

ALMA Cycle 3 Technical Handbook



www.almascience.org

ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada), NSC and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ.

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:

Version	Date	Editors
1.0	2015-03-20	A. Remijan, M. Adams, R. Warmels

Contributors

Mark Adams, Eiji Akiyama, Paola Andreani, Yiping Ao, Shin'ichiro Asayama, Denis Barkats, Andy Biggs, James Braatz, Crystal Brogan, James Chibueze, Stuartt Corder, Juan Cortes, Paulo Cortes, Carlos De Breuck, Itziar De Gregorio, Bill Dent, James Di Francesco, Jennifer Donovan Meyer, Darrel Emerson, Daniel Espada, Ed Fomalont, Roberto Galvan-Madrid, Antonio Hales, Bunyo Hatsukade, Cinthya Herrera, John Hibbard, Akihiko Hirota, Elizabeth Humphreys, Todd Hunter, Remy Indebetouw, Akiko Kawamura, Takeshi Kamazaki, Eelco van Kampen, Rüdiger Kneissl, Shinya Komugi, Yasutaka Kurono, Mark Lacy, Stephane Leon, Harvey Liszt, Carol Lonsdale, Andreas Lundgren, Brian Mason, Geoff Mathews, Jeff Mangum, Nuria Marcelino, Ivan Marti-Vidal, Rie Miura, Koh-Ichiro Morita, Erik Müller, Hiroshi Nagai, Lars-Åke Nyman, Sachiko Okumura, Rosita Paladino, Alison Peck, Dirk Petry, Neil Phillips, David Rabanus, Matias Radiszcz, Suzanna Randall, Mark Rawlings, Anita Richards, Masao Saito, Markus Schmalzl, Scott Schnee, Kimberly Scott, Kartik Sheth, Richard Simon, Hiroko Shinnaga, Giorgio Siringo, Thomas Stanke, Felix Stoehr, Kenichi Tatematsu, Ignacio Toledo, Baltasar Vila-Vilaro, Eric Villard, Rein Warmels, Al Wootten, Jorge F. García Yus, Martin Zwaan.

In publications, please refer to this document as:

A. Remijan et al., 2015, ALMA Cycle 3 Technical Handbook Version 1.0, ALMA
ISBN 978-3-923524-66-2

Contents

1	Introduction	5
2	Array Components	7
2.1	The ALMA Telescope	7
2.2	The 12-m Array	8
2.3	The Atacama Compact Array (The Morita Array)	9
2.3.1	The 7-m Array	9
2.3.2	The TP Array	10
3	Principles and Concepts of Interferometry	11
3.1	Introduction	11
3.2	Single-dish Response	12
3.3	Visibilities and Aperture Synthesis	15
3.4	The Visibility or <i>uv</i> -Plane	18
3.5	Fields-of-view and Mosaics	19
3.6	Spatial Filtering	19
3.7	Multi-configuration Observations	22
3.8	Units and Conversions	24
4	Receivers	27
4.1	Local Oscillators and IF Ranges	29
4.2	The Cycle 3 Receivers	29
4.2.1	Band 3 receiver	30
4.2.2	Band 4 receiver	33
4.2.3	Band 6 receiver	36
4.2.4	Band 7 receiver	39
4.2.5	Band 8 receiver	42
4.2.6	Band 9 receiver	45
4.2.7	Band 10 receiver	48
5	The Correlators	51
5.1	The 64-input Correlator	52
5.1.1	TDM mode	52
5.1.2	FDM mode	53
5.1.3	Correlation and realtime processing	54
5.2	The ACA Correlator	54
5.2.1	7-m Array	56
5.2.2	Total Power Array	58
5.3	Digitizers	58
5.4	Online WVR Correction	58
5.5	Capabilities of the Correlators	58
5.5.1	Polarization	58
5.5.2	Spectral Setup	59

5.5.3	Time Resolution	61
5.6	Practical Performance	62
5.6.1	Sensitivity	62
5.6.2	Linearity	63
5.6.3	Integration time intervals, channel-average and spectral data	63
5.6.4	Online WVR correction	64
5.6.5	Correlator speed and data rates	64
5.6.6	Final data product - the ASDM	65
6	Spectral Setups	67
6.1	Introduction	67
6.2	Frequency Definitions	68
6.3	Spectral Setup: Multiple lines and Continuum	69
6.3.1	Observing Frequencies for Continuum	69
6.3.2	Multiple Spectral Windows in the Same Baseband	71
6.4	Spectral Setups for Lines near the Edge of the Bands	72
6.5	Spectral Setup in Band 9 & 10 and DSB Considerations	73
6.5.1	Usable Bandwidth	73
6.5.2	Spurious Signals	74
6.5.3	Doppler Setting and Velocity Reference Frames	74
6.6	Limitations and Rules for Spectral Setups in Cycle 3	74
7	Imaging with ALMA	77
7.1	Introduction	77
7.2	Cycle 3 Configurations	77
7.3	Shadowing	78
7.4	Resolution and Beam Shape	82
7.5	Response and Snapshot	83
7.6	Spatial Scale Filtering	87
7.7	Mosaicing	89
7.8	Multi-array and Multi-configuration Observations	90
7.9	Multi-array and Multi-configuration Imaging	93
8	Observing Modes	97
8.1	Observing Project Structure	97
8.2	Program Execution	99
8.3	The Observing Process	100
8.3.1	The source selection algorithm	101
8.4	Single Field Interferometry	102
8.5	Pointed Mosaic Observations	103
8.6	Single Dish Observations	103
8.7	Polarization	104
8.7.1	Sessions	106
8.8	Multiple Region Modes	107
8.9	Observations of Ephemeris Objects	107
9	ALMA Sensitivity Calculator	109
9.1	Calculating the System Temperature	109
9.1.1	Sky temperature	109
9.1.2	CMB temperature	110
9.1.3	Receiver temperatures	110
9.1.4	Ambient temperature	111
9.1.5	DSB receivers	112
9.1.6	System temperature (OT version)	112
9.1.7	System Temperature (Applet and GUI version)	113

9.2	The Sensitivity Calculation	114
9.2.1	12-m and 7-m Arrays	114
9.2.2	Total Power Array	115
9.3	User Interface	115
9.4	Total Time Estimates	116
10	Calibration and Calibration Strategies	117
10.1	The Measured Visibility Function	117
10.1.1	The Data Format and the Correlator Model	117
10.1.2	The Visibility Function Dependence	117
10.2	Main Interferometric Calibration Strategies	118
10.2.1	Antenna-based Calibrations	118
10.2.2	Calibration Flow through ALMA system	118
10.2.3	A priori, Online and Offline Calibration Methods	119
10.3	ALMA-Supplied a priori Calibrations	119
10.4	Online Calibrations	121
10.5	Offline Calibrations	123
10.6	Appendix	125
10.6.1	ALMA Catalog Description	125
10.6.2	Phase Calibrator Sensitivity	125
10.6.3	Bandpass Calibrator Sensitivity	126
10.6.4	Band-to-Band Phase Calibration	126
10.6.5	Bandwidth Phase Transfer	126
10.6.6	Astrometry	127
10.6.7	Long Baseline Observations with Baselines > 3 km	127
11	Quality Assurance	129
11.1	Cycle 3 "Best Efforts"	130
11.2	QA0	130
11.3	QA1	131
11.4	QA2	131
11.5	QA3	132
11.6	The Quality Assurance Report	132
11.7	QA Criteria	133
12	Logical Data Structure and Data Flow	135
12.1	Data and Control Flow	135
12.2	The ALMA Science Data Model	137
13	Pipeline Processing	139
13.1	Introduction	139
13.2	Pipeline In ALMA Operations	139
13.2.1	Pipeline Triggering	139
13.2.2	Pipeline Execution	139
13.2.3	Pipeline Runtime Database-querying	140
13.2.4	Pipeline and the ALMA Quality Assurance Process	140
13.2.5	Standard and non-standard Observing Modes	144
13.3	Pipeline Processing of an ALMA OUS	144
13.3.1	Overview	144
13.3.2	Pause after Data Import	144
13.3.3	Output of the Pipeline Processing	144
13.4	Pipeline Initialisation & Calibration Table Editing	144
13.5	General Interferometric ALMA Data Processing	145
13.5.1	Data Import & Calibrator Flux Retrieval	145
13.5.2	Deterministic & Manual Flagging	145

13.5.3	Flagging lines in Solar System Object Flux Calibrators	145
13.5.4	Reference Antenna Selection	146
13.5.5	Offline WVR Calibration for 12m Array Data	146
13.5.6	Low gain flagging	146
13.5.7	Bandpass calibration & flagging bandpass solutions	146
13.5.8	Setting the fluxscale	146
13.5.9	Deriving the fluxes of the phase and bandpass calibrators	146
13.5.10	Time Gain Calibration	147
13.5.11	Applying the calibration	147
13.5.12	Imaging	147
13.5.13	Exporting the products	147
13.6	Low SNR Heuristics in Interferometric Pipeline Processing	147
13.6.1	Mapping the gain solutions from broad spectral windows to narrow ones	148
13.6.2	Using the gain solutions obtained from the aggregate of all spectral windows	148
13.6.3	Additional Bandpass Pre-averaging	148
13.7	Total Power Data Processing	148
13.7.1	Data Import	148
13.7.2	Deterministic & Manual Flagging	149
13.7.3	Tsys & Sky Calibration Tables	149
13.7.4	Applying the calibration	149
13.7.5	Automated line identification and Baseline subtraction	149
13.7.6	RMS Flagging	149
13.7.7	Applying Kelvin-to-Jansky factors	149
13.7.8	Imaging	149
13.7.9	Exporting the products	150
14	Data Archiving	151
14.1	Introduction	151
14.2	Data Flow and Archive	151
14.3	PI Data Delivery and Data Delegation	152
14.4	Archive Query	153
14.5	Request Handler	153
A	Antennas	155
A.1	Design and Properties	155
A.2	Antenna Foundations	157
A.3	Antenna Transportation	159
A.4	Cryostat	159
A.5	Amplitude Calibration Device	162
A.5.1	Atmospheric Calibration Procedure	163
A.6	Water Vapor Radiometers	164
B	The LO and IF System	167
B.1	Functions of the LO & IF system	167
B.2	Summary of Operation	168
B.3	Frequency Generation and Distribution in ALMA	169
B.3.1	Reference and LO signal generation	171
B.3.2	LO1 signal generation and distribution	171
B.3.3	Line length corrections	172
B.3.4	The First Local Oscillator (LO1)	173
B.3.5	LO2 and the IF processor units and IF switch	174
B.3.6	Digitization and signal transmission	175
B.4	Other functions of the LO/IF	175
B.4.1	Delay corrections	176
B.4.2	Sideband suppression - LO offsetting	176

B.4.3	Sideband separation - 90 degree Walsh switching	177
B.4.4	Interference rejection - 180 degree phase switching	177
C	Calibration Source Selection	179
C.1	Antenna-based Sensitivity Determinations	179
C.1.1	ALMA Sensitivities	179
C.1.2	Calibration Parameters and Derived Antenna-based RMS	179
C.2	Choosing the Calibrators	180
C.2.1	The Phase Calibrator	180
C.2.2	The Bandpass Calibrator	181
C.2.3	Other Calibrator Types	183
D	Acronym Dictionary	185

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 arranged in a series of different configurations. It operates over a broad range of observing frequencies in the millimeter and submillimeter regime. ALMA Early Science operations started with Cycle 0 in September 2011 and the official inauguration took place in March 2013. Cycle 3 is still within the ALMA Early Science period, in the sense that a reduced number of antennas, frequency bands, array configurations and observing modes will be available. Cycle 3 operations will include standard and non-standard modes, with non-standard mode observations being conducted on a best effort basis, similar to projects in previous Cycles. Users should refer to Appendix A (Capabilities) in the ALMA Cycle 3 Proposer's Guide, for the latest information and a description of standard and non-standard modes.

The Technical Handbook provides additional information on technical aspects and its limitation of the Cycle 3 setup for ALMA users, to a deeper level than what is described in the ALMA Early Science Primer. It should, however, not be necessary to use the Technical Handbook to prepare an ALMA proposal.

The handbook is divided in 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 10–14). It also includes a number of appendices with expanded information on specific ALMA hardware components and calibration, and an acronym list.

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 the 7-m and Total Power (TP) Arrays. A general description on the different array elements and their 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 seven receivers offered in Cycle 3. The general technical specifications and a brief explanation on local oscillators and IF range is presented. Plots with the atmospheric transmission and the typical system temperatures per receiver 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. Modes for continuum and spectral line observations are presented.

Chapter 6 describes how the spectral setup is done in the correlators. It covers a description of the signal path and local oscillator chain used between the frontends and the correlators, and how these are used to define spectral setups for the user.

Chapter 7 describes several aspects of imaging to consider in ALMA observations. A short description on the different configurations proposed for Cycle 3 is included. 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 3 and how projects are observed. Single field in-

terferometry, mosaics, single-dish observations, polarization, multiple region modes, 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 process. It also provides the criteria used for passing the quality assurance.

Chapter 12 describes the structure of the data.

Chapter 13 includes a description on the ALMA pipeline, its infrastructure, the pipeline heuristics and the use of pipeline in data calibration.

Chapter 14 describes how the data are stored, the data flow and the ALMA archive.

The Technical Handbook concludes with a number of Appendices, which contain supplemental material about concepts and hardware, such as the antenna, transporter (Appendix A), LO system (Appendix B), and the calibrator source selection algorithm (Appendix C).



Figure 1.1: ALMA antennas on the Chajnantor Plateau. Credit: ALMA (ESO/NAOJ/NRAO), O. Dessibourg

Chapter 2

Array Components

In this chapter we describe the main characteristics of each ALMA array. Unless otherwise noted, the description is appropriate for the fully completed ALMA. Note that not all capabilities may be available during the Early Science observing period.

2.1 The ALMA Telescope

Upon completion, ALMA will be composed of 66 high-precision antennas. Fifty of these antennas will be 12-meter dishes in the 12-m Array, used for sensitive, high-resolution imaging. These will be complemented by the Atacama Compact Array (ACA or 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 3, ALMA will cover most of the wavelength range from 3.6 down to 0.32 mm (from 84 to 950 GHz in frequency), and when ALMA is completed the coverage will be from 10 to 0.32 mm (31-950 GHz).

The array is located on the Chajnantor plain of the Chilean Andes (latitude = -23.022917° , longitude = -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 Joint ALMA Observatory (JAO) Central office in Santiago.

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 moved 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 3) will be published in Appendix A (ALMA Cycle 3 Capabilities) of the document *ALMA Proposer's Guide* (Appendix A). Complementary background information on ALMA and its capabilities for Early Science can be found in the document *A Primer for Early Science*. Both documents

¹Dedicated to the honour of Koh-ichiro Morita, Figure 2.2)

²<http://almascience.edu/about-almalweather>

can be found at: <http://almascience.org/documents-and-tools/>.



Figure 2.1: The ALMA 12-m Array in its compact configuration (left hand-side of the image). The ACA with all CM and 4 PM antennas are distributed in the right hand side of the image (circled). 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, antenna baselines in the 12-m Array will reach a maximum extent of about 16 km.

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 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 backend 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, and rapid switching between those bands is possible. Each receiver band (Chapter 4) detects two orthogonal linear polarizations and down converts the signals to an intermediate frequency with eight GHz of bandwidth per polarization. Bands 3-8 are dual sideband (2SB) and Bands 9 and 10 are double sideband (DSB).

The Local Oscillator (LO) signals (see Section B.2) 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 make simultaneous independent observations. Note that the sub-arrays feature of the 12-m Array is currently being developed and tested and is not yet a capability that is offered for Early Science observations. The 8 GHz total IF bandwidth from the selected receiver is divided into four 2 GHz-wide basebands which are digitized at four Gsamples/s, with three-bit resolution, and transmitted on optical fibers (Section B.3.6). 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 correlators. 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 is preceded by Tunable Filter Banks, which makes it a digital hybrid XF correlator or FXF. These can select

subbands 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 modes. The 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, although the mean data rate will be considerably less.

The 12-m Array configurations have been designed so that in the most extended configurations the spatial angular resolution will be as small as 5 milliarcseconds at 950 GHz. The specific 12-m Array configurations available for Cycle 3 as well as the configuration schedule will be published in the "Capabilities" Appendix of the *ALMA Proposer's Guide* and described in more detail in Chapter 7.

2.3 The Atacama Compact Array (The Morita Array)

Using an interferometer to obtain images of extended or large-scale structures leads to the well-known "zero spacing" problem (See Chapter 3). 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 recovered³. 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 or 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 0m up to 12-m spacings as auto-correlations. The 7-m Array samples baselines from 9m to 30m, 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 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 collection, and observing modes. When ALMA is complete, the standard observing modes for the TP Array will include spectral line and continuum observations with raster or Lissajous on-the-fly (OTF) scans, position switching and frequency 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). Similar to 12-m Array antennas, each antenna contains one front-end, including a cryostat, amplitude calibration devices, and one backend. Unlike the antennas of the 12-m Array, the 7-m antennas do not contain Water Vapor Radiometers (WVRs). 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 originate from and are identical to the ones sent to the 12-m Array (i.e., from the AOS building).

³Strictly 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 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 GB/s per baseband.

No baseline coverage from even the most compact configuration of the 12-m Array is obtained for spacings smaller than 15 m. The array configuration of the 7-m Array are designed to fill missing spacings in from about 9 m to ~ 30 m (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 (See Appendix A). The specifications of the TP antennas are almost identical to the 12-m Array elements. The TP antennas are located on stations surrounding the 7-m Array. The TP antennas will eventually be fitted with a nutating subreflector (also known as wobbler or chopper) in full operations. The ALMA nutator is a single axis device, used to alternatively point the primary beam of the antenna between two defined sky positions by tilting the subreflector along the azimuth direction, at constant elevation. A modulation of the observed signal at a well known frequency is introduced by rapidly switching between the celestial source and a nearby reference position. A synchronous demodulation technique is then used to extract the flux of the source. This method can greatly improve the quality of single dish observations by suppressing atmospheric and noise components at frequencies lower than the nutator frequency.

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 call materials describe the observing mode and capabilities offered for the TP Array for each cycle.



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.

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. In this Chapter, we describe the principles and concepts behind interferometry, so that ALMA users can plan and understand their observations better. If more information is desired, the topic of interferometry is covered in more detail in the following seminal texts:

- *Interferometry and Aperture Synthesis in Radio Astronomy - Second Edition*, by Thompson, A. R., Moran, J. M., & Swenson, G. W. (Wiley-VCH)
- *Tools of Radio Astronomy - Fifth Edition*, by Wilson, T. L., Rohlfs, K., & Hüttemeister, S. (Springer)
- *Synthesis Imaging in Radio Astronomy II*, PASP Conference Series, Vol. 180, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco - ASP)



Figure 3.1: The Plateau de Bure Interferometer (top) and the Submillimeter Array (SMA) (bottom) are the precursors to the ALMA telescope; both are still in full operation and pioneered the science of millimeter wave interferometry.

We first provide a very basic picture of the core concepts behind how interferometry works. Interferometry

involves the combination of signals received from the sky by two 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, called a visibility, consists of the brightness of the emission on the angular scale sampled, i.e., the amplitude of the sinusoid, and the relative position of that brightness on the sky, i.e., 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 angular scales have been sampled. Interferometry, however, works extraordinarily well for observing intrinsically compact targets.

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

3.2 Single-dish Response

As in all astronomy, we define *brightness*, or equivalently *specific intensity*, I_ν , 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.e.,

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

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, i.e.,

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

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, i.e.,

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

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. (The ALMA receivers have been designed to detect both modes, however.) As with all telescopes, antennas bring incident EM power to a focus after reflecting it off a primary surface. The antenna response, i.e., its sensitivity, is a summation of all EM power brought to the focus.

Antenna response is actually dependent on angle from the on-axis pointing direction of the antenna, due to self-diffraction. 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

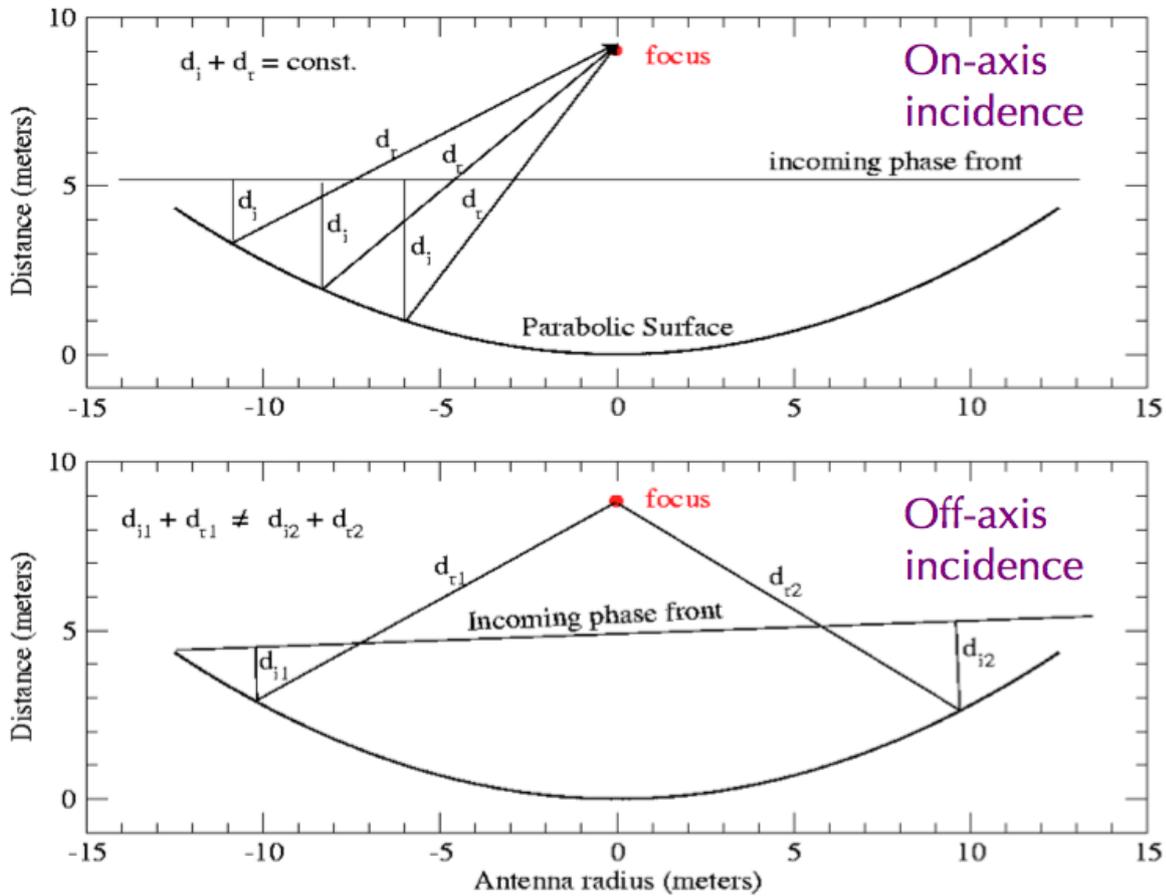


Figure 3.2: Schematic of an incoming plane-parallel wavefront reflecting off a antenna of diameter D and being brought to a focus. The top panel shows the case for a wavefront arriving on-axis. The bottom panel shows the case for a wavefront arriving off-axis. Note that the paths of incident EM power 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.

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 $1.22\lambda/D$ radians, 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 leads 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 ≈ 0.85 mm (350 GHz). The power response is largest on-axis but it declines to 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:

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

where λ is the wavelength of observation and D is the diameter of the antenna. For example, the HPBW of a uniformly illuminated antenna of 12 m diameter at $\lambda = 0.85$ mm is $14.95''$. HPBW is sometimes referred to as Full Width at Half Power or FWHP.

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 nulls is termed the Beam Width between First Nulls (BWFN), and is given by:

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

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.

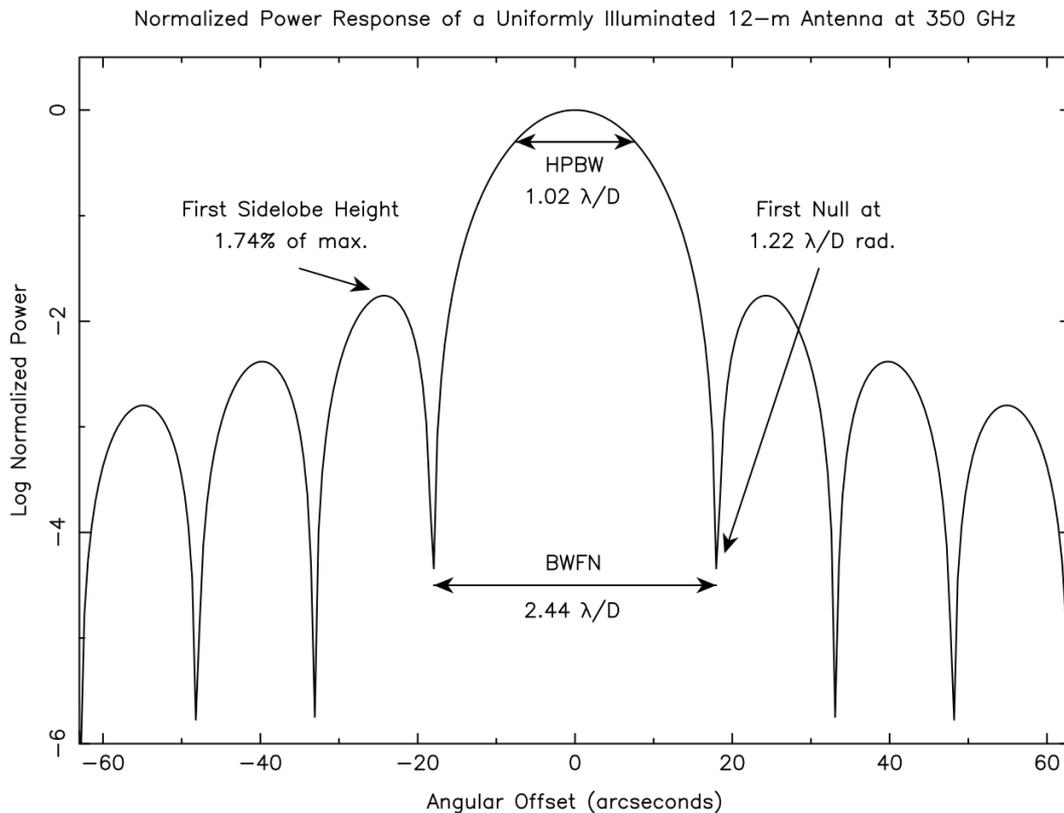


Figure 3.3: Normalized 1-D antenna power response for a 12 m antenna uniformly illuminated at 300 GHz. The power is in log units to emphasize the sidelobes. The HPBW of the primary beam is $\sim 1.02 \lambda/D$ and the angle of the first null, i.e., its resolution, is $\sim 1.22 \lambda/D$. Note that the HPBW measured from actual 12 m ALMA antennas is $\sim 1.13 \lambda/D$.

Up until now, we have described an idealized antenna. Actual antenna power response 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 (via an illumination pattern with a -12 dB taper) 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 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. Observing at millimeter/radio wavelengths is essentially diffraction-limited.

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

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

3.3 Visibilities and Aperture Synthesis

For millimeter/radio astronomy, the angular resolution of a single-dish observation is very low compared to those found at optical wavelengths since λ is larger by many orders of magnitude. Though the D of millimeter/radio 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 the angular resolutions of 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 $\sim 14''$, and the Arecibo 300-m diameter antenna has an angular resolution at 21 cm of $\sim 3'$.

To obtain higher angular resolution images than are possible with individual millimeter/radio telescopes, signals from physically separated antennas can be combined through interferometry. With this technique, sometimes called *aperture synthesis*, the resolution benefits of a large diameter aperture can be obtained. Observers, however, must contend with the reality that only certain angular scales, i.e., those determined by the projected spaces of each pair of antennas, will be sampled. In this section, we build on the concepts introduced previously to discuss aperture synthesis in more detail.

Earlier, we described how a plane-parallel wavefront arriving on-axis to an antenna is brought to a focus by a parabolic surface. Since there are no path differences in that case, 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 , i.e.,

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

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, we can rewrite the expression for received power in terms of the sum of time averages of the products of voltages from element pairs, i.e.,

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

Next, we can further rewrite this expression in terms of the sums of element pairs which are the same and those which are not, i.e.,

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

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, we see that 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, i.e., multiplying, receiver. Moreover, if the source properties are unchanging, there is no need to measure all pairs at the same time. A given large filled-aperture telescope 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 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*. We can measure this distance 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., with cables) so that the signals from both antennas arrive at the correlator with the same phase.

Moving slightly off-axis, we can describe a small angle from the axis as α , and its 1-D sky position as $l = \sin \alpha$. 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 $x = u \sin \alpha = ul$. Indeed, we can consider all distances in our 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, i.e.,

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

Expanding to two dimensions, we can introduce β , a direction on the sky orthogonal to α . Also, we define $m = \sin \beta$ as the small angle analog to l in this new direction, and $v = b \cos \varphi$ where φ is the angle of the position s_o on the sky from the reference position orthogonal to θ . Finally, we define y 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:

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

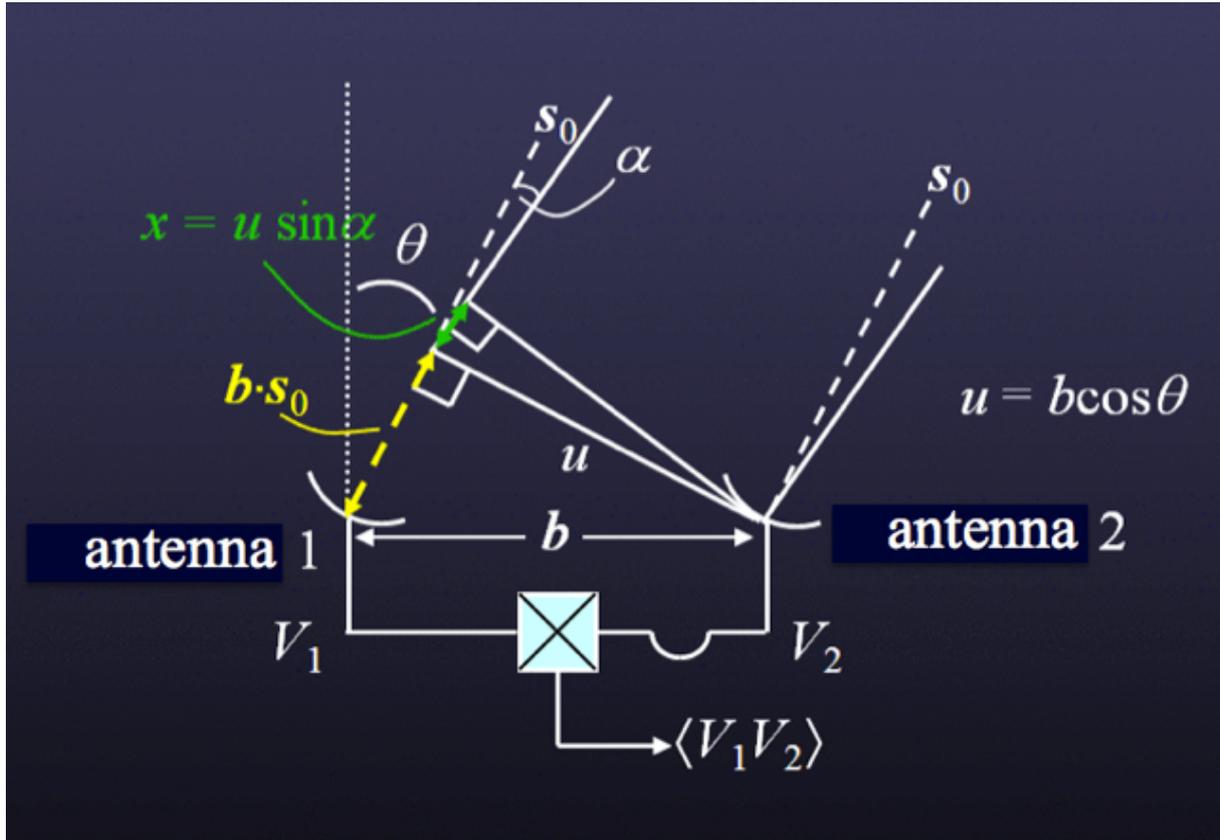


Figure 3.4: An ideal 1-D two-antenna interferometer consisting of two antennas, 1 and 2, separated by physical distance (i.e., a baseline) b . The antennas are both pointed towards a sky location given by s_o , which is at an angle θ from the meridian. The projected distance between the two antennas 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.

We identify u and v 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, we identify l and m 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, 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:

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

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

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

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

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 phase ϕ contains information about the location of brightness with spatial frequencies u and v relative to the phase center.

3.4 The Visibility or uv -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:

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

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

By measuring the distribution of complex visibilities (in the visibility or uv -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 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 complex 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 uv -plane is measured. Since visibilities are

actually Hermetian complex numbers, 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 uv -plane is needed. Improving coverage of visibilities over the uv -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 uv -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 uv -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 uv -plane. The combination of these schemes can reasonably sample the uv -plane, yielding an image that can resemble the true sky brightness distribution.

3.5 Fields-of-view and Mosaics

Note that we have described an interferometer with idealized antennas. Each antenna of an actual interferometer, however, has finite diameter. As described in the preceding section, such antennas have their own power response on the sky P_N (i.e., the Airy function = $\mathcal{A}(l, m)$). 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, $\mathcal{A}(l, m)$ is actually included formally in the correlator output:

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

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*, however, is only necessary if the resulting image contains extended emission or emission located far from the phase center.

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 uniform sensitivity if the observed positions are spaced by $\lambda/(\sqrt{3}D)$, where D is the diameter of the antenna. In this case, Nyquist sampling is achieved and 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. Mosaics can be made with adjacent positions that are spaced either closer or more distant than the Nyquist spacing, with non-uniform sensitivities. Mosaics provide increased areal coverage but come with a cost to observing time. For example, to obtain uniform sensitivity across the primary beam of just one pointing requires observations of six other pointings arranged in a hexagonal pattern on the sky and spaced at the Nyquist spacing around the one 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

Though the principles described in Sections 3.3-3.5 should enable the true sky brightness distribution to be recovered, it is impossible in practice to sample completely the uv -plane and obtain all visibilities. The incomplete uv -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 important for ALMA users to understand these limitations.

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 uv -plane has been sampled at $2M$ discrete points. We can then characterize the sampling distribution as an ensemble of $2M$ (Dirac) delta functions, i.e.,

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

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

$$I^D(l, m) = FT^{-1}\{B(u, v)\mathcal{V}(u, v)\}. \quad (3.24)$$

Using the convolution theorem, we can rewrite Equation 3.24 as:

$$I^D(l, m) = b(l, m) * I(l, m)\mathcal{A}(l, m). \quad (3.25)$$

In effect, the image obtained is the convolution of the true sky brightness distribution modified by the antenna power response with the point spread function, $b(l, m) = FT^{-1}\{B(u, v)\}$, the Fourier transform of the uv -plane sampling distribution. The point spread function is sometimes called the *synthesized beam* or the *dirty beam*. Again, the central feature of the antenna power response is called the primary beam. In addition, 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*.

The measure of how similar an image is to the true sky distribution is sometimes referred to as *image fidelity*. Formally, image fidelity is the ratio of the maximum brightness of an image to its 1σ rms noise level. Image fidelity depends on the specifics of coverage of the uv -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 uv -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 uv -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 uv -plane and the brightness of emission on sampled scales. If the uv -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 uv -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 we have been 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. 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 uv -plane, i.e., the uv -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 uv -plane seen in panel b. As an aside, note that two sets of uv -plane samples

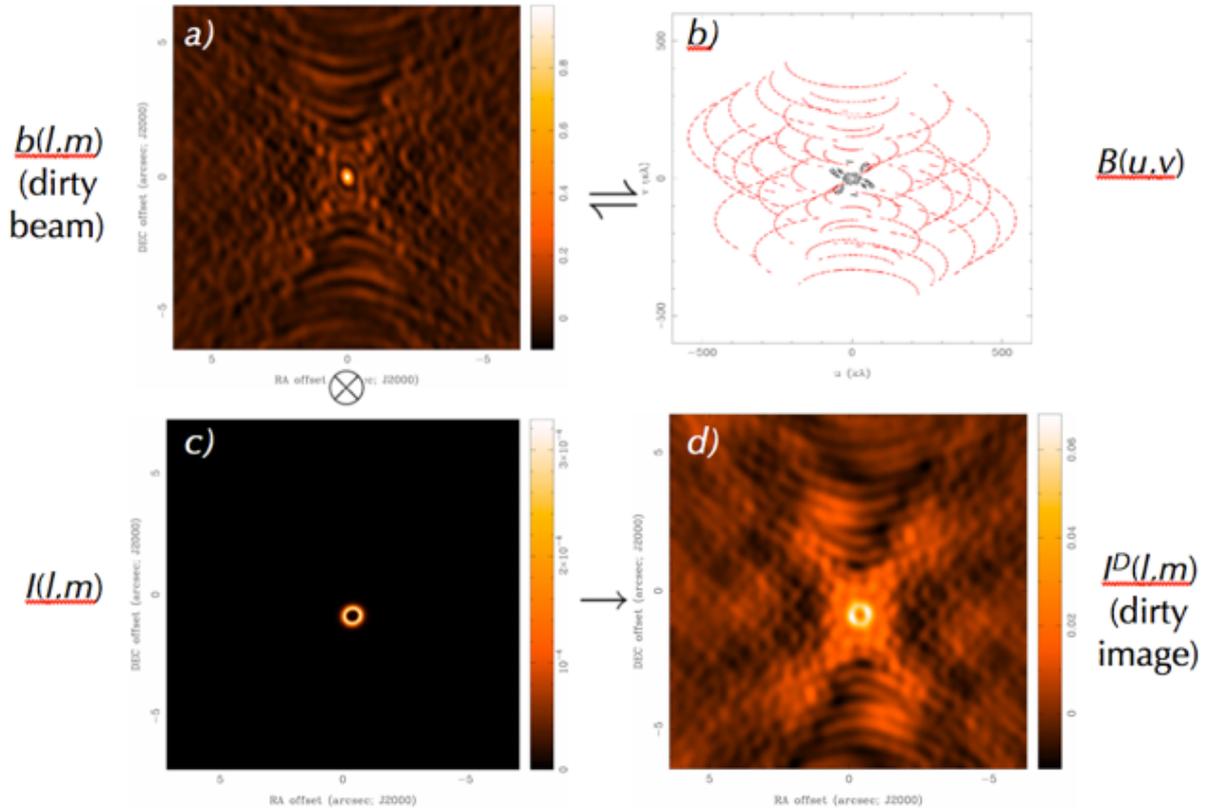


Figure 3.5: Imaging concepts. *Panel a (upper left)*: Example of a dirty beam, $b(l, m)$. *Panel b (upper right)*: The related ensemble of discrete points sampled in the uv -plane, $B(u, v)$. The black points were obtained from a compact configuration while the red ones were obtained from an extended 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.)

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 we can easily see how incomplete sampling of the uv -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 configurations. In general, distributions connected through 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 uv -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 uv -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:

$$\text{Interferometer Resolution} = \theta_{res} = k \lambda / L_{max}, \quad (3.26)$$

where k is a factor that depends on how the visibilities are weighted during inversion (typically ~ 1 ; see Figure 7.5) 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 that can be sampled by a baseline equal to an antenna diameter. In effect, visibilities at locations on the uv -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:

$$\text{Maximum Recoverable Scale} = \theta_{MRS} \approx 0.6 \lambda / L_{min}, \quad (3.27)$$

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 centre of an image made by an array with L_{min} to the same made with no central hole in visibility sampling is $1/e$; see Wilner & Welch 1994.) The smallest baseline possible in an array occurs when two antennas are adjacent to each other. Of 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 blocks the reception by the antenna in the rear, causing the latter to have reduced sensitivity. 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 actual ALMA configurations. Here, panel *a* shows an optical image of the galaxy M51 we use 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 32-antenna observations of the galaxy using the CASA task *simobserve* only in very extended, moderately extended, and compact configurations, respectively, and CLEANing. Angular resolutions of $0.55''$, $1.1''$, and $3.6''$ are obtained, respectively, and the corresponding maximum recoverable scales are $6.6''$, $14.4''$, and $36.4''$, respectively. Larger scale emission from the galaxy has been filtered out in the very 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 moderately 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 uv -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. 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.

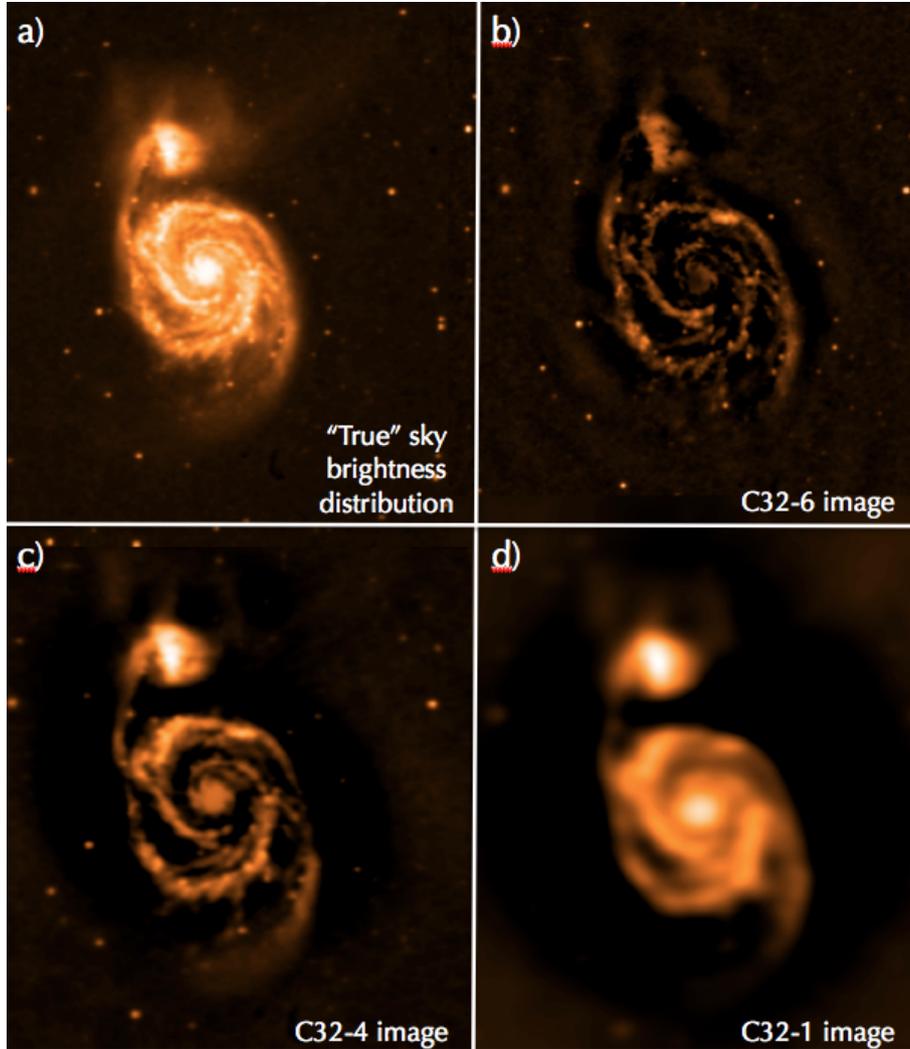


Figure 3.6: Examples of spatial filtering using the CASA task *simobserve* and actual ALMA configurations from Cycle 1. *Panel a (upper left)*: An optical image of the galaxy M51 used as a template for a true sky brightness distribution for the simulations. The frequency of the emission has been changed to 100 GHz, 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 39 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 an ALMA configuration with minimum and maximum baselines of 40.6 m and 1091.0 m, respectively (C32-6). The resulting synthesized beam is $\sim 0.55''$ and the maximum recoverable scale is $6.6''$. *Panel c (lower left)*: Medium-resolution image of the galaxy when observed in an ALMA configuration with minimum and maximum baselines of 20.6 m and 558.2 m, respectively (C32-4). The resulting synthesized beam is $\sim 1.1''$ and the maximum recoverable scale is $14.4''$. *Panel d (lower right)*: Low-resolution image of the galaxy when observed in an ALMA configuration with minimum and maximum baselines of 14.2 m and 81.4 m, respectively (C32-1). The resulting synthesized beam is $\sim 3.8''$ and the maximum recoverable scale is $36.4''$.

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. Otherwise, information on scales covered by lower SNR data is relatively less reliable, making interpretation of the images difficult. Assuming the SNR of the individual datasets is high, combination is best done in the visibility domain rather than the image domain to minimize the effect of artifacts produced by aliasing, i.e., incomplete uv -coverage, in either dataset. For example, interferometer data obtained from different configurations should be combined in the visibility domain and then the new ensemble should be Fourier transformed to produce a new dirty image.

Combining single-dish data and interferometer data also works best in the visibility domain, as long as both datasets have high SNRs. In this case, the single-dish image can be Fourier transformed into the visibility domain and the resulting visibilities added to the ensemble of those obtained by the interferometer. The new ensemble can be then Fourier transformed en masse to produce a new image. Such data combination works best if the single-dish and interferometric datasets have significant uv -coverages in common. For example, a reasonable overlap in uv -coverage can provide enough data to reveal amplitude calibration differences that can be minimized by re-scaling the single-dish visibilities relative to the interferometric ones. In general, a reasonable overlap of uv -coverage will occur if the single-dish data are obtained by an antenna that has a diameter twice the minimum baseline of the interferometer, e.g., approximately the interferometer antenna diameter. Multiple interferometer pointings, i.e., mosaics, can also partially recover missing low spatial frequency information.

3.8 Units and Conversions

Finally, we end this Chapter 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:

$$I_\nu(\theta, \varphi) = \frac{2k\nu^2}{c^2} T_B(\theta, \varphi). \quad (3.28)$$

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

$$S_\nu = \frac{2k\nu^2}{c^2} \int T_B d\Omega. \quad (3.29)$$

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

$$\left(\frac{T}{1 \text{ K}} \right) = \left(\frac{S_\nu}{1 \text{ Jy}} \right) \left[13.6 \left(\frac{300 \text{ GHz}}{\nu} \right)^2 \left(\frac{1''}{\theta_{max}} \right) \left(\frac{1''}{\theta_{min}} \right) \right]. \quad (3.30)$$

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, we see 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.

References

Wilner, D. J., & Welch, W. J. 1994, *ApJ*, 427, 898

Chapter 4

Receivers

The ALMA front end can accommodate up to 10 receiver bands covering most of the wavelength range from 10 to 0.3 mm (30–950 GHz). Each receiver 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. In Cycle 3, Band 3, 4, 6, 7, 8, 9, and 10 are available (see available frequency and wavelength ranges for these bands in Table 4.1). The receivers are described in more detail in the following sections as well as in the references listed in Table 4.2.

