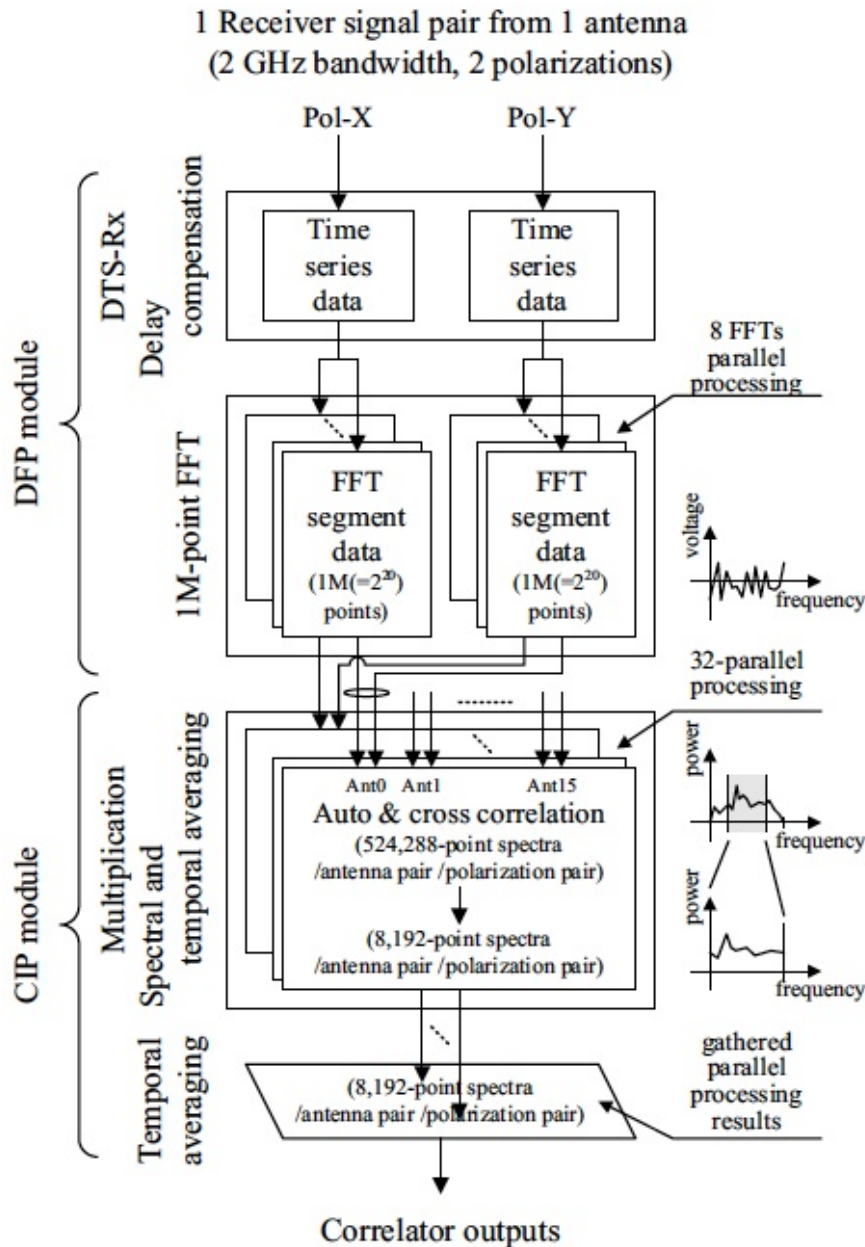


Figure 5.3: Block diagram of the ACA Correlator. One of four quadrants is shown. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. FFT is performed in an FPGA delivering signals in the spectral domain which are then multiplied. The auto and cross-correlation data are then accumulated in time, and the parallel streams averaged together in the Correlation and Integration Processor (CIP) module. The correlated spectra are then fed to the ACA-CDP for further accumulation and processing.

After receiving the digitized BB signals in the Digital Transmission System Receiver (DTS-Rx), the DTS-Rx and FFT Processor (DFP) modules compensate for geometrical delays between antennas and perform the 2^{20} -point FFT that produces a 2^{19} -point complex spectrum⁹ (hereafter, voltage spectrum) for every BB per

⁹ 2^{20} -point FFT produces a 2^{20} -point complex voltage spectrum of a double sideband including the image sideband that will be discarded.



Figure 5.4: Two quadrants of the ALMA Compact Array (ACA) correlator installed in the ACA correlator room. Credit : ALMA (ESO/NAOJ/NRAO), S. Okumura

antenna per polarization, with a channel separation of 3.815 kHz ($= 2 \text{ GHz} \div 524288 \text{ channels}$) in a 16-bit complex integer form. The 16-bit complex voltage spectra are re-quantized into a 4-bit complex integer and sent to the Correlation and Integration Processor (CIP) modules.

The CIP module trims the required frequency range and multiplies the antenna-based voltage spectra to generate baseline-based cross power spectra that correspond to cross-correlations. Antenna-based power spectra, corresponding to auto-correlation, are also generated in the same way. Cross-polarization power spectra can be optionally produced. The (cross) power spectra are channel-averaged and time-integrated as designated before they are sent to the ACA-CDP computers in the Computing subsystem through optical fibers.

The ACA-CDP performs further spectral processing such as non-linearity correction (see Section 5.6.2), FPS and temporal integration, before data are sent to the archive.

The overall hardware design of the ACA Correlator is shown in Figure 5.5. The ACA Correlator is equipped with high-speed FPGA chips rather than Application Specific Integrated Circuit (ASIC) chips used in the 64-input Correlator. For technical details, see Kamazaki et al. 2012, PASJ 64, 29.

5.3 Digitizers

A digitizer is a device converting continuous voltage waveforms into quantized voltage levels, sampled at discrete timings, and encoded in a digital format. The ALMA digitizers are located in the antenna back ends (BEs) where they convert BB signals into the ALMA digital format (Recoquillon et al. 2005, ALMA Memo No. 532). The Digital Transmitter (DTX) on the BE transfers the digital signals through optical fibers to the data receivers (DRXs) in the correlators.

Each antenna BE is equipped with 8 digitizers to accept 8 signal streams of $4 \text{ BBs} \times 2 \text{ polarizations}$. To sample and quantize the 2-4 GHz BB signals, the ALMA digitizer samples at the Nyquist frequency of 4 GHz and uses 8 quantization voltage levels (3-bit per sample). However, the 64-input Correlator reduces the quantization levels to 4 (2-bit). Signals through the TFBs are re-quantized in 4-bit before the correlation process. The ACA Correlator handles 3-bit digital form in FFT and the spectrum is requantized in 4-bit before cross multiplication. These quantization processes affect the sensitivity and linearity in the power spectra. Compared to an ideal analog system, the 64-input correlator has a sensitivity loss of 12% (TDM) or 15.5% (FDM) due to the quantization including re-quantization. Similarly the ACA Correlator has a 4.8% sensitivity




Figure 5.5: Overview diagram of the ACA Correlator. The signals from the BackEnd Subsystem are processed in the DFP and CIP modules and the auto- and cross-correlation products are sent to the Computing Subsystem (see text and Figure 5.3 for details of these). The modules are controlled by the Monitor and Control Interface (MCI), which communicates with the correlator control computer (ACA-CCC). The DTS-Rx is connected with the monitor and control computer (ACA-DMC) for the monitor and control compatibility with the 64-input Correlator.




Figure 5.6: Data capture sequence of the ACA Correlator. The design enables output with multiple spws with different channel spacings by averaging multiple channels (spectral binning). The frequency range of an spw is selectable from the 524288-point (512K) spectra across 2 GHz bandwidth.

loss. These numbers suppose that the digitizer input power is in the expected range and that the digitizer adjustments are optimal. See also Section 5.6.1 about the quantization loss.

5.4 Offline WVR Correction

The Water Vapor Radiometer (WVR; see Appendix A.6 for technical description) is a device mounted on each of the 12-m antennas to measure the amount of Precipitable Water Vapor (PWV) in order to correct for pathlength fluctuations in the troposphere.

While the WVR data were previously recorded together with the visibilities, since Cycle 7 the online software will only produce uncorrected visibilities, and the WVR correction will be applied offline in the CASA pipeline. Recording both WVR corrected and uncorrected results impacted data rates and storage, and producing only one data stream allows for shorter integration times.

The WVR correction will not be applied in the ACA Correlator because the array is so compact that PWV fluctuation is not significant.

5.5 Capabilities of the Correlators

5.5.1 Polarization

Since Cycle 8, ALMA adds the full-polarization capability with the 7-m array as well as with the 12-m array. However, polarimetry with the TP Array remain unavailable for Cycle 8.

Both correlators are capable to produce correlation pairs of $\langle XX^* \rangle$, $\langle XY^* \rangle$, $\langle YX^* \rangle$, and $\langle YY^* \rangle$ where X and Y stand for two linear orthogonal polarization components. These correlations relate to the Stokes visibilities of I , Q , U , and V as

$$\begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix} = \frac{1}{2} \begin{pmatrix} 1 & 0 & 0 & 1 \\ \cos 2\psi & -\sin 2\psi & -\sin 2\psi & -\cos 2\psi \\ \sin 2\psi & \cos 2\psi & \cos 2\psi & -\sin 2\psi \\ 0 & -i & i & 0 \end{pmatrix} \begin{pmatrix} \langle XX^* \rangle / (G_X G_X^*) \\ \langle XY^* \rangle / (G_X G_Y^*) \\ \langle YX^* \rangle / (G_Y G_X^*) \\ \langle YY^* \rangle / (G_Y G_Y^*) \end{pmatrix}, \quad (5.1)$$

where G_X and G_Y stand for complex antenna-based gains in X and Y polarizations and ψ is the parallactic angle. This formulation doesn't include cross talk, also known as 'D-terms', whose calibration will be discussed in Chapter 10.

Both correlators form cross-polarization products of $\langle XY^* \rangle$ and $\langle YX^* \rangle$ only when polarimetry is required. This is obtained at the expense of less spectral channels.

5.5.2 Spectral Setup

Both correlators are capable of accepting eight 2 GHz bandwidth signal streams consisting of four BBs and two polarizations, and to be configured to have multiple spws within a BB.

Each spw yields a continuous spectrum composed of uniformly spaced spectral channels. Table 5.1 summarizes the maximum number of spectral channels and the channel spacing of the correlators in dual linear polarization mode (XX and YY). The maximum number of channels per spw would be doubled (and the channel spacing halved) for single polarization (XX) observations. Conversely, for full polarization observations (XX , XY , YX , and YY), the maximum number of channels is halved and the channel spacing correspondingly doubled as compared to Table 5.1. The bandwidth and the number of channels available from the ACA correlator are identical through the application of FPS (Frequency Profile Synthesis). Note that the channel spacing is not the same as the spectral resolution because of the effects of the applied weighting function, as shown in Figure 5.7. See also Chapter 6 about multiple-spw setup.

The spectral profile which the correlators output is a convolution of the true spectrum with the spectral resolution function shown in Figure 5.7. The spectral resolution function is given by the Fourier transform of the weighting function applied to the correlation function in the lag domain. The spectral resolution is often characterized by the FWHM (full width at half maximum) of the spectral resolution function.

Uniform weighting in the lag domain results in the sinc function with the FWHM of $1.21 \times$ the channel spacing and yields spectral sidelobes. This 'ringing' phenomenon affects the spectra when a narrow line, interference spike, or strong edge channels are present. Alternative weighting functions¹⁰ are applicable to suppress the spectral sidelobes as shown in Figure 5.7. Note that the sidelobes do not matter in most of astronomical observations where a line profile spreads over several spectral channels. The choice of a weighting function is a trade off between resolutions and sidelobes. The default weighting function is the Hanning function, which gives the spectral resolution of $2 \times$ the channel spacing and the maximum sidelobe level of -2.6% . A proposer may request another weighting function if scientifically necessary and technically justified through the OT during the Phase 2 validation stage. The proposer must contact their supporting ARC to make this change at Phase 2.

Spectral channel averaging is available to bin or average spectral channels in the CDP. Channels can be averaged together; factors of $N = 2, 4, 8, \text{ or } 16$ are available. The main purpose is to reduce the data rate to the archive and the total data volume. It also provides a broader spread of correlator functionality between the current TDM (which has only 128 channels in dual polarization) and full FDM (with 3840 channels in dual polarization mode). As such, spectral averaging is beneficial to those projects requiring more resolution than TDM, but where the FDM channels at full resolution are unnecessary. Table 5.2 shows the resolutions (in kHz)

¹⁰For a full description, see <http://mathworld.wolfram.com/ApodizationFunction.html>.




Figure 5.7: (left table) Spectral resolutions (FWHM) and the maximum sidelobe levels for weighting functions. ALMA employs the Hanning window in default. (right figure) Spectral resolution functions corresponding to various weighting functions

for different values of N , using Hanning weighting, in the different bandwidth modes. The channel spacings are in parentheses.

Note that the default Hanning window function gives a resolution 2 times the channel spacing, so using $N = 2$ (cutting the number of spectral channels from 3840 to 1920) results in negligible loss of final resolution. It is recommended, unless the maximum spectral resolution is required by the observations, to reduce the number of channels when feasible. However, note that this is a non-reversible operation. For Cycle 8, the default value for spectral averaging will be $N = 2$. The proposer must specifically request and justify the use of $N = 1$ for spectral averaging.

The ACA Correlator is an FX correlator where the uniform weighting in the FFT segment is equivalent to the Bartlett weighting in the lag domain. FPS in the CDP can be applied to match the spectral resolution function to that of the 64-input Correlator. The bandwidth and the number of channels in the ACA correlator remain unchanged through FPS since Cycle 4.

5.5.3 Time Resolution

As shown in Figure 5.8 there are three stages where correlated data are accumulated or averaged within a correlator sub-system:

1. Accumulator#1: implemented in the correlator itself. Data are accumulated without applying a proper averaging factor. The accumulation period is called dump duration.
2. Accumulator#2: implemented in the CDP, this stage first truncates $\sim 5\%$ band edges for each spw and then averages the whole spectral channels into one. Those channel averages are then further averaged in time during a number of dump duration periods. This average period is called channel average duration, which must be an integer multiple of the dump duration.

Usable bandwidth (MHz)	$N =$ Channels =	Spectral resolution (channel spacing) [kHz]				
		1	2*	4	8	16
1875		3840	1920	960	480	240
		977 (488)	1129 (977)	1938 (1953)	3904 (3096)	7813 (7812)
937.5		488 (244)	564 (488)	969 (977)	1952 (1953)	3906 (3906)
468.8		244 (122)	282 (244)	485 (488)	976 (977)	1953 (1953)
234.4		122 (61)	141 (122)	242 (244)	488 (488)	977 (977)
117.2		61 (31)	71 (61)	121 (122)	244 (244)	488 (488)
58.6		31 (15)	35 (31)	61 (61)	122 (122)	244 (244)

Table 5.2: Spectral resolution and channel spacing (in parentheses) in kHz for different correlator bandwidth modes (left column) and for different spectral channel averaging factors (columns, $N = 1$ to 16), using Hanning smoothing. The number of channels can be reduced from 3840 (for the unaveraged case, $N = 1$) down to 240 (for $N = 16$). As N increases, the spectral resolution functions of adjoining channels are combined and then the spectral resolution approaches the channel spacing. Values are given for the dual polarization case. For Cycle 8, the default value for spectral averaging will be 2.

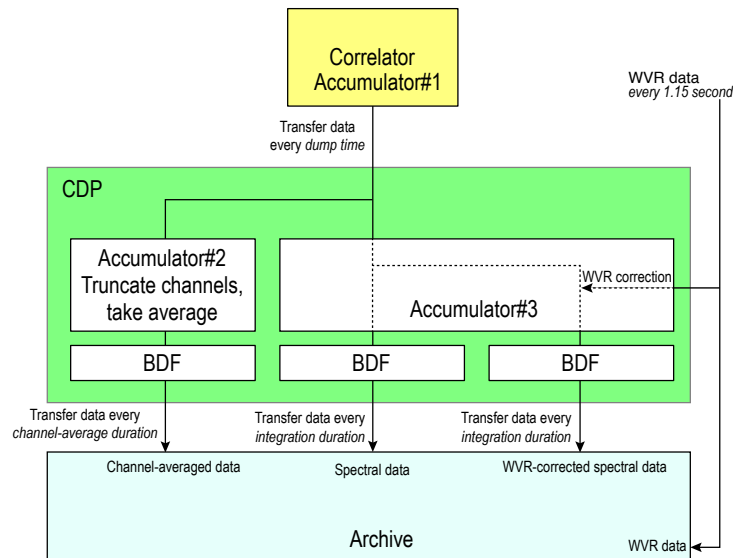


Figure 5.8: Basic data processing and accumulation/averaging steps between the correlator and archive. The CDP applies channel truncation, spectral averaging, and time averaging for the channel-averaged data. It also applies online WVR correction to produce the WVR-corrected spectral data. They are packed in Binary Data Format (BDF) and are archived together with the WVR data.

3. Accumulator#3: implemented in the CDP, this stage averages full resolution spectra during a number of dump duration periods. This average period is called integration duration, which must be an integer multiple of the channel average duration. This also outputs WVR-corrected spectral data.

Most of the data produced by the correlator corresponds to full resolution spectra, and its data rate can be approximated by the following expression:

$$R_{\text{int}} = N_{\text{ant}}^2 \times N_{\text{BB}} \times N_{\text{ch}} \times N_{\text{pol}} \times N_{\text{byte}}/T_{\text{int}}, \quad (\text{bytes/sec}) \quad (5.2)$$

where N_{ant} , N_{BB} , N_{ch} are the number of antennas, BBs, and total number of spectral channels, respectively. $N_{\text{pol}} = 1, 2$, and 4 for single, dual and full polarization products, respectively. N_{byte} (bytes/visibility) is the byte size of the real or imaginary part of a visibility, and T_{int} is the integration duration in the Accumulator#3. For example, an array configuration with $N_{\text{ant}} = 43$, a spectral setup with $N_{\text{BB}} = 4$, $N_{\text{ch}} = 3840 \times 2$ polarization products (XX and YY), $T_{\text{int}} = 6$ sec, and assuming $N_{\text{byte}} = 2$ yields a data rate of 18.9 MB/sec.

In Cycle 8, ALMA regulates R_{int} at 70 MB/sec. The ALMA Observing Tool (OT) estimates the data rate for each science goal and warns if it exceeds the limit. In this case, it is required to consider the optimal trade between time and spectral resolutions, number of polarization products, number of BBs, and number of antennas in the array.

Some specific considerations about different time duration follows:

- Integration duration of spectral data with the 64-input Correlator

The basic integration time of the 64-antenna Correlator is 16 and 1 ms for cross- and auto-correlation, respectively. The correlator integration duration can be, in TDM and FDM, any value that would keep the integration data rate R_{int} under the regulated limit for a given science cycle. Note that the 64-input Correlator sets $N_{\text{byte}} = 2$ in real-time, based on the actual magnitude of the data produced (correlation regime). For validation purposes, and only in FDM, a favorable condition of $N_{\text{byte}} = 2$ is assumed.

- Integration duration of spectral data with the ACA Correlator

The ACA Correlator uses $N_{\text{byte}} = 4$ and its dump time is a multiple of 1 ms and 16 ms for auto-correlation and cross-correlation, respectively. The maximum data rate from the ACA Correlator is 3.6 MB/s, independently of the number of antennas. However, the average data rate to the archive will be lower during typical observing modes, because of overheads and the use of TDM in calibration scans.

- Channel-averaged duration

The average power spectra across the whole bandwidth, in each spw, can be recorded in an independent spw consisting of only one spectral channel. This single channel spw, without bandpass calibration, is not used for astronomy, but for recording phase and amplitude variations through the observing time. Its duration should normally be equal to or shorter than the WVR integration duration.

Time-domain observing mode

A time-domain observing mode is dedicated for observing time-variable objects such as pulsars and transients. This capability is offered starting Cycle 8, using the 64-input correlator, as described in Section 8.11.4. For Cycle 8 this mode is limited to Band 3 only. This mode employs the VLBI phased array capability whereby signals from individual antennas are summed up after phase alignment (phasing) through adjustment of the TFBS. The phased sum signal is converted to VDIF format and recorded onto Mark 6 VLBI recorders. The disk modules containing the recorded baseband data will then be shipped to the VLBI correlation facility at the Max Planck Institut für Radioastronomie for playback and generation of PSR-FITS format data appropriate for timing and search analyses. The PSR-FITS file will have a user-specified time resolution of up to 8 μs and a spectral resolution of 32 channels/BB. See Section 8.11.4 for additional information. Beside generating the phased sum signal, the correlator also produces an ASDM data set that contains the standard interferometry visibilities from the array with a standard time resolution.

5.5.4 Phase Switching

The 180° phase switching capability is implemented to suppress spurious signals and IF cross talk between the LO1 and the digitizers¹¹. The 90° phase switching capability is offered to separate, in real-time, upper and lower sidebands of DSB receivers (Bands 9 and 10). The expedient is based on 90° phase switching at the LO1, and accumulation in different memory areas (bins) in the correlator¹². Walsh function series are employed to make antenna-based switching patterns orthogonal to each other. The Walsh sequences permutes every 16 ms, and the full cycle completes after 2048 ms to achieve orthogonality for all antennas in the array. The dump duration, during a side band-separation observation, must be an integer number of the complete Walsh cycle. See also Appendix B.4.3.

The 180° phase switching

The phase switching capability (see principles in Figure 5.9) is realized by phase modulation in the LO1 and demodulation in the DTS transmitter module after digitization.

State	Ant.	LO phase	IF signal	Demod.	Cross Corr.
0	1	ϕ_{LO}	$S_1 e^{-i\phi_{LO}} + N$	0	$\langle S_1 \cdot S_2^* \rangle + \langle N \cdot N^* \rangle$
	2	ϕ_{LO}	$S_2 e^{-i\phi_{LO}} + N$	0	
1	1	ϕ_{LO}	$S_1 e^{-i\phi_{LO}} + N$	0	$\langle S_1 \cdot S_2^* \rangle - \langle N \cdot N^* \rangle$
	2	$\phi_{LO} + \pi$	$S_2 e^{-i(\phi_{LO} + \pi)} + N$	π	
average					$\langle S_1 \cdot S_2^* \rangle$

Table 5.3: Cross correlation of a baseline under 180° phase switching. The received signal of antenna 2, S_2 , is modulated by 180° by the LO1 and demodulated in the DTS transmitter module at each antenna. Spurious signals, N , in the IF or cross talks are not modulated via the LO1 and thus are canceled via demodulation and accumulation of data in states 0 and 1.

The 90° phase switching

This function aims to separate sidebands for the DSB receivers of Bands 9 and 10. When the LO phase of antenna 2 is offset by 90°, the phases of USB and LSB signals in IF will be $-\phi_{LO} - \frac{\pi}{2}$ and $\phi_{LO} + \frac{\pi}{2}$, respectively. The 90°-demodulation is taken by swapping real and imaginary parts of the cross-correlations in the CDP.

¹¹Note that this scheme doesn't work for auto-correlations.

¹²See Thompson, Moran, and Swenson, "Interferometry and Synthesis in Radio Astronomy (2nd edition)", §6.1 and §7.5.

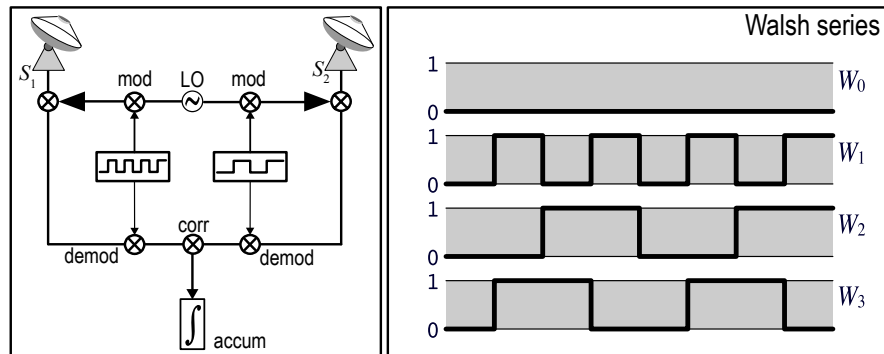


Figure 5.9: Schematic diagram of phase switching. The LO1 signal phase is modulated with the the Walsh series pattern and demodulated in the DTS before cross-correlation. The Walsh function series are employed to generate switching patterns that are orthogonal to each other.

After demodulation, the sign of LSB correlation will be inverted with respect to USB. Thus, the average and the difference of cross-correlations in two states gives the USB and the LSB visibilities, respectively.

State	Ant.	LO phase	IF signal	Demod.	Cross Corr.
0	1	ϕ_{LO}	$S_{1,U}e^{-i\phi_{\text{LO}}} + S_{1,L}e^{i\phi_{\text{LO}}}$	0	$\langle S_{1,U} \cdot S_{2,U}^* \rangle + \langle S_{1,L} \cdot S_{2,L}^* \rangle$
	2	ϕ_{LO}	$S_{2,U}e^{-i\phi_{\text{LO}}} + S_{2,L}e^{i\phi_{\text{LO}}}$	0	
1	1	ϕ_{LO}	$S_{1,U}e^{-i\phi_{\text{LO}}} + S_{1,L}e^{i\phi_{\text{LO}}}$	0	$\langle S_{1,U} \cdot S_{2,U}^* \rangle - \langle S_{1,L} \cdot S_{2,L}^* \rangle$
	2	$\phi_{\text{LO}} + \frac{\pi}{2}$	$S_{2,U}e^{-i(\phi_{\text{LO}} + \frac{\pi}{2})} + S_{2,L}e^{i(\phi_{\text{LO}} + \frac{\pi}{2})}$	$\frac{\pi}{2}$	
(State 0 + State 1)/2					$\langle S_{1,U} \cdot S_{2,U}^* \rangle$
(State 0 - State 1)/2					$\langle S_{1,L} \cdot S_{2,L}^* \rangle$

Table 5.4: Cross correlation of a baseline under 90° phase switching. The suffixes of 1, 2, U , L stand for antenna 1 and 2, and USB and LSB, respectively. The LO phase of antenna 2 is modulated by 90° . After 90° -demodulation and correlation, the cross-correlations of USB and LSB have opposite signs. Thus, the average and the difference of cross-correlations in two states gives the USB and the LSB visibilities, respectively.

5.6 Practical Performance

5.6.1 Sensitivity

A digitizer adds quantization noise to its input analog signal, with a consequent signal-to-noise reduction or sensitivity loss. The ALMA digitizer employs 3-bit (8-level) quantization¹³, and additional re-quantization processes are applied in the correlators.

In Cycle 8, the 64-input Correlator TDM mode makes use of 2-bit quantization that yields a quantization efficiency of $\eta_{\text{Q}} = 0.88$. In the FDM case, digital filtering is applied to 3-bit quantized signals followed by 2-bit re-quantization. The two stages of quantization would cause the net quantization efficiency of $\eta_{\text{Q}} = 0.96 \times 0.88 = 0.85$ if they were independent (see mode chart tables of Escoffier et al., ALMA Memo 556), nevertheless, the real efficiency can be slightly better than the multiplication (Iguchi et al. 2005, PASJ 57, 259)¹⁴.

The ACA Correlator feeds 3-bit quantized signals into the FFT processors where butterfly arithmetic is performed with 16-bit precision. The FFT output spectrum is re-quantized into 4-bit (16 levels) before cross multiplication. The combination of these quantizations results in a loss of 4.8% (i.e. $\eta_{\text{Q}} = 0.952$; Kamazaki et al. 2012, PASJ 64, 29).

The quantization efficiencies above are theoretical values under optimal conditions of the input signal level and threshold voltages. Although the signal level is adjusted for every scan, significant variation of power level (e.g. observation of strong sources, unstable weather, or a frequency band near atmospheric absorption lines) can violate the conditions required to optimize the quantization efficiency.

5.6.2 Linearity

Digital quantization also causes non-linearity in the visibility measurements. The relationship between the correlation coefficient of quantized signals with that obtained for analog signals (i.e. infinite quantization levels) is known as the Van Vleck relationship, which is used for the non-linearity correction.

Figure 5.10 shows the Van Vleck relationships in 2-, 3-, and 4-bit quantizations. The correlation coefficients of quantized signals keep adequate linearity for small correlation coefficients¹⁵. This indicates that for most sources the cross-correlation response is effectively linear.

However, auto-correlations are affected by the non-linearity because the correlation coefficient at zero lag must be unity by definition. The non-linearity in auto-correlation power spectra influences system temperature

¹³The quantization efficiency would be $\eta_{\text{Q}} = 0.96$ if optimal signal levels, threshold voltages, and perfect signal processing were performed (Thompson 1998, MMA Memo 220).

¹⁴It is reported that the combination of 2-bit quantization, digital filtering, and 2-bit re-quantization yields $\eta_{\text{Q}} = 0.81$, better than $0.88 \times 0.88 = 0.77$.

¹⁵The departure from linear relation is $\delta\rho/\rho < 10^{-3}$ for $\rho < 0.2$

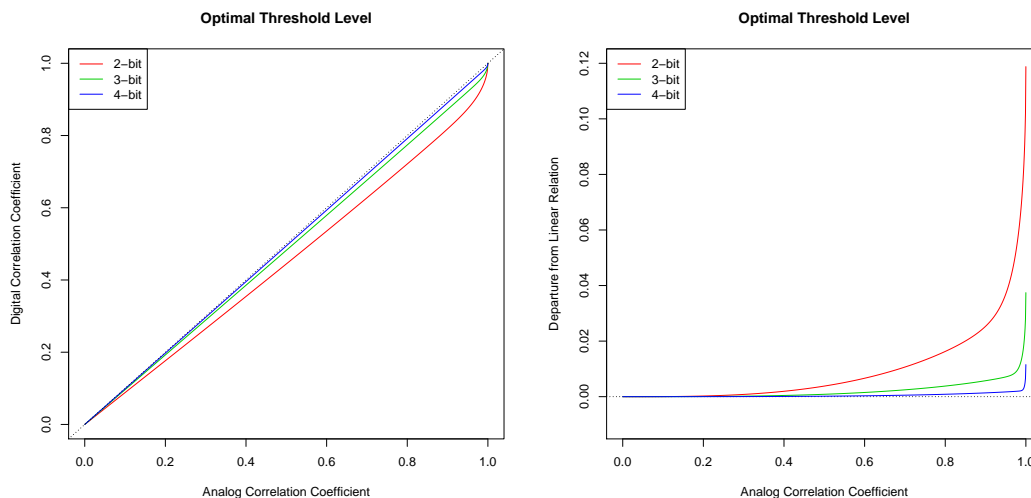


Figure 5.10: Van Vleck relationship of 2-, 3-, and 4-bit quantizations at the optimal signal power level and threshold voltages. (Left): Correlation coefficient of digitized signals as a function of analog correlation coefficients, ρ . (Right): Departures of the Van Vleck relationship from linear relation using the single factors of $\eta_4 = 0.881$, $\eta_8 = 0.963$, and $\eta_{16} = 0.988$ for 2-, 3-, and 4-bit quantizations, respectively.

measurements and the TP Array observations. In both correlators, the CDPs measure total power of digitized signals and determine the Van Vleck relationship to correct the (cross) power spectra.

5.7 Final data product - the ASDM

The final product from each observation in the archive is known as the ASDM (the ALMA Science Data Model), each of which has a unique hexadecimal name (e.g. `uid://A002/X2fed6/X3f`). The ASDM contains the metadata (headers, descriptions of the observation setup, ancillary data, etc), and the binary data (the raw data itself), and is described in more detail in Section 12.2. The following describes the spectral data in the ASDM and how its related to the correlator output.

In the ASDM, the binary data are saved in a structure of spws. All of the data in a single spw must share the same frequency setup, including the number and width of spectral channels, and the integration time. The observed spws will be a combination of the science spws set up by the proposers in Phase 1 of the OT, and additional spws from observations needed for calibration (pointing, and sometimes system temperature) set up during Phase 2. Additionally, the WVR data are stored in a spw with 4 channels around the water line at 183 GHz. In the ASDM, except for the WVR spw, each requested spw maps into two output spws in the data: one with the requested dimensions of N channels per polarization product (for example 128 or 3840), and a second channel averaged version with one channel per polarization product (this averaging is done in the correlator). The channel averaged data are used by the online telescope calibration system (TelCal) and for realtime diagnostic purposes (QuickLook), and are typically not used downstream in the data reduction. Overall, this can lead to ASDMs with a large number of spws. For example, a typical science observation in FDM mode can have more than 25 spws. Every scan/spw combination has an associated “intent” that indicates its purpose (pointing, system temperature, science, etc). The intent can be used in CASA to determine how to utilize each spw in the data reduction process.

Chapter 6

Spectral Setups

This chapter describes the frequency setup of ALMA. In particular it shows how spectral setups are defined from the user viewpoint and how these then set up the ALMA hardware. It summarises the functions and operation of the local oscillator (LO) and Intermediate Frequency (IF) systems, as well as outlining some of the other points pertaining to the frequency setup. For those interested in more details of the LO and IF components and how the hardware works, please see Appendix B.

To the ALMA system, a spectral setup includes the hardware LO, IF and correlator settings, such that each spectral window (spws) covers the desired lines and/or continuum frequencies. To the user, however, the spectral setup is normally defined in the Observing Tool (OT) just in terms of the desired lines or observing frequencies, spw bandwidths and spectral resolutions: there is normally no need for the user to worry about the details of each hardware setting. For full details of the OT and how to use it, see the user manual and reference manuals, available from the ALMA website¹ (and also in the OT itself). In this chapter, we try to reconcile these two viewpoints.

6.1 Introduction

ALMA uses two stages of heterodyne conversion to shift the signals from the sky or observing frequency down to a range where electronics can be used to perform the digital sampling and cross-correlations. In general, the signals of observing frequency f_{sig} are mixed to IF signals of frequency f_{IF} using a narrow Local Oscillator signal (LO, at f_{LO}). f_{IF} is the difference frequency between f_{LO} and f_{sig} , where f_{IF} is always, for ALMA, much lower than f_{sig} . The output signal is split into two separate sidebands at frequencies offset from f_{LO} , known as the upper and lower sideband (USB or LSB).

Figure 6.1 illustrates the ALMA signal path and basic principle of operation. The first heterodyne stage uses the SIS mixer and LO1 to mix (or downconvert) the astronomical signal to 4 GHz- to 8 GHz-wide IF bands, giving both USB and LSB. Up to four 2 GHz wide basebands (or BBs) can be placed in the available IF range in either or both the USB or LSB IFs using the second stage of downconversion. Within each of these basebands it is possible to place up to four spws. Each spw forms a final contiguous spectrum, with bandwidths from 58.594 MHz up to 1.875 GHz wide (see Section 6.3).

Using the OT during both the proposal preparation (Phase 1) and scheduling block (SB) creation (Phase 2), the user adds spws to each baseband by choosing the frequencies (or spectral lines) to be observed in the BB, along with the required bandwidth or spectral resolution. During the subsequent spectral tuning, there are effectively four different LOs set up by the system (Figure 6.1):

- LO1 which sets the frontend tuning frequency. This has continuous coverage.
- LO2 which positions the basebands within the frontend receiver IF (each baseband uses a different LO2, hence there are four LO2s). This does not have continuous frequency coverage, as it is generated by a 125 MHz frequency comb plus a 20–42.5 MHz continuous frequency synthesizer.
- LO3 which is the clock frequency of the digitizers (fixed at 4 GHz),

¹<http://almascience.org/documents-and-tools/>

- LO4 (also known as the tunable filterbank LO, or TFB LO), which is a digital LO synthesized in the correlators allowing positioning of the spws within each baseband. Each spws has effectively a separate LO4.




Figure 6.1: Summary block diagram of ALMA signal path and LO system for one spw, one baseband, one polarization. The Backend IF system is located in each antenna, and the digitized signal for each baseband/polarization is fed to the correlator in the central technical building through buried optical fibers. See Appendix B for a more detailed description of the components and the other tasks that the LO systems perform.

The OT and the real-time system have a tuning algorithm that attempts to find the best tuning solution for LO1, LO2 and LO4 based on the requested observing frequencies. Because the LOs themselves are generated by combinations of frequency synthesizers, several different tuning solutions may be possible, and the algorithm performs the requisite optimization.² If no solution is possible, the OT will notify the user. A more detailed description of the LO operation is given in Appendix B.2.

There are similar restrictions on the spectral setups and tuning as in recent cycles (see Table 6.2 in Section 6.8). The main limits are that the edges of the 2 GHz-wide basebands cannot lie outside the receiver tuning range listed in Table 4.1, and the edges of the (untruncated) spws cannot lie outside the individual 2 GHz-wide basebands. In addition, the OT will give validation warnings if a) any part of a spw is closer than 30 MHz to the baseband edge or b) more than half of the spw is within (30 MHz + the line width) of the baseband edge.

The settings of each baseband in the correlator are independent, so the bandwidth, resolution and correlator mode as well as the value of LO4 can be different for each baseband. For example, it is possible to have a 58

²Note that with tuning setups having multiple basebands, only the hardware LO1 and LO2, the requested frequencies may not appear in exactly the center channel in all the spws, because of the non-continuous coverage of LO2. Normally the tuning algorithm will use LO4 to compensate for this and re-center the lines. However in the widest bandwidth modes (FDM 1875 MHz and TDM) the spws cannot be shifted using LO4. Then the *sum* of the absolute frequency offsets in all basebands can be as much as 26.38 MHz for 4-baseband setups, giving a worst-case average baseband offset of 6.6 MHz; this is considered small compared with the bandwidth (and the frequency labelling will still be correct). See Section B.2 and “ALMA LO System Setup Algorithms” (Scott, 2009) for more details.

MHz-wide spws centered on a particular line in baseband 1, and simultaneously use a broadband time division multiplexing mode (or TDM mode - see Chapter 5) for all the other basebands. It is possible to have multiple FDM spws *within* each baseband, described in Section 6.3.2 below; however, currently the underlying resolution of all spws *within the same baseband* has to be equal. Different spectral averaging is allowed however within the spws. The spectral setups of the ACA Correlator and the 64-input Correlator are - to the user - effectively the same, with the ACA Correlator having the same allowable spectral functions as the 64-input Correlator.

6.2 Frequency Definitions

There are several frequency definitions used in the ALMA system, in the OT Phase 1 and Phase 2. Most important to the user are:

Center Frequency (Rest) (f_{spws}) This is set by the user in the OT Phase 1, and is the central frequency of the spws in the requested rest frame of the source. Note that the source rest frame is selectable, but is commonly set to the local kinematic standard of rest (LSRK), which is the conventional local standard of rest based on the average motion of the Sun with respect to the solar neighborhood (see Section 6.8). Note it is **strongly recommended** that the user input the appropriate source velocity or redshift in the Field Setup tab of the OT (Chapter 9).

Center Frequency (Sky) This is seen in the OT Phase 1, and is the actual central frequency of the spws in the local ALMA (i.e. sky) velocity frame, after including the velocity of the source. However, it does *not* include the velocity shift of the chosen coordinate frame with respect to ALMA (i.e. the extra velocity from Earth rotation and orbit) as this is not known until actual runtime.

Baseband desired Center Frequency (f_{BB}) This is the frequency center of the *baseband*, shown in OT Phase 2. Both the Sky and Rest Baseband center frequencies are seen. If the spws are centered on the basebands, these will be the same as the Phase 1 Center Frequencies (above). Otherwise they can be different because of the TFB LOs (see below).

Center Offset Frequency (f_{offset}) Used in the OT Phase 2, this is the offset due to the TFB LO. The center frequency of the spws, f_{spws} , is given in terms of this and Baseband Center Frequency by:

$$f_{spws} = f_{BB} \pm (f_{offset} - 3.0GHz) \quad (6.1)$$

where the sign depends on the observing sideband.

6.3 Spectral Setups

The wide IF bandwidth and tuning ability allows for routine simultaneous imaging of multiple lines at different spectral resolutions in different basebands. Some examples (with the approximate line frequencies in GHz) are shown in Table 6.1 (for a source of redshift zero). Note that in many cases the lines will not necessarily appear in the center of the spws (for example, in the Band 6 combination 6b). When <4 spws are required for the primary lines, the other BBs can be set up to cover fainter lines or to observe the continuum, potentially in TDM mode. If lines are the main goal in one spw, the other continuum spws frequencies need to be setup manually; one method of doing this is described in the next section. The selection of secondary lines to be observed can be done using the OT spectral interface. In the case of continuum spws, to maximize the sensitivity, the widest bandwidth mode should be chosen (i.e. in TDM, or 1.875 GHz in FDM - preferably with some channel-averaging, see Section 6.4). Also the continuum spws should cover as much of the IF band as possible. Not only will this maximize the continuum SNR on the science target, but these continuum data can be used to improve phase and amplitude calibration. For this reason it is recommended that all four available basebands be used in all observations. Note that there is no gain in sensitivity by overlapping spws in frequency, since the input signal is the same in the overlap region; the OT uses the unique aggregate bandwidth in calculating the sensitivity, taking any spw overlap into account. However, for sources with known high line density (one spectral feature per 10 MHz), PIs are encouraged to set up all the spws in FDM mode. This will allow a more robust determination of the line-free channels used to form the aggregate continuum during data processing and imaging.

6.3.1 Observing Frequencies for Continuum

For optimum sensitivity, continuum spws should be set to the frequencies with the lowest system temperature. Because the mixer frequency responses are fairly flat, normally this corresponds to the best atmospheric transmission. For full continuum observations (without any required lines, or with only low resolution) the OT has standard optimized frequencies for each band which are used in the OT Phase 1 when Single Continuum is selected; these are noted in Table 6.1 and illustrated by the red in Figure 6.2.




Figure 6.2: ALMA bands available in Cycle 8 (Bands 3 - 10), showing the frequency of the standard continuum settings as red shading. This gives the coverage of both USB and LSB, except for Bands 9 and 10, which have 8 GHz of bandwidth using the USB only. The available frequency coverage of each band is shown by the grey shading, and atmospheric transmission for 0.6 mm of PWV is given by the blue line.

If mixed line+continuum operation is desired, for example with a single line in BB1 observed at frequency f (GHz), the other BBs (2-4) can be set up for continuum observing using TDM mode³. For 2SB receivers, typically this would have BB2 in the same sideband (offset by 2 GHz from f) and BB3 and BB4 in the opposite sideband (offset from BB1 and BB2 by the front end center IF frequency, f_{IF}). To maximise sensitivity the continuum BBs should cover the maximum aggregate unique bandwidth, i.e. they should not overlap in frequency. If sb is the sign of the sideband of BB1 ($sb = +1$ for USB, $sb = -1$ for LSB), then the four BBs could be set up as follows:

- BB1: f (BB covering the primary line)

³Note - this case assumes one spws per BB

Band/setup	Species/transition	Freq. (GHz)	Sideb.	bandwidth	BB	spws	Notes
3 a	Standard cont.	97.5 ¹	dual	TDM			LO1=97.5
3 b *	HCO ⁺ (1-0)	89.188	LSB	58 MHz	1	1	HCO ⁺ /HCN/H ₂ CO
-	HCN (1-0)	88.632	LSB	58 MHz	2	1	
-	CH ₃ OH (7(-2,6)-7(1,6))	101.293	USB	125 MHz	3	1	
-	H ₂ CO (6(1,5)-6(1,6))	101.333	USB	125 MHz	3	1	
-	continuum spws	101.3	USB	TDM	4	1	
3 c *	CO (1-0)	115.271	USB	58 MHz	1	1	CO/CN/C ¹⁷ O
-	CN (N=1-0)	113.499	USB	58 MHz	1	2	
-	C ¹⁷ O (1-0)	112.359	USB	117 MHz	2	1	
-	continuum spws	102.5	LSB	TDM	3	1	
-	continuum spws	100.5	LSB	TDM	4	1	
4 a	Standard cont.	145.0 ¹	dual	TDM			LO1=145.0
4 b	CS (3-2)	146.969	LSB	117 MHz	1	1	CS/DCO ⁺
-	DCO ⁺ (2-1)	144.077	LSB	117 MHz	2	1	
-	SO ₂ (4(2,2)-4(1,3))	146.605	LSB	117 MHz	3	1	
-	H ₂ CO (2(0,2)-1(0,1))	145.603	LSB	117 MHz	4	1	
5 a	Standard cont.	203.0	dual	TDM			LO1=203.0
5 b	H ₂ O (3(1,3)-2(2,0))	183.31	USB	250 MHz	1	1	
-	continuum spws	181.4	USB	TDM	2	1	
-	H ₂ S (1(1,0)-1(0,1))	168.763	LSB	250 MHz	3	1	
-	continuum spws	170.8	LSB	TDM	4	1	
6 a	Standard cont.	233.0 ¹	dual	TDM			LO1=233.0
6 b *	¹² CO (2-1)	230.538	USB	937 MHz	1	1	J=2-1 CO isotopes
-	continuum spws	234.0	USB	TDM	2	1	
-	C ¹⁸ O (2-1)	219.560	LSB	937 MHz	3	1	
-	¹³ CO (2-1)	220.399	LSB	937 MHz	3	2	
-	continuum spws	216.5	LSB	TDM	4	1	
7 a	Standard cont.	343.5 ¹	dual	TDM			LO1=343.5
7 b *	continuum spws	343.0	LSB	TDM	1	1	CO/HCO ⁺ /HCN
-	¹² CO (3-2)	345.796	LSB	469 MHz	2	1	1/2-corr
-	HC ¹⁵ N (4-3)	344.200	LSB	234 MHz	2	2	1/4-corr
-	H ¹³ CN (4-3)	345.340	LSB	234 MHz	2	3	1/4-corr
-	HCN (4-3)	354.505	USB	469 MHz	3	1	
-	HCO ⁺ (4-3)	356.734	USB	117 MHz	4	1	1/4-corr
-	NHD ₂ (2(2,0)-2(1,2))	356.230	USB	234 MHz	4	2	1/2-corr
7 c	¹² CO (3-2)	345.796	USB	58 MHz	1	1	J=3-2 CO/ ¹³ CO
-	continuum spws	344.8	USB	TDM	2	1	
-	¹³ CO (3-2)	330.588	LSB	58 MHz	3	1	
-	continuum spws	331.6	LSB	TDM	4	1	
8 a	Standard cont.	405.0 ¹	dual	TDM			LO1=405.0
8 b	CI (³ P ₁ - ³ P ₀)	492.160	USB	117 MHz	1	1	CI/ ¹³ CI
-	CS (10-9)	489.751	USB	117 MHz	2	1	
-	continuum spws	477.5	LSB	TDM	3	1	
-	continuum spws	479.5	LSB	TDM	4	1	
9 a	Standard cont.	679.0	USB	TDM			LO1=671.0
9 b	¹² CO (6-5)	691.472	USB	469 MHz	1	1	CO/CS
-	CS (14-13)	685.436	USB	117 MHz	2	1	
-	H ₂ S (2(0,2)-1(1,1))	687.303	USB	234 MHz	3	1	
-	C ¹⁷ O (6-5)	674.009	LSB	469 MHz	4	1	
10 a	Standard cont.	875.0	USB	TDM			LO1=867.0
10 b *	CO (7-6)	806.652	LSB	469 MHz	1	1	CO/HCO ⁺
-	HCO ⁺ (9-8)	802.458	LSB	117 MHz	2	1	
-	continuum spws	805.5	USB	TDM	3	1	
-	continuum spws	803.65	USB	TDM	4	1	

Table 6.1: Examples of spectral setups possible with the existing hardware and software. This includes the standard continuum-only setups for each band, and some multiple line/continuum configurations. *Notes:* **1.** Frequency for standard continuum setups are the mean observing frequency of all spws, \approx LO1 in dual-sideband (2SB) receivers, this frequency isn't actually covered if both sidebands are used. See Figure 6.2.

- BB2: $f - 2.0$ (continuum BB in the same sideband, 2 GHz below the primary line)
- BB3: $f - (2.sb.f_{IF})$ (continuum BB in the opposite sideband)
- BB4: $f - (2.sb.f_{IF}) - 2.0$ (continuum BB in the opposite sideband)

where f_{IF} is the front-end IF frequency (6.0 GHz for Bands 3, 4, 5, 7 and 8, and normally 8.0 GHz for Bands 6, 9 and 10). For DSB receivers, it is common to keep all the BBs in the same sideband, so BB3 and BB4 would be at $f - 4.0$ GHz and $f - 6.0$ GHz. Note that this is an approximate rule and for the DSB receivers, the image sidebands are now turned on by default providing the added sensitivity for continuum observations (Section 6.3.4). When at all possible, adjustments of the continuum BBs (for example, choosing whether the opposite sideband is LSB or USB) should be done using the OT spectral display, to avoid deep atmospheric absorption features. This is particularly important at Bands 8, 9 and 10, and near the water lines around 183 GHz in Band 5 and 325 GHz in Band 7 (see Figure 6.2). Also, if contiguous spectral coverage of the continuum is desired, the BBs should be offset by 1.875 GHz rather than 2.0 GHz as suggested above. Adding TDM spws in otherwise unused BBs (at non-overlapping frequencies) will improve the overall sensitivity for calibration and it is recommended to use all four BBs whenever possible.

6.3.2 Multiple Spectral Windows in the Same Baseband

The system allows the capability to have up to four spws in each baseband, and each *baseband* can have completely different setups. This is useful for projects where several lines need to be observed simultaneously at high spectral resolution. However, in cases where four lines or fewer are observed in total, if the full bandwidth/resolution is needed for each line and they are well-separated, it may be preferable to place each in a different baseband. Alternatively, in cases where the lines are close to one another, or less spectral resolution/bandwidth is acceptable, it may be better to include multiple lines in the same baseband or even in the same spws, and use the other basebands in TDM mode to maximize the total (aggregate) bandwidth for calibration. An example of this is shown in Table 6.1 for spectral setup 6 b, where both $C^{18}O$ and ^{13}CO are observed in the same baseband (3.1 and 3.2). Channel-averaging can be applied to each spws, which simply bins spectral channels together in the correlator data processing (see Section 6.5).

Multiple spws per baseband does have some restrictions: each spws within the same baseband must have the *same* correlator channel width (i.e. before channel-averaging)⁴, and the sum of the correlator resources in each baseband should add up to 1.0 (or less). Given these restrictions, it is possible to set up different bandwidths and/or different channel-average values for each spws. Multiple spws are set up in the Phase 1 OT simply by adding more spectral lines in the same baseband, and setting the spws to have a correlator fraction less than 1.0 (so for 4 spws, each has a correlator fraction of 1/4). There may be up to 4 spws per baseband, however, the total number of spectral channels in each baseband is limited by the correlator resources. So doubling the number of spws in one baseband will result in half the number of channels per spws, i.e. a lower resolution.

Figure 6.3, which is adapted from the Phase 2 Spectral Editor of the Observing Tool (OT), illustrates a more complex spectral setup with multiple spws and basebands. The lines observed are given in Table 6.1, example 7 b. In this case, the frequency of LO1 is 350.0 GHz, and the upper and lower sidebands of the Band 7 receiver are shown as green shaded areas. The four basebands, illustrated in this case by the red horizontal bars, can be moved around, but only within the two sidebands. The spws are also shown, labelled by the name of the primary spectral line or continuum. In this example, baseband 1 is set for a TDM spw covering the whole 2 GHz baseband for a continuum measurement, whereas in other basebands, FDM is used. Baseband 2 has three spws with 0.56 MHz resolution and channel-average of 2; two of 234 MHz bandwidth and 1/4 of the correlator resources per spw, and one with 469 MHz and 1/2 of the correlator resources. At the same time, in BB3 one line is observed with 469 MHz bandwidth and 0.49 MHz resolution (with a factor of 4 channel-averaging). In BB4, one spw has 234 MHz bandwidth and 0.28 MHz resolution (with 2 channel-average) and one spw has 117 MHz bandwidth and 0.48 MHz resolution (with 4 channel-average). Basebands (and spws) are allowed to partially overlap, although the overlapping region only counts once in the aggregate bandwidth and estimate of continuum sensitivity. So it may be preferable to observe closely-spaced lines with a single baseband (e.g. $C^{18}O$ and ^{13}CO in example 6 b above) or better, with a single spw, as this can free an additional TDM spw for improved continuum sensitivity. Many combinations of spws and basebands can be set up in this way.