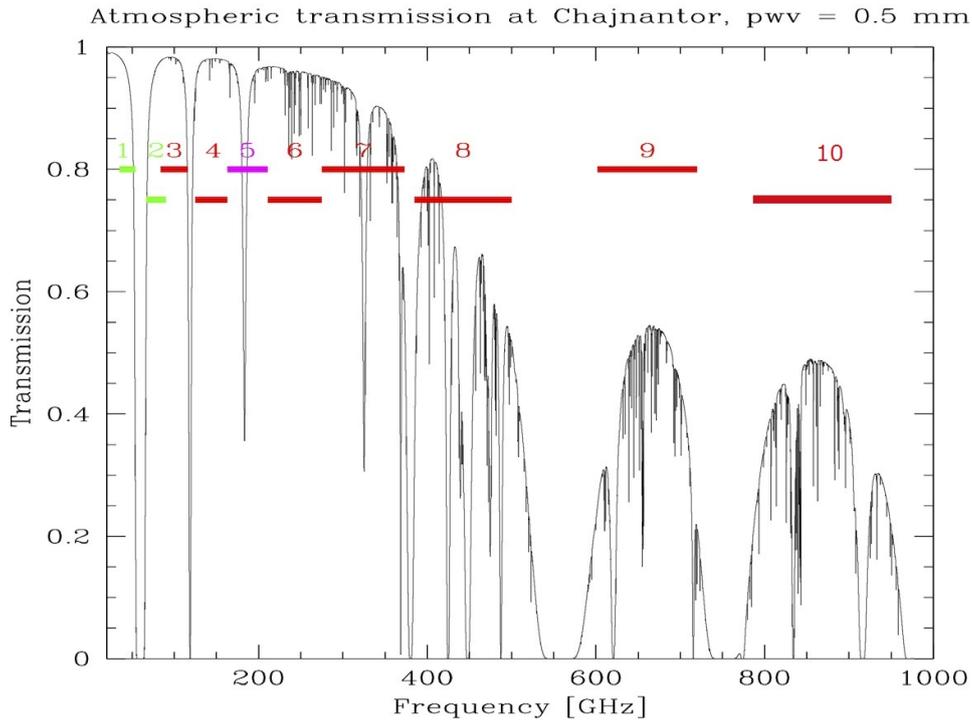


Figure 4.1: The ten ALMA receiver bands. Receiver bands for Cycle 3 are shown in red superimposed on a zenith atmospheric transparency plot at the AOS for 0.5 mm of PWV.

The ALMA receivers in each antenna are situated in a single frontend assembly (see Appendix A, Section A.4). The frontend assembly consists of a large cryostat containing the receiver cold cartridge assemblies (including SIS mixers and LO injections) and the 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) is mounted above the frontend (see Appendix A.5). Each

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	<41 ⁷	<45 ⁷
4	125.0-163.0/ 1.84 - 2.40	133 - 155	2SB	4-8	7.5	<51	<82
6	211.0-275.0/ 1.09-1.42	221 - 265	2SB	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: Receiver Characteristics. *Notes to Table:* **1.** Frequency range is the maximum available, at the extreme upper and lower limits of the IF passband. For FDM mode, the coverage is a bit smaller (See Section 6.5.1) **2.** Sideband modes: SSB means single sideband receiver, 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. However, the autocorrelation noise performance is degraded by a factor of up to about 1.5 below 5.5 GHz (Section 4.2.3) **4.** Maximum instantaneous IF bandwidth: 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 6.5.1). **5.** In future Cycles, the maximum bandwidth will double in cross-correlation mode, because both sidebands can be separated and correlated using 90-degree phase switching (see Section 6.5 and Section B.4.3). **6.** List of the minimum specification of the SSB receiver temperature (T_{rx}), unless otherwise noted, is shown. These values are the average over the IF band. The sections on individual receiver bands describe the real values measured, which in many cases are better than specifications. **7.** The specification for Band 3 receivers is $T_{rx}<41$ K at LO=104 GHz, and $T_{rx}<45$ K for any other valid LO setting. Both values should be the average over all four IFs and 4 GHz bandwidth. **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
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	2012ITST...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.

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.

To avoid overloading the cryostat cooler, only three bands can be switched on at a time. From a hardware point of view, it takes only about 1.5 seconds to switch between these bands, but in reality, switching between phase calibrator and science source can take up to three minutes. For bands that are not switched on, the time to fully thermally stabilize a receiver from an off state is up to 60 minutes - this is mainly to ensure the optimum flat bandpass shape. All of the receivers are mounted off axis in order to avoid extra rotating band selection mirrors, which necessitates an offset of the antenna to change band. This means that only one receiver can be used at a given time.

4.1 Local Oscillators and IF Ranges

The observed sky frequencies need to be down converted to frequency bands between 0-2 GHz in order to send the signals to the correlator. The frequency down conversion involves a set of Local Oscillators (LOs). The LO and Intermediate Frequency (IF) systems are described in detail in the Appendix (B).

The frontend 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, 7, and 8 have an output range of 4–8 GHz, Band 6 a range of 6–10 GHz and Band 9 and 10 a range of 4–12 GHz (Table 4.1). The possible sky frequency ranges covered by each receiver with the first Local Oscillator (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 $IF_{hi} - IF_{lo}$. 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 6.5.1).

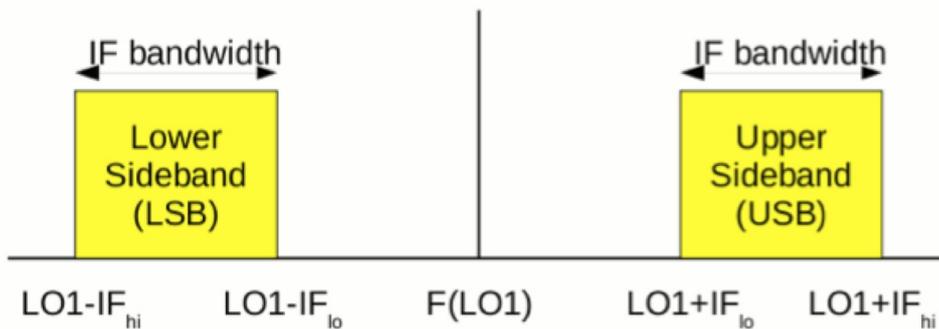


Figure 4.2: IF ranges for the two sidebands in a heterodyne receiver.

4.2 The Cycle 3 Receivers

The Band 3, 4, 6, 7, and 8 receivers are dual 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 (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 LO, TFB LO) by small and opposite amounts, which depend on the antenna, such that the signals from two antennas in the image sideband do not correlate.

The Band 9 and 10 receivers are double sideband (DSB) receivers, where the IF contains noise and signals from both sidebands. They only have two outputs, one per polarization. However, the IF effective bandwidth is 7.5 GHz per sideband (after passing through the IF processing units), so the total instantaneous bandwidth is the same as Bands 3, 4, 6, 7 and 8. In Cycle 3, only one sideband per spectral window is correlated, and the other rejected using LO offsetting, as mentioned above. This does not remove the noise from the rejected sideband. The noise of the sideband that is kept twice that of the DSB noise level. In the future, suitable phase switching will be introduced in the correlator, and both sidebands can be correlated and processed independently, thus doubling the effective system bandwidth.

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

4.2.1 Band 3 receiver

Band 3 is the lowest frequency band available in Cycle 3, covering a frequency range of 84.0–116.0 GHz (in 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 4.3). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the SIS mixers.

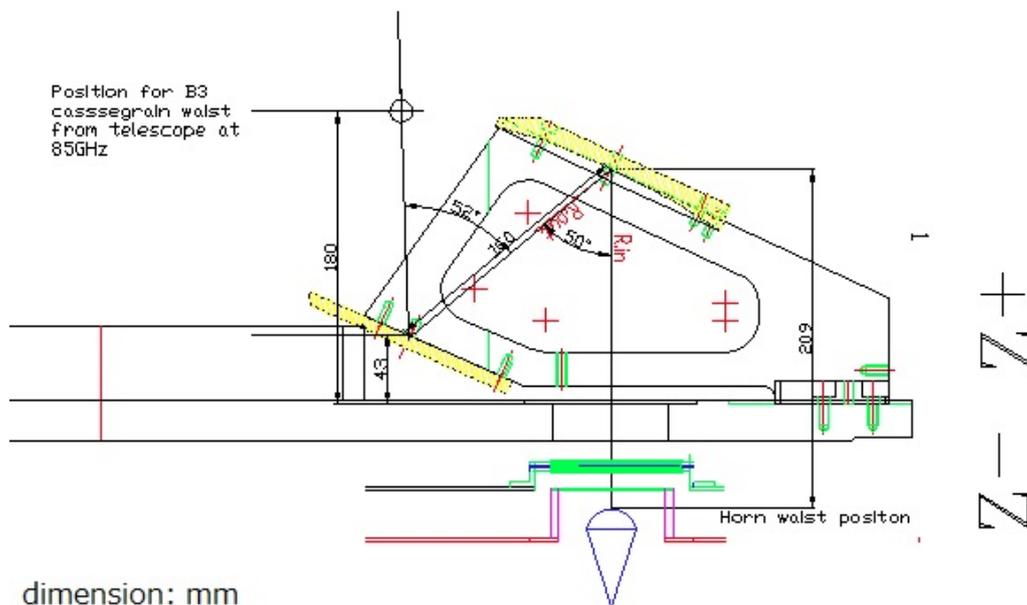


Figure 4.3: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the solid line, and the Cassegrain focus is shown by the small circle to the upper right.

A block diagram of the Band 3 receiver, including the cold cartridge and warm cartridge assembly, is shown

in Figure 4.4. The Cold Cartridge Assembly (CCA) contains the cold optics, OMT, SIS mixers and the low-noise HEMT first IF amplifiers. At room temperature, the Warm Cartridge Assembly (WCA) includes further IF amplification and the Local Oscillator covering 92–108 GHz.

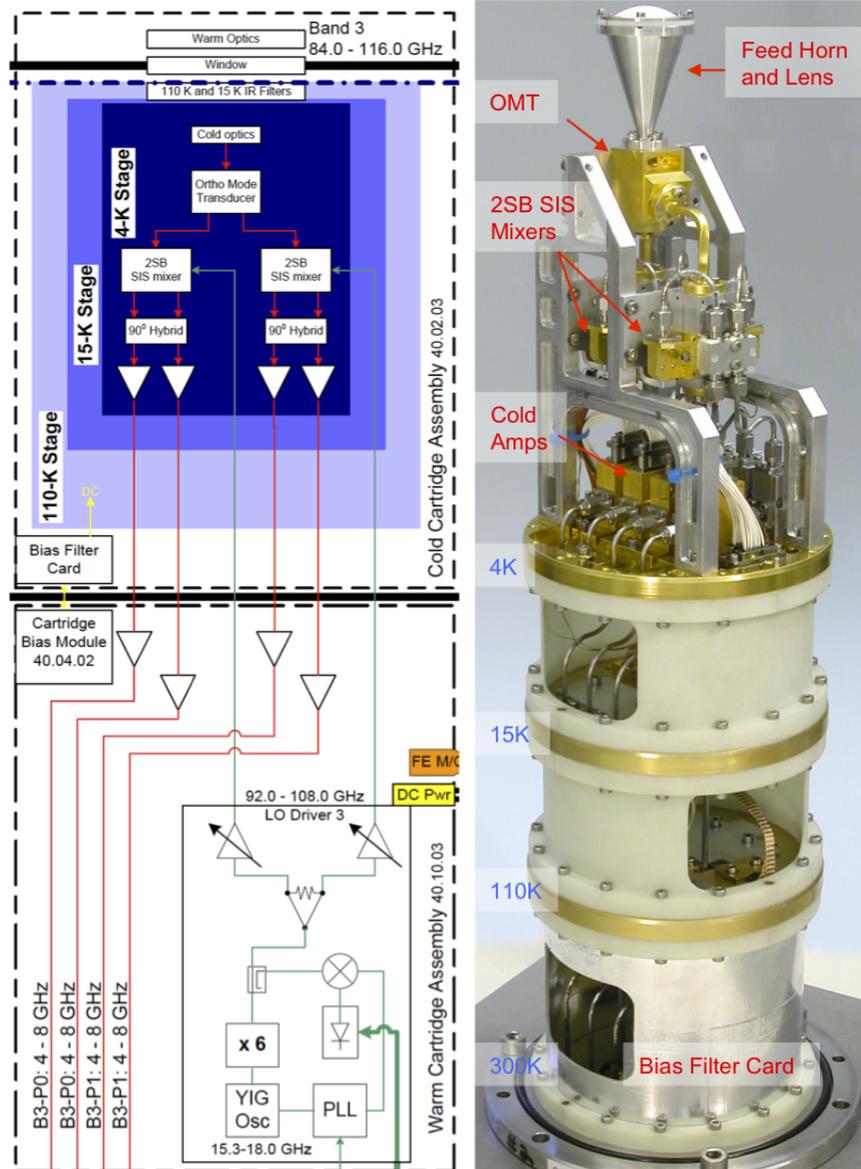


Figure 4.4: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges were constructed in Canada at NRC-HIA, Victoria.

The specification for the Band 3 receiver noise performance (T_{rx}) is <41 K at LO=104 GHz, and <45 K for any other valid LO setting. The atmospheric transmission over most of Band 3 is very high, even with a large PWV (Figure 4.5) which means observations in Band 3 can, in principle, take place with 10 mm or more of PWV. The resulting system temperature (T_{sys}) shows the expected rise at the higher end, due to an atmospheric oxygen line (Figure 4.6).

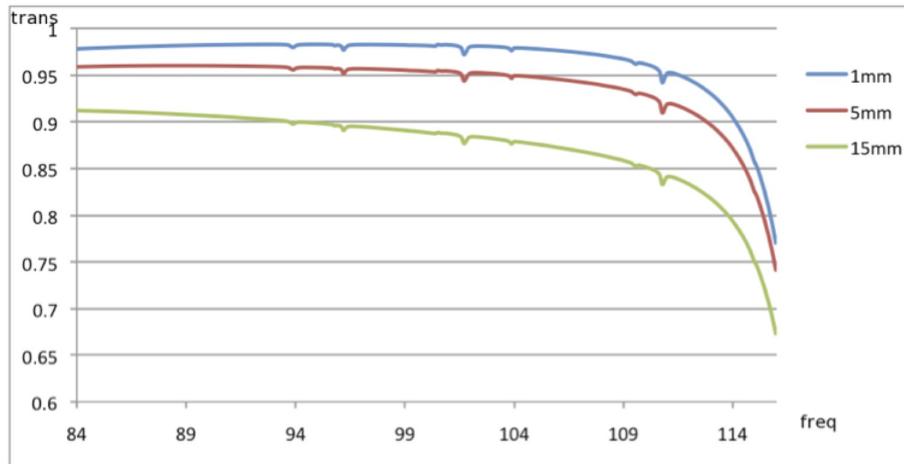


Figure 4.5: Band 3 zenith transmission for 1, 5 and 15 mm of PWV. Frequency is in GHz.

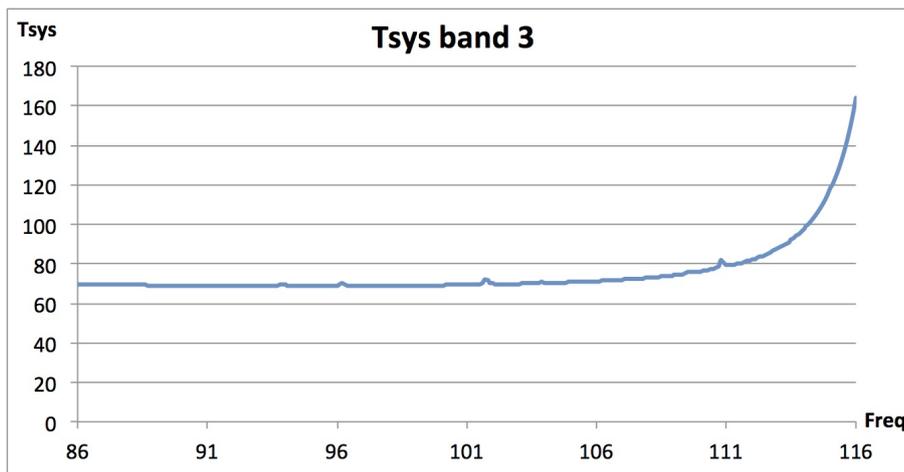


Figure 4.6: Typical system temperature (T_{sys}) at zenith for Band 3 with 1.262 mm of PWV. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.2 Band 4 receiver

The Band 4 receiver covers the 125 to 163 GHz spectral window (in the 2 mm atmospheric window). The signal collected by the telescope is focused to the Band 4 cartridge using a set of warm mirrors (Figure 4.7). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the 2SB SIS mixers.

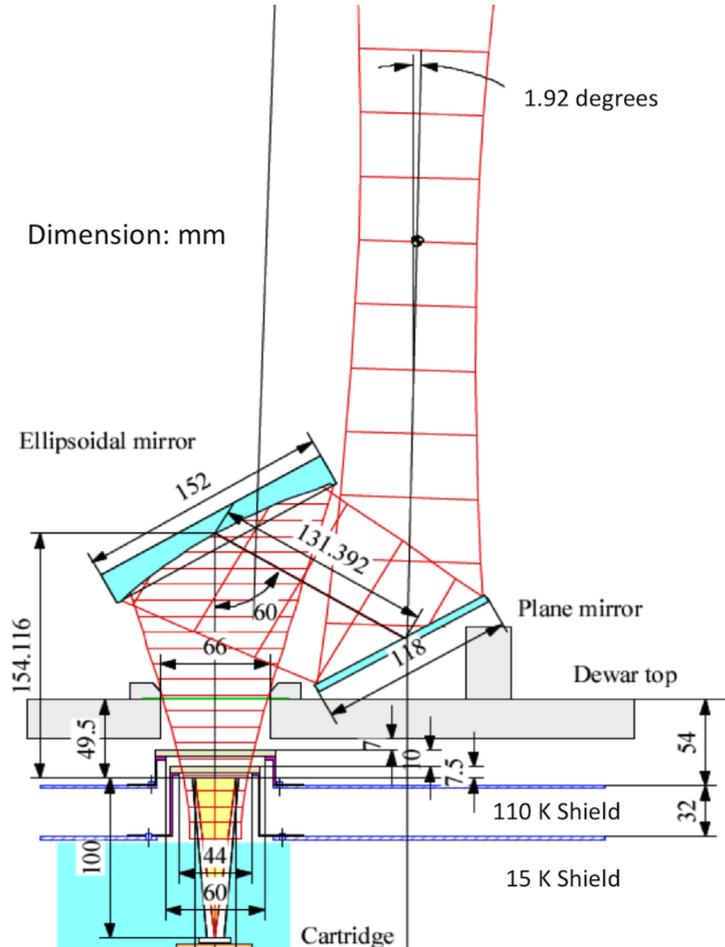


Figure 4.7: Optical layout of the Band 4. Red indicates Band 4 Gaussian beam of 5 times the beam width.

A block diagram of the Band 4 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.8. The Band 4 CCA contains a feed horn, an OMT as a polarization splitter, 2SB SIS mixer assemblies, cold IF amplifiers, isolators, and LO frequency doublers. The RF signal is down converted to 4–8 GHz using a 2SB mixer unit.

The atmospheric transmission in Band 4 is shown in Figure 4.9 for three typical PWV values. Most observations in Band 4 will be done with PWV < 5 mm. The specification for Band 4 receiver noise performance (T_{rx}) is <51 K over 80% of the band, and <82 K over the whole band (SSB T_{rx}). However, the performance of the receiver is considerably better than 50 K over the band. The resulting system temperatures (T_{sys}) for 1.262 mm PWV are shown in Figure 4.10.

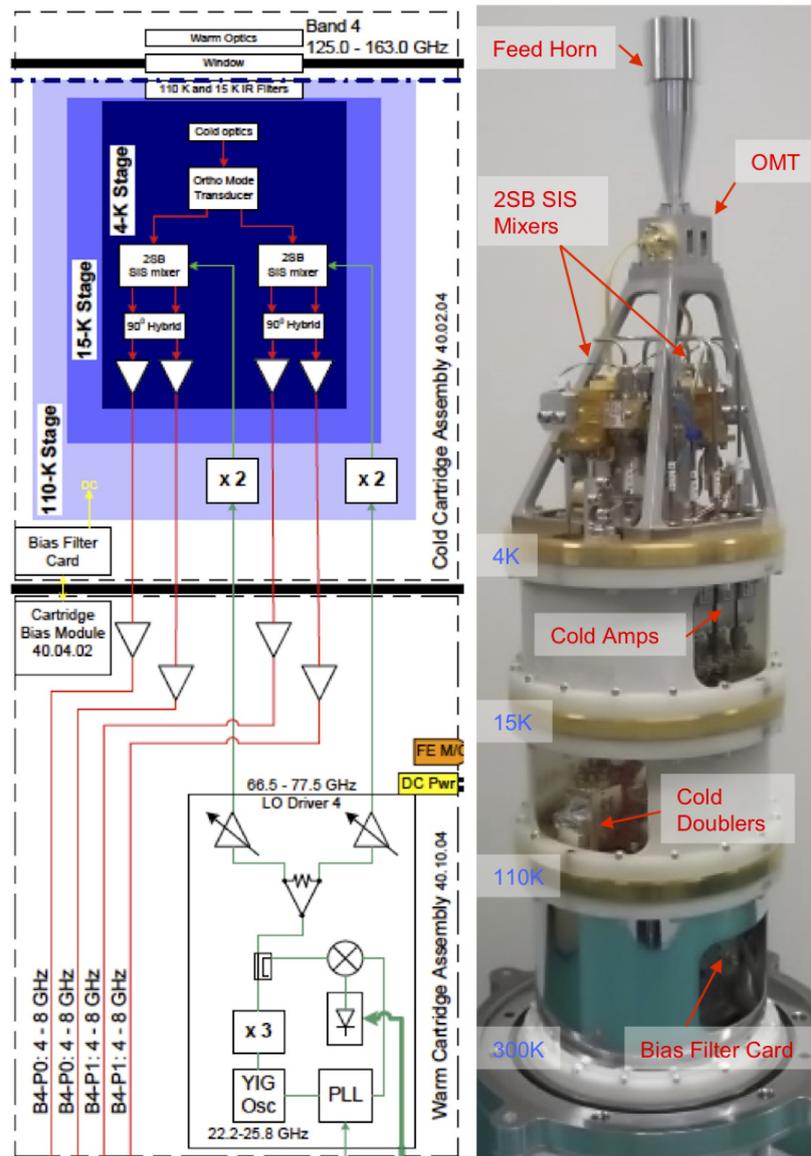


Figure 4.8: Block diagram of the Band 4 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 4 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 4 cartridges are constructed at the Advanced Technology Center (ATC).

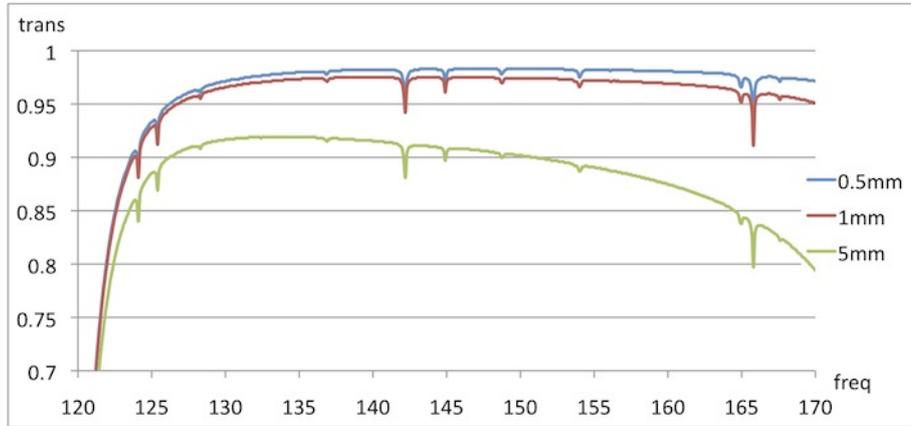


Figure 4.9: Band 4 zenith transmission for 0.5, 1 and 5 mm of PWV. Frequency is in GHz.

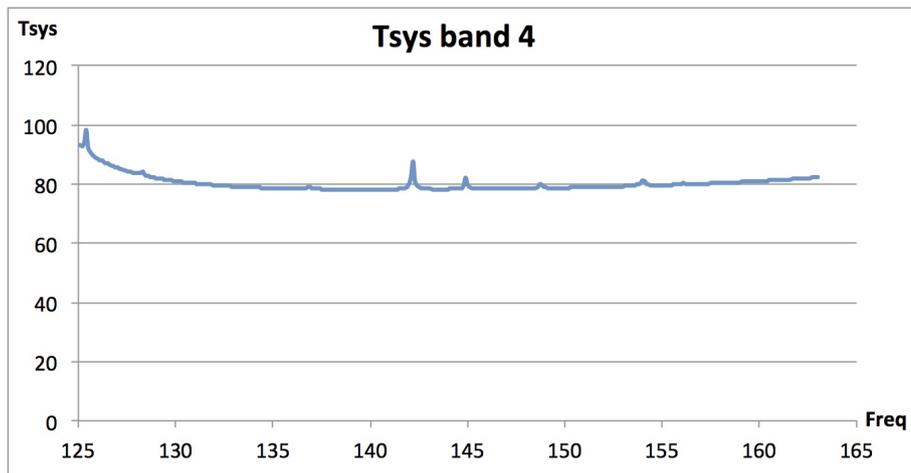


Figure 4.10: Typical system temperature (T_{sys}) at zenith for Band 4 with 1.262 mm of PWV. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.3 Band 6 receiver

The Band 6 receiver covers a frequency range of 211.0–275.0 GHz (the 1.3 mm atmospheric window). This receiver has a window with a pair of off-axis ellipsoidal mirrors inside the cryostat (Figure 4.11). A single feedhorn feeds an ortho-mode-transducer (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 4.12.

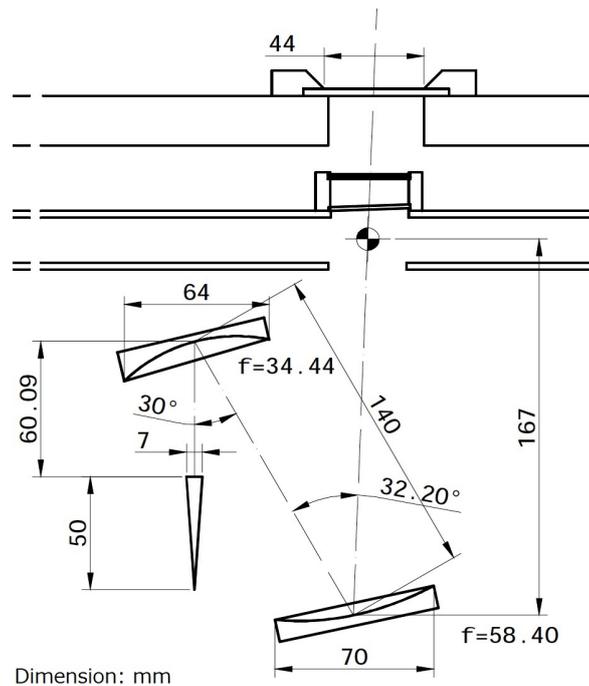


Figure 4.11: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocker. Note the slightly inclined inner window, designed to minimize standing waves.

The Band 6 IF frequency has been chosen to allow for multiple simultaneous line observations¹; it now covers the range 5.0–10.0 GHz. There is ~ 10 –25% excess noise below 5.5 GHz due to LO1, however this multi-transition setup is still considerably more efficient than observing each line separately. However, it is recommended that for continuum observations, the IF range 6–10 GHz is used. Also, it should be noted that the full range 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 4.13 for three typical 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 <138 K over the whole band (SSB T_{rx}). The measured results are considerably better, typically 50 K over most of the band. The resulting system temperatures (T_{sys}) for 1.262 mm PWV are shown in Figure 4.14.

¹ Specifically, 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 and requires the IF to reach to 5.0 GHz in order to cover all three lines

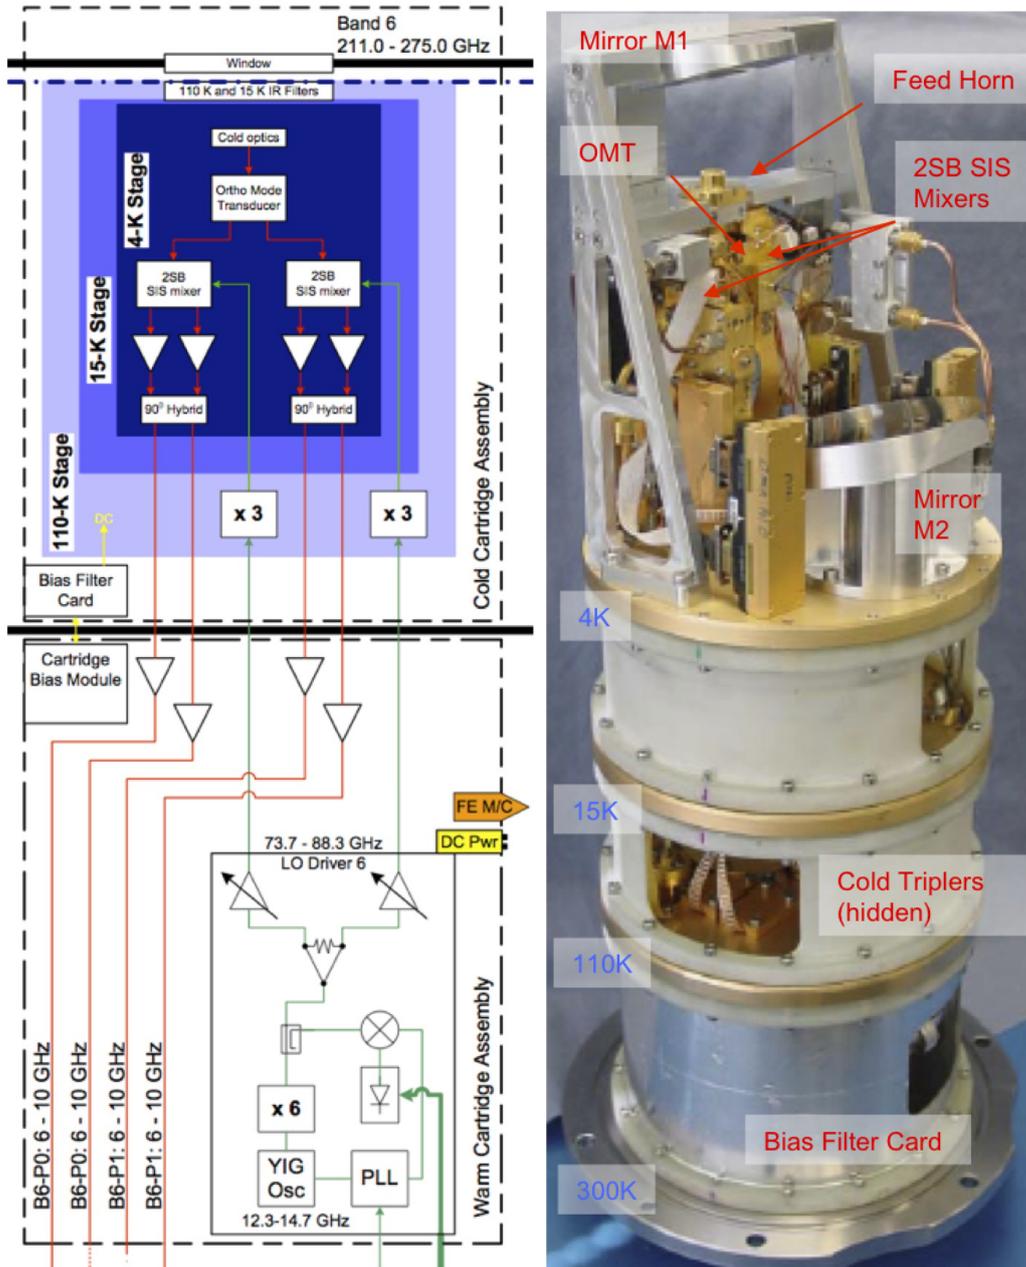


Figure 4.12: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra $\times 3$ multiplier inside the cryostat. The Band 6 cartridges were built at NRAO, Charlottesville. Note that the IF output range is actually 5-10 GHz. The range shown is the one recommended for continuum observations (see text).

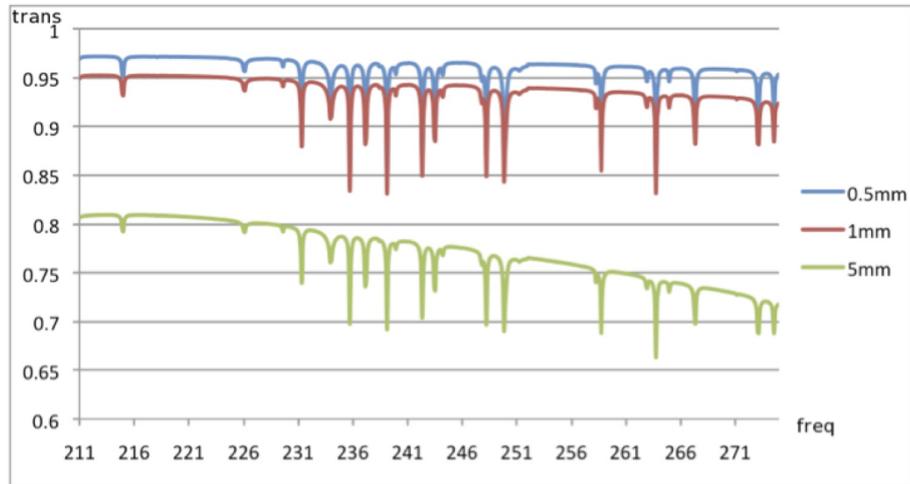


Figure 4.13: Band 6 zenith transmission for PWV=0.5, 1 and 5 mm. Frequency is in GHz. Most of the narrow absorption lines are from ozone.

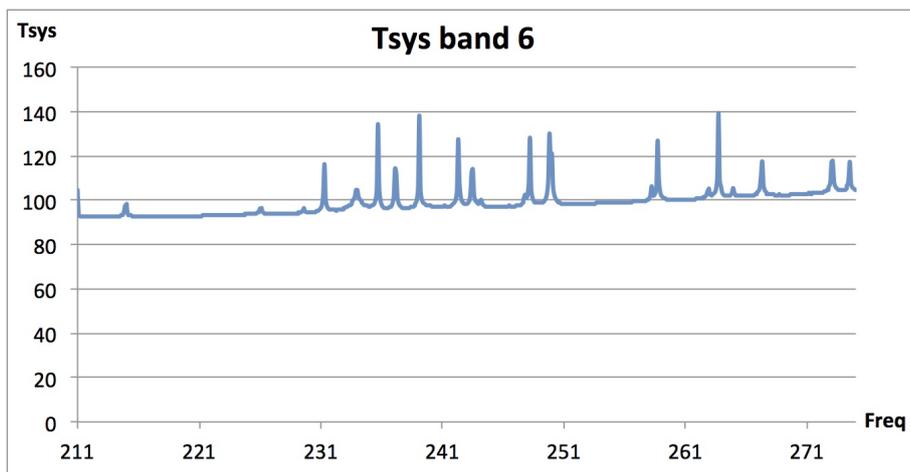


Figure 4.14: Typical T_{sys} at zenith for Band 6 with 1.262 mm PWV, based on measured values of the receiver temperatures. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.4 Band 7 receiver

The Band 7 receiver covers the frequency range 275–373 GHz (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 4.15). A block diagram of the Band 7 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.16.

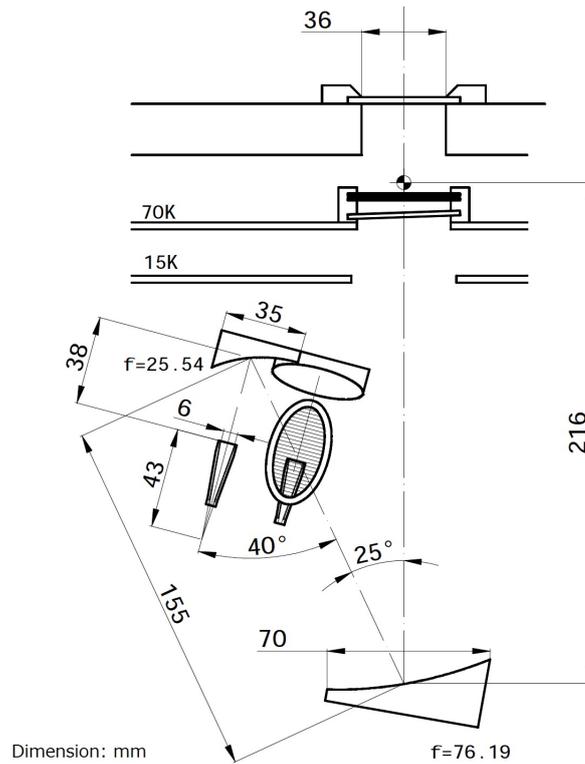


Figure 4.15: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid.

The atmospheric transmission in Band 7 is shown in Figure 4.17 for three typical PWV values. The specification of the Band 7 receiver noise temperature is $T_{rx} < 147$ K over 80% of the range and < 221 K over the whole tuning range, except at the upper end of the band (370–373 GHz), where the specifications are < 300 K SSB. However, the performance of the receiver as measured in the lab is considerably better than this. The resulting system temperatures (T_{sys}) for 1 mm PWV are shown in Figure 4.18. 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.

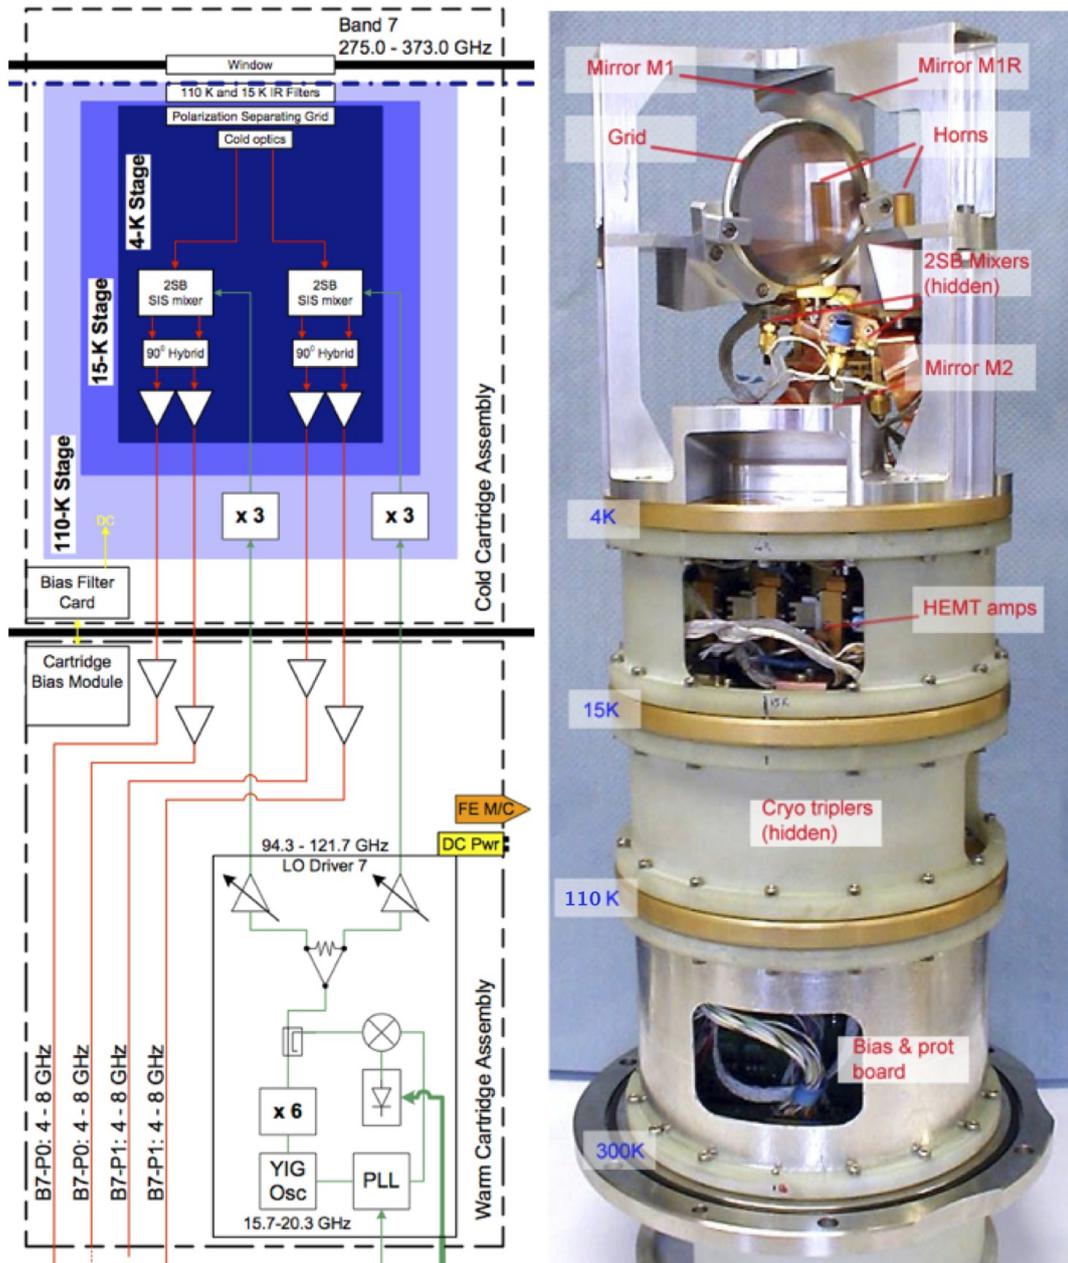


Figure 4.16: Band 7 frontend receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built at IRAM in France.

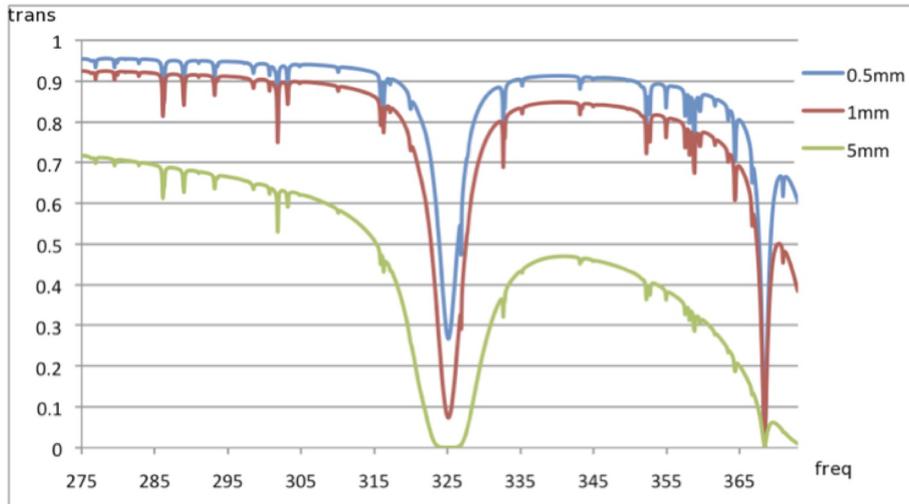


Figure 4.17: Band 7 atmospheric zenith transmission for PWV=0.5, 1.0 and 5.0 mm. Frequency is in GHz. The deep atmospheric absorption at 325 GHz is due to water, and the less prominent absorption feature at 369 GHz is due to oxygen.

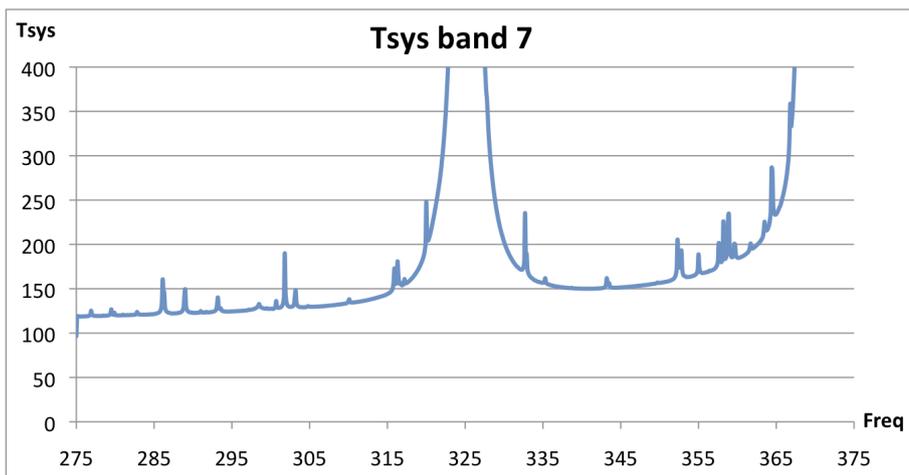


Figure 4.18: Typical T_{sys} at zenith for Band 7 with PWV=0.913 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.5 Band 8 receiver

Band 8 covers the frequency range 385-500 GHz (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 ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the 2SB SIS mixers (Figure 4.19).

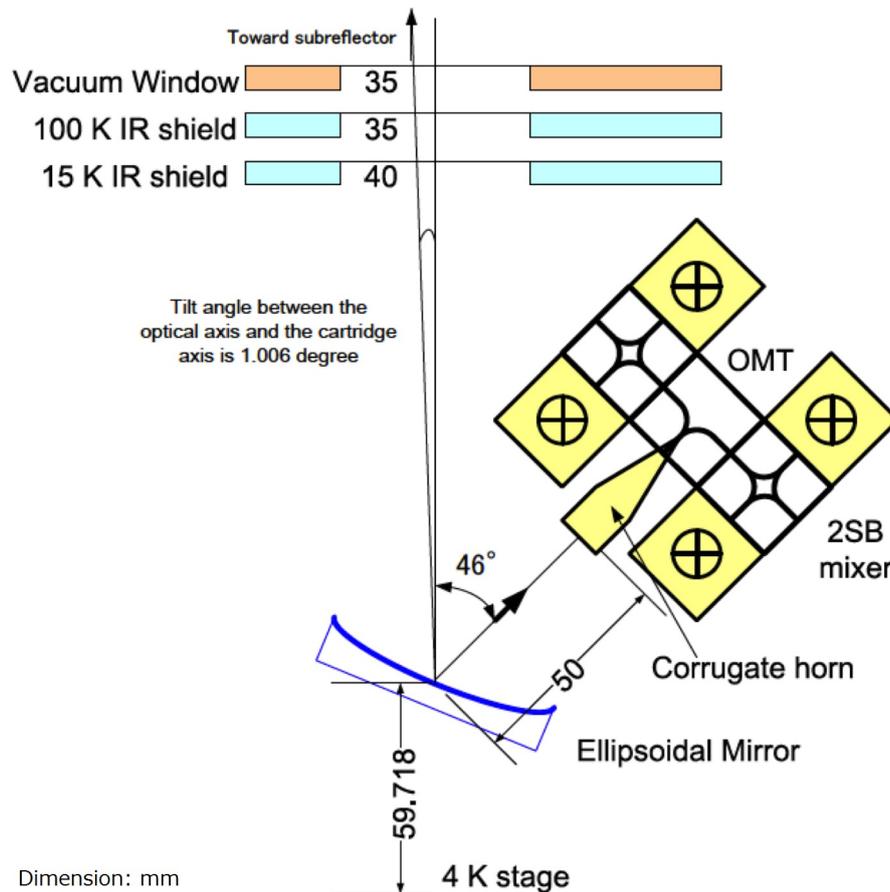


Figure 4.19: Optical layout of the Band 8.