⁴although, as before, spws in *different* basebands can have different channel widths

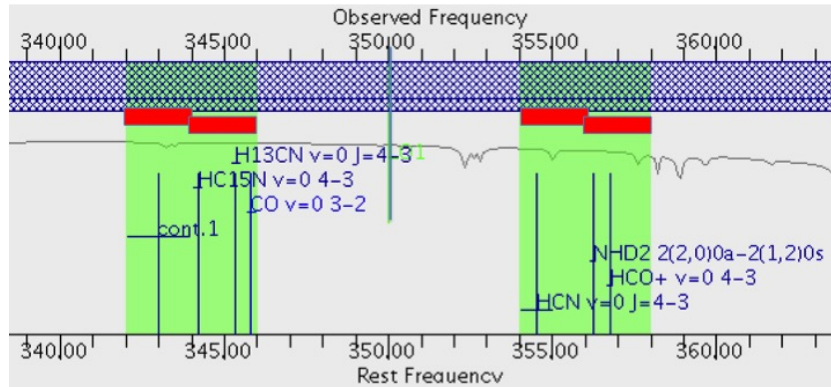


Figure 6.3: Illustration of a frequency setup, based on the OT spectral display. Green areas are the IF ranges, horizontal red bars are the 2 GHz-wide basebands (BB1–BB4), and smaller horizontal lines represent the spws (up to 3 per baseband in this example). The frequency of LO1 is shown by the central vertical line. The blue hashed area represents the tuning range of the front end (in this case a section of Band 7), and the curved line represents the nominal atmospheric transmission for the chosen PWV. The bandwidths of the spws are illustrated by the widths of the horizontal lines. The lines observed are given in example 7 b in Table 6.1.

6.3.3 Spectral Setups for Lines near the Edge of the Bands

The restriction that the baseband coverage cannot fall outside the maximum or minimum tuning range of the receiver is an issue for certain lines at the edge of the tuning range. One obvious example is the $^{12}\text{CO}(J=1-0)$ line at a redshift of zero (rest frequency of 115.271 GHz). This is close to the maximum tuning range of Band 3 (116.000 GHz, see Table 6.1). A setup using the wide bandwidth modes (TDM or FDM/1875 MHz) centered at the line frequency with zero redshift will not validate in the OT, because some of the basebands will fall outside the maximum Band 3 frequency (116.0 GHz). Narrower modes will not go over the edge and are allowed.

There are two possible solutions: If the full bandwidth is absolutely required, the center rest frequency can be set to the closest valid frequency, resulting in an offset of the line from the spws center (in this case, set to 115.0 GHz and the line will be offset by 0.271 GHz). Another solution is to choose a narrow spws bandwidth (e.g. 1 GHz or less); in this case the spws will be offset from the center of the baseband, and the line will be at the center of the spws.

6.3.4 Spectral Setup in Bands 9 and 10: DSB Considerations and 90-degree Phase Switching

Bands 9 and 10 use double-sideband (DSB) receivers. If only one sideband is required (the so-called 'signal' sideband) the other sideband (known as the 'image' sideband) can either be suppressed using LO offsetting (currently using LO1 and LO2). Alternatively, both sidebands can simultaneously be recorded using 90-degree phase switching, although this can only be used in interferometric mode and is now the default - see Section B.4.2.

Observing a single sideband

The choice of which sideband a particular baseband is configured to observe (and which sideband is suppressed) depends only on the relative sign of the LO1 and LO2 offset. This can be arbitrarily different for different basebands. So it is possible to set up two basebands at approximately the same LO1/LO2 frequency, but

observe one line in the upper sideband in one baseband, and another in the lower sideband using the next baseband, just by having different signs in the LO2 offsetting. However, in this case it might be better to use 90-degree switching (see next section) to record both using the same baseband. Note that different spws in the *same* baseband must have the same sideband, as image suppression uses LO2, which is common to the whole baseband.

Observing both sidebands - 90-degree switching

Starting in Cycle 6, it was possible to double the continuum bandwidth by simultaneously recording both sidebands from the DSB receivers using 90-degree phase switching. This is enabled by phase modulation of LO1, with ± 90 degree demodulation in the correlator, controlled with a Walsh sequence (see Chapter 5; also ALMA Memo 287). It has the advantage of doubling the continuum bandwidth to a nominal 16 GHz per polarization (2×2 GHz per baseband $\times 4$ basebands) - particularly helpful in these high-frequency bands as it can improve the continuum sensitivity by $\sim \sqrt{2}$, as well as making lines available in the opposite sideband. The actual improvement however, depends on the atmospheric transmission in the opposite sideband, which in Bands 9-10 can result in image system temperatures 20% (or more) higher or lower than the main signal sideband. Even though this doubles the number of spws there is no loss of correlator spectral resolution, so the correlator data rate is doubled. In rare cases, this may cause correlator data rate limitations in FDM mode, which could be mitigated by longer integration times, lowering resolution by increasing the channel-averaging factor, and/or using TDM in the continuum-only basebands. For Cycle 8, there is a restriction that this mode can only be used with TDM or the maximum bandwidth (1875 MHz) FDM configuration, and the spws can only be centered on each baseband even when the spectral scan mode is selected (Section 6.6). It also requires that the integration duration should be a multiple of 2048 ms.

6.4 Usable Bandwidth

The IF system contains an anti-aliasing filter which limits the bandwidth of the basebands. Nominally this filter has -1 dB points at 2.10 and 3.90 GHz, giving a maximum bandwidth of 1.8 GHz. However, the IF response is such that the usable bandwidth is slightly wider - i.e. closer to 1.9 GHz. In FDM mode, the correlator outputs a maximum bandwidth of 15/16 of the nominal bandwidth, and reduces the number of channels by the same factors, from 4096 to 3840 (or 2048 to 1920 after the default channel-average factor - see Chapter 5). So, for a nominal 2 GHz spectral setup, the correlator outputs a bandwidth of 1.875 GHz. Thus in FDM wide-band, the filters do not truncate the spectrum, and the full available correlator bandwidth in wide-band mode can be used. In TDM the correlator outputs a bandwidth of 2.000 GHz, but typically the edges of the spectra are affected by low power due to this filter and some ringing effects (see upper panel in Figure 6.4). It is recommended that 4 (in double-polarization) or 8 (single-polarization) channels are removed or flagged manually offline. This results in approximately the same usable bandwidth in both TDM and FDM modes and is illustrated in Figure 6.4. Note that if the centers of two basebands are separated by 1.875 GHz it is possible to set up spws that can be joined in data reduction to yield a contiguous 3.75 GHz spectrum.

6.5 Spectral Resolution and Channel-averaging

The spectral resolution of each spws is set by a combination of the inherent resolution/bandwidth setup in the correlator, the correlator weighting function (also known as the smoothing function) and the channel-averaging (also known as spectral-averaging) factor (which is carried out at a late stage of processing in the correlator). This is described in Section 5.5.2. The default FDM setup from the OT in Cycle 8 uses Hanning smoothing with a channel-averaging factor of $N=2$ (see Section 5.5.2). Hanning smoothing plus a channel average of 2 results in a spectral resolution equal to the *final* channel spacing $\times 1.15$ - only $\sim 15\%$ worse than Hanning without channel-averaging but has the advantage of halving the data rate. If the line is well resolved, with ≥ 3 channels covering the linewidth, then this will not cause a problem. However, marginally resolved lines should keep the maximum resolution. If possible it is recommended to use this default in order to reduce the data rate and size of the final product. If the extra 15% resolution is required or the line is expected to be marginally resolved, a justification for using a channel average of $N=1$ should be made. Note that for TDM mode, Hanning smoothing without channel averaging ($N=1$) is the default.




Figure 6.4: Comparison of TDM (left) and FDM (right) auto-correlation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in power in the upper and lower ~ 4 channels due to the anti-aliasing filter. In this case the FDM spectrum has 3840 channels covering only the central 1.875 GHz, and the drop in power at the edges of this bandwidth is negligible (and comparable with the variations in the bandpass). (The narrow spike in the center of the FDM data is a test signal).

6.6 Spectral scans

For coverage of frequency ranges greater than the IF bandwidth, a spectral scan observing mode is available. If spectral scan is selected in the OT, this will observe up to 5 different LO1 frequency tunings to cover the requested range, using all four BBs in USB and LSB. Tunings will be programmed to include an overlap between each setting to avoid issues with edge channels. Band 9 and 10 will also use the 90-degree phase switching to double to effective bandwidth per tuning only if the TDM or the maximum bandwidth (1875 MHz) FDM configurations are selected. If more than 5 tunings are required to cover the range, additional science goals need to be created.

Spectral scans are now observed more efficiently by using the bandpass calibrator as the flux density calibrator. This change, along with all tunings of a calibrator now being observed immediately after each other which reduces the amount of pointing calibrations considerably, should also make spectral scans more efficient. It is recommended that the user selects System-defined calibration in the Calibration Setup in the OT when choosing spectral scan mode. However, the user can select a separate flux calibrator from a solar system object if necessary. It is also possible to now set up a spectral scan observation on ACA only observations.

6.7 Spurious Signals

Most spurious signals in the cross-correlation data are suppressed using a combination of 180-degree phase modulation of LO1 and using different LO offsets for different antennas, both indexed with a Walsh pattern (see Section B.4.4). Not only does this effectively suppress signals generated after the front ends, but it also improves the image sideband rejection for 2SB receivers by an additional factor of ~ 20 dB. However, it does not reduce spurious signals coming in at the observing frequency. Results from tests using these Walsh functions show very few remaining spurious signals, and these are further reduced by the fringe tracking. However, harmonics of the LO in the WVRs (ie 91.66, 183.32, 274.98, 366.64, 458.3 ... GHz) are very bright and cannot be removed through these methods⁵. These are very narrow and may require a few spectral channels to be

⁵Although the WVR LOs have small offsets which are different for each antenna, which helps to reduce their cross-correlated signal

flagged out during data reduction. *Observing lines at these frequencies should be avoided if possible.*

6.8 Doppler Setting and Velocity Reference Frames

In most cases the system will be set up to provide online correction for the science target velocity in a particular reference frame and the Earth motion in that frame. The primary velocity reference frames recommended for use in ALMA are:

Topocentric In this case *no* correction for the source or Earth motion is made. The Center Frequency (both Rest and Sky) will be identical.

Barycentric This is with respect to the center of mass of the Earth-Sun system, and is very close to the heliocentric frame.

LSRK Velocity with respect to the Kinematic local standard of rest, at 20.0 km/s in the direction $18^h, +30^\circ$ [B1900.0]. This frame is based on the mean velocity of the stars in the Solar neighborhood.

If a target velocity and reference frame other than topocentric is selected in the OT, at the start of the observation, the velocity of the science target in the chosen reference frame *and* the velocity of the observatory relative to the chosen reference frame are used to set the center frequencies of the science spws⁶. For example, for a source of radial velocity 100 km/s in the LSRK frame, the difference between rest and sky frequency in the OT will be the equivalent of 100 km/s. At run-time, the extra velocity from the Earth orbit and rotation (up to ± 30.3 km/s and ± 0.46 km/s), plus the motion of the LSRK (20.0 km/s toward RA=18 hrs, $\delta = 30^\circ$ for epoch 1900.0) will be included. This means that the LO solutions are slightly different at run-time compared with those from the OT, although this is normally transparent to the user⁷. It is also possible to just make the correction for the motion of the observatory relative to the chosen reference frame, ignoring the source velocity in that frame. The frequency setting for the science target will normally also be used in the same execution for both the bandpass and amplitude calibration.

For sources with an external ephemeris file, the rate of change of distance between the target and the observatory taken from the ephemeris is used to compute the source velocity at the start of the execution. This velocity is used throughout the observation (like non-ephemeris observations, the velocity from the ephemeris is not updated). Combining executions with different velocities must be done in CASA offline.

The user can specify either the velocity or the redshift of the target. Again, it is strongly recommended that the user input the appropriate source velocity or redshift and reference frame in the Field Setup tab of the OT (Chapter 9). Note that velocities can be defined in three different ways, resulting in different conversions to frequency. These differences only become significant at high velocities.

The velocity formulae are:

Radio $v = c(f_o - f)/f_o$. This is the default.

Optical $v = c(f_o - f)/f$

Relativistic $v = c(f_o^2 - f^2)/(f_o^2 + f^2)$.

In these equations, v is the source velocity, and f and f_o are the observed and the rest frequencies of the line. If a redshift z is specified, the conversion to frequency uses $z = (f_o - f)/f$. For further details, see the Knowledgebase article "What are the frequency reference frames in CASA?"

⁶Note that ALMA does *not* do Doppler *tracking*, where the frequency would be continuously updated for Earth motion during the observation. Doppler corrections are only *set* once at the start of each execution. Compensation for the small changes during an observation (<0.1 km/s for an execution lasting 2 hours) are made during CASA offline reduction.

⁷The system allows for a small (< 1 MHz at Band 6) overhang of the spws outside the 2 GHz-wide basebands at runtime, to allow for the very slightly different Doppler corrections in different spws, which will shift the spws with respect to one another.

Rule	Description
1.	LO1 must lie within the LO tuning ranges given in Table 6.1.
2.	No part of the 2.0 GHz-wide basebands can extend over the edge of the IF passband. So, the baseband centers cannot be closer than 1.0 GHz to the IF passband edge. For example, with a 4.0–8.0 GHz IF range, the baseband center frequency must lie between 5.0–7.0 GHz. The system actually does allow a small extension of the edges of the basebands over the IF edges, to cope with the differential Doppler shifts in the different basebands, but this is small (< 1 MHz at Band 6) and is transparent to the OT user (see Section 6.8).
3.	For 2SB receivers (Bands 3–8), the number of basebands in one sideband can only be 0, 1, 2, 3 and 4. However, note that three basebands in one sideband and one baseband in the other sideband is not allowed whereas the combination three-zero is. For DSB receivers (Bands 9 and 10), there is no such restriction (the number can be 0, 1, 2, 3 or 4).
4.	For > 2 basebands, BB1 and BB2 should be in the same sideband, as should BB3 and BB4.
5.	No part of the full nominal bandwidth of the spws can extend over the edge of the 2 GHz-wide baseband. For a mode with nominal bandwidth B (e.g. 62.5 MHz), that means the spws center IF frequency (a.k.a. Center Offset Frequency in the OT, Phase 2) must be $>(2000+B/2)$ and $<(4000-B/2)$ MHz. The current version of the OT forces this restriction. For 2 GHz FDM and TDM modes, this means that the spws must be at the center of the baseband. However there is a further restriction on this, as noted in the next rule.
6.	The spws usable bandwidth (ie 15/16 of the nominal spws bandwidth of a multiple of 62.5 MHz) should be in an allowed region of the baseband. This is in addition to (5). In practice this means that the required range of the spws should normally be inside the range $\sim 2050 - 3950$ MHz (i.e. >50 MHz from the edges of the 2–4 GHz second IF) to ensure that the edge of the anti-aliasing filter does not significantly affect the IF power. In practice it is possible to extend some fraction of the spws to <50 MHz of the IF edges (although this is no longer so critical for the Band 6 setup 6 b in Table 6.1, as the allowable IF range has been extended down to 4.5 GHz).
7.	The line frequency does not need to be in the center of the spws in Phase 1 of the OT: if a line like ^{12}CO 1-0 is requested, it will generate an SB with the correct TFB LO offset, as long as rule 5 is obeyed (which may require using a narrower spws). This is mentioned in more detail in Section 6.2
8.	Only 2-bit, Nyquist sampling is allowed in the correlator.
9.	It is possible to have multiple targets with different redshifts within the same Science Goal in the OT. For SGs including sources with more than one redshift, all the observations must be achievable using five or fewer tunings within the same receiver band, considering the source redshifts and, in the case of spectral lines, the line widths and configuration of spws.
10.	The number of spws per baseband can be 1, 2, 3 or 4. For 3 spws, the correlator resources per spws should be set to 1/4, so only 3/4 of the available resources are used in this case.
11.	All the spws <i>within the same baseband</i> are required to have the same correlator channel width (before channel-averaging). An individual spws within a baseband may occupy 1, 1/2 or 1/4 of the resources available in the baseband and the sum of the fractional resources within one baseband must be ≤ 1 . The correlator resources are proportional to the number of correlator spectral channels.
12.	With DSB receivers (Band 9 and 10), 90-degree phase switching is allowed, giving both sidebands simultaneously. This is only available for TDM and FDM 1875 MHz bandwidth modes.

Table 6.2: Rules for spectral setups.

Chapter 7

Imaging with ALMA

7.1 Introduction

Imaging ALMA data consists of Fourier transforming calibrated interferometric visibilities from one or more configurations or arrays; de-convolving the interferometer point spread function (PSF); and, if Single-Dish data have been collected, merging them with the interferometric image or cube. We describe in this chapter some of the key considerations of this process that should be taken into account in order to obtain the best results from ALMA observations.

As described in Chapter 3, the van Cittert-Zernike theorem is the mathematical foundation of synthesis imaging. It describes a fundamental relationship between the sky brightness distribution (I), the beam pattern (A) and the visibility distribution \mathcal{V} :

$$A(l, m)I(l, m) = \int \int \mathcal{V}(u, v) e^{2\pi i(ul+vm)} du dv \quad (7.1)$$

A more detailed discussion of the van Cittert-Zernike theorem can be found, e.g., in Rohlfs & Wilson (2004). An interferometric observation can be represented by N discrete points in the uv plane $\mathcal{V}_k(u_k, v_k)$ where $k = 1, \dots, N$. The variables u_k and v_k represent the x and y components of a specific k^{th} baseline between a pair of antennas, measured in wavelengths λ . There are in general $N_{\text{ant}}(N_{\text{ant}} - 1) \times 0.5$ distinct baselines; factoring in distinct frequency channels and temporal integrations, N can reach a few million or more for ALMA. One implication of the van Cittert-Zernike theorem, which will be discussed in more detail elsewhere in this chapter, is that *the interferometer point-spread function (PSF) is the Fourier Transform of the set of (u_k, v_k) points measured*. In the parlance of radio astronomy this set of (u, v) points is often called the “ uv -coverage” of an observation, and the PSF is called the “synthesized” or “dirty” beam. In practice it is often useful to define a weighted visibility function \mathcal{V}^W (see Briggs, Schwab & Sramek, 1999) that enables the properties of the synthesized beam to be controlled to some extent, often at the expense of sensitivity. The weighted visibility function can be represented as:

$$\mathcal{V}^W(u, v) = \sum_{k=1}^N R_k T_k D_k \delta(u - u_k, v - v_k) \mathcal{V}_k(u_k, v_k). \quad (7.2)$$

In Equation 7.2, R_k , T_k and D_k are weights which control the relative weighting of individual visibilities in imaging. R_k is a noise-variance weight derived from the sensitivity of the k^{th} visibility measurement. It accounts for variables such as the integration time, the system temperature, and the bandwidth. These factors are determined by the instrument, technical set-up, and observing conditions, and are not controlled by the imaging process. T_k is a “taper” weight—usually taken to be a multiplicative, Gaussian factor in the uv plane—that can be used to down-weight the longest baselines, suppress small-scale sidelobes, and increase the synthesized beam width and surface brightness sensitivity. Finally, D_k is a uv “density” weight that can be used to offset the high concentration of measurements near the center of the uv plane, typically increasing angular resolution and reducing the large-scale sidelobes caused by abrupt gaps in the uv -coverage density.

The following sections describe: the ALMA configurations (Section 7.2); the effects of antenna shadowing (Section 7.3); the factors that affect angular resolution and beam shape (Section 7.4); the special case of

short “snapshot” observations (Section 7.5); the large spatial scale response of ALMA (Section 7.6); mosaicing (Section 7.7); multi-array and multi-configuration ALMA observations (Section 7.8); and, finally, multi-array and multi-configuration imaging (Section 7.9).

7.2 Configurations

Depending on the range of angular scales required, an ALMA Science Goal will obtain data from either: a single 12-m Array configuration; two 12-m Array configurations; one or two 12-m Array configurations plus the 7-m Array; or one or two 12-m Array configurations plus the 7-m and Total Power (TP) Arrays¹. The detailed logic by which specific combinations are chosen based on the PI Science Goals is described in Section 7.8. Generally, the PI’s desired angular resolution determines the most extended array, while the largest angular scale of interest (or LAS) determines the most compact array— whether that is a second 12-m Array, the 7-m Array, or the Total Power in order of progressively larger spatial scales. Intermediate configurations or arrays are added to connect these spatial scales as required. Note that, at present, total power continuum observations are not offered, limiting the largest angular scale of ALMA continuum observations to that provided by the 7-m Array.

During Cycle 8, the 12-m Array will be arranged in eight different configurations. As discussed below, operational factors impact the actual configuration achieved for a given observation, so the configurations presented in this chapter that are to be used for simulations or project planning are referred to as “notional configurations”. These notional imaging configurations of 43 12-m antennas (C43-X) are available on the Science Portal and in CASA. In real operations, configurations of all 50 antennas are used, out of which at any given time, and depending on the requested band, a number of antennas may not be available for observations. These notional configurations are identical to those used since Cycle 5. ALMA is committed to providing at least 43 antennas for each 12-m Array observation and on average provides more than 43 antennas; in exceptional circumstances, observations may occasionally be made with an array with up to two fewer antennas. In Cycle 8, the two most extended configurations (C-9 and C-10) are not offered. The two most extended configurations offered in Cycle 8 will be C-7 and C-8; the longest baseline offered will be 8.5 km for Bands 3-7 and 3.6 km for Bands 8-10. The 7-m Array will be available in only one configuration of twelve antennas, where ten antennas should be provided in practice, except in exceptional circumstances where observations may be carried out with eight or nine antennas.

Tables 7.1 and 7.2 give the basic properties of the 12-m Array and 7-m Array configurations. These properties are based on simulations which use the notional antenna configurations described above (*i.e.*, for the 12-m Array the C43-X configuration files available on the Science Portal were used). Each configuration is characterized by a *maximum recoverable* (spatial) *scale* θ_{MRS} and an *angular resolution* θ_{res} . Generally, these are taken to be determined by the shortest and longest baselines in the configuration, respectively. As discussed in Sections 7.4 and 7.6, ALMA uses the somewhat more robust 5th and 80th percentile baseline lengths instead. Figures 7.2 and 7.3 show the notional configurations of the 12-m Array, which are the 43 pads in each configuration that have the highest operational priority in terms of antenna availability. Figure 7.4 shows the notional 10-antenna configuration of the 7-m Array. The three TP Array single-dish 12-m antennas will be in fixed positions, although their relative spatial locations do not substantially affect their imaging characteristics.

12-m Array configurations in particular are impacted by a variety of “real-world” factors including, but not limited to, the availability of specific antenna pads², and necessary, ongoing antenna relocations. The Scheduling subsystem takes these changes into account alongside the performance margins that the ALMA QA2 process allows. Thus, from the user’s point of view, configurations are not directly selected: instead, the required angular resolution (θ_{res}) and largest angular scale (θ_{LAS}) are specified in the ALMA Observing Tool (OT), and projects are observed using array configurations that will achieve the specified science goals. Since Cycle 5 it has also been possible for PI’s to specify an acceptable *range* of θ_{res} . More details about array combinations are provided in Section 7.8.

As mentioned above, the antenna configurations used in real operations (denoted C-X) can differ from the C43-X notional configurations, and not all 50 antennas may be available for operations at any given time. These real-world configurations often use pads of the next more compact C43-X configuration and typically

¹The combination of the 7-m Array and the Total Power Array constitutes the ALMA Compact Array (ACA). The ACA is also known as the Morita Array in honor of Prof. Koh-ichiro Morita.

²ALMA’s 54 12-m antennas can be distributed among 192 “antenna pads” which provide power, IF signal & network connections, and a structurally stable foundation.




Figure 7.1: ALMA's high resolution images of nearby protoplanetary disks, which are results of the Disk Substructures at High Angular Resolution Project (DSHARP). The data provided by the observing team are available at <https://almascience.org/almadata/lp/DSHARP/> Credit: ALMA (ESO/NAOJ/NRAO), S. Andrews et al.; NRAO/AUI/NSF, S. Dagnello

	Band	3	4	5	6	7	8	9	10
	Frequency (GHz)	100	150	185	230	345	460	650	870
Config.									
7-m	θ_{res} (arcsec)	12.5	8.35	6.77	5.45	3.63	2.72	1.93	1.44
	θ_{MRS} (arcsec)	66.7	44.5	36.1	29.0	19.3	14.5	10.3	7.67
C-1	θ_{res} (arcsec)	3.38	2.25	1.83	1.47	0.98	0.735	0.52	0.389
	θ_{MRS} (arcsec)	28.5	19.0	15.4	12.4	8.25	6.19	4.38	3.27
C-2	θ_{res} (arcsec)	2.30	1.53	1.24	0.999	0.666	0.499	0.353	0.264
	θ_{MRS} (arcsec)	22.6	15.0	12.2	9.81	6.54	4.9	3.47	2.59
C-3	θ_{res} (arcsec)	1.42	0.943	0.765	0.615	0.41	0.308	0.218	0.163
	θ_{MRS} (arcsec)	16.2	10.8	8.73	7.02	4.68	3.51	2.48	1.86
C-4	θ_{res} (arcsec)	0.918	0.612	0.496	0.399	0.266	0.2	0.141	0.106
	θ_{MRS} (arcsec)	11.2	7.5	6.08	4.89	3.26	2.44	1.73	1.29
C-5	θ_{res} (arcsec)	0.545	0.363	0.295	0.237	0.158	0.118	0.0838	0.0626
	θ_{MRS} (arcsec)	6.7	4.47	3.62	2.91	1.94	1.46	1.03	0.77
C-6	θ_{res} (arcsec)	0.306	0.204	0.165	0.133	0.0887	0.0665	0.0471	0.0352
	θ_{MRS} (arcsec)	4.11	2.74	2.22	1.78	1.19	0.892	0.632	0.472
C-7	θ_{res} (arcsec)	0.211	0.141	0.114	0.0917	0.0612	0.0459	0.0325	0.0243
	θ_{MRS} (arcsec)	2.58	1.72	1.4	1.12	0.749	0.562	0.398	0.297
C-8	θ_{res} (arcsec)	0.096	0.064	0.0519	0.0417	0.0278	-	-	-
	θ_{MRS} (arcsec)	1.42	0.947	0.768	0.618	0.412	-	-	-
C-9	θ_{res} (arcsec)	0.057	0.038	0.0308	0.0248	0.0165	-	-	-
	θ_{MRS} (arcsec)	0.814	0.543	0.44	0.354	0.236	-	-	-
C-10	θ_{res} (arcsec)	0.042	0.028	0.0227	0.0183	0.0122	-	-	-
	θ_{MRS} (arcsec)	0.496	0.331	0.268	0.216	0.144	-	-	-

Table 7.1: Resolution (θ_{res}) and maximum recoverable scale (θ_{MRS}) values for different 7-m Array and 12-m Array configurations are shown for one representative frequency in each ALMA receiver band. The value of θ_{MRS} is computed using the 5th percentile baseline (L05) from Table 7.2 and Equation 7.7. The value of θ_{res} is the mean size of the interferometric beam obtained through simulations with CASA. Computations were done for a source at zenith; for sources transiting at lower elevations, the North-South angular measures will increase proportional to $1/\sin(\text{elevation})$. Calculations are based on the notional C43-X and 7m configurations. *Note that C-9 and C-10 are not offered in Cycle 8.*

have +4.3% (+/- 1.4%) larger beams than the notional configurations, apart from C-8 (+10.4%) and C-9 (-7.1%). A priority is set for the availability of the outer antennas defining the resolution. In practice, the actual configurations have a spread of resolutions (5-10%). The final beam resolution of a data set is further affected by the projected baselines during an observation, $1/\sin(\text{elevation})$ and the uv-weighting applied during data reduction, e.g. choice of the Robust parameter (20-40%). Both the OT and scheduling system take the effect of projected baselines on resolution into account, with the OT assuming a nominal elevation and the scheduling system using the actual parameters at the time of observation. Although the 7-m Array is not subject to all of these constraints, its configuration is typically somewhat elongated North-South to facilitate observing targets at a range of Declinations without excessive shadowing (§ 7.3).

In Table 7.1, angular resolutions for the 12-m array are calculated using the notional C43-X antenna configurations for a 2-hour observation centered on transit (at zenith) with $Robust = 0.5$ and $imsize = 100$. For the 7-m Array, imaging parameters are the same but angular resolutions are calculated using the notional 10-antenna 7-m Array and a 20-minute track centered on zenith.

7.3 Shadowing

During interferometric array observations at low elevations — particularly in compact configurations — one antenna’s view of the source can be partially blocked by another antenna, corrupting data from the first antenna. This phenomenon is known as shadowing. Given ALMA’s location (latitude = -23.02917°), targets as far North

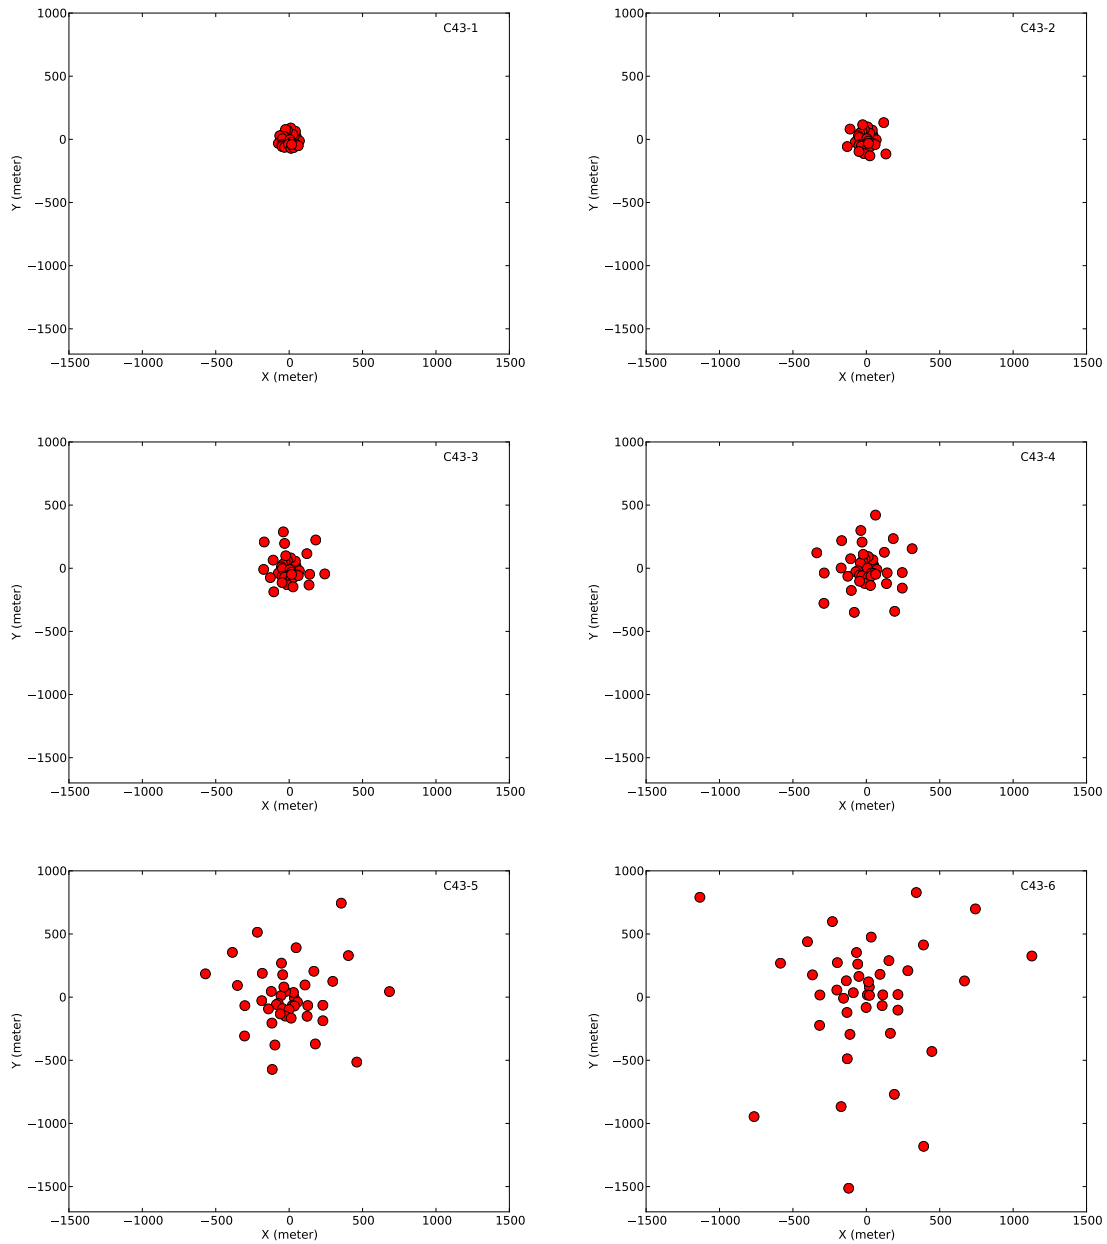


Figure 7.2: Notional (C43-X) 12-m Array compact configurations available for running simulations and planning observations. The actual array configurations available during a Cycle can be different than these notional configurations.

Configuration	7-m	C43-1	C43-2	C43-3	C43-4	C43-5
Minimum baseline (m)	8.7	14.6	14.6	14.6	14.6	14.6
5th percentile or L05 (m)	9.1	21.4	27.0	37.6	54.1	90.9
80th percentile or L80 (m)	30.7	107.1	143.8	235.4	369.2	623.8
Maximum baseline (m)	45.0	160.7	313.7	500.2	783.5	1397.9
Configuration	C43-6	C43-7	C43-8	C43-9	C43-10	
Minimum baseline (m)	14.6	64.0	110.4	367.6	244.0	
5th percentile or L05 (m)	148.6	235.2	427.3	746.9	1228.1	
80th percentile or L80 (m)	1172.5	1673.1	3527.3	6482.6	8685.9	
Maximum baseline (m)	2516.9	3637.8	8547.7	13894.2	16194.0	

Table 7.2: Basic parameters of the notional 7-m Array and 12-m Array configurations. The baselines are projected for a source transiting at zenith.

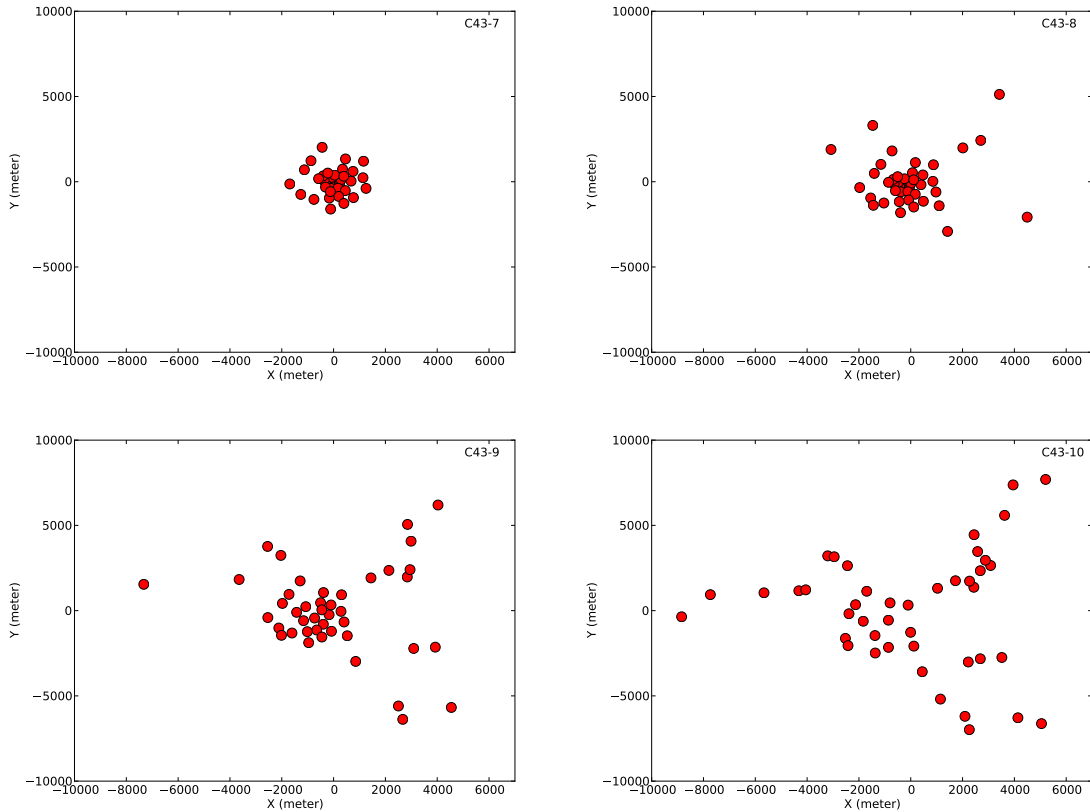


Figure 7.3: Notional (C43-X) 12-m Array extended configurations available for running simulations and planning observations. The actual array configurations available during a Cycle can be different than these notional configurations.

as declination $+47^\circ$ can in principle be observed, corresponding to a maximum source elevation at Chajnantor of $\sim 20^\circ$. However, shadowing by adjacent antennas can be a significant problem at such low elevations. Figure 7.5 shows the fraction of data that will be shadowed (*i.e.*, the “shadowing fraction”) when sources of various declinations are observed with the most compact 12-m notional ALMA configurations (C43-1 and C43-2) and the 7-m Array. As can be seen, the shadowing fraction can be as large as 55% for sources observed with the most compact 12-m Array configuration. Image quality and time on source will necessarily be limited for such northern sources. Given the short baselines in the 7-m Array, sources with declinations less than -70° or




Figure 7.4: Notional 7-m Array configuration available for running simulations and planning observations.

greater than $+25^\circ$ are subject to significant shadowing. For the 12-m Array, shadowing becomes significant ($> 5\%$) in the most compact configuration for sources with declination lower than -65° or higher than $+20^\circ$.

Note that independent of shadowing, observations at low elevation often result in elliptical uv -coverage and, therefore, elliptical synthesized beams. The OT will provide an estimate of the beam ellipticity for those observations requesting low elevation targets in order for the PI to properly plan these types of observations and anticipate the data products that will be provided given such an elliptical beam.




Figure 7.5: Shadowing fraction (%) for the most compact configurations as a function of declination (for a 1-hour observation at transit).

7.4 Angular Resolution and Beam Shape

7.4.1 Definition of angular resolution

The angular resolution θ_{res} of an array can in general be estimated with the following equation:

$$\theta_{res} = \frac{k\lambda}{L_{max}} \text{ [radians]} \quad (7.3)$$

where θ_{res} is the FWHM of the main lobe of the synthesized beam; k is a factor that depends on the uv plane weighting function; λ is the observing wavelength in meters; and L_{max} is the longest baseline in meters. For commonly used weighting schemes, k is typically in the range $0.7 < k < 1.2$. Since uv -coverage is in general non-uniform, however, the longest baseline is often not a representative value for the largest distance at which there is significant uv -coverage. Thus L_{max} is a somewhat fragile indicator of angular resolution and ALMA uses the 80th percentile of the uv distance as a more robust proxy for angular resolution. Via simulations using the ALMA notional configurations, the following equation was empirically determined:

$$\theta_{res} \approx \frac{0.574\lambda}{L_{80}} \text{ [radians]} \quad (7.4)$$

where again λ is the observing wavelength in meters, and L_{80} is the 80th percentile of the uv distance in meters (see also Table 7.2).

7.4.2 Effects of weighting scheme

As shown in Equation 7.2, the synthesized beam is affected by uv plane weights, including the user defined *taper weights* T_k and *density weights* D_k . Taper weighting schemes generally take the form of a Gaussian envelope function in uv distance, and are not used within the ALMA pipeline. Two commonly used, limiting forms of density weighting are *natural weighting* and *uniform weighting*:

- $D_k = 1$ is called natural weighting, and results in maximum sensitivity but a relatively large synthesized beam due to the typically high density of points in the inner region of the uv plane.
- $D_k = \frac{1}{N_s(k)}$ is called uniform weighting, where $N_s(k)$ is the number of visibilities in a region centered on the k^{th} visibility³. It removes the dependence of spatial-scale sensitivity on the density of visibilities (samples) in the uv plane. Uniform weighting increases angular resolution at the expense of sensitivity.

To bridge the extremes of natural and uniform weighting, Briggs (1995) defined a continuous scheme that uses a ‘‘robustness’’ parameter R (called *robust* in CASA). In CASA, uniform weighting is close to *robust* = -2 and natural weighting is close to *robust* = 2 . Figure 7.6 shows the angular resolution achieved for an observation with the notional C43-4 configuration at 100 GHz using *robust* values between -2 and 2 . As can be seen, the angular resolution varies from $0.7''$ (*robust* = -2) to $1.2''$ (*robust* = 2). Unless otherwise specified, the angular resolutions and time ratios presented in this chapter were computed using CASA simulations with Briggs weighting and *robust* = 0.5 . The ALMA OT also uses the C43-X configurations and Briggs weighting with *robust* = 0.5 for calculations of angular resolution and sensitivity. Real observations are processed in the ALMA Pipeline with an initial default of *robust* = 0.5 for 12-m Array observations, but a range of values between 0 and 2 are additionally considered and the value which best meets the desired θ_{res} and sensitivity is used. A value of *robust* = 0.5 is always used when processing 7-m Array observations.

It is important to note that the degree of impact that a given *robust* value has on the final beam size is dependent on the uv -coverage of the gridded data to be imaged. Snapshot data will have less uv -coverage than longer integrations, as will single channels of an image cube compared to continuum MFS images. When the fractional bandwidth of a spectral window being imaged is high enough, the full set of channelized visibilities from a given baseline in dense inner region of the uv plane may fall within a single grid cell, but span multiple grid cells in the outer regions of the uv plane. Thus, the gridded $N_s(k)$ used in calculating D_k will change depending on this fractional bandwidth and affect what resolution *uniform* weighting yields (along with the entire robust

³In gridded imaging such as that used by CASA, uniform weighting is normalized by the number of visibilities in a given grid cell, rather than a region directly centered on each visibility. This is the origin of the effects discussed for the *perchanweightdensity* parameter.




Figure 7.6: Angular resolution achieved using different values of the CASA *robust* parameter for a 1-hour observation at 100 GHz and a declination of -23 deg in the notional C43-4 configuration. Note that *robust* = -2 is close to uniform weighting and *robust* = 2 is close to natural weighting. The dotted line corresponds to $k = 1$.

trade-off curve). Similarly, the image size (*imsize* in CASA) can also have a few to 10% effect since the finer uv gridding associated with a larger field-of-view and number of image pixels changes the weighting scheme and its effect on the synthesized beam.

Additionally, the exact method used for the *uniform* weighting normalization has an impact on the resultant resolution for spectral cubes. CASA versions prior to 5.6 had historically used a method in which the full dataset was used to calculate $N_s(k)$, and then these weights were applied to the data in each channel. As the $N_s(k)$ calculated for the full dataset is a less good descriptor of the uv density distribution for edge channels than it is for central channels (due to the fractional bandwidth effects described above), this results in a ‘spectral curvature’ where edge channels have a different beam size and RMS noise compared to central channels.

In CASA 5.6, a new parameter *perchanweightdensity* was introduced to TCLEAN to alleviate this spectral curvature. The old behavior is equivalent to *perchanweightdensity* = *False*, while *perchanweightdensity* = *True* forces $N_s(k)$ to be calculated on a per channel basis. This removes the spectral curvature, but dramatically limits the effective range of beam sizes that Briggs weighting can provide and shifts the beam sizes for a given value of *robust* closer to naturally weighted beam. In CASA 5.6 and above, *perchanweightdensity* is set to *True* as the default. The ALMA project has retained *perchanweightdensity* = *False* for QA2 until such a time as the impact of setting it to *True* can be properly accounted for in all subsystems since it changes the beam size in a way that depends upon the uv -coverage and correlator configuration, and reduces the extent to which changing *robust* can mitigate mismatches between the achieved beam and the PI requested AR. **It is important to note that in order to achieve similar results between the OT and simulations with CASA \geq 5.6, *perchanweightdensity* must be explicitly set to *False*.**

7.4.3 Other factors affecting beam size and shape

Source declination directly impacts uv plane sampling during the observation(s), and thus the beam shape. In addition to shadowing, sources that must be observed at low elevations also have shorter *projected* North-South baselines, causing the shape of the uv plane sampling distribution and the beam shape become more elliptical. For example, Figure 7.7 shows the different beam shapes for sources observed at declinations of -70° (similar to the SMC and LMC) and -30° . Figure 7.8 shows the uv -coverage for these sources, revealing the projection




Figure 7.7: Beam shape for notional configuration C43-1 with a 2-hour observation of a transiting source at a declination of either -70° (left) or -30° (right).

and subsequent ellipticity of the uv plane sampling distribution, together with the large fraction of shadowing, for the more southern sources. To mitigate this effect, one can convolve the resulting images to a more circular (albeit larger) beam after deconvolution.

To demonstrate the range of these effects across the declinations observable with ALMA, Figure 7.9 shows the minor and major axes of the synthesized beam widths (θ_{res}) for each notional array configuration as a function of source declination for a 2-hour observation at 100 GHz, imaged with $robust=0.5$. Figure 7.10 shows the geometrical mean of the major and minor axes of θ_{res} at the same frequency. The beam width scales with λ , but bear in mind that not all configurations can be used with the higher-frequency bands. Also, note that these values are given for the notional configurations. As discussed in Section 7.2, the C-X configurations generally result in $\sim 4\%$ larger beams than the notional C43-X configurations, with real operations falling between these two limiting cases.

If the imaging requirements of a project are particularly scientifically critical, then for purposes of proposal preparation it is advisable to carefully evaluate the expected performance using CASA simulations that employ the notional configurations, target declination, and expected uv -coverage of the proposed observations. In this case it is also advisable to include realistic noise and to CLEAN⁴ the images carefully in order to realistically assess the results that can be expected. Special care should be given to assessing the range of acceptable θ_{res} , taking into account the range of possible achieved θ_{res} for a given request, including the applicable QA2 beam criteria and the in-practice deviations discussed in Section 7.2 (typically a $\sim 4\%$ enlargement due to the presence of additional antennas, although C-8 and C-9 are special cases).

7.5 Snapshots

The sampling function S of the visibility distribution is defined as:

$$S(u, v) = \sum_{k=1}^N \delta(u - u_k, v - v_k) \quad (7.5)$$

⁴Most interferometric imaging in CASA is now done using TCLEAN, which is an improved implementation of the CLEAN algorithm. The CLEAN *task* in CASA has been deprecated; in a near-future release, CLEAN and TCLEAN will be renamed and eventually the “old” CLEAN will be removed from the code.




Figure 7.8: uv plane coverage for notional configuration C43-1 with a 1-hour observation of a transiting source at a declination of either -70° (left) or -23° (right). For the source with a declination of -70° , 25.7% of visibilities are expected to be flagged due to shadowing.




Figure 7.9: Cycle 8 angular resolutions as a function of source declination. Each stripe shows, for a particular notional configuration, the range of the major and minor axes of the synthesized beams expected from a 2-hour observation at 100 GHz and at transit, imaged with $robust = 0.5$. The notional 12-m Array configurations are arranged with C43-1 at the top and C43-10 at the bottom.




Figure 7.10: Geometrical mean of the major and minor axes of the synthesized beams as a function of source declination. These correspond to a 2-hour observation at 100 GHz and at transit, imaged with $robust = 0.5$. The notional 12-m Array configurations are arranged with C43-1 at the top and C43-10 at the bottom. The 7-m Array is also shown.

If S were a continuous function, e.g., a Gaussian, the synthesized beam would also be a Gaussian i.e., a central peak with a smoothly decreasing response away from the center. Given the finite number of baselines, however, the sampling function is an ensemble of Dirac delta functions and its Fourier transform is a central peak surrounded by a complex pattern of sidelobes. This “dirty beam” response is a consequence of the gaps in the uv plane, the sidelobes becoming increasingly prominent as the gaps increase (following the Gibbs phenomenon). Conversely, as there are fewer gaps in the uv plane—i.e. better the uv -coverage—the PSF sidelobes become less prominent.

Because the baseline lengths and orientations as seen from the source change as a function of time due to the Earth’s rotation, the uv -coverage increases and the PSF sidelobes are commensurately reduced in proportion to the duration of the observation. Therefore, long integrations on a target can give better image quality. Short integrations can still be valuable if the science target has a relatively compact, simple structure and is not too faint. Such observations are sometimes called *snapshots*, and they produce obvious sidelobes which can be mitigated by applying a uv taper, i.e. down-weighting of the visibilities on longer baselines, during the imaging process. The uv taper is represented as the set of taper weights T_k in Equation 7.2.

Configuration	7-m	C43-1	C43-2	C43-3	C43-4	C43-5
natural	44.0%	5.5%	5.0%	6.2%	10.4%	6.3%
briggs ($R = 0.5$)	42.2%	7.8%	6.6%	6.7%	7.3%	8.0%
uniform	39.0%	7.7%	16.7%	12.4%	11.0%	13.1%
Configuration	C43-6	C43-7	C43-8	C43-9	C43-10	
natural	4.6%	9.6%	7.8%	11.1%	13.0%	
briggs ($R = 0.5$)	6.2%	5.6%	9.7%	9.6%	11.3%	
uniform	11.3%	11.2%	10.1%	8.5%	12.0%	

Table 7.3: Sidelobe levels for a 1-hour observation of an unresolved source at a declination of -23° with the different notional array configurations. The levels are indicated with three different weighting schemes used for imaging.

Table 7.3 gives the level of sidelobes for each notional configuration, for a 1-hour observation of a source at




Figure 7.11: uv plane sampling distributions of a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the notional C43-1 configuration to observe a source at a declination of -23° .

Dec= -23° . Figures 7.11 and 7.12 show examples of a uv plane sampling distribution and cleaned image for a snapshot of 1-minute duration and a longer integration of 1-hour duration using the notional configuration C43-1. As shown in Figure 7.11, the uv -coverage in 1 minute is quite uniformly sampled, but much less dense than that of the 1-hour integration. In addition, the cleaned images of the snapshot and the longer integration (see Figure 7.12) are similar in terms of angular resolution and the apparent differences are quite small. The main difference between the two (besides sensitivity) is in the dirty beam: the sidelobe level for the snapshot will be much higher than for the 1-hour integration. The snapshot image will then require more careful cleaning in order to avoid introducing spurious sources from the strong sidelobes. Condon et al. (1998) gave a very comprehensive description of how this issue impacted the 20-cm NRAO VLA Sky Survey (NVSS).

A useful practice to disentangle sidelobe effects from real point sources, especially in imaging cases with high dynamic range, is to perform CASA simulations using a component list of the strongest sources together with the notional array configuration. This test allows one to estimate the sidelobe fingerprint left by the strong point sources after deconvolution (cleaning), although the residuals will likely be higher in practice due to imperfect calibration. It is worthwhile to use the interactive mode during the deconvolution so that residuals can be monitored. Finally, careful “masking”—or identification of spatial regions with reliably real emission—can improve the dynamic range. Best results are obtained by carefully editing the mask at each (major) cycle of clean. Recent CASA versions include an automated clean masking heuristic which can improve image quality over what is achieved with a broad or static clean mask (Kepley et al., 2020).

7.6 Large Spatial Scale Response

As described in Chapter 3, an interferometer measures the Fourier components of the sky brightness distribution in an area of the uv plane defined by the array configuration used to observe the target. There is, in particular, some shortest baseline that a given array and configuration measures, ultimately limited by the physical size of the antennas. This shortest baseline corresponds to some *largest spatial scale* which can be usefully imaged by the interferometer.

This spatial filtering is a serious issue that must be considered carefully for each science case. To illustrate the concept, the annularly averaged amplitudes of the Fourier transforms of three uniform disks with sizes of $5''$, $10''$ and $20''$ are shown in Figure 7.13 for an observation at 100 GHz. The smallest uniform disk is closest to a point source and so it has power up to a baseline of 180 m. Meanwhile, the most extended disk has power only up to ~ 40 m. Therefore, an array with only baselines larger than 40 m will not be particularly sensitive to emission on spatial scales larger than $\sim 20''$, and will filter out most of the emission from such an extended disk.

An important consequence of such filtering is that an interferometer only detects a fraction of the total flux




Figure 7.12: Images obtained from a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the notional C43-1 configuration to observe a source at declination of -23° . Black contours at 100, 200, 300, 500, 700 and 900 mJy beam^{-1} are overlaid.

density for sources with emission on size scales larger than its shortest baseline. Indeed, if the source only has structures on size scales larger than the shortest observed baselines, one can “resolve out” the source entirely. In order to ameliorate the effects of spatial filtering on extended sources, the 7-m and TP Arrays are available to supplement the 12-m Array configurations if needed.

To quantify these effects, the maximum recoverable scale (θ_{MRS}) is defined as the largest spatial scale to which a given array is sensitive. It is in principle determined by the length L_{min} of the shortest baseline in the array, which measures spatial scales (*i.e.* Fourier modes with a period) of λ/L_{min} . In practice, the uneven uv coverage at the shortest baseline means that the sensitivity of this measurement to such large structures is not very good, and a smaller value of θ_{MRS} is typically adopted. The exact filtering depends on the details of the large-scale structure as well as the short-baseline uv -coverage, and it is best determined by simulations. ALMA has adopted a criterion of measuring 10% of the total flux density of a uniform disk, which for well-constructed array configurations yields:

$$\theta_{MRS} \approx \frac{0.6 \lambda}{L_{min}} \text{ [radians]} \quad (7.6)$$

where λ is the observing wavelength in meters, and L_{min} is the shortest baseline in meters. For reasons similar to those explained in Section 7.4, ALMA uses the 5th percentile shortest baseline L_5 instead of the very shortest baseline L_{min} to calculate θ_{MRS} ; using simulations of configurations the following equation was empirically determined:

$$\theta_{MRS} \approx \frac{0.983 \lambda}{L_5} \text{ [radians]} \quad (7.7)$$

where λ is the observing wavelength in meters, and L_5 is the 5th percentile of uv distance in meters. This is the equation used to assess MRS in the OT for purposes of choosing configurations during proposal preparation. Table 7.1 lists the θ_{MRS} for the different array configurations. If the scientifically required θ_{res} and θ_{LAS} cannot be achieved by a single array or configuration, multiple 12-m configurations and/or the 7-m and TP arrays, will be called for, as described in Section 7.8. Note that these expressions assume good uv -coverage. This is typically the case for ALMA 12-m Array observations, but may not be for very short 7-m Array observations due to the relatively small number of antennas.




Figure 7.13: Annularly averaged amplitudes of the expected visibilities of three model uniform disks (annular averages) at 100 GHz as a function of uv distance.

To emphasize the dependence of these effects on source structure, Figure 7.14 shows the annularly averaged visibility amplitudes versus uv distance of an example astronomical source, M51. This example is taken from an H- α image, but it is used as a model to indicate the spatial structure expected at 100 GHz. The amplitudes of the visibilities can be approximated as a power law of uv distance with a negative index. This distribution indicates that most of the power is located in larger scale structures and that power decreases rapidly at smaller scales. In this example, the visibility amplitude detected by the 7-m Array (at $k_{uv} \approx 5k\lambda$) would be ~ 23 Jy, whereas that detected by the notional configuration C43-6 ($k_{uv} \approx 200k\lambda$) would only be ~ 0.9 Jy. As each source will be different, ALMA users should ideally use simulations to estimate the power distribution and subsequent flux recovery for targets they wish to observe.

7.7 Mosaicing

The field of view of a single interferometer pointing is determined by the antenna primary beam. A uniformly illuminated, circular aperture will have a beam width at half maximum ($FWHM$) of $1.02 \frac{\lambda}{D}$. It will also have very high side lobes due to the abrupt truncation of the antenna illumination pattern; indeed, these side lobes would be particularly problematic for single-dish observations. Most radio and millimeter receivers illuminate their antennas with approximately Gaussian illumination patterns that smoothly go down to -10 dB to -15 dB response at the edges of the dish. The ALMA feedhorns were designed to illuminate the dish with a -12 dB edge taper to provide a nearly Gaussian primary beam with low sidelobes while preserving as much of the resolution and sensitivity as possible. This trade-off is a fundamental aspect of radio telescope design. The resulting $FWHM$ of the ALMA 12-m antennas is measured to be $FWHM = 1.13 \frac{\lambda}{D} = 58''$ at 100 GHz, and it varies inversely with frequency. For example, Figure 7.15 shows a Gaussian profile that approximates the shape of the primary beam at 112 GHz with $FWHM = 52''$. As described in Chapter 3, the primary beam attenuation can be corrected as the last step of imaging. In addition to correcting the sky distribution of the signal to its correct value, however, this correction also increases the noise with angular distance from the primary beam center. For example, at the radius of the $FWHM$, the noise will be increased by a factor of two compared to that at the primary beam center. Moreover, if a map size larger than approximately the $FWHM$ of the primary beam is required, then it is necessary to observe several adjacent pointings, i.e., a mosaic, to recover the sky emission.

The default mosaic pointing pattern used by ALMA is a fully sampled hexagonal grid with equilateral triangles whose vertices are separated by $\theta_{hex} = \frac{\lambda}{D\sqrt{3}} = 0.511 \times FWHM$. With this spacing, the mosaic will




Figure 7.14: Image of H α emission from M51 used as a model of emission at 100 GHz (*left*) and the expected visibility amplitudes (logarithmic scale) plotted versus uv distance (*right*). Note that M51 cannot be observed from ALMA.

sample the emission at the Nyquist spatial frequency. Note that a hexagonal mosaic has spacing θ_{hex} along a row (e.g., in right ascension) and $\frac{\sqrt{3}}{2}\theta_{hex}$ between rows (e.g., in declination). Figure 7.16 gives an example of such a mosaic. To estimate the number of pointings (N_p) necessary to cover an area of $L_X \times L_Y$ using this hexagonal pattern, the following expressions can be used:

$$\begin{aligned} N_X &= (int) \left(\frac{L_X}{0.511FWHM} + 1. \right) \\ N_Y &= (int) \left(\frac{2 \times L_Y}{0.511FWHM\sqrt{3}} + 1. \right) \\ N_p &\approx (2N_X - 1.) \times \frac{N_Y}{2}. \end{aligned}$$

For an observation at 100 GHz with $L_X = L_Y = 4$ arcmin, 85 pointings are defined, similar to that obtained by the above formulae. In the case of an odd number of rows, N_p may differ slightly from the value returned by the OT. Note that mosaicing with the 7-m Array is more efficient because with the smaller diameter of its antennas, the FWHMs are commensurately larger ($FWHM = 1.13 \frac{\lambda}{D}$ also holds for the 7-m antennas).

Like a single pointing, a mosaic has an analogous “mosaic primary beam response pattern” that is the combination of the individual primary beams of the different pointings. Near the mosaic center, a Nyquist sampled hexagonal mosaic pattern has a sensitivity about 1.58 times that of a single pointing, with the sensitivity decreasing with the fall off of the mosaic primary beam response pattern. Another frequently used mosaic pattern has pointings separated by $FWHM/\sqrt{2}$ (see for example the NVSS survey); this pattern covers an area more efficiently but with little gain in sensitivity over a single pointing. Generally, the Nyquist sampled hexagonal pattern (the default for ALMA) is a good choice for smaller mosaics and the sparser hexagonal pattern is most often used for larger mosaics. The ALMA Observing Tool (OT) can be used easily to set up a mosaic of adjacent pointings with a user-defined spacing, though it is not recommended to exceed spacings greater than the constant noise pattern ($FWHM/\sqrt{2}$) if a well-sampled mosaic image is desired.

7.8 Multi-array and Multi-configuration Observations

As shown in Table 7.1, different 12-m Array configurations provide different angular resolutions, θ_{res} , which are a function of the longest baselines of the configuration. Similarly, the maximum recoverable scales, θ_{MRS} , of each configuration depend on the shortest baselines present. To achieve the requirements entered by the PI in the OT, θ_{res} and θ_{LAS} , multiple configurations may be needed. Note that θ_{LAS} is a property of the science target, while θ_{MRS} is a property of an array configuration. In particular, the most compact configuration for a project must ensure that θ_{MRS} is larger than or equal to the θ_{LAS} of the science target.

Based on the user-entered values of θ_{res} and θ_{LAS} for a given Science Goal, the OT will automatically attempt




Figure 7.15: Approximation of the primary beam profile of an ALMA 12-m antenna at 112 GHz with FWHM = 52".




Figure 7.16: An example of mosaicing with a field of 2 arcmin extent at 100 GHz using a hexagonal pattern with Nyquist sampling (white crosses).

to choose configurations that will produce a final image with the requested scientific goals⁵. In order to provide good image quality, there are some restrictions on which arrays and configurations can be combined (based mostly on the uv -coverage each provides). For example, at most two 12-m Array configurations are allowed and the most extended configurations (C-10) cannot be combined with any more compact configurations. If such restrictions result in the requested θ_{LAS} not being achievable, the OT will return a validation error and θ_{LAS} must be modified. These allowable configuration pairings in the OT represent a set of designed, validated, and supported pairings for imaging, but a PI may also request any set of array configurations as separate Science Goals (for example to provide complementary shorter spacings for C-10 observations). With such an ad hoc setup, simulations of the observations are highly recommended to determine how imaging would be affected.

Since the middle of Cycle 4, a “stand-alone ACA” option has been offered for observations that use the 7-m Array, with the TP Array added if required, but without 12-m Array observations. It is currently not possible to request observations that only require the TP Array except as complements to existing 7-m Array data that suggest that it is scientifically needed.

In detail, the array configuration combinations are selected as follows:

1. An interferometric array configuration (12-m or 7-m) is selected that comes closest to achieving the entered value of θ_{res} .
2. If the θ_{MRS} of this array configuration is less than the requested θ_{LAS} , the largest structures in the source will not be well imaged. A more compact 12-m Array configuration will be added, if allowed by the set of combinations listed in Table 7.4.
3. If the θ_{MRS} of the most compact configuration included thus far is still less than the requested θ_{LAS} , a 7-m Array component will be added.
4. If θ_{MRS} is still $< \theta_{LAS}$, a TP Array component will be added if possible. Single dish continuum, Band 9, and Band 10 observations are not currently supported so there is a fundamental limit to how large a source can be reliably imaged in these cases.

This procedure is carried out automatically by the OT and is entirely dependent on the values of θ_{res} and θ_{LAS} that are entered by the user. A detailed list of the allowed combinations is given in Table 7.4. Figure 28 of the Cycle 8 2021 ALMA Primer gives visual representation of required array combinations as a function of θ_{res} and θ_{LAS} .

Note: if the PI specifies a range of θ_{res} there may be multiple 12-m array combinations which could achieve the Science Goal. For instance, C-5 or C-4 could serve for the extended 12-m configuration. This would imply C-2 or C-1, respectively, for the compact 12-m configurations; but C-5 plus C-1 is not an allowed combination. In this situation, the initial observations will narrow down the range of possibilities such that subsequent data will be taken in the most “compatible” array combinations. In this example, if C-5 data are collected first, then every effort will be made to collect *all* of the extended array data in C-5, and subsequent, compact configuration observations will be taken in a C-2-like configuration. *For purposes of total time calculations in Phase 1, the most expensive combination will be assumed.*

Table 7.4 also shows the relative integration times that will be used by each array/configuration. While the overall integration time is set by the specific requirements of the Science Goal a PI defines, the *relative* integration times between its component configurations/arrays are determined by general considerations as discussed in Mason & Brogan (2013). The relative integration times do not depend on the details of the Science Goal beyond the fact that the Science Goal definition determines which combination of configurations or arrays is needed. Broadly speaking, because of their smaller collecting areas, the TP and 7-m Arrays require more on-source integration time than the 12-m Array. Furthermore, it is evident that the higher surface brightness sensitivity of the more compact 12-m Array configurations requires longer 7-m Array integrations. These time ratios were calculated by simulating matched image sensitivities when the more extended of two arrays or configurations is uv tapered to give the same angular resolution as the more compact configuration. This approach, used since Cycle 6, is an incremental improvement to the approach used in previous Cycles and which was described in Mason & Brogan (2013). The actual time ratios adopted for ALMA observations, while guided by these calculations, are subject to further operational constraints. In particular, the ideal time ratios are

⁵The entire suite of potential multi-array/configuration observations that satisfies the Science Goal requirements is the *Group Obs Unit Set* or GOUS. The individual array, individual configuration components it comprises are the *Member Obs Unit Sets* or MOUSes.

θ_{res} (arcsec)	θ_{LAS} (arcsec)	Array combination	Time ratios	Total Time
0.042	< 0.496	C-10	1	$1.0 \times \Delta_{extended}$
0.042	> 0.496	-	-	-
0.057	< 0.814	C-9	1	$1.0 \times \Delta_{extended}$
0.057	0.814-4.11	C-9 + C-6	1 : 0.21	$1.21 \times \Delta_{extended}$
0.057	> 4.11	-	-	-
0.096	< 1.42	C-8	1	$1.0 \times \Delta_{extended}$
0.096	1.42-6.7	C-8 + C-5	1 : 0.22	$1.22 \times \Delta_{extended}$
0.096	> 6.7	-	-	-
0.211	< 2.58	C-7	1	$1.0 \times \Delta_{extended}$
0.211	2.58-11.2	C-7 + C-4	1 : 0.23	$1.23 \times \Delta_{extended}$
0.211	> 11.2	-	-	-
0.306	< 4.11	C-6	1	$1.0 \times \Delta_{extended}$
0.306	4.11-16.2	C-6 + C-3	1 : 0.25	$1.25 \times \Delta_{extended}$
0.306	16.2-66.7	C-6 + C-3 + 7-m	1 : 0.25 : 0.6	$1.8 \times \Delta_{extended}$
0.306	> 66.7	C-6 + C-3 + 7-m + TP	1 : 0.25 : 0.6 : 1.0	$2.3 \times \Delta_{extended}$
0.545	< 6.7	C-5	1	$1.0 \times \Delta_{extended}$
0.545	6.7-22.6	C-5 + C-2	1 : 0.26	$1.26 \times \Delta_{extended}$
0.545	22.6-66.7	C-5 + C-2 + 7-m	1 : 0.26 : 1.21	$2.5 \times \Delta_{extended}$
0.545	> 66.7	C-5 + C-2 + 7-m + TP	1 : 0.26 : 1.21 : 2.1	$3.3 \times \Delta_{extended}$
0.918	< 11.2	C-4	1	$1.0 \times \Delta_{extended}$
0.918	11.2-28.5	C-4 + C-1	1 : 0.34	$1.3 \times \Delta_{extended}$
0.918	28.5-66.7	C-4 + C-1 + 7-m	1 : 0.34 : 2.4	$3.7 \times \Delta_{extended}$
0.918	> 66.7	C-4 + C-1 + 7-m + TP	1 : 0.34 : 2.4 : 4.0	$5.3 \times \Delta_{extended}$
1.42	< 16.2	C-3	1	$1.0 \times \Delta_{extended}$
1.42	16.2-66.7	C-3 + 7-m	1 : 2.4	$3.4 \times \Delta_{extended}$
1.42	> 66.7	C-3 + 7-m + TP	1 : 2.4 : 4.1	$5.1 \times \Delta_{extended}$
2.3	< 22.6	C-2	1	$1.0 \times \Delta_{extended}$
2.3	22.6-66.7	C-2 + 7-m	1 : 4.7	$5.7 \times \Delta_{extended}$
2.3	> 66.7	C-2 + 7-m + TP	1 : 4.7 : 7.9	$8.9 \times \Delta_{extended}$
3.38	< 28.5	C-1	1	$1.0 \times \Delta_{extended}$
3.38	28.5-66.7	C-1 + 7-m	1 : 7	$8.0 \times \Delta_{extended}$
3.38	> 66.7	C-1 + 7-m + TP	1 : 7 : 11.9	$12.9 \times \Delta_{extended}$
12.5	< 66.7	7-m	1	$1.0 \times \Delta_{extended}$
12.5	> 66.7	7-m + TP	1 : 1.7	$2.7 \times \Delta_{extended}$

Table 7.4: Array/Configuration combinations with the corresponding $\{\theta_{res}, \theta_{LAS}\}$ conditions for an observation at 100 GHz. All time ratios are with respect to the time spent in the *most extended configuration*. Thus, 1:2:3 means 1h in the most extended array; 2h in the intermediate array; and 3h in the most compact array. Similarly the *most extended* configuration determines the angular resolution. The actual resolution obtained with combined configurations can be 50% lower due to different weighting (see text). Note that for the full array combination, the total time is not equal to the sum of the individual times because TP and 7-m Array observations are run in parallel.

often impractically large for the 7-m and TP Arrays being used with C-1 or C-2. In previous Cycles, the time ratio was “capped” at a maximum value— typically 5 for adjacent arrays or configurations in a given imaging combination (with imaging combinations defined by Table 7.4). Starting in Cycle 6, a less restrictive cap of 7 for adjacent arrays or configurations was adopted, but in practice this choice only affects the 7-m Array time ratio with respect to C-1 and C-2. Furthermore, *all* 7-m:12-m Array time ratios are rescaled by the factor by which the 7-m:C-1 time ratio was reduced from its ideal value. This approach has the effect of more uniformly applying the operational constraint across different arrays and configurations with a modest and consistent impact on image characteristics.

Several additional considerations can impact the relative observing times which the OT sets up for different

arrays in practice. First, the OT enforces a minimum on-source integration time in order to ensure that the telescope is used efficiently. The minimum allowed on-source time is taken to be the greater of 5 minutes and half of the integration time spent on the bandpass calibrator. Second, the 7-m Array needs integrations of at least 1 hour to provide sufficient *uv*-coverage for good image quality⁶. Snapshot observations with the 7-m Array are therefore strongly discouraged (see also Section 7.5). Third, for total power observations it is necessary to observe a slightly larger area than the region of interest specified in the Science Goal in order to obtain sufficiently high-quality information over the entire region. Finally, when a custom (12-m Array) mosaic is specified by the PI, the OT attempts to identify the appropriate corresponding 7-m and total power map coverages which are needed. The actual implementation of this heuristic will, however, give rise to some variation in the relative observing times, depending on the geometry of the PI-specified custom mosaic. In addition, when selecting "time simultaneous" mosaic observations in the OT for the 12-m and 7-m Arrays, the OT will not select the same number of pointings for both arrays but select the optimal number of pointings the in 7-m array that cover the specified 12-m array area.

7.9 Multi-array and Multi-configuration Imaging

If data have been collected in multiple arrays or configurations as described in Section 7.8, the data from each will be processed separately and delivered if they pass Quality Assurance (see Section 11). At present, PIs are responsible for combining these products together, although they are encouraged to contact their ARC for assistance if desired. PIs may also refer to the ALMA M100 Science Verification CASA Guide⁷ for detailed guidance on data combination. It is important to note that when combining data from multiple arrays or configurations, the resultant non-Gaussianity of the dirty beam may cause a bias in flux measurement. This effect is discussed by Jorsater & van Moorsel (1995), who present an analytic correction and also note that it can be mitigated by deeper cleaning.

The recommended combination procedure comprises either one or two steps, depending on whether or not total power data are collected. First, the interferometric data will be imaged together in a single, "joint" deconvolution. This step can be done in CASA by passing all interferometric measurement sets— 7-m and one or more 12-m Array configurations— directly to the TCLEAN task. All 7-m and 12-m data delivered during Cycle 8 can be directly combined by the PI⁸.

If total power data are part of a project, then there is a second step to the combination procedure, where the deconvolved interferometric (7-m+12-m) image cube is combined with the total power image cube. The recommended procedure is to do this by "feathering" them together. Feathering is a commonly used technique in radio imaging for combining two images together by forming a weighted sum of their Fourier transforms. The procedure is as follows:

1. The total power and interferometer images are Fourier transformed.
2. The beam from the total power image is Fourier transformed ($FTSDB(u, v)$), to be used as a weighting function. Alternatively, one can specify some smaller portion of the total power antenna aperture, corresponding to a wider (single-dish) beam.
3. The Fourier transform of the interferometer image is multiplied by $(1 - FTSDB(u, v))$. This step down-weights the large spatial scale components of the interferometer map which are poorly measured, if measured at all.
4. The Fourier transform of the total power image is multiplied by the ratio of the volumes of the interferometer restoring beam to the single-dish beam, thereby putting the maps in the same units.

⁶For a mosaic, this needed minimum integration time is applicable to each pointing in the mosaic plus its immediately adjacent pointings. Since the OT enforces a time-on-source minimum at the science target (not per mosaic pointing) level, simulations are advised if large, high dynamic range 7-m Array mosaics are sought.

⁷https://casaguides.nrao.edu/index.php/M100_Band3

⁸Manually calibrated data deliveries from early ALMA Cycles that were calibrated using CASA versions earlier than 4.3 require a special procedure to put the 7-m and 12-m Array data weights on an equal basis before the data can be combined. The required steps are described in the "Data Weights and Combination" CASA guide <https://casaguides.nrao.edu/index.php/DataWeightsAndCombination>. All ALMA data calibrated by the ALMA Pipeline have correct weights, irrespective of Cycle number or CASA version.

5. The results from 3 and 4 are added and Fourier transformed back to the image plane.




Figure 7.17: Expected uv -coverage for a notional C43-1 (1 hour) and 7-m observations (5 hours)




Figure 7.18: Radial density of uv -coverage with uv -distance using notional C43-1 (1 hour) and 7-m observations (5 hours)

The resulting cube provides high angular resolution due to the interferometric data, but will have more accurate large-scale features and total flux densities because the single dish data are present as well. A thorough

explanation of feathering is given by Cotton (2017), and a comparison to other techniques for combining single dish and interferometer data is found in Stanimirovic (2002).

An alternative method of combining interferometric and TP data is possible using a tool called *tp2vis*, which was developed as a NA ALMA Development Study⁹. This method, which can be run in CASA, converts a total power image into appropriate *uv* data (pseudo-visibilitys) which are then concatenated together with the much higher resolution interferometric data. This entire *uv* data set can then be cleaned to obtain an image that contains all angular scales.

As an example, the *uv*-coverage and *uv* plane sampling distributions for a 1-hour observation with the notional C43-1 configuration and 5-hour observation with the 7-m Array are shown in Figure 7.17 and Figure 7.18, respectively. The Cycle 8 7-m Array configuration with ten 7-m antennas provides *uv* measurements in the range 8-32 m and overlaps the 12-m Array configurations in the range $R_{uv} < 15$ m.




Figure 7.19: Simulated images obtained using notional C43-3 (*top left*; 1 hour), notional C43-3 + C43-6 (*top right*; 1 + 3.3 hours) and notional C43-3 + C43-6 + 7-m (*bottom left*; 1 + 3.3 + 1.3 hours) Array combinations, and the image model itself (*bottom right*).

Using the H α emission of M51 as a hypothetical model (see Figure 7.14), Figure 7.19 shows the simu-

⁹<https://github.com/tp2vis/distribute>

lated images resulting from observing with a notional C43-compact configuration (C43-3) only; a compact plus extended 12-m configuration (C43-3 + C43-6); and finally of also adding the 7-m Array. Natural weighting was used and deconvolution was performed using the CLEAN algorithm (TCLEAN now recommended). The recovery of larger scales is quite noticeable with the inclusion of the 7-m Array data. Using only the C43-6 configuration at 100 GHz provides an angular resolution of $0.27''$. Combined with the C43-3 configuration, the angular resolution is lowered to $0.31''$. Finally, the angular resolution with the full combination (C43-3 + C43-6 + 7-m) is $0.32''$, very similar to that of only the two 12-m Array configuration data but 17% larger than the angular resolution obtained with the extended 12-m Array data alone. That difference is due to the respective uv -coverage and weight of the array configurations in the combined dataset.

These simulations illustrate two general features of multi-configuration/array observations. First, they are effective in retrieving a wider range of spatial scales than are accessible to a single configuration or Array. Second, the addition of more compact configurations or arrays does tend to broaden the synthesized beam (PSF) slightly. This broadening can be mitigated by changing the uv weighting, e.g., via the *robust* parameter, potentially at the expense of sensitivity and θ_{RMS} . These effects can be explored via the use of simulations with the notional configurations.

References

- [1] Andrews, S. et al. 2018, Ap.J. 869, L41
- [2] Briggs, D. S. 1995, *High Fidelity Deconvolution of Moderately Resolved Sources* Ph.D. thesis, New Mexico Institute of Mining and Technology
- [3] Briggs, D. S., Schwab, F.R., Sramek, R. A. 1999, *Synthesis Imaging in Radio Astronomy II*, ASP Conference Series, 180, 127
- [4] Condon, J. J., R. W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J. 1998, AJ, 115, 1693
- [5] Cotton, W. D. 2017, PASP, 129, 094501
- [6] Jorsater, S. & van Moorsel, G. 1995, AJ 110, 2037
- [7] Mason, B. & Brogan, C., 2013, Relative Integration Times for the ALMA Cycle 1 12-m, 7-m, and Total Power Arrays, ALMA memo 598
- [8] Moriarty-Schieven, G. ed., 2021, Observing with ALMA: A Primer (Cycle 8 2021), ALMA Doc. 8.1, ver. 2
- [9] Stanimirovic, S. 2002, in *Single-Dish Radio Astronomy: Techniques and Applications*, edited by S. Stanimirovic, D. Altschuler, P. Goldsmith, & C. Salter, vol. 278 of *Astronomical Society of the Pacific Conference Series*, 375. astro-ph/0205329
- [10] Keply, A. et al. 2020, *Publications of the Astronomical Society of the Pacific*, Volume 132, Issue 1008, pp. 024505

Chapter 8

Observing Modes

The ALMA Observing Modes are the set of capabilities that ALMA offers each cycle to its user community. Before offering a particular mode to the community, it is implemented in the ALMA online software package and verified using the required hardware. Here, the software implementation is described along with a description of particular observing modes.

An observing proposal submitted to the ALMA archive will have an associated structure called the *Observing Project* that will accompany it along the whole length of its lifecycle. This structure is defined in the ALMA Project Data Model (APDM), which specifies all the relevant components and their contents needed for successful completion of a project. A summary view of the constituents of an Observing Project is shown in Figure 8.1.

8.1 Observing Project Structure

The organization of each science project is subdivided into a well-defined structure with clear hierarchical levels. At the bottom of this structure are the *Scheduling Blocks* (hereafter, SBs). An SB is the minimum set of instructions describing an ALMA observation that can be fully calibrated. Projects are broken down into a set of these fundamental units for flexibility, given the properties of the ALMA site and the continuously-varying status of the Observatory as a whole (including the weather), and to encapsulate the scientific objectives of the proposal into individual entities. After a proposal has been accepted for observation, the SBs to be executed during the observations are produced automatically by the ALMA Observing Tool (hereafter OT) during the Phase 2¹. The SB may be edited, using the OT, in Phase 2 after the proposal has been accepted for observation and also during the project execution if necessary².

An SB contains a large amount of information about what should be observed including, but not limited to, the positions and velocities of the science targets, details of the correlator setup and the integration and cycle times of the different calibrations. However, the SB itself does not control the observations as it is just a single set of XML instructions. Instead, the ALMA online software reads the SB and executes the *observing script* appropriate to the type of observation required (e.g., single-dish, single field interferometry, etc.). The SB has relatively little influence over the order in which the various sources are observed and does not describe all of the calibrations that should be performed (a prime example being the measurements of the system temperatures).

Each SB typically consists of set-ups, calibrations, and target observations that can be observed within 1.5 hours. The end of an SB execution may be specified in terms of a maximum amount of time or when certain well-defined science goals have been reached as specified in the *Science Parameters* section of the SB. An SB cannot be stopped and re-started. Therefore, an SB either runs to completion, fails, or is terminated by the *Astronomer on Duty* (hereafter AoD). Given the limited duration of the SBs, it is often necessary to observe them several times to achieve the sensitivity required by the PI. The current maximum execution time for a particular SB has been chosen based on statistical measurements of the stability of the system. A system failure may prevent the online software from storing the data in the archive, resulting in the loss of the observed

¹For a description of Phase 1 and Phase 2 see the Cycle 8 2021 Proposer's Guide stage of the project lifecycle and associated documents

²SBs can also be created in an ad-hoc manner for commissioning purposes



Figure 8.1: Block diagram of an Observing Project from the point of view of the observation preparation (top) and internal hierarchical structure of the Scheduling Block in actual executions (bottom). All projects have the same ObsUnitSet (OUS) levels, the Science Goal OUS level, the Group OUS level and the Member OUS level. Scheduling blocks are attached to the Member OUS. Each time a Scheduling Block is executed, Control creates a new Execution Block structure (see text).

project data. Thus, executing SBs for more than two hours becomes unsafe from the operational perspective. To optimize this scheme, the concept of a *session* was implemented, which is the continuous execution of a single SB until a certain goal has been reached (see detailed description later on). Sessions have the added benefit that some of the calibrations can be shared among the different executions, which increases the overall observing efficiency. For Cycle 8, sessions are primarily used for polarization and VLBI observations.

While the SB is the smallest entity used for observing, the Observing Unit Set (hereafter OUS) is the smallest unit for data processing³. The SB/Member OUS/Group OUS (hereafter MOUS and GOUS, respectively) are the smallest structures that hold science observations that need to be observed/processed/combined together. These data hierarchies therefore maximize observing flexibility and ensure that data gets processed as early and as fast as possible. In the OT, observations are divided into different Science Goals (see the ALMA Proposer’s Guide). In order to follow this structure, each GOUS is attached to a Science Goal OUS.

Each execution of an SB can be calibrated separately (except for sessions, where Calibrations can be shared). All executions of the same SB can then be merged at the imaging stage to produce calibrated science products (e.g. images or data cubes). Usually, there will be one SB in each MOUS, but the latter may hold multiple executions of this SB. For polarization observations, all the SBs that belong to the same session should be grouped together for data processing. Unfortunately, this is not explicitly described in the SB.

If the calibrated science observations of a MOUS have to be combined with science observations of another MOUS, then these are grouped together into one GOUS. Otherwise, they will be placed into different GOUSs. A typical example of MOUSs that belong to the same GOUS are observations of the 12-m Array, the 7-m Array and the TP array. According to the method above, these observations would be in three separate MOUSs, but in the same GOUS.

Pipeline processing happens at the MOUS level, as soon as all observations of a MOUS are complete. It is expected that in future Cycles, processing will also happen at the GOUS level, in the case that the GOUS contains several MOUSs which have already been processed. The only other event that triggers data reduction is the end of an observing Cycle, when all the MOUSs and, potentially, GOUSs are reduced, irrespective of their degree of completion. As polarization and other selected modes of observations do not currently fit this scheme, the data reduction may be done manually by the ALMA staff.

8.2 Execution Structure

Once a given SB has been selected for execution (by the Scheduler subsystem or by the AoD), it is read into the online Control software, and specifically by the Science Software Requirements subsystem (hereafter SSR). The SSR subsystem commands the lower level Control software to create an Execution Block (EB) structure that is attached to the SB. As an SB will be executed as many times as needed to fulfil the requirements, several EBs may exist for a given SB. Each EB contains a record of the parameters and conditions under which the SB was executed along with references to the acquired data. The internal hierarchical structure of the EB is also shown in Figure 8.1. The SSR subsystem constructs a sequential series of scans that Control executes for each of the required calibrations. Each scan execution is in fact carried out by breaking it down into a series of subscans, each of which is itself broken into a series of integrations⁴ (the correlator software only admits sequences of subscans). Although commands are issued at the scan/subscan level, the correlator output corresponds to a particular integration. In general, each calibration observation consists of a scan containing several subscans (e.g. the 5 data points of a pointing calibration constitutes a scan with five subscans (see Chapter 10). Similarly, the integration time on a single science source between phase calibrations consists of one scan comprised of a number of subscans. To optimize the execution, scans may be organized in scan sequences which are passed to Control for execution. Scans and scan sequences can be of arbitrary length, depending on the characteristics of a

³There are three types of OUSs, the Science Goal OUS (no processing happens at this level), the Member OUS and the Group OUS. Processing occurs usually at the Member OUS level, but the possibility to also process at the Group OUS level may be enabled in the future. A Member OUS is a collection of data (currently all data collected out of executions of a single SB) needed to achieve (part of) a Science Goal, and is usually associated with a specific array configuration. A Group OUS contains all Member OUSs needed to achieve the science goals of a complete science proposal, and can include Member OUSs corresponding to different ALMA configurations and/or arrays.

⁴A scan is a period of data acquisition that achieves a particular OBSERVER INTENT (e.g., TSYS measurement, pointing determination, science target observations, etc). A subscan is a distinct period of data acquisition collected in support of a scan (e.g., one pointing offset measurement in a pointing scan; one mosaic pointing of a mosaic scan, etc). Finally, an integration is the shortest period of data recorded in the data sets.

given observation but subscans are recommended to be 30 seconds or less in most cases. Integrations tend to be on the order of 1 to 10 seconds where the final value has to be an integer multiple of the correlator dump time. These values can be specified during the generation of SBs in Phase 2 with the OT; specifically, the subscan duration is specified in the SB target parameters section and the length of the correlator integration time and dump time is specified in the SB Spectral Setup. Calibration results from the Telescope Calibration subsystem (TelCal) and QuickLook (QL) pipeline are usually attached to a scan. (e.g. antenna pointing results). The SSR and the Control subsystem are responsible for the creation of all the metadata needed for data processing.

8.3 The Observing Process

For each of the SBs in a given project, a set of targets is specified (for more details, see the Cycle 8 2021 OT User Manual). These targets can correspond to either calibration or science executions. Targets are organized in observing groups where the first group (Group 1) is always the initial calibration group, with subsequent groups detailing the science observations⁵. All Science targets and relevant calibrators within a group are observed before the next group is started. An SB can have multiple groups.

All groups other than Group 1 are considered complete when all Science targets in the group have been observed for the requested time or have set below the elevation limit. After all groups are completed or the SB execution time limit is reached, the primary Phase calibrator for the group, which triggered the SB execution time limit as well as any deferred calibrators from Group 1, is observed.

Most observing modes, including single field interferometry, grouped source executions, pointed mosaics, and polarization use the same ALMA observing script, called the *standard interferometry* script. The set of necessary calibration measurements (e.g. flux, bandpass, etc) usually specified in Group 1 are performed at the beginning of the observing sequence. Unless previously specified by the user or ALMA staff, the sources are selected at run-time by the SSR query algorithm using the parameters defined in the SB as input. If sources of sufficient quality for calibration are found, the SB will execute.

Pointing, Atmospheric, and Sideband-Ratio calibrations are associated to the main (Bandpass, Amplitude, Polarization, and Phase) calibrations on an as-needed basis which is determined by the SSR at run-time. For example, a pointing calibration is usually done before the amplitude and bandpass calibrators are observed, and again before the main observations of the science target and phase calibrator cycle. Within a group, the Science targets are each observed in turn until observation of the primary phase calibrator (the calibrator with shortest cycle time in the group) is required. A typical cycle time for the phase calibrator may be 7–10 minutes for standard observations, and 1–2 minutes for long-baseline observations (see Chapter 10). This process is repeated until the observing requirements are met, or the SB reaches its maximum execution time limit. Any additional (*secondary*) phase calibrators are observed as specified in the SB. For a description of what each calibration entails, see Chapter 10.

The user has several options to select optimal calibrators:

- *System-defined Calibration* in the OT: let the OT set up default queries to the ALMA Calibrator Source Catalogue which will be used to select appropriate calibrators at run-time. This is the recommended mode for regular projects.
- *System-defined Calibration (force separate amplitude calibration)* in the OT: similar to the above System Defined Calibration except SSR will enforce a solar system object as the primary flux calibrator and not a grid source. This can provide additional accuracy for some projects.
- *User-defined Calibration* in the OT: manually enter specific calibration sources or set up the queries using alternative values for the parameters, but this carries some risk (for example, calibrators will not be observed during the execution of a group if they are not visible at the time of the observation), and thus must be fully justified in the Technical Justification of the proposal.

8.3.1 The Source Selection Algorithm

The ALMA system calibration for a given SB is typically done using astronomical sources selected at run-time. Exceptions to this process are long baseline or high frequency observations or SBs with user defined calibrations

⁵In some cases (e.g. polarization), additional calibrations may be added into subsequent groups and they will be interleaved according to their cycle times.

where calibrators are pre-selected and hard-coded into the SBs. To optimally select appropriate calibration sources, selection criteria are implemented in the SSR using the following steps:

- A list of sources is retrieved from the ALMA Calibrator Source Catalogue based on the query center coordinates and the search radius defined in the SB.
- For each source in the list, a search for flux measurements at the SB representative frequency is performed. If no flux measurements are found, the flux is extrapolated from other measurements using a spectral index of -0.7 to the SB representative frequency.
- The flux values are weighted based on the proximity to the representative frequency, science target and time since the last measurement present in the source catalog.
- An expected SNR is calculated from the flux and SB observing setup.
- The returned list of sources that passed the signal to noise criteria are ranked based on SNR and separation with respect to the query center and flux error.
- The final list is sorted based on ranking.

Using these rules, the criteria for all the ALMA calibrations were implemented adding specific requirements based on the nature of the target, such as:

- **Bandpass calibrator:** The bandpass calibrator integration time defined in the SB (typically 5-15 minutes, depending on e.g. band and spectral setup) and a number of dynamically determined parameters (e.g. the number of antennas in the array and the estimated system temperature) are used to determine a flux density threshold for the catalog search that would correspond to a minimum SNR per antenna of 50. If no suitable sources are found, the SNR limit is lowered. Other factors that are taken into account are elevation and shadowing. The search radius is typically 45 degrees from the query center, but if no suitable calibrator is found the search radius is increased.
- **Phase calibrator:** An SNR of 15 is used in the flux estimation. The required integration time on the phase calibrator is derived either using the widest spectral window (spw) bandwidth (usually 2 GHz), or the aggregate bandwidth when only narrower spectral set-ups are specified. Unless a value is set up in the SB, a search radius of 15 degrees from the query center is used as default (smaller cycle times and radii are used for the long baseline and high frequency observations, as described in Chapter 10).
- **Flux calibrator:** A solar system object is searched in the first instance for flux calibration purposes if this option was selected in the OT. Otherwise, a search is performed among grid sources and the first observable grid source is picked up. The solar system objects and grid sources used for flux calibration are discussed in Chapter 10.

To help checking the consistency of calibration results, the SSR has a logic to select at least two different sources for the sets of calibrators listed above.

8.4 Single Field Interferometry

Single field interferometry is the most basic form of observation that ALMA supports. It consists of standard calibration scans⁶ associated with the constituent calibration targets, and science observations of a single field (primary beam, see Chapter 3 for details). A typical observation will start with a bandpass calibration. The bandpass observation is executed to measure the spectral response of the system, and thus should be done on a bright source with simple spectral properties, such as a bright quasar with no emission or absorption lines and a reasonably flat spectrum. The flux scale calibration will be performed next, which is intended to obtain the observed flux of a well known source, such as a solar system object. The observed flux will be used to compare with the established flux model of this object to obtain the scaling factor to be applied to all other sources in the SB. Ideally, the flux scale calibrator sources should be small in angular size with respect to the synthesized

⁶The reader should refer to Chapter 10 for further details on all the different system calibrations.

beam so as not to resolve the source structure and add uncertainties in the flux calibration. In practice, many solar system objects may be moderately resolved. In such cases, only a subset of the baselines may be used to estimate the flux based on accurate models of the objects available in the analysis software. From this result, a scaling factor is derived which can then be applied to the data at the data reduction stage. If no solar system object is observable, the SSR will pick the first available grid source.

The phase information of the sources must be preserved and discerned from phase variations at the individual antenna elements. This is achieved with the phase calibration during the course of an observation. Since the phase is expected to change much more rapidly in time than the amplitude, phase calibrators will be observed more frequently than other calibrators. Because the phase varies on small scales on the sky, the calibrator must be as close as possible to the science target. The phase calibration observations are taken right before and after each observation of the science target, and the phase correction will be interpolated in time when applied to the science target. As the atmosphere fluctuates rapidly (especially at higher frequencies), radiometric observations of atmospheric water lines are also done to correct for these additional phase variations. These corrections can be applied online and offline (see Section 10.4.3 for further details).

Multi-source observations are observations of astronomical objects within a region of about 5° . They can be observed with one EB since the same phase calibrator and check source (see Chapter 10 for the description of a check source) can be used for all targets. Other aspects of the observations (flux calibration, bandpass calibration, phase-calibration, check source use) remain the same as for one target only that there are several targets to image in the final data product.

8.5 Pointed Mosaic Observations

Pointed mosaic observations enable a single SG to cover a field of view larger than the primary beam by making observations of multiple single fields that overlap in spatial coverage by an amount specified in the SB. Up to 150 pointings are possible in a single SB. This limitation is set by the maximum execution time for a single SB and the necessity to finish all fields at least once within an execution. In pointed mosaic observations, each of the fields will be assigned a different field ID in the data, but the same source ID. Thus, all fields will share bandpass, amplitude and phase calibrations which are done as single fields. For larger fields that cannot be covered by 150 pointings, multiple SGs may need to be defined. The mosaic is arranged as a single scan composed by a number of subscans corresponding to the individual rows in the mosaic (see Section 3.5). Both the science target's specific last mosaic pointing position and index within the science target list are handled by the SSR to ensure the next scan begins on the proper science target and proper offset position. The mosaic observations can be set up by specifying *1 Rectangular Field* in the OT.

8.6 Total Power (Single Dish) Observations

The purpose of adding data from single-dish observations using auto-correlations is to recover large scale emission from the science target that may have been spatially filtered-out by even the shortest baselines of the 7-m Array. For this reason, these observations are referred to as *zero spacing*, however, for convenience, these observations are also sometimes referred to as *total power* (hereafter TP), although in practice they are taken in auto-correlation mode rather than using a total power (square law) detector. Four 12-m antennas connected to the ACA Correlator are available for this purpose. For Cycle 8, only spectral-line observations are offered in this mode.

As with the *standard interferometry* script, (see Section 8.3), the *standard single dish* script uses the SSR capabilities to perform some of the calibration observations (e.g. pointing) interferometrically. The SB of a TP observation consists of a group of calibrations followed by an On-The-Fly (OTF) observation of a rectangular area on the science target for line mapping with periodic offsets to a certain reference position observed for calibration purposes. For the OTF and reference position integrations, only the auto-correlation data are written to the ASDM in order to minimize data rate and size, whereas the rest of calibrations use cross-correlation information for analysis.

Most of the calibration scans are executed in Group 2. The execution of Group 2 starts up with a pointing calibration scan. Then, the OTF mapping of science targets are made with periodic executions of atmospheric calibration scans inserted. The atmospheric calibration is made at the reference position of the science target.

The OTF map is observed as a series of raster rows, scanning in the coordinate system specified in the SB. In Cycle 8, scans are taken either in longitudinal or latitudinal directions as specified in the OT⁷. In future Cycles, scans may be taken in the two perpendicular directions in turn, to minimize scanning artifacts. The reference positions, assumed to be positions which are free of spectral-line emission, are specified either in absolute coordinates or as offsets from the map center. The recommendation for extragalactic objects that are above a few degrees in Galactic latitude, is to use horizon coordinates. However, in cases where high contamination is expected, e.g. Galactic CO, fixed offsets are recommended. A catalog of emission-free off positions is currently under development. By default, the OTF map will cover an area half a beamwidth larger than the interferometric observations on all sides of the map. This will ensure that undersampling at the map edges does not affect the data combination process with the interferometric data. A raster row in these observations is defined as a subscan, with a maximum length per scan (consisting of some number of subscans) of 600 seconds. The reference position is observed as specified by its cycle time during the science target scans. Pointing is calibrated on a bright calibrator near the Science target with a frequency indicated in the SB, and atmospheric calibrations are taken every 10 minutes at the reference position to measure the system temperature.

The calibrated TP map of the science target will be in units of Kelvin, on the antenna temperature (T_a^*) scale. Since all TP observations are expected to be combined with 12-m and 7-m Array data which come in calibrated units of Jy/beam, the TP data must also be converted into these units. This conversion from Kelvin to Jy/beam requires knowledge of the main beam efficiency η_{mb} and the beam size θ . These values are automatically derived and monitored by the observatory, using fits to the aperture efficiency of each antenna as a function of time including, but not limited to, elevation, outside temperature and surface accuracy. TP observations obtained prior to Cycle 6 instead had “amplitude calibration” maps measured separately for each project by obtaining a continuum map of a bright quasar or a planet with known flux.

8.7 Polarization

The ALMA antennas have receivers with linearly polarized feeds followed by a waveguide and a polarization splitter.⁸ In this way, the incoming radiation is separated into two orthogonal components (X and Y) which are down-converted and digitized independently. For each baseline, the digital signals are cross-correlated at the correlator where the outputs are the four cross-correlations visibilities XX, YY, XY, and YX (or V_{xx} , V_{yy} , V_{xy} , and V_{yx}). The relation between these four cross-correlations and the Stokes parameters, are ideally given by

$$\begin{aligned} V_{xx} &= I + Q \\ V_{xy} &= U + iV \\ V_{yx} &= U - iV \\ V_{yy} &= I - Q \end{aligned}$$

where I , Q , U , and V are the Stokes parameters. In an ideal world, one would be able to combine the observed cross correlated visibilities and recover the Stokes parameters as,

$$\begin{aligned} I &= \frac{V_{xx} + V_{yy}}{2} \\ Q &= \frac{V_{xx} - V_{yy}}{2} \\ U &= \frac{V_{xy} + V_{yx}}{2} \\ V &= \frac{V_{xy} - V_{yx}}{2i} \end{aligned}$$

⁷The minimum time per row in a map is 3 sec, and maximum speed at which data can be taken is usually limited by the correlator.

⁸For the ALMA receiver bands offered in polarization mode for Cycle 8, Bands 3, 4, 5 and 6 the polarization splitter is an *Ortho-Mode-Transducer* (OMT), and for Band 7 it is a polarized grid. See Chapter 4 for details.

From here, it is easy to see that the total intensity, given by Stokes I , is only a function of the *parallel hands* XX and YY , the linear polarization is given by Stokes Q and U while circular polarization is given by Stokes V . However, there are a number of additional quantities, which have to be considered and calibrated, that prevent us from directly using the measured cross-correlation in this simple way. We briefly enumerate some of these quantities here.

1. The splitting of the incoming radiation into the two orthogonal components is not perfect and small projections of one component into the other are produced. This is called the *instrumental polarization* or *D-terms*. The instrumental polarization is an antenna based quantity which also depends on the frequency and the band (cartridge design and external optics). Additionally, this quantity is measured in the frame of the antenna which is parameterized by the elevation and azimuth axes (Alt/Az) and thus, it rotates with respect to the frame of the sky. This rotation introduces an angular dependence into the visibilities which is given by the parallactic angle (ψ). By design, the instrumental polarization is small (a few percent), but not negligible.
2. The analog signal path followed by each of the individual X and Y polarizations is slightly different, which introduces a small delay in the signal that needs to be accounted for.
3. The offline calibration procedure arbitrarily sets the phase of an antenna to zero (the reference antenna). This results in the inverse of its X-Y offset being imprinted on the cross-hand phases of the other antennas. This phase offset between the cross-hands also needs to be determined.
4. Other effects, such as the position within the primary beam (off-axis polarization), are still under commissioning. Thus, only on-axis polarization is offered for Cycle 8 (see below for a detailed explanation.)

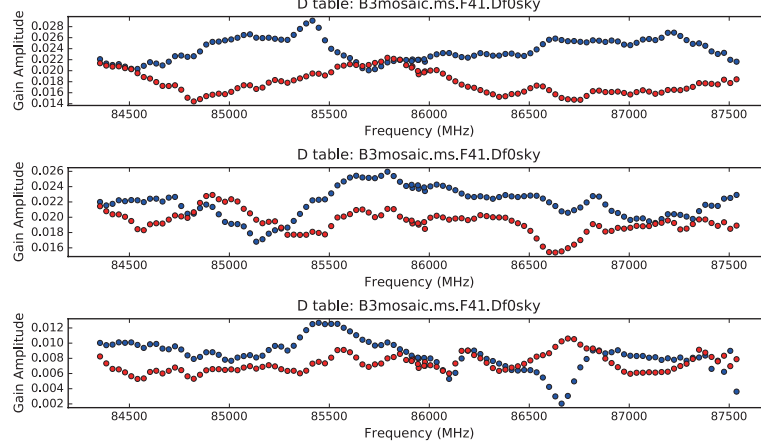


Figure 8.2: The D -term plots of antennae DA42 (top), DV03 (middle), and PM01 (bottom) for an observation in Band 3. The vertical axis is the fraction of the input signal voltage in one polarization that leaks into the output of the other polarization in voltage units and the horizontal axis is the frequency in MHz. The blue and red symbols represent D_X (the fraction of Y polarization signal that leaks into the X polarization) and D_Y (the fraction of X polarization signal that leaks into the Y polarization), respectively.

By taking into consideration the instrumental polarization and the parallactic angle dependence, the linearized equations 8.1–8.4 for the visibilities can be re-written as

$$\begin{aligned}
 V_{XX} &= (I + Q_\psi) + (U_\psi + iV)d_{X_j}^* + d_{X_i}(U_\psi - iV) + d_{X_i}(I - Q_\psi)d_{X_j}^* \\
 V_{XY} &= (I + Q_\psi)d_{Y_j}^* + (U_\psi + iV) + d_{X_i}(U_\psi - iV)d_{Y_j}^* + d_{X_i}(I - Q_\psi) \\
 V_{YX} &= d_{Y_i}(I + Q_\psi) + d_{Y_i}(U_\psi + iV)d_{X_j}^* + (U_\psi - iV) + (I - Q_\psi)d_{X_j}^* \\
 V_{YY} &= d_{Y_i}(I + Q_\psi)d_{Y_j}^* + d_{Y_i}(U_\psi + iV) + (U_\psi - iV)d_{Y_j}^* + (I - Q_\psi)
 \end{aligned}$$




Figure 8.3: Parallax angle (ψ) plot as a function of hour angle for the declination of 30, 10, -10, -30, and -50 degrees.

where d_{X_j} are the D -terms as a function of polarization and antenna, the asterisk denotes complex conjugates, and U_ψ and Q_ψ are the Stokes parameters as a function of the parallax angle. Figure 8.2 shows an example of the D -term spectra in Band 3 for selected antennas. The D -term level is typically a few percent at Bands 3, 4, 5, 6, and 7 on axis with some variations over frequency⁹. Without any D -term calibration, an unpolarized source may appear to be polarized at the 1% level. The most straightforward way to calibrate the D -terms is to observe an unpolarized source, where the cross-hands output will be purely D -terms since Stokes Q, U, and V will have no signal. However, bright and unpolarized astronomical sources are rarely found. Therefore, ALMA uses an unresolved quasar as a polarization calibration source. These sources are monitored to ensure they are bright enough ($> 2\%$), and the exact polarization is determined during data reduction provided that the source is observed over a wide enough range ($> 60^\circ$) of parallax angles to separate the effects of the source polarization, the D -terms and the cross-hands phase spectra and delay. Figure 8.3 shows that the rate of change of parallax angle is fastest near transit and slowest for low elevation sources. Thus, care should be exercised when choosing a polarization calibrator in order to maximize parallax angle rotation. The ALMA observatory will select appropriate calibrators for each project, but the users must take into account these limitations when planning the observation.

With the current calibration scheme, the expected minimum detectable degree of linear polarization, defined as three times the systematic calibration uncertainty, is 0.1% for compact sources (i.e., within the inner 1/3 of the primary beam FWHM) for both TDM and FDM observations. The minimum detectable degree of circular polarization is 1.8% of the peak flux for both TDM and FDM observations. This level of accuracy for circular polarization is currently on-axis, and thus the circular polarization imaging must be restricted within 1/10 of the primary beam FWHM to avoid beam squint effects. Note that the systematic calibration uncertainty can degrade by a factor of ≈ 2 depending on, for example, the choice of calibrator or parallax angle coverage.

The accuracy of absolute polarization position angle in a single-pointing image is 1° within the inner 1/3 of the FWHM of the primary beam, increasing to as much as 5° near the FWHM. These levels of accuracy will only be reached if there is sufficient signal-to-noise in the polarized emission.

A D -term component also arises from the off-axis geometry of feed horn, antenna illumination, and the alignment of optics. This D -term component will vary across the primary beam pattern (see Chapter 4). Generally, the D -term level becomes larger when increasing the offset from the beam center. This is the so-called *off-axis* instrumental polarization, which, for Cycle 8, will not be calibrated.