A block diagram of the Band 8 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.20. 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 4.21 for three typical 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. However, the performance of the receiver as measured in the lab is considerably better than this. The resulting system temperatures (T_{sys}) for 0.472 mm PWV are shown in Figure 4.22.

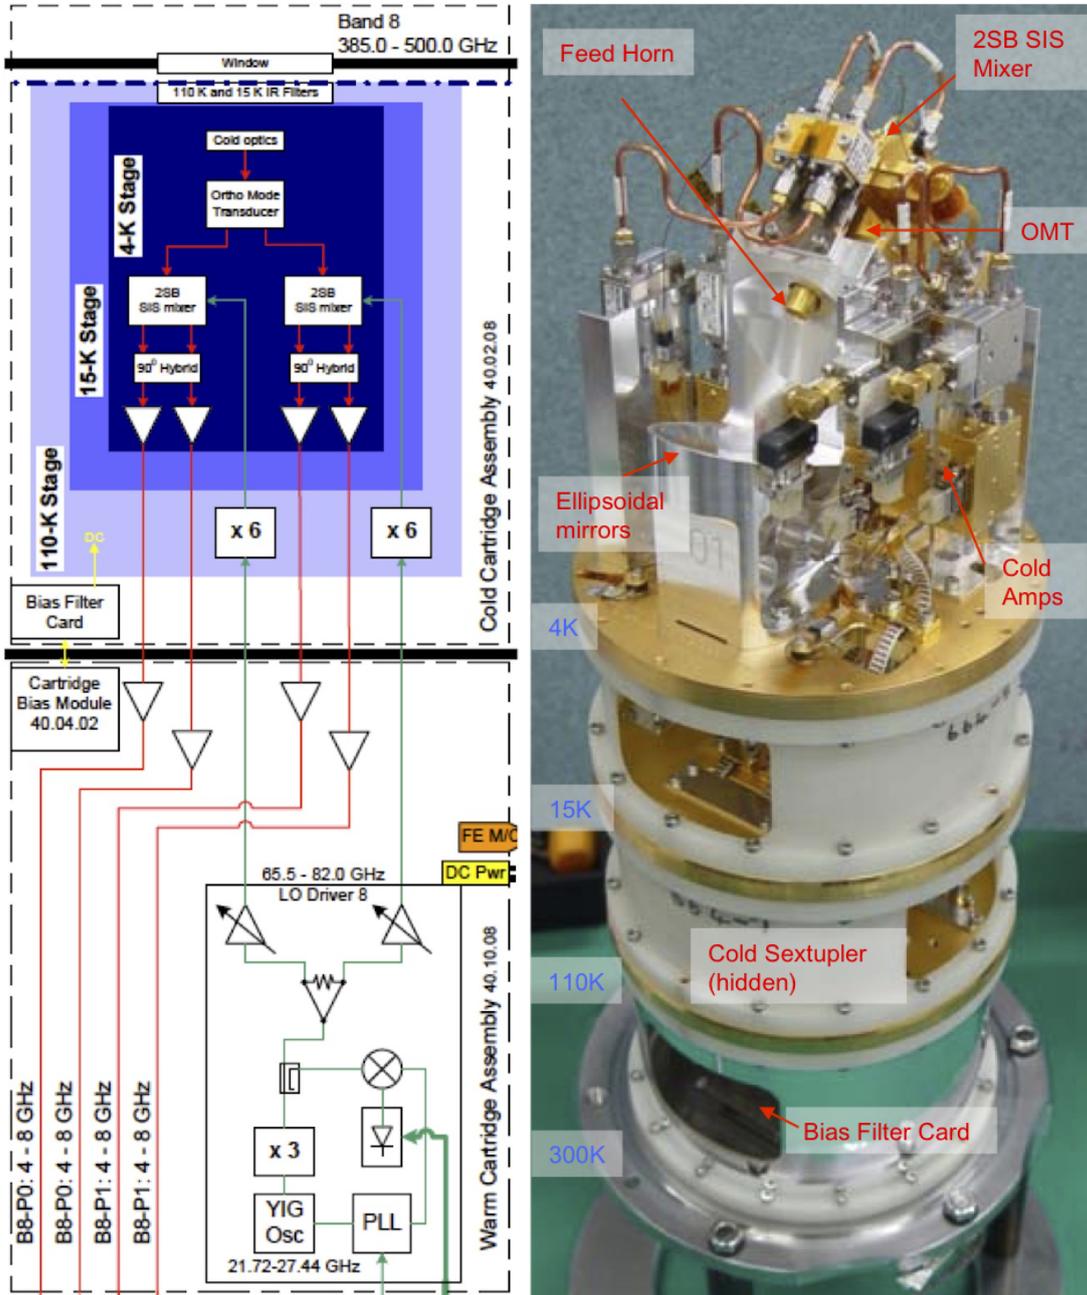


Figure 4.20: Block diagram of the Band 8 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 8 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 8 cartridges are constructed at the Advanced Technology Center (ATC).

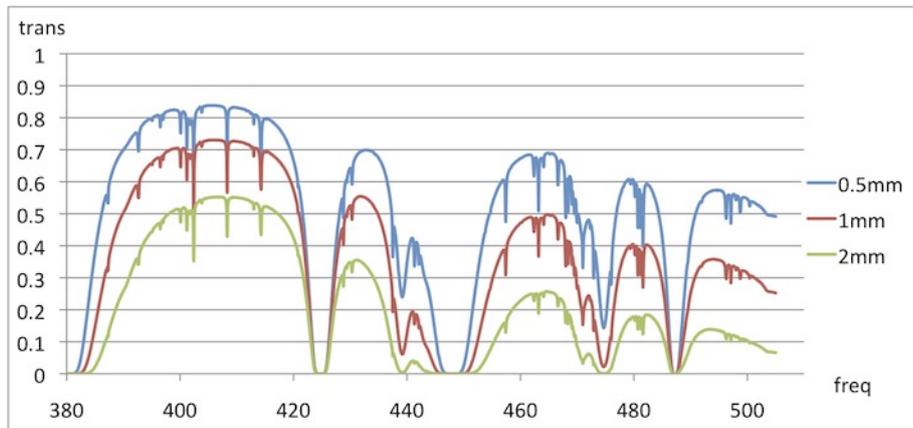


Figure 4.21: Band 8 atmospheric zenith transmission for PWV=0.5, 1.0 and 2.0 mm. Frequency is in GHz. The atmosphere in the Band 8 frequency range has some deep absorption by water and oxygen.

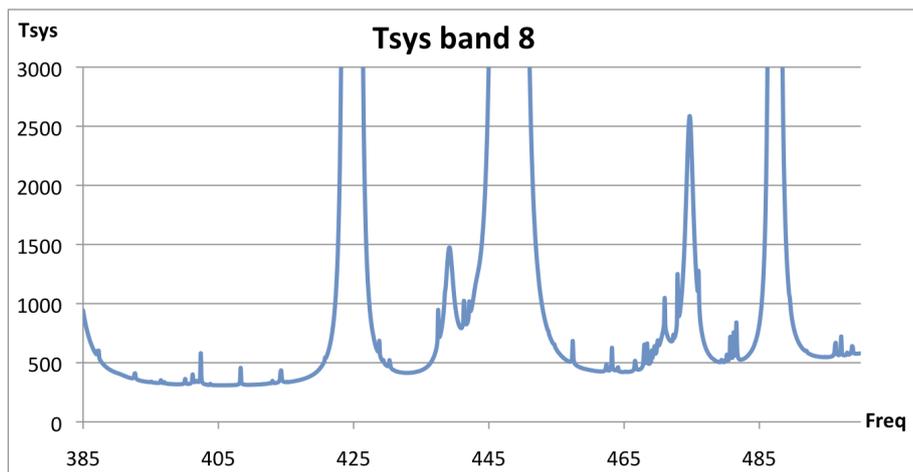


Figure 4.22: Typical T_{sys} at zenith for Band 8 with PWV=0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.6 Band 9 receiver

Band 9 covers the frequency range 602-720 GHz (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 4.23).

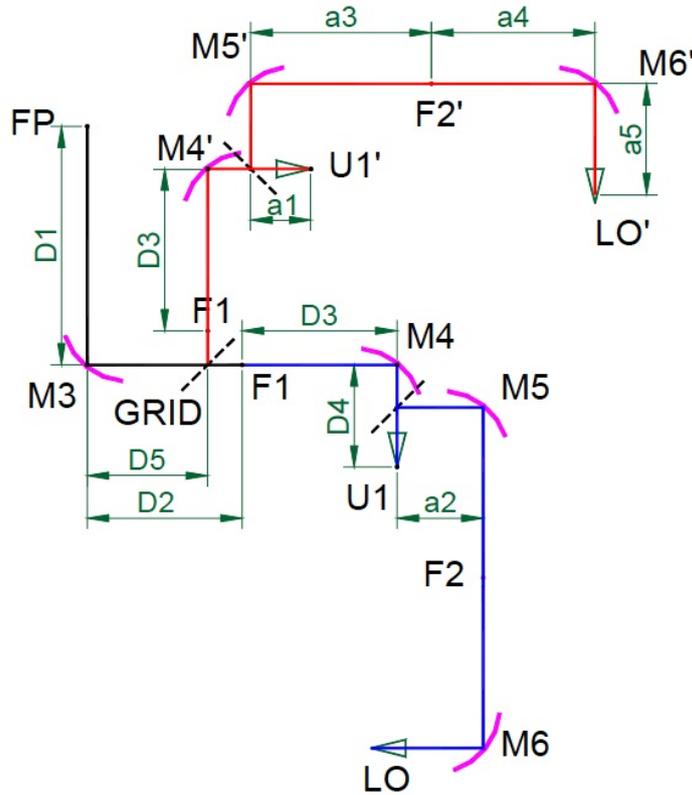


Figure 4.23: Basic Band 9 optics layout. The signal path is symmetrical for the vertical polarization (P0) and the horizontal polarization (P1). For P1 the path is as follows: a telescope focal point followed by mirror M3, grid, mirror M4, a beam splitter for LO insertion and finally the mixer horn U1. 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. In Cycle 3, LO offsetting is used to reject one of the two sidebands, which can be chosen independently for each spectral window. 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. The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 6.5.1), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.24.

The Band 9 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.25 for 3 low values of PWV. 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, and Figure 4.26 shows the expected T_{sys} for 0.472 mm of PWV, over most of the band given the expected receiver noise. Phase stability also limits when observations can be made. Therefore, most observations in Band 9 will be done at night during austral winter. As well as having a lower atmospheric transmission and a less stable atmosphere, Band 9 observing provides several challenges for observing: finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing for the relatively small primary beam, and the need for the highest level of stability in the rest of the system.

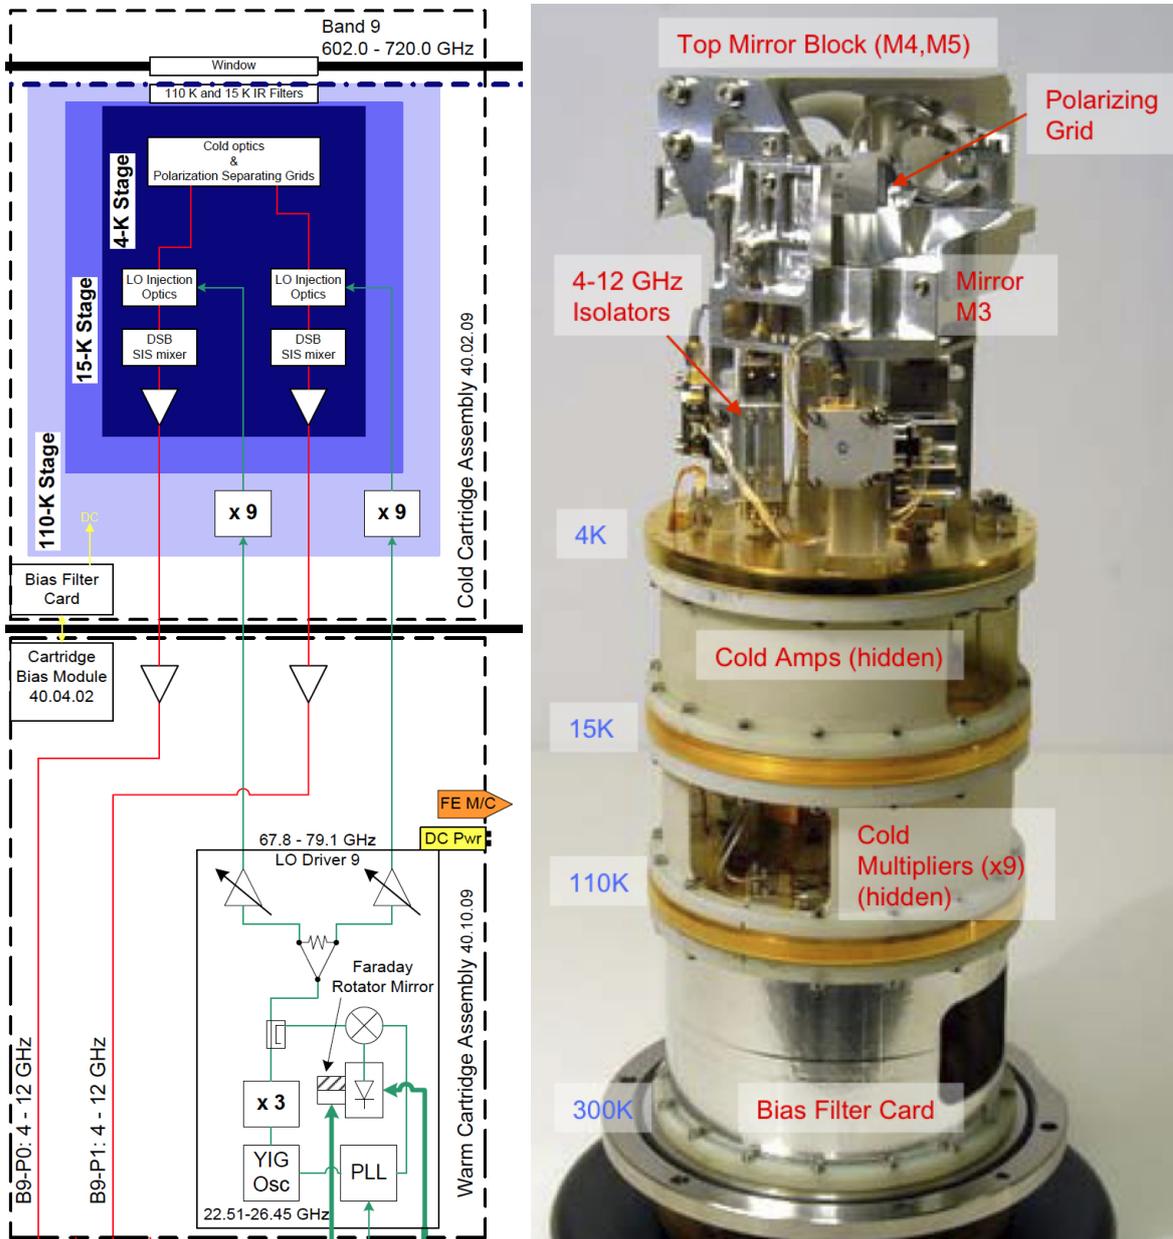


Figure 4.24: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 9 receiver was built at SRON in the Netherlands.

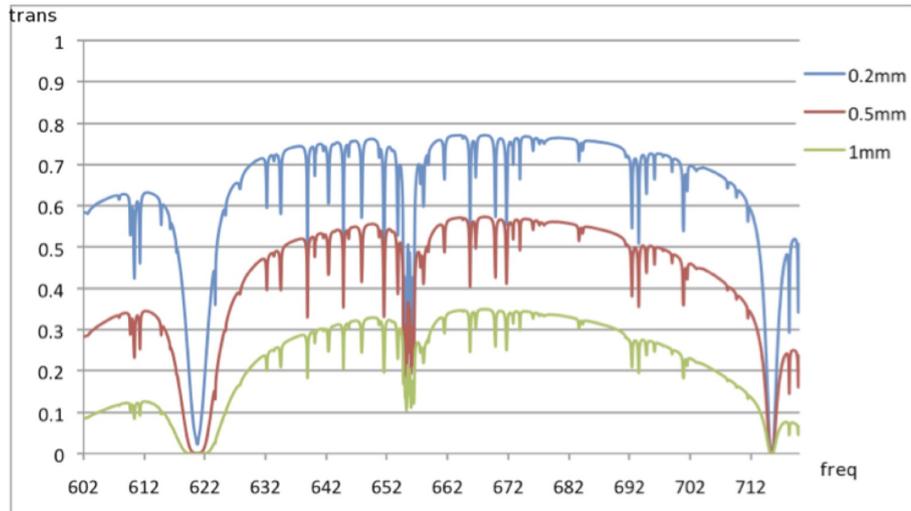


Figure 4.25: Band 9 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

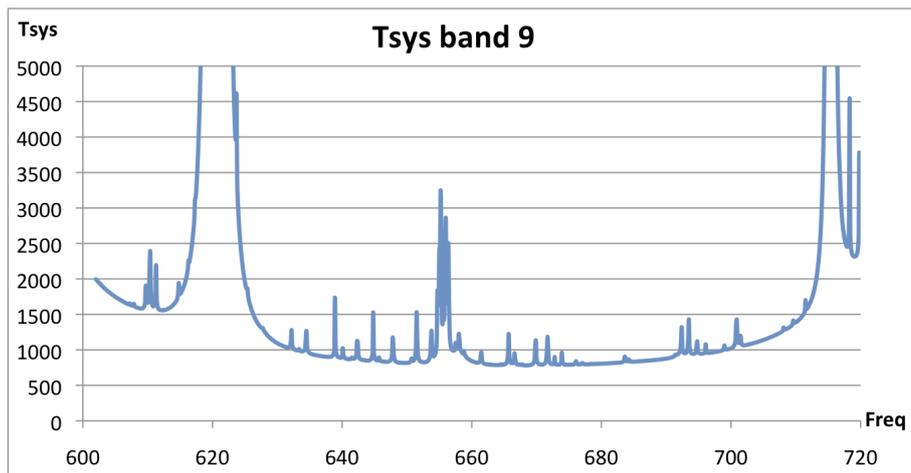


Figure 4.26: Typical T_{sys} at zenith for Band 9 with PWV = 0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.7 Band 10 receiver

Band 10 covers the frequency range 787-950 GHz (350 μm atmospheric window). Band 10 is the highest frequency receiver of the ten bands envisioned for the ALMA frontend 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 limits 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 4.27). The mixers are double sideband (DSB), and therefore LO offsetting is used to reject one of the two sidebands. The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 6.5.1), covering 4-12 GHz. A block diagram of the Band 10 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.28.

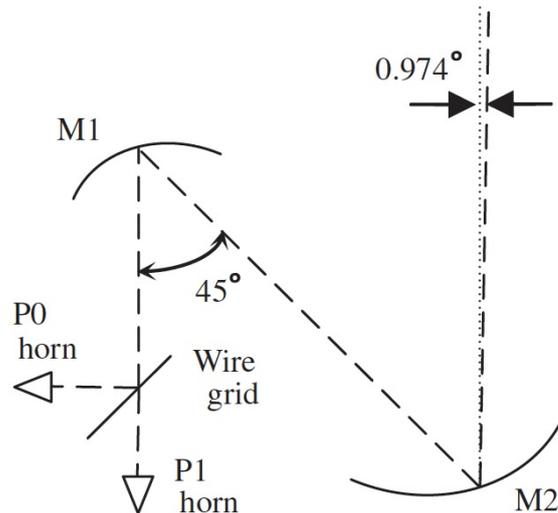


Figure 4.27: Schematic of ALMA Band 10 optics. ALMA Band 10 optics is composed of two elliptical mirrors, M1 and M2, a wire grid and two corrugated horns. The wire grid is used to separate the two linear polarizations, P0 and P1, and it is located after the two elliptical mirrors to minimize the number of optical components required.

The Band 10 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.29 for 3 low values of PWV. Figure 4.30 shows the expected T_{sys} for 0.472 mm of PWV, over most of the band given the expected receiver noise. Phase stability also limits when observations can be made. Therefore, most observations in Band 10 will be done at night during austral winter. As well as having a lower atmospheric transmission and a less stable atmosphere, Band 10 observing provides several challenges for observing: finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing for the relatively small primary beam, and the need for the highest level of stability in the rest of the system.

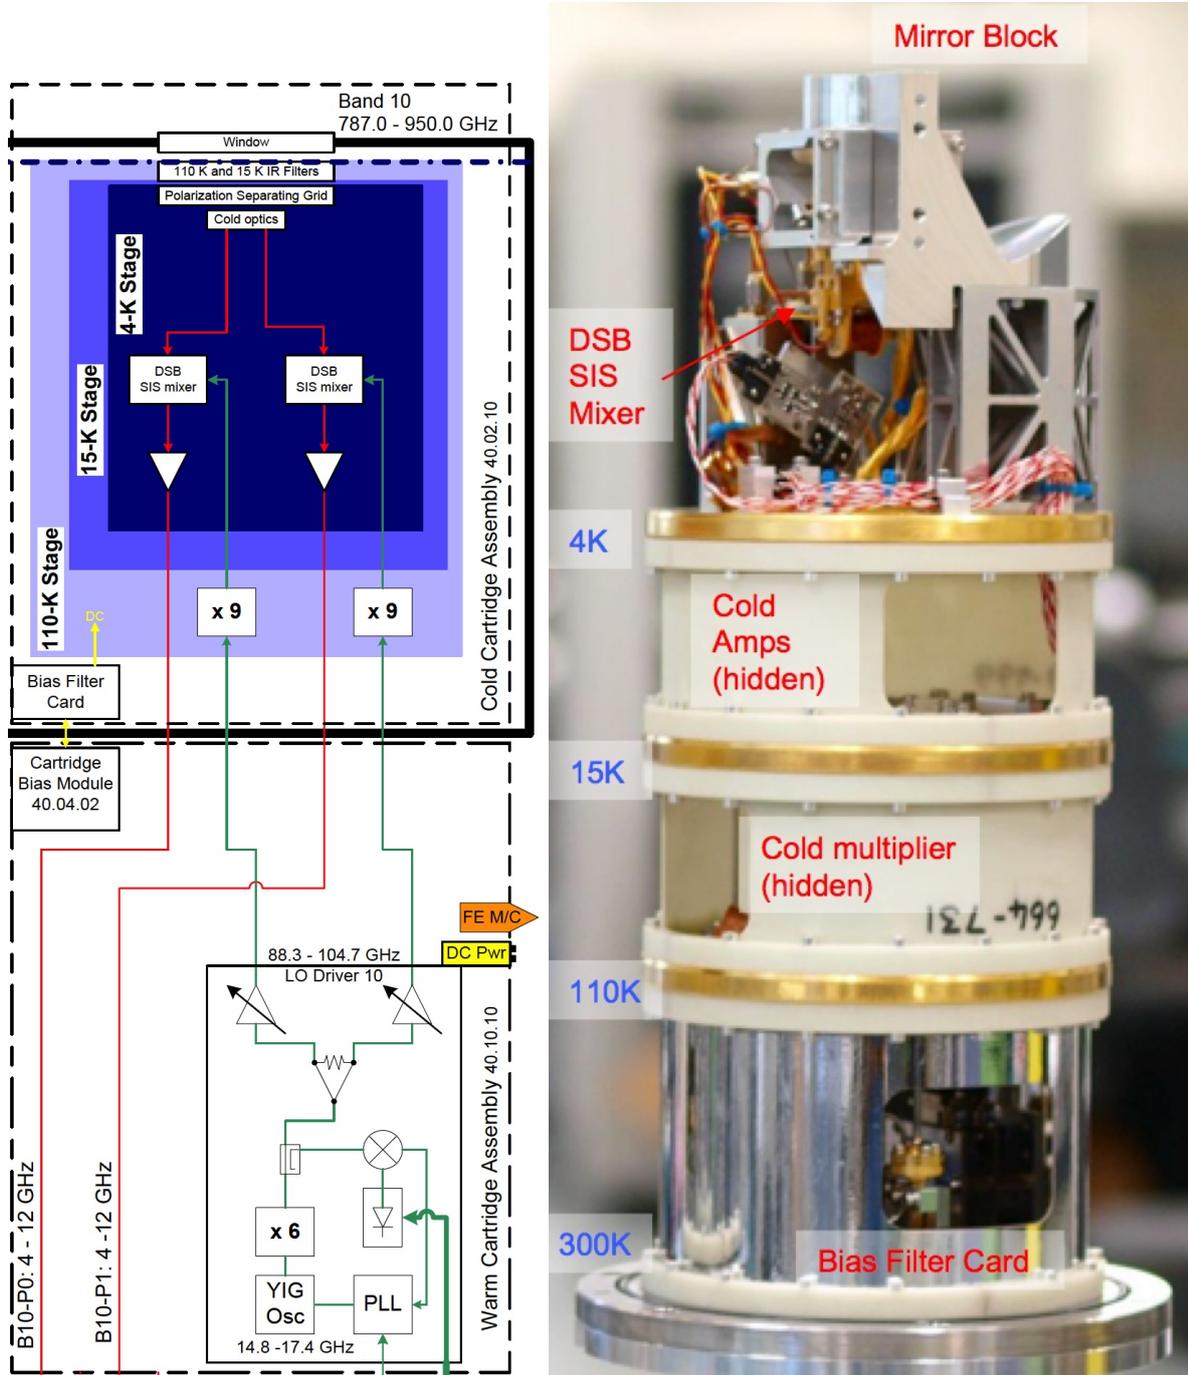


Figure 4.28: Block diagram of Band 10 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 10 cartridges are constructed at the Advanced Technology Center (ATC).

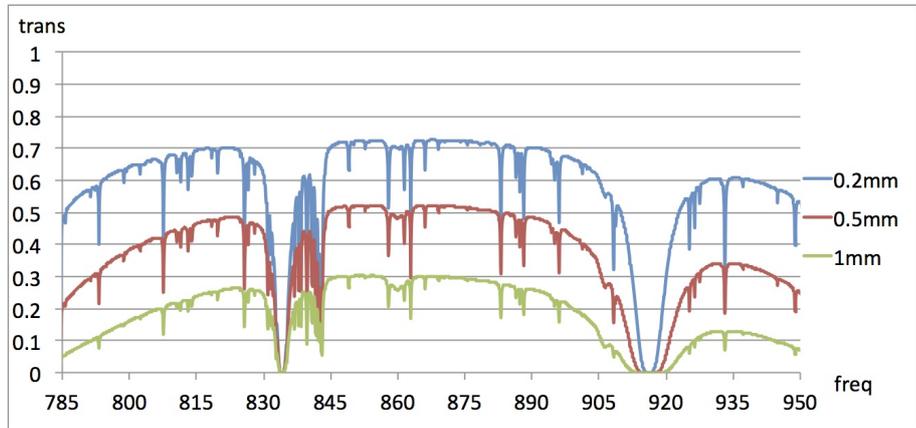


Figure 4.29: Band 10 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

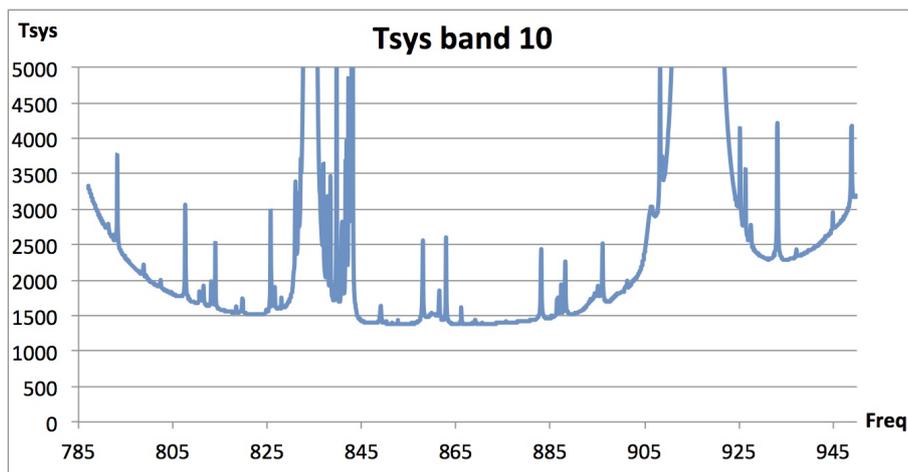


Figure 4.30: Typical T_{sys} at zenith for Band 10 with PWV = 0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

Chapter 5

The Correlators

A correlator is a virtual focal plane of an interferometer array. It accepts voltage-based signals from individual antennas, calculates cross-correlation and auto-correlations of them for each pair of antennas, and produces complex visibilities that users will receive to synthesize images. The correlator also provides pre-correlation delay and phase tracking functions to adjust the response to the wavefronts of received signals in order to maintain the coherence of the complex visibilities. Walsh-switching modulations for sideband separation (90° switching) and spurious-signal suppression (180° switching) are demodulated in the correlator¹, too.

All signals received by ALMA are processed in one of two correlators: the 64-input Correlator (also known as the Baseline Correlator) and the ACA Correlator. The 64-input64-input Correlator is used primarily for the main 12-m Array, while the ACA Correlator is used for the ACA 7-m Array and the Total Power Array (i.e. single-dish) observations, respectively². Both correlators run simultaneously and independently. Thus, while the 12-m Array observes an object using the 64-input Correlator, the ACA Correlator can be used with the 7-m Array or the Total Power Array observing either the same or a different object.

The received celestial signals 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 BaseBands (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 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 TP Array observations but also for normalization of cross power spectra and measurements of system noise temperatures. The cross-correlations are used for interferometry using the 12-m Array and the ACA 7-m Array, and also pointing and focus calibrations for all Arrays.

This chapter addresses capabilities of the correlators to be offered for Cycle 3 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, including online WVR correction, 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.

¹ Some 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° switching is modulated in the LO1 and demodulated in the digitizers. See also Section A.7.3 and Emerson 2005, ALMA memo No.537.

² Crossbar switching allows for some flexibility in this arrangement.

5.1 The 64-input Correlator

The 64-input Correlator employs hybrid design, also known as the FFX³ system (Escoffier et al. 2007, A&A 462, 801), that brings a $32\times$ spectral resolution of the traditional lag correlators (XF). 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 5.1. It consists of 4 quadrants, all of which are available for Cycle 3. Each quadrant can handle a 2-GHz dual-polarization BB for up to 64 antennas⁴. The full set of 4 quadrants is capable to accept 4 BBs to cover total 8-GHz bandwidth \times 2 polarizations.

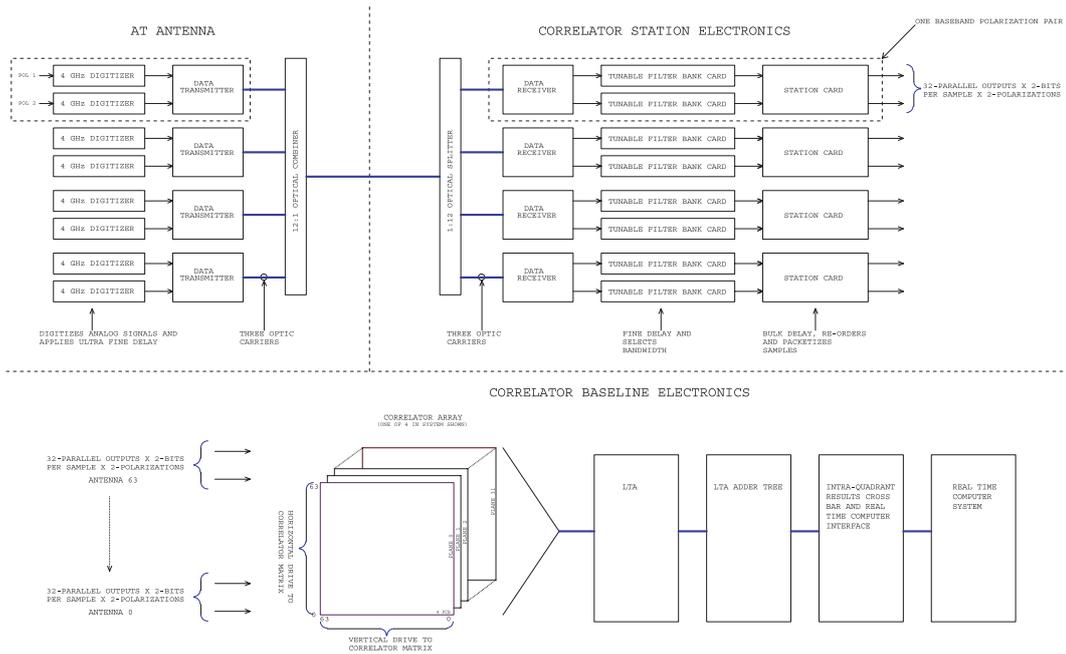


Figure 5.1: Overview diagram of the ALMA digitizers, data transmission system and 64-input Correlator (Escoffier et al. 2007, A&A 462, 801). The digitized data from the individual BBs are transferred to the Tunable Filter Banks (TFBs) and station cards (top right) for the parallelization and processing required on each antenna datastream. The correlator array (lower half) produce auto-correlations on the matrix diagonal and cross-correlations elsewhere. The correlations are then integrated in the Long-Term Accumulators (LTAs, lower right).

5.1.1 TDM mode

TDM is mostly used for continuum observations. The 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, sideband ratio measurements, etc.

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 (LSB) to reduce quantization levels from 3- to 2-bit per sample (see Section 5.3).

The TDM mode provides a SPectral Window (SPW)⁵ that consists of up to $256/N_{\text{pol}}$ channels per BB,

³ F, X and F stand for filtering, correlation and Fourier transform, respectively.

⁴ $\frac{N_{\text{ant}}(N_{\text{ant}} - 1)}{2} = 2016$ baselines and 64 auto-correlations for $N_{\text{ant}} = 64$.

⁵ An SPW is a contiguous spectrum whose frequency channels are uniformly spaced. See also Chapter 6 about the relation



Figure 5.2: 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

where N_{pol} is the number of polarization products per BB⁶ As the full 2000 MHz BB is covered, this requires some truncation of the band-edge channels in offline data processing – see Section 6.5.1.

5.1.2 FDM mode

FDM is used for spectral-line observations that require a higher spectral resolution than of TDM.

Each 2-GHz BB is split into 32 62.5-MHz sub bands via the TFBs that employ digital filtering with Field Programmable Gate Arrays (FPGAs). A TFB outputs digital 2-bit (4-level) re-quantized⁷ digital signals with the Nyquist sample rate of 125 MHz. The sub bands are trimmed by 15/16 to avoid aliasing and band-edge filter responses, and are set at intervals of 58.59375 MHz. The correlator stitches up multiple sub bands to output a seamless cross power spectrum to form an SPW with $\frac{15}{16} \times 8192/N_{\text{pol}}$ 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 sub bands to form an SPW is selectable from 32, 16, 8, 4, 2, or 1 to cover 1875, 937.5, 468.75, 234.375, 117.1875, and 58.59375 MHz, respectively. It is capable to have multiple (up to 4) SPWs in the same BB. While different spectral setups can be set for different SPWs, multiple SPWs in the same BB must have the same channel spacing.

between BB and SPW.

⁶ $N_{\text{pol}} = 2$ for standard Stokes-I observations that employ XX and YY products, and $N_{\text{pol}} = 4$ for full-Stokes observations.

⁷ Although TFB is capable at 2-, 3-, or 4-bit re-quantization, the cross-correlation processor accepts only 2 bits in Cycle 3. See also Section 5.3.

The center frequency of each SPW can be tuned over the 2 GHz-wide BB using a digitally-synthesized LO (the TFBLO, or LO4), with some constraints. For examples, the edges of the full bandwidth of the sub-bands cannot fall outside the 2 GHz BB range. The tuning must be at a 30.5-kHz step (see LO Appendix B in the Appendix). See also Section 5.5.2 for detail about spectral setting.

5.1.3 Correlation and realtime processing

The correlation module performs the multiply-and-add operations to produce 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 of 64 antennas⁸ with 8192 lags for a BB with a dual polarization. It is possible to set different modes in different quadrants (i.e. BBs); for example, while one BB set in TDM, other BBs can be set in FDM. For detail spectral setups, see Chapter chap:spectral-setup. The Long Term Accumulator (LTA; see Figure 5.1) takes short 1 ms or 16 ms integrations from the correlator circuits and provides longer term integration. Further time averaging over a dump period is performed in the CDP (Correlator Data Processor) computers. See also Section 5.5.3 about time resolution. Section 5.5 describes in detail about the processing of the data in the correlator.

5.2 The ACA Correlator

The ACA Correlator is dedicated for observations with the Atacama Compact Array (ACA - Morita Array) that consists of the 7-m Array with twelve 7-m antennas and the Total Power Array (TP Array) with four 12-m antennas⁹. The ACA Correlator employs the FX¹⁰ design, whereby the incoming time-domain data stream is converted into frequency-domain spectrum via Fast Fourier Transform (FFT) before cross multiplication to form power spectrum. The FFT part always accepts a full 2-GHz bandwidth and outputs 524288-ch spectra with a frequency resolution of 3.815 kHz. This design allows us flexible spectral handling as shown in Figure 5.6: 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, Frequency Profile Synthesis (FPS) is performed in the Correlator Data Processor (ACA-CDP) computer so that the outputs of both correlators are matched (Kamazaki et al. 2008, ALMA Memo 580).

A detailed diagram of 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.

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 performs the 2^{20} -point FFT that produces a 2^{19} -point complex spectrum¹¹ (hereafter, voltage spectrum) for every BB per antenna, with a channel separation of 3.815 kHz ($= 2 \text{ GHz} \div 524288 \text{ ch}$) 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 corresponds 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 output to the ACA-CDP computers in the Computing subsystem through optical fibers.

⁸2016 cross-correlations (real and imaginary) and 64 auto-correlations (real) totals 4096 real correlations.

⁹ Since the sub-array function will not be available in Cycle 3, the 7-m and the TP array observations cannot be operated simultaneously.

¹⁰ F and X stand for Fourier transform and correlation, respectively.

¹¹ 2^{20} -point FFT produces a 2^{20} -point complex voltage spectrum of a double sideband including the image sideband that will be discarded.

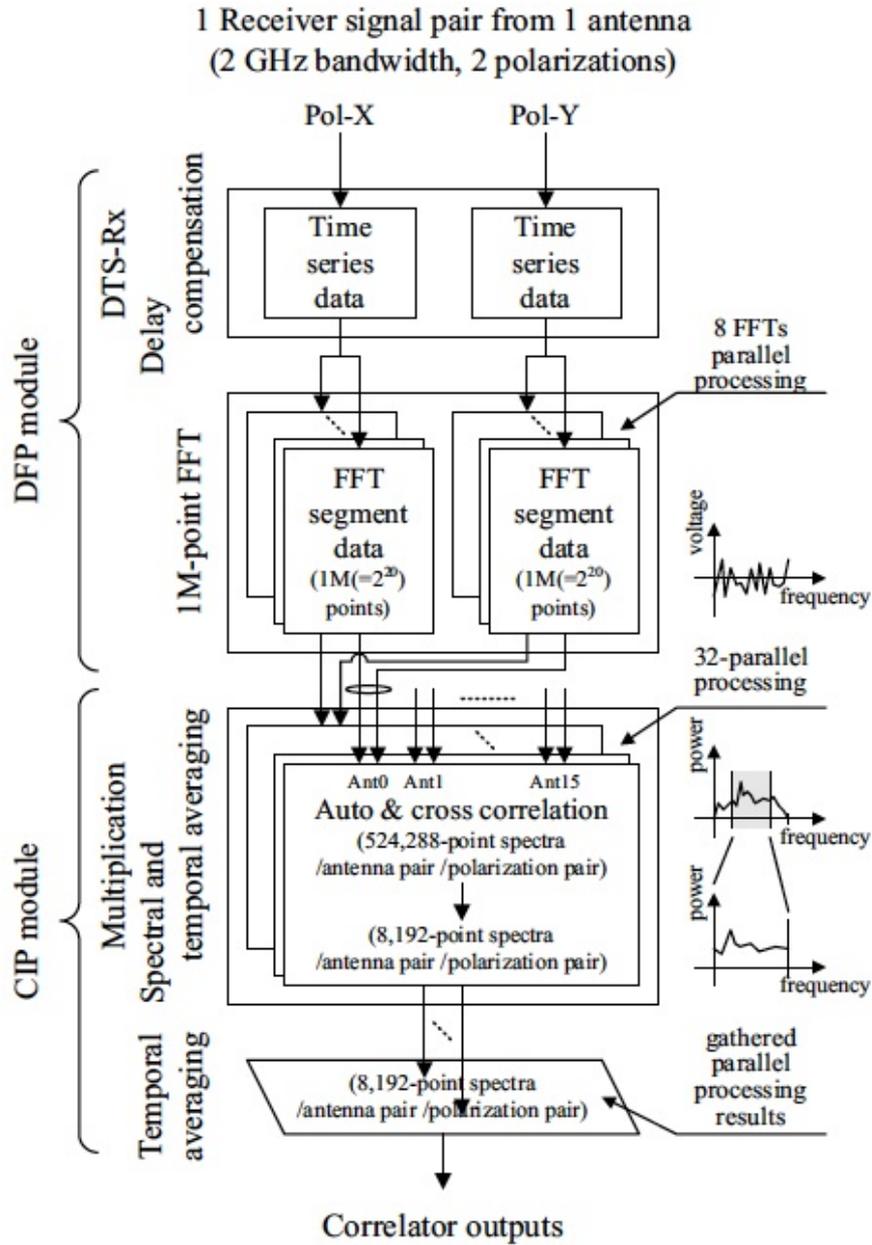


Figure 5.3: Block diagram of ACA Correlator. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. FFT is performed using FPGA, resulting in signals in the spectral domain which are then correlated. The auto and cross-correlation data are then accumulated 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.



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

The ACA-CDP performs further spectral processing such as FPS and temporal integration, before data is 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 detail, see Kamazaki et al. 2012, PASJ 64, 29.

5.2.1 7-m Array

The 7-m Array consists of twelve 7-m antennas whose correlation is dedicated in the ACA Correlator. The purpose of the 7-m Array is to cover low spatial frequencies in the (u, v) coverage and to enhance image fidelity in extended sources. Since the visibilities of the 7-m Array will be combined with those of the main array, these visibilities must be homogenized in terms of amplitude, phase, and spectral specifications. While amplitude and phase calibrations refer to standard flux and phase calibrators offline, the spectral profile will be matched via FPS in the CDP online processing.

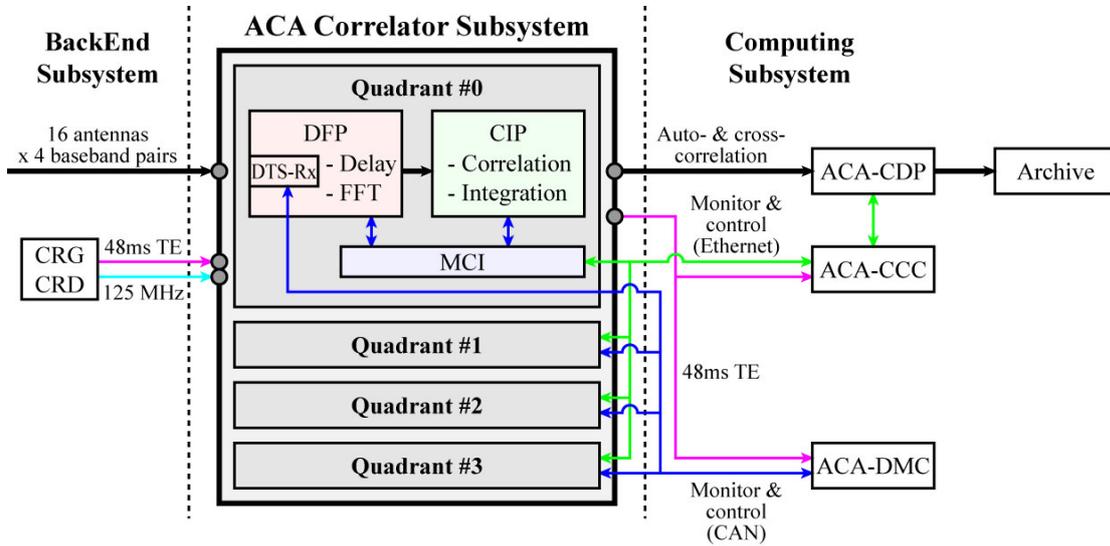


Figure 5.5: Overview diagram of the ACA Correlator. The signals from the BackEnd Subsystem are processed in the DFP and CIP modules and output to the Computing Subsystem (see text and Figure 5.3 for details of these). The modules are controlled by the Monitor and Control Interface (MCI), which communicates 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.

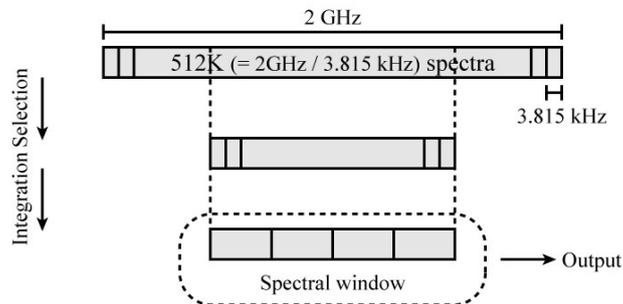


Figure 5.6: Data capture sequence of the ACA Correlator. The frequency range of an SPW is selectable from the 524288-point (512K) spectra across 2-GHz bandwidth. Additional channel-averaging within the selected frequency range before output.