Spectral line polarimetry is also offered in Bands 3, 4, 5, 6, and 7 at arbitrary frequency setups. Figure 8.4 shows the D -term solutions of a representative antenna (DV22) as obtained during commissioning. Figure 8.5 shows a zoomed-in view in frequency of the D -term solutions shown in Figure 8.4. These figures show that the spectral shape of the D -term solutions is smooth and no spur-like structure is seen even in highest frequency

⁹A resonance has been documented for Band 6 between the feed and the OMT. This resonance introduces a small ripple in the D -term solution which increases the leakage amplitude by a small amount, but still only a few percent.




Figure 8.4: A frequency vs real part of the D -terms plot for an antenna (DV22) obtained towards J0522-364. Blue circles and red crosses are showing the D -term obtained using a setup of FDM with 2 GHz bandwidth and using a setup with higher spectral resolutions (FDM 500 and 62.5 MHz bandwidth respectively). Four panels show the D -terms obtained from the spw 0, 1, 3, 2 (from top left, clockwise).

resolution mode (30.5 kHz per channel at the right bottom panel of Figure 8.5). Figures 8.4 and 8.5 clearly demonstrate that the instrumental polarization is mostly spectral resolution independent for these datasets.

For Cycle 8, ALMA will offer linear polarization continuum mosaicking in Bands 3, 4, 5, 6, and 7. For this new mode, the spectral range will be fixed to the current default continuum frequency setups (see Table 8.1). Therefore, the P.I will not be able to freely choose the frequency tuning of a given spectral window. However, the PI will be able to choose between the TDM and FDM modes through the ALMA-OT interface using a restricted frequency resolution when FDM mode is chosen. This scheme has been implemented to allow for better continuum identification when mapping some sources of interest (e.g. high mass star forming regions). The limits in the spectral resolution that can be selected, for FDM polarization mosaicking, do not consider spectral averaging and are shown in Table 8.1. The maximum spectral resolution listed in the Table applies only to continuum mosaicking during cycle 8. Regular, on-axis, ALMA polarization observations are not subjected to this restriction. The allowed area to be mapped is subject to the same 150-point restriction that applies to total intensity (Stokes I) mosaics. The mosaic pattern can be arbitrary, but ALMA recommends a hexagonal grid. While Nyquist sampling (half a beam overlap) is sufficient, a sampling of less than Nyquist (i.e., a more “loosely packed” mosaic) must be justified. A proposal requesting a mosaic sampling rate of less than Nyquist might be rejected on technical grounds. The average error estimates for linear polarization mosaics are 1 degree in polarization position angle, and 0.1% in polarization fraction in the regions of the mosaic that correspond to the inner $\frac{1}{3}$ FWHM of a given pointing. Near the FWHM of a given pointing, the estimated upper limits for the errors increase to 5 degrees and 0.5%. These errors were determined from a grid of offset pointing observations toward highly linearly polarized point sources (e.g., 3C 279) that spanned the region of the primary beam out to the FWHM. Although increased sampling is allowed with the required justification, we did not find significant improvements in the errors.¹⁰

Also for Cycle 8, ACA stand-alone polarization observations will be offered. Single-field interferometric polarization observations with the ACA are available in Bands 3 to 7 with arbitrary frequency setups. The measurements are accurate for the central one third of the FWHM beam for linear polarization. However, circular polarization with the 7 meter antennas has not been explored. It should be important to note, that effects such as beam squint can produce spurious circular polarization as we move across the antenna beam (the $1/10$ FWHM in band 3 for a 7 meter antenna is $\sim 10''$ for an array resolution of $12''$ at 100 GHz). Multiple

¹⁰See Hull et al. 2020, PASP, submitted.




Figure 8.5: A close-up view in frequency space of the real part of D-term shown in Figure 8.4. Blue circles and red crosses are showing the D -terms obtained from setup #1 (FDM, 2 GHz bandwidth) and setup #2 (FDM, 500 and 62.5 MHz bandwidth respectively).

pointings are permitted, however mosaics are not supported at this time. The same calibration scheme employed for the main array will be used for the ACA as well as current imaging techniques. However, data combination in full Stokes between the ALMA main array and the ACA is not supported in CASA at this time. Because the ACA is always in a compact configuration, a restriction of 40° in elevation for the polarization calibrator is implemented in the dynamical scheduling system to avoid shadowing.

Band	spw 1 (GHz)	spw 2 (GHz)	LO 1 (GHz)	spw 3 (GHz)	spw 4 (GHz)	Maximum Spectral Resolution (MHz)	Velocity Resolution (km/s)
3	90.5	92.5	97.5	102.5	104.5	1.953	5.6
4	138.0	140.0	145.0	150.0	152.0	1.953	3.9
5	196.0	198.0	203.0	208.0	210.0	1.953	2.8
6	224.0	226.0	233.0	240.0	242.0	1.953	2.4
7	336.0	338.5	343.5	348.5	250.5	1.953	1.7

Table 8.1: Default frequency setup for polarization continuum observations including maximum spectral resolutions allowed for mosaicking.

8.7.1 Sessions

Although not directly related to polarization itself, sessions are the observing scheme designed to execute polarization projects with ALMA. In order to allow for the observation of large projects and to avoid the execution of long SBs (more than 2 hours), the concept of a session was implemented in the the online ALMA software. A session is defined as the continuous execution of the same SB until the scientific criteria are met. The session preserves the integrity of the observation by keeping the information from the previous SB execution in the run-time memory. Thus, the session will manage the calibrator’s cycle time in order to avoid unnecessary observations giving an additional level of optimization for ALMA. By default, calibrations such as the bandpass will be done every hour, thus saving observing time while the default cycle time for the polarization calibration is set to 35 minutes. The phase calibration-science target loop will be interrupted when an additional calibration is needed based on its cycle time (e.g. polarization)¹¹. This is done in such a way that the phase-science bracket is preserved. The session will remember the last time the polarization calibrator was observed and interleave the calibration when needed.

In general, the required parallactic angle coverage can be achieved with two to three executions of the SB, which results in about 3 to 4 hours of observation. The actual number of EBs is set by the ALMA dynamic scheduling system in order to cover the required parallactic angle coverage. Other cases in which the session scheme might be useful are large mosaics, surveys (multi-target), and large single dish raster maps, but those

¹¹The cycle time parameters are user controlled and can be explicitly specified in the OT target parameters section.

cases are not yet offered in session mode for Cycle 8.

8.8 Multiple Region Modes

In the frequency division modes (FDM) of the correlator (see Section 5.1.2), the final spectrum is synthesized using individual filters 62.5 MHz wide. When the total bandwidth is between 125 MHz and 1 GHz, it is possible to move these individual filter positions to create spws covering a number of disjointed spectral regions. This is called the Multiple Region Modes. The constraints are:

- The spw width must be a multiple of 62.5 MHz.
- The aggregate width of these spws must be equal to that of the original bandwidth selected.
- The spws must all fit within the 2 GHz baseband used.
- The other parameters (resolution, polarization and sensitivity options) must be the same for all spws.

This mode is useful when a number of line features which require high spectral resolution are spread across the IF bandwidth. Since the filters have a unit width of 62.5 MHz, if the user chooses a mode with 250 MHz total bandwidth, it is possible to place four separate windows, each with 62.5 MHz width, anywhere within the 2 GHz baseband. In Cycle 8, a maximum of four spws per baseband will be offered.

8.9 Observation of Ephemeris Objects

Observation of solar system objects (with the exception of the Sun) is supported as in previous Cycles. Several well known solar system objects including planets, satellites and asteroids can be selected from a pulldown menu in the Observing Tool. For other sources, including non-sidereal objects, an external ephemeris file can be supplied as an input. The JAO is not responsible for observing ephemeris objects using old or outdated ephemerides. It is the responsibility of the PI and the observing team of observations of ephemeris targets to provide a valid target ephemeris file during the Phase 2 process and any update during the Cycle if necessary. The ephemeris file must be in JPL Horizons format. A typical ephemeris file may consist of the date (time), Right Ascension, Declination, range and range rate, for example:

```
*****
Date__(UT)__HR:MN      R.A._(ICRF/J2000.0)_DEC      delta      deldot
*****
2012-Jun-26 13:00      06 22 57.33 +23 16 11.23 1.01653182506561 -0.2593303
2012-Jun-26 13:01      06 22 57.49 +23 16 11.14 1.01653170585632 -0.2583150
2012-Jun-26 13:02      06 22 57.66 +23 16 11.05 1.01653158664703 -0.2572937
2012-Jun-26 13:03      06 22 57.82 +23 16 10.96 1.01653146743774 -0.2562665
2012-Jun-26 13:04      06 22 57.98 +23 16 10.84 1.01653134822845 -0.2552333
2012-Jun-26 13:05      06 22 58.14 +23 16 10.74 1.01653122901917 -0.2541942
2012-Jun-26 13:06      06 22 58.30 +23 16 10.66 1.01653122901917 -0.2531491
```

More information on the format and precision needed is available in the ALMA Observing Tool documentation.

8.10 Solar Observations

Continuum mode solar observing was first made available for ALMA Cycle 4. It will continue to be supported, albeit with limitations, in Cycle 8. In particular, the following conditions apply to solar observing in Cycle 8.

- Only Band 3, Band 5, Band 6, and Band 7 continuum observations using the following default frequencies (GHz) are offered:

```
*****
Band      LO Freq.      LSB              USB
          BB1         BB2         BB3         BB4
*****
3          100         92-94         94-96         104-106        106-108
5          198         190-192       192-194       202-204        204-205
6          239         229-231       231-233       245-247        247-249
7          346.6       338.6-340.6   340.6-342.6   350.6-352.6    352.6-354.6
```

- Interferometric observations must be done using the TDM mode (see section 5.1.1); frequencies are fixed to 2 GHz-wide spws centered on the frequencies shown in the table above. The spectral-line FDM observing mode is not offered for solar observations (Section 5.1.2).
- Simultaneous observations in different frequency bands is not offered. Each execution block can only include one band.
- Observations will be available in each band and configuration as follows:

```
*****
Config    Band 3    Band 5    Band 6    Band 7
*****
C-1       X         X         X         X
C-2       X         X         X         X
C-3       X         X         X
C-4       X
```

- Because the WVR receivers are saturated when the dishes point at the Sun, online WVR phase correction will not be applied and offline WVR correction for on-source (solar) data is not possible (see Section 8.10.2).
- Interferometric observations will use a combined array comprising both 12-m and 7-m antennas and will be processed with the 64-element correlator (Section 5.1).
- To minimize shadowing of 7-m antennas, observations will be carried out when the Sun is above an elevation of 40°.
- Mosaicing of larger FOVs can be carried out with up to 150 different pointing offsets relative to the phase center specified in the ephemeris.
- As for the rest of ALMA observations, a single observation (an Execution Block) cannot exceed 2 hours, which will include the time overheads for bandpass and flux calibration. These calibration overheads amount to about 25 mins.
- The minimum integration time of interferometric observations is fixed to 1 second.
- Observations may be performed using dual linear polarization (XX, YY) or single polarization (XX) correlations; full polarization measurements are not currently offered for solar observations.
- Total-Power full-Sun fast-scanning single-dish observations are offered as context and to recover the largest angular scales for interferometric observations; proposals requesting only Total-Power single-dish observations will not be accepted in Cycle 8.
- Single-dish Total-Power mapping will be “monomode” (double-circle scan pattern) with a 2400 arcsec-diameter circular FOV centered on the solar disk (see below).
- Single-dish full-Sun will not be executed when the Sun is at elevations above 70° because the required fast-scan azimuth slew speeds are too high.
- The time cadence of full-sun images obtained from total power observations is fixed to about 10 minutes for Band 3, 12 minutes in Band 5, 14 minutes for Band 6, and 25 min for Band 7.

8.10.1 Solar Observing Modes

Solar observations with ALMA are possible because the surface of the antennas is designed to scatter the optical and IR radiation to an extent that the subreflector and other elements in the optical path are not damaged or degraded. However, additional steps must be taken to allow useful observations of the Sun to be made. ALMA receivers are designed for a maximum RF signal corresponding to an effective brightness of about 800 K at the receiver input. Since the quiet Sun has a temperature of $\sim 5000\text{--}7000$ K at ALMA frequencies, the solar signal must be attenuated or the receiver gain must be reduced to ensure that receivers remain linear, or nearly so.

The initial solution adopted by ALMA was the use of a “solar filter” (SF) that is mounted on the Amplitude Calibration Device (ACD) of each antenna (Section A.5). When placed in the optical path, the solar filter is required to attenuate the signal by $4+2\lambda_{mm}$ dB with a return loss of -25 dB (-20 dB for $\nu > 400$ GHz) and a cross polarization induced by the filter of -15 dB, or less. While the use of solar filters has been demonstrated to work, their use introduces several disadvantages, not the least of which are cumbersome calibration procedures.

Yagoubov (2013b, 2014) pointed out that the ALMA SIS mixers could be de-biased to reduce the mixer gain and effectively increase the saturation level to a degree that allows solar observations without the use of the solar filters, at least for non-flaring conditions on the Sun. These produce lower conversion gain and since the dynamic range scales roughly inversely with gain, these settings can handle larger signal levels before saturating. In addition to the SIS bias voltage, the local oscillator (LO) power can be altered in order to further modify the receiver performance. However, LO power settings have not yet been fully optimized for solar observing.

Two so-called “Mixer De-Tuned” or “Mixer De-Biased” settings have been adopted for solar observations in Band 3, Band 5 and Band 6. These are referred to as solar observing modes MD1 and MD2.

- **Band 3:** MD1 mode uses a bias voltage that sets the SIS mixer to the 2^{nd} photonic step below the voltage gap whereas MD2 mode employs a bias voltage corresponding to the 2^{nd} photonic step above the gap. For MD1 mode, the Band 3 receiver temperature suffers a modest increase, to ~ 50 K, and receiver compression is limited to $\sim 10\%$. For MD2 mode, however, the Band 3 receiver temperature increases significantly, to 800 K, but receiver performance is believed to be essentially linear.
- **Band 5:** MD1 mode settings are nominal for ALMA; i.e., the bias voltage is that used under normal observing conditions. MD2 mode employs a bias voltage corresponding to the 1^{st} photonic step above the voltage gap. The receiver temperatures are higher, in the 500-600 K range, but the receiver is essentially linear.
- **Band 6:** MD1 mode settings are nominal for ALMA; i.e., the bias voltage is that used under normal observing conditions. Receiver compression is again on the order of $\sim 10\%$. MD2 mode employs a bias voltage corresponding to the 1^{st} photonic step above the voltage gap. Again, the receiver temperature is significantly higher, 800 K, but the receiver is essentially linear.
- **Band 7:** MD1 mode settings are nominal for ALMA. Here, too, the bias voltage is that used under normal observing conditions. Receiver compression is somewhat higher than it is for Band 3 and Band 6, of order 15%. No stable MD setting has been established for an MD2 mode.

The MD1 mode is considered to be a “quiet Sun” mode (coronal holes, the solar limb, quiescent filaments, prominences, quiet areas outside of active regions) whereas the MD2 mode is recommended for the “active Sun” (active regions, active filaments, science objectives that require accurate photometry).

8.10.2 Array Observations

The Sun is an extremely large source compared with the Primary Beam of either the 7-m or 12-m antennas (see Chapter 2). The Primary Beam is filled with complex emission when pointing at the Sun, as are the beam sidelobes. The ALMA array ultimately measures the brightness temperature contrast relative to the background Sun, which is resolved out by the array. As noted earlier, in order to recover the absolute brightness temperature of solar targets, it is necessary to include not only interferometric observations (by the 7-m and 12-m antennas), but also Total Power measurements made with a single dish. Single dish fast-scan mapping of the Sun in Total Power mode is addressed in the next subsection.

An advantage to using MD mode observing is that the water vapor radiometers (WVRs), which are used to correct differential phase errors introduced by precipitable water vapor over the array, are not blocked by the ACD. They can therefore be used, in principle, to make such corrections to solar data. Unfortunately, unless the optical depth of the sky is ~ 2.5 or more, which would represent highly non-optimum observing conditions, the WVRs saturate on the Sun. Until the WVRs are modified or replaced to increase their dynamic range to accommodate the Sun, phase corrections based on WVR measurements will not generally be possible when pointed at the Sun. For this reason, solar observations are currently restricted to compact array configurations to minimize such phase errors.

Another consideration, again regardless of whether SFs or the MD modes are used (see Section 8.10.1), is the system IF attenuator settings. The input power changes significantly as the antennas move from the (solar) source to a calibrator and back. The IF chain has two variable attenuators (in steps of 0.5 dB) to ensure that signal levels remain within nominal limits: one in the IF Switch and one in the IF Processor. A concern is whether the variable attenuators themselves introduce unacceptable (differential) phase variation between source and calibrator settings, thereby corrupting phase calibration referenced against suitable sidereal calibrators; and whether there are differences between the spw bandpass response between source and calibrator scans as a result of attenuator settings. Careful testing has shown that this should not be a significant concern.

Bandpass calibration is carried out in the usual manner using MD modes: i.e., a strong calibrator is observed in an MD mode and the bandpass solution is obtained. Bandpass shape and stability were checked for MD modes and attenuator states in Bands 3 and 6. It was found that perturbations to bandpass amplitudes and phases were small. For the IF Switch and IF Processor settings adopted for MD mode observing it was found that the RMS difference between bandpass phases for an MD attenuator state and the nominal attenuator state was generally a fraction of a degree for both Band 3 and Band 6, the maximum being 1.2 deg. Similarly, the normalized amplitude difference was typically a fraction of 1%. No explicit correction for differential bandpass is needed.

In the non-solar case, the antenna temperature (T_{ant}) is small compared to the system temperature, and T_{ant} can therefore be neglected for amplitude calibration (see Chapter 10). In contrast, unlike most cosmic sources, the antenna temperature of the Sun is large (~ 7000 K at 100 GHz). It is therefore necessary to measure both the system temperature and the antenna temperature when pointing at the Sun in order to compute the System Equivalent Flux Density (SEFD) to correctly scale visibility amplitudes.

To estimate the antenna temperature T_{a^*} on the Sun, “single-dish” measurements must be performed using all antennas of the array. Specifically, the standard observing sequence for solar interferometric observations will include the following measurements:

- a “sky” observation P_{sky} , offset from (by typically 2°) and at the same elevation as, the target (Sun)
- a “cold” load observation P_{cold} (also known as the “ambient” load), in which an absorber at the temperature of the thermally-controlled receiver cabin (nominally 20°C) fills the beam path
- a “hot” load observation P_{hot} , in which an absorber heated to $\sim 70^\circ\text{C}$ fills the beam path
- a “zero” level measurement P_{zero} , which reports the levels in the detectors when no power is being supplied

Then the telescope moves to the target (Sun) where the IF attenuation levels are set appropriate to the input power. After the target scan, the telescope again moves to the “sky” position and takes another measurement, called the “off” measurement P_{off} , without changing the IF attenuation.

The antenna temperature of the science target is then given by:

$$T_{a^*} = (P_{sun} - P_{off}) \frac{(P_{sky} - P_{zero})}{(P_{off} - P_{zero})(P_{hot} - P_{cold})} (T_{hot} - T_{cold})$$

The auto-correlation data output from the baseline correlator cannot be used for this measurement because it has insufficient dynamic range to measure P_{zero} . Instead, the necessary measurements rely on Total Power data obtained by the baseband detectors.

8.10.3 Single-Dish Mapping

Fast scanning observations are ideal for recovering the flux or brightness distribution on angular scales ranging from the ALMA primary beam width to the scale of the target in question (typically a few arcminutes), or up

to the full disk of the Sun. Briefly, fast-scan mapping entails making Total Power (and more recently, auto-correlation measurements) as the telescope pointing is driven continuously and smoothly through a sampling pattern on the target that avoids sudden acceleration or deceleration of the antenna drive motors. A major advantage of fast scanning is that it minimizes the impact of atmospheric variation, and the full solar disk can be mapped in as short as 7 minutes.




Figure 8.6: Left: Schematic illustration of a double-circle scan pattern used to map the full disk of the Sun. Actual patterns sample the source more densely to ensure that it is no less dilute than Nyquist sampling. Right: Example of a fast-scan map of the Sun obtained in Band 6.

While various types of scan patterns have been developed and tested for ALMA dishes, Cycle 8 supports the use of a “double circle” pattern, which maps a circular region on the sky. The double-circle pattern is particularly well suited for full-disk mapping because its coverage matches the shape of the solar disk and it repeatedly revisits the region of the center of the disk, allowing atmospheric opacity variations to be corrected for (see Figure 8.6). Standard observing procedures include focus and pointing checks on suitable sources prior to the fast-scan mapping.

8.10.4 ALMA Solar Ephemeris Generator Tool

The tool (URL <http://celestialszenes.com/alma/coords/CoordTool.html>) was developed by the Czech ARC Node to be used together with the OT in preparing solar observations and was tested during the December 2015 solar campaign. It is a javascript based application which runs in any modern browser and on many different operating systems.

User interface consists of several panels. In the input panel, it is possible to select the latest SDO/AIA image in several bands or to upload the users’ own FITS files. Currently, only uncompressed FITS with defined solar WCS keywords using CROTA2 formalism are supported. Visualization and display panels enable the user to pan and zoom in/out the region of interest or to show/hide coordinate grids and tweak the image display by false coloring and level scaling functions typical of other astronomical software packages and the OT.

The actual pointing is done by clicking the desired feature with a green cross marking the position. The coordinates of the pointing are displayed in several coordinate systems inside the pointing panel, where it is also possible to manually define the pointing.

Finally, in the observation panel, the user defines start and end times of the observation and differential rotation profile which will be used for generation of an OT-compatible ephemeris file. There are several rotation profiles to choose from or the user can define his/her own. Clicking *Generate ephemeris file for OT* will display the generated file which can be downloaded by following the *Download data below* link and then imported into the OT.

The ephemeris file is generated from the JPL Horizons file which the tool queries directly from Horizons website. However, during Cycle 4 regular observing sessions there was a period when JPL Horizons site went offline and hence the Ephemeris Tool was unable to function properly. Work is currently underway to enable precise ephemeris calculations in the Ephemeris Tool even when JPL site goes down.

The ALMA Solar Ephemeris Generator Tool comes with a user manual and it is served on several sites for backup purposes. Currently, there are no plans to include the tool functionality into the OT.

8.11 VLBI Observing Mode

The Very Long Baseline Interferometry (VLBI) Observing Mode (VOM) is a variant of the standard interferometry mode with some additional capabilities to allow ALMA to participate in global VLBI networks operating at millimeter wavelengths. Since Cycle 4, ALMA VLBI mode observing has been offered in Band 3 in conjunction with the Global Millimeter VLBI Array (GMVA)¹² and in Band 6 in conjunction with the Event Horizon Telescope (EHT) network¹³. For details about proposing for these opportunities, see the current ALMA Call for Proposals (<http://almascience.org/proposing/call-for-proposals/>).

For all VOM observations in Cycle 8, ALMA will be operated as a phased array of up to 41 12-m antennas, depending on the configuration and weather conditions. The typical expectation is for ~ 37 12-m antennas, making phased ALMA equivalent to a parabolic antenna of diameter $d_{\text{eq}} \sim 73$ m, neglecting efficiency losses. From Table 9.3, the aperture efficiencies in Band 3 and Band 6 are 0.71 and 0.68, and T_{sys} is 70 K and 100 K, respectively. Defining $\textit{Gain} \equiv A_{\text{eff}}/(2k_B)$ where k_B is the Boltzmann constant and $\textit{SEFD} \equiv T_{\text{sys}}/\textit{Gain}$, the corresponding estimates for these quantities are:

$$\begin{aligned} \textit{Gain} &= 0.71\pi(0.5d_{\text{eq}})^2/(2k_B) \times 10^{-26} = 1.05 \text{ K/Jy} \\ \textit{SEFD} &= 67 \text{ Jy} && \text{(Band 3)} \\ \textit{Gain} &= 0.68\pi(0.5d_{\text{eq}})^2/(2k_B) \times 10^{-26} = 1.0 \text{ K/Jy} \\ \textit{SEFD} &= 100 \text{ Jy} && \text{(Band 6)} \end{aligned}$$

Note that the VOM is *not* compatible with the use of subarrays.

When ALMA is operated as a VLBI station, all standard interferometry data products are also output by the ALMA correlator. Therefore, during any VOM execution an observer simultaneously obtains data equivalent to a standard ALMA interferometric observation of the science target, in addition to the VLBI data products. The latter are recorded independently on Mark 6 VLBI recording systems and are shipped to common sites for correlation with the other VLBI site data before delivery to the PI (see below). Some important details of the VOM are elaborated in the next subsections.

8.11.1 General VLBI Considerations

The supporting VLBI networks (i.e., the GMVA for Band 3 and the EHT for Band 6), together with ALMA, place certain restrictions on the observing process for VLBI experiments involving ALMA. These are largely driven by the complexity of orchestrating the simultaneous, reliable execution of a common VLBI schedule, as well as the calibration requirements of the individual telescopes that comprise the global array. In practice, planning an ALMA VLBI observation is similar for Band 3 and Band 6, aside from details such as differing bandwidths and frequencies.

Successful VLBI proposals will be passed to a network Scheduler who builds the common VLBI schedules (encoded in a so-called VLBI EXperiment, or VEX file built with SCHED¹⁴) for a particular observing campaign. In these schedules, some number of contiguous hours will be devoted to each set of science targets and necessary VLBI calibrators.

The schedule will be worked out between the Scheduler and the PI well in advance of the observation. Most sites participating in the VLBI network also have a “friend” of VLBI who assists the Scheduler in working out details specific to that site. Sample VLBI schedules from previous campaigns are available from the aforementioned GMVA website.

¹²See <http://www3.mpifr-bonn.mpg.de/div/vlbi/globalmm/>

¹³See <http://www.eventhorizontelescope.org/>

¹⁴<http://www.aoc.nrao.edu/software/sched>

In general, the required time on the VLBI science target(s) and VLBI calibrators will be broken into a number of VLBI scans of several minutes duration, separated by gaps of the order of several minutes to allow time for local calibrations (pointing, system temperature measurements, etc.) and antenna slew time sufficient to meet the needs of each observatory. VLBI observations that include ALMA will require that the gaps between VLBI scans are of sufficient length to allow for the performance of the calibrations normally required for standard interferometric observing, including flux, gain, bandpass, and polarimetric calibrators (see below). In Cycle 8, targets that are sufficiently bright to permit phase self-calibration (≥ 500 mJy) shall be observed in a so-called “active” phasing mode as has been used in previous cycles. Cycle 8 will also introduce a new VLBI observing mode that allows weaker sources to be observed via “passive” phasing as described in Section 8.11.3. The amplitude calibration is anchored to a Solar System Object when possible, and the flux solutions are bootstrapped to a strong source (≥ 500 mJy) included in the VLBI schedule.

Most VLBI stations currently record *circular* polarization (either left, right, or dual), in contrast to ALMA which records dual linear polarizations. Consequently, *every VLBI observation with ALMA must be treated as a polarization observation* in order that a transformation from linear to circular polarization can be made correctly during post-correlation processing. A special software tool, `PolConvert`, is available for this purpose. Robust execution of `PolConvert` requires a minimum session duration of 3 hours (session length, not time on science target) to achieve adequate parallactic angle coverage on the polarization calibrator sources for the computation of polarimetric “leakage” or D terms (see Section 8.7).

Other VLBI observatories have receiver capabilities that differ from ALMA. One complication is that the sample rate used at ALMA is non-traditional for VLBI, so care must be taken in the experiment set-up to ensure that the VLBI data can be correlated with the other stations. The VOM spectral set-ups chosen for each band have been selected to address these constraints and therefore are fixed. For Band 3 (at 3mm) there are four 1.875 GHz wide-bands centered at 86.268, 88.268, 98.328 and 100.268 GHz, respectively. In Band 6 (at 1.3 mm) the bands are centered at 213.1, 215.1, 227.1 and 229.1 GHz, respectively. A second complication is that not all millimeter VLBI sites are capable of matching ALMA’s full recording bandwidth. The EHT currently matches ALMA’s maximum recording rate at all sites (64 Gbps); the GMVA records what its sites can provide, up to the maximum of ALMA’s 2-GHz, dual-pol recordings, but in all cases at least 4 Gbps (512 MHz per polarization), centered in ALMA’s first baseband at 86.268 GHz. (See <https://www3.mpifr-bonn.mpg.de/div/vlbi/globalmm> for current GMVA capabilities.)

Calibration of VLBI data generally requires observations of VLBI calibration targets interleaved with the main science target. At ALMA, these VLBI calibration measurements are carried out in a manner analogous to the observations of the science targets (i.e., using the VOM and active phasing of the ALMA array). The VLBI calibrators are generally chosen for their utility across the entire global array to meet a particular calibration need (e.g., fringe-finding, bandpass calibration, polarization calibration) and need not be the same calibrators as observed for calibration of the (ALMA-only) standard interferometric data.

During VLBI observations, the observatories are usually not interconnected by a fast network, so the VLBI recordings (onto disk modules) of the antenna baseband signals are physically transferred to common correlation facilities managed by the appropriate network (MPIfR Bonn and MIT Haystack Observatory for the Band 3 and Band 6 data, respectively). These sites are responsible for the correlation of the VLBI data and delivery of the correlated VLBI data products to the observer, as discussed in Section 8.11.2. Correlation is performed using the DiFX correlation software.¹⁵

8.11.2 ALMA Considerations for VLBI

As mentioned in the previous section, all VOM observations should be viewed as polarimetric observations. The reader is referred to Section 8.7 for guidance on this.

In order to participate in a VLBI observation, there are two special requirements at ALMA. The first is that the ALMA control system must tune receivers appropriately and point to the targets specified in the VEX schedule at the appropriate times. The second is that the signals from the ALMA antennas must be coherently summed up, decimated, and recorded on media suitable for delivery to the common VLBI correlator. In order to form this coherent sum, the (arbitrary) phases of the signals at the antennas need to be adjusted during the observation so that they can be coherently added. The necessary signal processing for this takes place within special cards and circuitry in the ALMA correlator, and using special software within the ALMA control system.

¹⁵<http://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/start>

These components were developed by the ALMA Phasing Project (APP).¹⁶ Further information on the APP and the ALMA Phasing System (hereafter APS) is described in Matthews et al. (2018).¹⁷ Key aspects of the APS are also presented in the following sections, with particular emphasis on those aspects which are essential for the VLBI proposer or data analyst. The discussion begins with a description of the VOM Scheduling Block in the next section. The subsections immediately below discuss, respectively, the scan sequences used by the VOM, some important details about how the APS works, and finally, a few comments about the analysis of VLBI data (a full discussion is outside the scope of this document).

VLBI Observing Mode Schedule Block and Execution

As with other observing modes, VOM observations are executed with a Scheduling Block which provides observation-specific details to an SSR observing script. In the case of the VOM, the script is referred to as `StandardVLBI.py`.

The Scheduling Block is created from the observing project and the VEX file using a special tool called `VEX2VOM`. This tool reconciles the Phase 2 Scheduling Block with the schedule and targets specified in the VEX file and thus provides the detailed operating instructions for the SSR observing script. The preparation of the final Scheduling Block with `VEX2VOM` is done by the Friend of VLBI or the AoD prior to execution, once all the necessary details of the observing array are known.

Like the `StandardInterferometry.py` script, the `StandardVLBI.py` script ensures that the necessary calibrations as specified by the OT are performed with the desired cadence. Unlike the `StandardInterferometry.py` script, it must give precedence to the scheduled VLBI scans which need to occur at the appointed times. As mentioned above in Section 8.11.1, the VEX file will include gaps of various durations which are created to allow the observing script to make ALMA-specific calibration observations. Thus the script organizes its work by first noting when the next VLBI scan will occur and then noting which calibrations can be executed within the time remaining until the VLBI scan, and finally executing those calibrations that fit. The VLBI scans in the VEX file include a start time specified in so-called VEX time format, `YYYYyDOYdHHhMMmSSs`, so the execution of the `StandardVLBI.py` script by the SSR will look similar to the following:

```
perform initial ALMA calibrations
VLBI Scan 2016y089d07h00m00s on Target-X for 300s
perform some ALMA calibration
VLBI Scan 2016y089d07h10m00s on Target-X for 300s
perform some ALMA calibration
VLBI Scan 2016y089d07h20m00s on Target-X for 300s
perform some ALMA calibration
VLBI Scan 2016y089d07h30m00s on Target-X for 300s
perform some ALMA longer calibrations
VLBI Scan 2016y089d07h50m00s on Target-X for 300s
perform some ALMA calibration
VLBI Scan 2016y089d08h00m00s on Target-X for 300s
. . .
perform final ALMA calibrations
```

VLBI Scan Sequence

The VEX file specifies the observations in terms of VLBI scans which begin and end at appointed times. The VOM translates these requests into a sequence of (ALMA) correlator scans which typically start prior to the specified time in order to allow the phasing system to converge on a good solution. Each of these correlator scans appears in the ALMA meta-data, and is processed by the telescope calibration system (TelCal) according to the specified intents. The relevant ones for this discussion are intents for ‘*phasing*’ and for ‘*WVR correction*’.

The *phasing* intents direct TelCal to calculate the phases of the signals (i.e., the signals of each polarization in each band for each antenna) relative to those of a designated “reference” antenna. These phases are calculated

¹⁶The APP (S. Doeleman, Principal Investigator) was funded by a Major Research Instrumentation award from the National Science Foundation and an ALMA North America Development Project award.

¹⁷Matthews, L. D., Crew, G. B., Doeleman, S. S., et al., PASP, 130, 5002.

from “channel-averaged” data products produced within the correlator. Presently there are eight channel averages per baseband calculated during one of these correlator subscans (see Fig 8.7). With those phases known, commands are issued to the tunable filter banks (TFBs) within the station cards in the correlator to adjust the antenna phases by exactly the computed values so as to bring all signals into phase with those of the reference antenna. Once these adjustments have been made, a “residual” phase correction typically needs to be made on the following correlator scan, and indeed, such corrections continue to be made every scan until the end of the recording. The phasing loop is closed in the sense that small errors in the phase adjustments may be corrected on subsequent scans. This is the so-called “slow” loop which is illustrated schematically in Figure 8.7.




Figure 8.7: VLBI Scans in the Phasing System. Each VLBI scan is partitioned into “subscans” for correlation and “slow” timescale processing (of order seconds) in TelCal. WVR adjustments are made in the CDP on a “fast” timescale (every second). See text for full discussion and explanation.

While the correlator scans are shown as time-contiguous in Figure 8.7 (green bar), there are in fact small gaps to allow for potential reprogramming of the correlator processing. Even so, the solutions from TelCal (purple) do not arrive at the TFBs (red) until after the correlator scan has started—the time of application is noted and the early part of the scan is excluded from the subsequent phase calculation. Usually one phase application is needed to get an acceptable solution, so the recording should start no sooner than the start of the third correlator subscan. These parameters are all programmable, and judicious choices for them are made when the Scheduling Block is created. For Cycle 8, the subscans are 16 s long with 1 s (channel average) integrations; there are then about 30s of observations prior to the start of the recording. Note that signals from the antennas flow continuously to the summation logic and are recorded independently of the correlator subscan timing, so there are no comparable gaps in the VLBI data. In addition to the “slow” loop, there is also a “fast” loop that optionally may be enabled and which uses the WVR data from each antenna to correct the antenna phases in real time in response to the wet component of the atmosphere. These corrections are performed at a frequency of ~ 1 Hz (blue in Figure 8.7). This loop is open (corrections are continuously made, but feedback is only achieved through the “slow” loop). It may be used if conditions warrant. There is no limit to the duration on the operation of these loops other than general ALMA limitations with regard to total data collected.

VLBI Phasing System

Moving beyond the timing of the phasing system, there are several other aspects of the system that an ALMA VLBI observer should be familiar with. All ALMA observations operate with an array of antennas (i.e., those available to a scheduled observation). Of these, a portion of the antennas are controlled by the phasing system to form the coherent sum signal which is used for the VLBI recordings. The previously mentioned reference antenna is merely a designated member of this “phased array”. At least two antennas from the active array (designated “comparison antennas”) are held outside the phased array. These comparison antennas are thus

available for diagnosing the performance and efficiency of the phased array system. Finally, after construction, the summed signal is decimated to two bits per sample (as is the case with all of the ALMA signals in the VOM), so it is possible to have the ALMA correlator correlate the sum signal during online processing. This is done by co-opting the Correlator Antenna Input (CAI) of some antenna that is not used by the VOM (usually CM10) and inputting the sum signal into the correlator as if it had come from this “sum antenna”. The sum signal therefore appears as antenna on pad “APP001” in any ASDM file containing ALMA data acquired with the VOM. There are a few things to note about the data and metadata for this antenna:

- there are no useful metadata, as this is not a physical antenna
- there are no WVR data, for the same reason
- the data are **highly** correlated with other antennas in the phased array since the same receiver noise is present in both.
- this antenna should generally be flagged during analysis of the (ALMA-only) interferometric array data

The selections of antennas for the phased array and the comparison array are made when the Scheduling Block is created, although these may be adjusted if necessary shortly prior to or during the observation. Normally, the antennas to be phased are chosen to lie within a certain radius from the array center (up to several hundred meters, depending on the band and the observing conditions). The efficiency of the phasing system decreases on longer baselines, so in practice, maximum baseline lengths in the phased array will typically be $\lesssim 1$ km. The reference antenna is in general selected to be centrally located, and typically at least one of the comparison antennas is chosen to be within the radius of the phased array. Because of the decimation to two-bits, the number of phased antennas must always be odd.

A final important note concerns delays between antennas. The ALMA control system estimates the total delays between all antennas and adjusts the signals through a variety of techniques (in hardware and software) so that the visibilities found in the dataset have essentially zero delay on each baseline, (i.e. there is no slope in a plot of visibility phase as a function of frequency on any baseline.) However, a (significant) component of the delay is removed with a frequency dependent phase rotation in the online correlation processing, not in the hardware, so there is some significant delay present in the signals at the TFBs where the APS phase corrections are made and correspondingly in the logic which creates the “sum antenna” signal. Thus in order for the phasing system to work properly, it is necessary to turn off this component of software delay correction and to instead correct it in the phasing calculations. To do this, the observing band is subdivided into several channel averages (at least 8) and the normal phase-solving procedure produces different, independent phases that can be applied to the TFBs corresponding to each channel average. While this removes most of the delay, it does leave a small residual phase-slope within each channel average. It also leads to the restriction that the source needs to be bright enough in each of the channel averages so that a usable phase solution can be found. A flux density limit for actively phased VLBI targets of ≥ 500 mJy is therefore imposed (but see Section 8.11.3).

VLBI Observing Mode Analysis

As pointed out in Section 8.11, any VOM observation actually involves two concurrent interferometric observations: one on the ALMA scale that results in correlated data products between all participating ALMA antennas and one on a global scale, where ALMA serves as a station in a global VLBI array. A description of the reduction and analysis of ALMA VLBI data is beyond the scope of this document. This process is, however, informed by the analysis of the ALMA-scale observation with regard to observatory performance (SEFD, system temperature, etc.), which for the most part may be obtained in the usual way. However, since the VLBI signal that is recorded is that of a “sum” antenna whose properties in turn depend on the phasing performance, a few details should be pointed out.

One of the most significant is that the correlated amplitude of the signal measured by an ideal phased array should scale as the square root of the number of phased antennas, assuming the antennas are identical (at ALMA, only 12-m dishes are typically used in the phased sum and they have similar effective areas; if the 7-m dishes were included, the relation would become more complicated). In practice, a number of effects conspire to lower performance from this ideal relation¹⁸ to $\sim 80\%$ of what would be expected in a perfect system. Part of

¹⁸ALMA-05.11.63.03-0001-A-REP.pdf

this efficiency loss stems from the fact that the sum antenna signal is decimated to two bits from the individual antenna signals (about 7%) in addition to the original decimation of to two bits (about 12%) for a total reduction of about 0.8 (this is based on analyses subsequent to those reported in Matthews et al. 2018, PASP, 130, 5002). To this must be added phasing efficiency losses caused by decorrelation, which are weather-dependent. This loss is typically $\sim 10\%$ and its effect on the visibility amplitudes is calibrated during QA2. For a specific observation, the net effective area of the ALMA “sum” dish may be estimated directly from the ASDM data on a per-scan basis.

The performance of the phasing system is captured in the metadata of `CalAppPhase` table of the ASDM file. There is one table entry per scan per baseband. In each table, there are temporal bounds on validity (i.e., the range of time within the scan of stable phase) and entries to list the disposition of the antennas amongst the different categories (phased antennas, reference antenna, and comparison antennas) and whether this represents a change from the previous scan. The table primarily reports a number of phase values, N_v , which is dependent on the “packing mode”. There is one value per channel average per polarization per antenna. Additionally, the `CalAppPhase` table reports lower-level details of the phasing system, including online per-antenna assessments from TelCal of the antennas drawn from the quality of the phasing solution. This table may be used in conjunction with the ALMA data for a particular scan to calculate the phasing efficiency and thus to make an appropriate correction to the absolute VLBI correlation amplitudes.

As indicated in Section 8.11.1, the VLBI correlation products are delivered to the PI for analysis, separate from the ALMA ASDM data. Traditionally for the GMVA the former has been in the form of FITS-IDI files. A complication for work with ALMA data is that the polarization conversion is performed with a software tool (`PolConvert`) that requires input from a CASA analysis of the ALMA data set (see Goddi et al. 2019, PASP, 131, 5003). This conversion can be performed on the FITS-IDI file or at the correlator using the raw correlator output if the CASA analysis is available. Arrangements for doing this will be worked out between the PI of each project and the supporting observatory collaborations (GMVA and EHT) following the observation. During the (`PolConvert`) process, a detailed calculation is performed from the calibration tables to provide a traditional ANTAB format file for amplitude calibration of the visibilities on baselines that include ALMA.

8.11.3 Observing Targets Fainter than 500 mJy with Passive Phasing

To insure adequate SNR for the calculations necessary to phase up the ALMA array, (“active”) phasing is currently restricted to sources that are point-like and have correlated flux densities ≥ 500 mJy on short (< 1 km) intra-ALMA baselines. However, beginning in Cycle 8, VLBI science targets weaker than 500 mJy may be observed using a “passive phasing” technique. In this mode, a phase calibrator close in angular distance to the science target (termed a “phasor”) is used by the APS to phase the antennas immediately prior to a slew to the science target. The phasor must be point-like and brighter than 500 mJy on intra-ALMA baselines and ideally should be within ~ 6 or 3 degrees of the science target for Bands 3 or 6, respectively (see section B.6.1 in the Cycle 8 2021 Proposer’s Guide). When selecting a phasor, closer is better than brighter. If there are no suitable phasor candidates within the suggested limits, the proposer should justify in the proposal the need to use a more widely separated calibrator (see section B.6.1 in the Cycle 8 2021 Proposer’s Guide). In addition to higher overheads for slewing, losses in correlated amplitude as a result of phase decorrelation are likely to occur on the science target when using a phasor more than a few degrees from the target, particularly in Band 6.

In practice, there is no lower limit to the flux density of science targets that can be observed with passive phasing, aside from SNR considerations on the VLBI baselines. However, while the 500 mJy flux density threshold is removed from the *science target* in the passive phasing mode, the *phasor* source itself is actively phased and must adhere the same requirements for flux density and source structure as any standard, actively phased VLBI target.

With passive phasing the need to periodically observe the phasor will result in some reduction of observing efficiency compared with the case when it is possible to actively phase the array directly on the science target. There are also trade-offs to consider when choosing the phasor/science target duty cycle. Typical passive phasing observing sequences will consist of a short (~ 1 – 3 min) VLBI scan on the phasor, immediately followed by a longer scan on the science target itself (up to ~ 5 min). For accurate absolute flux calibration, care must be taken not to exceed the expected atmospheric coherence time during the science target scans, since correlated amplitude will diminish with time as phase coherence is lost. In general, shorter scan sequences will have higher overhead but will be more robust against atmospheric variability, which cannot be predicted precisely in

advance. The proposer should indicate (for the scheduler’s benefit) their preferences regarding this trade-off.

8.11.4 Phased-Array Observing

A special case of the passive phasing technique (hereafter the Phased-array Mode or PM) is now available for general use in Cycle 8. This mode has been tested successfully in Band 3 via observations of the Vela pulsar.¹⁹ The PM utilizes phased ALMA as a standalone array. Although this is *not* a VLBI observation in the traditional sense, ALMA is operated as a phased array and VLBI technology is used to record the baseband data for subsequent use in pulsar analyses or searches.

In Cycle 8, the PM is offered only in Band 3 for pulsar science projects. Extension of the PM to other bands in future cycles is expected to be straightforward, but it will require additional testing and commissioning. As with the VOM, the PM requires ALMA to be operated as a phased array using the APS. Thus, a schedule of recordings is required, which is analogous to the situation with VLBI observations. The PM will be supported by the GMVA collaboration, headquartered at the Max Planck Institut für Radioastronomie (MPIfR) in Bonn, Germany. Therefore for successful proposals, the schedule will be generated by the GMVA Scheduler in consultation with the principal investigator (PI) well in advance of the observation. In preparation, the PI will be required to specify various observing considerations using the OT. The total potential recording bandwidth available for Band 3 will be 4×2 GHz (for both polarizations) and each baseband/polarization will contain 32×62.5 MHz frequency channels.

Phase-up of the array for the PM will broadly follow the same mechanism used by the standard VLBI passive phasing mode as described in Section 8.11.3. In practice, the total required integration time will be broken into a number of scans, each of up to several minutes in duration, separated by ~ 1 –3-minute actively phased exposures on a phasor target selected from the ALMA Calibrator Source Catalogue. Provided the weather is stable, the phase coherence on a pulsar target may be maintained for as long as several minutes in Band 3. Requirements for the properties of the phasor are the same as for non-pulsar passive phasing experiments (see Section 8.11.3), although the separation to the phasor may be relaxed with justification in the proposal (see section B.6.1 in the Cycle 8 2021 Proposer’s Guide).

After the observation, the recordings will be sent to MPIfR for subsequent processing. This post-processing involves the generation of FITS-format data of a special type used by the pulsar community (known as PSR-FITS²⁰). These data are suitable for pulsar and transient searching, pulsar timing, or any other relevant scientific objectives (see the next section).

Pulsar Mode Post-processing and Data Products

Every PM project will generate two types of data products: pulsar data in the form of a high time-resolution time series (in PSRFITS format) and the ALMA standard interferometric data (in ASDM format). A sketch summarizing the overall data flow is shown in Figure 8.8. To generate the pulsar data product, the baseband data are first read off the disk modules into a collection of VLBI Data Interchange Format (VDIF) data files (one per scan per sub-band per polarization).

Next, a software routine²¹ is used to convert the VDIF files into the search-mode PSRFITS format. To obtain one time sample for a single frequency channel, the routine extracts the time series in both polarizations (with identical time stamps) from one “packet” in the VDIF file (also referred to as a frame), performs a Discrete Fourier Transform, makes conversions from voltage to intensity in the frequency domain, and collapses the power over the whole Fourier spectrum. There are 125,000 frames per second for each recorded data stream (one polarization from one baseband). Detections can then be outputted in the form of total intensity only, intensity of dual polarizations, coherence products, or full Stokes products.

The time samples (in 32-bit floating point) are then packed together following the structure of the PSRFITS specification, retaining the original frequency resolution. The time resolution of the output file may be as high as 8 microseconds, but multiples of that (smaller files) can be produced upon request. The final data product will be on a per scan and per subband basis and will be suitable for direct processing by standard pulsar software

¹⁹Liu, K., Young, A., Wharton, R., et al. 2019, ApJ Letters, 885, L10.

²⁰https://www.atnf.csiro.au/research/pulsar/psrfits_definition/Psrfits.html

²¹Available at: <https://github.com/xuanyuanstar/vdif2psrfits>

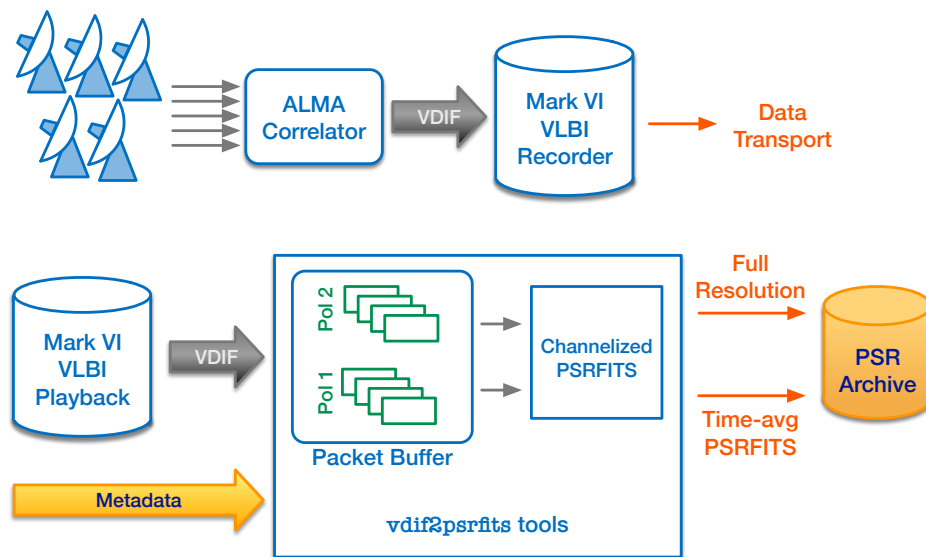


Figure 8.8: Diagram showing the data flow during a pulsar mode (PM) observation. Baseband data recorded at ALMA will be sent to MPIfR (Bonn, Germany) for playback and PRSFITS generation. Reproduced from Figure 2 of Liu et al. (2019).

such as PRESTO²², PSRFITS_UTIL²³, and DSPSR²⁴ to perform pulsar and transient searching, pulsar timing, or any other desired scientific investigations. The data rate (per second) for each spectral window is given by 4 bytes \times number of polarizations \times 32 channels \times 125,000; for example, for 8- μ s resolution and 4 polarizations, it is $4 \times 4 \times 32 \times 125,000 = 64$ MB/s.

For PM observations of a known pulsar, the post-processing will not include construction of the integrated pulse profile (which requires the ephemeris of the pulsar). However, this can be done later by the PI.

Technical Justifications for Phased-Array Observations

Upon the submission of a proposal requesting PM observations in Cycle 8, the PI will be required to provide the following information via the OT. Note that there will be one science target per goal (per SB), but multiple goals per project are allowed, if justified.

- Name and position (RA and DEC) of the science target.
- Total observing time requested. In PM projects, approximately one third of the total requested time will be spent on the science target and the rest of the time on the VLBI calibrators (including the phasing calibrator and other calibration targets in the schedule). Additional calibrations for the ALMA interferometric data are performed during schedule gaps.
- Phase calibrator (phaser). This should generally be chosen from the ALMA Calibrator Source Catalogue. However, a phaser may be chosen from outside of this Catalogue, if justified.
- Selection of the spectral window. In Cycle 8, Band 3, there are four 2-GHz spectral windows, centered at 86.268 GHz, 88.268 GHz, 98.328 GHz, 100.268 GHz, respectively. The PI can choose between one and four of these by declaring it in the justification box.
- Time resolution of the PSRFITS data product. This needs to be an integer multiple of 8 μ s.
- Polarization of the PSRFITS data product. The PI can choose any one from the following: (1) Total intensity; (2) dual polarizations (XX & YY); (3) coherence product; (4) All four cross correlations (XX, XY, YX, & YY).
- Any additional request, if applicable.

8.12 High Frequency Observing

High-Frequency Observations (specifically Bands 8, 9, and 10) require some special observing techniques, mostly because of the combined ill effects of lower atmospheric transmission and fainter and more widely spaced calibrators at these frequencies. These techniques, which are just addenda to regular observing modes, are being progressively implemented in the ALMA observations since Cycles 6, after some extensive testing is carried out. It is expected that all of them will improve not only the quality of the final data products, but also the overall observing efficiency at the ALMA highest frequencies. For reference, a brief description of the nomenclature that could be used by Contact Scientists and the ALMA Observatory in general when matters regarding high-frequency projects are discussed, reported or communicated to PIs is given below:

- **High-Frequency Cone-Searches and Calibrator Surveys:** Observations in the high-frequency ALMA bands require that the availability of calibrators near the target source is assessed near the time of the planned observations. Searches for phase calibrators and check sources (see Chapter 10 for details) within 10 degrees of the science targets (i.e., the so-called ‘‘Cone-Searches’’) in each SB are carried out by ALMA staff and Allegro, the EU ARC Node in the Netherlands, within 90 days of the high-frequency observations. The results of those measurements are subsequently ingested into the ALMA Calibrator Catalogue, so that they are available for queries at the time the actual observations take place.