5.2.2 Total Power Array

The TP (Total Power) Array consists of four 12-m antennas in the single-dish observation mode, for the purpose of obtaining zero-spacing visibilities to recover the spatially-integrated flux densities that are not obtained with cross-correlations. In the TP Array observations, the ACA Correlator generates only parallel polarization auto-correlations of XX and YY . Cross-polarization products are not taken in Cycle 3. Single-dish dedicated observing modes such as position switching, raster scan, and OTF will be employed for the TP Array. Fast switching (~ 10 Hz) using the nutating subreflector will not be capable in Cycle 3.

5.3 Digitizers

A digitizer is a device to convert continuous waveforms (analog) into digital form sampled at discrete timings and quantized voltage levels. The ALMA digitizers are equipped in the antenna back ends (BEs) to convert BB signals into the ALMA digital format (Freund 2002, ALMA Memo No.420). 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 $4 \text{ BBs} \times 2$ polarizations. To cover the analog BB signals at the frequency band of 2–4 GHz, the digitizer samples at the Nyquist frequency of 4 GHz with the voltage quantization of 8 levels (3 bit per sample). However, the 64-input Correlator reduces the quantization levels to 2-bit per sample at the station card. The ACA Correlator handles 3-bit digital form in FFT and the spectrum is re-quantized in 4-bit before cross multiplication. These quantization processes affect sensitivity and linearity in power spectra. The quantizations in the 64-input and the ACA Correlators yield sensitivity loss of 12% (TDM) or 16% (FDM) and 4.8%, respectively, if input power of analog signals is adjusted at the optimal level. See also Section 5.6.1 for detail.

5.4 Online WVR Correction

Water Vapor Radiometer (WVR; see Appendix A.6 for technical description) is a device to measure the amount of Precipitable Water Vapor (PWV) in order to correct pathlength fluctuations in the troposphere. The CDP of the 64-input Correlator applies a correction of the pathlength fluctuation, for each dump duration, based on the measurements of WVRs equipped on 12-m antennas. The online WVR correction reduces variance in visibility phases and enables longer-time integration of the visibilities (thus slower output rate) keeping coherence. Although the time integration is a irreversible process, the WVR data are recorded together with the visibilities.

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

5.5 Capabilities of the Correlators

5.5.1 Polarization

Both correlators are capable of polarimetry by producing correlations of $\langle XX^* \rangle$, $\langle XY^* \rangle$, $\langle YX^* \rangle$, and $\langle YY^* \rangle$. These correlations relate to the Stokes visibilities of I , Q , U , and V as

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

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

Chapter 8.

The 64-input Correlator forms cross-polarization products of $\langle XY^* \rangle$ and $\langle YX^* \rangle$ only when polarimetry is required. Although the ACA Correlator is capable of polarimetry, the cross-polarization products of the 7-m Array or the TP Array observations will not be offered for Cycle 3.

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 output multiple SPWs within the bandwidths. Each SPW corresponds to a continual spectrum composed by uniformly-spaced spectral channels. Table 5.1 summarizes the spectral performances of the correlators. See also Chapter 6 about multiple-SPW setup.

Correlator/Mode	Bandwidth (MHz)	Number of Ch.	Ch. spacing (MHz)
TDM	2000 ^a	128	15.625
FDM	1875	3840	0.488
FDM	938	3840	0.244
FDM	469	3840	0.122
FDM	234	3840	0.061
FDM	117	3840	0.0305
FDM	58.6	3840	0.0153
ACA/FPS	2000/1937.5 ^b	128/124	15.625
ACA/FPS	2000/1992.2 ^b	4096/4080	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. TDM and FDM stand for the Time Division Mode and Frequency Division Mode of the 64-input Correlator. For the ACA Correlator, spectral performances with/without FPS (Frequency Profile Synthesis) are described. FPS doesn't reduce available bandwidths in narrower bandwidth modes. NOTE: The channel spacing is *not* the same as the spectral resolution because of the applied weighting function; for calculating spectral resolution. *a*: Usable bandwidth except band edges is ~ 1875 MHz. *b*: Application of FPS trims both band edges for the case of 2-GHz bandwidth but other narrower cases.

The spectral profile which the correlators output is a convolution of the true spectrum with the spectral resolution function shown in the figure of Table 5.2. The spectral resolution function is given by 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.

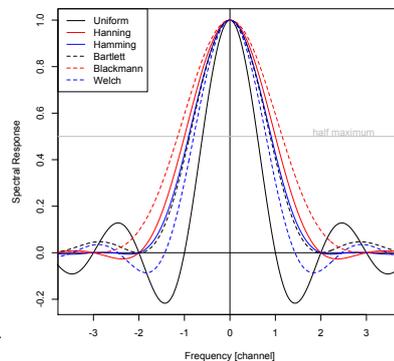
The uniform weighting in the lag domain results in the sinc function with the FWHM of $1.21 \times$ the channel spacing and yields spectral sidelobes that causes 'ringing' in the spectra when a narrow line, interference spike, or strong edge channels are present. Alternative weighting functions¹² are applicable to suppress the spectral sidelobes as listed in Table 5.2. 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 $2.0 \times$ the channel spacing and the maximum sidelobe level of -2.6% . If a different weighting function is required in a specific science goal, the proposer should justify it in the technical case, and select the setup in the Phase 2 SB.

Channel averaging is available to bin or average spectral channels on the CDP. Channels can be averaged

¹²For a full description, see <http://mathworld.wolfram.com/ApodizationFunction.html>.

Weighting	FWHM (ch)	Max sidelobe
Uniform	1.21	-0.217
Hanning (default)	2.00	-0.026
Hamming	1.82	+0.007
Bartlett	1.77	+0.047
Blackmann	2.30	+0.001
Welch	1.59	-0.086

Table 5.2: Spectral resolutions (FWHM) and the maximum sidelobe levels for different weighting functions. Corresponding spectral resolution functions are illustrated in the right plot.



together in factors of $N=2, 4, 8,$ or 16 . The main purpose is to reduce the data rate to the archive and the total data volume. It 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). It might be quite acceptable for those that need something with more resolution than TDM, but where the FDM channels at the full resolution are unnecessary. Table 5.3 shows the resolutions (in kHz) for different values of N , using Hanning weighting, in the different bandwidth modes. The channel spacings are in parentheses. $N=1$ is the default unbinned case, where the resolution is $2\times$ the channel spacing.

Usable bandwidth (MHz)	N = Channels =	Spectral resolution (channel spacing) (kHz)				
		1	2	4	8	16
		3840	1920	960	480	240
1875		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 5.3: Spectral resolution and channel spacing (in brackets) in kHz for different correlator bandwidth modes (left column) and for different channel averaging factors (columns, $N=1$ to 16), using Hanning smoothing. The number of channels can be reduced from 3840 (for the un-averaged case, $N=1$) down to 240 (for $N=16$). Values are given for the 2-polarization case.

Note that the default Hanning window function gives a resolution 2 time 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 that unless the maximum spectral resolution is required by the observations, then the number of channels be reduced when feasible. This is selected in Phase 2 of the SB creation. However, note that this is a non-reversible operation!

The ACA Correlator is an FX correlator where the uniform weighting in the FFT segment is equivalent of the Bartlett weighting in the lag domain. The FPS

In the case of the ACA Correlator, the spectral resolution function in voltage spectrum is given by a sinc function, so that of the (cross) power spectrum is a sinc^2 function, without additional weighting. FPS is used in the CDP in order to match the spectral resolution function to that of the 64-input Correlator. Since FPS is a convolution process with an artificial spectral profile with a limited-length support, this process reduces the bandwidth at both band edges, in the case of 2-GHz SPW, as listed in Table 5.1.

5.5.3 Time Resolution

As shown in Figure 5.7 the correlated data are processed in the CDP and accumulators before archived. The time resolution is limited by the shortest integration period in the correlators and also the maximum data dump rate. The data dump rate, R_d , given by

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

must be lower than the acceptable data rates described below. Here, N_{ant} , N_{BB} , N_{ch} are number of antennas, BBs, and spectral channels (sum of polarization products), respectively. N_{byte} (bytes/visibility) is the byte size of a real or imaginary part of visibility. T_{int} is the integration time in the correlator. For typical instance, a configuration of $N_{\text{ant}} = 32$, $N_{\text{BB}} = 4$, $N_{\text{ch}} = 3840 \times 2$ polarization products (XX and YY), $N_{\text{byte}} = 2$ and $T_{\text{int}} = 6$ sec yields 10.5 Mbytes/sec, that is within the acceptable data rate of the 64-input Correlator. The ALMA Observing Tool (OT) estimates the data dump rate for each science goal. It is required to consider the trade off between time resolution, spectral resolution, polarization, number of BBs, and number of antennas.

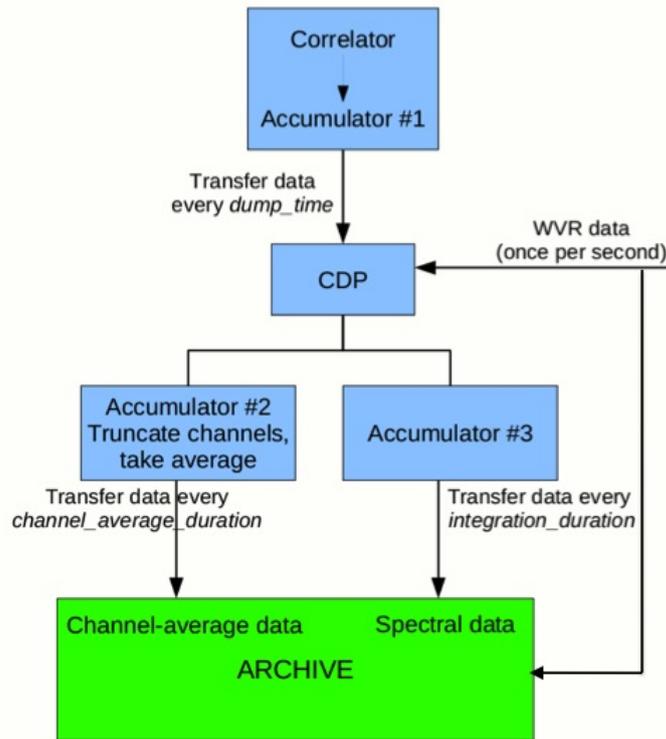


Figure 5.7: Basic data processing and accumulation steps between the correlator and archive. The CDP applies WVR correction before time integration of the spectral data and channel-averaged data. The set of WVR, channel averaged data, and spectral data are packed in Binary Data Format (BDF) and archived.

Pulsar-gating function is not implemented in the ALMA correlators. Timing-sensitive analysis must be taken offline under the time resolution limited by the condition of the integration duration.

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

The correlator dump time is 32 ms and 48 ms for TDM and FDM, respectively, and the integration duration of spectral data in the CDP must be a multiple of the dump time. And the data rate must accommodate the maximum data rate. Although the specification of the maximum data rate from the 64-input Correlator is designed to be 62.5 Mbytes/s, the verified data rate is 26 Mbytes/sec to date. As in Cycle 2, in Cycle 3 the maximum data rate is presumed to be regulated at 17 Mbytes/sec. The 64-input Correlator employs $N_{\text{byte}} = 2$ or 4, depending on a required dynamic range of visibilities.

- Integration duration of spectral data with the ACA Correlator

The ACA Correlator employs $N_{\text{byte}} = 4$ and the dump time 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, independent of the number of antennas. However, the average data rate to the archive will be lower during typical observing, because of overheads and the use of TDM modes during some observing.

- Integration duration of online WVR Correction

When online WVR correction is employed, the integration duration of the 64-input Correlator must be equal or longer than that of the WVR system, which is 1.152 sec. For Cycle 3, neither the 7-m Array nor the TP Array uses online WVR.

- Channel-averaged duration

The average of power spectra across whole spectral channels in each SPW can be recorded in another SPW with a channel-averaged duration that must be multiple of the correlator dump time and must be equal or shorter than the integration duration of the WVR. This single-channel SPW, without bandpass calibration, is not used for astronomy, but for recording time variations of phases and amplitudes.

5.6 Practical Performance

5.6.1 Sensitivity

Quantization of digital signals involves reduction of information and causes sensitivity loss. The ALMA digitizer employs 3-bit (8-level) quantization¹³, and additional re-quantization processes are applied in the correlators.

In Cycle 3, TDM of the 64-input Correlator makes use of 2-bit that yields the quantization efficiency of $\eta_Q = 0.88$. For the case of FDM, digital filtering is applied to 3-bit quantized signals followed by 2-bit re-quantization. The two stages of quantization would cause the multiplied quantization efficiency of $\eta_Q = 0.96 \times 0.88 = 0.85$ if they were independent (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 taken in 16-bit precision. The spectrum of FFT output is re-quantized in 4-bit (16 levels) before cross multiplication. The combination of these quantizations results in loss of 4.8% (i.e. $\eta_Q = 0.952$; 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 level adjustment is taken for every scan sequence, significant variation of power level (e.g. observation of strong sources, unstable weather, or frequency band near atmospheric absorption lines) can violate the condition to reduce the quantization efficiency.

¹³The quantization efficiency would be $\eta_Q = 0.96$ 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 $\eta_Q = 0.81$, better than 0.88×0.88 .

5.6.2 Linearity

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

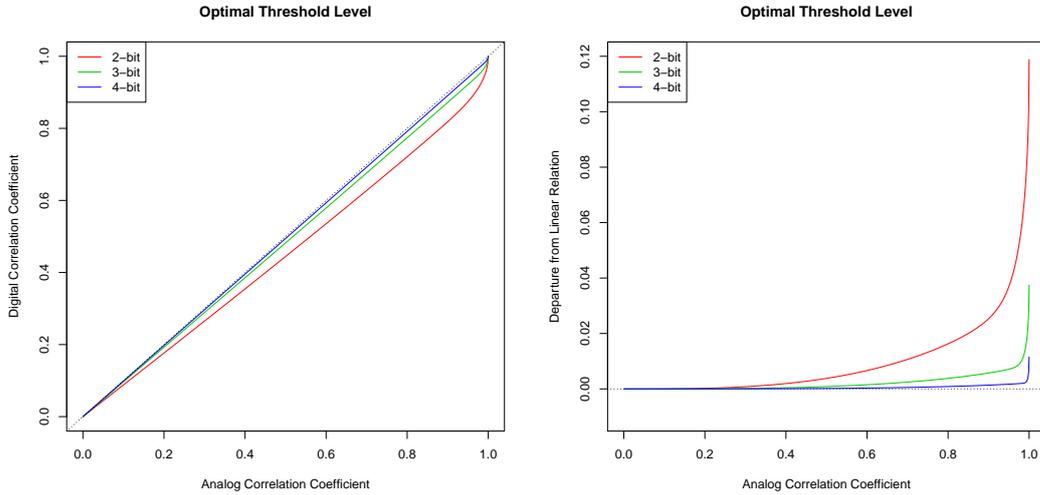


Figure 5.8: Van Vleck relationship of 2-, 3-, and 4-bit quantizations at the optimal signal power level and threshold voltages. (Right): Correlation coefficient of digitized signals as a function of analog correlation coefficients, ρ . (Left): Departures of the Van Vleck relationship from linear relation employing single factor of $\eta_4 = 0.881$, $\eta_8 = 0.963$, and $\eta_{16} = 0.988$ for 2-, 3-, and 4-bit quantizations, respectively.

Figure 5.8 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 measurements and the TP Array observations. In the case of 64-input Correlator, the 2-bit Van Vleck correction will be applied in the CDP or every integration period. For TP Array observations, a single constant correction factor is applied to the spectra to mediate the non-linearity. This subtle difference from the Van Vleck correction and the simple correction yields up to $\pm 5\%$ systematic error in amplitude of spectral lines.

5.6.3 Integration time intervals, channel-average and spectral data

The correlators produce two datasets: the spectral data, and the channel-averaged data. These are written to the archive from the CDP at different rates (see Figure 5.7). The channel-averaged data is a single complex number for each SPW. In the case of FDM, the data is truncated in the CDP before averaging (only the central 15/16 are used). For TDM no truncation is applied. The primary purpose of the channel-averaged data is to provide a smaller dataset for TELCAL, the ALMA online processing software, to compute real-time telescope calibration corrections. It is stored in the ALMA Science Data Model (ASDM) dataset as a separate SPW (see 5.6.6). Note that this should not be used as science continuum data in normal observing, but the continuum should be constructed offline in Common Astronomy Software Applications (CASA) using the appropriate portions of the spectral data. The spectral data contains the requested number of channels for each polarization product. The time intervals involved at this stage of the observing and data acquisition are also shown in Figure 5.7. In detail, they are:

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

Dump duration: The internal time period inside the correlator (16 ms for cross-correlations and multiples of 1 ms for auto-correlations) over which data is accumulated before sending to the CDP. Data can be corrected for atmospheric phase fluctuations using the WVR correction once every dump duration in the CDP, or offline in CASA once every Spectral Integration Time (see below). During observations with the 64-input Correlator in Cycle 3, the dump time should be a multiple of 48 ms in FDM and 32ms in TDM. A dump duration of 480 or 960 ms is typically used, and it should not be more than the dump time of the WVR system. It is always 16ms on the ACA Correlator.

Channel-average duration: This is the time interval between channel average data being written to the ASDM. The channel average time must be a multiple of the dump duration. This should be also kept small (i.e. ≤ 1 second), because this is the dumptime of the WVR system.

Spectral integration duration: The spectral integration duration is the shortest integration time before the full-resolution spectral data is written to the ASDM. It is also the shortest time interval that can be selected (and potentially self-calibrated) in CASA. The spectral integration time must be a multiple of the channel average time. For Cycle 3 observations using FDM, a Spectral Integration time of several seconds is used to avoid data rate problems (see 5.6.5). For TDM observations, which have fewer channels, a faster rate can be used.

Subscan duration: A further time interval, not shown in Figure 5.7, is the subscan duration or time. In an observation, this will effectively be the shortest time interval where no parameter is changed in the system. For example, this could be a single integration on a source, one point of a mosaic, one scan of a raster, integration at one frequency, or on one of the loads, etc. The subscan duration must be a multiple of the spectral integration time. Typically a subscan might be 30 or 60 seconds; the ASDM might therefore contain 5-10 spectral integration times per subscan, resulting in 5-10 rows of data per baseline/polarization/BB product.

Scan time: Total time per scan. Must be a multiple of the subscan duration. This might be as long as 5 minutes or more in the case of an on-source integration in a single field of mosaicing observations.

5.6.4 Online WVR correction

Semi-realtime correction of pathlength fluctuations are done in the CDP using the WVR data from each antenna. The advantage of doing WVR correction in realtime is that it can be performed relatively rapidly (up to the largest value of the dump duration or the WVR chopper wheel rotation and readout period, approximately 1 second), thus enabling corrections to be made for faster variations in atmospheric transmission. At the same time, corrected data can be integrated in the CDP, allowing data to be transferred from the correlator to the archive at a slower rate. Note that once the data are combined, the online WVR correction is non-reversible.

For some time during Cycle 3, it is expected that both the online corrected and the uncorrected data will be written to ASDM once every Spectral Integration Time (although note that this will double the overall data rate; see Section 5.6.5). So, as in previous cycles, it will be possible to do the WVR correction offline on the slower uncorrected data (or even not to do any WVR corrections). But it is expected that the online WVR corrections will track the sky more accurately, and this data will eventually be used directly. The optimal method of WVR correction, potentially including a combination of online and offline corrections, is still under investigation.

7m-antennas do not have WVRs and the intention is to use the WVR correction derived from the TP antennas, although more testing is required before this is fully operational.

5.6.5 Correlator speed and data rates

Although the specification (correct the spelling - specificaiton) of the maximum data rate from the 64-input Correlator is 62.5 Mbytes/s, it has been verified up to 26 Mbytes/sec to date. In Cycle 3 it will again be regulated to about 17 Mbytes/sec; the same as as what was offered in Cycle 2. For the ACA Correlator the

maximum data rate is 3.6 Mbytes/s. These rates set limitations on the shortest value of the spectral integration time in the correlator, particularly in FDM mode, where the number of channels without spectral averaging is 32 times larger than TDM mode. The FDM raw data rate (in Mbytes/sec) is approximately

$$N_A(N_A - 1)/2 \times N_{bb} \times 8192 \text{ channels} \times (4/10^6) \text{ Mbytes}/I_T \quad (5.3)$$

where N_A is the number of antennas, N_{bb} is the number of BBs (normally 4) and I_T the spectral integration time. Note that this rate will double if both WVR online-corrected and uncorrected data is required from the correlator. As in Cycle 1, the usual FDM spectral integration times set in the SBs are typically expected to be 6 seconds; this may be increased if the on-line WVR correction proves to give accurate and reliable path-length corrections. So under most circumstances the maximum data rate limitations should not have any impact on the science that can be done in Cycle 3, as it is expected that these high values will seldom be reached.

Note that the *average* data rate to the archive will be considerably lower than the maximum correlator data rate, because of observing overheads and the fact that some data taking within an SB will be in TDM rather than FDM mode (eg pointing and atmospheric calibration). The maximum average data rate to the archive is currently 6.6Mbyte/s (see Chapter 14).

5.6.6 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 an unique hexadecimal name (eg uid://A002/X2fed6/X3f). The ASDM contains the meta-data (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 it related to the correlator output.

In the ASDM, the binary data are saved into data structures called SPectral Windows (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 spectral windows 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 spectral window with 4 channels around the water line at 183 GHz. In the ASDM, except for the WVR spectral window, each requested spectral window 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 "hannel averaged version with one channel per polarization product (this averaging is done in the correlator). The channel averaged data are used by the on-line telescope calibration system (TELCAL) and for real time diagnostic purposes (QuickLook), and are typically not used downstream in the data reduction. Overall, this can lead to ASDMs with a large number of spectral windows. For example, a typical science observation in FDM mode can have upwards of 25 spectral windows. Luckily every scan/SPW combination has an "intent" associated with it that indicates its purpose (pointing, system temperature, science, etc). This intent can be used in CASA to decode how to utilize each spectral window in the data reduction process.

Chapter 6

Spectral Setups

In this chapter we outline the ALMA local oscillator (LO) system and the various functions that it performs, how to define spectral setups from the user viewpoint and how these are used to set up the LO system, as well as some of the other points pertaining to the frequency setup. In the later sections, we go into more detail about the hardware, outlining the signal path between the frontends and correlators, the LO chain and the LO distribution system, and describing some of the major LO components.

To the ALMA system, a spectral setup effectively consists of the settings of the local oscillators and correlator, 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, spectral window bandwidths and spectral resolutions: 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 user manual and reference manuals, available from the ALMA website¹ (and also in the OT itself).

The following sections describe how to use the LO system. For those interested in the details of the components and how the hardware works, please jump to Appendix B.

6.1 Introduction

ALMA allows up to four 2-GHz wide *Basebands* to be placed within the IF range of the frontends using LO2. Furthermore within each baseband it is possible to place up to 32 *Spectral Windows* (SPWs). Each Spectral Window (SPW) forms a final contiguous spectrum, with bandwidths in FDM mode from 58.6 MHz up to 1.875 GHz wide (see Section sec:ubw). Using the OT during both proposal preparation by the PI (Phase 1) and scheduling block (SB) creation by ALMA staff (Phase 2), the user chooses the frequencies (or transitions) to be observed and the required spectral window bandwidths or spectral resolutions, and the OT will look for the best solution for the settings of the LOs, correlator etc. ALMA uses heterodyne conversion to shift the signals from the frontend 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 frequency f_{sig} are mixed to Intermediate Frequency (IF) signals of frequency f_{IF} using a Local Oscillator (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 the other signals. The output signal may correspond to an input frequency above or below f_{LO} , known as the upper and lower sideband (USB or LSB).

During the spectral tuning, there are effectively four different LOs which are set up by the system:

- LO1 which sets the frontend tuning frequency
- LO2 which positions the basebands within the frontend receiver IF (each baseband uses a different LO2, hence there are four LO2s)

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

- LO3 which is the clock frequency of the digitizers (fixed at 4 GHz)
- LO4 (also known as the tunable filterbank LO, or TFB LO), which is a digital LO synthesised in the correlators allowing positioning of the spectral windows within each baseband. Each SPW has effectively a separate LO4

The OT and the realtime system have a tuning algorithm that attempts to find the best tuning solution for all these LOs based on the required observing frequencies. Because the LOs themselves are generated by combinations of frequency synthesizers, several different tuning solutions may be possible, and the algorithm picks the best one. Alternatively, no solutions may be possible, and the OT will flag this. A more detailed description of the LO operation is given in Section B.2. Cycle 3 has similar restrictions on the spectral setups and tuning as Cycle 2 (see Table 6.2 in section 6.6). The main limits are that the edges of the 2 GHz basebands cannot lie outside the receiver tuning range listed in Table 4.1, and the edges of the (untruncated) spectral windows cannot lie outside the 2 GHz-wide basebands². The settings of the SPWs in each baseband are independent, so the resolution and correlator mode as well as the value of LO4 can be different in each baseband. So for example it is possible to have a 62.5 MHz-wide SPW centered on a particular line in baseband 1, and simultaneously use a broadband time division multiplexing mode (or TDM mode - see correlator chapter) for all the other basebands. Cycle 3 also allows for multiple FDM spectral windows within each baseband, described in Section 6.3.2 below. The spectral setups of the ACA Correlator and the 64-input Correlator are - to the end 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_{SPW}) This is set by the user in the OT Phase 1, and is the central frequency of the SPW 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 average velocity of stars in the solar neighborhood.

Center Frequency (Sky) This is seen in the OT Phase 1, and is the actual central frequency of the SPW in the local ALMA (ie 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 (ie the extra velocity from Earth rotation and orbit) as this is not known until actual runtime. For example, for a source of radial velocity 100km/s in the LSRK frame, the difference between rest and sky frequency in the OT will be the equivalent of 100km/s. At runtime, the extra velocity from the Earth orbit and rotation (up to ± 28 km/s and ± 1 km/s), plus the motion of the LSRK (20.0 km/s toward RA=18 hrs, $\delta = 30^\circ$ for epoch 1900.0) are included.

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 SPW, f_{SPW} , is given in terms of this and Baseband Center Frequency by:

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

where the sign depends on the observing sideband.

²Actually in Cycle 3 the system allows for a small overhang of the spectral windows, to enable differential Doppler corrections to be made at runtime; ie slight differences in the Doppler correction for different SPWs.

Note that in certain circumstances (e.g. with complex tuning setups having multiple SPWs per baseband - see below) it may not be possible to set the requested center frequencies to exactly the center channel in all the SPWs. This is because of limitations in the tuning of the LOs. However, the tuning software will keep these offset to a minimum, and in any case they will be a small fraction of the SPW bandwidth. See section B.2 for more details.

6.3 Spectral Setup: Multiple lines and Continuum

The wide IF bandwidth and tuning ability allows for simultaneous imaging of multiple lines. 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 given). When <4 spectral windows are required for the primary lines, the others can be set up to cover fainter lines or to observe the continuum, potentially in TDM mode. 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 in these SPWs (i.e. 1.875 GHz in FDM, or TDM - see Section 6.5.1). Also the continuum SPWs should cover as much of the IF band as possible, so they should ideally not overlap in frequency. 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. The continuum SPW frequencies need to be setup manually, and one method of doing this is described below.

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 can be selected in the OT Phase 1; these are noted in Table 6.1 and illustrated by the red in Figure 6.1.

If mixed line+continuum operation is desired, for example with a single line in SPW1 observed at frequency f (GHz), the other SPWs (2-4) can be set up for continuum. For 2SB receivers, normally this would have SPW2 in the same sideband (offset by 2 GHz from f) and SPW3 and 4 in the opposite sideband (offset by the frontend center IF frequency, f_{IF}). The continuum SPWs should cover the maximum aggregate unique bandwidth. If sb is the sign of the sideband of SPW1 ($sb = +1$ for USB, $sb = -1$ for LSB), then the four SPWs could be set up as follows:

- SPW1: f [SPW covering the primary line]
- SPW2: $f - 2.0$ (in the same sideband, with a continuum SPW 2 GHz below the primary line)
- SPW3: $f - (2.sb.f_{IF})$ (in the opposite sideband)
- SPW4: $f - (2.sb.f_{IF}) - 2.0$ (in the opposite sideband)

f_{IF} is the frontend IF frequency (6.0 GHz for Band 3, 4, 7 & 8, and normally 8.0 GHz for Band 6, 9 and 10). For DSB receivers, it is common to keep all the SPWs in the same sideband, so SPW3 and SPW4 would be at $f - 4.0$ GHz and $f - 6.0$ GHz. Note that this is an approximate rule - if possible, adjustment of the continuum SPWs (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 Band 8, 9 and 10, and near the water lines around 325 GHz in Band 7 (see Figure 6.1). Also, if contiguous spectral coverage of the continuum is desired, the SPWs should be offset by 1.875 GHz rather than 2.0 GHz as suggested above.

Band	Species/transition	Freq. (GHz)	Sideb.	bandwidth	bb	SPW	Notes
3a	Standard cont.	97.5 ¹	dual	2 GHz/TDM			LO1=97.5
3b *	HCO ⁺ 1-0	89.188	LSB	62.5 MHz	1	1	HCO ⁺ /HCN/H ₂ CO
-	HCN 1-0	88.632	LSB	62.5 MHz	2	1	
-	CH ₃ OH/H ₂ CO	101.293/101.333	USB	125 MHz	3	1	
-	continuum SPW	101.3	USB	2 GHz/TDM	4	1	
4a	Standard cont.	145.0 ¹	dual	2 GHz/TDM			LO1=145.0
4b	CS 3-2	146.969	LSB	125 MHz	1	1	CS/DCO ⁺
-	DCO ⁺ 2-1	144.077	LSB	125 MHz	2	1	
-	SO ₂ 4(2,2)-4(1,3)	146.605	LSB	125 MHz	3	1	
-	H ₂ CO 2(0,2)-1(0,1)	145.603	LSB	125 MHz	4	1	
6a	Standard cont.	233.0 ¹	dual	2 GHz/TDM			LO1=233.0
6b *	¹² CO 2-1	230.538	USB	125 MHz	1	1	J=2-1 CO isotopes
-	continuum SPW	231.6	USB	2 GHz/TDM	2	1	
-	C ¹⁸ O 2-1	219.560	LSB	125 MHz	3	1	
-	¹³ CO 2-1	220.399	LSB	125 MHz	3	2	
-	continuum SPW	218.5	LSB	2GHz/TDM	4	1	
7a	Standard cont.	343.5 ¹	dual	2 GHz/TDM			LO1=343.5
7b *	¹² CO 3-2	345.796	LSB	62.5 MHz	1	1	CO/HCO ⁺ /HCN
-	HC ¹⁵ N	344.200	LSB	62.5 MHz	1	2	
-	HC ¹⁵ N	345.340	LSB	62.5 MHz	1	3	
-	continuum SPW	343.200	LSB	2 GHz/TDM	2	1	
-	HCO ⁺ 4-3	356.734	USB	62.5 MHz	3	1	
-	H ¹⁵ CN 4-3	355.440	USB	62.5 MHz	3	2	
-	HCN 4-3	354.505	USB	125 MHz	4	1	
7c	¹² CO 3-2	345.796	USB	62.5 MHz	1	1	J=3-2 CO + isotope
-	continuum SPW	344.8	USB	2 GHz/TDM	2	1	
-	¹³ CO 3-2	330.588	LSB	62.5 MHz	3	1	
-	continuum SPW	331.6	LSB	2 GHz/TDM	4	1	
8a	Standard cont.	405.0 ¹	dual	2 GHz/TDM			LO1=405.0
8b	CI 3P1-3P0	492.160	USB	125	1	1	CI/13CI
-	CS 10-9	489.751	USB	125 MHz	2	1	
-	continuum SPW	477.5	LSB	2 GHz/TDM	3	1	
-	continuum SPW	479.5	LSB	2 GHz/TDM	4	1	
9a	Standard cont.	679.0	USB	2 GHz/TDM			LO1=671.0
9b	¹² CO 6-5	691.472	USB	500 MHz	1	1	CO/CS
-	CS 14-13	685.436	USB	500 MHz	2	1	
-	H ₂ S	687.303	USB	500 MHz	3	1	
-	C ¹⁷ O 6-5	674.009	LSB	500 MHz	4	1	
10a	Standard cont.	875.0	USB	2 GHz/TDM			LO1=867.0

Table 6.1: Examples of spectral setups possible in Cycle 3. This includes the standard continuum-only setups, 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, so this frequency is not actually covered if both sidebands are used. See Figure 6.1. * Template spectral setup released with Cycle 3 version of the OT.

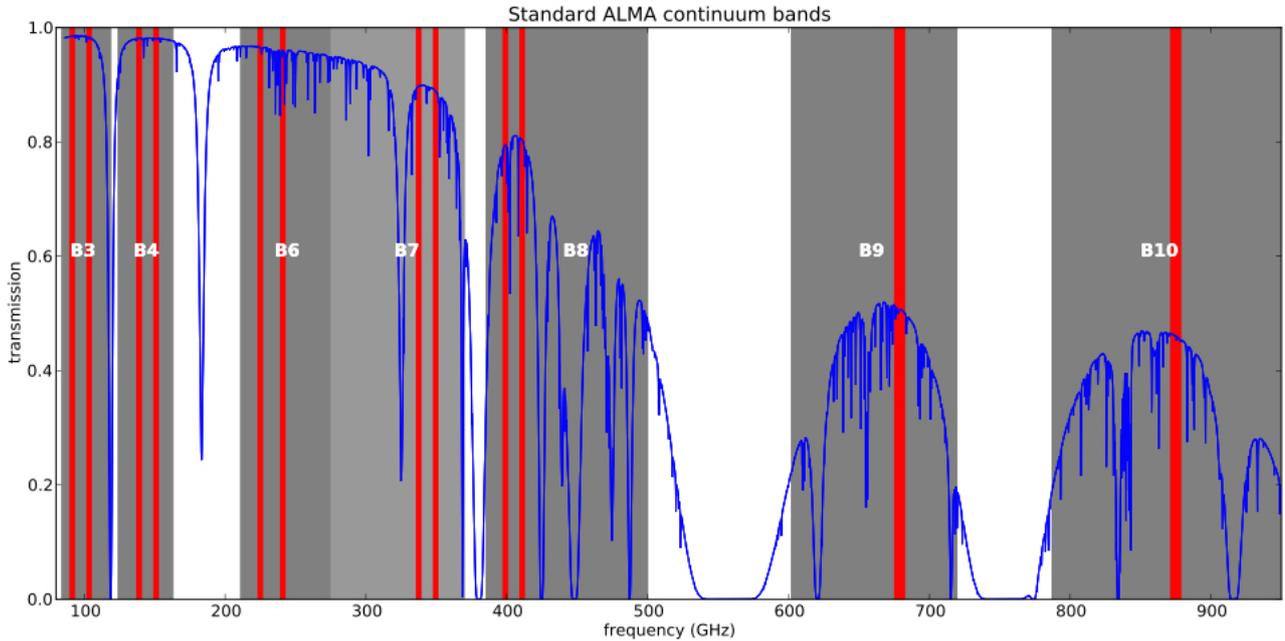


Figure 6.1: ALMA bands available in Cycle 3 (bands 3, 4, 6, 7, 8, 9, & 10), showing the frequency of the standard continuum settings as red shading. This gives the coverage of both USB and LSB, except for Band 9 & 10, which have 8 GHz of bandwidth using the USB only. The available frequency coverage of each band is shown by the grey shading, and atmospheric transmission for 0.6 mm of PWV is given by the blue line.

6.3.2 Multiple Spectral Windows in the Same Baseband

Cycle 3 allows the capability to have up to four SPWs in the same baseband. This is useful for projects where several lines need to be observed simultaneously at high spectral resolution. 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 is often better to place each line in a different baseband. In other 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 SPW, and use the other SPWs in TDM mode for maximizing the bandwidth for phase calibration. There are some restrictions: each SPW within the same baseband must have the *same* resolution³, and the correlator resources used for one baseband (effectively the number of spectral channels, before channel-averaging) will be distributed among the SPWs. An example of this is shown in Table 6.1 for spectral setup 6b, where both $C^{18}O$ and ^{13}CO are observed in the same baseband (3.1 and 3.2). Multiple SPWs are set up in the Phase 1 OT simply by adding more than one spectral line in the same baseband. This will create SPWs with a correlator fraction less than 1.0 (so for 4 SPWs, each has a correlator fraction of 1/4). The number of SPWs may be increased up to 4 in Cycle 3, however, the total number of spectral channels in each baseband is limited by this correlator resources rule. So doubling the number of SPWs in one baseband will result in half the number of channels per SPW, so means either lower resolution or a reduced bandwidth per SPW.

Figure 6.2, which is adapted from the Phase 2 Spectral Editor of the Observing Tool (OT), illustrates a more complex spectral setup with multiple SPWs/baseband which is possible in Cycle 3. The lines observed are given in Table 6.1 example 7b. The blue hashed area represents the receiver tuning range (in this case, a section of Band 7), and the curved line represents the nominal atmospheric transmission for the chosen PWV. The LO1 is 350.0 GHz, and the upper and lower sidebands of the Band 7 receiver are shown as yellow shaded areas. The four basebands, illustrated in this case by the green horizontal bars, can be moved around, but only within the two sidebands. The spectral windows are also shown. In this example, baseband 2 is set for a

³although, as before, SPWs in *different* basebands can have different resolutions

TDM SPW covering the whole 2 GHz baseband for a continuum measurement. In some of the other basebands, multiple spectral windows are used. Baseband 1 has 3 SPWs of 58MHz bandwidth and 960 channels per SPW. At the same time, we are observing two lines in bb3 with 58MHz/1920 channels per SPW and one in bb4 with 125MHz bandwidth and 3840 channels. The basebands (and SPWs) may partially overlap, as seen in this case with bb3 and bb4. Many combinations of SPWs and basebands can be set up in this way.

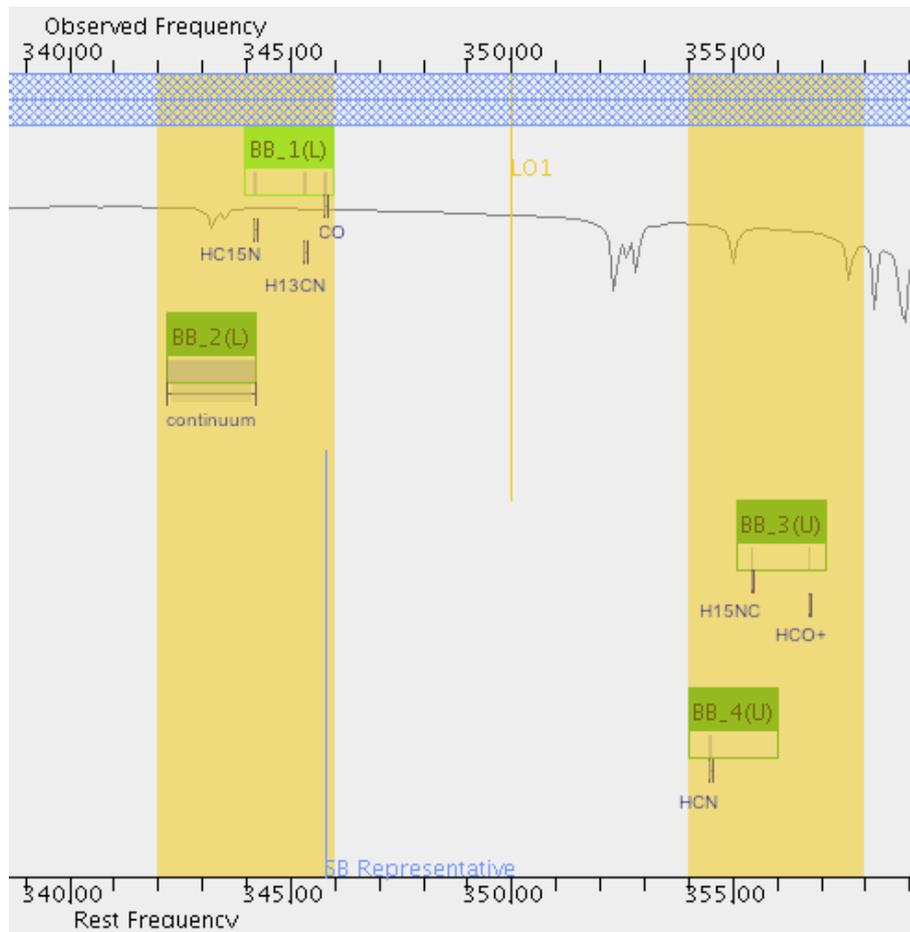


Figure 6.2: Illustration of a frequency setup, based on the OT spectral display. Yellow areas are the IF ranges, green bars are the 2 GHz-wide basebands (bb1-4), and smaller horizontal bars represent the spectral windows (1, 2 and 4 per baseband in this example). The frequency of LO1 is shown by the central vertical line. The blue hashed area shows the possible tuning range of the frontend, and the curved line gives an indication of the atmospheric transmission. The bandwidths of the SPWs are illustrated by the widths of the horizontal bars. The lines observed are given in example 7b in Table 6.1.

6.4 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 4.1). 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 baseband 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 SPW 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 SPW bandwidth (e.g. 1 GHz or less); in this case the SPW will be offset from the center of the baseband, and the line will be at the center of the SPW.

6.5 Spectral Setup in Band 9 & 10 and DSB Considerations

The spectral setup in Band 9,10 is effectively the same as other bands. The only difference is that these use double-sideband (DSB) receivers, so the unwanted (or image) sideband is suppressed using LO offsetting (currently using LO1 and LO2). LO offsetting can only be done in interferometric mode - see section B.4.2. 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 a 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. If this sounds confusing, don't worry - just treat each baseband independently as far as sideband choice is concerned. 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.

6.5.1 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 - see correlator chapters). So for 2 GHz nominal, the correlator outputs a bandwidth of 1.875 GHz. Thus in FDM wideband, the filters do not truncate the spectrum, and the full available correlator bandwidth in wideband 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 6.3). 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 6.3.

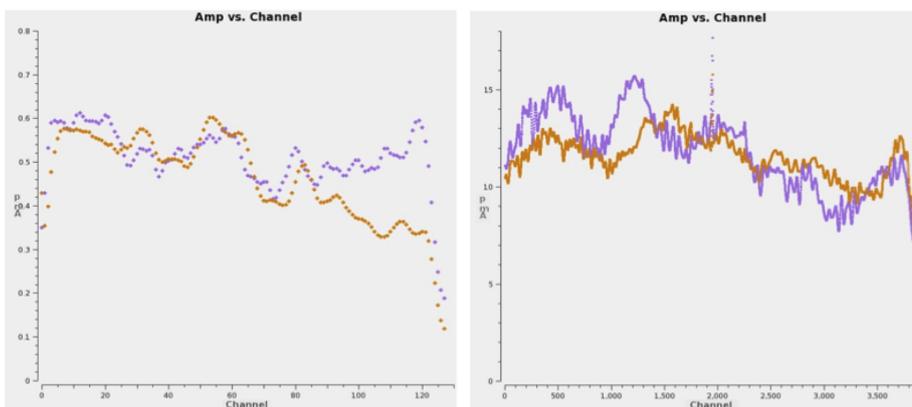


Figure 6.3: Comparison of TDM (left) and FDM (right) autocorrelation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in power in the upper and lower 4 channels due to the anti-aliasing filter. 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).

6.5.2 Spurious Signals

Most spurious signals in the cross-correlation data are suppressed using Walsh-switching (see Section B.4.3). This effectively suppresses signals generated after the frontends, but will not reduce those coming in at the observing frequency. Results from tests with 180-degree phase switching using 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 methods. These are very narrow and will need to be flagged out during data reduction. *Observing lines at these frequencies should be avoided if possible.*

6.5.3 Doppler Setting and Velocity Reference Frames

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

Topocentric In this case *no* correction for the source or Earth motion is made. The Center Frequency - 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 20.0kms^{-1} in the direction $18^h, +30^o$ [B1900.0].

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 and the velocity of the observatory relative to the chosen reference frame are normally used to set the center frequencies of the science SPWs⁴. 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. This frequency setting for the science target will normally also be used in the same observation for 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.

6.6 Limitations and Rules for Spectral Setups in Cycle 3

Although the correlators will eventually allow a broad flexibility of SPWs in a single observation, these abilities are gradually being introduced and tested by the observatory before release. The current rules for spectral setups are given in Table 6.2).

⁴Note that ALMA does *not* do Doppler *tracking*, where the frequency would be continuously updated for Earth motion during the observation. In ALMA it is only *set* once at the start of each execution.

Rule	Details
1.	LO1 must lie within the LO tuning ranges given in Table 4.1.
2.	No part of the 2.0 GHz-wide basebands can extend over the edge of the IF passband. This means that the baseband centers cannot be closer than 1.0 GHz to the IF passband edge. For example, for 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.
3.	For 2SB receivers (Bands 3,4,6,7,8), the number of basebands in one sideband can only be 0, 1, 2 or 4 (i.e. not 3). For DSB receivers (Band 9,10), there is no such restriction (the number can be 0, 1, 2, 3 or 4).
4.	No part of the full nominal bandwidth of the SPW 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 SPW 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 SPW must be at the center of the baseband. However there is a further restriction on this, as noted in the next rule.
5.	The SPW usable bandwidth (ie 15/16 of the nominal SPW bandwidth of multiple of 62.5 MHz) should be in an allowed region of the baseband. This is in addition to (4). In practice this means that the required region of the SPW should normally be inside the range $\sim 2050 - 3950$ MHz (i.e. >50 MHz from the edges of the 2-4GHz second IF) to ensure that the edge of the anti-alias filter does not significantly affect the IF power. In practice it may be possible to extend some part of the SPW to <50 MHz of the IF edges; this is necessary for the Band 6 setup 6b in 6.1, where the ^{12}CO and ^{13}CO lines are only 70 MHz from the IF band edges.
8.	The line frequency does not need to be in the center of the SPW in Phase 1 of the OT: if a line like ^{12}CO 1-0 is requested, it will generate an SB with the correct TFBLO offset. This is mentioned in more detail in Section 6.2
9.	Only 2-bit, Nyquist sampling is allowed in the correlator.
10.	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 spectral windows.
11.	Multiple spectral windows are allowed in a single baseband. However all the SPWs in the same baseband are required to having the same bandwidth and resolution. An individual SPW within a baseband may occupy 1, 1/2 or 1/4 of the resources (i.e. spectral channels) available in the baseband and the sum of all the fractional resources used within one baseband must be ≤ 1 .

Table 6.2: Rules for spectral setups.