²²<https://www.cv.nrao.edu/~sransom/presto/>

²³https://github.com/demorest/psrfits_utils

²⁴<http://dsp.sr.sourceforge.net/>

- **Bandwidth Switching (BWSW):** For high-frequency observations (Bands 8, 9, 10) with spectral set-ups including only narrow-band spws, the PI will be asked to add an additional wide-band spw. This approach increases the accuracy of the gain calibration, allowing a solution across the combined spws and/or transferring of the wide-window solutions to the narrow spws. When adding extra spws is not possible (i.e., projects that have occupied all the spws with narrow-band set-ups), it is possible to use bandwidth switching, in which the phase calibrator is observed with wider spws for increased signal-to-noise. The wide-to-narrow spw complex-gain offsets are measured from observations of the bandpass calibrator, allowing the phase cal solutions to be transferred to the target field. For more details, see Chapter 10.
- **Sideband separation by 90-degree Walsh phase-switching (Bands 9 and 10):** The Band 9 and 10 receivers are of DSB design, i.e. do not separate or suppress either sideband, and signals from both sidebands are superposed on a single IF output per polarization (see Chapter 4). Prior to Cycle 6, the only option available for interferometric observing in these bands has been to use orthogonal LO frequency-switching sequences to suppress the contributions of one sideband to the visibilities. Since Cycle 6 a means of sideband separation is also offered, resulting in twice the output bandwidth and number of channels. The separation uses orthogonal sequences of 90-degree phase shifts applied at the first LO in each antenna, combined with synchronised accumulation over the sequence duration of the three phase combination states of the cross-correlations of each baseline by the correlators. The result is subsequently combined to produce the two sideband spectra in the correlator data processing software. There are 128 patterns in a switching sequence, each pattern switching on 16 ms intervals, and therefore total length of 2048 ms. This mode therefore adds a constraint: the dump and integration durations of the correlator must be a multiple of 2048 ms. Due to an unrelated requirement that subscans must start and end on 48 ms Timing Event (TE) boundaries, and 2048 ms not being a multiple of 48 ms, the allowed values of subscan durations are somewhat limited, as will be seen when adjusting times in the OT. In Cycle 8, only a single pair of spws (LSB+USB) in each baseband is offered, and either the TDM (2 GHz, 256/Npol channels) or widest bandwidth FDM (1.875 GHz, 7680/Npol channels) modes must be used (additionally for FDM, the spw should not be moved from the center of the baseband). The modes in the four basebands remain independent, e.g. one baseband can use FDM to observe a spectral line, while the other basebands can use TDM to detect continuum.

Chapter 9

ALMA Sensitivity Calculator

The main tool for calculating the sensitivity of ALMA is the ALMA Sensitivity Calculator (ASC). This is an application contained within the ALMA Observing Tool (OT) that is used to perform various conversions between sensitivity (in either flux or temperature units) and time. Its main use is to take the user-requested sensitivity and calculate the necessary observing time for each source in a Science Goal using parameters such as the representative frequency, source declination, number of antennas and the expected precipitable water vapour (PWV). Many of the inputs cannot be changed by the user and, for example, the OT will always assume a fixed number of antennas for the particular cycle and will always use the standard PWV octile that is appropriate to the frequency of observation.

In addition, a separate version of the ASC is available in the OT which allows a user (via a GUI) to experiment with various sensitivity options such as the number of antennas and PWV octile. An almost identical version is also available as a web application that can be found in the ALMA Science Portal¹. This differs from the GUI version incorporated in the OT in that it is possible to enter some parameters that are outside the ranges that are applicable to the current observing cycle.

9.1 Calculating the System Temperature

When determining the time required to achieve a particular sensitivity, the system temperature, T_{sys} , is a fundamental parameter as it takes into account various sources of noise that make it difficult to detect the very weak astronomical signals that ALMA is trying to detect. The most prominent sources of noise are from the receivers and from the atmosphere. The latter is highly variable, both in time and frequency, and thus dynamic scheduling and careful placement of spectral windows (spws) are crucial.

9.1.1 Sky Temperature

The OT's estimate of both the atmospheric zenith opacity, τ_0 , and the sky temperature, T_{sky} , are calculated using the Atmospheric Transmission at Microwaves (ATM) code². This provides values of the opacity and the atmospheric "output radiance", in steps of 100 MHz, for the seven different octiles of PWV. The sky temperature is converted from the radiance using the Planck function and includes the contribution due to the CMB.

The ATM code only provides measurements of the sky temperature at the zenith, $T_{\text{sky}}(z = 0)$, and therefore the OT must account for the greater atmospheric emission at lower elevations. It does this by assuming that the emission from the atmosphere can be approximated as

$$T_{\text{sky}} = T_{\text{atm}}(1 - e^{-\tau_0 \sec z}). \quad (9.1)$$

Inserting the ATM values of $T_{\text{sky}}(z = 0)$ and τ_0 into Equation 9.1 allows the mean physical temperature of the

¹<http://almascience.org/>

²See Pardo, J. R., Cernicharo, J., Serabyn, E., 2001, ITAP, 49, 1683. This calculates the sky temperature by integrating the atmospheric temperature profile, this having been formed from the average of 28 radiosonde measurements taken at the ALMA site during November 1999.

Octile	PWV (mm)
1	0.472
2	0.658
3	0.913
4	1.262
5	1.796
6	2.748
7	5.186

Table 9.1: Octiles of PWV measured at the ALMA site from years of monitoring data and used in the ASC. The first octile corresponds to the best weather conditions and shows that 12.5% of the time, PWV values at least as good as 0.472 mm can be expected. Subsequent octiles give the corresponding value for 25%, 37.5% etc.

atmosphere, T_{atm} , to be measured i.e.

$$T_{\text{atm}} = \frac{T_{\text{sky}}(z = 0)}{(1 - e^{-\tau_0})}. \quad (9.2)$$

This can be reinserted into Equation 9.1 to calculate the sky temperature at any zenith angle as

$$T_{\text{sky}}(z) = T_{\text{sky}}(z = 0) \frac{(1 - e^{-\tau_0 \sec z})}{(1 - e^{-\tau_0})} \quad (9.3)$$

$T_{\text{sky}}(z)$ is then corrected for the fact that the required noise temperatures (T_n) are defined assuming $P_\nu = kT$ and thus a correction for the Planck law is required, i.e.

$$T_n = T \times \left(\frac{h\nu/kT}{e^{h\nu/kT} - 1} \right) \quad (9.4)$$

The octiles characterize the amount of PWV that can be expected at the ALMA site, i.e. a value of PWV *at least as good* as the first octile value can be expected 12.5 per cent of the time, a value at least as good as the second octile 25 per cent of the time, and so on. The octiles corresponding to the ALMA site (determined from many years of monitoring) are shown in Table 9.1.

When estimating the time for a project, the OT will always select a PWV octile that is appropriate to the frequency being observed. It does this by calculating the time required for each octile and then choosing (and reporting) the highest (worst) octile for which the increase in time relative to the first is less than 100 per cent. A consequence of this definition is that the octile also depends on source declination, i.e. sources at low elevations will require better weather conditions. The resulting curve of octile versus frequency is shown in Figure 9.1, for a source declination of zero degrees. A user can override this choice in the GUI or web application versions, but submitted projects will always use an automatic choice.

9.1.2 CMB Temperature

The temperature of the Cosmic Microwave Background is included in T_{sky} and thus does not feature further in this document.

9.1.3 Receiver Temperatures

The receiver temperatures of the various receivers have been updated in Cycle 7. For Bands 3 to 8, these are now based on actual measurements at the telescope and are generally conservative i.e. refer to the worst value seen at any part of the band. The values used in the ASC are given in Table 9.1.3. Note that single sideband noise temperatures are reported for Bands 1-8 and double sideband temperatures for Bands 9 and 10.

At the moment, no attempt is made to incorporate the frequency dependence of the receiver temperature, i.e. only a single value is used per band. Ultimately, it is the intention to use the actual measured values for all receivers and to incorporate the frequency response across the band.




Figure 9.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown at the top of the plot. The water line at 183 GHz (Band 5) is particularly prominent. In general, higher frequencies require drier observing conditions.

Note that the ASC doesn’t concern itself with the so-called “zero-point fluctuations” as the requisite half photon of noise ($h\nu/2k$) has already been included in the noise measurements provided by the various receiver groups (A. Kerr, private communication).

The receiver temperatures are already expressed in terms of the Planck expression and thus do not require the correction given in Equation 9.4.

9.1.4 Ambient Temperature

Ambient temperature is essentially spillover from the sidelobes of the antenna beam corresponding to emission from the ground and the telescope itself. This is held constant at 270 K (median value as measured from many years of monitoring data at the ALMA site). However, the value used by the ASC is converted to a noise temperature according to Equation 9.4 and thus its total contribution is frequency dependent and can vary between the different sidebands.

9.1.5 DSB Receivers

Due to the way that the astronomical signals are down-converted to an intermediate frequency, every heterodyne radio receiver simultaneously detects radiation from two sidebands. This means that, if nothing were done to prevent it, a spw processed by the correlator would contain two sets of astronomical signals mixed in with one another, one from the “signal” sideband and an undesirable one from the “image” sideband. In the case of 2SB

ALMA Band	Receiver Type	$T_{\text{rx,ASC}}$ (K)
1	SSB	25
2	SSB	30
3	2SB	40
4	2SB	42
5	2SB	50
6	2SB	50
7	2SB	72
8	2SB	135
9	DSB	105
10	DSB	230

Table 9.2: Receiver temperatures assumed in the ASC as a function of ALMA band. For Bands 3–8, these are all based on actual measurements made at the telescope and correspond to the worst value seen within each band. For many bands, these are considerably better than the original specifications in Table 4.1.

receivers, the contribution from the image sideband (emission from the source and noise) is suppressed to a very high degree. However, for DSB receivers (Bands 9 and 10) it is only possible to remove the source contribution (either by LO offsetting or 90-degree Walsh switching) and so the noise cannot be neglected.

One important consequence of this is that, if a spw has its image counterpart in an area of very poor atmospheric transmission, it can greatly increase the system temperature and lead to very long on-source times. Therefore, it is important to avoid areas of bad atmospheric opacity in the image spws and the OT therefore shows the location of these in the spectral visual editor (Figure 9.2).

One subtlety is that the tuning software will always place spws in the upper sideband if possible. Therefore, to take a simple case, a single spw centered at 637 GHz will find its image equivalent in the middle of the zero transmission feature at ~ 621 GHz. Where possible, this situation can be avoided by defining dummy spws such that the line of interest is forced into the other sideband (Figure 9.2).

9.1.6 System Temperature

The OT version uses two distinct formulas depending on whether a double sideband receiver is being used, or not. The DSB equation is the following:

$$T_{\text{sys,dsb}} = \frac{1}{\eta_{\text{eff}} e^{-\tau_0 \sec z}} \left(2 \times T_{\text{rx}} + \eta_{\text{eff}} (T_{\text{sky,s}} + T_{\text{sky,i}}) + (1 - \eta_{\text{eff}}) \times (T_{\text{amb,s}} + T_{\text{amb,i}}) \right) \quad (9.5)$$

where

- T_{rx} – receiver temperature
- $T_{\text{sky,s}}$ – sky temperature at the requested frequency in the signal sideband
- $T_{\text{sky,i}}$ – sky temperature in the image sideband
- $T_{\text{amb,s}}$ – ambient temperature in the signal sideband
- $T_{\text{amb,i}}$ – ambient temperature in the image sideband
- η_{eff} – the coupling factor, or forward efficiency. This is equal to the fraction of the antenna power pattern that is contained within the main beam and is currently fixed at 0.95
- $e^{-\tau_0 \sec z}$ – the fractional transmission of the atmosphere, where τ_0 is equal to the zenith atmospheric opacity and $\sec z$ is the airmass at transit (the ASC always assumes that the source is being observed at transit).




Figure 9.2: A Band 9 spectral setup as displayed in the spectral visual editor in the OT. The top figure shows an example of how the image equivalent of a single spw can fall into an area of poor atmospheric transmission, leading to much higher T_{sys} than necessary. Placing a dummy spw at a higher frequency can remedy the situation (bottom).

The terms η_{eff} and $e^{-\tau_0 \sec z}$ both attenuate the source signal and thus we must divide through by them in order to obtain a measure of the system noise that is relative to the unattenuated source. Note that this is always done at the *signal* frequency.

For SSB and 2SB receivers, the equation is the following:

$$T_{\text{sys,ndsb}} = \frac{1}{\eta_{\text{eff}} e^{-\tau_0 \sec z}} \left(T_{\text{rx}} + \eta_{\text{eff}} T_{\text{sky,s}} + (1 - \eta_{\text{eff}}) \times T_{\text{amb,s}} \right) \quad (9.6)$$

where the terms are all the same as in Equation 9.5.

9.1.7 System Temperature (web application and OT GUI)

The equations to calculate the system temperature used by the web application and the OT's GUI are similar to the above, however, both are unaware of the details of the tuning setup and so do not know the location of the image sideband and cannot perform a rigorous calculation of its contribution to the system temperature. In this case, the same equation is used for all receiver types and the DSB noise contribution is simply double the single sideband case. This is controlled via the sideband gain ratio, g :

$$T_{\text{sys}} = \frac{1 + g}{\eta_{\text{eff}} e^{-\tau_0 \sec z}} \left(T_{\text{rx}} + \eta_{\text{eff}} T_{\text{sky,s}} + (1 - \eta_{\text{eff}}) \times T_{\text{amb,s}} \right). \quad (9.7)$$

For Bands 1 and 2 (Single Sideband; SSB) and 3-8 (Sideband Separating; 2SB), $g = 0$. For DSB receivers, $g = 1$.

9.2 The Sensitivity Calculation

Once T_{sys} has been determined it is possible to calculate the point-source sensitivity given a requested amount of on-source observing time or vice versa. At no point is any account made for the expected level of loss in sensitivity due to problems with the telescopes such as residual pointing and focus error.

Under certain circumstances it will not be possible to achieve the theoretical sensitivity, particularly if the source being observed is very bright and/or the u, v coverage relatively sparse. In these situations the image

can be dynamic-range limited due to residual source flux scattered across the map as a result of calibration errors and imperfect source deconvolution.

9.2.1 12-m and 7-m Arrays

When dealing with the 12-m and 7-m Arrays, the point-source sensitivity, σ_S , is given by the standard equation:

$$\sigma_S = \frac{w_r 2k T_{\text{sys}}}{\eta_q \eta_c A_{\text{eff}} (1 - f_s) \sqrt{N(N-1)} n_p \Delta\nu t_{\text{int}}}. \quad (9.8)$$

The various parameters are:

- w_r – robust weighting factor. ALMA Pipeline imaging and subsequent QA2 assessment is performed assuming that the visibilities are weighted using robust weighting, specifically a Briggs robustness factor of 0.5. Simulations have shown that this factor is equal to 1.1.
- A_{eff} – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency $\eta_{\text{ap}} = R_0 \exp(-16 \pi^2 \sigma^2 / \lambda^2)$ where σ is the rms surface accuracy of the antenna. The latter is set to the design goal of 25 μm and 20 μm for the 12-m and 7-m antennas respectively³. R_0 is the product of a number of different efficiencies and is equal to 0.72. See Table 9.3 for values of antenna efficiencies in various ALMA bands.
- f_s – shadowing fraction. For the more compact 12-m configurations and the ACA 7-m Array, antennas can block the field-of-view of other antennas in the array and thus reduce the total collecting area. The shadowing fraction is a function of source declination as shown in Fig. 7.5.
- η_q – quantization efficiency. A fundamental limit on the achievable sensitivity is set by the initial 3-bit digitization of the baseband signals. This is equal to 0.96.
- η_c – correlator efficiency. This depends on the correlator (64-input or ACA) and correlator mode, although the efficiency of all 64-input Correlator modes is equal to 0.88. The ACA efficiencies *do* depend on the mode, but this is only taken into account in the OT, not by the web application or OT GUI, which therefore assume a value of 0.88.
- N – number of antennas. This defaults to 43 for the 12-m Array and 10 for the 7-m Array.
- n_p – number of polarizations. $n_p = 1$ for single polarization and $n_p = 2$ for dual and full polarization observations.
- $\Delta\nu$ – resolution element width. As already mentioned, this should be equal to 7.5 GHz for continuum observations. This is due to the maximum usable bandwidth of a spw being limited to 1.875 GHz by the anti-aliasing filter through which the baseband signal passes. $n_p \Delta\nu$ is often referred to as the effective bandwidth.
- t_{int} – total on-source integration time.

The associated surface brightness sensitivity (K) is related to the point-source sensitivity (Jy) by

$$\sigma_T = \frac{\sigma_S \lambda^2}{2k \Omega} \quad (9.9)$$

where Ω is the beam solid angle. This is related to the user-entered spatial resolution, θ , by

$$\Omega = \frac{\pi \theta^2}{4 \ln 2}. \quad (9.10)$$

This assumes that the telescope beam is a circular Gaussian with a half power beamwidth of θ .

³Note that not all antennas might achieve this specification. The performance of a given antenna will also vary with the thermal conditions and the length of time between surface realignments.

Band	Frequency (GHz)	$\eta_{\text{ap},12 \text{ m}}$ (%)	$\eta_{\text{ap},7 \text{ m}}$ (%)
3	100	71	71
4	145	70	71
5	183	69	70
6	230	68	69
7	345	63	66
8	405	60	64
9	690	43	52
10	870	31	42

Table 9.3: Aperture efficiencies at typical continuum frequencies for both the 12- and 7-m antennas. The effective area, A_{eff} , is equal to η_{ap} multiplied by the physical area of the dish i.e. 113.1 m² and 38.5 m² for the 12-m and 7-m antennas respectively.

9.2.2 Total Power Array

In the case of the TP Array, a different equation is used

$$\sigma_{\text{TP}} = \frac{2kT_{\text{sys}}}{\eta_{\text{q}}\eta_{\text{c}}A_{\text{eff}}\sqrt{N}n_{\text{p}}\Delta\nu t_{\text{int}}}. \quad (9.11)$$

This is similar to Equation 9.8, but there is only a factor of \sqrt{N} in the denominator and there is no need to take shadowing into account.

Particularly for continuum observations, the above equation is likely to be too optimistic due to rapid fluctuations of the receiver gain and atmospheric opacity. These require extremely demanding calibration strategies and, as these have not yet been commissioned, only spectral line total power projects are currently possible.

9.3 User Interface

The main way that a user interacts with the Calculator is through a GUI in the OT or via the web application in the ALMA Science Portal – both are essentially identical. By entering various parameters, the time required to achieve a particular sensitivity (in either Jy or K) can be calculated, or vice versa. The inputs that affect the sensitivity or time are given below; a screenshot of the OT’s GUI version is shown in Figure 9.3.

- Source declination – this is used to calculate the maximum elevation of the observation and thus the minimum airmass i.e. the ASC assumes that the source is transiting. It is also used to calculate the amount of shadowing that is likely to affect the ACA 7-m and the smaller 12-m configurations.
- Observing frequency – this sets the receiver temperature, antenna efficiency and the PWV octile.
- Bandwidth per polarization – this otherwise straightforward parameter should be set to 7.5 GHz for continuum observations (see Section 9.2.1). For spectral line observations, it is usually set to the frequency/velocity resolution that one requires in one’s spectrum.
- Water column density (Precipitable Water Vapor; PWV). The user is able to enter one of the seven octile values, or the calculator will set this automatically depending on the frequency entered.
- Number of antennas – the ASC currently assumes 43 from the 12-m Array, 10 from the 7-m Array and three from the TP Array.
- Angular resolution – this affects the time estimates when sensitivities are specified in temperature units. The calculator will not perform any calculations when Kelvins have been specified, unless a non-zero value for angular resolution has been entered. The calculator will also issue a warning if the angular resolution falls outside of the range corresponding to 125-m and 1-km baselines.

Common Parameters

Dec	-21:03:39.900	
Polarization	Dual	
Observing Frequency	86.00000	GHz
Bandwidth per Polarization	7.50000	GHz
Water Vapour Column Density	5.186 mm (7th Octile)	
Trx, tau, Tsky	45 K, 0.047, 12.693 K	
Tsys	77.736 K	

Individual Parameters

	12m Array	7m Array	Total Power Array
Number of Antennas	43	10	3
Resolution	0.8 arcsec	0.00000 arcsec	67.7 arcsec
Sensitivity (rms)	10.00000 uJy	1.00000 mJy	1.00000 mJy
(equivalent to)	0.00258 K	Unknown K	0.00004 K
Integration Time	1.24903 h	1.29120 min	4.51151 min

Integration Time Unit Option: Automatic
Sensitivity Unit Option: Automatic

Buttons: Calculate Integration Time, Calculate Sensitivity, Close

Figure 9.3: Screenshot of the GUI version of the ALMA Sensitivity Calculator as implemented in the ALMA Observing Tool. The white area at the bottom is for displaying error messages i.e. parameters out of bounds. The example here shows the achievable sensitivity for all three arrays for an on-source time of 10 minutes.

The calculator reports the values of T_{rx} , τ_0 , T_{sky} (including the correction for the source elevation) and T_{sys} that correspond to the entered frequency and PWV.

9.4 Total Time Estimates

Note that the time calculated by the ALMA Sensitivity Calculator does not account for telescope overheads (calibration, software and hardware latencies, etc.) and is only used to calculate the expected on-source time. However, the time estimates calculated by the OT do include these various sources of overhead and Chapter 5 of the OT User Manual should be consulted for details on how the total time required to observe and calibrate an ALMA project is calculated. The web application and OT GUI will only provide on-source times.

Chapter 10

Calibration and Calibration Strategies

This chapter describes the methods and philosophy used by ALMA in order to calibrate the correlated visibility function data. Since calibration provides a central and important part of ALMA’s production of images, this chapter contains many references to other chapters that describe other aspects of the ALMA systems. In particular, the reader of this chapter will want to be familiar with the Principles and Concepts of Interferometry described in Chapter 3 and the ALMA Observing Modes described in Chapter 8.

Most PI observations are calibrated and imaged using the ALMA Pipeline (packaged with CASA), while a small number modes are still “manually” processed in CASA using a Python generated script (see Chapter 11). In either case, an understanding of the current calibration process is needed by the PI in order to evaluate the received data package.

10.1 Fundamental Synthesis Relationship

The relationships between the visibility function $\mathcal{V}(u, v)$ and the sky emission $I(l, m)$, embodied in the van Cittert-Zernike theorem (see Chapters 3 and 7), are:

$$\mathcal{V}(u, v) = \iint A(l, m)I(l, m)e^{2\pi i(ul+vm)} dl dm = Ae^{i\phi} \quad (10.1)$$

$$A(l, m)I(l, m) = \iint \mathcal{V}(u, v)e^{-2\pi i(ul+vm)} dudv \quad (10.2)$$

The visibility function, \mathcal{V} , is thus the summation of the emission distribution of the source $I(l, m)$ (where (l, m) are its direction cosines), convolved by the exponential term that is the delay difference of the signal from the source to each antenna. The primary beam term, $A(l, m)$, describes the relative sensitivity of the antennas and is only a few arcminutes in size. The (u, v) *spatial coordinates* are not the physical separation of the two antennas, but their separation perpendicular to the source direction, since it is this difference that contributes to the phase offset between the two antenna signals at the correlator (see Figure 3.4). The visibility function (real and imaginary part) is often separated into the amplitude, A , and a phase, ϕ , since many calibrations apply to amplitude and phase independently.

10.2 Calibration Fundamentals

10.2.1 Calibrated and Measured Visibility Data

The visibility function that is obtained at the output of the correlator for any baseline antenna pair (i, j) is averaged and quantized over a short interval of time (t_k), with the output as an array of values associated with many frequency channels (ν_f). The use of the quantized form of Equation 10.2 to produce images using the discrete Fourier Transform, after the calibration and editing described in the present chapter, is discussed in Chapter 7.

10.2.2 Calibration Simplification

The complex multiplier correction required to change *measured* visibility function \mathcal{V}_m to the *true* visibility function \mathcal{V}_0 called the complex baseline gain, $G^{i,j}(t_k, \nu_f)$. It is a function of antenna pair (i, j) , time (t_k) , and frequency, (ν_f) , and is defined by the equation:

$$\mathcal{V}_0^{i,j}(t_k, \nu_f) = \mathcal{V}_m^{i,j}(t_k, \nu_f) * G^{i,j}(t_k, \nu_f). \quad (10.3)$$

Only two total intensity polarizations (XX and YY) are considered here. See Section 8.7 for the additional calibrations associated with all four (XX,XY,YX,YY) correlation products.

Although the baseline by baseline calibration approach is valid, the number of complex gain calibration terms for any time is equal to the number of the baselines, which is about 1000 for a typical ALMA observation. If every time-stamp and frequency channel required a calibration term, then the files containing the calibration solutions would become as large as the visibility data. However, there are three factorizations of the complex baseline calibration gain that are valid and generally used:

Antenna-based: Nearly all of the changes that corrupt the visibility function (e.g. atmosphere, system noise, amplitude changes, delay changes) can be decomposed into the two complex antenna-based gain factors associated with any baseline. This reduces the number of gain correction terms for an N -element array from $N(N - 1)/2$ baselines to N antennas. Each antenna calibration has considerably more sensitivity since it is determined from a type of averaging over large number of visibility baselines. This antenna-based gain separation property is called *closure*¹. Some *non-closing calibrations* are associated with certain correlator problems such as the effects of interference and large non-linear antenna-based corrections.

Amplitude and Phase-based: The corruption of the visibility function is often associated with effects that independently change the visibility amplitude or its phase. For example, the dry atmosphere can produce a phase error caused by unmodeled delay above an antenna, and also absorb the signal that causes an amplitude error. However, these corrections will have different time-scales and frequency dependence, and can be estimated and corrected independently. A counter example would be if some part of the system produced independent variations of the real-part or imaginary-part of the visibility function.

Dispersive and Non-Dispersive: Many antenna-based changes are associated with delay variations along the antenna signal paths. These calibration errors produce antenna-based phase changes that vary linearly with frequency, and are called *non-dispersive*. The linear phase-frequency relationship implies that these delay-type corrections need to be measured at only one frequency and then scaled to other frequencies. Dispersive changes are any other phase changes, for example a phase jump at all frequencies at any time or the propagation effects of the ionosphere.

The above calibration assumptions lead to the following calibration terms for baseline (i, j) ,

$$\begin{aligned} G^{i,j}(t_k, \nu_f) = & \\ & A^i(t_k, \nu_0) * A^j(t_k, \nu_0) \text{ Temporal amplitude} \\ & + A^i(t_0, \nu_f) * A^j(t_0, \nu_f) \text{ Bandpass amplitude} \\ & \phi^i(t_k, \nu_0) - \phi^j(t_k, \nu_0) \text{ Temporal phase} \\ & + \phi^i(t_0, \nu_f) - \phi^j(t_0, \nu_f) \text{ Bandpass phase} \\ & + \Delta G^{i,j}(t_k, \nu_f) \text{ Any additional correction} \end{aligned} \quad (10.4)$$

where

$A^i(t_k)(\nu_0)$ is the temporal amplitude variation of antenna i
 $A^i(t_0)(\nu_f)$ is the frequency amplitude variation of antenna i
 $\phi^i(t_k, \nu_0)$ is the temporal phase variation of antenna i
 $\phi^i(t_0, \nu_f)$ is the frequency phase variation of antenna i

The last line represents any residual non-closing (dispersive) calibration errors.

¹See e.g. Taylor et al., 1999, Synthesis Imaging in Radio Astronomy, ASP Conference Series, Vol. 180, 1999

It is assumed that the sum of all antenna signal changes produced by many different effects are linear in their amplitude product or in their phase sum. For example, the antenna signal from an amplitude change of an amplifier and that from an atmosphere absorption change will combine linearly.

10.2.3 Reference Antenna

The phase calibration for any baseline depends on the phase difference between the two relevant antennas. Thus, the same increase in the antenna phase for all antennas will not change the baseline phase correction. To bring specificity to the antenna-based phase corrections, the most common convention is to choose one antenna as the reference antenna, which is defined to have zero phase at all times over a specified set of frequency channels.²

This convention then determines unambiguous antenna-based phases with time and frequency. Care should be used in the choice of reference antenna. Any anomalous phase changes with time or frequency that affect the reference antenna will be reflected (negatively) in the phases of the other antennas. Particularly harmful is a phase jump in the reference antenna that will cause a phase jump of opposite sign in all of the other antenna phases. While not affecting the value of the corrected visibility phases, these antenna-based calibration phases are often used for system monitoring and other more sophisticated calibration methods (see Section 10.5.3), so an unstable antenna reference will make additional use of the antenna phase solutions difficult. For the more extended configurations, this concern motivates the use of a reference antenna near the center of the array where the atmospheric phase changes are likely to be some sort of average over the entire array. For a similar reason, it is preferable, if possible, to keep the same reference antenna for polarization observations and for other observations that span many hours across independent Execution Blocks (EBs), so that the antenna phases retain some continuity. However, this may not be possible in some cases.

Since the calibration amplitude corrections are multiplicative, a reference antenna amplitude, while possible, is not needed and is usually counterproductive. Finally, for calibration types for which the amplitude and phase corrections are coupled (e.g. determining cross-talk in polarization calibration), a reference antenna should not be used. See Section 8.7.

10.2.4 Calibration Measurements Using Hardware Devices

Accurate calibration parameters can be obtained from many monitoring devices associated with aspects of the ALMA system. Examples are the gain and noise measurements of the receivers and measured changes in the data path lengths from the antennas to the correlator. Environmental effects such as the delay from the water vapor above each antenna can be estimated from water vapor radiometers (see Section 10.4.3). The dry air delay above each antenna is estimated by measuring or computing the pressure, temperature and humidity from weather stations located throughout the array.

10.2.5 Calibration Measurements Using Astronomical Sources

This calibration technique is simple: a steady calibration signal from an unresolved calibrator radio source of known properties (point source, known flux density, frequency spectrum, and accurate celestial position) is observed with one or all of the antennas. The difference between the measured visibility function and that expected from the known properties of the calibrator source gives the calibration amplitude and phase as a function of time and frequency. For some calibrations (e.g. antenna pointing check or bandpass determination; see Sections 10.4.4 and 10.4.6, respectively), the precise strength or exact position of the calibrator is not needed; what is important is that it is a stable beacon.

One of the more serious calibration errors is produced by the fast and systematic delay changes in the troposphere over the ALMA array. In order to remove these delay changes, an observation method known as *phase referencing* is used. This method simply alternates observations between the science target and a nearby astronomical calibrator source (phase calibrator). The faster the switching and the closer the calibrator to the science target, the more representative will be the interpolated calibration values derived from the calibrator to that of the science target. More details are given in Section 10.4.9.

²Most ALMA frequency set-ups contain four or more independent frequency channel groups, called spectral windows (spws). The same reference antenna should be used for all of the spws in the dataset.

10.2.6 Source Catalog

To aid in the selection of calibrators for any observation, an ALMA source catalog has been compiled and contains parameters for over 3000 calibrator sources south of declination 40° . The approximate minimum flux density of a source at 90 GHz included in the catalog is about 10 mJy. Most are unresolved at long ALMA baselines, have little or no large-scale structure, and usually have a radio position more accurate than about $0.003''$. The source density is about one source per nine square degrees, so that one of the calibrators will be on average about three degrees from any random science target position.

The location and strength of the calibrator source chosen for the calibration will depend on details of the calibration purpose. For example, if a calibrator source is needed to determine the pointing or focus parameters of an antenna at any frequency and elevation, a series of calibrator observations at a range of elevations is needed. Each short observation will measure the source visibility amplitude changes with pointing direction or focus settings. In these cases, the calibrator is used simply as a constant beacon whose signal changes as the antenna is moved away from its nominal setting.

About 40 bright calibrators (called "grid sources") are observed with the 7-m Array every two weeks in order to track their variability; this is typically $\sim 10\%$ per month. While most calibrators are largely constant during this time, occasional flaring does occur. The observations are currently made at 97.5 and 343.5 GHz, and occasionally at 233 GHz. Included is one scan of a solar system object with an accurate flux density model: Uranus, Neptune, Mars (less accurate during storms), Callisto and Ganymede, Titan, Ceres, Pallas and Vesta. The resultant absolute flux density accuracy for grid sources is about $\sim 5\%$ to 10% between Band 3 and Band 7, based on a relative calibration uncertainty from measurement statistics between observations of 3% to 6% (1 sigma), and a typical 5% uncertainty in the models. The polarization variation of these quasars are also monitored and is discussed in Chapter 8.

10.2.7 Flagging Data Based on Calibrations

Some flagging of the data are made before calibrations are done, e.g. an antenna out of lock, large incorrect pointing of an antenna, or shadowing of an antenna by another. However, a significant part of the data flagging is inferred from calibration values that are anomalous in some way. This discrimination is made relatively simple for ALMA data since most calibration values should be similar among more than 40 antennas that are used for many observations.

During most forms of calibration that use a relatively strong source, the consistency of the calibration values among similar antennas is expected. Some cases are obvious, such as the loss of correlation amplitude associated with baselines that contain one particular antenna. The ALMA Pipeline task `lowgainflag` will flag antennas with gain < 0.5 or > 1.5 times the global median because such large deviations are can rarely be calibrated. Intermediate cases require assessing how well the calibrated data matches the model visibilities for the calibrator (point source). This philosophy is implemented in the ALMA Pipeline tasks `bandpassflag` and `gfluxscaleflag`, which search for outliers in the scalar difference of the corrected amplitude and model amplitude on the bandpass calibrator and the flux and phase calibrators, respectively. Antenna-based, baseline-based and time-based issues are searched for and flagged above a rigorously-developed set of thresholds. Antennas or baselines that have been fully flagged on the phase calibrator are propagated to the science target via either the caltable (for the case of antennas) or explicit flags (for the case of individual baselines). See the ALMA Pipeline documentation, available off the ALMA Science Portal, for further details. Issues that affect only a small subset of channels are not easily detectable by current ALMA Pipeline heuristics, and are usually due to hardware issues, such as errors in the correlator. One of the goals of ALMA Quality Assurance (AQUA) and Quality Assurance (QA) 0 (see Section 11.2) is to find these kinds of frequency-dependent issues earlier so that, where necessary, their cause can be investigated and fixed.

10.3 Observatory Calibrations

This section summarizes the fundamental ALMA observatory calibrations that are carried out by the observatory outside of PI science observations. Some of these calibrations are done infrequently and may only be redone if an antenna has been through a period of a major change or overhaul, or if an antenna has been moved during re-configuration of the array.

10.3.1 Antenna Primary Beam Pattern

To first order, the relative response of the antenna sensitivity as a function of angular offset from the pointing center varies with the antenna diameter and frequency. Assuming that the antenna surface and focus are accurately determined, the field of view, θ (defined as the angular width between half-power sensitivity), in radians, is inversely proportional to the number of wavelengths λ in the antenna diameter D , such that:

$$\theta = 1.13\lambda/D \quad (10.5)$$

More details on the shape and characteristics of the primary beam are given in Chapter 3 and in Baars (2003)³. The beam pattern shape is important for constructing accurate mosaic images.

10.3.2 Antenna Surface Efficiency

The large-scale and small-scale surface irregularities of each antenna surface were measured (mostly using interferometric holographic techniques on a bright quasar) and improved to meet specifications⁴. The antenna efficiencies for each ALMA band are given in Table 9.3. For the 12-m antennas, the efficiency ranges from about 70% at Band 3 to about 30% at Band 10. For the 7-m antennas, the efficiency ranges from about 70% at Band 3 to about 40% at Band 10.

10.3.3 Antenna Pointing Accuracy

The antenna pointing correction is the difference between the true sky location of the peak sensitivity of the antenna primary beam and the location shown on the azimuth and elevation encoders. Pointing offset observations are made for about 40 bright quasars over the sky⁵. Using an analysis system called *TPOINT*⁶, the all-sky pointing offsets are fit in elevation and azimuth to 19 parameters. The residual all-sky point error is about 2". There is also a *fast switching mode* in which the relative pointing difference between two sources separated by less than two degrees is about 0.6". The switching time, including settling of the antennas and electronics is about 2 seconds.

10.3.4 Antenna Focus

The emission reflected from the main antenna surface is directed toward the feed for each band using a subreflector that can be tilted and moved. The appropriate focus curves (subreflector offsets and tilt) are generated and stored for each antenna, and these remain fixed for weeks. Additional information is given in Appendix A.

10.3.5 Antenna Integration Calibrations

Scheduled array configuration changes are made approximately every 3 to 4 weeks, during which typically about 8 antennas are moved over a period of about 10 days. After a change in the configuration of the array, a set of *antenna integration* calibrations are carried out before a recently moved antenna is integrated back into an array used for PI science.

Improved Antenna Pointing

Additional pointing observations are needed after an antenna has been moved to a new pad. Of the 19 total fitted pointing parameters, five are expected to change: the azimuth and elevation offset, two antenna tilt terms and one overall rotation term. This is part of the antenna integration process and uses a short *all-sky-pointing* observation, described in Section 10.3.3. These short observations are needed to obtain the a nominal pointing accuracy over the sky of about 2".

³<http://legacy.nrao.edu/alma/memos/html-memos/alma456/memo456.pdf>, (ALMA Memo 456)

⁴Mangum et al, 2006, PASP, 118,1257

⁵<https://safe.nrao.edu/wiki/pub/ALMA/CalExamples/PointingCalStepByStep.pdf>

⁶<http://www.tpointsw.uk/index.htm>

Antenna Offset Delay and Focus

The signal path length between an antenna (excluding the variable delay of the source signal to each antenna during an observation) receiver to the correlator will change considerably after an antenna is moved. If these offset path length differences are not removed, the correlated visibility between antennas will show a large phase slope over the frequency band. For example, if the signal path length from an antenna to the correlator is offset by 10 cm, this will produce a phase slope of 100° over a 2 GHz spw and produce significant decorrelation over the the spw. The delay offset is measured by a short observation of a strong quasar using a reference antenna that has not been modified or recently moved. A more accurate removal of a delay offset and the frequency characteristics are determined from the bandpass observation in most PI observations.

The focus parameters associated with an offset in the azimuth and elevation with a possible small rotation are also checked during the antenna integration process.

Antenna Positions

In order to make high quality images using phase referencing techniques, the relative antenna locations must be known to an accuracy of 0.3 mm. This limits the systematic phase errors between the calibrator and the target during phase reference observations. If uncorrected, the error produces distortions in the target image.

From Equation 10.1, the maximum phase difference, $\Delta\phi$ in degrees, between a calibrator and a science target, that is produced by an error in the antenna position, ΔB (mm), is given by:

$$0.02 \times \Delta B \times \nu \times Sep = \Delta\phi \quad (10.6)$$

where Sep (deg) is the calibrator-target angular separation and ν (GHz) is the observing frequency.

Occasionally, an initial antenna position error may be as large as 5 mm. In Band 7 (350 GHz), for example, this would result in a phase offset of 50° between the calibrator and target even in the fortunate situation of finding a phase calibrator within 3° of the science target. This offset remains relatively constant over one hour, so its effect on the image quality would be worse than that of short-term atmospheric variations of comparable size.

The accurate antenna positions are determined from a sequence of many short observations of bright quasars that are observed over the sky, called *all-sky delay* observations. Many strong quasar calibrators are observed, each for 5 seconds, with about 40 seconds needed to move between each quasar, often more than 40° away in order to sample the entire visible sky in about 20 minutes. The observations are made at Band 3, where most quasars are strongest, and using four spws in TDM mode for a total bandwidth of 8 GHz. During the antenna integration process at the beginning of a configuration or when an antenna is moved and is added back to the array, the observations may be limited to a small number of antennas over about 30 minutes to obtain about 1 mm accuracy before PI observations are started.

However, near the end of an antenna move period when the array pattern has been set, a more accurate all-sky delay observation is run, often using the entire 12-m Array to improve the SNR of the solutions. For a typical observing time of about 45 minutes, the accurate antenna-based delay for about 50 scans are obtained, and the systematic delay for each antenna versus azimuth and elevation are fit to an improved antenna position, often with an accuracy of < 0.1 mm. An analysis script in the TelCal system, called *tc_antpos*, is used to determine the fit. An example of the fit for one antenna of the scan delays before and after the correction of the fitted residual antenna offsets is shown in Figure 10.1.

After the all-sky delay run, the updated antenna positions are placed in the online Telescope Monitor and Control DataBase (TMCBD) that is used by the delay server. The updated antenna positions are also currently placed in a file which can be accessed by the ALMA Pipeline and other offline reduction systems to correct the observations that were made before the better antenna positions were updated in the system.

The antenna position uncertainties depend on the distance from the array center (the reference antenna is usually chosen near the array center), because the atmospheric phase variations, caused by large-scale turbulence, generally increase with baseline length. The typical antenna position errors are about 0.05 mm per kilometer of baseline length. Hence, at the longest baselines, a typical antenna position uncertainty could be as large as 1 mm.




Figure 10.1: **Antenna Position Determination:** The antenna position fit for antenna DA46, at 3.5 km from the reference antenna, from a 50-minute all-sky-delay observation on Dec 25, 2017. Each of the 45 points show the delay solution from one scan, separated by about 1 minute. The observations were made at Band 3 where a delay change of 0.001 ns is equal to a phase change of 40° . The red points show the original delay measurements and the scatter is caused by antenna position offsets as a function of quasar elevation and azimuth, The black points show the delay after correction of the DA46 antenna position error by about 4 mm. After correction, the spread of phases between sources over the sky is about only 20° , so the expected phase difference from a residual antenna position offset between a target and calibrator that are just a few degrees apart will be considerably less.

10.4 Execution Block Calibrations

This section describes how an ALMA science Execution Block (EB) is calibrated. Each EB contains the minimum set of necessary calibration observations, which are automatically defined in the Scheduling Block (SB).

Several EB's (each limited to about 1 hour integration on a target) may be needed to obtain sufficient SNR or (u, v) coverage to meet the science goals as defined by the PI. When there are multiple EBs, the calibrator sources observed may differ between EBs if, as in most cases they are selected at run-time by the SSR query algorithm (see Section 8.3.1); only in a few cases will calibrator sources be fixed in the SB by ALMA staff. The specific antennas in the array may also change between EBs. It is useful to keep the same reference antennas when calibrating each EB even though the EBs are calibrated separately. The quality of the EBs may differ, mostly because of the different phase stability conditions. The AQUA QA0 and QA0+ assessments (see Chapter 11) are carried out after each EB is observed in order to assess the quality of the data.

10.4.1 Making the Measurement Set

Data from ALMA are stored at the ALMA site, at the JAO in Santiago, Chile, and mirrored to each of the ALMA Executives in North America, Europe and East Asia (See Chapter 13). These data are transformed into an ASDM (ALMA Science Data Model) data structure using *asdmExport* (see Chapter 12). This format is then converted into a measurement set (ms) using the CASA task *importasdm*. These steps will not be executed by most PIs, but could be useful for special tests and calibrations in certain cases.

10.4.2 Amplitude Calibration Procedure (System Temperature Measurement)

An important ALMA online calibration determines the antenna receiver sensitivity and the radiative sky temperature. These corrections are obtained from a 15-second scan where a hot load, ambient load, and sky are alternately placed in front of the feed. Both 12-m and 7-m antennas are equipped with these Amplitude Calibration Devices (ACDs). The design, use and accuracy of these ACDs are discussed in Section A.5.1.

At frequencies below about 400 GHz, where the system temperatures are stable, an ACD scan is made every 10 to 20 minutes. However, at higher frequencies, and wherever the opacity is large and more variable, all scans will have an associated ACD measurement. The online measurements of T_{sys} and T_{rx} are stored in the CalAtmosphere table in the ASDM datasets, and are displayed in real-time at the telescope to detect problems immediately for QA0 assessment (see Section 11.2).

An example of the T_{sys} measurements from a Band 9 observation is shown in Figure 10.2. The two antennas show the typical range of the T_{sys} measurements at this high frequency. Notice that the T_{sys} in each spw contains narrow lines that are associated with tropospheric emission. For example, the additional line emission of 20% in the first spw near channel 60 is the emission from the ozone line at 673.914 GHz.




Figure 10.2: **System Temperature at Band 9:** The measured system temperatures at Band 9 for a target source for two antennas with different properties. Top row is for DV05 for four spws, bottom row is for DV09 for four spws. The ordinate is the system temperature, the abscissa is the channel number in each spw. The average frequency for each spw is 673.9, 675.9, 677.9, 679.9 GHz, and the channel separation is 1.56 MHz. The two polarization T_{sys} differs significantly for DV05. The finite width of each T_{sys} curve is caused by the small time variation over the 50-minute EB.

10.4.3 Water Vapor Measurements

Variations from water vapor (Precipitable Water Vapor; PWV) in the antenna line-of-sight cause significant delay variations, up to 0.5 mm/second, corresponding to 50° of phase/second at 90 GHz. The changes are driven by wind turbulence and cloud structure, and cover a large range of spatial scales from 5 meters to many kilometers. All 12-m ALMA antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures the emission from atmospheric water vapor every 1.1 seconds at four frequencies near 183 GHz, in nearly the same direction as the ALMA pointing. The phase correction method using the radiometers was developed by B. Nikolic et al. (2013)⁷, and information about their design and use is given in Section A.6.

⁷Nikolic, B. et al., 2013, A&A, 552, 104

These WVR corrections are made by the online calibration software system in TelCal⁸. An offline script using the CASA task `wvrgcal` also converts the water vapor emission to a predicted antenna delay. At present, only the uncorrected visibility data are archived and the WVR correction is applied offline as part of the calibration process. In addition to allowing a global solution (across the observation) for the brightness temperature to delay conversion coefficients, the offline application of the delay correction enables small edits or scaling of the WVR emission to delay to be made if needed⁹. The online corrected data, however, will be available for use of the AQUA and QA0 (Chapter 11) analysis of each observation.

The WVR correction often reduces the water vapor-induced phase fluctuations on time scales of 2 seconds to many minutes by 50% or more, enabling more observations to meet the rms and image quality success level. An example of the antenna-based phases at Band 9 for an antenna before and after WVR correction is shown in Figure 10.3. The PWV was ~ 0.45 mm for this example, and it is the typical improvement for relatively low WVR content. The ALMA Pipeline task `hifa_wvrgcalflag` determines the WVR correction, generates flags for WVR data that are out of range or noisy, and automatically interpolates from up to three neighboring antennas (if near enough).




Figure 10.3: **WVR correction:** The improved temporal phase by using the WVR correction is shown in the four panels for different antennas at Band 3 in the long baseline configuration over a period of three minutes during a strong bandpass source observation. The red lines show the temporal phase before correction; the green lines after correction. The short-term variations are removed as is most of the longer term drifts. The offsets from zero phase are arbitrary.

10.4.4 Antenna Pointing Improvement

Additional pointing observations are sometimes included in an EB (see Section 8.2). For observations > 350 GHz, the pointing between sources more than about 20° apart may differ by more than $2''$ from the nominal pointing solution. For observations that demand the highest pointing accuracy, for example $< 0.5''$, a pointing

⁸<http://adsabs.harvard.edu/abs/2011ASPC..442..277B>

⁹L. Maud, <https://arxiv.org/abs/1707.03506>

every 10 minutes on a source near the phase calibrator or science target can be done (known as pointing referencing).

10.4.5 Antenna Position Improvement

The determination of the most accurate antenna positions (typically 0.2 mm accuracy), from a comprehensive all-sky delay observation, usually occurs after all antennas have been moved during an array re-configuration. The shorter observations made just after one or more antennas have been placed on a new pad (see Section 10.3.5) are accurate to about 1 mm. PI observations executed before the most accurate antenna positions have been made are not processed by the ALMA Pipeline until the final position updates are placed in an offline file which is used during calibration to apply the most accurate antenna position corrections. However, the AQUA and QA0 analysis of these observations are not impacted by any remaining $\approx 1\text{mm}$ antenna position offsets.

10.4.6 Bandpass/Frequency Calibration

The goal of the bandpass calibration is to determine the *relative* amplitude and phase of the frequency channels across each spw. Because of the excellent temporal stability of the ALMA bandpass (Kameno 2015)¹⁰, the *bandpass calibrator* is observed only once at the beginning of an EB. It is generally a bright quasar within about 30° of the science target, and is chosen from the list of grid sources (see Section 10.2.6). The scan length is between 2 minutes and 30 minutes and depends on the strength of the calibrator and the bandpass accuracy needed.

The ALMA online software selects a bandpass calibrator based on the latest flux measurement, so that a SNR of 50σ per antenna and 128 MHz bandwidth can be achieved. If a suitable calibrator cannot be found, the online software may relax the threshold down to a SNR of ~ 14 . With these criteria, spectral dynamic range can be calculated as $\text{SNR} \times \sqrt{2} \times \sqrt{N(N-1)/2} \times \sqrt{\Delta\nu/128 \text{ MHz}}$, where a factor of $\sqrt{2}$ is arisen by two orthogonal polarizations, N is the number of antennas, and $\Delta\nu$ is the desired frequency resolution. For example, we get a spectral dynamic range of ~ 830 with SNR of 50σ , 20 km s⁻¹ velocity resolution at observation frequency of 300 GHz, and 43 antennas. Please see Table 2 of the ALMA Cycle 8 2021 Science Primer (available off the Science Portal) for the relation between velocity resolution and channel width.

There are two steps in obtaining the relative bandpass calibration. First, the temporal phase variations over the bandpass scan, (see Equation 10.4) is determined by averaging the bandpass source over most of each spw, for each time stamp t_k . An example of these phase changes a scan is shown in Figure 10.4.

After temporal phase variations over the bandpass scan are removed, the second step determines the bandpass amplitude and phase dependence for all channels in each spw/polarization. The bandpass solutions are usually normalized to average amplitude 1.0 and phase 0.0 for each spw. An example showing the bandpass solution for four antennas at Band 9 is given in Figures 10.5 and 10.6. This typical bandpass shape is relatively flat to about 3% and 8° over the 2 GHz width, often with a channel to channel rms of about 1% and 1°. Occasionally (however, not in the example shown), there may be a significant phase slope that was not removed during the antenna integration calibrations which will be removed with the bandpass calibration. The 200 MHz-width ripples in the bandpass have several possible causes related to small signal path reflections in the antenna structure and in the electronics. They are stable with time, however. Finally, at the edges of many spws, significant amplitude departures occur and are generally flagged. In this example, the ALMA Pipeline task `hifa_bandpassflag` flagged channels 0-9 and 123-127 for all antennas/spws.

10.4.7 Flux Density Scale

The correlated visibility amplitude scale obtained from the correlator output is the ratio of the correlated signal with the total system noise, and is usually much less than about 1%. Since each antenna receiver noise level plus the emission from the atmosphere above each antenna are accurately measured in units of their effective temperature by the ALMA calibration devices, the percentage correlation is converted to Kelvin (K) by application of the system temperature (Section 9.1.6).

The most straightforward conversion of the visibility amplitude from units of K to units of Jansky (Jy) is to include in the EB an observation of a source with a known flux density. Since most quasars calibrators are variable (typically 10% per month, with the stronger ones more variable), the inclusion of a solar system object

¹⁰https://almascience.nrao.edu/documents-and-tools/documents-and-tools/alma-technical-notes/ALMATechnicalNotes15_FINAL.pdf




Figure 10.4: **Bandpass Phase-up**: The phase variation for a Band 9 bandpass calibrator. (Top) Xpol and Ypol phase and their difference for 675 GHz spw, (Middle) Xpol and Ypol phase and their difference for 679 GHz spw, (Bottom) 679 GHz - 675 GHz spw for Xpol only and their difference. The WVR correlation has been applied to all data. The phase variations over time are extremely well correlated between polarization and spws, and these phase offsets are virtually constant over an EB.