Chapter 7

Imaging with ALMA

7.1 Introduction

As described in Chapter 3, the van Cittert-Zernike theorem (e.g. Rohls & Wilson 2004) describes a fundamental relationship between the sky brightness distribution (I), the beam pattern (A) and the visibility distribution \mathcal{V} :

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

An ALMA observation can be represented by N discrete points in the uv -plane $\mathcal{V}_k(u_k, v_k)$ where $k = 0, \dots, N$ and N can reach a few million. In addition to the visibilities \mathcal{V}_k obtained, a weighted visibility function \mathcal{V}^W can be defined (see Briggs, Swab & Sramek, 1999) that enables the resulting synthesized beam size and the sensitivity to be controlled:

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

In Equation 7.2, R_k , T_k and D_k are weights. R_k indicates the sensitivity of the k^{th} visibility measurement based on certain factors, such as the integration time, the system temperature, or the bandwidth. (Note that these factors are given by the instrument, technical set-up, and observing conditions, but not controlled by the imaging process.) T_k can be used to taper the outer edge of the uv -coverage, suppress small-scale sidelobes and increase the beam width. Finally, D_k can be used to offset the high concentration of measurements near the center of the uv -plane and to reduce the sidelobes caused by gaps in the uv -coverage.

For the D_k weight, the following definitions are used commonly during the imaging process:

- $D_k = 1$ is called natural weighting, and is based primarily on the density of visibilities in the uv -plane. Natural weighting yields maximum sensitivity and a relatively large synthesized beam.
- $D_k = \frac{1}{N_s(k)}$ is called uniform weighting, where $N_s(k)$ is the number of visibilities in a region centered on the k^{th} visibility. It removes the dependence of spatial-scale sensitivity on the density of visibilities (samples) in uv -plane. Uniform weighting allows highest angular resolution at the expense of sensitivity.

7.2 Cycle 3 Configurations

In Cycle 3, depending on the range of angular scales required, an ALMA project will obtain data from one or two 12-m Array configurations, and potentially also from the 7-m Array. If yet larger angular scale information

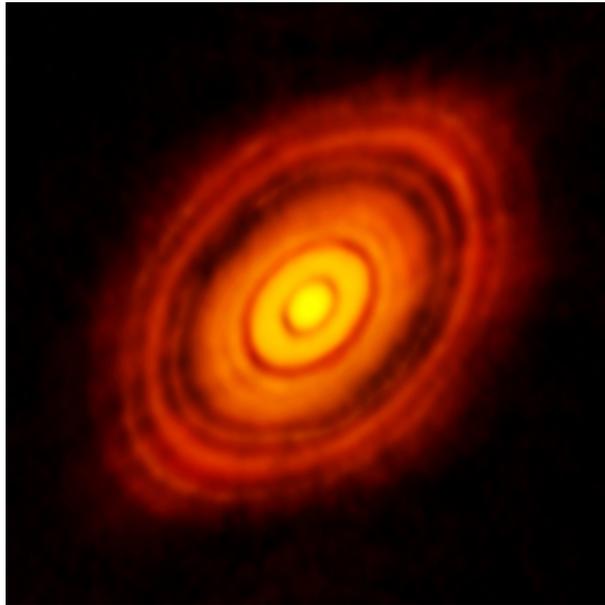


Figure 7.1: This is the sharpest image ever taken by ALMA. It shows the protoplanetary disc surrounding the young star HL Tauri. These new ALMA observations reveal substructures within the disc that have never been seen before and even show the possible positions of planets forming in the dark patches within the system. Credit: ALMA (ESO/NAOJ/NRAO)

is required, single-dish data from the Total Power (TP) Array will be obtained as well when possible (for instance, the TP Array does not currently support continuum science observing). Imaging consists of combining the interferometric data together using an appropriate weighting function; de-convolving the interferometer synthesized beam, point spread function (PSF) or “dirty beam” which results from finite sampling of the uv -plane; and if TP Array data exist, gridding them into an image cube and combining this with the de-convolved interferometer image cube using the feathering technique. We describe in this chapter some of the aspects of imaging that should be considered in order to obtain the best results from ALMA observations.

During Cycle 3, the 12-m Array, with 36 antennas, will be arranged in 8 different configurations. The maximum baselines achievable will be 10 km for Bands 3-6, 5 km for Band 7 and 2 km for Bands 8-10. The 7-m Array, with 10 antennas, will be available in only one configuration. Due to insufficient baseline overlap, the two most extended configurations, C36-7 and C36-8, will be offered alone and not in combination with any other configuration. Tables 7.1 and 7.2 give the basic properties of the 12-m Array and 7-m Array configurations. In addition, Figures 7.2 and 7.3 show antenna locations for the configurations of the 12- and 7-m Arrays. The two TP Array single-dish 12-m antennas will be in fixed positions. **It should be noted that these configurations (especially for the 12-m configurations) and hence the numbers derived for them, are representative of the actual configurations that will be in use** due to a variety of factors including hybrid arrays that exist during antenna reconfiguration periods, precise antenna pad availability, etc... From the users point of view configurations are not directly selected, instead the required angular resolution (θ_{res}) and the largest angular scale (θ_{LAS}) that needs to be recovered are entered in the OT. Projects will then be observed using array configurations chosen to achieve the specified goals. More details about array combination are provided in Section 7.8.

7.3 Shadowing

ALMA is located at latitude=-23.022917°, longitude=-67.754649°. During observations, particularly in very compact configurations like those of the 7-m Array or C36-1, one antenna’s view of the source can be partially blocked by another, making the data from that antenna unusable. This phenomenon is known as shadowing

	Band	3	4	6	7	8	9	10
	Frequency (GHz)	100	150	230	345	460	650	870
Configuration								
7-m	θ_{res} (arcsec)	15.0	10.0	6.5	4.3	3.3	2.3	1.7
	θ_{MRS} (arcsec)	42.8	28.6	18.6	12.4	9.3	6.6	4.9
C36-1	θ_{res} (arcsec)	3.4	2.3	1.5	1.0	0.7	0.5	0.4
	θ_{MRS} (arcsec)	25.3	16.9	11.0	7.3	5.5	3.9	2.9
C36-2	θ_{res} (arcsec)	1.8	1.2	0.8	0.5	0.4	0.3	0.2
	θ_{MRS} (arcsec)	25.2	16.8	11.0	7.3	5.5	3.9	2.9
C36-3	θ_{res} (arcsec)	1.2	0.8	0.5	0.4	0.3	0.2	0.1
	θ_{MRS} (arcsec)	25.2	16.8	10.9	7.3	5.5	3.9	2.9
C36-4	θ_{res} (arcsec)	0.7	0.5	0.3	0.2	0.15	0.1	0.08
	θ_{MRS} (arcsec)	9.6	6.4	4.2	2.8	2.1	1.5	1.1
C36-5	θ_{res} (arcsec)	0.5	0.3	0.2	0.14	0.1	0.07	0.06
	θ_{MRS} (arcsec)	7.8	5.2	3.4	2.2	1.7	1.2	0.9
C36-6	θ_{res} (arcsec)	0.3	0.2	0.1	0.08	0.06	0.04	0.03
	θ_{MRS} (arcsec)	4.8	3.2	2.1	1.4	1.0	0.7	0.5
C36-7	θ_{res} (arcsec)	0.1	0.08	0.05	0.034	-	-	-
	θ_{MRS} (arcsec)	1.5	1.0	0.65	0.43	-	-	-
C36-8	θ_{res} (arcsec)	0.075	0.05	0.03	-	-	-	-
	θ_{MRS} (arcsec)	1.1	0.7	0.5	-	-	-	-

Table 7.1: θ_{res} and θ_{MRS} for the 7-m Array and 12-m Array configurations available during Cycle 3 as a function of a representative frequency in a band. The value of θ_{MRS} is computed using the minimum baseline from Table 7.2 and equation 7.5; the value of θ_{res} is the mean size of the interferometric beam obtained through simulation with CASA, using Briggs uv-plane weighting with $robust=0.5$. (This value of $robust$ offers a compromise between natural and uniform.)

Configuration	7-m	C36-1	C36-2	C36-3	C36-4	C36-5	C36-6	C36-7	C36-8
Minimum baseline (m)	8.7	14.7	14.7	14.7	38.6	47.9	77.3	248.3	346.5
Maximum baseline (m)	32.1	160.7	376.9	538.9	969.4	1396.4	2299.6	6074.2	9743.7

Table 7.2: Basic parameters of the 7-m Array configuration and the eight 12-m Array configurations offered during Cycle 3. The baselines are projected for a transiting source ($HA = \pm 0.5h$) at a declination of -23° . Note that C36-7 will not be available for Bands 8-10, and C36-8 will not be available for Bands 7-10.

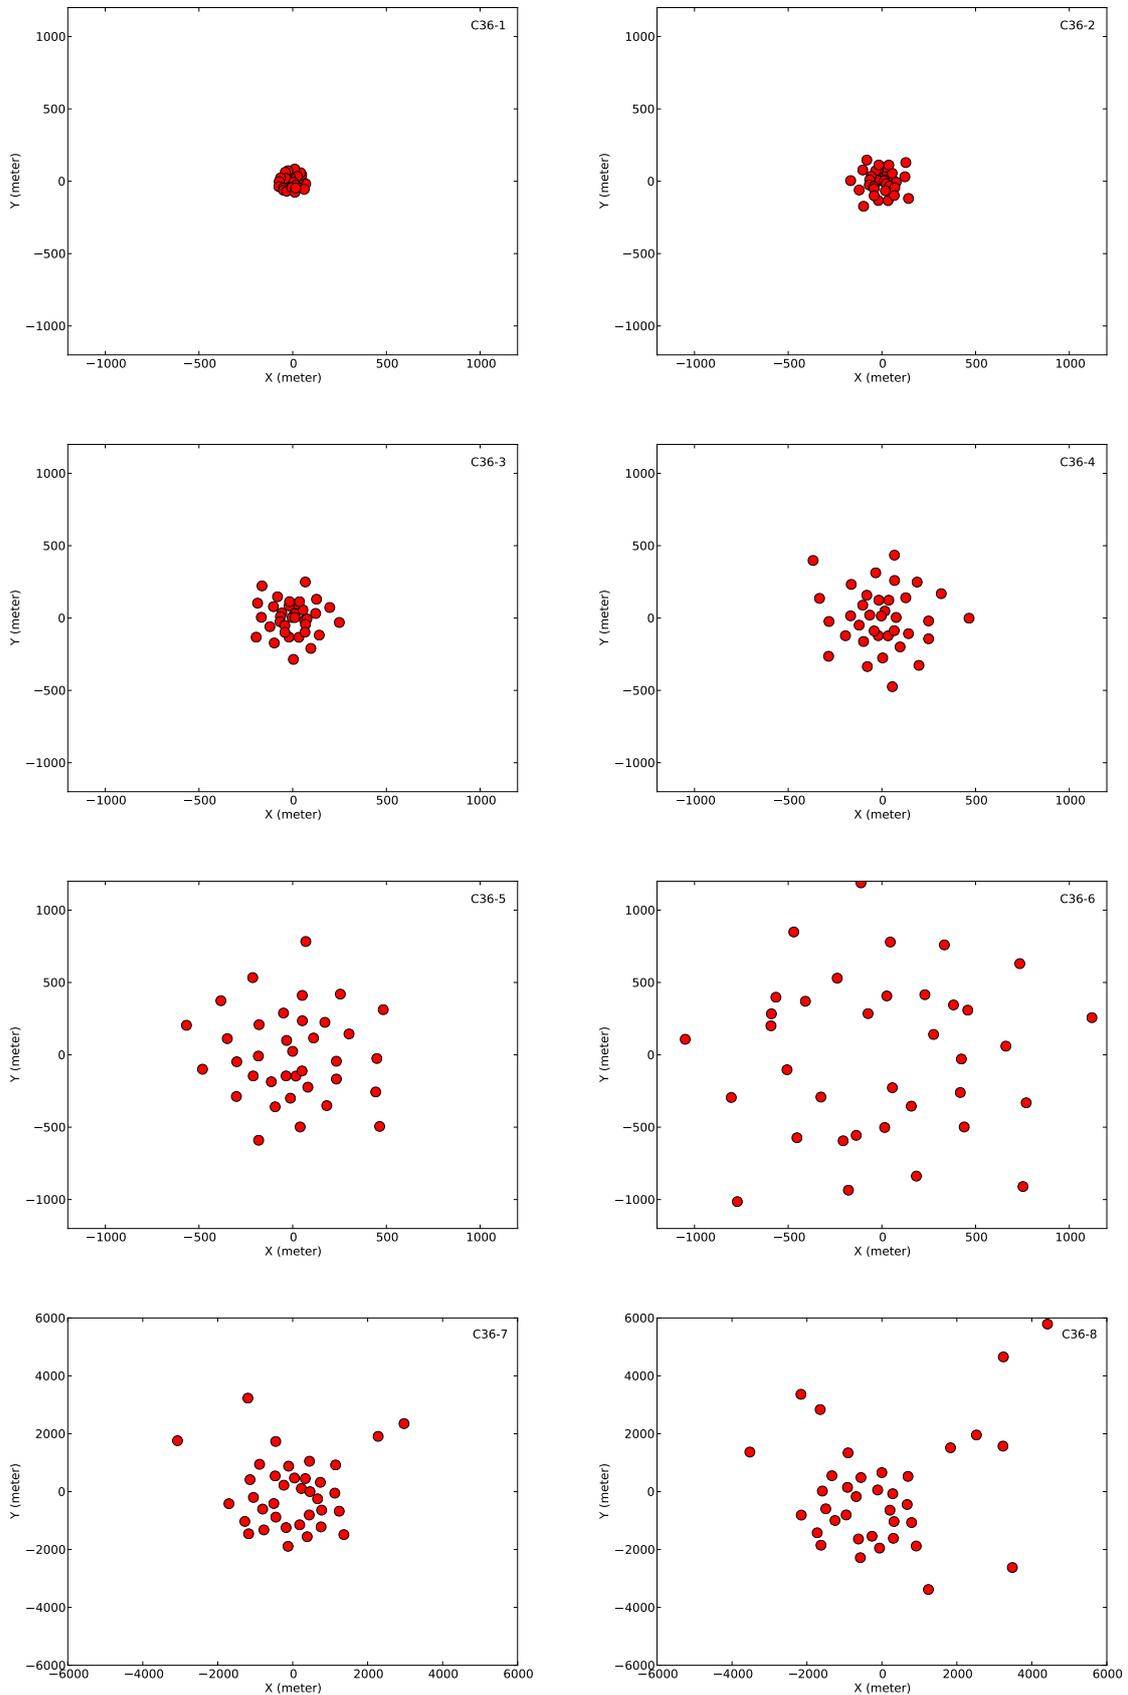


Figure 7.2: Representative 12-m Array configurations for Cycle 3. Note that the scale changes for the final two panels.

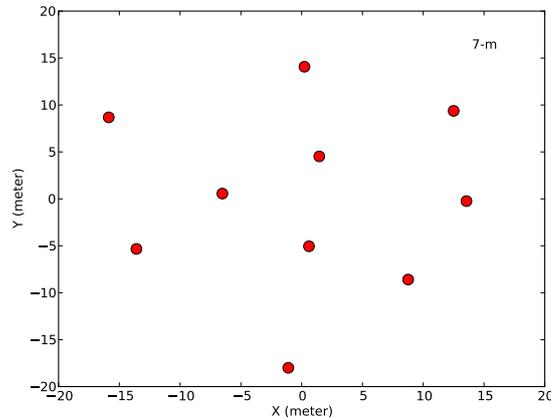


Figure 7.3: Representative 7-m Array configuration for Cycle 3.

and can be quite severe for sources that are observed at very low elevation. Targets as far North as declination $+40^\circ$, corresponding to a maximum source elevation at Chajnantor of $\sim 25^\circ$, can in principle be observed from the ALMA site, but shadowing by adjacent antennas becomes an increasing problem at low elevations. For example, Figure 7.4 shows the percentage of data that will be shadowed (the “shadowing fraction”) when sources of various declinations are observed in the most compact ALMA configurations. As it can be seen, the shadowing fraction can be as large as 40% for sources observed with the most compact 12-m Array configuration (C36-1). The imaging capability, as well as the time on source, will necessarily be limited for such northern sources, especially at the higher frequencies. Shadowing depends on the antenna configuration. Given the short baselines in the ACA configuration, sources with declinations less than 60° or greater than $+20^\circ$ 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 75° or higher than $+25^\circ$.

Note that another issue with observations of sources at low elevations, whether there is shadowing or not, is that the uv coverage, and therefore the synthesized beam, will not be symmetric.

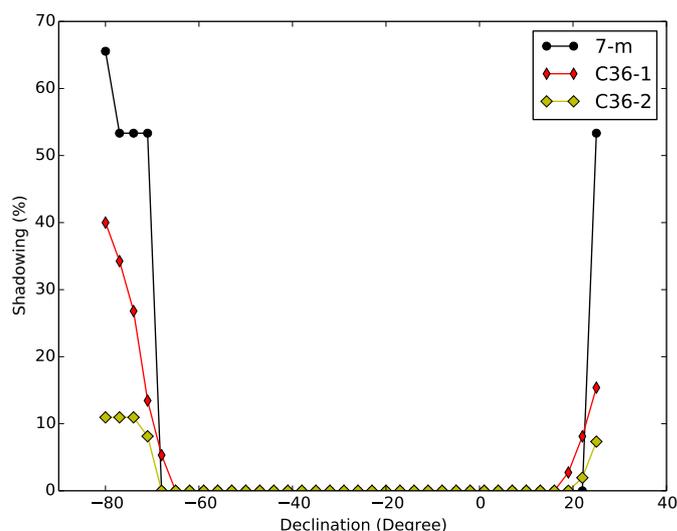


Figure 7.4: Shadowing fraction (%) for the most compact configurations as a function of declination.

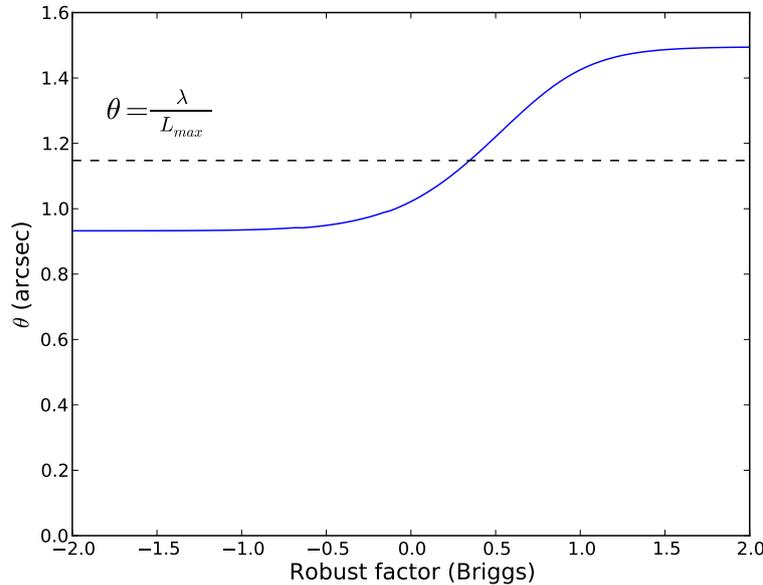


Figure 7.5: Angular resolution achieved using different values of the CASA *robust* parameter for a 1-hour observation at 100 GHz in the C36-3 configuration. Note that *robust* = -2 is close to uniform weighting and *robust* = 2 is close to natural weighting. The dotted line corresponds to $\frac{\lambda}{L_{max}}$.

7.4 Resolution and Beam Shape

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

$$\theta_{res} = \frac{k\lambda}{L_{max}} [\text{radians}] \approx \frac{61800}{L_{max}\nu} [\text{arcseconds}] \quad (7.3)$$

where k is a factor that depends on the uv -plane weighting function, L_{max} is the longest baseline in meters, λ is the observing wavelength in meters, and ν the observing frequency in GHz. An important consideration when imaging is the uv -plane weighting scheme (k) used when Fourier transforming the visibilities. Natural weighting gives the highest-sensitivity image as each visibility is assigned a weight based on its intrinsic uncertainty. Uniform weighting, on the other hand, modifies the weights in such a way that each part of the uv -plane contributes equally to the final image. If this were not done, the shorter baselines would dominate as they have a higher density in the uv -plane. Therefore, uniform weighting degrades the sensitivity (the intrinsic weights are no longer used) but increases the angular resolution as the longer baselines have a higher weight.

In order to bridge the extremes of natural and uniform weighting, Briggs (1995) defined a continuous scheme that uses a “robustness” parameter R (this parameter is called *robust* in CASA). In CASA, uniform weighting is close to *robust* = -2 and natural weighting is close to *robust* = 2. Figure 7.5 shows the angular resolution achieved for an observation with the C36-3 configuration at 100 GHz using *robust* between -2 and 2. As can be seen, the angular resolution varies from 0.9'' (*robust* = -2) to 1.35'' (*robust* = 2). The angular resolutions presented here for ALMA were all computed using CASA simulations with Briggs weighting and *robust* = 0.5.

The synthesized beam shape, which is the Fourier transform of the uv -plane sampling during the observation(s), is a function of the source declination. In addition to shadowing, sources that must be observed at low elevations also have shorter *projected* North-South baselines and thus the shape of the uv -plane sampling distribution becomes more elongated and the beam shape more elliptical. For example, Figure 7.6 shows the different beam shapes for sources observed at declinations of -70° (similar to the SMC and LMC) and -30° . Also, Figure 7.7 shows the uv -coverage for these sources, revealing the elongation of the uv -plane sampling distribution, together

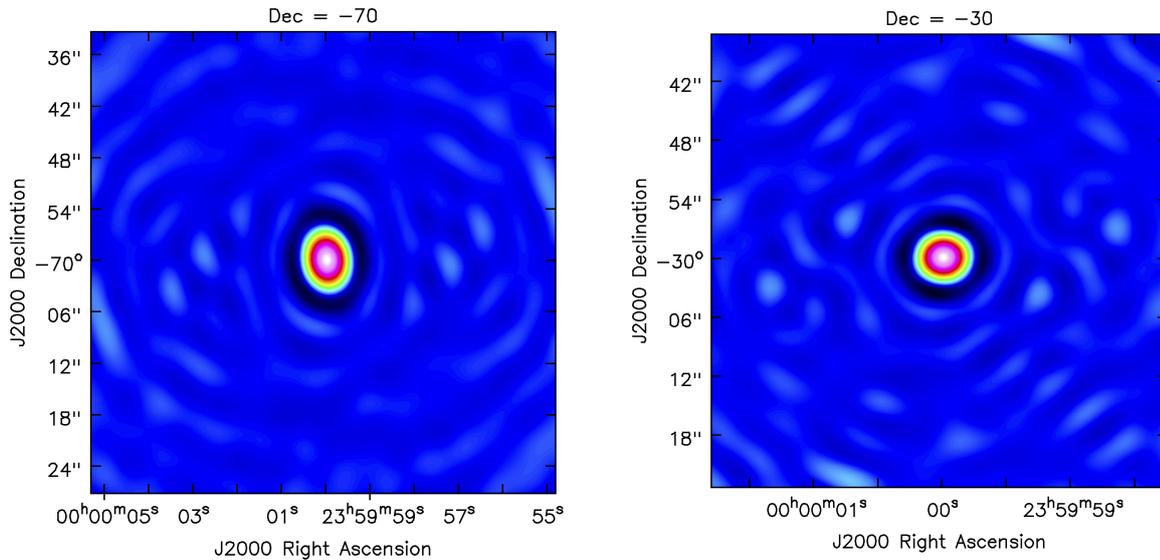


Figure 7.6: Beam shape for configuration C36-1 with a 2-hour observation of a transiting source at a declination of either -70° (*left*) or -30° (*right*).

with the large fraction of shadowing, for the more southern sources. In order to mitigate this effect, one can convolve the resulting images to a more circular (albeit larger) beam after deconvolution.

Figure 7.8 shows the minor and major axes of the synthesized beam widths (θ_{res}) for each array configuration as a function of source declination for a 1-hour observation at 100 GHz. Figure 7.9 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 e.g., C36-8 is not available at bands 7-10.

Note that, in the OT, angular resolution is computed using Briggs weighting with $robust = 0.5$, while sensitivity is computed assuming natural weighting. Thus, the final angular resolution achieved for a project with nominal sensitivity may be slightly degraded by the use of natural weighting. The most appropriate value of $robust$ to use when imaging a dataset must be determined based on the goals of the scientific case (required angular resolution, sensitivity, etc). In cases where more precise information is required, CASA simulations using the representative configurations, actual target declination, uv-coverage, and clean parameters can be used to predict the resolution and sensitivity for a specific science use case. You should clean the images carefully, and you should also add noise (and realistic systematic errors, if possible) to your simulations, to really see what kind of residual sidelobes are expected. (Cleaning perfect data does not show the limitations that one might obtain with real data.)

7.5 Response and Snapshot

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

$$S(u, v) = \sum_{k=1}^N \delta(u - u_k, v - v_k) \quad (7.4)$$

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 centre. However, given the finite number of baselines, the sampling function is an ensemble of Dirac functions and the Fourier transform of this produces a central

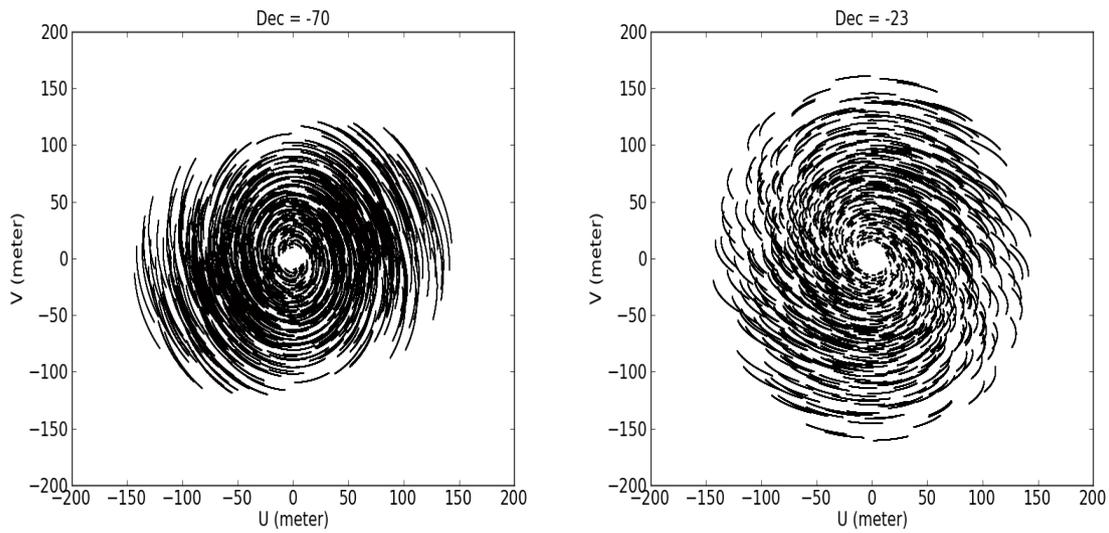


Figure 7.7: uv -plane coverage for configuration C36-1 with a 2-hour observation of a transiting source at a declination of either -70° (left) or -23° (right). For the source with a declination of -70° , the shadowing fraction represents 19% of the total.

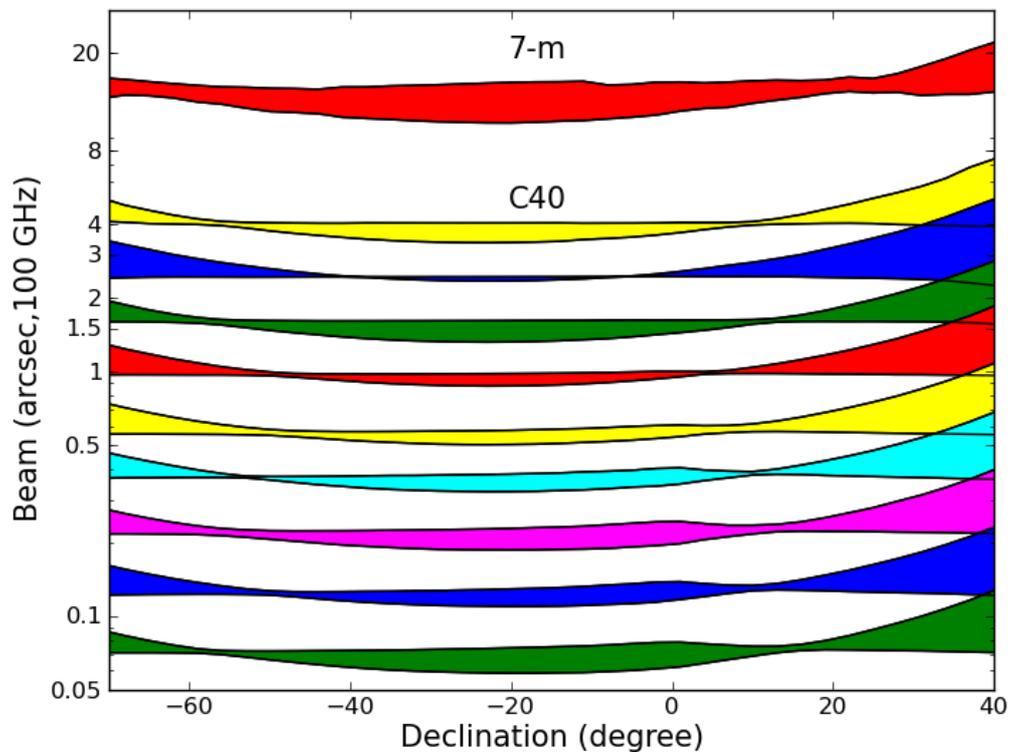


Figure 7.8: Cycle 3 angular resolutions as a function of source declination. Each color shows, for a particular configuration, the range of the major and minor axes of the synthesized beams expected from a 1-hour observation at 100 GHz. Configurations are arranged with C36-1 at the top and C36-8 at the bottom.

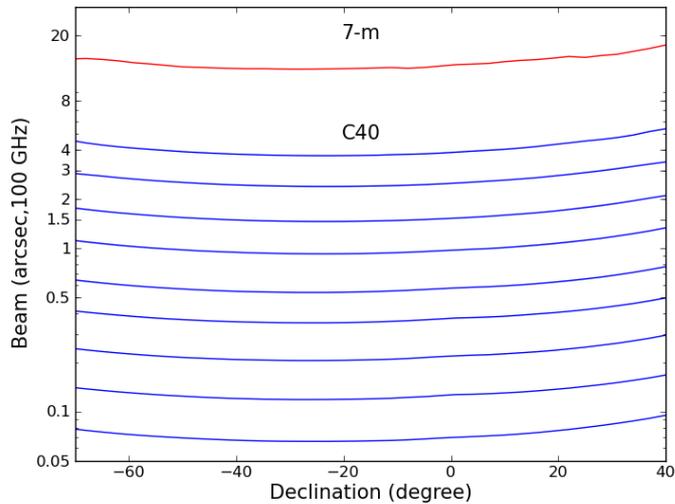


Figure 7.9: Geometrical mean of the major and minor axes of the synthesized beams as a function of source declination. These correspond to a 1-hour observation at 100 GHz. 12-m configurations are arranged with C36-1 at the top and C36-8 at the bottom. The 7-m Array is also shown.

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 (Gibbs phenomenon). As the projected baselines (i.e., the baseline lengths and orientations as seen from the source) change as a function of time as the Earth rotates, the longer a source is observed for, the greater the uv coverage and the smaller the sidelobes. Short integrations are still valuable, particularly when using many of the 50 ALMA antennas, however, if angular structure and sensitivity are less important. Such integrations are sometimes called *snapshots*, and they produce obvious sidelobes which can be mitigated by applying a *uv-taper*, a down-weighting of the visibilities on longer baselines, during the Fourier transform. Table 7.3 gives the level of sidelobes for each configuration, for a 1-hour observation of a source at Dec= -23° .

Configuration	7-m	C36-1	C36-2	C36-3	C36-4	C36-5	C36-6	C36-7	C36-8
Natural	62.6%	11.7%	9.0%	10.1%	9.2%	12.5%	11.0%	13.3%	14.3%
Briggs ($R = 0.5$)	62.7%	9.7%	7.3%	9.7%	9.2%	9.8%	9.4%	9.1%	10.4%
Uniform	62.1%	9.3%	12.2%	13.0%	12.2%	12.0%	10.1%	8.1%	11.4%

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

Figures 7.10 and 7.11 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 expected with the configuration C36-4. As shown in Figure 7.10, 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 7.11) 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 with relatively strong point sources, is to perform CASA simulations using a component list of the strongest sources together with the actual array configuration. This test allows one to estimate the sidelobe fingerprint left by the strong point sources after deconvolution, although the residuals will likely be higher in practice due to imperfect calibration. It is

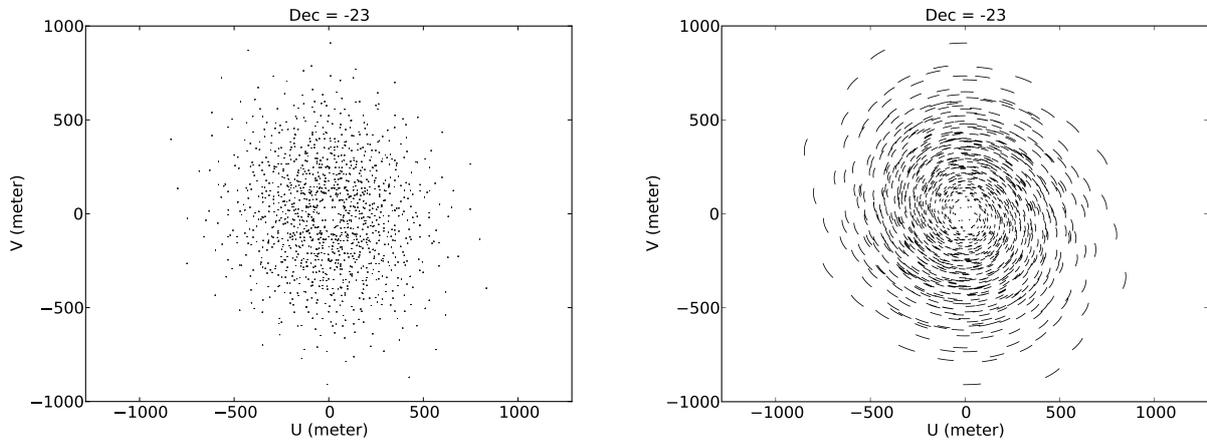


Figure 7.10: uv -plane sampling distributions of a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the C36-4 configuration to observe a source at a declination of -23° .

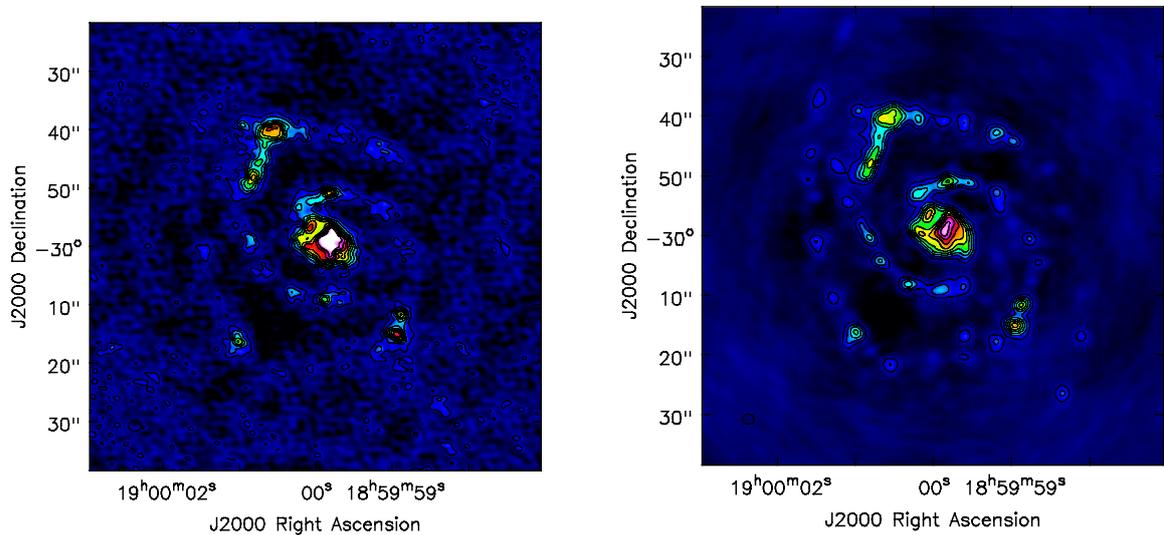


Figure 7.11: Images obtained from a model ALMA observation with 20-s integration (*left*) and 2-hour integration (*right*), using the C36-2 configuration to observe a source at declination of -30° . Black contours at 100, 200, 300, 500, 700 and 900 mJy beam^{-1} are overlaid.

worthwhile to use the interactive mode during the deconvolution so that residuals can be monitored.

7.6 Spatial Scale Filtering

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. Since the uv -plane is a representation of the power of the sky brightness distribution per angular scale, an interferometer observation acts like a passband filter in the spatial domain. An astronomical source is then filtered if its Fourier transform has substantial power on angular scales outside the region of the uv -plane sampled by a given configuration. To illustrate this concept, the Fourier transform of three uniform disks with sizes of $5''$, $10''$ and $20''$ is shown in Figure 7.12 for an observation at 100 GHz. The smallest uniform disk is closest to a point source and so it has large amplitudes up to a baseline of 180 m. Meanwhile, the most extended disk has large amplitudes only up to ~ 40 m. Therefore, an array with baselines larger than 40 m will not be sensitive to emission on angular 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 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. It is to ameliorate the effects of spatial filtering for extended sources that the 7-m and TP arrays are available in addition to the 12-m configurations.

The maximum recoverable scale (θ_{MRS} , see Equation 3.27) is the largest angular structure to which a given array is sensitive; it is determined by the length L_{min} of the shortest baseline in the array. In principle this baseline measures spatial scales with a period of λ/L_{min} . In fact the sensitivity of this measurement to such large structures is not very good, so a smaller value of θ_{MRS} is typically adopted. Indeed, the exact filtering depends on precisely how the structure of the large scale emission maps to the (lack-of) short-baseline uv -coverage and is best determined by simulations. ALMA has adopted a criterion of measuring 10% of the total flux density of a uniform disk, which yields:

$$\theta_{MRS} \approx \frac{0.6 \lambda}{L_{min}} [\text{radians}] \approx \frac{37100}{L_{min}\nu} [\text{arcseconds}] \quad (7.5)$$

where L_{min} is the shortest baseline in meters, λ is the observing wavelength in meters, and ν the observing frequency in GHz. Table 7.1 lists the θ_{MRS} for the various Cycle 3 array configurations. In addition, Figure 7.13 shows the fractions of flux from a uniform disk lost according to angular scale for two array configurations, together with the analytical expectation for θ_{MRS} . The most compact configuration, C36-1, can measure up to 50% of the total flux for angular scales larger than $20''$ at 100 GHz. 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.

The spatial filtering of an interferometer observation is a serious issue that must be considered carefully for each science case. Ideally, the range of spatial frequencies over which the source has interesting structure will be known *a priori* so that the appropriate configuration can be used during observations. Selecting a specific range of angular scales has the effect of discarding all emission from those angular scales not in that range. If very large spatial scales (relative to those measured by the 12-m and 7-m Arrays) are important, the combination of the TP Array, 7-m Array, and 12-m Array should be used, as discussed in Section 7.8.

Figure 7.14 shows the visibility amplitudes *vs.* uv -distance of an example astronomical source, M51. This is taken from an H- α image, but we use it to indicate the flux expected at 100 GHz. The amplitudes of the visibilities can be approximated as a power law of uv -distance with a negative index. This indicates that most of the power is located in larger scale structures and that power decreases rapidly at smaller scales. This result shows that the flux that would be received by the 7-m Array would be much higher than that received by the 12-m Array with extended configurations (e.g., C34-6). The only case where flux is independent of the sampling in the uv -plane is for point sources, which have the same amplitude for all visibilities (i.e., the Fourier transform of a Dirac function). In our M51 example, the visibility amplitude detected by the 7-m Array (at $k_{uv} \approx 5k\lambda$) is

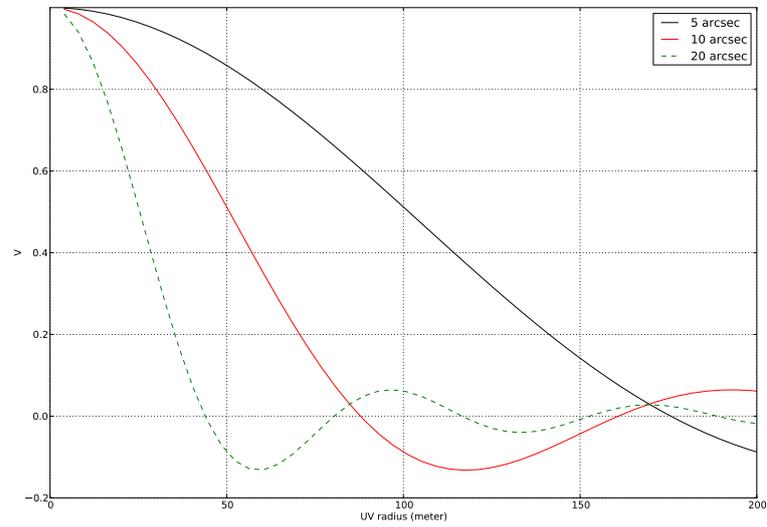


Figure 7.12: Expected visibilities of three model uniform disks (annular averages) at 100 GHz as a function of uv -distance.

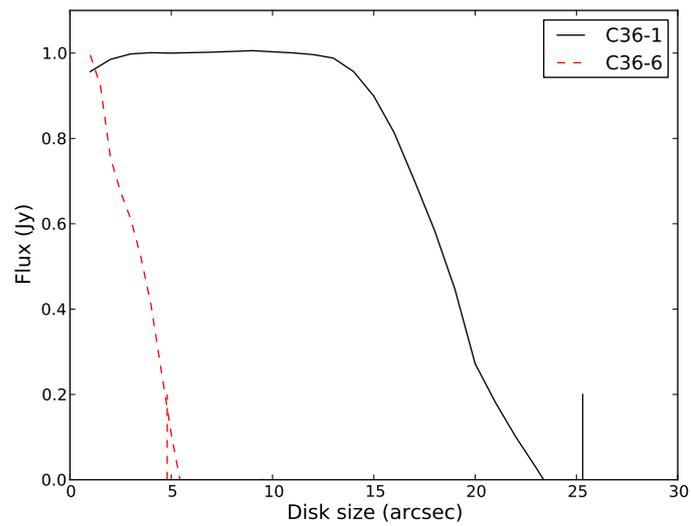


Figure 7.13: Fractions of flux expected from a uniform disk observed in two different array configurations, C36-1 and C36-6, at 100 GHz. The vertical ticks indicate the respective analytical expectation of θ_{MRS} .

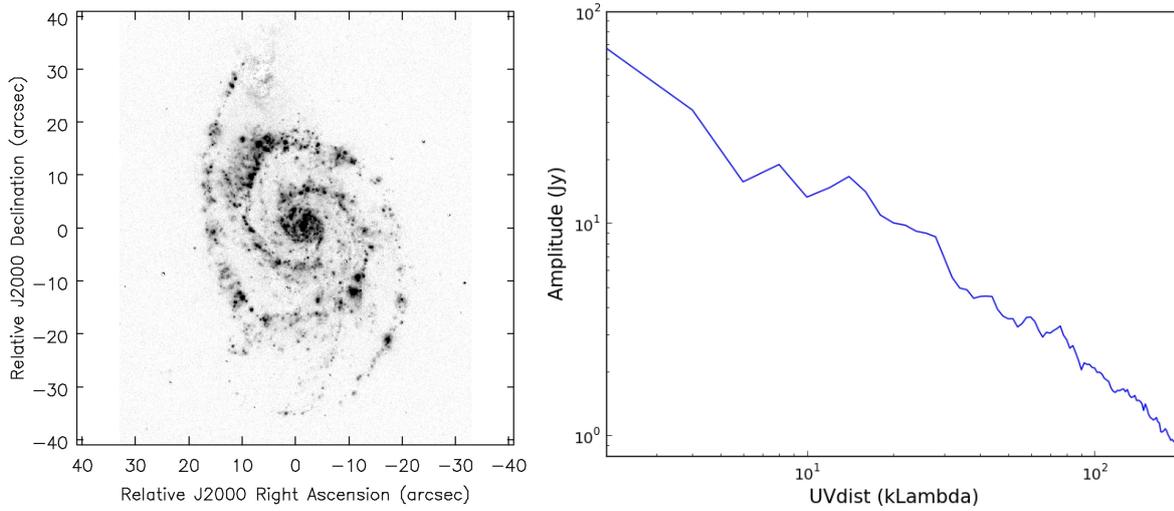


Figure 7.14: Image of $H\alpha$ emission from M51 used as a model of emission at 100 GHz (*left*) and the expected visibility amplitudes with uv -distance (*right*).

~ 23 Jy, whereas that detected by the configuration C36-6 ($k_{uv} \approx 200k\lambda$) is only ~ 0.9 Jy. Ideally, ALMA users should use simulations to estimate the distribution of power at different length scales 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 power ($HPBW$) of $1.02 \frac{\lambda}{D}$. It will also have very high side lobes due to the abrupt truncation of the antenna illumination pattern; 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 -10dB to -15dB response at the edges of the dish. The ALMA feedhorns were designed to illuminate the dish with a -12 dB edge taper in order to provide a nearly Gaussian primary beam with low sidelobes while preserving as much of the resolution and sensitivity as possible. This is one of the fundamental tradeoffs of radio telescope design. The resulting $HPBW$ of the ALMA 12-m antennas is measured to be $HPBW = 1.13 \frac{\lambda}{D} = 58''$ at 100 GHz, and it varies inversely with frequency. For example, Figure 7.15 shows a Gaussian profile which approximates the shape of the primary beam at 112 GHz with $HPBW = 52''$. As described in Chapter 3, the primary beam attenuation can be corrected as the last step of imaging. However, in addition to correcting the signal to its correct value it also increases the noise (i.e., the signal-to-noise is unchanged). For example at the radius of the $HPBW$, the noise will be increased by a factor of two compared to that at the phase center. Moreover, beyond a certain map size, the antenna is not sensitive and it is necessary to observe several adjacent pointings, i.e., a mosaic, to recover the sky emission.