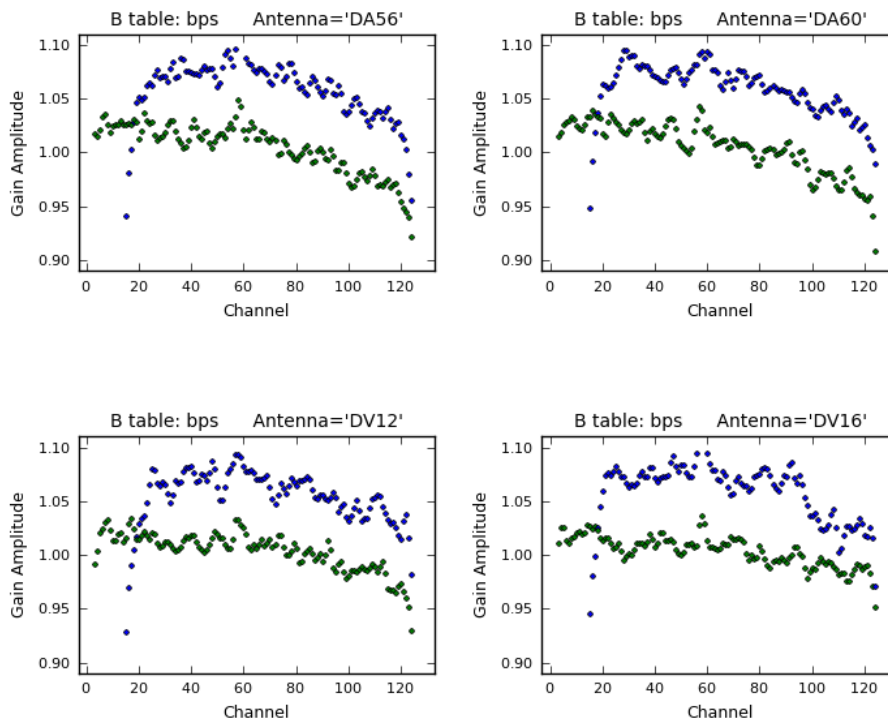


Figure 10.5: **Bandpass Amplitude:** The bandpass amplitude at Band 9 for four antennas, Xpol are shown as blue dots and Ypol as green dots. Xpol channels below 10 and Ypol channels above 120 should be be flagged. Notice that the line emission near channel 60, seen in the T_{sys} solution (Figure 10.2), has been significantly reduced after multiplication of the raw visibility amplitude by the T_{sys} factor.

(see Section 10.2.5) at the higher frequencies is recommended. The flux density models are determined between 30 and 2000 GHz¹¹, and are accurate to about 5% at the lower frequencies. This increases at frequencies above Band 7, however, since the spectral extrapolation of the grid source flux densities is not accurate; in Bands 9 and 10, the conversion from K to Jy is accurate to about 20%. Thus, one scan of an object can determine the K to Jy factor. Since the solar system objects all have significant angular sizes so the K to Jy factor determination should be limited to the shorter baselines. With the success of the grid source monitoring observations every 10-14 days to follow their variable flux density, most PI observations now include one scan of a grid quasar to use as both the bandpass and and flux density calibrator.

Assuming that the conversion of the visibility amplitude from percentage correlation to Kelvin is accurately made, the further conversion from K to Jy can be estimated solely from the antenna effective area at the observation frequency. The effective area is equal to the antenna surface collecting area times the antenna efficiency (see Table 9.3 and Section 10.3.2 for details of the efficiencies of the ALMA antennas).

For frequencies up to 400 GHz (12-m antenna efficiencies between about 60% and 70%; 7-m antenna efficiencies between about 65% and 70%) this leads to an SEFD (Source Equivalent Flux Density) = 35.0 K/Jy for the 12-m Array and an SEFD = 96 K/Jy for the 7-m Array. The SEFDs increase at the higher frequencies. However, these conversion values may be more accurate than those from quasars and solar system objects, especially at the higher frequencies (Section 9.2.1). This SEFD flux calibration method is currently being tested.

10.4.8 Amplitude Temporal Variations

The amplitude scale factor determination from K to Jy is obtained only at the time of the scan used to determine this ratio. At low elevations and for frequencies above 350 GHz, there are small temporal variation of gain with

¹¹<https://library.nrao.edu/public/memos/alma/memo594.pdf>




Figure 10.6: **Bandpass Phases:** The bandpass phase at Band 9 for four antennas, Xpol blue dots and Ypol green dots.

time related to antenna efficiency changes with elevation or from the differential heating of the antenna surface support legs or small systematic pointing changing. These changes can produce amplitude changes up to 10%, especially during daylight hours. If the phase calibrator used (see Section 10.4.9) is relatively strong, then its flux density can be determined from comparison with the bandpass strength. Then, during the observation, the amplitude solution part of the phase-referencing solution can be used to estimate the temporal gain variations over the observation.

10.4.9 Phase Referencing Calibration

One of the most important calibrations needed for obtaining good quality images is *phase referencing calibration*. It is usually the final calibration after the others, discussed above, have been made. The temporal amplitude variations during most PI observations, discussed in Section 10.4.8, are rarely as critical as the phase calibration. The EB structure consists of many cycles of visits to the phase calibrator, with visits every ~ 80 seconds for the most extended configurations or up to every 600 seconds for the most compact configurations.

The first reduction step uses the CASA task *gaincal* to determine the antenna-based phase correction from all phase calibrator scans. The phase solutions are usually made separately for each spw and polarization, but they all have virtually the same temporal variation since the instrumental phases between the spws and polarization are steady to a few degrees over an hour or more. For this reason, when the phase calibrator is relatively weak (less than 30:1 SNR per antenna solution), the spw and polarization channels can be summed after the phase offset is removed to obtain one temporal solution for all spws and polarizations.

The second step of phase referencing interpolates these antenna-based phases to the other scans in the EB. Some interpolation options are discussed below. As discussed in Section 10.2.3, the choice of reference antenna is important, and a non-flagged and stable antenna near the array center should be chosen, even if several iterations of phase calibration are needed to determine this a satisfactory reference.

Some comments on efficient and accurate phase referencing are described below.

Phase Transfer from Calibrator to Target(s): The simplest method is to linearly interpolate the calibrator scan-averaged antenna-based phases between two consecutive phase calibrator scans to that of the science target or other sources. This is done with the CASA task *applycal*. More complicated interpolation methods including the use of several phases over a scan, or several scans over a longer period of time, are possible, but the simple linear interpolation is usually satisfactory.

Calibrator Cone Search and Detectability: An *acceptable* phase calibrator for a given science target is the closest source that is above the necessary SNR limit and point-like. For most science targets, the best calibrator choice from the ALMA catalog is clear. In some cases, often if one or several possible calibrators have $SNR < 100$ and have no flux density measurement in the last three months, calibrators are included in a short flux-checking observation, called a *cone-search*, that is made about two months before the science observation is to take place. Checks for an extended structure or resolution of the calibrators are also made. Cone-searches are also used to select the appropriate check source for an observation (see Section 10.4.10).

Phase Calibration Sensitivity: The strength of the phase calibrator should be sufficient to produce a scan solution per antenna with a minimum antenna-based SNR of 15. This SNR leads to an antenna-based solution of a 7% amplitude error and 4° phase error. With the typical number of antennas and scans used for most PI observations, this noise-like calibration error will average over the science target scans and antennas, and limit the image accuracy to $< 0.01\%$ of its peak intensity if it is not otherwise noise limited.

Combining spws: Since the phase difference among the spws and polarizations are relatively constant over one to two hours, the calibration sensitivity can be increased by a factor of 2.8 if visibility data from the four spws and both polarizations are first combined. The phase differences among the channels can be determined from the bandpass observation (Sec 10.4.6), and are stable to a few degrees.

An example of the phase calibration of two antennas from Band 9 observations for a 1-km configuration is shown in Figure 10.7. For nearly all ALMA calibration, the X- and Y-polarized phases follow each other to a few degrees since nearly all of the variations are caused by tropospheric delay changes over each antenna. A good check on the signal-to-noise of the calibration scan and other problems is given by the X-Y phase difference easily obtained from CASA task `plotcal` (`poln='/'`). The difference should be relatively flat with variations that are consistent with the expected rms noise. Hence, unstable periods of calibration from an antenna can be easily recognized with the appropriate data flagging made.

In some cases, a close inspection of the phase referencing results may show suspected antennas or time periods where additional flagging may be useful. Some considerations include:

Scan Decoherence: When the phase variations over a calibrator scan are large, there is considerable decoherence of the visibility data. Such decoherence can be verified for the stronger phase calibrators using solution times less than the scan length to see the fast phase changes. This decoherence is larger at longer baselines and effectively decreases the resolution of the observation. No corrective action is recommended. It is better to acknowledge the decoherence, rather to make amplitude corrections that often depend on the relative strength of the calibrators and the targets.

Calibrator Image: An image of the phase calibrator is provided as part of the calibration products. It should be a high quality point source (after cleaning) with a peak intensity and integrated flux density that agrees with the flux density assumed in the calibrations. The residual sidelobes or rms noise should be a small fraction of its peak intensity. If the phase calibrator image does not meet these expectations, then the calibration and editing process is suspect and should be re-investigated.

10.4.10 Check Source Observation

For any observation with maximum baseline length greater than 5 km, and/or at frequencies greater than 400 GHz, *check source* scans are included in an SB. The criteria for choosing a check source from the ALMA catalog are that the calibrator should be: (1) about the same distance from the phase calibrator as the phase calibrator is from the science target and (2) strong enough to produce an image with $SNR > 15$ using all scans and spws. This flux limit is about a factor of three less than that for the phase calibrator used for the observation. As with the phase calibrator selection, most chosen check sources are unresolved even at the longest ALMA baselines,




Figure 10.7: **Calibrator Phases:** The calibrator temporal phase at Band 9 for two antennas. (Left) 1.0 km from the reference antenna. (Right) 0.3 km from the reference antenna. Over the 4000-second EB, 10 phase calibrator scans were made, each of 60-second length with three 20-second points connected to show the internal scan phase variation. The X and Y polarizations show virtually the same temporal changes, as expected. Within each scan, the phase change is about 20° . Between scans separated by 7 minutes, the difference is about 50° and is a little larger for the longer baseline. The X-Y phase difference (green (points/line) shows the approximate SNR of each solution, and is much less than the atmospheric variations. The offset is arbitrary for clarity.

but exceptions do occur. The check source is typically observed two or three times per EB with a scan length the same as the phase calibrator scan.

The check source is sufficiently strong and its image is routinely made in the QA0+ analysis procedure after each execution. An image of the check source is also made by the ALMA Pipeline during data processing. An analysis of the check source image quality from many long baseline observations in 2017 suggests that its peak to total flux density ratio is a good indication of the observation quality (see Chapter 11). This ratio varies with the rms phase difference between phase calibrator scans, and with its angular separation from the phase calibrator. These check source distortions (including a position offset from its known position) are a good indication of the target image quality which is often too weak or too extended to be used on its own to estimate the quality. Generally, the science target distortion and its position offset is similar to that of the check source, but approximately scaled by their relative distances from the phase calibrator.

10.4.11 Estimate of Image Quality in QA0

An initial evaluation of an EB's quality is made within hours of an observation, and is discussed in Section 11.2. While many observing and data properties affect the image quality, it is significantly affected by the properties of the tropospheric phase variations over time and over the array, and these properties are readily obtained from the phase calibration (Section 10.4.9). For that reason, a QA0 metric, based on phase variations in phase-referencing is described here.

The tropospheric phase variations, to first order, decrease the effective resolution of an observation because the longer spacings have more phase variations than that of the shorter spacings. For all sources, this will decrease the peak flux density by smearing out the structure, and is often called *phase-referencing coherence*. Using phase-referencing observation tests over a range of tropospheric phase variations and calibrator-target angular separations, the expected coherence, as a function of calibrator phase rms between scans and a function of the calibrator-target angular separation, are illustrated in Figure 10.8. These results are relatively independent of the maximum baselines and observing frequency.

The explicit peak/integrated flux densities in the following roughly equate to the QA0 status described in Chapter 11. If the peak/total flux density of the object is > 0.85 , the observation is considered to be satisfactory, i.e. QA0 Pass. For example, from the tests, a calibrator-target separation of $< 2^\circ$ with a scan-to-scan phase calibrator rms of 45° should provide a good quality image. A calibrator-target separation of 5° needs a phase




Figure 10.8: **Phase Referencing Coherence**: Simulation results to determine the relationship between image quality and phase referencing parameters. (Top): The peak/integrated image quality versus the PhaseCal-Target separation. Curves for different PhaseCal scan phase rms's are shown. (Bottom): the peak/integration image quality versus PhaseCal scan phase rms. Curves for difference PhaseCal-Target angular separations are shown.

calibrator rms of $\sim 25^\circ$ in order to pass. For a Peak/Total < 0.6 , the observation should probably fail and be re-observed, (QA0 Fail). Thus, for a calibrator-target separation of 5° , the rms scan phase difference should be less than about 40° to pass. The intermediate quality images with peak/integrated between 0.6 and 0.8 generally have a QA0 Semi-pass designation and QA2 processing is needed to determine if the EB is successful. Of course, the image quality using the shorter baselines only will produce a more coherent image, but see Section 11.2 for the QA0 rules for the number of antennas with sufficiently low phase rms that can be used for QA0 Pass.

When there are multiple EBs in an MOUS, the image quality metrics should be compared for the check source (if available) and the science target for each EB. Sometimes, the image quality/phase stability will differ significantly among the EBs, but the image weighting does not consider the phase coherence properties. Hence, down-weighting or even removing a poorer quality EB could be investigated to determine which set of EBs produces a better quality image.

In addition, the position registration among the EBs is best evaluated by comparing the check source position peak among the EBs, or from a strong target compact feature if one exists. Offsets at the level of 0.2 times the resolution are reasonable and of no concern. Larger offsets between the peak emission may be seen if one or more EBs has relatively poor phase conditions and if the target is somewhat complex. If some EBs have good agreement in the check source position while others do not, an image made from the first EBs can be used as the initial model image for self-calibration (see Section 10.5.1) of the other EBs.

10.4.12 Calibration for Specific Observing Modes

The above sections apply to most ALMA interferometric observing modes (see Chapter 8) including, but not limited to, single or multiple EBs, single pointings or mosaics, non-moving or moving (Ephemeris; 8.9) targets, single or multiple configurations or single or multiple sources per EB. However, there are certain observing modes where special calibration is needed at the time of observation. These are described below.

Polarization Calibration

For ALMA polarization observations, the current calibration scheme, systematic calibration uncertainty, and imaging of the four Stokes parameter data, $I, Q, U, (V)$, are described in detail in Section 8.7. After the normal XX and YY gain calibrations are made, three additional calibrations are needed: determination of the X and Y relative delay in the receiver analog path, the determination of the cross hands phase (XY -phase), and the determination of the “D-terms”, or the lack of purity in the polarization decomposition into a linear base X and Y . Most polarization observations require two or three EBs to cover a sufficient range of parallactic angle ($\geq 60^\circ$).

Total Power Calibration

Details of the Total Power observing mode, including the current calibration scheme, are described in Section 8.6.

Solar Calibration

Details of the Solar observing mode, including the current calibration scheme, are described in Section 8.10.

10.5 Additional Calibrations

This section contains suggestions for additional analyses and calibrations beyond those normally done for standard ALMA data reduction. Depending on the science goals, the observing conditions, and the complexity of the science target image, further processing by the PI may improve the image quality. Several common techniques are described, but additional information and support can be obtained from the local ARC.

10.5.1 Self-Calibration

The phase referencing technique makes periodic observations of a known point source to measure the temporal antenna-based amplitude and phase calibrations during an observation. These calibrations are ‘transferred’ to the target from which an image is made. If the target is bright and sufficiently compact, further calibration

of the target data can be obtained because of closure property of interferometric data (Sec 10.2.2). For the 12-m Array, any further calibration of the ≈ 1000 baseline visibility functions can be made by determining the approximately 50 antenna-based corrections. This redundancy permits an iterative process that 1) improves the visibility data calibration, and 2) improves target image accuracy.

A description of the philosophy, criteria and methods for self-calibration (hereafter, self-cal) is given in Brogan, Hunter & Fomalont (2018)¹². Because of its iterative nature and the need to have sufficient SNR of the target over an reasonable solution period, self-cal methods vary depending on many properties of the data. Based on current tests (QA0+ automatic attempts to self-cal ALMA observations; see Section 11.3), it is estimated that for about 25% of ALMA science targets, self-cal can improve the quality of the image, sometimes by a factor of five or more as measured by the peak flux density to the rms noise in the image.

Suggested Self-Cal Steps with an Example

The first step towards successful self-cal is to determine if self-cal is possible for a given science target.

Is self-cal improvement likely? As a rule of thumb, the peak/rms of the initial phase-referenced image should be at least 20. Whether the initial image is from the continuum data or a spectral line feature does not matter, but the image/data link must be maintained.

First try: For the initial attempt to determine if self-cal improvement is likely, try a phase-only self-cal on the normally calibrated target data with the cleaned target model in the visibility model column¹³. A solution time of one scan (one to five minutes) is a reasonable first step. A good indication of the self-cal success is to compare the X or Y phases versus scan for many antennas using `plotcal` with `poln = "` and `poln = '/'`. An example is given in Figure 10.9.

The next steps are typical if the previous step is successful with good SNR (small X-Y phase difference).

Further phase-only self-cal: Repeat the phase-only self-calibration with a more optimum `solint` interval. If the continuum is used, combine the spws using the bandpass source to align the phases to improve the SNR if a solution is needed. Make sure that the improved image model is used for the next iteration. After application of this phase self-cal, clean the target relatively deeply, especially for a target that contains both large-scale and small-scale emission. It is better to include as much flux in the clean components as possible, even with some increased noise components. If in doubt, keep cleaning around the known emission regions. Another phase-only self-cal iteration can be considered, perhaps with a shorter solution interval.

Amplitude and phase self-cal: If the peak/rms ratio is more than about 100, additional self-cal iterations that include an amplitude correction is recommended. Use a model from the previous phase-only self-cals that have been cleaned deeply so that most of the flux density is in the model image. The usual timescale of intrinsic amplitude changes are about 5 to 10 minutes so use this solution interval even if the phase self-cal scale is much less. Use the option `calmode='ap'` rather than `calmode='a'`. This will produce a better solution in the amplitude stability even though the phase corrections should be minimal. The amplitude self-cal solutions often show outlier points associated with brief antenna drop-outs that should be flagged. For strong sources, a peak/rms > 1000 is often possible but requires several additional self-cal iterations with careful editing of the data.

The self-cal example shown here was from a one-hour observation in November 2017 at 250 GHz using 43 antennas with a maximum baseline length of 11 km. For this example, the image improvement from the normal phase-referenced image to that made after two phase self-cals with a solution interval of 18 seconds (1/3 of a target scan) is shown in Figure 10.10. For additional details, see Brogan, Hunter & Fomalont (2018).

¹²Brogan, Hunter & Fomalont, 2018, <https://arxiv.org/abs/1805.05266>

¹³Make sure the model column exists and contains data corresponding to the image clean components, otherwise use the task `FT` on the clean model to generate this model column data.




Figure 10.9: **Self-cal Phases:** The self-cal phases for four antennas. The blue points are the X-pol self-cal phases. The green points are the X-Y pols self-cal phases, offset by 100° for clarity. The scan length was 55 seconds and the frequency was 240 GHz. The X-Y difference is a function of the solution SNR and shows that the 'X' or 'Y' self-cal phases are not noise limited.

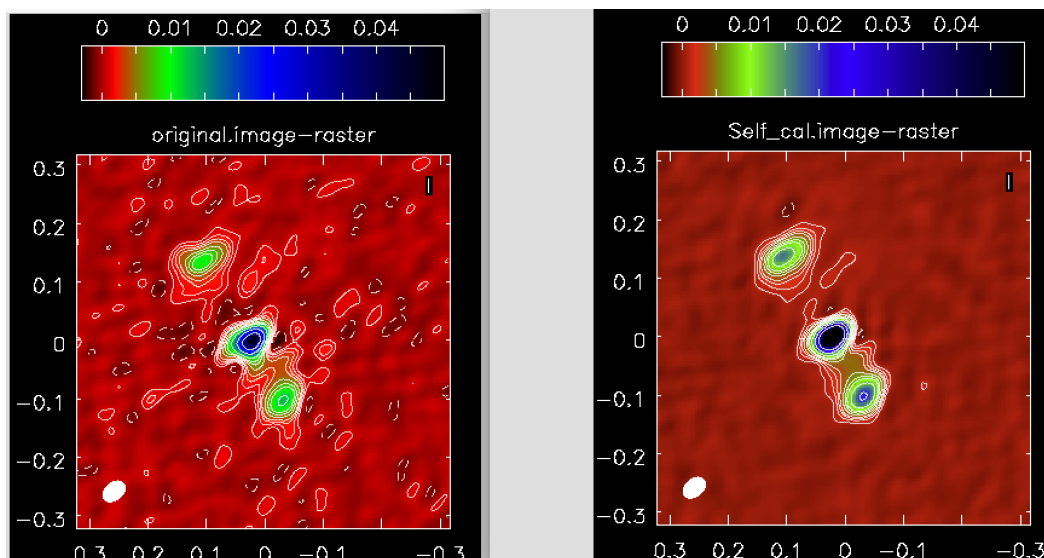


Figure 10.10: **Self-cal Image Improvement:** (Left): The image derived from the normal calibration using 48 12-m antennas at 227 GHz, with a phase calibrator–science target separation of 3.8° and a cycle time of 70 seconds. (Right): The image after two self-cal phase iterations, with a solution time of 18 seconds (1/3 scan). For the normal and self-cal images, respectively, the peaks are 48 and 80 mJy and the rms noise values are 0.35 and 0.22 mJy, respectively. This gives a peak/rms of 137 and 364, respectively, or about a factor of three improvement. The theoretical noise level for the integration time and frequency is 0.20 mJy.

10.5.2 Astrometric Observations/Calibrations

General Information

An astrometric-type project is one where measuring the accurate celestial position of the science target is one of the main scientific goals. Examples are: identifying the radio emission associated with that of an optical target with a positional accuracy better than the nominal one or measuring the positional change of a target over a period of hours to years using appropriate multi-epoch experiments.

The choice of the observing Band, array configuration(s), and target integration time depends on the properties of the science target: angular size, flux density and spectrum, expected celestial motion and the precision needed. The choices may be influenced by the availability of configurations during the cycle. Although the more extended configurations and higher frequencies might potentially provide more accurate astrometric precision, these astrometric gains are often outweighed by the limited time available at high frequencies, due to weather, and at configurations >3 km with good phase conditions.

Beginning with Cycle 8, an astrometry panel is present in the Calibration Setup editor of the ALMA Observing Tool. As discussed in the Cycle 8 2021 Proposer’s Guide, Section A.9.5, if the astrometric goal can be achieved using a standard imaging experiment, as mentioned below, then no action is required by the PI and the default option should be selected. Otherwise, see "Enhanced Positional Accuracy" below.

Nominal positional accuracy: In standard experiments the positional accuracy is limited by the SNR of the image and the configuration used. The relationship between the rms positional accuracy, the resolution, and the peak SNR of the target is, $\text{pos}_{\text{acc}} = \text{beam}_{\text{FWHP}} / \text{SNR} / 0.9$ where pos_{acc} is the position rms accuracy on the target, $\text{beam}_{\text{FWHP}}$ is the Full Width Half Power Resolution, SNR is the SNR of the image target peak, and the 0.9 is a factor to account for a nominal 10% signal decorrelation.

In nominal ALMA observing conditions, the atmospheric phase fluctuations limit the $(\text{pos}_{\text{acc}} / \text{beam}_{\text{FWHP}})$ ratio to 0.05 at best so that an image with $\text{SNR} > 20$ will not improve the target positional accuracy, and a smaller SNR will degrade the accuracy. Hence, it is suggested that the spectral energy distribution (hence observing frequency) and source structure (if any) of the target is considered, so the requested sensitivity at the representative frequency allows for an SNR ratio of 20. If the synthesized beam is finer than $0.15''$,

the positional accuracy is a factor of ~ 2 poorer than the one given by the above equation. This is induced by the phase stability degradation with increasing baseline length.

Enhanced Positional Accuracy: Enhanced positional accuracy is suggested for PIs who, need a positional accuracy better than nominal, have a problem with the distance to the default phase calibrator, or need multi-epoch observations to determine target motion.

In case “Enhanced Positional Accuracy” is selected, a check-source will be added by default to the observations to be used for assessment of astrometric accuracy and phase transfer quality. Both the phase-calibrator and check-source will be automatically selected at the start of the observation. In case the PI decides to identify and fix their own choice of calibrators, that is possible via the “user-defined calibration” in the top panel of the “Calibration Setup” editor. Instructions to do so are described in the Cycle 8 2021 Proposer’s Guide.

Tests pursued thus far have shown that if the separation between target and phase-calibrator is more than four degrees, the astrometric accuracy may worsen beyond the nominal (see Maud et al. 2020, ApJS, 250, 18). The proposed enhancement method is to choose a second phase calibrator (via the “user-defined calibration” option) that is not much further away from the target than the first calibrator and ideally in the opposite direction (see instructions below). By combining the responses of both phase calibrators, the effective target-phase-calibrator separation can be effectively reduced.

Types of Astrometry goals:

Absolute Positioning: Phase referencing between a target and a phase calibrator measures the angular separation between the two objects generally to an accuracy of about 5% if the synthesized beam is larger than 0.15” (10% if smaller; see equation above). Previously, astrometric observations were most commonly done at Band 6 or 7 and in configurations C-4 to C-7 because these array configurations are regularly scheduled and good phase stability with them is relatively common. If the absolute position of the target is needed to an accuracy below the nominal one, then the reference phase calibrator should have an J2000-ICRF VLBI absolute positional error < 2 mas, it should also be sufficiently bright and close to the target (< 4 deg). To check these, please check the instructions below.

Target Temporal Motion: To measure the relative motion of a target over time, the same phase calibrator should be used for all epochs (fixed in the "user-defined calibration"), even if its J2000-ICRF VLBI absolute position error is above 2 mas. Both the absolute position and temporal motion accuracy is not improved with phase self-calibration iterations.

Relative Motion Between Two Close Targets: If two or more objects are in the same ALMA field of view (nominally within the 40% primary beam power area), then the accuracy of their relative separation depends on their SNR, not on the position accuracy of the phase calibrator. The use of self-calibration may also improve the image quality (see Section 10.5.1) and thus the accuracy of the relative positions of different compact objects in the field.

Calibrator availability

By accessing the ALMA Calibrator Source Catalogue, one can search for calibrators around the target position. By filling in those fields, the PI should get in return a list of the closest calibrators and the observed fluxes at different frequencies (the most frequently observed Bands are 3 and 7, with occasional measurements at Band 6). The characteristics that should be checked are:

- Is the positional error smaller than 2 mas? Add “Measured RA Err.” and “Measured DEC Err.” in quadrature and multiply by 1000 since these columns are in units of arc-second.
- Is the calibrator bright enough at the frequency of interest? The value in the "Flux" column should provide a minimum SNR of 15 within 30 seconds (Section 8.3.1 in Technical Handbook; use the Exposure Time Calculator for this purpose or its web version). If the frequency of interest is not available, consider interpolation between Band 3 and Band 6 or 7 values (consider using the task “getALMAFlux” within the CASA extension Analysis Utils package). If only one measurement exists at a Band that is not the one of interest, please select “Enhanced positional accuracy”.

- Is the calibrator resolved at the requested spatial resolutions? In this case, you need to use the values in “UvMin” and “UvMax”. These should comply with $UvMin < 2.5E5/MRS$ and $UvMax > 2.5E5/beam$, where MRS and beam are the Maximum Recoverable Scale and the synthesized beam achieved by the selected configuration.

Note that due to flux variability and lack of intra-month/year monitoring of fainter calibrators, if the candidate calibrator barely complies with the flux criterion, the user is suggested to select “Enhanced positional accuracy”. In this way, the observatory might check the calibrator flux prior to the observation during routine observatory calibration observations.

J2000 ICRF quasar Frame and ICRS stellar frame and Ephemeris Motion

Because nearly all ALMA calibrators are distant quasars with J2000 epoch position on the FKN ICRF quasar frame, the normal phase referencing to an experimental target will provide its J2000 ICRF position after calibration and production of its image. However, most solar system and stellar objects (some of which have a moving phase center defined by an ephemeris track) are on the International Coordinate Reference System (ICRS) position grid with an origin at the barycenter of the solar system. This is the coordinate system used by Hipparcos and for many objects in our Galaxy. The difference between the ICRF and ICRS coordinate systems is about a 20 mas between the target positions. Hence, make sure that the coordinate system of the field source is specified correctly.

10.5.3 Band-to-Band Phase Calibration

As noted in Section 10.4.9, the closer a phase reference calibrator is to the target, the more accurately the calibrator scans will remove the tropospheric and instrumental antenna-based phase errors from the target scans. For frequencies above 350 GHz, where the density of quasar calibrators in the sky is relatively low, ALMA observations with 12-m Array configurations longer than 3 km and for 7-m Array observations that have low sensitivity, a relatively close calibrator to the target at the experiment frequency is often not available.

The goal of band-to-band phase transfer is to calibrate a target observation at high frequency ν_h by observing the phase calibrator at lower frequency ν_l where there is a greater chance of finding a detectable calibrator close to the target. The accuracy of this calibration method requires that the measured antenna-based spw phase of the low frequency calibrator scan $\phi_l(t)$ is converted accurately to what would have been measured at the target frequency, $\phi_h(t)$. The conversion has two parts: a simple frequency scaling between the low and high frequency (so called non-dispersive, delay-like behavior); plus a nearly constant instrumental phase offset ψ among all of the spws that must be measured for each 60 min experiment. This relationship is illustrated in the following equation:

$$\phi_h(t) = \phi_l(t) * (\nu_h/\nu_l) + \psi \quad (10.7)$$

The instrumental phase difference ψ is determined by quickly alternating observations of a bright quasar, called the *Differential Gain Calibrator* (DGC), between the low and high frequencies with a cycle time of about 30 seconds. From a 5-minute set of DGC scans near the beginning and near the end of a 60-minute experiment, the value and stability of the instrumental phase, ψ , for each low-frequency to high-frequency spw pair is measured. The DGC phase accuracy is about $5-10^\circ$, often limited by the short-term variable atmospheric phase variations, and it is generally constant over a 60-minute period. Occasionally, an antenna/spw DGC phase will have a small phase slope over the experiment, or a discrete phase jump. These rare occurrences can be detected in the DGC scans and corrected, or the offending antenna/spw data removed.

In Cycle 8, it is anticipated that band-to-band calibration may be used for 7-m Array observations in Bands 8, 9 and 10 in cases where there is no suitable high-frequency phase calibrator. A detailed investigation of the optimum scheduling of the DGC is underway using test data obtained in Bands 3 and 7 in the Cycle 7 long baseline configurations.

10.5.4 Band-Width-Switching

If all of the spws are narrow bandwidth (< 1 GHz total bandwidth), it is likely that a reasonably nearby phase calibrator cannot be detected within a scan. In this case, the strategy for phase calibration is to observe the

phase calibrator with wide-band spws and the science target in the desired narrow bandwidth by switching the spectral setup between the relevant scans. In a method similar to band-to-band calibration, a strong DGC calibrator is initially observed at both the wide and the narrow bandwidths in order to determine their phase difference among the spws that should be constant with time. There will be a slight delay change among the different spws, but they will be much closer in frequency than that for band-to-band observations. Tests are now in progress to determine the stability of the wide-to-narrow bandwidth phase difference.

Chapter 11

Quality Assurance

The goal of ALMA Quality Assurance (QA) is to ensure that the data products delivered to the PIs meet the expected quality standards. That is, the delivered products have reached the desired control parameters outlined in the science goals (or is as close to them as possible), they are calibrated to the desired accuracy, and calibration and imaging artifacts are mitigated as much as possible.

The QA analysis will be based on a calibration plan that specifies which observations must be acquired, and at which intervals, to monitor system performance and environmental factors as they evolve with time. Furthermore, it will also be used in (a) assessing the merging of data within each science goal taken with different configurations, (b) the inclusion of 7-m Array and TP Array data, and (c) the final image quality. Errors introduced by user supplied parameters, such as incorrect source coordinates, inadequate frequency setting (e.g. an incorrect redshift) or inadequate sensitivity limits (leading to an inadequate integration time or inadequate uv -plane coverage) are outside the scope of the ALMA QA, unless the error occurred due to faulty information or tools provided by the Observatory.

To be more efficient in detecting problems, ALMA QA has been divided into several stages that mimic the main steps of the data flow. The broad classification of this multi-layered QA approach is:

QA0: The first check and monitoring of the calibration and overall performance during and just after an observation, performed at Execution Block level.

QA0+: A simplified and fast scripted imaging QA. Typically completed 30–90 minutes after each observation is taken, this is mainly used when the results from QA0 are unclear or marginal and will also be available in the QA0 report in Cycle 8.

QA1: This involves the measurement of performance parameters and telescope properties by the observatory.

QA2: Full calibration and generation of imaging products.

QA3: Issues found with the data by the PI or ALMA staff after data delivery.

The QA process is handled by the Program Management Group (PMG) and the Data Management Group (DMG) at the JAO with significant contributions from ARC personnel. The process is primarily implemented using the ALMA Quality Assurance (AQUA) Tool (see Section 11.7) and other software necessary for the assessment of data quality. Responsibility for data quality assurance rests with the Data Management Group Head within the Department of Science Operations, drawing upon the resources of the Program Management Group and the Data Reduction Managers (DRMs) at the JAO and ARCs.

The final output of the ALMA QA0 and QA2 processes is a QA Report per ObsUnitSet (Member or Group¹) that summarizes all the relevant QA information for each of the different QA stages up to, and including, the final imaging. This report is included in the data package delivered to the PI. The QA3 process is handled differently (see Section 11.6). A more detailed description of the different stages of QA is given below.

¹Group ObsUnitSet processing, that is, the combination of datasets taken in different configurations/arrays. It is expected that group processing will be implemented as a regular data reduction strategy in later ALMA cycles.

11.1 Cycle 8 Quality Assurance Goals

The ultimate goal for ALMA QA is that the delivered products are considered "Science Ready", suitable for publishing with little need for the user to reprocess. However, ALMA cannot guarantee that all the scheduled Cycle 8 projects are completed (especially those with grades B and C), or that specific criteria like sensitivity or angular resolution are precisely met, given the restrictions of the array configurations or system performance at the time of the observations. The observatory will therefore attempt to meet the sensitivity and resolution stated by PIs in their Science Goals, within certain tolerances, as described later in this chapter. In principle, each project component ("Scheduling Block" or SB) is scheduled for the number of executions that are expected to reach these goals (based on nominal Array performance values). Additional executions may be needed during the observing cycle if specific executions do not pass QA0 or QA2 (see below).

In Cycle 8 users may need to invest their own time and expertise to ensure that the data products are suitable to meet their science goals. In rare cases, users may need to re-reduce the raw data if the quality is not satisfactory. At any time, users may request to visit their ARC or ARC node to obtain help and to be assisted with the data processing and analysis.

11.2 QA0

The first stage of QA, known as QA0, is a near real-time verification of data quality for each SB execution, or execution block (EB). It deals with performance parameters on timescales of an SB execution length or shorter, and thus is performed at the time or immediately after each execution. Initial QA0 assessment is performed by AoDs (Astronomers on Duty) at the OSF, using the software tools AQUA, AosCheck, and QuickLook, based on the near real-time output from the calibration scans obtained by the ALMA TelCal software. This information is complemented with reports derived using Monitor and Control display tools to monitor specific parameters not directly tracked by the calibrations (e.g., total power level variations, weather parameters, etc).

QA0 metrics/parameters have been selected to check the health of the whole signal path from the atmosphere down to the correlators, as well as issues connected with the observation itself. AosCheck examines the calibration data, looking for significant outliers or out-of-spec results, and other problems (for example, due to the correlator). Together with additional checks in AQUA itself, these are combined to provide a QA0 assessment for each EB. The parameters checked at QA0 can be grouped into the following general categories²:

Atmospheric Effects: Weather Parameters, Sky Opacity, System Temperature, Phase Fluctuations, WVR Outputs.

Antenna Issues: Antenna complex Gains, Antenna/pad delays, Relative/Offset Pointing, Geometric Shadowing, Antenna/Pad positions.

Front-End Issues: Bandpass shape & instability, Sideband Ratios, Receiver Temperatures, Phase variations, variations in T_{sys} spectral shape.

Connectivity Issues: IF Delay Measurements, System Temperatures, unusual relative phase or amplitude variations between spws, interference spikes.

Correlator Issues: Bandpass spectral shapes, Delay Measurements.

Observation Issues: Calibrator fluxes, incomplete datasets or incomplete mapping.

Tolerances for individual parameters in the above list (for example the maximum range of T_{rx} or the cutoff for low antenna gain) are defined in AosCheck. These are combined to assess the QA of individual sub-units of the EB, such as antennas, spws, scans, periods of time, or targets. An overall assessment of QA0 of the EB is then made based on combining these results. The overall tolerances for this QA0 assessment that have been adopted by ALMA for this Cycle are listed in Section 11.2.1. Additional checks are made in AQUA on other observatory calibration data, including a recent measurement of the flux of the amplitude calibrator, and a relevant measurement of the antenna/pad position.

Based on these results, each SB execution is classified into three main categories:

²Some additional hardware problems can occur which are automatically flagged out by the system before they reach the QA0 stage, and may not be seen in the data or at QA0 unless this hardware flagging has failed. Examples of these are LO lock problems, antenna not on-source, or the calibration device not in position.

- **QA0_Pass** are datasets that comply with all the QA0 criteria and will be used for the final imaging.
- **QA0_Semi-pass** are datasets that do not fulfill all the QA0 criteria, but contain data that is deemed of some scientific value (for instance, the presence of usable bandpass/amplitude calibrator data). QA0_Semi-pass data are not included in the final data products, but PIs can access those data from the ALMA Science Archive.
- **QA0_Fail datasets** are those that are not included in the other two categories. Since they represent unusable data or data that cannot be calibrated, they are not made available to the PIs through the ALMA Science Archive nor visible in SnooPI.

In addition, other temporary states can be set at this time, which require further intervention before they can be re-set to one of the above main states:

- **QA0_Pending_Other** is used if the state based on the initial assessment is unclear; it requires additional QA, using QA0+ imaging as well as possibly further checks (see Section 11.3).
- **QA0_Pending_Fluxcal** is used if the EB is a Pass but the amplitude calibrator (which is usually also the bandpass calibrator) has no usable flux measurement within the last 7 days. Once a flux measurement is available in the archive, these data will automatically be set to Pass.
- **QA0_Pending_Correct_ANTPOS** is used if an antenna was used in the array which had not been fully integrated - usually because the antenna positions in the observatory Telescope Monitor and Control DataBase (TMCDB) have not yet been updated after an antenna move. Once the positions are measured and updated, these data will be set to Pass.

11.2.1 QA0 criteria

Of the antennas available in each interferometric dataset, at least 50% of them (or 22 out of 43 in a 43 element 12-m Array and 5 out of 10 in a 10 element 7-m Array) must be ‘usable’ for the dataset to be considered for a QA0 Pass. If fewer than 50% of the antennas are usable, the EB is normally set to QA0_Semi-Pass (Section 11.3). Datasets with more than 50% but less than 100% usable antennas will normally be QA0 Pass, but will count less than datasets with 100% usable antennas: this is done through the use of a weighting factor called the Execution Fraction (section 11.2.2). The antenna restriction does not apply to the TP Array (where the minimum number of usable antennas is 1). The decision as to whether an antenna is ‘usable’ or not depends on several criteria - both system and weather related. The QA0 criteria that have been adopted by ALMA during Cycle 8 are the following:

- **Antennas:** Situations that will be considered to render an antenna or spw as not usable include issues with the antenna or frontend itself, such as T_{sys} value $>$ a factor of 2 higher than the median of others in the same spw, an absolute delay in that spw of > 0.5 ns, relative amplitudes on the calibrators $< 50\%$ of the median of other antennas/spws, gain calibrator amplitude or phase fluctuations which are significantly higher (factor of > 3) than other antennas at a similar distance from the array center, or scan-based phase noise on the bandpass calibrator significantly higher ($> 3\sigma$) than other antennas. These criteria are designed to identify outliers due to antenna or system problems, as distinct from the effects due to poor weather conditions noted below.
- **Weather:** If the atmospheric phase fluctuations on the complex Gain calibrator measured over the whole execution are such that the antenna-based phase noise³ is larger than 1 radian for a given antenna, then that antenna is regarded as ‘unusable’ for QA0 accounting purposes in calculating the Execution Fraction. Note that these high phase noise antennas are not flagged out at QA0, and will still be reduced in the pipeline.
- **Bandpass:** A high SNR (> 30) must be attained per channel on observations of the bandpass calibrator, after binning to no fewer than 4 channels. The dataset will be set to either QA0_Semi-pass or Fail if the

³The antenna-based phase noise is determined in the following way: data from spws and the two polarizations are phased up and combined to obtain a single phase solution per antenna per Gain calibrator scan. The phase noise for that antenna is then determined from the mean value over the whole execution of the difference in phase solutions between adjacent scans.

bandpass calibrator signal is too weak to attain this SNR depending on the outcome of QA0+ (Section 11.3). The QA0 bandpass results (both amplitude and phase) are used to check that the passband shape does not show statistically-significant spikes, high spectral rms, or correlator platforming issues. Also the ATM (T_{sys}) calibration observations are used to check for significant passband shape changes during the execution. If any of these effects are seen, either the relevant channels or the whole spw on that antenna can be flagged, and will be regarded as ‘unusable’.

- **Gain/phase:** A SNR (for each antenna) of ≥ 10 per scan must be attained after averaging all spws. The dataset will be set to either QA0_Semi-pass or Fail if the gain/phase calibrator signal is too weak to achieve a SNR ≥ 10 depending on the outcome of QA0+ (Section 11.3). Also note that individual scans on the Gain calibrator are checked for statistically-significant ($> 5\sigma$) amplitude drop-outs.
- **Execution:** More than 20% of the expected time is required on the science target for a given SB execution, and the overall Execution Fraction (see 11.2.2) must be greater than 0.2 to be considered as a QA0 Pass. The dataset will be set to QA0_Semi-pass if it does not satisfy the above criteria.
- **Calibrations:** Critical calibrations (bandpass, ATMcal, at least one pair of Gain/phase scans bracketing the science observation) must be completed. There may also be additional calibration requirements, such as polarization with > 60 degrees of parallactic angle rotation during the session. Datasets which contain only usable calibrator data shall be set to QA0_Semi-pass.
- **Storage:** Data must be able to be read from the Archive. Any dataset that cannot be retrieved from the Archive will be set to QA0_Fail.
- **Time Sensitive Observations:** For Target of Opportunity (ToO) or time critical observations, the dataset must be collected within the time range specified by the PI. For time simultaneous observations between the ACA and 12-m Arrays, the executions must start at the same time. Observations taken where the data are not collected simultaneously will be set to QA0_Semi-pass.

11.2.2 Execution Fraction

For operational and scheduling purposes, individual SB executions during Cycle 8 will effectively be weighted by measuring their data quality. This weighting is quantified by an Execution Fraction (EF), where 1.0 is the value predicted per execution by the OT, based on typical observing conditions and number of antennas. An EF of 0.0 means that the data are not at all usable and the observations should be set to QA0_Fail or semipass, while an EF > 1.0 means that the EB is better than predicted. The EF is a single number representing the normalized fraction of the theoretical sensitivity (based on the radiometer equation) that a given execution should have reached, combined with a measure of the expected image quality (see Section 11.3). It is calculated using the number of ‘usable’⁴ antennas compared with the number that should be in an array in Cycle 8⁵, the actual time on-source compared with expected, the mean T_{sys} at the representative frequency compared with that assumed by the OT, and the fraction of antennas with phase noise below the set limit (of 1 radian)³.

Executions with *lower* data quality than the reference will have fractional executions between 0 and 1.0, while those with better quality (such as observations carried out in unusually good weather with lower T_{sys} , or with more than 43 antennas), will have an EF value above 1.0. For Cycle 8, the EF on the main array is capped at 1.5; this allows extra execution time in the case of very good weather (where the EF would otherwise be significantly higher), to avoid the situation where uv coverage may be important. The scheduler then sums the EFs of executions to determine how many more are required for SB completion. Note that datasets with $\text{EF} < 0.2$ will be set to QA0_Semi-pass. For Mosaic observations, if the execution is incomplete (for example if the SB finished unexpectedly early) and $< 90\%$ of positions have been observed, the EF will be dropped to 0.2. The data will then still be reduced, but the execution will not count significantly against the total time.

⁴see 11.2.1

⁵The requirement for Cycle-8 is forty-three 12 m antennas on the main array, ten 7-m antennas in the 7-m Array, and three 12 m antennas in the TP Array.

11.3 QA0+ and imaging quality

QA0+ is an additional stage of QA0 which runs a simplified and (relatively) rapid scripted calibration and imaging pipeline of the channel-averaged data, producing continuum-only images of the science target, phase calibrator, and check source (when available). The idea behind QA0+ is that if an execution is left in a QA0 "Pending Other" state (perhaps because the phase noise is above the QA0 limit, or the QA0 status is unclear for any reason), the QA0+ images can be used as an additional level of data quality assessment. The EB can then potentially be set to "Pass" if the images meet the QA specifications. Typically the QA0+ scripted calibration and imaging process takes 30-90 minutes to complete. To speed up calibration, it does not run bandpass or the full Tsys calibration, but uses WVR-corrected data. If QA0 recommends the remcloud option (for wet cloud removal) it will also run this. It uses flags from the QA0 process (flagging out bad antennas, but not those with high phase noise) and combines both polarizations and all spws with bandwidth >400 MHz. For imaging it will run a single iteration of CASA tclean with a clean box centered on the first science target. For science targets with SNR >20, it will try self-calibration on a scan-based time interval and will generate a self-calibrated image.

As well as the images, also provided by QA0+ are the beam size and shape for different Robustness parameters (from the phase calibrator), and the science target image, peak and integrated flux, and image RMS. Note that these are very approximate values, obtained from the mean Tsys, and taking the mean antenna gain (Jy/K) to obtain the fluxes. As well as the results from the current execution, the data are combined with previous executions of the same SB, to give the running results so far. Starting in Cycle 8, these values will be available to the PI through the QA0 report for each execution in SnooPI.

For EBs requiring QA0+, checks of image quality are performed where possible. The metric used is the peak to integrated flux ratio for a point source (either a check source or the science target if known to be compact). Also the peak to residual ratio after a single tclean iteration can be used. Both indicate how much flux from a point source is distributed over the map due to phase decoherence. Whereas completely random phase fluctuations simply reduce the target flux and scatter power over the whole map, semi-systematic or large-scale atmospheric phase changes over the array will distort the image in a semi-coherent way. These can be more egregious during long-baseline or high-frequency observing, and their effect can be monitored by observing a check-source as part of these observations⁶; QA0+ is then used to assess the image quality.

A perfect image has a peak/integrated value of 1.0; a 'good' image has a peak/integrated value > 0.8 and images with values < 0.5 are regarded as "marginal". The QA0+ criterion is that datasets with a peak/integrated flux ratio on the check source of >0.5 are set to Pass. For those without a check source, in some cases, the science target is bright enough that the images can be compared with previous executions of the same EB. Those judged of similar quality are also set to Pass. For executions where the check source peak/integrated ratio is <0.5, or there is no check source or clear science target, the QA0 status is maintained at (or set to) Pending Other. These are unclear cases, and require a further stage of checking, usually offline by the full ALMA data reduction pipeline, before setting the dataset to either QA0 Pass or Semi-pass.

11.3.1 Relationship between image quality and phase noise criteria

The quality of images is strongly affected by the phase stability of the atmosphere⁷, so measuring the phase noise provides a good proxy for the final image quality. However, there are cases where it is really necessary to check the image directly. ALMA uses nearby phase calibrators to measure and correct for residual sky phase fluctuations (phase referencing). But this correction is not perfect. The image quality of these phase-referenced observations is affected by spatial differences in the line-of-sight delays to the target and phase calibrator, as well as baseline errors thought to be due to large-scale, slowly-varying atmospheric structures⁸. The spatially-dependent differences between target and phase calibrator are linked to the time-dependent phase fluctuations through the turbulent structure of the atmosphere, so monitoring just the time-dependent fluctuations can still give a reasonably good measure of the overall data quality, even without checking the images. Figure 11.1 confirms that image quality (measured by the ratio of peak-to-integrated flux on a point source, taken from QA0+ results on the check source during the 2017 long-baseline campaign) is correlated with the phase fluctuations (measured by the median phase difference between antennas³). The phase is measured in degrees,

⁶Check sources are QSOs within a few degrees of the phase calibrator, and are normally observed for a short time during long baseline or high frequency executions.

⁷Although many of these effects can be removed using self-calibration when the target is bright enough - see Section 10.5.1

⁸These baseline errors only become significant on the longer baseline observations (of several km or more)

and therefore just depends on the changing pathlength compared with the wavelength, independent of the actual observing band. As stated in Section 11.3, a perfect image has a peak/integrated value of 1.0; a "good" image has a peak/integrated value > 0.8 and images with values < 0.5 are "marginal". Datasets with phase fluctuations less than 1 radian almost always give reasonably good image quality and can be set to QA0 Pass immediately. The QA0+ imaging stage is normally used for the marginal datasets; some can then be set to Pass as the images are deemed acceptable, but a significant fraction (perhaps 20%) are of very poor image quality.

Time-dependent phase noise is not the only parameter that affects image quality, however, Figure 11.1 indicates that observations where the target-calibrator separations is more than 3deg on the sky tend to have worse image quality even if the time dependent phase fluctuations are low (< 1 radian). This is illustrated further in Figure 11.2, where the image quality is plotted against the phase calibrator separation. For reasonable conditions (phase RMS < 1 radian), generally there is a slow dependence on separation; "good" image quality is generally achieved when the separation is small (< 4 deg) and phase fluctuations are low (< 1 radian).




Figure 11.1: Check source image quality (measured by peak to integrated ratio) shown as a function of phase fluctuations (the mean antenna-based scan difference on longer baselines, in degrees). Data are taken from 2017 long-baseline observations in Bands 3-7 (no dependence on the band was found). Blue points are QA0 pass, red crosses are QA0 Semi-pass. Encircled symbols indicate that the phase calibrator has a more than 3deg separation from the check source. The highest image quality (peak/integrated > 0.8) generally occurs for phase fluctuations < 40 deg and for the closest calibrators. The cutoff for an immediate QA0 Pass from AosCheck is < 57 deg (1 radian); higher phase noise datasets are passed on to QA0+ for assessment of image quality.

11.4 QA1: Observatory calibration and QA

In addition to data taken during the SB execution, the observatory tracks array and antenna performance parameters which vary slowly (typically on timescales longer than a week) and which can affect data quality. They are measured by AoDs and System Astronomers at predefined periods as "Observatory Tasks", and after major interventions including, but not limited to, antenna moves, front-end swaps and subreflector repairs, or if significant deterioration of performance is detected during operations. Typically after such interventions, antennas are set to a non-integrated state and not used for regular science observing. Antenna integration into




Figure 11.2: Check source image quality (measured by peak to integrated ratio) shown as a function of phase calibrator separation (in degrees on the sky). Data are taken from 2017 long-baseline observations. Blue points are low phase noise (< 1 radian) and red points are high phase noise datasets (> 1 radian). The blue line shows a linear fit to the low noise datasets weighted by the inverse of the phase fluctuation. The phase-referencing image quality deteriorates somewhat with larger separations for low noise data, but for high noise data the separation is less important. Datasets with peak/integrated less than 0.5 are regarded as low quality and are passed on to QA0+ for assessment of image quality.

the array then requires measurement and ingestion of specific array and antenna parameters (described in more detail below).

Observations and reduction of such data are done jointly by the AoDs and specialized groups within DMG and PMG. The end products are a set of parameters that are ingested into observatory databases, in particular the TMCDB which records hardware status and history. The TMCDB gives the up-to-date view of hardware parameters which are used by the system during observations. Problems that would significantly downgrade the quality of the data are solved by, for example, dropping the offending antenna from the active array and fixing, re-measuring and updating the problematic parameters. Particular parameters which are derived from these observatory calibrations include:

Array Calibrations: Baseline and antenna position measurements and pad delays.

Antenna Calibrations: Pointing models, focus model curves, surface measurements, beam patterns (including polarization observations), beam squint (offset between polarizations), IF delays (relative delays between basebands and polarizations), front-end delays (relative delays between front ends).