Observing a mosaic with ALMA is needed if a map size larger than approximately the $HPBW$ of the primary beam is required. The default mosaic pointing pattern used by ALMA is a “fully sampled” hexagonal grid with equilateral triangles whose vertices are separated by $\theta_{hex} = \frac{\lambda}{D\sqrt{3}} = 0.511 \times HPBW$ which will 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:

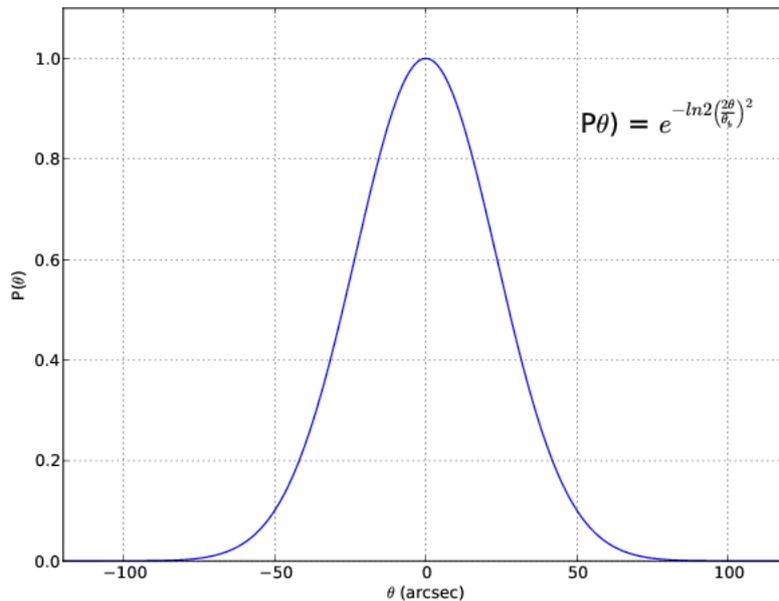


Figure 7.15: Approximation of the primary beam profile of an ALMA 12-m antenna at 112 GHz with HPBW = 52".

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

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 OT. Note that mosaicing with the 7-m Array is more efficient because with its smaller diameter, the HPBW is commensurately larger ($\text{HPBW} = 1.13 \frac{\lambda}{D}$ also holds for the 7-m antennas).

Like a single pointing, a mosaic has an analogous “mosaic primary beam response pattern” that is the convolution of the individual *HPBW* 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 $\text{HPBW}/\sqrt{2}$ (see for example the NVSS survey), this pattern covers area more efficiently but with little gain in sensitivity over a single pointing. Generally the Nyquist sampled hexagonal pattern is a good choice for smaller mosaics and the constant noise hexagonal pattern is most often used for larger mosaics. The OT can be used to easily setup a mosaic of adjacent pointings with a user-defined spacing, though it is not recommended to exceed spacings greater than the constant noise pattern ($\text{HPBW}/\sqrt{2}$) 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} , depending on the maximum baseline; and different maximum recoverable scales, θ_{MRS} , depending on their minimum baselines. In order to achieve the requirements entered by the PI in the OT, θ_{res} and θ_{LAS} , one or more configurations will 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 smallest configuration for a project must deliver a θ_{MRS} that is larger

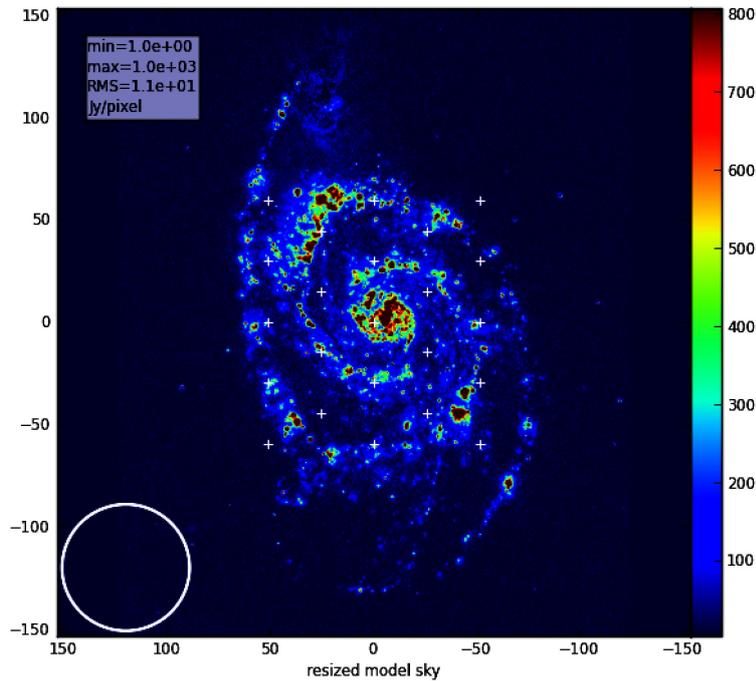


Figure 7.16: An example of mosaicing with a field of 2 arcmin at 100 GHz using an hexagonal pattern with Nyquist sampling (white crosses).

than or equal to the θ_{LAS} of the science target.

The PI should be very careful to set the angular properties $\{\theta_{res}, \theta_{LAS}\}$ of their proposed ALMA observations since they have a strong impact on the way they will be planned and executed. Below is a summary of how the requested ALMA arrays and array configurations are determined from these parameters. Note that in addition to the guidelines below, there are some restrictions. For example, at most two 12-m configurations are allowed and the two longest baseline configurations (C36-7 and C36-8) cannot be combined with smaller configurations for Cycle 3. Additionally, for Cycle 3 the ACA is not available as a stand-alone array and its use should be considered only in combination with the 12-m Array. **Note that the 12-m Array will be available for angular resolutions up to twice that of C36-1.** The complete list of allowed combinations is given in Table 7.4.

1. If $\{\theta_{res}, \theta_{LAS}\}$ is covered by the $\{\theta_{res}, \theta_{MRS}\}$ of a single 12-m Array configuration listed in Table 7.1, only one 12-m Array configuration is needed.
2. If $\{\theta_{res}, \theta_{LAS}\}$ is not covered by the $\{\theta_{res}, \theta_{MRS}\}$ of a single 12-m Array configuration, but is covered by two distinct 12-m configurations, one compact and one extended, then those two 12-m Array configurations will be used provided it is a valid combination described in Table 7.4.
3. If θ_{LAS} is larger than the θ_{MRS} measurable by *any* 12-m Array configuration at the wavelength of interest, then 7-m Array data will also be obtained. Because the 7-m Array is designed to provide good uv -coverage in *combination* with relatively compact configurations of the 12-m Array (C36-1 – C36-3), two possibilities arise:
 - (a) The angular resolution θ_{res} is small enough to be achievable with a compact 12-m Array configuration, overlapping with the 7-m Array in the shorter baselines. In this case, data will be obtained from the 7-m Array and a single configuration of the 12-m Array.
 - (b) The angular resolution θ_{res} is too high to be achievable with a compact 12-m Array configuration. In this case, data will be obtained from the 7-m Array and *two* configurations of the 12-m Array,

“C36-compact” and “C36-extended” (i.e., one of the more compact and one of the more extended 12-m Array configurations). Here C36-extended provides the required angular resolution, and C36-compact provides the sensitivity to intermediate angular scales needed to smoothly match up with the 7-m Array data. Which specific 12-m Array configurations are used for C36-compact and C36-extended depend on the range of $\{\theta_{res}, \theta_{LAS}\}$; allowed combinations are shown in Table 7.4.

4. Finally, if θ_{LAS} is larger than the θ_{MRS} of the 7-m Array, the Total Power Array will also be included.

It is important that you provide an accurate estimate of θ_{LAS} for your target, as the quality of the final image will depend very much on the observation with the proper set of array configurations.

θ_{res} (arcsec)	θ_{LAS} (arcsec)	Array combination	Time ratios	Total Time
0.075	< 1.1	C36-8	1	$1.0 \times \Delta_{extended}$
0.075	> 1.1	-	-	-
0.1	< 1.5	C36-7	1	$1.0 \times \Delta_{extended}$
0.1	> 1.5	-	-	-
0.3	< 4.8	C36-6	1	$1.0 \times \Delta_{extended}$
0.3	4.8-25.2	C36-6 + C36-3	1 : 0.5	$1.5 \times \Delta_{extended}$
0.3	25.2-42.8	C36-6 + C36-3 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.3	> 42.8	C36-6 + C36-3 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
0.5	< 7.8	C36-5	1	$1.0 \times \Delta_{extended}$
0.5	7.8-25.2	C36-5 + C36-2	1 : 0.5	$1.5 \times \Delta_{extended}$
0.5	25.2-42.8	C36-5 + C36-2 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.5	> 42.8	C36-5 + C36-2 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
0.7	< 9.6	C36-4	1	$1.0 \times \Delta_{extended}$
0.7	9.6-25.3	C36-4 + C36-1	1 : 0.5	$1.5 \times \Delta_{extended}$
0.7	25.3-42.8	C36-4 + C36-1 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.7	> 42.8	C36-4 + C36-1 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
1.2	< 25.2	C36-3	1	$1.0 \times \Delta_{extended}$
1.2	25.2-42.8	C36-3 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
1.2	> 42.8	C36-3 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$
1.8	< 25.2	C36-2	1	$1.0 \times \Delta_{extended}$
1.8	25.2-42.8	C36-2 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
1.8	> 42.8	C36-2 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$
3.4	< 25.3	C36-1	1	$1.0 \times \Delta_{extended}$
3.4	25.3-42.8	C36-1 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
3.4	> 42.8	C36-1 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$

Table 7.4: Array combination with the corresponding $\{\theta_{res}, \theta_{LAS}\}$ conditions for an observation at 100 GHz. As in the OT, the angular resolution is computed from the most extended configuration. The actual one obtained with combined configurations can be 50% lower due to different weighting (see text). NOTES: a) for the full array combination, the total time is not equal to the sum of the individual times because TP Array and 7-m Array observations are run in parallel; b) for intermediate values of θ_{res} , please see text. **The largest angular resolution allowed for a project will be twice the resolution offered by C36-1.**

As a user, you will not be able to directly choose the number of 12-m Array configurations to use for your project, it will be decided automatically by the OT, based on the $\{\theta_{res}, \theta_{LAS}\}$ that you have indicated, subject to the guidelines above and the range of valid options described in Table 7.4). The ALMA staff will then schedule the appropriate array configuration(s) to achieve the angular properties requested by the project. As discussed above, the angular resolution depends on the weighting scheme used during imaging. For similar reasons, the combination of data from different configurations can alter the angular resolution and provide a slightly worse one than the most extended configuration taken alone. The quality of the data will be assessed per observation on an array. For the observation with the most extended 12-m configuration, the goal will be to reach the angular resolution θ_{res} you have specified; for the most compact configuration, the goal will be to reach the θ_{LAS} you have specified; for the configurations in-between, if any, the goal will be to reach the

range $\{\theta_{MRS}, \theta_{res}\}$ of that configuration. A detailed description of the quality assurance process can be found in Chapter 11. For the representative Cycle 3 12-m configurations described in Tables 7.1 and 7.2: C36-1, C36-2, and C36-3 are considered compact, C36-4, C36-5 and C36-6 are considered extended (C36-7 and C36-8 are also extended, but will be offered alone, not in combination with any other configuration). The extended configurations cannot be combined directly with the 7-m Array configurations because doing so would lead to large “holes” in the uv -coverage. For cases 3 or 4, the ALMA user will be advised by the OT to add the ACA: it is then the choice of the user to accept this or not; if a user chooses to go against the recommendation by the OT, he/she will need to adequately justify this decision in the proposal.

In addition to the different array combinations available during Cycle 3, Table 7.4 also summarizes the relative integration times that will be used by each array: because of their smaller collecting areas, the TP Array and 7-m Array require more integration time than the 12-m Array in order to match sensitivity at similar spatial scales. A time ratio of 2 between 7-m and 12-m Array will be used in Cycle 3. For TP Array, a ratio of 4 will be used, but the overall increase of time for the project will be less as the TP and 7-m Arrays will be run in parallel. In the end, a typical time ratio for a full combination [TP : 7-m : C36-compact : C36-extended] will be [4 : 2 : 0.5 : 1]. It is very important to note that θ_{LAS} can have a large impact on a proposal, since it can multiply by a factor of 5.0 the total time needed for the project with respect to the 12-m Array due to the longer integration times the TP Array and 7-m Array require. Thus a strong scientific justification for θ_{LAS} should be provided when multiple configurations are required. An additional consideration when invoking the 7-m Array (with only 10 antennas) is that it provides limited uv coverage for snapshot observations. For Cycle 3, it is recommended that at least one hour of 7-m Array hour angle coverage is necessary to provide sufficient uv -coverage for good image combination with the 12-m Array, and snapshot observations are strongly discouraged. A comprehensive plot about the array combination is shown in the Figure 15 of the ALMA Primer for Cycle 3 (Moriarty-Schieven et al., 2015).

Table 7.4 indicates the possible array combinations, assuming the angular resolution of each representative ALMA configuration. However, it takes time to move antennas, which means that the transition from one configuration to the next is progressive, thus providing intermediate angular resolutions. For the purpose of determining the total time required by a project, the OT will assume only the set of 8 array configurations described in Tables 7.1 and 7.2, using for the most extended configuration the one that provides an angular resolution just below the θ_{res} specified by the PI, and if necessary, for the most compact configuration the one that provides a θ_{MRS} just above the θ_{LAS} specified by the PI (within the allowed combinations). Using this information, the project will be broken down into one or more observations, each covering a different, but complementary, range of angular scales. Later during scheduling, the range of angular scales available to the current 12-m configuration will be compared to those required by the project and if they match that portion will become queued for observation (assuming that other important parameters, such as the weather, the project ranking, the executive balance, are appropriate). Note that, as mentioned before, the values of θ_{MRS} and θ_{res} for a particular antenna array are intrinsically linked to the weighting scheme used at the imaging stage; in particular, Figure 7.5 shows that the final angular resolution of an image can be tweaked continuously thanks to the Briggs weighting scheme: this gives some flexibility at the scheduling stage, as it allows running observations whose angular scale requirements, though slightly different, are close to those covered by the antenna array available at the time.

7.9 Multi-array and Multi-configuration Imaging

If your project requires observations with different arrays and/or configurations, we will process the data for each of them, and deliver it to you if it passes the quality assurance. Then, when all the observations have been obtained, we will combine them to make the final image, and deliver it to you, again, if it passes the quality assurance.

The combination procedure will include two steps: first, the imaging of all interferometric data together, then the actual combination of the single-dish map with the interferometric map. The imaging of all interferometric data together can now be done directly with clean in CASA, as of the current release (4.3.0), **assuming the entire calibration was also done with that version (or more recent)**. In Cycle 3, the combination

of the single-dish and interferometric maps will be done with the feathering technique. The term “feathering” is used in radio imaging to describe how to combine or “feather” two images together by forming a weighted sum of their Fourier transforms in the (gridded) uv -plane. Intermediate size scales are down-weighted to give interferometer resolution while preserving single-dish total flux density.

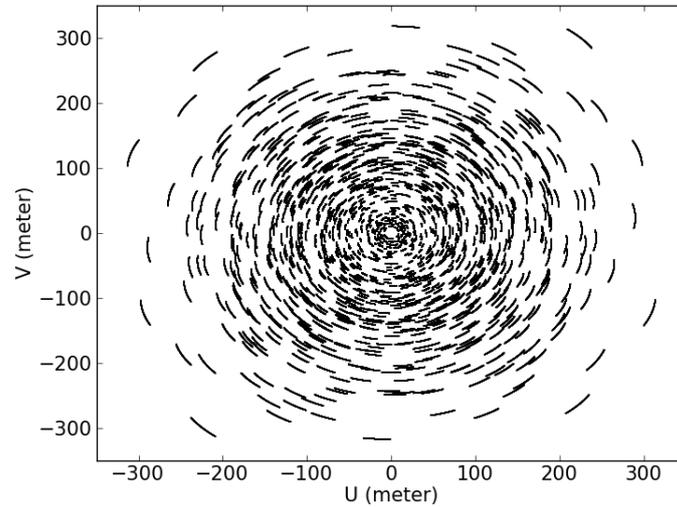


Figure 7.17: Expected uv -coverage for C36-2 (1 hour) and 7-m observations (2 hours)

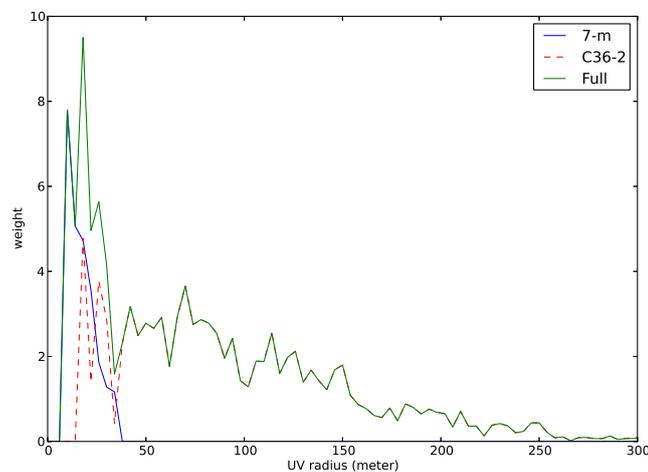


Figure 7.18: Radial density of uv -coverage with uv -distance using C36-2 (1 hour) and 7-m observations (2 hours)

The feathering technique does the following:

1. The single-dish and interferometer images are Fourier transformed.
2. The beam from the single-dish image is Fourier transformed ($FTSDB(u, v)$), (alternatively, one can specify some smaller portion of the single dish aperture, which corresponds to a wider beam).
3. The Fourier transform of the interferometer image is multiplied by $(1 - FTSDB(u, v))$. This basically down weights the shorter spacing data from the interferometer image.

4. The Fourier transform of the single-dish image is scaled by the volume ratio of the interferometer restoring beam to the single dish beam.
5. The results from 3 and 4 are added and Fourier transformed back to the image plane.

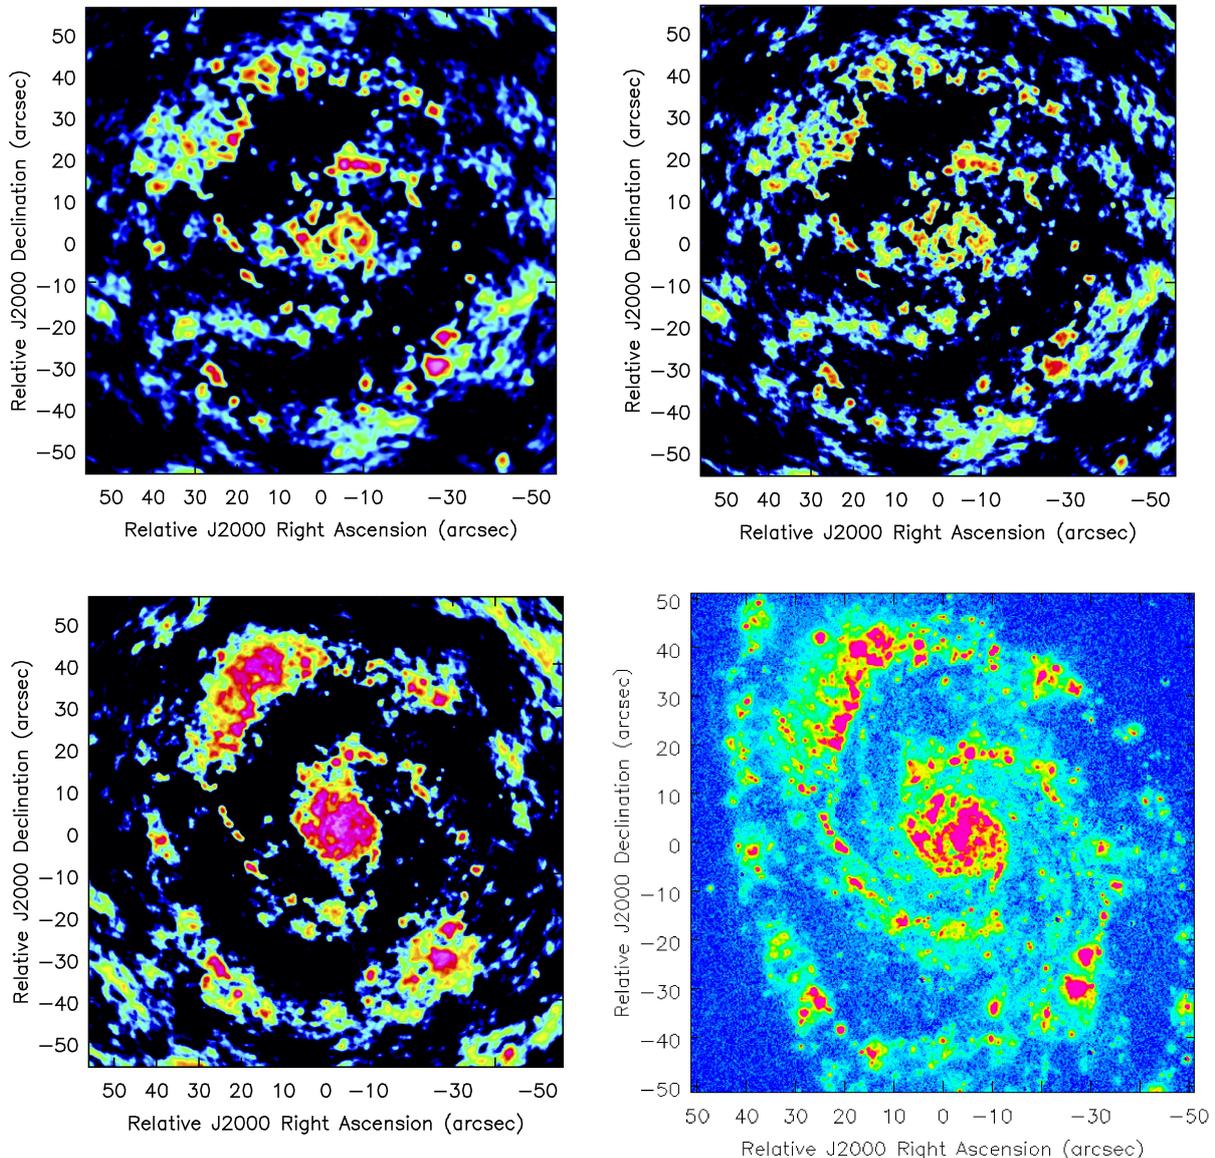


Figure 7.19: Images obtained using C36-3 (*top left*; 1 hour), C36-3 + C36-6 (*top right*; 1 + 2 hours) and C36-3 + C36-6 + 7-m (*bottom left*; 1 + 2 + 4 hours) array combinations, and the image model itself (*bottom right*).

A 1-hour observation with the C36-2 configuration and 2 hours with the 7-m Array is shown in Figure 7.17 (*uv*-coverage) and Figure 7.18 (*uv*-plane sampling distribution). The Cycle 3 7-m Array configuration with ten 7-m antennas provides *uv* measurements in the range 8-32 m and takes over the 12-m Array configurations in the range $R_{(u,v)} < 15$ m. Using the H α emission in M51 as a model, and given the spread of its emission over different length scales (see 7.14), Figure 7.19 shows the images resulting from observing with a C36-compact configuration (C36-3), a compact and extended configuration (C36-3 + C36-6), and adding the 7-m Array. Natural weighting was used and deconvolution was performed using the CLEAN algorithm. The recovery of larger scales is quite noticeable with the use of the 7-m Array. Using only the C36-6 configuration at 100

GHz provides an angular resolution of $0.27''$. Combined with the C36-3 configuration, the angular resolution is lowered to $0.31''$. Finally, the angular resolution with the full combination (C36-3 + C36-6 + 7-m) is $0.32''$, very similar to that of only the two 12-m Array configurations but 17% larger than the angular resolution obtained with the extended 12-m Array alone. That difference is due to the respective uv -coverage and weight of the array configurations in the combined dataset. Changing the weighting scheme (e.g. to uniform) can restore the highest angular resolution at the expense of sensitivity.

References

- [1] Briggs, D. S. 1995, *High Fidelity Deconvolution of Moderately Resolved Sources* Ph. D. thesis, New Mexico Institute of Mining and Technology
- [2] Briggs, D. S., Schwab, F.R., Sramek, R. A. 1999, *Synthesis Imaging in Radio Astronomy II*, ASP Conference Series, 180, 127
- [3] 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
- [4] Moriarty-Schieven, G. ed., 2013, *Observing with ALMA: A Primer for Early Science*, ALMA Doc. 2.1, ver. 1
- [5] Rohlfs, K, Wilson, T. L. 2004. *Tools of Radio Astronomy*, Fourth Edition, ISBN 3-540-40387-6 Springer

Chapter 8

Observing Modes

The ALMA Observing modes are the set of capabilities that ALMA offers each cycle to its user community. For example, single field interferometry, single dish observations, polarization, etc. are observing modes. In order to execute an approved ALMA observing project, the particular observing mode needs to be implemented in the ALMA online software package and verified using the required hardware¹. Here, we describe the ALMA observing modes along with the project structure and a brief description of the software.

8.1 Observing Project Structure

An observing proposal created using the ALMA Observing Tool (hereafter OT) and 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 the successful completion of a project. A summary view of the constituents of an Observing Project is shown in Figure 8.1.

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 SB). They are the minimum set of instructions that describe an ALMA observation. The SBs are produced by the OT which is the tool used during the Phase 1² stage of the project life cycle. The SB may be edited using the OT in Phase 2 after proposal acceptance and before/during project execution³. The SB is the fundamental unit of the observing project and allows for flexibility given the properties of the ALMA site and the continuously varying status of the observatory as a whole (including the weather). It also encapsulates the scientific objectives of the proposal into a single entity.

An SB contains a large amount of information about what should be observed, such as the positions and velocities of the science targets, details of the correlator setup, the integration and cycle times of the different calibrations, etc. 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 observing mode 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 may be performed (e.g. measurements of the system temperatures).

Each SB has an execution time of typically one hour and all the set-ups, the calibrations, and target observations are executed within this period. 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 is atomic in the sense that it cannot be stopped and re-started from the last scan. Therefore, an

¹The implementation and verification of ALMA observing modes was part of the Commissioning and Science verification (CSV) phase as well as also part of the Extension and Optimization of Capabilities (EOC)

²For a description of Phase 1 and Phase 2 see the Cycle 3 Proposer's Guide

³SBs can also be created in an ad-hoc manner for commissioning purposes

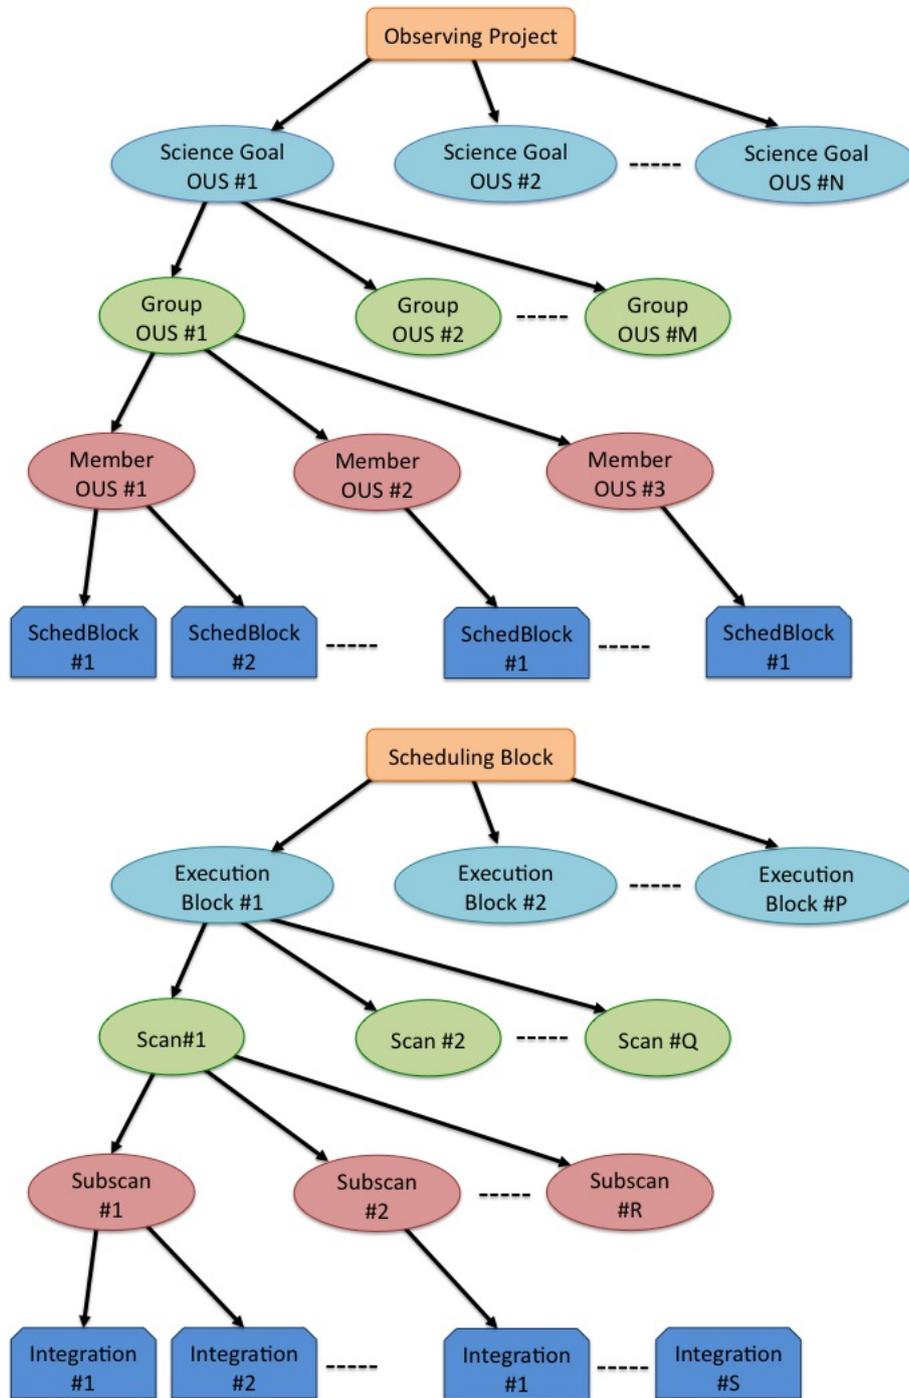


Figure 8.1: Block diagram of an Observing Project from the point of view of the observation preparation as done in the OT (top) and the internal hierarchical structure of the SB in the actual executions (bottom). All projects have the same ObsUnitSet (OUS) levels, the Science Goal OUS level, the Group OUS level and the Member OUS level. Scheduling blocks are attached to the Member OUS.

SB runs to completion, fails, or is terminated by the *Astronomer on Duty* (or AoD). Given the limited duration of the SBs, it is often necessary to observe it several times to achieve the required sensitivity. The current maximum time (or end execution) for a particular SB, has been chosen based on statistical measurements of the stability of the system. A system failure will prevent the online software from flushing the data into the archive losing the observed project data. Thus, executing SBs for more than 2 hours becomes unsafe from the operational perspective. To optimize this scheme, we implemented what we call a *session*, which is the continuous execution of a single SB until a certain goal has been reached (see detailed description later on). For this cycle, sessions are primarily used for polarization observations. If the SB is the smallest entity used for observing, the *Observing Unit Set* (hereafter OUS) is the smallest unit for data processing.

An OUS comprises all SBs that have to be processed together to produce calibrated science data products (e.g. images or data cubes). In most cases, there will be exactly one SB in each Member OUS, but the latter may hold multiple executions of this SB. The actual execution of an SB is called *Execution Block* or EB. For polarization observations, all the EBs that belong to the same session should be grouped together for data processing, which for this cycle will be done manually by the ALMA staff (i.e. no pipeline processing).

If the calibrated science observations of a Member OUS have to be combined with science observations of another Member OUS, then these are grouped together into one Group OUS. A typical case of several Member OUSes that belong to the same Group OUS are observations of the 12-m Array, the 7-m Array and the Total Power (TP) Array. According to the method above, these observations would be in three separate Member OUSes, but in the same Group OUS as one final image will be produced.

Pipeline processing happens at the Member OUS level, as soon as all observations of a Member OUS are completed. In addition, processing happens also at the Group OUS level, in case the Group OUS contains several Member OUSes which have already been processed. The only other event that triggers data reduction is the end of an observing season, when all the Member OUSes and potentially Group OUSes are reduced, regardless of their degree of completion.

The SB/Member OUS/Group OUS are the smallest structures that hold science observations that need to be observed/processed/combined together. This data hierarchy therefore maximizes observing flexibility and ensures that data get 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 Group OUS is attached to a Science Goal OUS. Each Science Goal contains only one Science Goal OUS.

8.2 Program Execution

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, particularly by the *Science Software Requirements* subsystem (hereafter SSR). An overview of the SSR, including communication with other software subsystems, is shown in Figure 8.2. The SSR subsystem commands the lower level CONTROL software to create an EB object that is attached to the SB. As many SBs will be executed several times, a number of EBs may exist for a given SB. Each EB contains a record of the parameters and conditions under which SB was executed along with references to the acquired data. The archived data identified by each EB uses the *Alma Science Data Model* (ASDM)⁴. The internal hierarchical structure of the EB is also shown in Figure 8.1. The SSR subsystem constructs a sequential series of scans for each of the required calibrations and commands CONTROL to execute them. 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 integrations. 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 points of pointing calibration constitutes a scan with 5 subscans). Similarly, the integration time on a single science source between phase calibrations can consist 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. A typical example of a scan sequence is as follows:

⁴The ASDM uses the relational databases paradigm to define the tables containing the data and metadata along with the relationships between them

[CALIBRATE_SIDEHAND_RATIO, CALIBRATE_ATMOSPHERE, CALIBRATE_BANDPASS]

where each of the items in the list corresponds to ALMA calibration intents⁵. Scans and scan sequences can be of arbitrary length, depending on the characteristics of a given observation, but subscans are recommended to be 30 seconds or less. Integrations tend to be on the order of one to 10 seconds where the final value has to be an integer multiple of the correlator dump time. These values can be specified during Phase 2 SB construction in the SB target parameters section. Calibration results from the *Telescope Calibration* subsystem (TelCal) and *QuickLook* (QL) pipeline are usually attached to a scan where one example is an antenna pointing result. The SSR and the Control subsystem are responsible for the creation of all the metadata needed downstream for data processing.

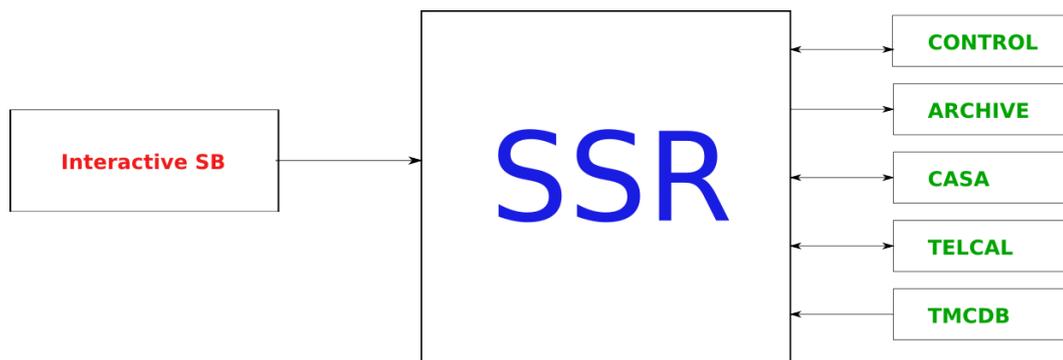


Figure 8.2: Overview of the SSR in term of the ALMA software components that the subsystem talks to.

8.3 The Observing Process

The execution of an ALMA observation is done by means of an observing script. The script will parse the SB and, for a given project, the set of targets specified in the SB are created. These targets can correspond to either calibration or science executions. Targets are organized in observing groups⁶ where the first group (group 1) is always the initial calibration group, with subsequent groups detailing the science observations⁷. The associated calibrations are attached to a particular calibration right after the SB is parsed and the groups created (see below). All Science targets and relevant calibrators within a group are observed before the next group is started. A list of observing scripts is shown in Table 8.1.

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 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 are observed.

Most observing modes, including single field interferometry, multi-target executions, pointed mosaics, and polarization use the same ALMA observing script, called the *standard interferometry* script. The set of necessary calibration measurements (e.g. amplitude, bandpass, etc) usually specified in group 1 are performed at the beginning of the observation sequence if appropriate astronomical sources are available. 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.

Calibrations such as pointing, atmospheric, and sideband-ratio are associated to the main calibrations (bandpass, amplitude, polarization, and phase) on an "as-needed" basis which is determined by the SSR at run-time⁸, but

⁵This intents are identifiers that the ALMA pipeline uses to process the data

⁶These groups are different from the Group OUS described in Section 8.1

⁷Additional calibrations can be added into subsequent groups to be interleaved according their cycle times. One example is the polarization calibration

⁸On some cases the atmospheric calibration might get associated to the science target

Script Name	Purpose
StandardInterferometry.py	Performs a regular interferometric ALMA observation
StandardSingleDish.py	Performs a regular ALMA single dish observation (Total Power)
AmpliCalSurvey.py	This script executes the observation of calibration sources (routine monitoring of quasars)
TowerHolography.py	Performs a surface calibration observation using the holography transmitter

Table 8.1: The list of the most common observing scripts in the SSR are listed here along with brief description.

with little control from the user. Usually, pointing is verified 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 sequence until the 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. This process is repeated until the observing requirements are met, or the SB reaches its time limit. Any additional (*secondary*) phase calibrators are observed as specified in the scheduling block. For a description of what each calibration entails, see Chapter 10.

The user has several options to select optimal calibrators. He/she can let the OT set up default queries to the calibrator database in the ALMA archive which will be used to select appropriate calibrators at run time (*system-defined calibration* in the OT). This is the recommended mode. He/she can also 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 (*user-defined calibration* in the OT).

8.3.1 The source selection algorithm

The ALMA system calibrations for a given SB is done using astronomical sources selected at runtime. To correctly select the appropriate calibration sources, a selection algorithm has been implemented in the SSR using the following steps:

- A list of sources is retrieved from the ALMA catalog using the query center and the search radius defined in the SB field-source parameters as inputs.
- For each source in the list we check whether there are flux measurements at the SB representative frequency. If no flux measurement is present at the desired frequency, we extrapolate the flux using a spectral index of -0.7 to the SB representative frequency from the closest measurement available in frequency⁹.
- The flux is weighted based on proximity to the representative frequency, target, and time since the last measurement present in the source catalog.
- A signal-to-noise criteria is applied to the returned list of sources and ranked based on source strength, separation respect to the query center and flux error.
- The final list is sorted based on ranking and returned.

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

⁹Higher frequency bands such as Band 9 or 10, may not have current measurements from the ALMA source monitoring observations due to the weather requirements, which increase the chance that the flux used by the query is extrapolated

- Bandpass calibrator. A signal to noise of 50 given a 15 minutes integration is used in the flux estimation. The smallest spectral window defined in the SB (bandwidth) is used for this calculation. The default search radius is defined to be 45° from the query center, but the user can overwrite this value as needed. A minimum elevation is calculated to avoid antenna shadowing.
- Phase calibrator. A signal to noise of 15 given a 2 minutes integration is used in the flux estimation. The widest spectral window bandwidth (usually 2 GHz) is used for this calculation. The default radius used in the search is 20° from the query center, which can be modified in the SB as needed. As with the bandpass a list of sources is assembled that passed the signal to noise criteria, but the closest one to the query center is returned. Thus, we privilege proximity to the science target rather than flux for the phase calibration.
- Flux calibrator. Solar system objects are queried first, then grid sources¹⁰. The first visible grid source is used. The following list corresponds to both solar system objects and grid sources used for flux calibration,
 - Solar system objects: Mars, Venus, Jupiter, Uranus, Neptune, Titan, Ceres, Ganymede, Callisto, Pallas.
 - Grid sources J0237+288, J0238+166, 3c84, J0334-401, J0423-013, J0510+180, J0519-454, J0522-364, J0538-440, J0635-7516, J0750+125, J0854+201, J1037-295, J1058+015, J1107-448, J1146+399, 3c273, 3c279, J1337-129, J1427-421, J1517-243, J1550+054, J1613-586, 3c345, 3c454.3, J1733-130, J1751+096, J1924-292, J2025+337, J2056-472, J2148+069, J2232+117, J2258-279, J2357-5311

8.4 Single Field Interferometry

Single field interferometry is the most basic form of observation that ALMA supports. It consists of standard calibration scans associated with the constituent targets, and science observations of a single field (primary beam). 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 amplitude calibration will be taken 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 model fluxes of these objects to obtain the scaling factor to be applied to all other sources in the SB. Ideally these amplitude calibration sources should be small with respect to the synthesized beam so as not to resolve the source structure and create uncertainties in the flux calibration. In practice many solar system objects may be moderately resolved. In such cases only a subset of the antennas 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 from the catalog of routinely monitored sources by ALMA.

As the interferometer is a device for measuring the spatial coherence function, the phase calibration is used to calibrate the phase variation at the individual antenna elements during the course of an observation. Since the phase is expected to change much more rapidly in time than the amplitude, these sources will be observed more frequently than the other calibrators. Because the phase varies on small scales on the sky, the calibrator must be quite close to the science target. The phase calibration observations are taken before and after observations of the science target, and the phase correction will be interpolated in time when applying to the science target. As the atmosphere can fluctuate rapidly (specially at higher frequencies), radiometric observations of atmospheric water lines are done to correct for these additional phase variations. These corrections can be applied online and offline (See Chapter 5).

¹⁰Grid sources correspond to a set of sources (usually quasars) that ALMA monitors periodically for flux calibration purposes

8.5 Pointed Mosaic Observations

Pointed mosaic observations are also offered for this Cycle. This mode enables a single *Science Goal* to cover a field of view larger than the primary beam by making observations of multiple single fields that overlap by an amount specified in the proposal. 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 execution¹¹. In pointed mosaic observations, each of the fields will be assigned different field identifiers in the data, but the same source identifier. Thus all fields will share their bandpass, amplitude and phase calibrations which are done as single fields. For larger fields that cannot be covered by 150 pointings, multiple SBs may need to be defined. The mosaic is arranged as a single scan composed of a number of subscans corresponding to the individual positions in the mosaic (see Section 8.2). Both the science target specific last mosaic pointing position and index within the science target list are handled by the SSR to ensure that the next scan begins on the proper science target and proper offset position. Regular mosaic observations (including multiple small mosaics) should be specified using a *Rectangular Field* in the OT; for irregular mosaics *custom-mosaic* under *Multiple Pointings* should be selected.

8.6 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*. For convenience, these observations are also sometimes referred to as *total power* (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.

The observing script that executes this mode is called *standard single-dish* script. As with *standard interferometry*, this script uses the SSR capabilities to perform some of the calibrations observations (e.g. pointing) interferometrically. The SB of a TP observation consists of a group of calibrations followed by an *On-The-Fly Mapping* (OTF) observations of a rectangular area on the science target for line mapping with periodic offsets to reference position for calibration observations. For the OTF and reference position integrations, only the auto-correlation data will be written to the data set or ASDM¹² and to minimize data rate and size, whereas calibrations use cross-correlation information for analysis.

The OTF map is observed as a series of raster rows, scanning in the coordinate system specified in the SB. In Cycle 3, scans are taken either in longitudinal or latitudinal directions as specified in the OT. In later cycles, scans may be taken in the two perpendicular directions in turns to minimize scanning effects. The reference positions, assumed to be positions which are free of emission, are specified as offsets relative to the map center. By default, the OTF map will cover an area one half of the full beam larger than the interferometric observation on all sides. 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, and the length of a scan in time (consisting of some number of subscans) is defined in the SB. The reference position is observed as specified by its cycle time during the science target scans. The pointing is calibrated on a bright calibrator near the science target with a frequency indicated in the SB, and atmospheric calibrations are taken 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 a Kelvin-to-Jansky conversion factor, which will be delivered to the project principal investigator (See Chapter 10). These are measured separately for each project by obtaining a continuum map of a bright quasar or a planet with known flux. These *amplitude calibration* maps are reduced similarly to the

¹¹In future cycles this restriction may be lifted by using sessions, as the mosaic field index will be remembered from one SB execution to the next

¹²The ASDM it is also loosely used to name the actual raw data that is archived after an observation is executed

science observations, but keeping the considerations required to process continuum data (e.g. no continuum subtraction). The emission is compared to a known model or observed flux (taken from the most recent calibrator survey measurements of grid sources in the case of quasars) to obtain the conversion factor, which will then be applied to the science observations.

8.7 Polarization

The ALMA antennas have receivers with linearly polarized feeds followed by a waveguide with a polarization splitter¹³. Thus, the incoming radiation is separated into two orthogonal components (X and Y) which are downconverted and digitized independently. The digital signals are cross-correlated for each baseline at the correlator and the output are four cross-correlations (XX, YY, XY, and YX). These four cross-correlations, as a function of the Stokes parameters, are ideally given by

$$V_{xx} = I + Q \quad (8.1)$$

$$V_{xy} = U + iV \quad (8.2)$$

$$V_{yx} = U - iV \quad (8.3)$$

$$V_{yy} = I - Q \quad (8.4)$$

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

$$I = \frac{V_{xx} + V_{yy}}{2} \quad (8.5)$$

$$Q = \frac{V_{xx} - V_{yy}}{2} \quad (8.6)$$

$$U = \frac{V_{xy} + V_{yx}}{2} \quad (8.7)$$

$$V = \frac{V_{xy} - V_{yx}}{2} \quad (8.8)$$

From here, it is easy to see that the total intensity is in Stokes I (function of the *parallel hands* XX and YY), the linear polarization is described by Stokes Q and U ; while circular polarization is described by Stokes V . However, there are number of issues that prevent us from directly using the measured cross-correlation in this simple way.

1. The splitting of the incoming radiation into 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 frequency and band (cartridge design and external optics). Additionally, this quantity is measured in the frame of the antenna which is an elevation and azimuth frame (Alt/Az) and thus, it rotates with respect to the frame of the sky. This rotation introduce an angular dependance into the visibilities as a function of the paralactic angle. By design, the instrumental polarization is small (a few percent), but not negligible.
2. The 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. An XY-phase offset is introduced by using a reference antenna. The XY-phase offset is a spectral phase-only bandpass relating the X and Y systems of the reference antenna.

¹³ In most of the antennas this device is an *Ortho-Mode-Transducer* (OMT) with the exception of Band 7 that has a polarized grid

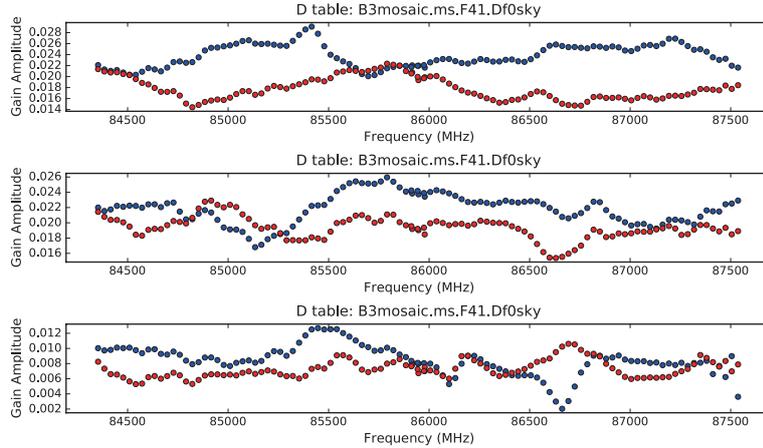


Figure 8.3: The (D – term plots of DA42 (top), DV03 (middle), and PM01 (bottom). The vertical axis is the fraction of the input signal voltage in one polarization that leaks into the output of the other polarization in voltage unit and the horizontal axis is the frequency in MHz. The blue and red symbols 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.

4. Other effects, such as position within the primary beam (off-axis polarization), are still under commissioning. Thus, only on-axis polarization is offered for Cycle 3 (see below for a detailed explanation.)

By taking into consideration the instrumental polarization and the parallactic angle dependence, we can re-write the equations for the visibilities as

$$V_{XX} = (I + Q_\psi) + (U_\psi + iV)d_{X_j}^* + d_{X_i}(U_\psi - iV) + d_{X_i}(I - Q_\psi)d_{X_j}^* \quad (8.9)$$

$$V_{XY} = (I + Q_\psi)d_{Y_j}^* + (U_\psi + iV) + d_{X_i}(U_\psi - iV) + d_{Y_j}^* + d_{X_i}(I - Q_\psi) \quad (8.10)$$

$$V_{YX} = d_{Y_i}(I + Q_\psi) + d_{Y_i}(U_\psi + iV)d_{X_j}^* + (U_\psi - iV) + (I - Q_\psi)d_{X_j}^* \quad (8.11)$$

$$V_{YY} = d_{Y_i}(I + Q_\psi)d_{Y_j}^* + d_{Y_i}(U_\psi + iV) + (U_\psi - iV)d_{Y_j}^* + (I - Q_\psi) \quad (8.12)$$

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 parallactic angle. Figure 8.3 shows an example of D – terms at Band 3. The D – term level is typically few percent at Bands 3, 6, and 7 on axis with some variations over frequency. 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. The cross-hands output will be purely D – terms since Stokes Q, U, and V will have no signal. However, bright unpolarized sources are rarely found. Alternatively, it is possible to separate the source polarization, the D – terms, and the cross hands phase spectrum and delay by observing an unresolved polarized source over a wide range of parallactic angles, where many quasars suffice this condition. The polarization calibrator should be observed at least to cover a wide range (> 45 degrees) of parallactic angle variation. The rotation of the parallactic angle mostly occurs right before and after the transit, but low declination sources do not rotate much (see Figure 8.4). The D – terms may have a slight elevation dependence, and therefore it is favorable for the polarization calibrator to be close to the science target. The ALMA observatory will choose appropriate calibrators, but the users must take into account these limitations when planning the observation. The users can check the calibrator candidates from the ALMA polarization calibrator database.

With the current calibration scheme, linear polarization imaging of a compact source, on-axis, at the level of 0.1% in fractional polarization is feasible. The accuracy of absolute polarization position angle will be nominally 6 degrees, which will scale down with the number of antennas in final image. The user will be encouraged to see

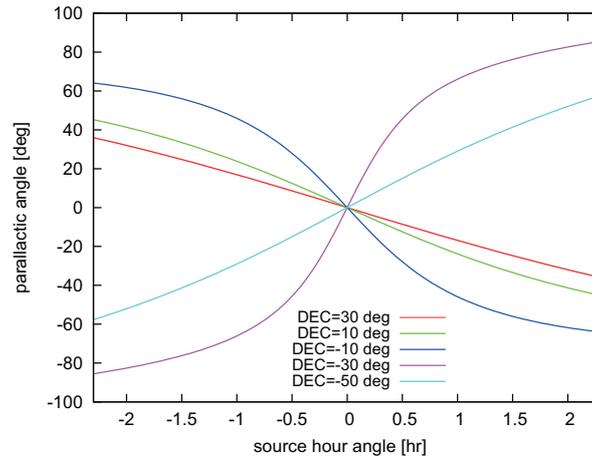


Figure 8.4: Parallactic angle plot as a function of hour angle for the declination of 30, 10, -10, -30, and -50 degrees.

the polarized flux and polarization position angle of the calibrator in the data reduction process and check the consistency comparing with the results obtained by the ALMA calibrator monitoring survey. Although a Stokes V (circular polarization) image can be produced with reduction and analysis software (CASA), this capability is still under commissioning. The ALMA observatory cannot guarantee the accuracy of Stokes V images at the current cycle. Thus, the users may investigate the Stokes V emission at their own risk.

As the D – term component also arises from the off-axis geometry of feed horn, antenna illumination, and the alignment of optics, it will also vary across the primary beam pattern. Generally, the D – term level becomes larger when increasing the offset from the beam center. This is so-called *off-axis* instrumental polarization. In Cycle 3, the off-axis instrumental polarization calibration will not be employed. Therefore, the acceptable angular size for the observing source is limited where the off-axis instrumental polarization level does not significantly become large. For this reason, the largest angular size of the source must be less than $\text{FWHM}/3$ of the primary beam. Polarization imaging with accuracy better than 0.3% can be achievable in this area. The users must justify that the source polarization level of science target is well higher than these accuracies.

For this cycle, polarization observations are limited to TDM correlation mode (or 64 channels per polarization product with 2 GHz bandwidth), but at arbitrary frequencies within the offered bands (3, 6, and 7). No higher spectral resolution, mosaic, ACA, and TP Array observations are offered. Further, a minimum execution time of 3 hours will be imposed to ensure that sufficient coverage in parallactic angle for the calibration.

8.7.1 Sessions

Though not directly related to polarization itself, sessions is the observing scheme designed to observed polarization projects with ALMA. In order to allow the execution of long programs and to avoid the execution of long scheduling blocks (more than 1.5 hours), we implemented the concept of a session into the SSR. Given the current stability of the system, it was considered an unnecessary risk to have more than 1.5 hours scheduling blocks, because in the event of a crash all data are lost. Thus, a session scheme was designed as an alternative. A session is defined as the continuous execution of the same SB until the scientific criteria are met. A session will manage the cycle time of each of the calibrators from the starting point of the session, i.e. the first execution of the SB. In this way, calibrations are only executed when needed avoiding unnecessary observations giving an additional level of optimization for ALMA. By default, calibrations such as flux calibration will be done once per project and bandpass every hour saving observing time. Also, the phase calibration science target loop will be broken (preserving phase calibration) when an additional calibration is-needed (e.g. polarization)¹⁴. The current mode for sessions in ALMA is the observation of polarization projects. Because for an Alt/Az antenna

¹⁴The cycle time parameters are user controlled and can be explicitly specified in the OT target parameters section

the frame of the sky rotates, the calibration of the of instrumental polarization requires sampling a strong compact source as a function of parallactic angle approximately every 30 minutes. The session will remember the last time the polarization calibrator was observed and interweave the calibration when needed. We can achieve the parallactic angle coverage by running the polarization SB 3 times, which will give between 3 to 4 hours of observation. This constitutes a perfect use for the session scheme. Other cases where the session scheme might be useful are large mosaics, surveys (multi-target), and large single dish raster maps, but those cases are not offered in session mode for Cycle 3. The implementation of sessions is done by keeping information in memory about the previous execution of the SB. This is done through an interface provided by CONTROL software to access a persistent component which we use to store and retrieve this information.

8.8 Multiple Region Modes

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

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

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 4 separate windows, each with 62.5 MHz width, anywhere within the 2 GHz baseband. In Cycle 3, a maximum of 4 spectral windows per baseband will be offered.

8.9 Observations of Ephemeris Objects

Observation of solar system objects (with the exception of the Sun) is supported in Cycle 3. Several well known solar system objects including planets, satellites and asteroids can be selected from a pulldown menu in the OT. For other non-sidereal objects an external ephemeris file can be supplied as an input. 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. An enhancement for Cycle 3 are high precision ephemeris which are now supported (See Figure 8.5). The ephemeris file should contain at least one row for every hour and, except for fast objects, one row every 20 minutes is sufficient. The online system will linearly interpolate between rows in an ephemeris when determining the antenna pointing & signal propagation delay. A general example of an ephemeris file is:

```
*****
Date__(UT)__HR:MN      R.A.__(ICRF/J2000.0)___DEC      delta      deldot
*****
$$SOE
2014-Dec-01 00:00 Am  18 51 01.8305 -20 41 18.087 33.5888476523329  17.7176571
2014-Dec-01 00:20 Am  18 51 01.9382 -20 41 18.057 33.5889897735456  17.7169995
2014-Dec-01 00:40 m   18 51 02.0461 -20 41 18.027 33.5891318772715  17.7132925
2014-Dec-01 01:00 m   18 51 02.1542 -20 41 17.997 33.5892739389694  17.7065199
2014-Dec-01 01:20 m   18 51 02.2624 -20 41 17.967 33.5894159340629  17.6966891
```

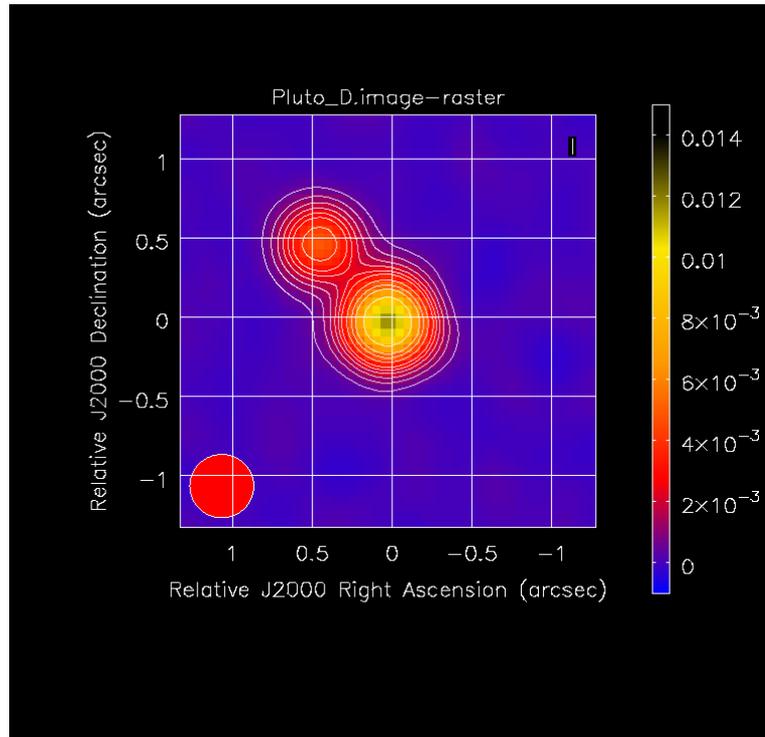


Figure 8.5: Example of a high precision ephemeris test experiment to track the position of Pluto. Charon is also seen at the top left.

More information on the format and precision needed is available in the ALMA Observing Tool documentation. Finally, solar observations are not allowed for Cycle 3 but will be available at a later date.

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 allows a user (via a GUI) to experiment with various sensitivity options, but which is also internally used by the OT to calculate its time estimates based on the parameters entered into a project’s Science Goals. The GUI version is also available as a Java applet in the ALMA Science Portal¹.

Although the user may experiment with various sensitivity options (PWV octile, number of antennas, etc.) in both the Java applet and the OT’s GUI, the final time estimate for a project cannot be influenced to the same extent. For instance, the OT will always assume a fixed number of antennas for the particular Cycle, and will always use the PWV octile that it calculates as being appropriate to the frequency of observation.

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 are crucial.

9.1.1 Sky temperature

The OT’s estimate of both the atmospheric opacity and the sky temperature are calculated using the Atmospheric Transmission at Microwaves (ATM) code². This provides values of the 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 assumes that the source is at the zenith and therefore the OT has to account for the greater atmospheric emission at lower elevations. The emission from the atmosphere is often approximated as

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

where T_{atm} is the mean physical temperature of the atmosphere. This can be calculated using Equation 9.1 and

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.

from there the sky temperature as a function of zenith angle:

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

T_{sky} 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.

$$T_n = T \times \left(\frac{h\nu/kT}{e^{h\nu/kT} - 1} \right) \quad (9.3)$$

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 9.1.1.

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 50 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 version of the ASC, 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

For many of the ALMA bands, the calculator currently only uses the specifications for the receiver temperatures (required over 80% of the receiver bandwidth) and not the actual measured values. However, for five of the Early Science Bands (3, 6, 7, 8 and 9), typical values measured in the laboratory are used as these are usually significantly better than the specifications. The measured values are somewhat conservative and so are in between what we might expect at the middle and edges of the bands. 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.

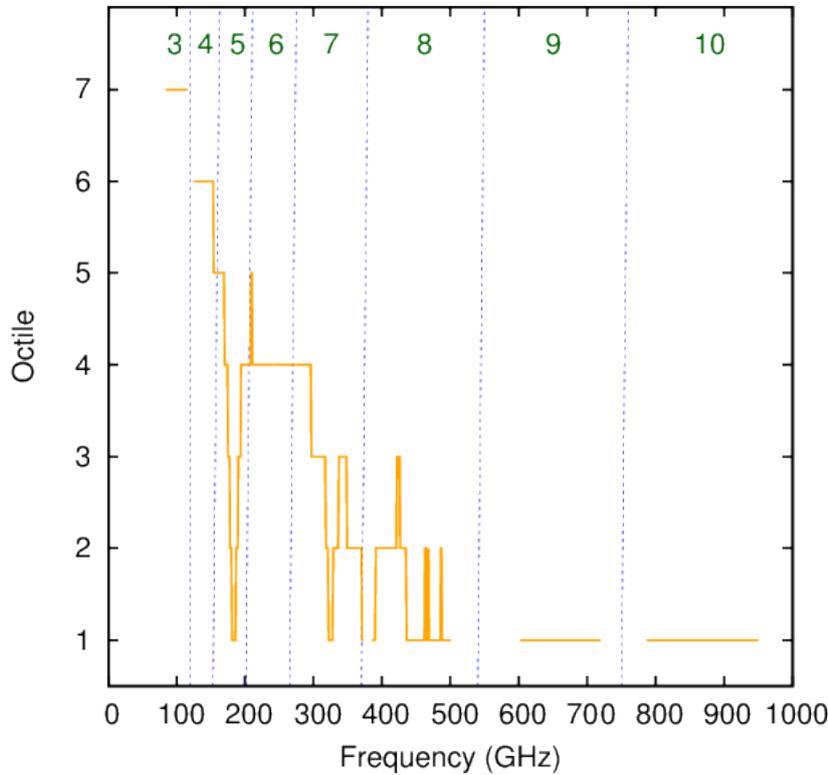


Figure 9.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown at the top of the plot. The water line at 183 GHz (Band 5) is particularly prominent. In general, higher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

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.

Note that the calculator 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 9.3.

9.1.4 Ambient temperature

This 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 9.3 and thus its total contribution is frequency dependent and can vary between the different sidebands.

ALMA Band	Receiver Type	$T_{\text{rx,spec}}$ (K)	$T_{\text{rx,ASC}}$ (K)
1	SSB	17	17
2	SSB	30	30
3	2SB	37	45
4	2SB	51	51
5	2SB	65	65
6	2SB	83	55
7	2SB	147	75
8	2SB	196	150
9	DSB	175	110
10	DSB	230	230

Table 9.2: Receiver temperatures (and their specifications) assumed in the ASC as a function of ALMA band. For most of the bands we are currently assuming the ALMA specification for the receiver temperature that should be achieved across 80% of the band, $T_{\text{rx,spec}}$. In practice, the receivers actually outperform the specification and for Bands 3, 6, 7, 8 and 9 the ASC uses “typical temperatures measured in the laboratory” (highlighted in bold text).

9.1.5 DSB receivers

Every heterodyne radio receiver detects radiation from two sidebands simultaneously, this being due to the way that the astronomical signals are down-converted to an intermediate frequency before further processing. This means that, if nothing were done to prevent it, a spectral window 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 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 and so the noise cannot be neglected.

One important consequence of this is that, if a spectral window 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. This is the case even though no data is being produced from the image sideband! Therefore, it is important to avoid areas of bad atmospheric opacity in the image spectral windows and the OT therefore shows the location of these in the spectral visual editor (Figure 9.2).

One subtlety is that the software that takes the input spectral window frequencies and calculates the tuning solutions from these will always try and place these in the upper sideband if possible. Therefore, to take a simple case, a single spectral window centred at 637 GHz will find its image equivalent in the middle of the zero transmission feature at ~ 621 GHz. This situation should be avoided if at all possible, including defining dummy spectral windows such that the line of interest is swapped (Figure 9.2).

9.1.6 System temperature (OT version)

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:

$$T_{\text{sys,dsb}} = \frac{1}{\eta_{\text{eff}} e^{-\tau_0 \sec z}} \left(2 \times T_{\text{rx}} + \eta_{\text{eff}} (T_{\text{sky,s}} + T_{\text{sky,i}}) + (1 - \eta_{\text{eff}}) \times (T_{\text{amb,s}} + T_{\text{amb,i}}) \right) \quad (9.4)$$

where

- T_{rx} – receiver temperature
- $T_{\text{sky,s}}$ – sky temperature at the requested frequency in the signal sideband

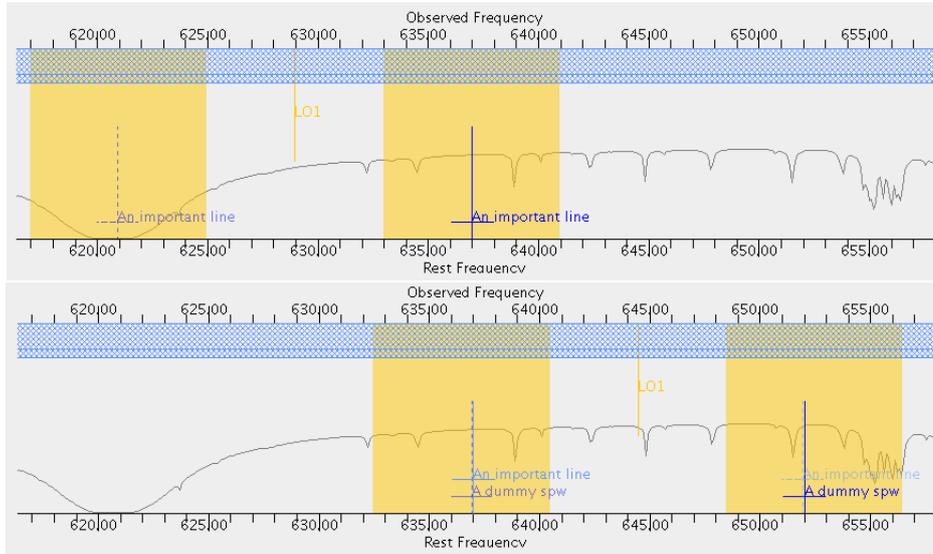


Figure 9.2: A Band-9 spectral setup as displayed in the spectral visual editor in the OT. The top figure shows an example of how the image equivalent of a single spectral window can fall into an area of poor atmospheric transmission, leading to much higher T_{sys} than necessary. Placing a dummy spectral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

- $T_{\text{sky},i}$ – 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).

The terms η_{eff} and $e^{-\tau_0 \sec z}$ both attenuate the source signal and we thus 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 receivers that aren't DSB, the equation is the following:

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

where the terms are all the same as in Equation 9.4.

9.1.7 System Temperature (Applet and GUI version)

The equation used for the Java applet and the GUI that is accessible from within the OT is similar, but these are stand-alone applications that do not know the location of the image sideband and thus cannot perform a rigorous calculation of the contribution of this 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 :

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

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 residual pointing and focus error.

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:

$$\sigma_{\text{S}} = \frac{2kT_{\text{sys}}}{\eta_{\text{q}}\eta_{\text{c}}A_{\text{eff}}\sqrt{N(N-1)}n_{\text{p}}\Delta\nu t_{\text{int}}}. \quad (9.7)$$

The various symbols are

- A_{eff} – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency (η_{ap}). The latter is given by the Ruze formula i.e. $\eta_{\text{ap}} = R_0 \exp(-16\pi^2\sigma^2/\lambda^2)$ where σ is the rms surface accuracy of the antenna – the specification of 25 μm and 20 μm for the 12-m and 7-m antennas respectively is currently used³. R_0 is equal to 0.72. See Table 9.3 for values of antenna efficiencies and effective areas in various ALMA bands
- 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 efficiency *will* depend on the mode, but this is only taken into account in the OT, not by the Java applet (see below) which therefore also assumes a value of 0.88
- N – number of antennas. This defaults to 36 for the 12-m and 10 for the 7-m Array
- n_{p} – number of polarizations. $n_{\text{p}} = 1$ for single polarization and $n_{\text{p}} = 2$ 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 spectral window being limited to 1.875 GHz by the anti-aliasing filter through which the baseband signal passes. $n_{\text{p}}\Delta\nu$ 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

$$\sigma_{\text{T}} = \frac{\sigma_{\text{S}}\lambda^2}{2k\Omega} \quad (9.8)$$

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
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 the efficiency multiplied by the physical area of the dish i.e. 113.1 m² and 38.5 m² for the 12 and 7-m antennas respectively.

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

$$\Omega = \frac{\pi\theta^2}{4 \ln 2}. \quad (9.9)$$

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

9.2.2 Total Power Array

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

$$\sigma_{\text{TP}} = \frac{2kT_{\text{sys}}}{\eta_{\text{q}}\eta_{\text{c}}A_{\text{eff}}\sqrt{N}n_{\text{p}}\Delta\nu t_{\text{int}}}. \quad (9.10)$$

This is the same as Equation 9.7, apart from there only being a factor of \sqrt{N} in the denominator, where N is the total number of total power antennas.

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 a Java applet on a web page – 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.
- 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 (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.

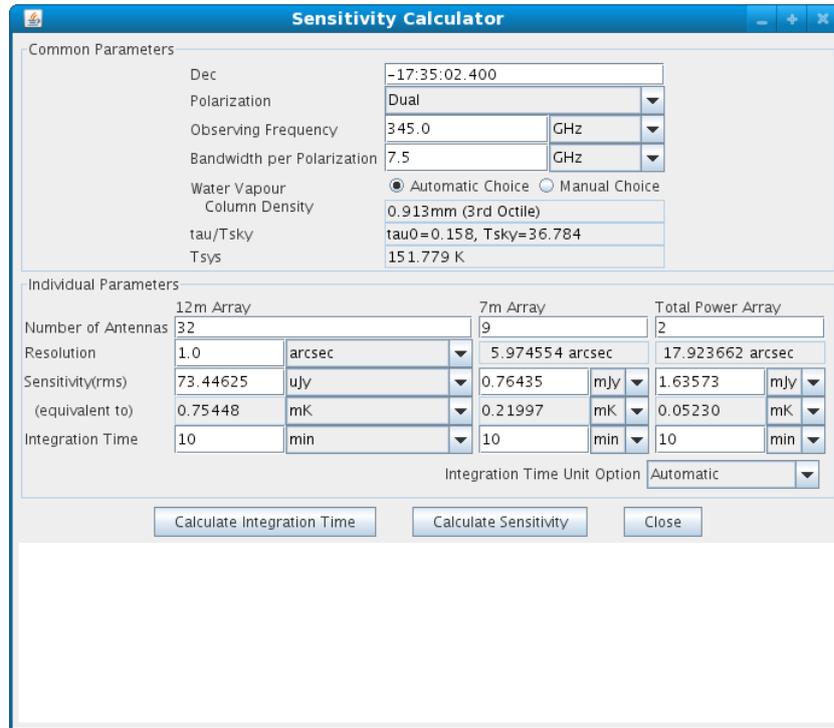


Figure 9.3: Screenshot of the GUI version of the ALMA Sensitivity Calculator as implemented in the ALMA Observing Tool. The white area at the bottom is for displaying error messages i.e. parameters out of bounds. The example here shows the achievable sensitivity for all three arrays for an on-source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

- Number of antennas – the ASC currently defaults to the values for Early Science, namely 36 from the 12-m Array, 10 from the 7-m Array and two 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.

The calculator reports the values of 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 therefore the time is always assumed to be the true on-source time. 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.

Chapter 10

Calibration and Calibration Strategies

10.1 The Measured Visibility Function

10.1.1 The Data Format and the Correlator Model

The primary output of the ALMA correlator is a complex quantity, called the *observed complex visibility function*, \mathcal{V}^o . Its amplitude measures the fractional correlation of the radio signals between any two antennas, and its phase measures the time-of-flight phase between the signals. These data, together with meta-data that describe relevant parameters associated with the observations, are stored in the ALMA archive (See Chapter 14). Details of the flow of the astronomical data from the antennas through the correlator to the archive are given in Chapter 14. These data are stored in the ALMA archive in the form of the ALMA Science Data Model (ASDM). For ASDM details see Chapter 12.

Because most observations are associated with celestial objects, their diurnal motion produces up to kilohertz phase rate changes in the visibility phase. For this reason a *correlator delay model* is included in correlator processing in order to remove the a priori signal delays that are a function of baseline length, observing frequency, and source position (See Chapter 5). An important part of the correlator delay model is the estimated atmospheric path delay to each antenna from the radio source with an assumed position, σ , called the *phase center* of the observations. With this correlator model delay/phase rotation removed, the visibility function can then be averaged coherently over many seconds. Other parts of the a priori model can be included in the correlator processing or applied directly in the antenna signal paths.

10.1.2 The Visibility Function Dependence

The fundamental observational piece is the *scan*. It is an observation with one field name, associated with a phase center, and one band, b , and many other parameters to specify the observation details. The scan has many time samples. Since virtually all ALMA observations use phase referencing as the major calibration method, scans associated with calibrator sources are usually interleaved with scans associated with the science target.

The visibility function associated with a scan/band has the following dependencies: time stamps, t_k , antenna-pair indices (i, j) , separation of the observing band into four SPW's s with further channelization, c . This frequency specification is called a *tuning* (See Chapter 6). Finally, each sample contains up to four polarization correlations, p_q . The observed complex visibility function, with all its dependencies, is written as

$$\mathcal{V}^o(\sigma, \nu_b)_{i,j}[t_k, \nu_{s,c}, p_q] \quad (10.1)$$

The complex visibility function is often described by its amplitude, \mathcal{A}^o and its phase ϕ^o components since virtually all of the calibrations operate on amplitude and phase, rather than real and imaginary parts of the visibility function.

The *autocorrection* data, where $i = j$ is also contained from the correlation of the data. The calibration and use of this data are given in Section 8.6 and Chapter 13.

10.2 Main Interferometric Calibration Strategies

10.2.1 Antenna-based Calibrations

Nearly all of the modifications of the measured visibility function are antenna-based since virtually all effects occur in the signal path are associated with one particular antenna, from the troposphere above the antenna through all of the electronic processing along any antenna path, up to the correlator. In general the calibration equation is:

$$V_{(i,j)}^o(t_k, \nu_c) = G_i(t_k, \nu_c) * \tilde{G}_j(t_k, \nu_c) V_{(i,j)}^t(t_k, \nu_c) + noise + g_{(i,j)}(t_k, \nu_c) V_{(i,j)}^t(t_k, \nu_c) \quad (10.2)$$

where $G_i(t_k, \nu_c)$ is the antenna-based complex gain calibration as a function of time and frequency. This is the normal calibration equation. The second line contains the effect of noise and a pure baseline-calibration, and these will be ignored for the most part.

The complex antenna gain is also separated into the antenna-based amplitude and phase gains,

$$G_i(t_k, \nu) = a_i(t_k, \nu) \exp[i\psi(t_k, \nu)] \quad (10.3)$$

where $a_i(t_k, \nu)$ is the antenna-based amplitude correction and $\psi(t_k, \nu)$ is the antenna-based phase corrections, both as a function of time t_k and frequency ν_c . These are the major calibration quantities that are described in this chapter.

10.2.2 Calibration Flow through ALMA system

The modifications of the radio waves from a celestial source occur at many locations along the path from the source to the data archive. A relatively complete list of signal changes, based on G Moellenbrock's 2014 Socorro Summer School Lecture,¹, occur roughly in the propagation order of the radio signal is:

I = Ionosphere: Opposite rotation of linear polarization caused by ionosphere. Unimportant at ALMA frequencies.

T = Troposphere: Absorption of signal and delay of signal caused by several atmospheric components.

P = Antenna Voltage pattern: The relative response of the antenna from the reference position, caused by focus and pointing offset.

p = Parallax angle: The rotation of the antenna feed orientation with respect to the sky (alt-azimuth mounts).

X = Linear pol angle: The orientation of the feed with respect to the antenna structure.

d = Polarization leakage: The non-orthogonality of the two orthogonal polarized feeds.

T = System Temperature: Amplitude calibration factor from system sensitivity.

M = Model used in the correlator to remove the major delay temporal and angular variations.

D = Delay offset of the signals between antennas.

B = Bandpass response: The relative amplitude and phase of an antenna system as a function of frequency.

¹https://science.nrao.edu/science/meetings/2014/14th-synthesis-imaging-workshop/lectures-files/MoellenbrockCalibration2014_FINAL.pdf

G = Electronic gain: The temporal amplitude and phase changes within the antenna system.

F = Flagging: Editing of data based on calibration problems.

Do not be intimidated by the above list. The main ALMA calibrations of interest to most users deal with the temporal electronic gain G and the bandpass response B . A good approximation that simplifies the calibration process is that the temporal dependence and frequency dependence of nearly all instrumental changes are only *lightly coupled*, so their variations can be determined independently, or at least iteratively.

10.2.3 A priori, Online and Offline Calibration Methods

There are three stages of calibrations in ALMA.

A priori Calibrations: These are calibrations that are associated with relatively fixed properties of the array, for example the antenna response patterns. They are measured periodically by the ALMA staff using methodologies that best suit the calibration. These calibrations are briefly outlined below.

One of the more important a priori calibrations in cycle 3 is the *Go/noGo* 2-min observation/analysis procedure on a strong calibrator. It is run periodically in order to assess the current ALMA phase stability. Within minutes the 30-second phase rms as a function of baseline length at any frequency will be determined. If the phase rms at the longer configuration baselines are larger than 30° at the frequency of the anticipated observations, then observations are not recommended. Unless conditions are changing rapidly, this Go/NoGo evaluation need be done only every few hours.

Online calibrations: Measurement of relevant instrumental parameters are made during the observations to determine calibration values that can be applied as the data are taken. These include antenna pointing calibration (applied as a collimation offset relative to the master antenna pointing models which are updated on approximately weekly timescales), measurement of $T_{\text{sys}}/T_{\text{RX}}$ using the Amplitude Calibration Devices (ACDs), and determining Water Vapour Radiometer channel coefficients to allow rapid online wet path delay variation to be compensated in the correlator data processing software. Application of some online calibrations such as T_{sys} are deferred to the offline calibrations, however.

Offline calibrations: The offline calibrations are those done after the completion of the experiment. These are usually associated with relatively fast-changing calibrations that cannot be made online. These correction are applied by the reduction scripts (pipeline or script generator) using measured system properties, such as the receiver operation temperatures, or by the analysis of of special-purpose *calibrator* observations that are embedded within the science observing blocks (SB). Two examples are the removal of the frequency dependence across the measured bandwidth, B , and the removal of the amplitude and phase changes during the experiment, G . The accurate calibration of these terms have the largest impact on the image quality, and are of most concern to the user.

10.3 ALMA-Supplied a priori Calibrations

The ALMA staff will provide calibrations that are long-term and/or require specialized observations and monitoring. These are outlined below:

Antenna Beam/Surface Characteristics (P): The relative sensitivity of each antenna as a function of azimuth and elevation from the beam center (called primary beam pattern) is used to correct the images of extended sources, and multi-pointings observations (mosaicing) with overlapping fields. The measurements are made by mapping the antenna beam using a bright radio source, either single dish scanning around a strong source or with interferometric observations at many points within the antenna sensitivity area. The measurements are processed in special-purpose software to determine the accurate antenna surface, and adjustments of the antenna panel supports can be made if the surface is not within specifications. The

measured primary beam patterns are available in the Casapy package and applied during the image-make step (see Chapters 3 and 7 for more details).

Pointing Models (P): Blind pointing of the array elements at a nominal accuracy of $2''$ RMS is achieved by a combination of metrology systems inside the antennas themselves to account for variable deformations, and static models maintained by the observatory. The static models are dominated by six terms: azimuth encoder/pad offset ("IA" term, order $1000''$), elevation encoder/collimation offset ("IE" term, few hundred arcsec), azimuth collimation offset ("CA", few tens of arcsec), tilt in North-South and East-West directions ("AN" and "AW" terms, order $10''$), and elevation sag ("HECE" term, around $10''$). The full models have 17 to 19 terms, plus collimation offsets for each receiver band relative to a reference band (traditionally Band 3, but this may be changed during cycle-3 in favour of a band without warm optics i.e. band 6 or 7). The terms in addition to those already mentioned, and the deviation of receiver band collimation offsets and antenna-type dependant elevation sag from a known average, are at the arcsecond level.

The pointing models are checked and adjusted on typically weekly timescales, primarily by making "all-sky pointing runs", which are executions of around an hour in duration that make typically 50-100 interferometric pointing measurements on a well-distributed set of bright quasars of accurately known position. Each pointing measurement scan consists of a set of subscans at different antenna pointing offsets around the expected pointing direction (with fixed delay centre). The scalar-averaged visibilities at each pointing offset are used to produce antenna pointing solutions, with typical accuracy of around $0.1''$. Ideally all antennas observe together on the Baseline (64-input) correlator, and the pointing offsets are appropriately scaled for different antenna diameters. Such pointing measurements in different bands on a subset of bright sources are used to determine the band collimation offsets for the models. The models are stored in the Telescope Monitor and Configuration Data Base (TMCDB), which includes change history. Pointing measurements over many weeks, and including measurements made in other observations, including science executions, can also be used in deriving models when appropriate. Currently only night time pointing measurements are considered in these *static* models, with daytime pointing modelling and metrology system improvements an ongoing effort.

When an array element is relocated from one pad to another, such as during array reconfiguration, the changes are primarily confined to the azimuth offset and tilt terms (IA, AN, AW). The tilt is measured after the relocation by spinning the antenna in azimuth and measuring the amplitude and phase of the outputs of inclinometers located above the azimuth bearing. This gravitational tilt is converted to an astronomical tilt by adding a local gravity vector (which in long-baseline configurations varies significantly across the array due to the nearby mountains). With the tilt updated in the pointing models, the azimuth offset is determined observationally by fitting total power data recorded while the antenna performs azimuth strokes across a planet or bright quasar. After this the pointing model is refined as described in the previous paragraph.

In all science executions pointing measurements are made near the science targets and calibrators at the time of execution, from which residual collimation offsets at typically arcsecond level are applied in an *auxiliary pointing model*, which is added to the master static TMCDB pointing model for each antenna. With this residual offset pointing correction, the pointing accuracy for science target and calibration scans is intended to be within $0.6''$ RMS. The total pointing model in-use at any time during an observation (sum of primary and auxiliary models) is stored in a `PointingModel` table in the ASDM datasets.

Focus Models (P): The focus of the ALMA antennas is adjusted by moving the subreflectors in three translational axes during observations according to a six parameter model, plus offsets for each receiver band relative to the reference band (Band 7). The models contain an offset for each axis which describes the optimum subreflector position for Band 7 at elevation 50° and temperature 0°C ("XR", "YR", "ZR"). In the axial (Z) and elevation-oriented direction (Y) there is a coefficient of sag with elevation ("ZS", "YS"). For the Z axis there is a temperature coefficient ("ZT"). These models are maintained by long-term compilation of focus measurements in all receiver bands for the antennas, with care taken with regard to hardware changes which may change the focus characteristics. Focus measurements are traditionally made in a similar way to as described for pointing measurements, except the subreflector is offset (in one axis per scan) relative to the nominal model position. The subreflector position in the axis being offset which yields maximum gain is solved for each antenna from the cross-correlation data. During Cycle-3 a

faster technique, deriving the 3 axis focus offsets from holographic beam maps (which don't require the subreflector to move), may be used in some cases. These measurements are made by the observatory as needed to maintain the models. The models are stored in the TMCDB in the same manner as the pointing models.

Typically at night time these models, including temperature dependence (a new feature for Cycle 3), can maintain the focus to within a few tens of microns, which is adequate for most observations. Particularly the X and Y position is generally better described by the models than can be determined from an individual focus measurement. Generally when switching to a new receiver band it is a useful cross-check to make a focus calibration in the Z axis, for which standard Scheduling Blocks are available, which allow to check for pointing and delay errors as by-products. These are especially important for Band 9 and 10 where focus accuracy is most critical, but this currently requires an extremely bright quasar to be available, or the atmospheric conditions to be excellent. In daytime the focus can deviate from the models by several hundred microns. In daytime a focus calibration observation is generally executed between every science execution for Band 6 and above (typically Band 7 is the absolute maximum for daytime observing) in order to try to track these variations, but the poor atmospheric stability means that the focus measurements themselves have large uncertainties. Generally it must be expected to have focus errors at the 100 micron level in daytime ALMA data. At night errors below 20 microns are achievable. As for pointing, modelling and correction of focus during daytime is ongoing work.

The focus model in-use at any time during an observation is stored in a `FocusModel` table in the ASDM datasets.

Antenna Delay Offsets (D): Short observations are often made at the beginning of each day's run in order to check the delay offsets of each signal path of each array element. These are made with short interferometric observations at the relevant tunings as needed. Typically if offsets of more than around 100 ps are seen then the delay model in the TMCDB will be adjusted. Delay measurements for a given tuning are typically repeatable to better than 10 ps, but the measurements vary by a few tens of ps from tuning to tuning due to receiver RF and IF phase response variations. PIs can expect online delay errors in their data to be below 100 ps, so it is possible to use channel-average data for quick continuum analysis. The primary motivation for maintaining accurate delays is for pointing, focus and beam measurement calibrations, which use the channel-average data to reduce processing time and allow very short integrations where needed.

Antenna Positions (M): The relative location of each antenna (essentially the relative positions of their focal points) must be known to within a fraction of the observing wavelength, about 50 microns. These values are among the most critical for the determination of the accurate correlator delay model, M . For the main array, the specification is 65 μm and for the 7-m array it is 33 μm . The baseline runs (as they are called) are executed by the ALMA staff at least once per week, but also whenever one or more antennas change pad locations. Then antenna positions are determined from about 50 to 100 one-minute interferometric scans at 100 GHz of quasar calibrators that are unresolved and have milli-arcsec accurate positions. The observations cover a wide range of elevation and azimuth to determine the three spatial offsets of the antenna position from that listed in the TCMBD used in the correlator model. The typical position offset associated with an antenna that has been moved is 1 mm even for pads that have already been occupied, well above the specification, so that baseline runs must be made after any configuration change. Depending on the weather conditions and other array operation constraints, the position of a recently moved antenna may not be updated in the TCMBD for several days. Thus, the correlator delay model will be incorrect and the visibility phase for data associated with this antenna can have significant phase variations with time and between sources. There are offline procedures in the data reduction scripts to correct for these known antenna position offsets from available antenna location/pad files.

10.4 Online Calibrations

Water Vapor Radiometer Corrections (M): Fluctuations in the line-of-sight precipitable water vapor (PWV) of an antenna cause significant delay corrections, up to 0.3 mm/sec which is 30° of phase/sec

at 90 GHz, and scales in phase linearly with frequency. The changes are driven by wind, and cover a large range of spatial scales. All ALMA 12-m antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures the emission of the atmospheric water line within four spectral bands near 183 GHz at a rate of about 1 Hz. The radiometers are described in Sections 5.4 and A.6. From these measurements, atmospheric models are used to convert the emission into delay changes in the line-of-sight of each antenna, and this estimated delay difference between any two antennas is removed from the visibility phase. For cycle 3 the PWV correction will be applied online by inserting the appropriate value (M) into the correlator model, M . Less than 10% of the times, the application of the WVR corrections increases the phase rms because of clouds or ice above the array. In these cases it is best not to use the WVR correction. The water vapor emission in the four channels are also stored in the archived visibility data base, so the corrections can be determined and applied offline using a similar algorithm. The 7-m array with a maximum baseline of 50 m does not have water vapor radiometers since the phase corrections between these antennas are small. The PWV measurements on the four PM antennas that surround the 7-m could be interpolated to apply an approximate WVR correction to each of the 7-m diameter antenna, but this is not planned.

System and Receiver Temperature (T): At millimeter and submillimeter wavelengths, the atmosphere both attenuates astronomical signals and acts as a black body emitter that adds additional noise to any measurements, T . This effect is a strong function of frequency, elevation, the column of wet and dry constituents of the atmosphere, and the temperature of the atmosphere. The noise from the receivers can also vary for a given tuning with factors such as cryostat temperature and pointing direction with respect to the local magnetic field.

To measure these effects, the ALMA front ends are equipped with an Amplitude Calibration Device (ACD) which can move hot and ambient blackbody loads in front of the receiver feeds. Consecutive autocorrelation (or total-power) integrations on these loads and the sky are used to measure the T_{sys} and T_{RX} in each spectral channel. This depends on linearity of the autocorrelation data with input power, which currently limits these observations to the lower spectral resolution TDM mode (256 total channels, divided across polarisation products) of the Baseline correlator^{2,3}. For the ACA correlator in Cycle-3 it is planned to use the same spectral configuration (resolutions) for temperature calibration as the science setups, as it does not have distinct operating modes and produces linear output⁴.

These are naturally total temperature measurements, which include the receiver-weighted contributions from both sidebands. In cross-correlations the image sideband only contributes noise but not signal (LO switching sequences are used to highly suppress image sideband cross-correlations), so the natural T_{sys} and T_{RX} are scaled-up by a sideband gain factor to produce single sideband temperatures which can be directly applied to visibilities to put them in temperature units. In Cycle-0 to Cycle-2 scans were included in every science execution to measure baseband-averaged sideband gains, however, at present it appears these are not actually helpful overall in terms of calibration accuracy⁵ and the approximately 70 seconds of time taken per execution is not justifiable. For Cycle-3 it is planned to only use default sideband gains that are a constant per receiver band (nominally 0.99 for 2SB and 0.5 for DSB bands). This policy may be changed as deemed necessary by the observatory.

These measurements should be sampled every 5-15 minutes depending on frequency. The specifications for T_{sys} calibrations are to reach 1% repeatability. The online measurements of T_{sys} and T_{RX} are stored in the `CalAtmosphere` table in the ASDM datasets, and are displayed as they are produced during the

²In TDM mode the autocorrelation zero lags can be used to fully compensate for quantisation effects on the autocorrelations, however, in FDM mode which adds a filter bank before the correlations are produced an external measure of the total digitised power is needed. The current plan is to use the baseband square-law detector data to provide this extra input for Cycle-4.

³At present the online processing time for FDM mode calibrations is also unfeasibly long for large arrays of antennas.

⁴in previous Cycles ACA temperature spectrum resolution was arbitrarily limited in the same way as Baseline correlator observations, although there was not technical reason for this. Opening up the resolution for the ACA in Cycle-3 is intended to ease the route to processing FDM temperature spectra from the Baseline correlator in future cycles. Note that this includes multiple spectral windows per baseband.

⁵The main problem is that when you really care about the sideband gain is when strong atmospheric emission features are present in either sideband, but this unfortunately is the worst case to measure as you have low S/N from the calibrator source at the relevant frequency. Obtaining sufficient S/N also means that we can generally only derive baseband-averaged sideband gains, which are of limited applicability to the atmospheric lines which are to be accounted for. Generally the measurements are error-prone, and the situation gets worse with fewer antennas (i.e. ACA arrays) and at higher observing frequencies.

observations (baseband-averaged) to allow problems to be immediately seen by AoD or Operator. The calibrations can also be re-derived offline from the raw data if there are problems online or it is desired to use a different or updated algorithm. The T_{sys} spectra, suitably interpolated, are applied offline to the correlation data. Assuming that the correlated data is in units of percentage correlation, multiplication by the T_{sys} will change the correlated data units to Kelvin.

Additional Antenna Pointing Checks (P) As mentioned previously under A Priori Pointing Models, pointing corrections are made during executions in order to reach ALMA's 0.6'' RMS pointing accuracy specification. These account for small variations with respect to the static models which are updated on weekly timescales. To some extent these per-execution pointing calibrations compensate for shortcomings in the antenna metrology systems, and errors in the static models, for instance due to hardware changes since the last update. These pointing measurements are managed by the observatory – there is no PI input. At present a pointing calibration is triggered when moving to a target more than 20 degrees from the previous pointing calibration.

Generally it is desirable to point using the receiver band used for science in an execution. This is mostly possible up to Band 8, at least for large arrays of antennas. At Band 9 and 10 the combination of high T_{sys} , low quasar flux, low antenna aperture efficiency ($\sim 0.2\text{--}0.3\%$ for 12m antennas) and inadequacy of the blind pointing accuracy (static model plus metrology correction) means that pointing must be done in a lower frequency band and the collimation offsets between the two bands be relied on. Future improvements may allow pointing in-band for Bands 9 and 10 in future Cycles. The choice of pointing tuning will be made by the observatory in order to maximise available array elements by avoiding bands or frequencies affected by hardware problems, and minimise pointing error.

The results of the pointing measurements in each execution are stored in the `CalPointing` table in the ASDM datasets, and are displayed as they are produced during the observations to allow problems to be immediately seen by AoD or Operator. These results can be used for offline data flagging of antennas e.g. if a correction is deemed too large, the measurement uncertainty too high, or the pointing result is missing e.g. because there was a problem with the antenna or receiver during the measurement.

10.5 Offline Calibrations

The offline calibrations are those that are applied by python/casapy scripts after completion of the SB observations. The first step in the process is the conversion of the archive ALMA Science Data Model (ASDM) format (the visibility data in binary form and the meta data in xml form) into a Measurement Set (MS) which contains one-dimension tables for all of the data products (Chapter 12). A combination of python scripts and casapy are then used for the calibration and display of the data.

An initial offline reduction at the OSF is called *QA0*. It is made on an MS that is been recently completed in order to determine if there are serious errors in the MS (See Chapter 12). If it passes this inspection, the SB is then queued for the complete offline calibration using the ALMA Pipeline for *standard* experiments, or the Script generator code for *non-standard* experiments.

Flagging (F) The online calibration system uses the results of many monitoring devices to indicate periods when data should be flagged; for example when an antenna pointing or focus setting is offset more than a specified amount from the calculated settings or when relevant local oscillators are not locked. These flags are stored in the SB archive data base in two ways, xml and binary flags. One of the initial offline calibration steps are to apply these flags.

Online vs Offline Calibrations?: Some calibration steps that could be applied online are applied during the initial offline reduction step using the appropriate data tables written by the online system. One example the WVR delay calibration. The main reason for postponing the online corrections is that occasionally the relevant data are not complete or partially wrong, so must be modified before incorporation of the relevant calibrations.