Source Calibrations: Monitoring of solar-system flux standards, and secondary quasar flux standards (e.g. grid sources, see Sections 10.2.6 and 10.4.7).

11.5 QA2

The QA2 concept refers to the Quality Assurance performed at data reduction level. The ALMA data are processed (i.e. calibration and imaging) either with the ALMA Data Reduction Pipeline, or manually, with the Script Generator. The ALMA Pipeline and the Script Generator are built on CASA, which is the official software for ALMA data processing and analysis. The ARCs and ARC nodes do not provide support for

any data processing or analysis done by users on platforms outside of CASA. The ALMA Pipeline and Script Generator can be run at either the JAO or the ARCs.

The entire data reduction workflow is managed by the Data Management Group (DMG) Head at the JAO, with support from the Data Reduction Managers (DRMs) at each Executive. It is only during this stage of data reduction that the science goals set by the PI can be compared with the actual values in the data products (i.e., noise RMS, angular resolution, SNR and dynamic range).

Like in previous Cycles, during Cycle 8 it is expected that, for the majority of the observing modes, the ALMA Pipeline (see the ALMA Science Pipeline Reference Manual on the Science Portal⁹) will be used for the calibration and imaging of the data.

The ALMA Pipeline provides a Weblog page with detailed information on the quality of the calibration, which provides the basis for QA2. Weblogs are included in the data packages sent to the PIs and are available to all users through the ALMA Archive. For manually-reduced datasets, the QA2 will be carried out with special purpose scripts whose outputs will be included in the data packages delivered to PIs.

Some datasets can be calibrated by the pipeline, but need to be imaged manually, or with manual intervention in the pipeline. In these cases, a weblog will be provided for the calibration portion, and an imaging script, or imaging weblog provided for the imaging. The QA2 metrics which determine the success of an observation are given below.

A complete description on what the ALMA Pipeline checks is available in the Cycle 8 Pipeline User Guide at the ALMA Science Portal. A short list of the QA2 parameters or issues checked during data reduction (calibration of individual Execution Blocks and joint imaging) are listed below:

Calibration Issues

- Bandpass quality: is the spectral profile flat, well behaved and devoid of spectral lines?
- Flux scale calibration: flux accuracy better than 5% for Bands 3, 4 and 5; 10% for Bands 6, 7 and 8; and 20% for Bands 9 and 10
- Phase transfer and astrometry: Cycle time and sky separation between phase reference and target; typical and extreme (unflagged) phase differences between phase reference scans.

QA2 criteria

For Cycle 8, the beam and noise RMS criteria used for the QA2 assessment are the following:

- Noise RMS in target images: The predicted value (after data flagging, if necessary) is compared to the value requested by the PI. The achievement of the requested noise RMS in the images must be within 10%, 15% and 20% of the goal for Bands 3/4/5/6, 7/8, and 9/10, in flux units per beam, respectively. For Solar Observations, this criteria does not apply as there is no blank sky in most of the images.
- Size of synthesized beam (spatial resolution) for specified weighting scheme, or specific imaging requirements (if stated by the PI). In order to satisfy the QA2 criteria, for Cycle 8, the beam area, not the individual beam axes, shall be the figure of merit. The PI requested angular resolution range shall be interpreted as the geometric mean of the major and minor axes at the PI specified declination. The beam area must be within the range requested by the PI, taking into account the modifications applied by OT to ensure schedulability of the SB (See Cycle 8 2021 Proposer's guide). In addition, those Member Observing Unit Sets (MOUS) that do not satisfy the QA2 requirement of beam area will have the execution fraction of the associated Execution Blocks (EB) reduced to 75% of its current value if failed due to erroneous beam, or by the corresponding fraction if not reached the required sensitivity.
- Image used for QA2 assessment: The image used to derive both the noise RMS and beam size values mentioned above shall be the image whose channel width is the one closest to the representative

⁹<https://almascience.nrao.edu/processing/science-pipeline>

bandwidth selected by the PI in ALMA OT, in the spectral window closest in frequency to the representative frequency.

Other criteria used for the final data characterization

In addition to the beam and noise RMS criteria, the following factors will be considered for assessing the quality of the data products:

- *uv*-coverage and time on target (after flagging).
- Spectral resolution and channel ranges used to make sample images.
- Dynamic range limitations.
- Amplitude scale consistency between executions of a multi-execution MOUS.
- Contamination by bright sources outside FOV or aliasing in the ALMA Pipeline task called clean.
- Minimum parallactic angle coverage for full polarization projects.

The data reduction scripts sent to the PI will include any modifications to the standard scripts which were found to be essential during QA2.

As in QA0, there are three possible QA2 states: QA2_Pass, QA2_Semi-pass or QA2_Fail. In detail:

- **QA2_Pass** implies that the scientific goals, as defined by measurable parameters such as noise RMS, LAS and angular resolution, have been achieved at the representative spw within the specifications listed above.
- **QA2_Semi-pass** refers to those datasets that fall short of meeting the PI requested science goals, but are otherwise of good quality (or as good as possible with the observations that could be achieved).
- **QA2_Fail** is a temporary state used during an observing cycle when an observation fails to meet the PIs goals by a significant margin and needs to be scheduled for additional observations. Although the Observatory will try to take a similar quality of data in all spws, those measurable parameters in non representative spws will not be used for the judgement of QA2_Pass or QA2_Fail.

If there are not enough additional executions of an SB to pass QA2 by the end of an observing cycle and the SB is not a component of a grade A proposal, the QA2 state is changed from QA2_Fail to QA2_Semi-pass and delivered to the PI.

11.6 QA3

QA3 is post-delivery evaluation of the data products delivered to the PIs. It is advisable that PIs check the data products themselves very soon after delivery. The QA3 process is often triggered by a PI who initially finds some issue in their data. which may reflect an underlying problem with the data, observing procedure or calibration. The issue is then reported to the corresponding ARC for support, typically through a Phase 2 ticket via the ALMA Helpdesk (<https://help.almascience.org/>) which is then reviewed by the project's contact scientist or other ARC personnel. After this initial triage and review, the contact scientist will request the user to open a formal QA3 ticket through a dedicated department in the Helpdesk only after it has been determined not to be a user error. The ARC receiving the Helpdesk QA3 ticket will retrieve the data from the archive and evaluate the nature and extent of the problem. The evaluation by the ARC should include an assessment on whether the problem is present only in a particular dataset or is a more global issue affecting data taken under similar set-ups and observing conditions.

If the identified problem reflects a problem with the performance of the array, the calibration or data reduction processes, or the QA process, the ARC will communicate their findings to the JAO, who will work on solving the problem in collaboration with the ARCs. The result will be communicated back to the reporting investigator.

An extension of the proprietary period of delivered datasets still within their proprietary period with QA3 issues will be granted based on the policies in the Cycle 8 Users' Policies document. There will be no extension

to the proprietary period of data already publicly available. During the QA3 evaluation process, which may include taking additional data for an observing program, the data will not be available through the Archive. The data will only become available after the issue has been resolved and new data products are ingested into the Archive.

11.7 The Quality Assurance Reports

The Quality Assurance Reports of QA0 and QA2 will be accessible to the users using SnooPI. These reports will be generated using the AQUA tool. For Pipeline calibrated datasets, additional reports (Weblogs) will be sent to the PI as part of the data delivery. For observing modes that require manual data reduction, script-generated QA2 reports will be delivered.

The QA0 report, which is a summary of the QA0 and QA0+ results, will be available from AQUA (through SnooPI) after each execution. This report can be used to check the status of ongoing observations. Note that the QA0+ image results are only very approximate, and only apply to the continuum (channel-averaged) data.

The basic unit of a QA2 Report is the ObsUnitSet, which represents part of the scientific goals stated by the PI during Phase 1 (project generation and review). An ObsUnitSet will typically contain several executions of an SB. Each execution only passes to data processing if it has passed QA0. There will be only one QA2 report for the whole Member ObsUnitSet¹⁰ generated by the JAO/ARCs at the end of the data reduction process. This report also has to be approved by the DRMs before the data products are delivered to the PI. An interface between the Pipeline and AQUA permits to generate a complete QA2 report, which is delivered as part of the data products package (including the Weblog information). The final report per MOUS delivered to the PI would in that case be a concatenation of all the relevant QA0 reports per execution with the QA2 report. Comments on each stage of the QA process (with supporting images, if required) would be added to the Report.

The standard policies for QA0 failures (or QA0 semi-passes) are that the observations of those Execution Blocks that failed have to be repeated. Failures to pass QA2 may trigger additional observations if the achieved RMS or imaging parameters do not fulfill the QA2 pass criteria. Additional observations may not be possible in circumstances, such as projects with very tight weather, very specific configurations, or time constraints. If the available data are insufficient to reach the required sensitivity, but are otherwise of good quality, they will be released to the PI at the end of the cycle.

11.8 QA in Solar Projects

Section 8.10 describes the difference between solar observing modes and those employed for sidereal sources. Those with implications for quality assurance are as follows:

- No WVR observations are possible during solar observing and therefore no antenna-based phase corrections can be made.
- Observations of the calibrators are made using the so-called MD modes (Bands 3, 5, & 6), which may yield intrinsically lower SNR. The stepped attenuators are set and held fixed at appropriate levels for both source and calibrator scans.
- The system temperature may be significantly higher for an MD mode. The antenna temperature is of course much higher as a result of the solar input.

For QA0, the criteria are largely the same as those itemized in 11.2.1. In the case of weather, the phase RMS criterion is the same for solar observations as they are for non-solar observations (1 rad). Note, however, that this criterion is applied to data that have not had WVR corrections applied and are consequently more stringent. For fast total power mapping, QA0 criteria are that the calibration sequence and map were completed for each execution.

Solar observations are currently a mode that cannot be pipeline processed. In the case of QA2 (cf. 11.5), calibration issues include evaluation of the antenna temperature, which is much higher than non-solar sources.

¹⁰As stated above, processing of group ObsUnitSets will not be done during Cycle 8. PIs should assume that all statements here apply to Member ObsUnitSets (MOUSs).

In particular, T_{ant} should be $6000 \pm 1500\text{K}$ for Band 3 and 5, and $5000 \pm 1500\text{K}$ for bands 6 and 7, except for limb observations. For final data characterization the criteria follow those in 11.5 with the following exceptions and caveats:

- Synthesized beam: check the phase calibrator map for loss of coherence, if any, due to phase variance.
- Spectral resolution and channel ranges: solar modes are fixed - no checks needed.
- RMS noise in target images: n/a for source dominated solar data.
- Combinations of array configurations: n/a for solar observations in Cycle 8.
- Mosaicing and/or contamination by bright sources: Pass if 90% of the mosaic has been completed, Fail if completion is below 20%, otherwise semi-pass.
- Polarization purity: n/a for solar observations in Cycle 8.

For total power mapping, there should be no artifacts in the map that can be attributable to the (double-circle) scan pattern or to transient sky absorption.

11.9 QA in VLBI Projects

As discussed in Section 8.11, VLBI observations are executed as polarimetric observations and a portion of the above-described procedures are applicable with an important distinction. Since VLBI observations are performed with ALMA's peer observatories according to a schedule worked out months in advance, it is not possible to re-observe failed observations (regardless of reason). Thus, the diagnostics provided as part of the normal QA0 process are examined by the staff, but their utility is restricted to making the limited adjustments to the observation which are practical. The principal adjustment is to remove antennas which were performing acceptably at the start of the observing project but develop issues serious enough to degrade the observation. Such serious issues are relatively rare - antennas that are merely not performing optimally are generally retained.

A version of the QA0+ online reduction is then made and it provides prompt feedback on the performance of the phasing system. Briefly, it provides time histories of the visibility amplitude and phase for baselines between the phasing reference antenna or phased sum and the comparison antennas. These reductions are made principally to verify that the sum signal provides a $\sim \sqrt{N}$ improvement in performance relative to the Reference Antenna. If for some reason there are issues, this provides an opportunity to investigate before the schedule finishes.

Finally, the QA2 process for VLBI is somewhat different than for standard interferometry (polarimetric) projects. Because VLBI projects will not be re-observed as a result of any imperfections, the focus is on doing the best reduction possible with the data that were taken. The fruit of this effort is delivered to the PI as with non-VLBI observations. The second key difference in the QA2 for VLBI is a polarimetric reduction specifically designed to allow the linear to circular (`PolConvert`) transformation of the VLBI-correlated visibilities and a normal VLBI analysis in the circular visibility basis. The details go far beyond the scope of this document, but the interested reader may consult a full description available elsewhere. (See *Calibration of ALMA as a Phased Array*, C. Goddi, I. Martí-Vidal, H. Messias, et al., *PASP*, 131, 5003.)

11.10 QA in Pulsar Projects

Pulsar data will not undergo the ALMA Quality Assurance process.

Chapter 12

Data Flow and Logical Data Structure

This chapter describes the data flow process from the observations until raw data is ingested into the ALMA Archive. It includes a brief description of all the main software subsystems involved in the data acquisition and archiving, as well as a summary of the data structure adopted by ALMA.

12.1 Data and Control Flow

This section describes the overall control of the ALMA system and the flow of data during observations. A summary of the main actors and operations involved in the observations is shown in Figure 12.1. Each of the light blue boxes in Figure 12.1 represent an ALMA subsystem involved in the observations. The rest of the boxes, colored according to the actor involved, include labels for the actions performed either by external agents (actors) or by those subsystems.

The software subsystems run within a cluster environment known as a Standard Test Environment (STE), with the STE used for the real array operation referred to specifically as an ALMA Production Environment (APE). The STE incorporates central server nodes, embedded computers for hardware control, dedicated correlator data processing nodes, and console nodes for user interaction.

A typical observing session would be started by the Telescope Operator interacting with the Executive subsystem. The Executive subsystem is in charge of starting up the ALMA Common Software (ACS) and its Common Object Request Broker Architecture (CORBA)-based services and then initializing all of the various software subsystems involved in the observing and data storage process. Once all the components are ready, the Executive also handles asynchronous events from several of the subsystems and responds to them accordingly. Among the events, the Executive also publishes a list of error conditions to the attention of the operator and the requests for status of the Control, Telescope Calibration, Correlator and Scheduling subsystems.

The actual observations start by the operator creating an array, which means selecting the antennas, photonic reference and correlator (if needed by the observations¹) to be used. The antenna selection is largely automated by filtering based on antenna types, connection to either correlator, and status of the antennas in the ALMA Dashboard². Up to six independent arrays can exist and operate at one time. Once an array is successfully created, Scheduling Blocks (SBs) can be queued and run on it. SBs can either be manually chosen, e.g. for array calibration or test purposes, or selected from a ranked list produced by the dynamic scheduler for regular science operations.

An SB is the smallest, calibratable element of a project. SBs are XML documents which are part of the ALMA Project Data Model (APDM), and are kept in the archive along with the proposal and other ancillary information. SBs have an associated state which among other things controls whether they can be executed. For science SBs to be executable the Phase-2 process³ should have been completed for them, and they should not have been fully observed or cancelled/timed-out at the end of an observing cycle.

The dynamic scheduling software considerations include the target elevations as a function of time, observing

¹Most ALMA observations use a correlator, either for interferometry or single dish spectroscopy, but single dish continuum observations can be made using only power detectors in the antennas.

²Pietriga et al. 2014, SPIE, 9152, 1B

³For a description of the observing preparation processes, see the Observer's Guide




Figure 12.1: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

conditions (PWV and estimated path length variability), actual array configuration, project completion status, proposal grade and rank, among other things. The dynamic scheduler will also suggest array calibration SBs at appropriate times. It is intended that by Full Operations the algorithm will be optimized to the point that the dynamic scheduler can fully automatically queue SBs during normal operation periods. At present the Astronomer on Duty will select SBs in the dynamic scheduler GUI to be added to the array’s scheduling queue for execution.

When an SB reaches the head of the array’s scheduling queue it starts an execution. The SB contains the name of a Python observing script, which is run by the Control subsystem and passed a Python representation of the SB content. The observing scripts, which reside in the Science Software Requirements (SSR) module of the software, use the SB content to determine the necessary execution sequence. The observing script commands the Control subsystem with sequences of scans and subscans. The scan and subscan specifications specify for each subscan the tuning, phase center and antenna pointing, calibration device position, intent metadata (to convey the purpose of the scan/subscan), and other parameters that may need controlling from the observing script level. Control executes the scan/subscan sequences by commanding all relevant hardware, the relevant correlator subsystem (Baseline (BL) or Atacama Compact Array (ACA)) and the Total Power Processor (TPP). Subscans, and most actions performed by the Control and correlator subsystems, start and end on 48 ms (Timing Event) boundaries which are accurately synchronised among the whole system. The subscans result in raw data files being sent to the archive subsystem incrementally during the observation, and the metadata rows being sent directly to the online Telescope Calibration (TelCal⁴) subsystem through CORBA-based services. TelCal publishes results from calibration scans it reduces by sending them back to the

⁴<http://adsabs.harvard.edu/abs/2011ASPC..442..277B>

Control subsystem (Data Capture, specifically) as calibration tables, where they are displayed by GUIs to allow Operator and AoD evaluation of the observation progress, as received by the observing script to apply online corrections e.g. for antenna pointing and focus. The Control subsystem writes all the metadata tables to the archive at the end of the observation.

Each run of an SB produces an Execution Block (EB) that is stored into the archive through two parallel paths: metadata tables compiled by the "Data Capture" component of the Control subsystem, and binary data files which are streamed directly from the data producers (ie, correlators) via the "Bulk Data" transmission system. The metadata contains the relevant references to the binary data files to allow them to be exported from the archive along with the metadata. The metadata tables are defined by the ALMA Science Data Model (ASDM, see Section 12.2)⁵. The tables are primarily stored as XML documents, but larger tables are written in a binary format and stored in the same way as the raw data binary data files. Both the metadata and binary data are received by an online archive component running on the STE, which provides buffering and handles the communication between the STE and the archive for the data storage. For more information on the archiving, please read Chapter 13 of this Technical Handbook.

The Data Capture software is as an interface between the real-time domain of the data taking and the storage side (see e.g., 2006 ADASS contribution by Hafok, Caillat and McMullin). It receives many inputs from other parts of the Control subsystem, as well as input from the Correlator and TelCal subsystems for each subscan and scan, which it compiles together into the ASDM tables. The Control software itself is also compiling values from many hardware devices into the information it provides to Data Capture. At the end of each scan the current table contents pertaining to the scan (an ASDM "slice") are passed to TelCal via CORBA-based services to allow it to produce any necessary calibration results (combined with binary data received directly from the Bulk Data streams), and when TelCal finishes it passes back results for Data Capture to append to calibration tables. When the execution finishes, Data Capture sends the finished ASDM tables to the online archive component for storage (in future the storage may be made more incremental to improve scalability), waiting for the last scan results from TelCal if needed. Data Capture is also responsible for monitoring the storage of the binary data to the archive at the end of the execution and reporting when the data is fully archived or if there was an error.

The archiving at the end of an execution can be carried-out in parallel to further executions, so the next SB from the queue can start shortly after the last scan of the previous execution. This is achieved by having an independent Data Capture component run for each execution, which remains alive until all the data from the execution is archived. At present the array is not allowed to be destroyed until all Data Capture components for executions in the array have completed archiving.

A summary plot of the main elements involved in data flow is shown in Figure 12.2.




Figure 12.2: Data flow components.

⁵see e.g. <https://casa.nrao.edu/casadocs/latest/reference-material/the-science-data-model>

12.2 The ALMA Science Data Model

The ALMA Science Data Model (ASDM; Viallefond, F. 2006, *Astronomical Data Analysis Software and Systems XV*, 351, 627) defines the collection of information recorded during an observation that is needed for scientific analysis. The ASDM is a common standard shared by ALMA and the EVLA. As described above it contains both binary data files (BDFs) and metadata tables. The metadata tables contain links to other tables and references to the binary files in the Archive.

The ASDM contains 16 core tables that are common to all observing modes, and up to 23 additional tables that are only created for specific observations. On top of these, TelCal also creates associated tables whenever it processes any calibrations. All tables are organized with a similar structure, with the columns listing the contents and the rows including the actual values. The core tables have been defined to outline some of the following: hardware characteristics, array configuration, antenna tracking, targets, tuning setups and resulting spectral windows, scan and subscan timing and intents, and flagging information.

The tables produced by TelCal have names prefixed by “Cal” (there is also the `CalDevice` table which is unrelated to TelCal). The list of these calibration tables expands as new observing modes and calibrations become available (see Chapter 10).

The ASDM metadata implementation in the ALMA software uses automatic code generation to provide common libraries for Python, Java and C++ to read and write ASDMs. These are used in the various parts of the software, both online and offline, which need to interact with ASDMs.

The BDF format is a mix of XML headers and actual binary data, encapsulated in multipart MIME. Data recorded in each file can include cross-correlations, auto-correlations, and flags (“BDF flags”). For a typical subscan with a correlator, four BDFs are produced for different types of data: spectral correlator integrations, channel-average correlator integrations, Water Vapour Radiometer (WVR) integrations, and baseband Square Law Detector (SQLD - See Chapter 5) integrations. The BDF structure is designed to support incremental writing. Currently all but the WVR BDFs are written incrementally. Typically one incremental write includes data from one integration or all integrations within a second when the integration time is shorter than a second. BDF file sizes can be up to ~ 100 GB at present.

Most users interact with the ASDMs via the Common Astronomy Software Applications (CASA) data reduction package. CASA uses a different internal storage format called the Measurement Set (MS)⁶. A piece of software distributed with CASA known as the “Filler” is responsible for producing an MS from an ASDM. It is invoked either from the command line as `asdm2MS` or by the CASA task `importasdm`. The filler uses the same ASDM libraries as the ALMA software for parsing.

⁶see e.g. <https://casa.nrao.edu/casadocs/latest/reference-material/measurement-set>

Chapter 13

Data Archiving

13.1 Introduction

The ALMA Archive is composed of three subsystems, the ALMA Frontend Archive (AFA), the ALMA Science Archive (ASA) and the Calibrator Source Catalogue (SC).

The ALMA Frontend Archive is at the center of the ALMA data flow (Figure 13.1). It provides the core persistence functionality. The AFA is a combined database and binary data storage system that is accessed by the different software subsystems through the same software layer. The AFA stores all metadata and data of ALMA including the user accounts, proposals, antenna configurations, monitoring, raw data, and reduced science data. The storage architecture is based on the Next Generation Archive System (NGAS) with Oracle technology for replicating the metadata.

The ASA makes the science data and metadata available to PIs and archival researchers for query and download following ALMA's data access policy (<http://almascience.org/documents-and-tools/latest/alma-user-policies>). An ALMA Science Archive Manual is available from the ALMA Science Portal¹. The ASA holds a small subset of metadata in a relational database and provides access to external interfaces like the Archive Query interface and Virtual Observatory (VO) tools.

The SC provides access to measurements of calibrator observations.

Each of the three ALMA Regional Centers (ARCs) in North America, Europe and East Asia holds a copy of the entire ALMA Archive subsystems for backup, user support and data distribution to the ALMA PIs and archival researchers. The ARCs provide a completely identical user experience to their communities and a user can download data from any ARC.

13.2 Data Flow and Archive

Data from the correlator, together with monitoring and weather data, are sent via dedicated optical fiber links to the OSF, where they are archived. A peak data rate of 66.6 MB/s can be sustained for short periods of time, i.e. days. This peak rate is a technical limitation of the data capture and data flow systems and will be imposed by the Observing Tool at the proposal validation stage.

The Pipeline processing system and Archive storage system have been designed to cope with an average data rate of 10% of the peak rate, i.e. 6.6 MB/s leading to a yearly amount of 200 TB of data. The volume of data products from the pipeline are expected to approximately equal the raw data volume, resulting in about 400 TB/yr being ingested into the ALMA archive.

As soon as data are taken and have passed (or semipassed) the QA0 quality control step (see Section 11.2), their metadata are harvested from the AFA into the ASA and made available for search. This allows archival researchers to see which data will become public in the future. Harvested metadata include all the information needed to describe the observations, including date and time, source coordinates, frequency settings for each spectral window, and spectral and spatial resolution.

The ALMA Archive at the OSF is designed to provide up to a year of temporary storage for the instrumental data (in the form of files in the "ALMA Science Data Model", or ASDM, format) and the monitoring data. The

¹<http://almascience.org/documents-and-tools/latest/science-archive-manual>

instrumental data are then transferred to the main archive at the SCO, where the pipeline is run and from where the data and pipeline products are distributed to the three ARCs. At this stage, a sizeable fraction of the data are processed at the ARCs and then mirrored back to SCO. The process of copying the data to any of the archives involves a replication of the metadata (support data) and of the bulk data (ASDM and FITS files). All data transfer is done over the network. Metadata replication from SCO to the ARCs happens within a few seconds, transfer of bulk data can take longer (up to several hours), depending on the amount of data to transfer to the individual ARC.

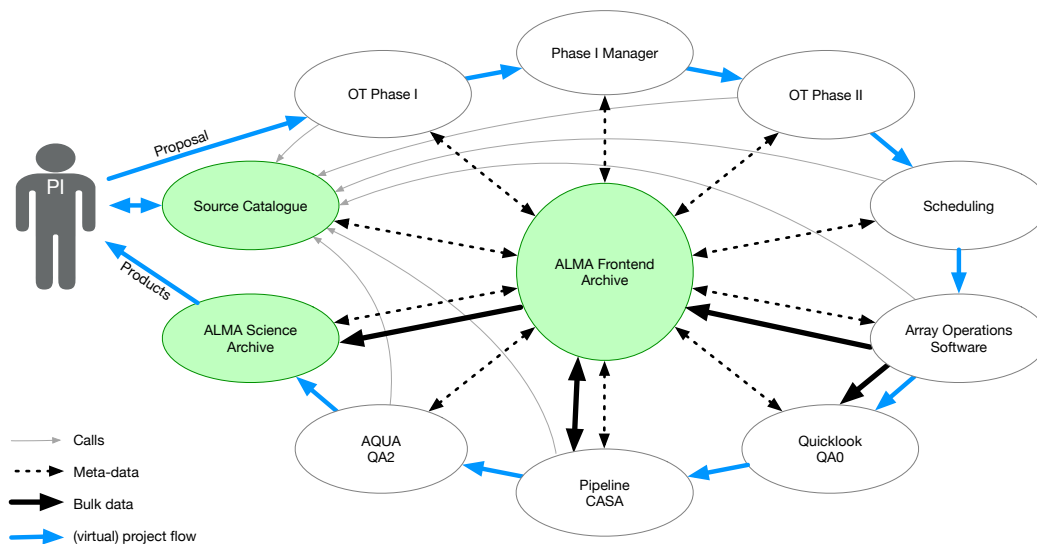


Figure 13.1: The ALMA Frontend Archive is in the center of the ALMA data flow. The ALMA Archive subsystems are shown in green.

Once the pipeline has processed an ObsUnitSet (hereafter OUS, see Chapter 8 for details on the OUS structure) and the science products have passed QA2 they are ingested into the ASA.

PIs of large programs are expected to return processed data products to the project as described in the Cycle 8 2021 Proposer’s Guide. These data will be transferred to JAO from where they will be ingested into the ASA.

13.3 Calibrator Source Catalogue

The Calibrator Source Catalogue (SC) contains information about more than 3000 ALMA calibrators, including position, flux density, structure and polarization. It is updated with measurements from Calibrator Survey observations, and used mainly for the selection of calibrators for bandpass, flux, phase, and polarization calibrations, either through the OT in project preparation or the Online Software during observations.

At <https://almascience.org/alma-data/calibrator-catalogue> the SC database can be accessed through a web-based user interface (see Figure 13.2).

Combined search options are available for sky position (radius or RA/Dec ranges), band or frequencies, flux density, time or polarization to facilitate the selection of calibrators. The search results are presented in a table and a downloadable format, and zoomable views are available for the sky positions and flux density history of the selected calibrator sources.

13.4 PI Data and Data Delegation

The unit of data delivery to the PI is the OUS. As soon as an OUS has passed QA2 and the data products have been replicated to the user’s home ARC, PIs will receive an email notification containing a link to their data and



Figure 13.2: Example of a search for polarized calibrators in an RA/Dec range with recent measurements and a flux density limit for a set of ALMA bands

supporting information. This sending of the notification to the PI triggers the start of the proprietary period of 12 months for standard proposals and 6 months for Director’s Discretionary Time (DDT) proposals. Extensions of the proprietary period can be granted under certain circumstances (see Section 8.4.4 of the ALMA Users’ Policy). When a Group OUS consisting of several Member OUSs is processed (for example, a combined TP, 7-m and 12-m Array observation), the Group OUS products are released separately, with their own 12-month proprietary period.

The data deliveries consist of one or more bundles of raw (ASDM-format) data and one "products" bundle tar file which includes the FITS files, logs, scripts, QA information and calibration tables.

Often PIs want to make their proprietary data available to collaborators, e.g. Co-Is. To this end, a data delegation service is available so that PIs do not have to give away their Science Portal password to anyone. Instead, PIs can give access rights to the data of a project to any registered ALMA user. To do so, PIs need to log into the Science Portal, go to their user profile page in the top right corner of the Science Portal page and then add delegates in the "Project delegation" tab.

13.5 Archive Query

At <http://almascience.org/alma-data/archive> users can query the holdings of the ASA (Figure 13.3). They then can download the data corresponding to their queries, if those data are public or if the users are authenticated and have the proper access rights to those data.

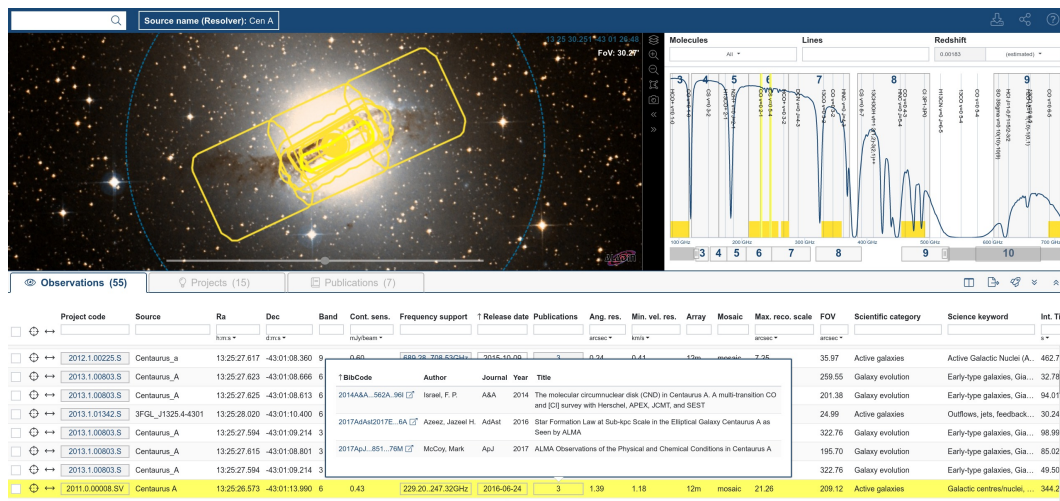


Figure 13.3: The ALMA Science Archive Query Interface

Queries can be made by physical quantities along the Position, Energy, Time and Polarisation axes. Help on how to query is provided in the tooltips of the query fields as well as through the "Query Help" link. Query constraints using standard operators for strings (*, ?) and numbers (>, <, ..) can be placed into the form fields. There is also a name resolver available for non-solar system objects (Sesame), which queries the Simbad, NED and Vizier databases. No operators can be used in the name resolver field.

By default, the Archive Query Interface (Figure 13.3) will present users with the metadata of the observations, although they can choose to also see the metadata per project or per publication. Users can sort and subfilter the results and add or remove columns from the result table to narrow their search. The sky coverage and the spectral coverage of the observations are visible in a graphical windows.

The interface also offers to download the results in VOTable, TSV (Tab-separated values) or CSV (Comma-separated values) format for further processing in tools like topcat². Astroquery³ can be used to encapsulate and simplify programmatic access to the ASA. Neither Topcat nor Astroquery are part of the official ALMA software, and support of users of these tools will be limited to best efforts.

ALMA offers Virtual Observatory (VO) access through the TAP (Table Access Protocol) and SIAPv2 (Simple Image Access Protocol Version 2) at the locations <https://almascience.org/tap> and <https://almascience.org/siap2>.

13.6 Request Handler

Once data of interest are defined, they can be selected via checkboxes on the interface and submitted to the ALMA Request Handler for download. The data are displayed in the full OUS hierarchy. This allows users to check immediately if there are data from other Member OUS or higher-level Group OUS products available. By default, only data products (i.e. images and cubes as well as ancillary processing documentation) will be selected for download. Users who want to download the raw data as well are asked to select "include raw" before hitting the Download button or to select the desired data manually. The sources row can be expanded to show all the source names that are part of the package. Furthermore, a readme file is provided summarising the structure of the data and providing the full links to the OUS. For recent observations, data files can be downloaded either as complete tar files or individually, file by file. A small arrow in front of some tar-files allow the user to list the contents.

²<http://www.star.bris.ac.uk/~mbt/topcat/>

³<https://astroquery.readthedocs.org/en/latest/alma/alma.html>

Five possibilities exist for the download itself:

- The first option is to use the download script. This script runs under Linux and MacOS and downloads files in parallel streams and is also adapted for downloads to a processing environment where no web browser is available.
- The second option is to use the download manager applet with the FireFox browser. This method is very convenient and allows for parallel downloads, too. It requires a Java Browser Plugin.
- It is also possible to use the download manager through Java Webstart technology.
- Moreover, a page with the links to all selected files can be displayed which then can be conveniently downloaded, e.g. using a browser plugin like "DownThemAll".
- Finally, it is possible to just use the right-click and "Save as ..." functionality of the web browser.

Only the data deliveries that users have permission to download can be selected. If they are not authorised to access any data delivery in the request, no "Download selected" button appears.

If the user is authenticated before requesting data for download, the request will be stored. This allows users to go back to previous requests. Note that these requests are stored only at the ARC the user is currently accessing. If data should be downloaded from a different ARC, then a new request has to be issued. For very large data requests PIs or archival researchers have the possibility to ask via the ALMA Helpdesk for data delivery on hard media i.e. USB hard-disks. The ARCs may have different policies regarding the details of the shipping of the hard-disks.

Appendix A

Antennas

A.1 Design and Properties

ALMA is composed of a total of 66 antennas, 54 with a diameter of 12 m and 12 with a diameter of 7 m. The four 12-m antennas used for total power observations and the twelve 7-m antennas together form the Atacama Compact Array (ACA). The ALMA antennas are manufactured by three different contractors. These are VertexRSI (North America) which provided 25 12 m antennas, Alcatel Alenia Space European Industrial Engineering MT Aerospace (AEM, Europe), which also provided 25 12 m antennas and Mitsubishi Electric Corporation (MELCO; East Asia), which provided the four 12 m total power antennas and the twelve 7 m antennas (Figure A.1).

All antennas have been designed to meet very stringent ALMA performance criteria, and to successfully operate under the extreme environmental conditions at the Array Operation Site (AOS), i.e. strong winds, large temperature ranges and gradients, solar irradiation and snow. The primary operating conditions are the following:

- Range of Ambient Temperatures: $-20\text{ }^{\circ}\text{C} \leq T_{amb} \leq +20\text{ }^{\circ}\text{C}$
- Gradient of temperature: $\Delta(T_{amb}) \leq 0.6/1.8\text{ }^{\circ}\text{C}$ in 10/30 minutes
- Wind Velocities $\leq 6/9\text{ m/s}$ (day/night)
- Full solar loading

The antennas have the following specifications within the Primary Operating Conditions:

Antenna Surface: RMS deviation of 25 (20) microns or less for 12-m antennas (7-m antennas) relative to an ideal parabola.

Pointing: Absolute pointing ≤ 2.0 arcsec all-sky. Offset pointing ≤ 0.6 arcsec within a 2 degree radius on the sky.

Primary Beam: The total power pattern response of each ALMA antenna shall be determined to a measurable and repeatable precision better than 1% at frequencies < 400 GHz and 2% at frequencies > 400 GHz.

Subreflector: 6 degrees of freedom to allow for alignment with the corresponding receiver beam.

Subreflector Motion: Maximum horizontal (X) and vertical (Y) displacements of ± 5 mm. Maximum focal displacement (Z) of ± 10 mm. The maximum rotation around the axes is 1.2 degrees. Positioning must be accurate to 5 microns.

Antenna Location: The phase center position of the ALMA antenna shall be determined to a radial precision of 65 microns (including the antenna structure and pad), stable over two weeks.

Configuration: The ALMA antennas shall be relocatable.

Lifetime: a minimum of 30 years.

Antennas used by ALMA have both 12 meter and 7 meter diameters, with the receivers mounted at the secondary (Cassegrain) focus. The 12-m antennas have a focal length of 4.8 meters, but the distance from the secondary focus to the plane of the subreflector of the 12-m antennas is 6000 mm, giving an effective focal ratio $f/8$, with an effective secondary focal length of 96 m and a plate scale of 2.15 arcsec per mm. The subreflector has a diameter of 750 mm. The 7-m antennas have a focal length, to the primary focus, of 2.572 meters. Given an effective focal ratio $f/8$, an effective secondary focal length is 56 m. The subreflector has a diameter of 457 mm.

The main reflectors of the ALMA 12-m and 7-m antennas are composed of individual panels. The size and number of panels varies between the different types of antennas:

VertexRSI: 264 panels spanning 8 rings with 12 (rings 1 and 2), 24 (rings 3 and 4), and 48 (rings 5 through 8) individual panels which are roughly a half-meter-square in area.

AEM: 120 panels spanning 5 rings with 8 (ring 1), 16 (ring 2), and 32 (rings 3 through 5) individual panels which are roughly one-meter-square in area.

Melco 12-m: 205 panels spanning 7 rings with 5 (ring 1), 20 (rings 2 and 3), and 40 (rings 4 through 7) individual panels which are roughly one-meter-square in area.

Melco 7-m: 88 panels spanning 5 rings with 4 (ring 1), 12 (ring 2), and 24 (rings 3 through 5) panels which are each roughly one-meter-square in area.




Figure A.1: The four different ALMA Antenna designs: Vertex 12-m, MELCO 12-m, AEM 12-m, and MELCO 7-m (from left to right).

Each panel has up to 5 adjustment screws, which can be used to optimize the surface accuracy of the individual antennas (based on holographic measurements). The surface of the panels are etched to scatter optical and near infrared solar radiation.

The antennas are equipped with a movable aluminium subreflector. Subreflector adjustment is used to maximize the transfer of power into the receivers by compensating for changes in the focus position due to gravitational- and temperature-induced deformations. The backplane of the subreflector is attached to a hexapod that controls its position and orientation. The hexapod has six degrees of freedom, displacement and tilt around the three axes, horizontal (X), vertical (Y) and along the optical axis (Z).

All antennas have a Cassegrain cabin that is kept at a constant temperature of 20 degrees Centigrade and contains the receivers, the amplitude calibration device and associated electronics.

A shutter protects the inside of the Cassegrain cabin when the antenna is not operating. A membrane transparent to the frequencies that can be observed with ALMA is located below the shutter to prevent airflow from the cabin to the outside when the shutter is open. The current design uses a 0.5 mm thick Goretex membrane.

	BUS	Number Rings/ Panels	Panel Mate- rial	Quad type ¹	Cabin	Drive System ²	Metrology System ³
Vertex	CFRP Al Invar	8/264	Al	+	Steel	Gear	4 linear displacement sensors + 1 two-axis tiltmeter (above the azimuth bearing)
Melco 12-m	CFRP	7/205	Al	+	Steel	Direct	Reference Frame metrology
Melco 7-m	Steel	5/88	Al	+	Steel	Direct	Thermal (main dish), Reference Frame metrology
AEM	CFRP Invar	5/120	Nickel Rhodium	x	CFRP	Direct	86 thermal sensors + 2 tiltmeters in yoke arms

Table A.1: Design Properties of the Different ALMA Antennas.

Notes: **1** Shape of the quadrupod supporting the subreflector as seen looking along the optical axis of the antennas when they are pointed to the viewer. **2** A gear drive consists of a main motor driving a series of connected reduction gears (i.e., gearbox) that do the actual precision work. A direct drive system does not require such gears and takes the power directly. The direct drives used in ALMA antennas are magnetically supported. **3** Jointly used to correct in semi-real time the pointing of the antennas, under a wide range of environmental conditions, to meet the ALMA specifications.

The different antennas use a combination of steel, aluminium, Carbon Fiber Reinforced Polymer (CFRP) and Invar to achieve the best compromise between stiffness, robustness, smoothness, and low thermal expansion (see Table A.1 for a summary of properties). Common to all antennas is that they have a steel pedestal.

All antennas have built in metrology systems which allow thermal and wind deformations to be computed and corrected. For these purposes, the antennas are fitted with thermal sensors, linear sensors and inclinometers (tiltmeters).

The Vertex antennas have a drive system that is gear-driven whereas the AEM and MELCO antennas have magnetically supported direct drives.

The antennas are controlled using the ALMA Control Software (ACS). ACS sends instructions to the Antenna Bus Master (ABM) computer, which are then sent to the Antenna Control Unit (ACU) through a CAN bus.

A.2 Antenna Foundations

The antennas are placed on specially-designed concrete pads to guarantee stable orientation and location (Figure A.2). All antennas are attached to the pads at three points at the vertices of a triangle. The three points (inserts) are located on a circle centered at the antenna pad with a spacing of 120 degrees.



Figure A.2: Structure of an antenna pad (actual pad at the OSF) (left) and detail of antenna anchored to a pad (right).

This interface guarantees a position repeatability error of the antenna, considered as a rigid body, not exceeding the values below:

- X/Y plane < 2 mm (peak to peak)
- Rotation around Z < 30 arcsec (peak to peak)
- Parallelism with respect to Z ± 10 arcsec with respect to Zenith

The minimum stiffness which the foundation must exhibit at each insert is:

- Vertical stiffness (Z) $> 13 \times 10^9$ N/m
- In X/Y plane $> 9 \times 10^9$ N/m

Stiffness includes the inserts, the concrete pad and the soil. This does not include the kinematic mount lower part nor the foot of the antenna. The position of the pads are measured to a precision of 65 microns, and then monitored for stability for over two weeks. The pads are equipped with two vaults that contain the power, communication, Local Oscillator (LO) and data transmission cables that are connected once the antenna is placed on the pad.



Figure A.3: The ALMA array with eight 12-m antennas (left), and an antenna being transported to the AOS (right).

A.3 Antenna Transportation

Antennas are moved from one pad to another using a specially-designed transporter (Figure A.3, righthand panel). ALMA has two of these vehicles. They are 20 meters long, 10 meters wide and 6 meters high, and each has 28 tires. The transporter positioning system performs a fine positioning of the antenna before setting it down on the foundation in the 3 in-plane degrees of freedom (x, y, rot-z) and 2 in-tilt (rot-x, rot-y). Adjustment in each of the 5 adjustment axes can be done independently. The adjustment range of the antenna positioning system compensates for the inaccuracy of the vehicle position with respect to the antenna foundation (which must be smaller than 10 cm) to achieve the required antenna positioning accuracy. The antennas can be positioned to within a few millimeters, ensuring accurate placement on the antenna foundation pads. More information on the transporters can be found on the ALMA EPO pages¹.

¹<http://www.almaobservatory.org/en/technology/transporters>




Figure A.4: Side view of ALMA front end showing cryostat assembly, with room temperature unit below.

A.4 Cryostat

The ALMA front end consists of a large closed-cycle 4 K cryostat containing individual cold cartridge assemblies (CCA) with mixers and LO injection for each band, along with room temperature electronics for the IF and LO for each band (the warm cartridge assembly, WCA) and fore-optics and entrance windows for each band. The water vapor radiometer (WVR) is mounted to one side of the cryostat using a pickoff mirror to direct the antenna beam into the WVR. The Amplitude Calibration Device (ACD) is mounted above the front end, and is described in Section A.5. Figures A.4 and A.5 show overviews of the front-end unit, with the cylindrical cryostat on top and the room temperature electronics beneath.

All of the receiver cartridges are in the same cryostat, with the mixers thermally-coupled to the same 3-stage Sumitomo cryocooler (Figure A.6). The three stages have nominal temperatures of 4 K, 15 K and 110 K. To avoid overloading the cooler, only three bands can be switched on at a time. It takes about 1 minute to switch between any of the bands that are switched on at a given time. For bands that are off, the time to fully thermally-stabilize them from an off state is 15 minutes – this is mainly to ensure a flat bandpass shape. All of the receivers are mounted off-axis to avoid extra rotating band-selection mirrors, which necessitates a pointing offset of the antenna to change band. The band pointing offsets are known and well-measured; the reference band for pointing is Band 6, and all offsets are with respect to this band. The four higher-frequency bands (Bands 7-10) are mounted close to the central boresight to minimize aberrations.




Figure A.5: Bottom view of ALMA front end, showing WCAs.

A.5 Amplitude Calibration Device

The ALMA specification for relative amplitude calibration repeatability² has been set to be better than 1% for frequencies below 300 GHz and better than 3% for all other frequencies covered by the ALMA Front End. To achieve this goal, ALMA has adopted a two-load amplitude calibration approach.

The Amplitude Calibration Device (ACD) is located above the cryostat. It consists of a robotic arm attached to the top plate of the front end (Figure A.7). The arm holds two calibration loads, one at ambient (i.e., receiver cabin) temperature and the other one maintained at 80 °C (353 K). In addition, this arm also holds a solar filter to attenuate solar radiation during observations of the Sun. The arm is designed to allow the two loads to be placed in the path of any of the receiver beams (Figure A.8). Typically it takes 2 seconds to move the arm from the park position to the position where one of the loads is in the beam, and also 2 seconds to change between loads.

To accurately calibrate radio astronomical data to a temperature scale, the actual brightness of the two loads has to be precisely known. Critical to this calibration precision is the coupling of the load to the beam of a given band. This coupling must be very good at any telescope elevation and free of reflections of the load emission. This is because any reflection from the loads back into the cryostat would be terminated at a different temperature and would cause standing waves. Both loads have thus been designed so that the actual effective brightness temperature and that computed from the measured physical temperature (with sensors embedded

²“Calibration Repeatability” means being able to make repeated measurements of the same flux densities (or brightness temperatures) for the same source under different conditions (weather, telescope elevations, front-end status, etc.).




Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view).




Figure A.7: Lateral view of the ACD on top of the ALMA front ends.




Figure A.8: Top view of an ALMA front end showing the robotic arm of the ACD retracted during normal observations or on top of one of the front-end inserts for calibration. The current design has been improved by placing all the loads in a wheel.

in the loads) using known emissivities differ by, at most, ± 0.3 K and ± 1.0 K for the “ambient” and “hot” loads, respectively. This requirement also sets a limit to the fluctuations and departure from the set temperature that are allowed for the “hot” load. Furthermore, the return loss specifications for these loads are -60 dB and -56 dB, respectively.

A.5.1 Atmospheric Calibration Procedure

The ACD is used to measure the receiver temperature and the sky emission by comparing the signals on the sky, ambient and hot loads. This is known as atmospheric calibration (ATM calibration), and is required to correct for differences in the atmospheric transmission between the science and the celestial amplitude calibrators. Normally ATM calibration is done during observations, both near the science target, as well as near the amplitude calibrator.

Traditionally, most mm and submm observatories have used the single-load calibration method, but several simulations have shown single-load calibration is not capable of reaching the relative amplitude calibration accuracies required by ALMA at all of its observing frequencies. However, that method has the very desirable feature that it is only weakly dependent on the opacity of the sky at the time of the observations. A method, using the two calibration loads within the ACD, has been devised in the past to try to achieve the same weak dependence on the opacities at the time of the observation. This method (“the α method”) uses the voltage outputs from the observations of both loads to simulate a single load with a brightness temperature close to that of the atmosphere at the observing frequency. This fictitious single load is defined as a weighted sum of the voltages of the “hot” and “ambient” loads so that the temperature calibration factors are almost independent of the optical depth. The fictitious load voltage output, V_L , is defined as:

$$V_L = \alpha V_{L_1} + (1 - \alpha) V_{L_2} \quad (\text{A.1})$$

where α is the weighting factor, and V_{L_1} , V_{L_2} the output voltages when the two loads are measured. From this definition and some algebra, one can find the optimum weighting factor needed to minimize opacity dependency, and the corresponding resulting calibration factors are:

$$\alpha = \frac{\eta J_M + (1 - \eta) J_{SP} - J_{L_2}}{J_{L_1} - J_{L_2}} \quad (\text{A.2})$$

$$T_{Cal} = (J_{M_s} - J_{BG_s}) + g\eta e^{\tau_s - \tau_i} (J_{M_i} - J_{BG_i}) \quad (\text{A.3})$$

where η is the forward efficiency of the antenna, g the sideband ratio, τ the opacity, and J_M , J_{SP} , J_{L_1} , J_{L_2} and J_{BG} are the emissivity temperatures of the average sky, the spill-over, the two loads and the background radiation, respectively. The subscripts s and i represent the signal and image bands, respectively. The system temperature is then derived using the formula:

$$T_{Sys} = T_{Cal} \frac{V_{Sky}}{V_L - V_{Sky}} \quad (\text{A.4})$$

For ALMA it has been found that with the current system, the non-linearities are the dominant source of error for this calibration. The system electronics and SIS mixers are not fully linear and dominate the relative amplitude calibration accuracy.

A.6 Water Vapor Radiometers

In the mm and submm regions, variations in the water vapor distribution in the troposphere that move across an interferometer cause phase fluctuations that degrade the measurements. ALMA uses the so-called “Water Vapor Radiometry” technique to correct for these phase fluctuations. Water Vapor Radiometry involves estimating the excess propagation path amount due to water vapor along a given line-of-sight by measuring the brightness temperature of the sky at frequencies near the atmospheric water vapor resonances. These temperatures can then be transformed into a path length and the difference between any pair of antennas in the array gives the final phase fluctuations to be corrected for a given baseline. ALMA has implemented this technique by placing a Water Vapor Radiometer (WVR) on each 12-m antenna (the 7-m antennas do not have WVRs). For the

WVRs to be effective, the measurements have to be taken with a cadence that is fast enough to map the actual variations in the atmosphere. The relevant shortest timescale is the antenna diameter divided by the wind speed as the path delay is averaged over the whole antenna beam and cannot therefore be corrected at any finer time resolution than that. The effective diameter is about 10 m for the ALMA antennas and the relevant windspeed is usually 10 m/s or a bit less so the fastest necessary sampling speed is 1 Hz. On timescales shorter than this 1 Hz timescale, the water vapor path fluctuations are expected to lead to small apparent pointing fluctuations which are analogous to the seeing effects in single-aperture optical telescopes. ALMA selected the 183 GHz line because it is quite bright and allows a more compact design than would the 22 GHz water line. It was decided to measure the temperature of the 183 GHz line in four regions offset from the center using filters of different bandwidths. The positions of the filters are indicated as blue boxes superimposed on the profile of the water vapor line in Figure A.9. The sensitivity specification for the WVRs is 0.08–0.1 K per channel RMS.




Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

It is very important that the WVR illuminates the same area of the sky as the ALMA band receivers in the near-field region. This is because the origin of the water vapor fluctuations is usually located in the lower troposphere (i.e., near the observatory), with one to several layers of water vapor clumps encompassing a wide range of sizes. Since the ALMA backends are located at the Cassegrain focus, an offsetting optical system (see Figure A.10) had to be designed to allow the WVR to measure along the optical axis of the antennas.

The WVRs are only able to detect the variations in atmospheric brightness temperatures due to the “wet” atmosphere (i.e., the precipitable water vapor). There are also variations due to the changes in bulk ambient temperature at different heights above the observatory. It is expected that these could become significant during day time and some techniques are being currently studied to try to measure them (including thermal sounders of the atmosphere that use the profiles of the emission of the oxygen molecules). The brightness temperature variations of the sky that the WVRs have to detect are sometimes quite small, so the quality of the receiving system becomes very important. In fact, the current specification for the ALMA WVRs is that they need to allow corrections of the path fluctuations (in μm):

$$\delta L_{corr} \leq \left(1 + \frac{w}{1mm}\right) 10\mu\text{m} + 0.02\delta L_{raw}. \quad (\text{A.5})$$

where w is the precipitable water vapor (PWV) content in mm along the line of sight, and L_{raw} the total fluctuations observed at any given time. Therefore, this formula includes the expected error of about 2% in measuring the total fluctuations, and states the total resulting path errors after correction (L_{corr}). For a 1 mm PWV, the residual term in the formula would be 20 μm . The stability specification for the WVRs is very stringent (0.1 K peak-to-peak over 10 minutes and 10 degree tilts). To achieve this, a Dicke-switching-radiometer approach was adopted. The input into the mixer is switched periodically (5.35 Hz) between two calibrated loads




Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

(the “cold” and “hot” loads at 293 K and 351 K, respectively), and the sky using a rotating vane embedded in the light path as shown in Figure A.11.

Calibration of the measurements is done following the usual method for a 2-load system. The ratios of the output powers when observing the “hot” and “cold” loads can be used to determine the receiver temperatures. Furthermore, these output powers from the loads are also used to extrapolate to a virtual load that has a brightness temperature similar to that of the atmosphere. The specification for the absolute accuracy of the calibration is 2 K (maximum error). The mixer system is an un-cooled DSB Schottky diode pumped by an LO at 15 GHz that undergoes 2 stages of multiplication. The receiver noise temperature is about 1000 K. After amplification, the IF signal is split into four complete chains (one per filter) and a bandpass filter is applied to select the four desired sampling regions in the profile of the water vapor emission line. In each IF chain, the signal is detected with diodes and after a Voltage-to-Frequency conversion, sent to the Control section for accumulation and control. There is a possibility of LO leakage out of the WVRs that could affect the ALMA receivers in the same antenna and others nearby. To avoid coherence, all the WVRs are tuned to a frequency slightly different (offsets by consecutive integer multiples of 10 kHz up to the total number of WVRs available). The final products sent to the ALMA Control system are time-stamped, calibrated measurements of the brightness temperatures in the 4 filter regions. The path length error due to the PWV can be calculated from these brightness temperature measurements and used to correct the data. Corrections at the scales of the sampling rates of the WVRs are possible at the correlator and refinements for longer timescales can also be done offline in CASA using the `wvrgecal` tool. Currently both streams of data are being recorded (WVR phase corrected and WVR non-corrected) in order to rigorously assess effectiveness.




Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

Appendix B

The LO and IF System

In this Appendix, the signal path, LO chain, and the relation between these systems and the observer's spectral setups are described. To the system, a spectral setup effectively consists of the settings of the local oscillators and correlator in the system such that each spectral window (spw) covers the desired lines and/or continuum frequencies. To the end-user, the spectral setup is normally defined in the Observing Tool (OT) just in terms of the observing frequencies and spectral resolutions, and there is no need to worry about the details of each LO setting. For full details of the OT and how to use it, see the OT User and Reference Manuals, available from the ALMA website¹ (and also in the OT itself).

The following sections show how the LO system works. For those only interested in the spectral setups and not the details of the components in ALMA, please see Chapter 6.

B.1 Functions of the LO and IF System

In the signal path from the front end to correlator, ALMA uses three frequency conversions. The associated LO and IF systems perform multiple functions:

1. Down-conversion of the sky frequencies to basebands in the range 2–4 GHz, which then alias down to 0–2 GHz for digitization.
2. Amplification and adjustment of the correct power levels into the digitizers.
3. Adjustment of the spw center frequencies (in the Correlator FDM modes) within the basebands. This is actually done in the correlator using the TFB LO, but can effectively be treated as a 4th stage of the LO system.
4. Application of frequency corrections for fringe rotation, and compensation for the slight differences in the Doppler shifts at each antenna due to the differential line-of-sight velocities with respect to the target.
5. Provision of geometric delay corrections.
6. Separation of the signal and image sideband. In the case of DSB receivers, selection of the wanted sideband(s). This is done through frequency offsets and phase modulation at each antenna using Walsh patterns.
7. Suppression of spurious signals and reduction of the effects of DC drifts in the samplers. This is done using phase modulation of the LOs using Walsh patterns.

Frequency down-conversion therefore effectively occurs in four stages: two hardware Local Oscillators (LO1 and LO2), a 4 GHz sampler/LO and a digital LO synthesised in the tunable filterbanks (TFBs) in the Correlator². Chapter 6 shows how to setup the system to observe spectral lines (particularly multiple spectral lines)

¹<http://almascience.org/documents-and-tools/>

²Note that the ACA correlator is designed to appear like the 64-input Correlator to the end user, although it does not use TFBs in the same way as the BLC

and continuum. Some other aspects of frequency setups are then discussed, including the usable bandwidth, spurious signals, and rules and limitations pertaining to this observing Cycle. An overview of the LO and IF operation in ALMA is given in Sections B.2 and in Section B.3, the hardware and how the LO frequencies are synthesised and distributed around ALMA are described.

B.2 Summary of Operation

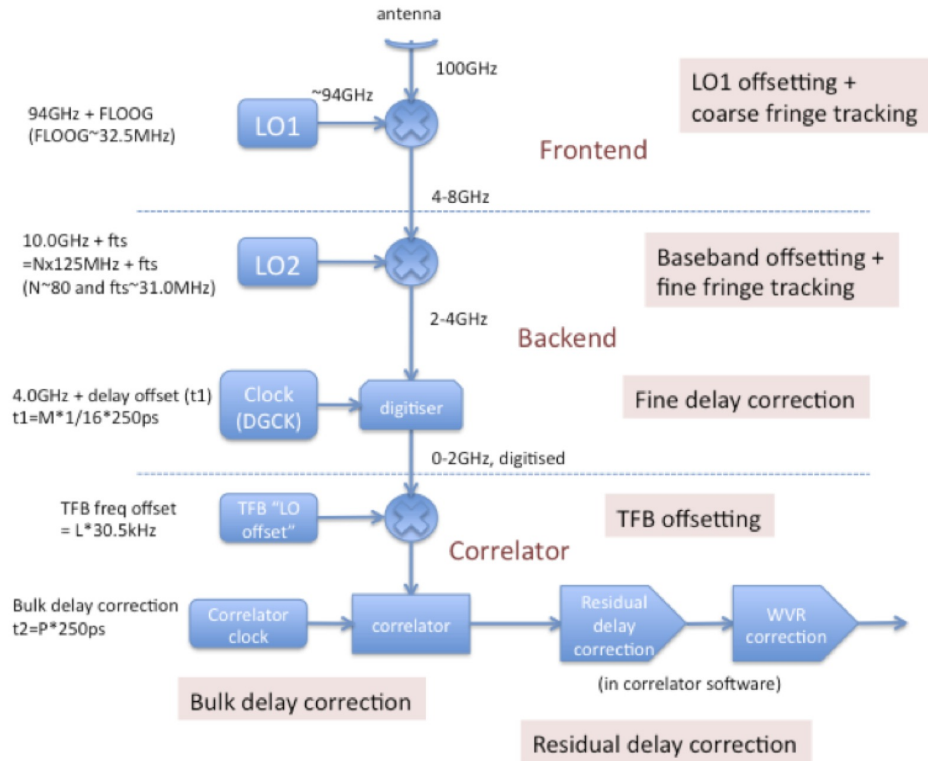


Figure B.1: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the front end, back end, and correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz seen in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized tuning steps, such as LO2 (which can be changed in a multiple (“N”) of 125 MHz plus or minus a finely-adjustable offset of “ f_{ts} ”), the TFB LO (which uses a multiple “L” of 30.5 kHz) and the Bulk Delay Correction (which has steps of 250 ps, with a factor of “P”). See text for descriptions of each stage.

Figure B.1 shows a simplified block diagram of the ALMA LO/IF system, showing example setups for an observing frequency centered on 100 GHz. Referring to this diagram, the system operates in the following way:

1. The front-end mixer uses LO1 to downconvert the observing frequency into an IF range covering up to 4-12 GHz. This wide range is needed to cover the IFs of all the ALMA bands, since the mixers for Bands 3, 4, 7 and 8 have an output IF of 4-8 GHz, Band 6 a range of 5-10 GHz and Bands 9 and 10 a range of 4-12 GHz. Over most of the front-end tuning range, LO1 and the front-end mixer can be used in upper or lower sideband; although at the edges of the tuning band, only one sideband is possible. LO1 consists of a common reference component for all antennas, plus a smaller offset component generated in the FLOOG (First LO Offset Generator) which is different for each antenna (see LO1 Section B.3.4). The FLOOG is used to perform coarse fringe tracking (i.e. rough correction for the small offsets in the observing frequency at each antenna), to offset the LO1 frequencies slightly to suppress internally-generated interference, and for sideband separation or selecting the sideband. It is also used to offset the LO1 phase in conjunction

with a Walsh switching pattern on the antennas. A 180 degree phase offset is used to remove spurious signals and DC offset while a 90 degree phase offset is used for sideband suppression and, in the future, for sideband separation on the DSB receivers (Bands 9 and 10).

2. In the back end (BE), the IF processor (IFP) splits the IF into basebands, each with a frequency range of 2-4 GHz, via a set of filters and tunable second LOs (LO2) (see IFS/IFP in Section B.3.5). LO2 is used to offset the individual baseband frequencies within the IF range. The LO2 and second mixer only operates in LSB, with a possible LO2 tuning range of 8-14 GHz. The LO2 signal itself is generated by a coarse synthesiser which can be set only in steps of 125.0 MHz, plus or minus a second fine-tuned synthesiser (fts) which provides an offset in the range 20.0-42.5 MHz (marked as "fts" in Figure B.1). The limited fts range and the 125 MHz quantization means that LO2 setting is not fully contiguous; consequently there can be up to ~30 MHz difference between the desired and the set value. With a single baseband, this can be compensated by a suitable offset of LO1, but with multiple basebands this is not always possible. So without additional correction, setups with multiple basebands could have the requested lines offset from the spw center by up to 30 MHz. However, the remaining differences in the different spws are compensated by applying an opposite offset to the TFBLOs (LO4 - see below.)³. An algorithm used by the OT and the realtime system generates the best LO tuning "solution" for LO1, LO2 and LO4 which minimizes the offset of the requested observing frequencies from the centers of the spws. Other uses of LO2 are that the finely-tunable fts is used for fine fringe tracking and LO2 can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband.
3. The 2-4 GHz analog IF signal from the second mixer in the IFPs is digitised (or sampled) with a 4.0 GHz clock (DGCK). A fine delay (or time) offset is applied to this clock in units of 1/16 of the clock period (250 ps) (the "fine delay correction").
4. In the FDM correlator mode, up to 32 digital filters (known as TFBs, or "Tunable Filterbanks") are applied to each digitised baseband signal, each of which can be individually adjusted across the baseband frequency (the TFB offsetting). This is effectively applying a digital LO (the TFBLO, or LO4), which is adjustable in steps of 30.517578125 kHz⁴ and allows the spws to be moved around within the basebands. At Phase 2, the TFB is centered on the baseband if the TFB "offset" is set to the default of 3000.0 MHz; it can be moved up to +/-900 MHz from that frequency, the range depending on the spw bandwidth. The TFB outputs are resampled and sent to the correlator. The TFBLO can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband. Finally the correlator software is used to perform the finest level of residual delay correction.

B.3 Frequency Generation and Distribution in ALMA

Figure B.2 shows a summary of the main units involved in the LO generation and distribution. The LOs are generated by the Central LO (CLO) (Section B.3.1) in the AOS Technical Building (lower half of diagram). A fiber-optic system is used to distribute these signals out to the antennas (Section B.3.1) incorporating a realtime path length correction system (Section B.3.3). In the antennas, the important outputs are LO1 (FE 1st LO) (Section B.3.4), LO2 in the IF Processor (Section B.3.5) and the digitizer clock (DGCK). All of these are required in each antenna, shown in the upper part of Figure B.2. In the following subsections, some of these components are described.

B.3.1 Reference and LO Signal Generation

The Central Local Oscillator (CLO) generates and distributes the reference, timing and LO signals to all ALMA components to ensure that antenna movement, electronics, and data acquisition are synchronized. These signals are distributed to the antennas through optical fiber using the light of three infrared lasers. The ALMA frequency and phase standard is a Hydrogen maser (installed in 2014), known as the Master Frequency Standard (MFS, lower right), which produces a signal at a frequency of 5 MHz. This is fed into the Central Reference Generator

³This is done automatically when the OT generates a spectral setup in an SB from a proposal. However, it is repeated at runtime. See <https://safe.nrao.edu/wiki/pub/ALMA/AlmaLamaMemos/lamaMemo808.pdf> for more details on the tuning algorithm

⁴The Phase 2 OT has an "adjust" button which quantizes the value entered by the user by this unit

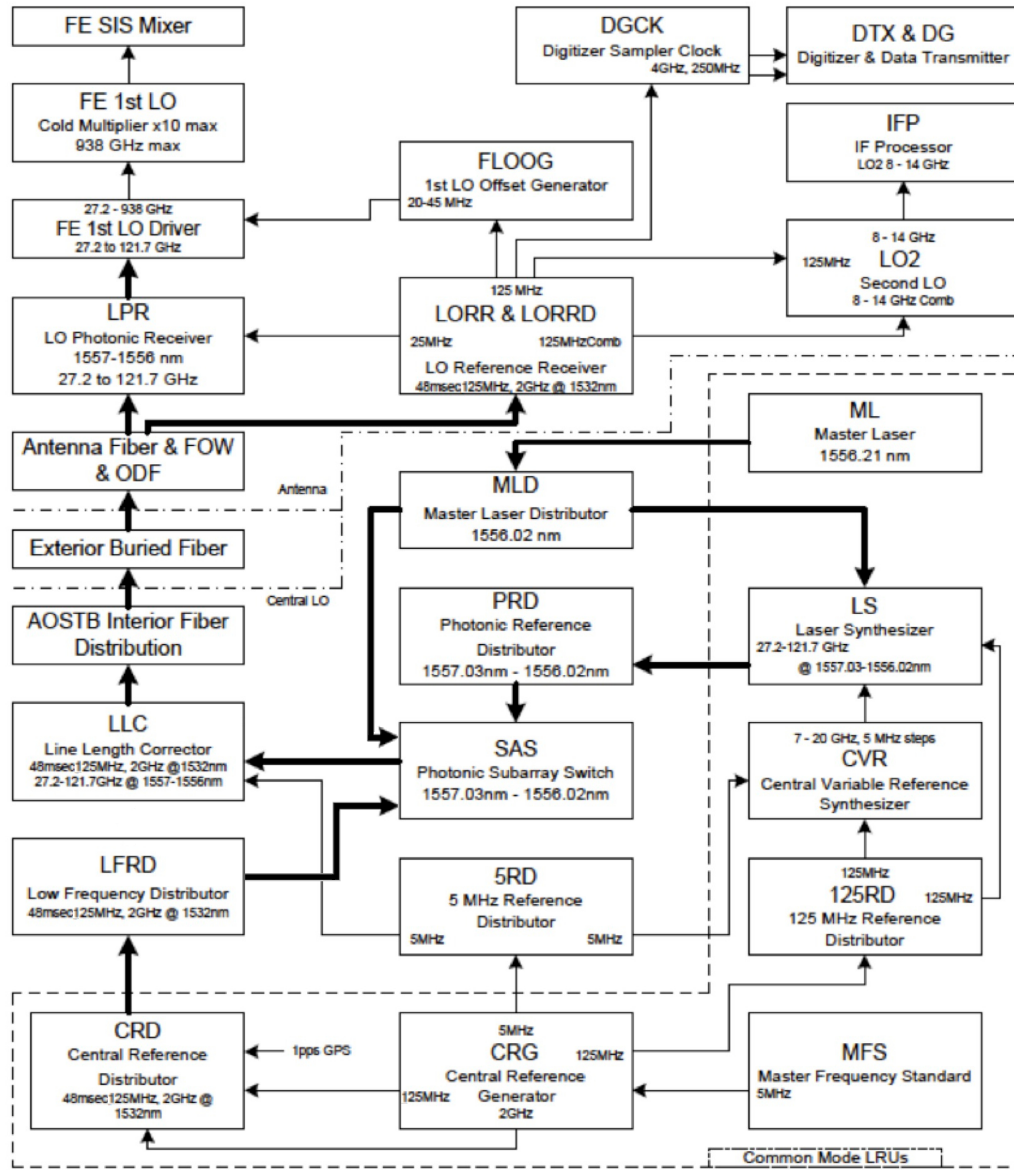


Figure B.2: Summary block diagram of the LO distribution system. The components in the lower section of the diagram (below the dashed-dotted line) form the Central LO (CLO), located in the AOS Technical Building. The components above are located in each Antenna (only one antenna is depicted in this diagram). These are linked by the exterior buried fibers linking the Technical Building with the antenna pads (shown middle-left of the diagram). Thin lines with arrows represent cable distribution, thick lines represent fiber-optic distribution.

module, which produces several signals at multiples of the 5 MHz signal. The 125 MHz signal becomes one of the standards used by many components in the ALMA system. At the AOS Technical building it is used by the Slave Lasers in the Laser Synthesizer modules (Section B.3.2). At the antennas, it is fed into the FLOOG, the Digital Clock (DGCK, see Section B.3.6), the IFPs, and the LO2 Synthesizers. The 2 GHz signal goes into the DTS cards at the antennas. All the reference signals are modulated onto a 1532 nm IR laser in the Central Reference Distributor (CRD) module. The CRD has an internal 48 ms (known as a TE, or Timing Event) clock that is also modulated into the same signal, and is used to synchronize many of the hardware events in the observatory. The modulated 1532 nm signal is sent to an optical distributor (with 80 outputs), the Low Frequency Reference Distributor (LFRD), that feeds it into the Sub Array Switch (SAS) modules, where it is merged with the signals from the Master and Slave lasers (see Section B.3.2). The reference and LO signals are fed to the antennas through the fiber-optic distribution system (see Section B.3.3).

B.3.2 LO1 Signal Generation and Distribution

LO1 is distributed through fiber-optics and regenerated photonically in each antenna front end by mixing the two infrared laser carriers (known as the Master (ML) and Slave Lasers (SL)) to produce a fixed frequency for all the antennas. Figure B.3 shows a block diagram of this part of the LO system. There are five usable (production grade) Slave Lasers (Laser Synthesizers, LSs), each containing four tunable units, that can produce six different LO1 frequencies which allow simultaneous observations at different frequencies with different subsets of the array (multiple arrays or sub-arrays), and may be used for observing modes that will benefit from rapid switching between frequencies (e.g., band-to-band transfer or spectral scans).




Figure B.3: LO block diagram, showing the Central LO (CLO) and the LO section in the WCA in each front end. Yellow lines represent common synchronization and reference signals fed to all antennas, blue lines are LO signals which are individual to each antenna. For a description of acronyms, see text.

The laser frequencies are generated in the CLO in the following way:

- The Master Laser (ML) generates a 1556 nm fixed optical reference signal, which feeds the Master Laser Distributor (MLD) – essentially a 6-way splitter.
- The Central Reference Generator (CRG) produces the 125 MHz IF, and the Central Variable Reference modules produce the tunable reference signals that are fed to the five Laser Synthesizers (LS) used to generate the phase locked laser tone offset by 27 – 121.7 GHz with respect to the Master Laser, and combine the two into a single fiber. The offsets between the ML and SL signals provides the beat note which is used to generate the LO1 frequency in the photomixers in the Warm Cartridge Assemblies (WCAs) in the front end (B.3.4). It is used by the software to set up the front-end observing frequency. With five LSs it is possible to generate five separate LO1 frequencies.
- The Photonic Reference Distribution (PRD) feeds the optical signals to the Sub Array Switch (SAS) module (one per antenna), which chooses the sub array to connect to.

Figure B.4⁵ shows the three laser signals after combination in the Sub Array Switches (SAS). The Master and Slave laser signals have wavelengths of about 1556/1557 nm and the laser carrier signal for the reference signals from the CRD has a wavelength of 1532 nm. The signals are distributed via a single-mode fiber optic line to each of the antennas. The fibers are distributed in buried trenches, and fed into the Cassegrain cabin on each antenna through dedicated Az and El fiber wraps. All are fed through Line Length Correctors (LLCs), which are used to correct for changes in the optical fibers. The LLCs are described in B.3.3 below.




Figure B.4: The ALMA fiber signals. The 1532nm carrier contains the frequency reference signals, and the 1556/1555 nm carriers are used to remotely generate LO1. Within each antenna, the optical signals are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing signals, and the LO Photonic Receiver for the LO Reference signals.

B.3.3 Line Length Corrections

The LO Reference signals are generated at the AOS Technical Building and need to be distributed via optical fibers to all the antennas. In order to guarantee that the phase of the LO signals is stable during the observations for fibers of up to 15 km in length, compensation for changes has to be done in real time. The method adopted by ALMA is based on a round-trip optical interferometer. Phase fluctuations for an optical fiber transmission system are mainly caused by thermal expansion of the fiber and mechanical stresses, which produce birefringent effects and changes in the absolute polarization of the signals. These changes, in turn, cause differential group propagation delays (DGD) that show up as LO phase jitter. The method implemented by ALMA to correct for this is known as the Line Length Corrector (LLC). Part of the LLC can be seen in Figure B.3, and a more detailed block diagram of the system is shown in Figure B.5.

The two-wavelength laser synthesizer signal (master and slave lasers) is adjusted in polarization, combined at the Laser Synthesizer (LS) and then passed through a 3-port polarizing beam splitter assembly (PBS). The polarization is aligned so that all the light passes through the beamsplitter. It then passes through a piezo-driven fiber stretcher assembly and the fiber to the antenna. At the antenna end there is a 3-dB coupler, so that half of the light goes to the turnaround assembly and half to the photomixer in each WCA. The turnaround assembly consists of a fiber frequency shifter (located in the LO Photonic Receiver module) and a Faraday Rotator mirror located within the WCA of specifically the Band 9 cartridge in each front end. The frequency of the signal traveling back to the AOS technical building thus receives a frequency shift of 25 MHz twice, therefore it comes back offset by 50 MHz from the original. The Faraday rotator reflects the signal but turns its polarization angle by 90 degrees to the incident polarization. This means that the outgoing and returning light is orthogonal everywhere along the fiber between the PBS and the Faraday Mirror. Back at the PBS, the returning signal is sent to a third port where it is mixed with a sample of the Master Laser reference signal in a low-frequency photodetector. This results in an output at the 50 MHz offset frequency. This output is compared in a phase detector with a 50 MHz reference signal and the phase of the whole loop is kept constant by a servo driving the fiber stretchers.

The current stretchers can cover ranges up to 5 mm in two modes. A “slow” mode (about 10Hz) copes with the large deformations (about 3 mm, allowing for some headroom at the ends of the ranges) and a “fast” response mode (about 1 kHz) copes with the small range variations (about 0.1 mm). The LLCs are reset to

⁵This figure needs an update. The Slave Laser (variable) wavelength is ~ 1557 nm.




Figure B.5: Block diagram of the Line Length Corrector system for ALMA. The FRM (Faraday Rotation Mirror), shown upper right, is located in the Band 9 cartridge of every front-end. It rotates the polarisation of the incoming light, and the resulting reflected signal is fed back through the buried optical fiber, to be compared with the outgoing signal. This allows the optical path length to be adjusted in a closed loop using the fiber stretchers.

mid-range at the start of every SB execution.

B.3.4 The First Local Oscillator (LO1)

The reference signal required to tune LO1 in the receivers is obtained as the difference of the wavelengths of two phase-locked infrared lasers, the Master and Slave lasers. The Master Laser (ML) has a fixed wavelength of 1556 nm and the tunable Slave Laser (SL) is offset from this; both are generated in the CLO (see Section B.3.2). The offset frequency can be anywhere in the range 27 to 38 and 65 to 122 GHz. The beat note from the two lasers constitutes the Photonic LO Reference; the LO1 reference signal is generated from this by photomixers located in the Warm Cartridge Assembly (WCA) of each receiver. This reference signal is used to drive a YIG (Yttrium Iron Garnet) oscillator operating at frequencies around 10–30 GHz (the exact range depending on the band), via a Phase Locked Loop (PLL) circuit. This produces LO1 for the SIS mixers via two sets of multipliers (see example Figure B.6 for Band 7). The same photonic reference signal is distributed to all antennas in the same sub-array. However, to correct for different delay rates required in different antennas, the First LO Offset Generator (FLOOG) in each antenna generates a small but variable (and different) offset frequency in the range 20–45 MHz which is also fed into each PLL. The FLOOGs for all the antennas are continuously tracked during an observation.

B.3.5 LO2 and the IF Processor Units and IF Switch

The output of each front-end cartridge is connected to an IF Switch unit (IFS) situated in the front end, which selects between bands, provides some amplification, and has variable attenuators to set the output levels. The four (or two) outputs from the IF switch unit are fed into two IF Processor units (IFP), one per orthogonal polarization. Figure B.7 shows a basic block diagram of one IF Processor (only one polarization channel is shown). The Bands 3, 4, 6, 7 and 8 receivers are dual-sideband (2SB), where both the upper and lower sideband signals are provided separately and simultaneously. So there are four outputs from each receiver cartridge in these bands, two sidebands times two polarizations. Each output has an IF bandwidth of up to 4 GHz. For Bands 9 and 10, the receivers are double-sideband (DSB), where the mixer produces a downconverted output

from signals in *both* USB and LSB. These bands have only two outputs, one per polarization, but the signal IF bandwidth of these DSB receivers is 8 GHz per output.




Figure B.6: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier, not shown here, is used to generate the LO1 frequency at 282.9 – 365.1 GHz (in this case, an x3 multiplier). The photonic LO signal (green) feeds a photomixer which creates a beat signal between the ML and SL frequencies. This is mixed with a fraction of the LO from the YIG (x6), and the difference frequency is used in the PLL. The FLOOG also generates a small offset frequency for the PLL, which is different for each antenna. See text for details.

The IF processors divide the incoming IF bands from both sidebands into four 2 GHz basebands and downconvert them to the 2-4 GHz range using the second LO (LO2). Since each baseband is fed by a separate LO2, it is possible to locate them at different frequencies within the IF bandwidth of the receiver (see Chapter 6 and Table 6.2 for limitations). The LO2s are common to both mixer polarizations which means that both polarizations will have the same spectral setups.

The LO2s are digitally-tuned YIG oscillators with a range of 8-14 GHz. LO2 is generated from a harmonic of 125 MHz, plus a fine-tuned synthesiser (fts) of range 20-42.5 MHz, added or subtracted depending on the lock sideband selected by the software. Note that this does not give continuous LO2 coverage, and has to be compensated elsewhere in the LO system.

The IFP unit has 0.5 dB stepped attenuators and Total Power detectors for tuning/optimization of the IF power levels into the digital samplers; these levels are set up at the start of each scan. It is important to note that the switch network layout in the IFP means it is NOT possible to select IF configurations with one baseband in one sideband and three in the other (except for DSB receivers, where this is done using sideband selection). The IFP has anti-alias filters, one set of which is switchable depending on whether the IF range in use is in the upper or lower part of the IF band. As well as downconversion, the LO2s can also be used for sideband separation when combined with the first LO (Section B.1).

The IF processor also has anti-aliasing filters, which define the 2 GHz baseband width and remove out-of-band signals (Section B.3.6). This results in the higher noise levels on the upper and lower 50-100 MHz of channels in the TDM correlator mode (see Section 6.4). These filters cause a decrease in the effective IF range to approximately 1.875 GHz.




Figure B.7: Block diagram of one polarization channel of the IF Processor. This has two IF inputs (at left), and feeds 4 baseband (BB) outputs to the digitizers (to the right of the diagram).

B.3.6 Digitization and Signal Transmission

The outputs of the IF Processor units are fed into the Data Transmission System Transmitter (DTX), that includes digitizers and formatters to convert the signals to optical wavelengths for transmission via optical fibers. There are four DTX units per antenna, each one handling data for a given baseband pair (i.e., the same 2 GHz baseband from each of the two orthogonal polarizations). Each baseband is digitized by a separate digitizer at 4 GHz (i.e., Nyquist sampling for a 2 GHz bandwidth), quantizing each sample into 3 bits (8 levels) per polarization, so that a total of 6 bits must be transferred per baseband pair. The digitized signal is then transferred to the formatter part that packages the data in frames of equal size. The output of each DTX module is fed to three optical fibers, each transporting 2 bits, and the signal leaves the antenna after passing through a Fiber Optic Multiplexer (FOM). All DTX modules are fed with reference/timing signals from an associated Digital Clock (DGCK), which is also used to do the fine delay tracking.

The outputs of the DTX are sent, via the optical fibers, to the AOS Technical Building where the process is inverted (conversion from optical to digital signal) at the DRXs (Data Transmission System Receiver modules), before the signals are sent to the correlator. Delay corrections due to changes in the length of the optical fibers are done using metadata information to realign the frames sent from the transmitting side at the antenna (DTX) and the receiving side at the Technical Building (DRX). Figure B.8 shows a block diagram of a single DTX module.

B.4 Other functions of the LO/IF

In addition to frequency downconversion, the LO/IF performs several other tasks, detailed below.

B.4.1 Delay Corrections

ALMA handles delay corrections via the “Delay Server” software package. It computes the corrections for all the different components involved with a cadence of one minute and distributes them buffered. The three main components along the data flow chain where the corrections are applied are: the First LO Offset Generator (FLOOG), the Digital Clock (DGCK) and the correlator (see Figure B.1). Fringe tracking is done at the FLOOG by slightly offsetting the frequency of the LO1 signal. Currently, the delay handled by the FLOOG is in steps of 250 ps. The FLOOG is also used for phase and frequency switching for suppression and separation of sidebands, and for rejection of internally-generated interference, described in the next subsections.

Fine delay corrections are handled by the DGCK that feeds the corrections into the four DTS modules in each antenna. The delay correction resolution of the fine delay is 1/16 of the 4GHz ADC clock. The bulk delay correction is handled by the Correlator in integer multiples of the 250 ps units. On top of these corrections, the correlator also handles the “residual” delay corrections at much higher temporal resolution ($<250 \text{ ps}/16$) by applying a linear phase gradient across the passband after correlation. Also, the correlator applies relative

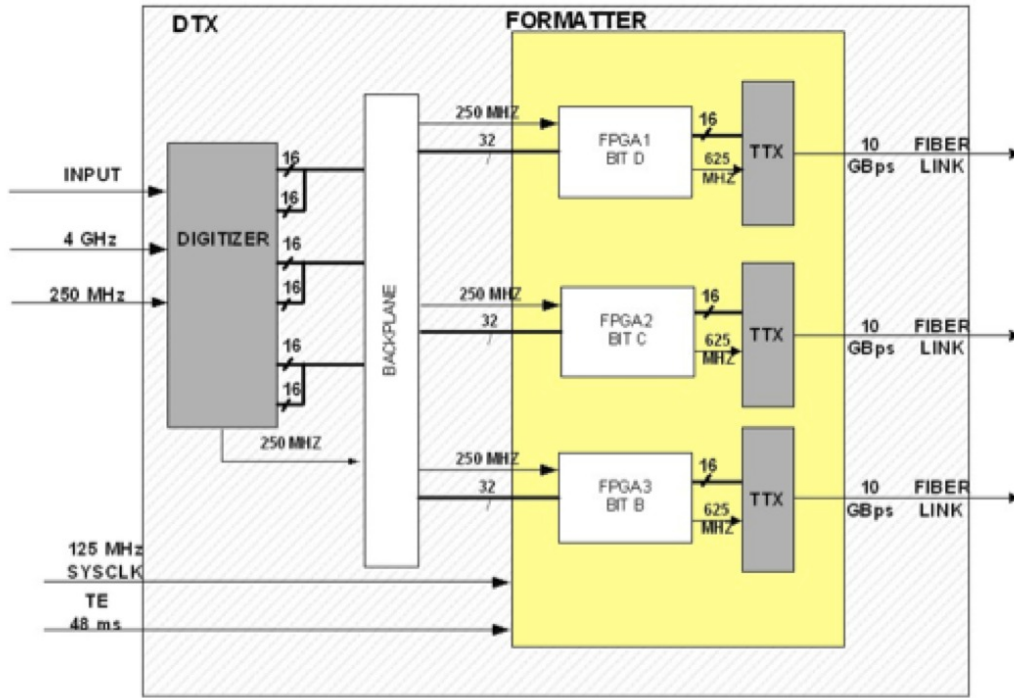


Figure B.8: The Data Transmission System Transmitter (DTX) module, as used in each antenna.

delay corrections between all the basebands and polarizations of a given ALMA band receiver. Currently, the first baseband of the X polarization is used as reference.

B.4.2 Sideband Suppression - LO Offsetting

Some of the ALMA receivers (e.g. Bands 3, 4, 6, 7, and 8) are inherently single sideband (SSB), either through having a mixer or quasioptic design which rejects the unwanted sideband⁶. Their intrinsic sideband rejection is typically only about 10-15 dB, which, although adequate for rejection of the unwanted sky noise, is not enough to remove strong lines from the other (image) sideband. Others receivers (Bands 9, 10) are double sideband (DSB), and the relative response of the two sidebands may not be equal, significantly affecting calibration. Accordingly, additional schemes are necessary for more effective removal of the unwanted sideband (known as sideband suppression), and for correlation of both sidebands independently (or sideband separation - see next section). Sideband suppression in ALMA is done using the FLOOG, and either LO2 (2LO offsetting) or a combination of LO2 and LO4 (3LO offsetting). A small frequency offset F_o is added to LO1 and subtracted from the other LOs, so that while the signal sideband remains at the same frequency, the image sideband is shifted $2F_o$ away from its nominal value. A different value of F_o is applied at every antenna (the offsets are defined using a Walsh pattern), so that all signal sidebands are at the same frequency, but all image sidebands are at slightly different frequencies and no longer correlate.

Note that each of the basebands has an independent LO2 and LO4. So by setting the sign of the offset in (LO2+LO4) differently, each baseband can be set up to observe in a different sideband.

For single-dish observing, such interferometric sideband rejection methods cannot be used, and a frequency scanning method is under development which will allow image rejection.

B.4.3 Interference Rejection - 180 degree Phase Switching

The FLOOG is additionally used to reject spurious signals prior to digitization by applying 180 deg phase switching according to orthogonal Walsh function patterns, with a pattern cycle time of 16 ms. The Walsh

⁶Although they have mixers to allow both sidebands to be observed separately and simultaneously

pattern is different on each antenna, and is demodulated by a sign change within the DTS; as a result, the wanted signals correlate, and the unwanted signals are canceled out. This rejects spurious signals generated in the system between the receiver and the sampler, and also suppresses sampler DC offsets. 180-Walsh is a default setup for all observations. For further details please see Section 5.5.4.

B.4.4 Sideband Separation - 90 degree Walsh Switching

It is possible to apply a 90 deg phase switch in the FLOOG and in the correlator processing, allowing correlation of the upper and lower sidebands separately. For Bands 9 and 10 (DSB) it will effectively double the bandwidth from 8 GHz to 16 GHz per polarization from the DSB receivers. For further details please see Section 5.5.4.

Appendix C

Acronym Dictionary

2SB	Dual Sideband
AIA	Atmospheric Imaging Assembly (SDO)
ACA	Atacama Compact Array
ACD	Amplitude Calibration Device
ACS	ALMA Common Software
AFA	ALMA Frontend Archive
ALMA	Atacama Large Millimeter/Submillimeter Array
AoD	Astronomer on Duty
AOS	Array Operation Site
APDM	ALMA Project Data Model
APE	ALMA Production Environment
APP	ALMA Phasing Project
APS	ALMA Phasing System
AQUA	ALMA Quality Assurance software
ARC	ALMA Regional Center
ASA	ALMA Science Archive
ASC	ALMA Sensitivity Calculator
ASDM	ALMA Science Data Model
ASIC	Application Specific Integrated Circuit
ATC	Advanced Technology Center, NAOJ
ATM	Atmospheric Transmission at Microwaves
AZ	Azimuth
BB	Baseband
BDF	Binary Data Format
BE	Backend
BL	Baseline
BLC	BaseLine Correlator
BWFN	Beam Width between First Nulls
BWSW	Bandwidth Switching
CASA	Common Astronomy Software Applications package
CCA	Cold Cartridge Assemblies
CCC	Correlator Control Computer
CDP	Correlator Data Processor
CFRP	Carbon Fiber Reinforced Plastic
CIP	Correlator and Integration Processor
CLO	Central Local Oscillator
CLT	Chilean Local Time

CORBA	Common Object Request Broker Architecture
CRD	CentralReference Distributor
CRG	Central Reference Generator
CSV	Commissioning and Science Verification
CW	Continuous Wave
DC	Direct Current
DEC	Declination
DFP	DTS-Rx and FFT Processor
DGC	Differential Gain Calibrator
DGCK	Digital Clock
DMG	Data Management Group within DSO
DRX	Data Receiver module
DSB	Double Sideband
DSO	Division of Science Operations
DTS	Data Transmission System
DTS-Rx	Data Transmission System Receiver
DTX	Data Transmitter module
EB	Execution Block
EHT	Event Horizon Telescope
EL	Elevation
EM	Electromagnetic
EPO	Education and Public Outreach
ES	Early Science
ESO	European Southern Observatory
FDM	Frequency Division Mode
FE	Frontend
FFT	Fast Fourier Transform
FITS	Flexible Image Transport System
FLOOG	First LO Offset Generator
FOM	Fiber Optic Multiplexer
FOV	Field of View
FP	Focal Plane
FPGA	Field-Programmable Gate Array
FPS	Frequency Profile Synthesis
FT	Fourier Transform
FWHM	Full Width Half Maximum
FWHP	Full Width to Half Power
FWBN	Full Width Between the Nulls
FX	Fourier transform correlation type correlator
FXF	Filtering, Correlation, and Fourier transform type correlator
GMVA	Global Millimeter VLBI Array
GOUS	Group Observing Unit Set
GPS	Global Positioning System
GUI	Graphical User Interface
HA	Hour Angle
HEMT	High Electron Mobility Transistor
HPBW	Half Power Beam Width
IF	Intermediate Frequency
IFP	Intermediate Frequency Processor
IRAM	Institut de Radioastronomie Millimetrique
JAO	Joint ALMA Observatory
JCMT	James Clerke Maxwell Telescope
JPL	Jet Propulsion Laboratory

LAS	Largest Angular Resolution
LFRD	Low Frequency Reference Distributor
LLC	Line Length Corrector
LO	Local Oscillator
LO1	First LO
LO2	Second LO
LO3	Digitizer Clock LO, Third LO
LO4	Tunable Filterbanks LO, Fourth LO
LORR	LO Reference Receiver
LS	Laser Synthesizer
LSB	Lower Sideband
LSRK	Local Kinematics Standard of Rest
LTA	Long Term Accumulator
MCI	Monitor and Control Interface
MD	Mixed De-Tuned or Mixed De-Biased mode
MFS	Master Frequency Standard
MIT	Massachusetts Institute of Technology
ML	Master Laser
MLD	Master Laser Distributor
MOUS	Member Observing Unit Set
MRS	Maximum Recoverable Scale
MS	Measurement Set
NAOJ	National Astronomy Observatory Japan
NGAS	New Generation Archive System
NRAO	National Radio Astronomy Observatory
NVSS	NRAO VLA Sky Survey
OMC	Operator Monitoring and Control
OMT	Ortho-mode Transducer
OSF	Operations Support Facility
OST	Observation Support Tool
OT	Observing Tool
OTF	On the Fly
OUS	Observing Unit Set
PBS	Polarization Beam Splitter
PDM	Propagation Delay Measure
PI	Principal Investigator
PLL	Phase Lock Loop
PMG	Program Management Group within DSO
PRD	Photonic Reference Distribution
PSF	Point Spread Function
PWV	Precipitable Water Vapor
QA	Quality Assurance
QA0	Quality Assurance Level 0
QA1	Quality Assurance Level 1
QA2	Quality Assurance Level 2
QA3	Quality Assurance Level 3
QL	QuickLook pipeline
RA	Right Ascension
RF	Radio Frequency
RMS	Root Mean Square
SAS	Sub Array Switch
SB	Scheduling Block
SCO	Santiago Central Office
SD	Single Dish

SDO	Solar Dynamics Observatory
SED	Spectral Energy Distribution
SEFD	Source Equivalent Flux Density
SG	Science Goal
SIS	Superconductor-Insulator-Superconductor mixer
SL	Slave Laser
SNR	Signal-to-Noise Ratio
spw	Spectral Window
spws	Spectral Windows
SQLD	Square Law Detector
SRON	Stichting Ruimte Onderzoek Nederland (Netherlands Institute for Space Research)
SSB	Single Sideband
SSR	Science Software Requirements
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable Filter Banks
TFB LO	Local Oscillator at the Tunable Filterbanks; LO4, Fourth LO
TMCDB	Telescope Monitor and Configuration DataBase
TP	Total Power
TPP	Total Power Processor
Tsys	System Temperature
T_{rx}	Receiver Temperature
USB	Upper Sideband
VLA	Very Large Array
VLBI	Very Long Baseline Interferometry
VOM	VLBI Observing Mode
VO	Virtual Observatory
WCA	Warm Cartridge Assembly
WVR	Water Vapor Radiometer
XF	Correlation-Fourier Transform Type Correlator
XML	eXtensible Markup Language
YIG	Yttrium-Iron Garnet Oscillator

Index

- (u, v) plane, 22
- uv taper, 94
- uv -coverage, 89
- 12-m Array, 9, 11
- 64-input Correlator, 9, 12, 57, 58
- 7-m Array, 9, 11, 13

- ACA Correlator, 9, 60
- Airy function, 18
- Alcatel Alenia Space European Industrial Engineering
MT Aerospace (AEM), 187
- ALMA
 - Calibrator Source Catalogue, 181, 182
 - Common Software (ACS), 177
 - Frontend Archive (AFA), 181
 - location, 11
 - Science Archive (ASA), 181
 - Science Data Model (ASDM), 179, 180
- ALMA Common Software (ACS), 11
- ALMA Project Data Model (APDM), 107
- ALMA Regional Center (ARC), 181
- ALMA subsystem, 177
 - Control, 177
 - Executive, 177
 - Telescope Calibration, 177, 178
- Amplitude, 15, 21, 141
- Amplitude Calibration Device, 31
- Amplitude Calibration Device (ACD), 192
- Angular resolution, 84, 90
- Antenna
 - delay offset, 146
 - integration, 145
 - pointing, 145, 149
 - position, 146, 150
 - positions, 146
 - power response, 18
- Antenna(s)
 - beam size, 16
 - CM, 12
 - design and properties, 187
 - foundations, 11, 189
 - pads, 189
 - panels, 188
 - performance specifications, 12, 187
 - PM, 12
 - power response, 18

 - response, 16
 - specifications, 187
 - surface, 187
 - transportation, 190
 - transporters, 11
 - types, 188
- AOS Technical Building, 201
- Aperture, 16
 - synthesis, 19
- Application Specific Integrated Circuit, 62
- Array
 - performance, 173
- Array Operations Site (AOS), 11
- Astrometry, 160
- Astronomer on Duty, 107, 166
- Atacama Compact Array, 9, 11
- Atacama Large Millimeter/Submillimeter Array, 9
- Atmosphere
 - opacity, 133
- Atmospheric calibration (ATM), 195
- Atmospheric transmission, 29
- Auto-correlation, 19

- band-to-band, 162
- Bandpass
 - calibration, 150
 - calibrator, 150
- Barlett weighting, 65
- Baseband(s), 12, 13, 57, 71
 - desired center, 73
- Baseline, 20
- Baseline calibration, 142
- Baseline Correlator, 57
- Beam shape, 90
- Bessel function, 18
- Binary Data File (BDF), 180
- Binary Data Format, 66
- Brightness, 15
 - temperature, 28
- Brightness temperature, 28

- Calibration
 - atmospheric, 112
 - bandpass, 111, 150
 - flux, 111
 - phase, 111

- polarization, 157
 - system-defined, 110
 - user-defined, 110
- Calibration observation, 109
- Calibration-astrometry, 160
- Cartridge assembly, 31
- Cassegrain focus, 188
- Chajnantor, 11
- Chajnantor Plateau, 10
- Cittert-Zernike theorem, 22
- Clock rate, 60
- Complex visibility, 21
 - distribution, 22
- cone-searches, 154
- Configurations, 84
- Contact scientist, 173
- Correlator
 - 64-input, 12
 - ACA, 12, 57
 - capabilities, 64
 - FX, 13, 60
 - XF, 12
- Correlator and Integration Processor, 62
- Correlator Data Processor, 60
- Correlator mode
 - FDM, 201
- Cosmic Microwave Background, 134
- Cross-correlation, 19
- Cryostat, 191
- Cycle 0, 9
- Data
 - delegation, 183
 - delivery, 183
 - download, 185
 - replication, 182
- Data and Control Flow, 177
- Data Capture module, 179
- Data Management Group, 165
- Data rate, 12
 - average, 181
 - peak, 181
- Data receivers, 62
- Deconvolution techniques, 25
- Delay
 - artificial, 20
 - geometrical, 20
- Digital Transmission System Receiver, 61
- Digital Transmitter, 62
- Digitizers, 62
- Director's Discretionary Time, 183
- Dirty
 - beam, 24
 - image, 24
- Dirty beam, 83
- Doppler shift, 199
- Doppler tracking, 80
- DTS and FFT processor, 61
- Dump rates, 12
- Early Science
 - Operations, 9
- Earth rotation, 22
- EB calibrations, 147
- Electromagnetic power, 15
- Ephemeris file, 80, 118
- Execution Block (EB), 109
- Feathering, 102
- Field of view, 23, 97
- Field programmable gate array, 59
- Flux density, 15, 28
- Flux-scale, 150
- Fourier transform, 15
- Fraunhofer approximation, 18
- Frequency
 - center offset, 73
 - definitions, 73
 - generation, 201
 - rest, 73
 - sky, 73
 - synthesizer, 72
- Frequency and phase standard, 201
- Frequency Division Mode (FDM), 58
- Frequency Profile Synthesis, 60
- Fringe, 22
 - rotation, 199
 - spacing, 22
 - tracking, 200
- Frontend, 29, 191
- Full Width Between Nulls, 17
- Full Width Between the Nulls, 17
- FXF, 58
- Gibbs phenomenon, 94
- Grid source, 112
- Group Observing Unit Set (GOUS), 109
- Half Power Beam Width, 16
- Hanning function, 64
- Helpdesk, 173
- Hot load, 148
- Image fidelity, 24
- Imaging
 - simultaneous, 73
- Integration time, 100
- Interferometer
 - resolution, 25

- two-antenna, 20
- Interferometry
 - single field, 111
- Intermediate frequency, 31, 71
- Jansky, 16
- Joint ALMA Observatory (JAO), 11
- JPL Horizon format, 118
- Koh-ichiro Morita, 14
- Lasers, 201
- LO and IF systems, 199
 - operations, 200
- Local Oscillator, 31, 57, 71
 - offsetting, 77
- Local Oscillator (LO), 12
- Long Term Accumulator, 60
- Lower sideband, 71
- Maximum recoverable scale, 13, 25, 96, 98
- Measurement set, 147
- Member Observing Unit Set (MOUS), 109
- Metadata, 180
- Metrology, 189
- Mitsubishi Electric Corporation (MELCO), 187
- Morita Array, 9
- Mosaicing, 23, 97, 112
- Multiple Region Modes, 118
- NED, 184
- Next Generation Archive System, 181
- non-dispersive, 142
- Observations
 - on-the-Fly Mapping, 112
 - total power, 112
- Observing
 - project, 107
 - proposal, 107
 - script, 107, 110
 - session, 109
- Observing mode(s), 9
- Observing session, 177
- Observing Unit Set, 174
- Observing Unit Set (OUS), 109
- On-the-fly scan (OTF), 13
- Operation Support Facility (OSF), 11
- Ortho mode transducer, 33
- Output radiance, 133
- Pampa la Bola, 11
- pcal-sens, 154
- Phase, 15, 21, 141
 - zero, 143
- Phase referencing, 153
- Phase 1, 71, 107
- Phase 2, 71, 107
- Phase calibration, 76
- Phase calibrator
 - image, 154
- Phase center, 21
- Phase slope, 146
- Phase switching, 79
- phase-frequency relationship, 142
- Pointing referencing, 150
- Polarization, 16, 78, 113
 - instrumental, 114
 - off-axis, 114
- Position switching, 13
- Precipitable water vapor, 63
 - octiles, 133
- Precipitable water vapour
 - octiles, 134
- Primary beam, 97
- Primary beam, 16
 - correction, 23
- Primary focus, 188
- Program Management Group, 165
- Proposer's Guide, 9, 109
- Proprietary period, 183
- PWV, 148
- Quality Assurance, 165
 - categories, 166
 - classification, 165
 - report, 174
- Quantization, 57
 - efficiency, 69
- Radiometer
 - Dicky-switching, 196
- Rayleigh resolution, 17
- Rayleigh-Jeans limit, 28
- Receiver(s)
 - 2 sideband, 136
 - double sideband, 136
 - heterodyne, 135
- Rectangular field, 112
- Reflector
 - secondary, 18
- Robustness, 90
- San Pedro de Atacama, 11
- Scan, 109, 178
- Scheduling Block, 166
- Scheduling Block (SB), 107
- Science
 - goal, 107, 112
 - parameter(s), 107
- Science Pipeline, 171

- Science Software Requirements, 109
- Self-calibration, 157
- Sensitivity, 74, 133
 - point-source, 137
 - surface brightness, 138
- Shadowing, 26, 86
- Shutter, 188
- Sideband
 - double, 30, 77
 - double (DSB), 12
 - dual, 12, 30
- Sidelobes, 17, 25, 83, 94
- Simbad, 184
- Sky brightness distribution, 15, 22
- Sky frequency, 71
- Snapshot, 94
- Solar observations, 120
- Solar system objects, 118
- Spatial
 - filtering, 23, 95
 - frequency, 15, 21
- Specific intensity, 15, 28
- Spectral setup, 71, 73
- Spectral window (spw), 58
- Spectral window (spws), 71
- Stokes parameters, 113
- Sub-arrays, 12
- Submillimeter Array (SMA), 16
- Subreflector, 145, 188
- Subscan, 178
- Subsystems
 - QuickLook, 110
 - Control, 110
 - Scheduler, 109
 - Telescope Calibration, 110
- Synthesized beam, 24

- Technical Handbook, 9
- TelCal, 146
- Telescope
 - filled-aperture, 20
- Telescope operator, 177
- Temperature
 - sky, 133
 - system, 133
- Tiltmeters, 189
- Time Division Mode (TDM), 58
- Tools
 - AosCheck, 166
 - AQUA, 166, 174
 - Monitor and Control, 166
 - QuickLook, 166
 - TelCal, 166
- Total Power Array (TP Array), 11
- Total Power Array (TP Array), 9, 13
- Transporter, 190
- Tunable Filter Bank (TFB), 12, 58, 59, 199
- Tuning range, 77

- Upper sideband, 71
- User Policies, 173

- Van Vleck relationship, 69
- Velocity definition
 - Optical, 80
 - Radio, 80
 - Relativistic, 80
- Velocity reference, 80
 - Barycentric, 80
 - LSRK, 80
 - Topocentric, 80
- VertexRS, 187
- Virtual Observatory, 184
- Visibility, 15
- Visibility function
 - weighted, 83
- Vizier, 184
- Voltage response, 18
- VOTables, 184

- Walsh-switching, 57, 79
- Water Vapour Radiometer (WVR), 12, 63, 195
- Weblog page, 172
- Weighting
 - natural, 90
 - uniform, 90
- WVR Correction, 63

- Zero phase, 143
- Zero spacing, 112
- Zero-point fluctuations, 135
- Zero-spacing problem, 13, 25



The Atacama Large Millimeter/submillimeter Array (ALMA), an international astronomy facility, is a partnership of the European Organisation for Astronomical Research in the Southern Hemisphere (ESO), the U.S. National Science Foundation (NSF) and the National Institutes of Natural Sciences (NINS) of Japan in cooperation with the Republic of Chile. ALMA is funded by ESO on behalf of its Member States, by NSF in cooperation with the National Research Council of Canada (NRC) and the Ministry of Science and Technology (MOST) in Taiwan and by NINS in cooperation with the Academia Sinica (AS) in Taiwan and the Korea Astronomy and Space Science Institute (KASI).

ALMA construction and operations are led by ESO on behalf of its Member States; by the National Radio Astronomy Observatory (NRAO), managed by Associated Universities, Inc. (AUI), on behalf of North America; and by the National Astronomical Observatory of Japan (NAOJ) on behalf of East Asia. The Joint ALMA Observatory (JAO) provides the unified leadership and management of the construction, commissioning and operation of ALMA.