Bandpass Calibration (B): The remaining frequency response associated with the SB will not be flat in amplitude and phase, although the delay errors that produce a large slope of phase with frequency should have been removed by the a priori and online calibrations. For each SB and tuning, the bandpass is determined from a scan of a bright calibrator source with a *known radio spectral index* in order to not introduce an amplitude slope in the bandpass or between SPWs. After removing any significant phase variations over the bandpass scan (see next item), the casapy task *bandpass* determines the amplitude and phase for each antenna/SPW/pol as a function of frequency, and places the results in a table. The length of the scan and frequency resolution needed to reach the required bandpass accuracy is discussed below in the sensitivity section. The stability of the bandpass with time and with angle is extremely good⁶, so that one bandpass determination per tuning per SB (about an hour) will not change by more than 0.4%, so that this bandpass can be applied to all of the data in the SB.

Electronic Gain Calibration (G): Virtually all ALMA observations use the phase referencing technique to calibrate the target source visibility data. The largest temporal changes, especially in phase, should have been removed with incorporation of the a priori and online calibrations, discussed above.

The casapy task *gaincal* is used to determine the antenna-based amplitude and phases for the phase calibrator scans. If the calibrator is relatively strong, antenna-solutions will be generated for each SPW and polarization. For a weak phase calibrator, the bandpass results can be applied first and then the combined SPW/pol antenna solutions for the phase calibrator can be determined.

The plots of the results should be carefully inspected and some guidelines are: (1) The antenna-based amplitudes should not vary by more than about 10% over the observation period, and differ by more than about 20% among the antennas. For antennas with low gain or obvious drop-outs, check on the data quality, Tsys application, and the observing logs for any mention of a problem with the antenna. (2) The interpolation of the antenna-based phase determinations between calibrator scans should be obvious, since this is the phase that will be applied to the target, even for the longest baselines. The use of the *Go/Nogo* observational check in Cycle 3 should avoid observations with unacceptably large phase variations that would give poor image quality. Data flagging should be made for antennas and time ranges as appropriate.

Absolute Flux Density Calibration (T): The Calibration of the data using the Tsys measurements converts the raw correlation amplitude into Kelvin. The conversion of the correlation units of Kelvin to Jy is made using the system equivalent flux density (SEFD) that is measured for each antenna. For example, the 12-m antennas at Bands 3, 6 and 7 have an SEFD of about 35 K/Jy. This a priori conversion is accuracy to about 5% at Band 3, increasing to 15% at Band 9.

In order to derive a more accurate absolute flux calibration scale, a measurement of a source with known flux density and structure is included in the SB. It is possible derive reliable Kelvin to Jy conversion by using any of the 10 solar system object for which very accurate visibility models are available. Also about 30 bright quasars (grid sources) at Bands 3 and 7 are monitored by ALMA every few weeks, and the derived flux densities (as interpolated in time and frequency) also provide the K to Jy conversion of the visibility.

Polarization Calibration (p,X): For Cycle 3, the polarization calibration is made on a bright polarized quasar that is polarized by at least 3%. The observation period needed for the calibration is about three hours during which the calibrator rotates at least 90° in parallactic angle. Two steps are needed to obtain calibrated cross-hand (XY,YX) data using the polarization calibrator. First, after the XX and YY data are calibrated (bandpass and gaincal), the XY and YX frequency dependence for each SPW will not be zero, but reflect the (X-Y) polarization delay and phase difference. This phase frequency dependence is the same for all antennas, but with opposite phases in the XY and YX data. These global phase-frequency terms are then removed.

The second calibration setup determines the *leakage* between the X and Y polarizations (called D-terms) for each antenna and SPW There are two leakage terms: one related to the non-orthogonality between the X and Y feeds, and one related to the orientation offset for the feeds. Both of these calibrations are done in the offline script and then applied to the entire data set in order to calibrate the polarization properties for all of the other sources. Because the XY,YX antenna beam sensitivity is considerably more complicated

⁶See http://almascience.org/documents-and-tools/alma-technical-notes/ALMATechnicalNotes15_FINAL.pdf

than that for the XX and YY polarizations, only sources with an angular extent less than 1/3 full-width half-power can be accurately imaged in linear polarization.

10.6 Appendix

10.6.1 ALMA Catalog Description

The ALMA catalog of sources contains information for about 10,000 calibrators south of $\delta = 45^\circ$. About 1100 are high quality calibrators with a positional accuracy $< 0.002''$, and flux density at 100 GHz greater than 30 mJy. The catalog is accessed in (<https://almascience.org/sc/>). The appropriate bandpass calibrator and phase calibrators for an observation are automatically chosen when the observing scheduling block is prepared, but these calibrators can be selected by the PI when the proposal is submitted or during discussions with the ALMA liaison before the observations.

On the average, the separation between a target and the 1100 good quality calibrator in the catalog is about 4° . A weak calibrator survey is now in progress at ALMA in order to find additional calibrators that are sufficiently compact and strong. The goal is to double the number of available calibrators for the start of Cycle 3 observations.

10.6.2 Phase Calibrator Sensitivity

The phase calibrator chosen for an SB is picked by the calibrator query (See Chapter 8) when the observing schedule is prepared about a month before the nominal observation date. In most cases, the optimum calibrator can be determined months ahead of time. If its flux density is less than about 1.3 times the lower flux density limit, there is a slight possibility its strength may fall below the sensitivity limit at the observation time. The flux limits should be scaled by $(N_{ant}/36)^{-0.5} (T_{sys}) [(t_{scan}/120) (\nu_{BW}/7.5GHz)]^{-0.5}$

The lower flux density limits for a phase calibrator for all of the observing bands are given in Table 10.1 below.

Flux Density Limit for Phase Calibrator				
Band	Frequency (GHz)	Assumed T_{sys} (K)	Assumed PWV (mm)	Flux Density (mJy)
3	86.243	83	5.186	4.6
3	115.271	173	5.186	9.7
4	146.969	86	1.796	4.9
6	230.538	102	1.262	6.0
7	345.796	151	0.658	9.5
8	461.041	431	0.472	30.0
8	492.161	638	0.472	46.0
9	624.208	4302	0.472	364.4
9	658.007	1721	0.472	152.9
9	691.473	1231	0.472	114.9
10	806.652	2405	0.472	271.3

Table 10.1: The minimum phase calibrator flux density in mJy needed for an antenna based phase determination with 4 deg rms (15 SNR) with a 120 seconds integration time, assuming 36 antennas and elevation of 50 degrees. The total bandwidth in both polarizations of 7.5 GHz is assumed.

The scaling of the flux density varies as T_{sys} , elevation, bandwidth, number of antennas in the usual way.

The recommended phase calibrator is the one closest to the target that meets the minimum sensitivity limit of an $\text{SNR} = 15$ per antenna solution. Since the nominal scan to scan atmospheric phase changes are larger than 5σ , using a stronger but more distant calibrator is not recommended.

10.6.3 Bandpass Calibrator Sensitivity

The bandpass calibrator must be significantly stronger than the phase calibrator since the frequency dependence versus channel for each SPW and polarization separately must be determined. However, the bandpass calibrator does not need to be close to the target (within $< 45^\circ$) so that a sufficiently bright source should be normally available. The main sensitivity concern is obtaining sufficient SNR for FDM channels. Testing has shown that bandwidths < 31 MHz, the bandpass is flat to 1%; so bandpass calibration over narrower bandwidths is not necessary.

The minimum flux density for a calibrator to meet the 50 SNR specification for each antenna gain determination (2% rms in amplitude for each of the observing bands are given in the table below. The flux limits should be scaled by $(N_{ant}/36)^{-0.5} (T_{sys}) [(t_{scan}/900) (\nu_{BW}/128\text{MHz})]^{-0.5}$

Flux Density Limit for Bandpass Calibrator				
Band	Frequency (GHz)	Assumed T_{sys} (K)	Assumed PWV (mm)	Flux Density (mJy)
3	86.243	66	5.186	60.7
3	115.271	173	5.186	128.0
4	146.969	86	1.796	64.1
6	230.538	102	1.262	78.8
7	345.796	151	0.658	125.2
8	461.041	431	0.472	395.7
8	492.161	638	0.472	605.8
9	624.208	4302	0.472	4801.4
9	658.007	1721	0.472	2014.5
9	691.473	1231	0.472	1514.2
10	806.652	2405	0.472	3575.0
10	850.000	1817	0.472	2922.5

Table 10.2: The minimum bandpass calibrator flux density in mJy needed for an antenna based gain solution with SNR of 50, assuming integration time of 900 seconds, 36 antennas, and elevation of 50 degrees. The resolution bandwidth is taken to be 128 MHz.

10.6.4 Band-to-Band Phase Calibration

The closer a calibrator is to the target, the more accurate will be the phase-referenced the image. For imaging at the higher frequencies (> 400 GHz), the technique of band-to-band phase transfer has been developed in order to measure the phase of the calibrator in a lower ALMA Band, then suitably scale the phase solutions to the higher frequency target. This calibration is effective if the low frequency calibrator is significantly closer than the best high frequency phase calibrator (See Figure 10.1). This calibration method should be available in Cycle 3 for high frequency (Bands 8, 9 and 10) observations and will be discussed with PIs during the generation of the SBs at Phase 2. Details are given in http://almascience.org/documents-and-tools/alma-technical-notes/ALMATechnicalNotes8_FINAL.pdf.

10.6.5 Bandwidth Phase Transfer

When observing a target within a narrow SPW bandwidth, a close phase calibrator that is detectable with this narrow bandwidth is often not available. In this technique, a strong calibrator (not necessarily near the

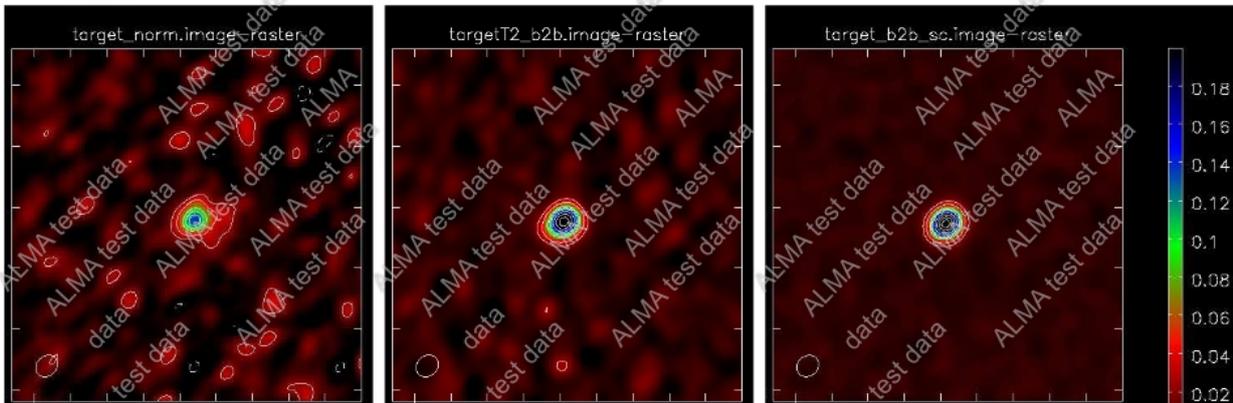


Figure 10.1: Band-to-band imaging of target source J1743-0350 at Band 9. *Left:* The target image calibrated using a strong source 14° away. *Middle:* The target image calibrated using a weaker calibrator 4° away with the solutions derived in Band 6. *Right:* The actual target image after self-calibration. The target peak flux densities from left to right are 0.17, 0.27, and 0.31 Jy. The rms noise on the images are 7.2, 3.6, and 1.6 mJy/beam. Credit: Ed Fomalont, Violette Impellizzeri (NRAO/JAO)

target) is observed at both a wide and narrow bandwidth in order to determine the amplitude ratio and the phase difference between the two bandwidths. The phase calibrator is then observed only with the wide bandwidth, and the wide band to narrow band amplitude and phase corrections are applied in order to calibrate the target observed only with the narrow bandwidth. This calibration method should be available in Cycle 3 for high frequency (Bands 8, 9 and 10) observations and will be discussed with PIs during the generation of the SBs at Phase 2. Details are given in http://almascience.org/documents-and-tools/alma-technical-notes/ALMATechnicalNotes11_FINAL.pdf.

10.6.6 Astrometry

The goal of astrometric experiments are to determine the accurate position of a target source. If the position accuracy needed is $> 0.05''$, then normal phase referencing is sufficient if the calibrator position accuracy is $< 0.01''$ and it is within about 5° of the target. Band 7 is probably the best band to determine the accurate position for most thermal objects.

For position accuracy to about $< 0.002''$ rms, it is recommended to include several check sources (secondary calibrators) in addition to the target and main phase calibrator. A scheduling block is available, and the PI should choose the main calibrator and several secondary calibrators. For more information, please contact the ALMA Helpdesk (<https://help.almascience.org/>)

10.6.7 Long Baseline Observations with Baselines > 3 km

The calibration of long baseline observations is similar to that at the shorter spacings. Several considerations are:

The possibility of self-calibration of the target is more important for long baselines, because the phase variations will be larger. Thus, the image quality with phase referencing alone may not reach the intended signal-to-noise and image fidelity. To determine if self-calibration is possible, the correlated flux density of the target at the longer spacing must be comparable to that needed to detect a phase calibrator.

Does the source have sufficiently small-scale structure for long baseline observations? For many thermal sources, this requires an emission temperature of a few hundred degrees or more.

The choice of a close calibrator to the target is important in order to removed systematic phase errors at the longer baselines. If a calibrator is not available within about 4° of the target, requesting ALMA to search for a closer calibrator, using the weak calibrator survey program, is recommended.

If the source is extended more than $2''$ at band 3 (scaling inversely with frequency), then the long baseline configuration will not be able to include all of the emission. **For Cycle 3, an additional configuration must be proposed as a second science goal in the proposal.**

Chapter 11

Quality Assurance

The goal of ALMA Quality Assurance (QA) is to ensure that a reliable final data product is delivered to the PI, that is, the product has reached the desired control parameters outlined in the science goals, it is calibrated to the desired accuracy and it is free of calibration or imaging artefacts. For Cycle 3, there are several restrictions that may make it a bit more difficult to achieve this goal. It has therefore been decided that QA will still be done on a “best effort” basis, covering all the major issues affecting data quality (see below). 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 tackle issues related to the merging of data within each science goal taken with different configurations, the inclusion of 7-m Array and TP Array data, and 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: Monitoring of calibrations and overall performance during observations

QA1: Measurement of performance parameters and telescope properties by the observatory

QA2: Full Calibration and generation of science products

QA3: Issues found with the data by the PI or the ALMA contact scientist after data delivery

The QA0, QA1 and QA2 stages will be handled by the Program Management Group (PMG) and the Data Management Group (DMG) (with contributions from ARC personnel) using the ALMA Quality Assurance (AQUA) Tool (see Section 11.6). Responsibility for data quality assurance in Chile rests with the Data Manager and his Deputy, within the Department of Science Operations, drawing upon the resources of the Program Management Group and the Data Management Group. The final output of the ALMA QA0-QA2 process is a “QA Report” per ObsUnitSet (Member or Group) that summarizes all the relevant QA information for each of the different QA stages up to, and including, the final imaging. This report will be included in the data package delivered to the PI. The QA3 stage will be handled separately, by the ARCs, via JIRA¹ tickets created by the ARC personnel (see below). A more detailed description of the different stages of QA is given below.

¹JIRA is a proprietary issue tracking product, commonly used for bug tracking, issue tracking, and project management

11.1 Cycle 3 "Best Efforts"

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, during ALMA Early Science (including Cycle 3) the Observations are subject to "best efforts", meaning that the observatory will attempt to meet the sensitivity and resolution stated by PIs in their Science Goals, within certain tolerances as described later in this chapter. Each project component ("Schedule Block" or SB) is scheduled for the number of executions that are expected to reach these goals (based on nominal Array performance values), but the resulting products may ultimately fall short of the requested goals, due to limits imposed by the available configurations or system performance. In particular, observers may need to invest their own time and expertise to ensure that the data products are of the appropriate quality and to re-reduce the raw data if the quality is not satisfactory. This may include the need to visit the relevant ARC or ARC node to get help and to assist with quality assurance and potential data reduction.

11.2 QA0

QA0 is a near-real-time verification of data quality. It deals with rapidly varying performance parameters (on timescales of an SB execution length or shorter) and thus has to be performed at the time of data taking. Assessment is performed by AoDs (Astronomers on Duty) at the OSF, with support from the Science Archive Content Managers, using the AQUA software tool, based on semi-real time output of the calibrations (obtained by the real-time TelCal ALMA software) as displayed by QuickLook and the "Calibration Summary" files that are produced at the end of each SB observation or sequence of SB repeats. 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 back-ends. These parameters can be grouped into the following categories:

Atmospheric Effects: Weather Parameters, Sky Opacity, System Temperature, Phase Fluctuations, Total Power Levels, WVR Outputs.

Antenna Issues: Antenna Gain, Relative/Offset Pointing, Focus, Antenna Tracking, Geometric Shadowing, Nutators.

Front-End Issues: RF Bandpass, Sideband Ratios, Receiver Temperatures, LO Lock Status.

Connectivity Issues: Total Power levels, Delay Measurements, System Temperatures, RF Bandpass, LO Lock Status.

Back-End Issues: Total Power levels, RF Bandpass, Delay Measurements.

The tolerances for these parameters that have been adopted by ALMA for this Cycle are listed in Section 11.7. Apart from these, the metadata is checked for inconsistencies and the amplitude calibrator for recent flux measurements. Each SB execution is classified into three categories, i.e. QA0_PASS, QA0_SEMIPASS and QA0_FAIL, based on the criteria described in that Section. QA0_PASS are datasets that comply with all the QA0 criteria and that will be used for the final imaging. QA0_SEMIPASS are datasets that do not fulfill all the QA0 criteria, but contain data that is deemed of scientific value. QA0_SEMIPASS data are not included in the final images, but PIs can access those data and reduce them if they wish. QA0_FAIL datasets are those that are not included in the other two categories. Since they represent unusable/uncalibratable data, they are not visible to the PIs in the Project Tracker.

For operational purposes, individual SB executions during Cycle 3 will be counted as a normalized fraction of the theoretical sensitivity that a given execution should have reached based on the number of antennas that should be in an array in Cycle 3 (i.e., 36), and the average system temperatures expected at the frequency of the observations. Executions with *higher* rms noise levels than the reference will have fractional execution

weights between 0 and 1, while those with lower rms noise levels, which can be achieved, for instance, with a larger number of antennas in the array, or with lower system temperatures, will have a fractional execution value above 1.

11.3 QA1

QA1 tracks array performance parameters which vary slowly (on timescales longer than a week). They are measured by AoDs and System Astronomers executing standard calibration SBs created as specified by the Calibration Plan. The QA1 related parameters will, in general, be measured at predefined periods during the month as “Observatory Tasks”, or if significant deterioration of performance is detected during operations. Currently, the various tasks to measure these parameters are done by different software packages. This situation will change in the near future by including some of the packages within TelCal and/or CASA. Reduction of “Observatory Tasks” output is done jointly by the AoDs and System Astronomers within the DMG. The product is a set of parameters with errors that are ingested into the TMCDB (up-to-date view of the parameters used in any observation), so that they can then be used during observations.

The tasks that fall into this category are:

Array Calibrations: Baseline measurements, Delays

Antenna Calibrations: All-sky pointing, Focus curves, Surface measurements, Beam patterns (including polarization observations), Relative delays between polarizations of same band

Source Calibrations: Monitoring of solar-system flux standards, and secondary quasar flux standards

11.4 QA2

QA2 deals with QA at the level of data reduction and imaging using the Science Pipeline or performed interactively by the ALMA Data Reducers Team. It is only at the stage of data reduction and analysis of data products that some of the science goals set by the PI can be compared with the results (i.e., rms, angular resolution, SNR, dynamic range, etc).

During Cycle 3, it is expected that, for the ALMA standard observing modes, the automated Pipeline (see Chapter 13) will be used for the calibration of the data and for some basic imaging. Additional imaging will be done by data reducers that will then perform the QA2 assessment. The Pipeline provides a Weblog page with detailed information on the quality of the Calibration, which will provide the basis for Pipeline QA2. For manually-reduced datasets, and the imaging of pipeline-calibrated data, the QA2 will be carried out with special purpose scripts whose outputs will be included in the data packages delivered to PIs. The QA2 metrics which determine the success of an observation are given in Section 11.7.

A summary list of QA2 parameters/issues checked during data reduction (calibration of individual ExeBlocks and joint imaging) are:

Calibration Issues:

- Flux scale calibration: is the absolute flux scale accurate enough?
- Phase transfer and astrometry: Cycle time and sky separation between phase reference and target; typical and extreme (unflagged) phase differences between phase reference scans.
- Calibration consistency if multiple arrays and/or single-dish data are combined.

Final Data Characterization:

- Longest baseline, visibility coverage and time on target (after flagging).
- Synthesised beam (spatial resolution) for specified weighting scheme; specific imaging requirements (if stated by the PI).
- Spectral resolution and channel ranges used to make sample images.
- rms noise in target images: standard (all-data) imaging; in images of specified spectral regions. These values are compared with those predicted from data after flagging and with those requested.
- Residual imaging artifacts: Sidelobe Levels, effects of “missing spacings” and possible dynamic range limitations.
- Combinations of array configurations and/or total power, e.g. amplitude scale consistency.
- Mosaicing and/or contamination by bright sources outside FOV or aliasing in clean.
- Polarization purity: minimum fractional polarization detectable, and polarization angle accuracy (if relevant).

The PI will also receive any modifications to standard scripts which were found to be essential during QA2.

There are three possible QA2 states a reduced dataset can be placed into. QA2_PASS implies that the scientific goals, as defined by measurable parameters such as noise RMS, LAS and angular resolution, have been achieved within a 10% of the specification. QA2_SEMIPASS refers to those datasets that fall short of meeting the PI requested science goals, but are otherwise of good quality. QA2_FAIL is a temporary state used during an observing cycle when an observation fails to meet the PIs goals by a significant margin (30% worse) and needs to be reschedule for additional observations. If these data fail to obtain enough additional observations to pass QA2 by the end of an observing season, their QA2 state is changed to QA2_SEMIPASS and they are delivered to the PI.

11.5 QA3

QA3 is post-reduction evaluation of the data products delivered to the PIs. It is advisable that PIs check the data products themselves and report any problems that they find to their Contact Scientist via the ALMA Helpdesk (<https://help.almascience.org/>). The QA3 process will be triggered by PI, the contact scientist of other ARC personnel. They will open an ALMA Helpdesk ticket reporting a problem with the data products which may reflect an underlying problem with the data, observing procedure or calibration. The ARC receiving the Helpdesk ticket will retrieve the data from the archive and evaluate the nature of the problem. The evaluation by the ARC should include an assessment on whether the problem is present only in a particular dataset or whether others taken under similar set-ups and conditions also show it. If the problem is deemed to reflect 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 observatory, which 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 with QA3 issues will be granted based on the policies in the Cycle 3 User Policies document.

11.6 The Quality Assurance Report

During Cycle 3 the Quality Assurance Reports of QA0 and QA1 will be accessible to the PIs using the Project Tracker. These reports will be generated using the AQUA tool. For the QA2 reports, for Pipeline calibrated datasets, the PI will receive, as part of the data delivery, the Weblogs, and the outputs of scripts that measure the critical QA2 metrics. For non-standard observing modes that require manual data reduction, script-generated QA2 reports will be delivered instead.

The basic unit of a Report is the ObsUnitSet, which represents a scientific goal stated by the PI during Phase 1 (project creation and review). An ObsUnitSet will typically contain several executions of SBs. For

each execution QA0 and QA1 reports are generated by the AoDs using the information available at the time of the observations, which includes TelCal outputs and other monitoring data (weather, total power levels, Corr GUI outputs, etc). A given execution is only cleared for reduction if it has passed both QA0 and QA1. There will be only one QA2 report for the whole ObsUnitSet generated by the DMG/ARCs at the end of the data reduction process; this also has to be approved before the data products are delivered to the PI. It is expected that sometime during Cycle 3, an interface between the Pipeline and AQUA is made available. From that point onwards, complete QA Reports will be generated using the AQUA software and delivered as such. The final report per ObsUnitSet delivered to the PI would in that case be a concatenation of all the relevant QA0 and QA1 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 and QA1 failures are that the observations of those ExecBlocks that failed have to be repeated. Failures to pass QA2 may trigger additional observations if the achieved RMS or imaging parameters are quite far from what has been requested by the PIs. Additional observations may not be possible in exceptional circumstances, such as projects with very tight weather or time constraints. If the available data are insufficient to reach the required sensitivity, but are otherwise of an acceptable standard, they may be released to the PI.

11.7 QA Criteria

For an execution to be considered QA0_PASS it must have a phase RMS (post-calibration) of ≤ 0.5 rad, enough calibrator data to be able to calculate necessary calibration terms (bandpass, amplitude, complex gain, and, if relevant, polarization), and enough science data to be useful. Furthermore, the execution has to be "significant" in terms of the fractional execution (i.e., more than 0.2). If it does not meet these criteria, but has some useful data (calibrator or source), or has a fractional execution that is less than 0.2, it is classified as QA0_SemiPass. This makes the data available through the ALMA archive, but it will not be used in the generation of the imaging data products. If neither of these hold, the execution is deemed to have no useful data and declared QA0_FAIL.

The QA0 pass/semipass/fail criteria that have been adopted by ALMA during Cycle 3 are based on the following:

- **Antennas:** QA0_SEMIPASS for fewer than 17 antennas available in the 12-m Array or fewer than 5 antennas in the 7-m Array or fewer than 2 Total Power Array antennas. States that will be considered to render an antenna unavailable include issues with the antenna itself (including large pointing scatter of $> 1/10$ HPBW, and bad band focusing with offsets from model of $> 1/5\lambda$ in Z and $> \lambda$ in X/Y), as well as issues with the receiver bands not being available at the antenna for a given observation.
- **Bandpass:** QA0_SEMIPASS if bandpass calibrator signal too weak, with amplitude wiggles across the spectral window of > 3 dB on the autocorrelations (not due to atmospheric features) or strong CW signals, and the phase/gain calibrator cannot be used instead.
- **System Temperatures:** QA0_SEMIPASS if more than 9 antennas in the 12-m Array, or more than 3 antennas in the 7-m Array, or more than 1 Total Power antenna with T_{sys} values $> 50\%$ of the others, with differences between polarizations $> 50\%$ for a given antenna or with $T_{\text{sys}} > 2000$ K. QA0_FAIL in case more than 15 antennas in the 12-m Array, more than 5 antennas in the 7-m Array, or more than 2 antennas in the TP Array have T_{sys} in the ranges specified.
- **Gain:** QA0_SEMIPASS Phase RMS > 0.5 rad, on scans of the gain calibrator, for a significant fraction of the baselines after WVR phase correction (using all possible bandwidth in a given spw)
- **Execution:** QA0_SEMIPASS All datasets whose execution failed at some point and contain less than 20% of the expected execution time of a given SB execution.
- **Calibrations:** QA0_FAIL Datasets missing critical calibrations that render them uncalibratable

- **Storage:** QA0_FAIL Data that could not be read from the Archive

For any other situation, the data will be accepted, although it may require some additional flagging for misbehaving antennas, baselines, etc. Any problems with QA1 that would significantly downgrade the quality of the data will be solved by the observatory by stopping the science observations and re-calibrating the problematic parameters of the array.

For QA2, the main criteria are the achievement of the requested noise RMS in the images (it must be within 10%, 15% and 20% of the goal for bands 3/4/6, 7/8, and 9, respectively), the synthesized beam shape (i.e., 20% of goal) and the calibration quality (phase RMS and absolute accuracy).

Chapter 12

Logical Data Structure and Data Flow

This chapter describes the the data flow process from the observations until 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 metadata 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. A summary of the main actors and operations involved in the observations is shown in Figure 12.1.

Each of the gray 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.

A typical observing session would be started by the Telescope Operator interacting with the Executive subsystem via one of the dedicated control computers in the Control Room (the so-called “Standard Test Environment” or STE). 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. This is done in several cycles to solve interdependencies between the different software components. 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 display of the Control, Telescope Calibration and QuickLook subsystems.

The actual observations start by manually creating an array, which means selecting all the antennas that will be involved in the observations. A Scheduling Block (SB) is then selected from the list provided by the Scheduler subsystem, and the execution is started. All the SBs for a given observing Cycle are stored in the Archive after successful Phase 2 completion. The Scheduler subsystem keeps a local up-to-date database of all the SBs (including their status) for that Cycle. The Telescope Operator and Astronomer on Duty (AoD) have two options for using the Scheduler during Early Science observations. They can either let the Scheduler suggest possible SBs to execute, or they can carry out targeted searches of the local database. For Early Science, the Scheduler can produce a ranked list of optimum SBs to execute next based on weather conditions and forecast, hardware and configuration status, project completion status, representative source position on sky, proposal rank and score and Executive percentages. The Telescope Operator and AoD can follow the suggestion of the Scheduler and select one of the top-ranked SBs or something else for execution. Selections that do not follow the advice of the Scheduler must be fully justified by the AoD and will be used to improve the selection algorithm within the Scheduler. It is expected that by Full Operations the algorithm will be optimized to the point that the Scheduler can run an automatic sequence of SBs which would have been manually selected by an AoD.

Once the execution of the SB has been selected from the Scheduler, it is dispatched to the Control subsystem.

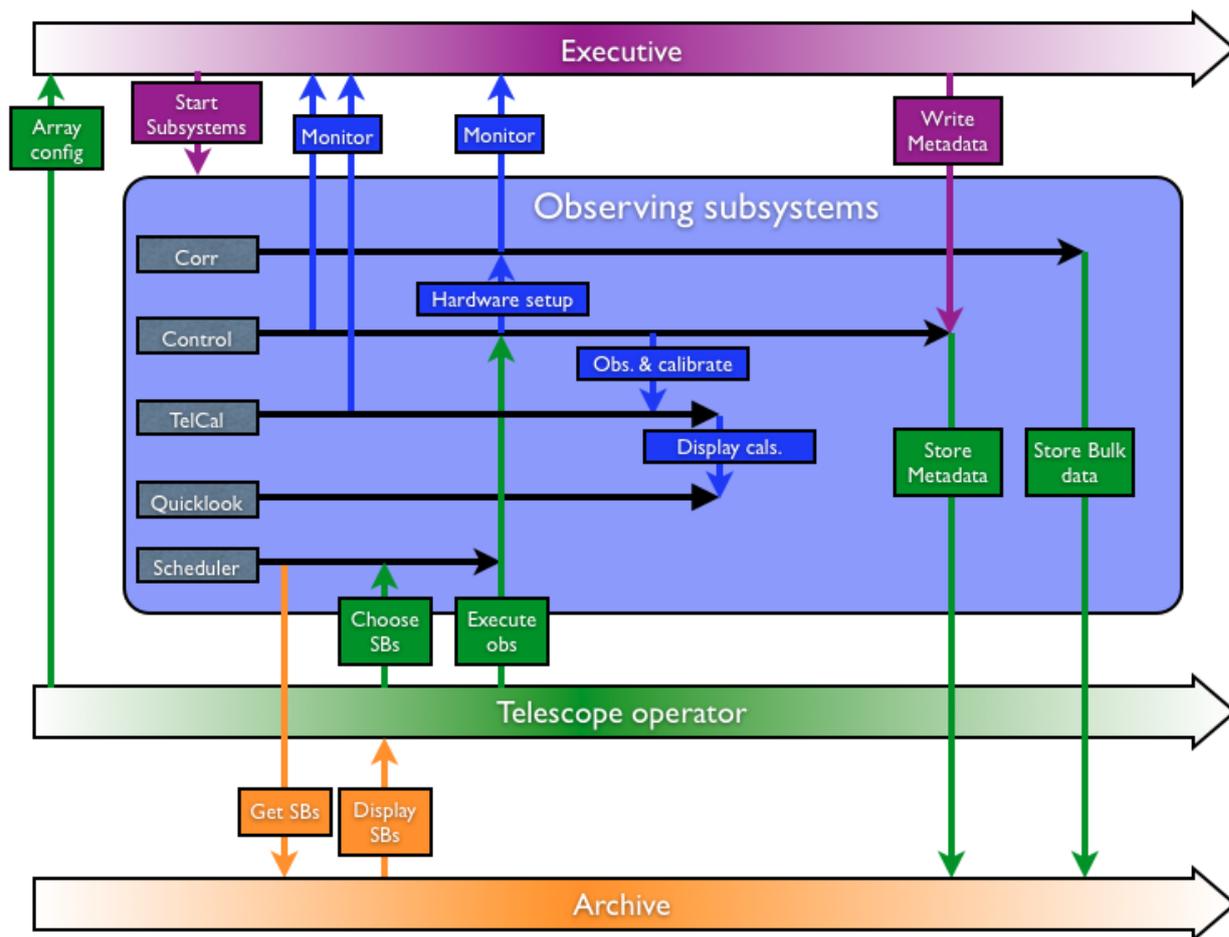


Figure 12.1: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

Control parses the SB XML document into python objects, and executes the observing script selected in the SB, which reside in the Science Observing Scripts module. The observing script uses the information from the SB as input to determine the necessary execution sequence, which is commanded to Control as 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), resulting in raw data and metadata being sent to the archive subsystem and made available to the online Telescope Calibration (TelCal) subsystem. TelCal publishes results from calibration scans it reduces, which are sent to the archive as calibration tables, and received by the QuickLook GUI and other software to display results to the Operators and AoD to evaluate the observation progress. To carry out its function, Control has many interfaces to the instrumental hardware. It is in fact one of the truly real-time subsystems within ALMA because it is in charge of synchronization of the actions of all antennas (scanning, source acquisition, etc) and the correlator to within 48ms (Timing Event (TE)). Control is also in charge of storage of data from all monitor points set in the hardware of the ALMA array.

Each run of an SB produces an Execution Block (EB) that is stored into the archive through two parallel paths and is handled by a part of the Control software called the “Data Capture” module. The best way to describe the Data Capture module is as an interface between the real-time domain of the data taking and the

storage side. As an interface, it captures and stores into the Archive all the relevant meta-data information pertaining to a complete description of the data and their supporting calibration and monitor datasets and condenses that information into a set of XML tables. The contents of all these tables are defined in the ALMA Science Data Model (ASDM). Together with these tables, Data Capture also creates the relevant links of these metadata to the actual bulk data that is directly stored into the archive. Furthermore, it also provides calibration data to the TelCal and QuickLook subsystems for calibration reduction and display in semi-real time. Finally, when the SB is finished, Data Capture is in charge of checking that all products representing the raw data have been produced and stored in the archive, and announces the completion of the SB to the Scheduler subsystem. It is clear from the list of roles above that Data Capture is a very complex module; it has to handle Correlator and backend data in addition to supporting (source information, spectral set-up, etc), and monitoring data which are needed for data reduction (weather, pointing, etc). All this information originates in different hardware and software elements, each of which can be sampling at different rates and with limited view of the behavior and state of the observing system.

A summary plot of the main elements involved in data flow is shown in Figure 12.2. As indicated in the figure, all components passing through the Data Capture module contribute in the generation of the metadata associated to a dataset, which is in XML format. The bulk of the data (binary) is sent to the Archive and linked to the metadata using requests to the Archive.

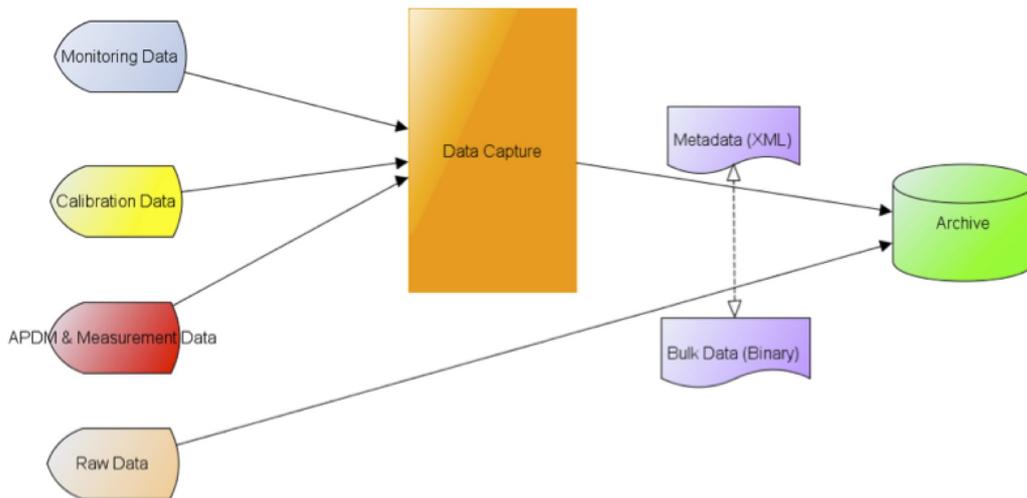


Figure 12.2: Data flow components.

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. As described above it contains both bulk binary data and metadata (XML) organized in tables. The tables contain links to other XML tables and addresses pointing to the actual bulk data 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 of the calibrations that can be done online. 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, auxiliary monitoring data, overall project and post-processing. A list of the core tables is shown in Figure 12.3. The term *referenced* in the tables means that a given table is linked to other tables via some references, that is, some of the information is shared between tables.

SDM Tables		
Referenced:		
Main		
Antenna	Field	SpectralWindow
ConfigDescription	Pointing	State
DataDescription	PointingModel	Station
ExecBlock	Receiver	Subscan
Feed	Scan	SwitchCycle
Not referenced:		
AlmaRadiometer	Focus	SBSummary
Annotation	FocusModel	Source
CalDevice	FreqOffset	SourceParameter
DelayModel	GainTracking	SpaceCraftOrbit
Doppler	Holography	SysCal
Ephemeris	Polarization(<i>required in MS</i>)	WVMCal
Sometimes referenced:		
Beam	required for single dish or mosaicked data	
CorrelatorMode	required for correlators; not allowed for others	
SquareLawDetector	required for total power or noise detectors; not allowed for others	
Mandatory:		
SDM		

Figure 12.3: ASDM Tables. Outlined set of tables are the core ones (i.e., present in all ASDMs).

The associated tables produced by TelCal all have a name starting by “Cal” and then a self-explanatory string on the type of calibration they are associated with. The list of these associated tables is being upgraded as new observing modes and calibrations become available (see Chapter 10). The current list is shown in Figure 12.4. In the figure *CalDM* stands for Calibration Data Model, which has been implemented by TelCal.

CalDM Tables		
CalAmpli	CalFocus	CalPointingModel
CalAtmosphere	CalFocusModel	CalPosition
CalBandpass	CalGain	CalPrimaryBeam
CalCurve	CalHolography	CalSeeing
CalDelay	CalPhase	CalWVR
CalFlux	CalPointing	
CalData		CalReduction

Figure 12.4: Current list of ASDM associated tables generated by TelCal.

Most users interact with the ASDMs via the Common Astronomy Software Applications (CASA) data reduction package. At that point, the original ALMA data format has been converted to a format more convenient to CASA, the Measurement Set (MS). Information on the MS content is described in the CASA manuals and description documents available at <http://casa.nrao.edu>.

Chapter 13

Pipeline Processing

13.1 Introduction

The ALMA Science Pipeline is used for the automated processing of ALMA interferometric and single-dish data. It is released as part of CASA, where the current Pipeline tasks are denoted by the prefixes: h, hif, hifa, hifv, and hsd (see Table 13.1). The Pipeline is data-driven, such that different Pipeline processing can automatically occur for different ALMA Observing Modes and different data quality. The algorithms employed to influence how automated data processing proceeds are referred to as the Pipeline Heuristics, and the outcomes of the heuristics (e.g. whether to apply a calibration table or not) are stored in the Pipeline Context. At any given point during Pipeline processing, the Pipeline State is therefore recorded and known by the Context. The Pipeline has been designed to be highly flexible. It has been developed such that Pipeline tasks can be removed, added or rearranged from the processing. It is possible to override the calibration tables created by the Pipeline and introduce self-made calibration tables into the processing, by editing the Pipeline State/Context at the appropriate processing point. This enables "mix and match" of standard CASA tasks and Pipeline tasks.

The ALMA Pipeline has been used in Science Operations since October 2014, during Cycle 2, and continues to evolve to include new capabilities and improved methods. At present, the Pipeline can be used in the processing of ALMA interferometric 7-m and 12-m Arrays. Processing of single-dish (total-power) data is also possible, and the techniques will be optimised as Cycle 3 proceeds. Standard Pipeline processing recipes are used in Science Operations. In the future the Pipeline will also have the capability to combine the data taken using multiple ALMA arrays.

13.2 Pipeline In ALMA Operations

13.2.1 Pipeline Triggering

The Observing Unit Set (OUS) structure of an ALMA Project determines the number of times, when and how the Pipeline is triggered for that project and which data products are output. The Pipeline triggers at the completion of each ALMA OUS. Figures 13.1, 13.2 and 13.3 indicate the different number and types of Pipeline processing needed for 3 different Science Goals. The Pipeline is not commissioned to have more than one SB in an OUS.

13.2.2 Pipeline Execution

The Pipeline is executed in ALMA Operations via the Pipeline Processing Request (PPR). The PPR is an xml file generated by executing the pipeline infrastructure command `pipelineMakeRequest`. The pipeline infrastruc-

task	task type	description
h_	Common tasks	Pipeline tasks used in the calibration and imaging of both interferometry and single-dish datasets
hif_	Interferometry common tasks	Pipeline tasks used in the calibration and imaging of both ALMA and EVLA interferometry datasets
hifa_	Interferometry ALMA tasks	Pipeline tasks used in the calibration and imaging of ALMA interferometry datasets only
hifv_	Interferometry EVLA tasks	Pipeline tasks used in the calibration and imaging of EVLA interferometry datasets only
hsd_	Single-dish tasks	Pipeline tasks used in the calibration and imaging of single-dish datasets only

Table 13.1: Pipeline Task Types

ture is a set of commands for the automated running of the pipeline within the Observatory. As it is only needed for this internal purpose, the pipeline infrastructure is not part of publicly-released software. In the Observatory, the PPR is executed within CASA and by default it will process a completed ALMA OUS. During the execution of the Pipeline, a python script of Pipeline tasks, which is equivalent to the PPR, is written. This script can be used to perform Pipeline calibration of that dataset, and it is provided with each ALMA delivery of Pipeline processed data.

13.2.3 Pipeline Runtime Database-querying

The ALMA Pipeline can query the ALMA Calibrator Catalog at runtime to obtain the fluxes of regularly monitored quasars used as flux calibrators.

13.2.4 Pipeline and the ALMA Quality Assurance Process

Only data labeled QA0_PASS are processed by the ALMA Pipeline during ALMA Operations. The Pipeline infrastructure command pipelineMakeRequest by default picks up all the QA0_PASS data for a given OUS to include in the processing, but QA0_FAIL and QA0_SEMIPASS data will not be processed. This is because the Pipeline has a certain minimum requirement on the calibrators available to it during an Execution Block or an Observing Session and on the data quality. ALMA Execution Blocks and Sessions that do not meet these criteria are given QA0 labels such that these data do not enter the Pipeline.

The Pipeline performs the data processing need to perform ALMA Quality Assurance 2 (QA2), i.e. it calibrates the data and eventually will perform all imaging. This results in data products that can be delivered to PI and stored in the archive. However, these products will only be delivered and ingested into the Archive if they pass QA2. In order to determine whether a Pipeline-processed ALMA OUS passes QA2 or not, the Pipeline produces Quality Assessment (QA) scores which reflect the data quality and how well Pipeline processing has proceeded. The details of how these QA scores are determined are described in the ALMA Quickstart Guide. The Pipeline QA scores, and the algorithms used to determine them, will in future be made available to the ALMA Quality Assurance tool (AQUA) for to aid in assessment of QA2 status. The Pipeline QA scores are additionally displayed to staff and users via the Pipeline Weblog, a set of html pages that provides plots and information on the Pipeline processing of each OUS.

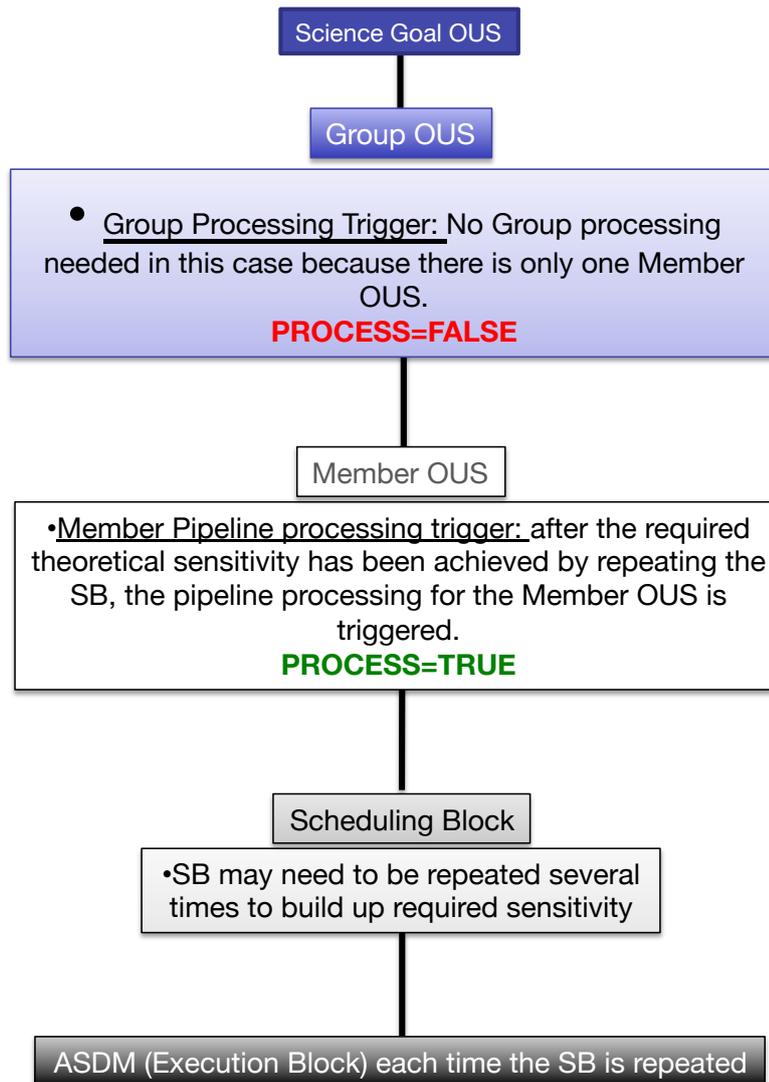


Figure 13.1: Pipeline processing of a Science Goal for single epoch, single array data. The same SB may need repeating N times to build up the requested sensitivity. This Science Goal would result in 1 Pipeline processing, with 1 set of science products delivered to PI.

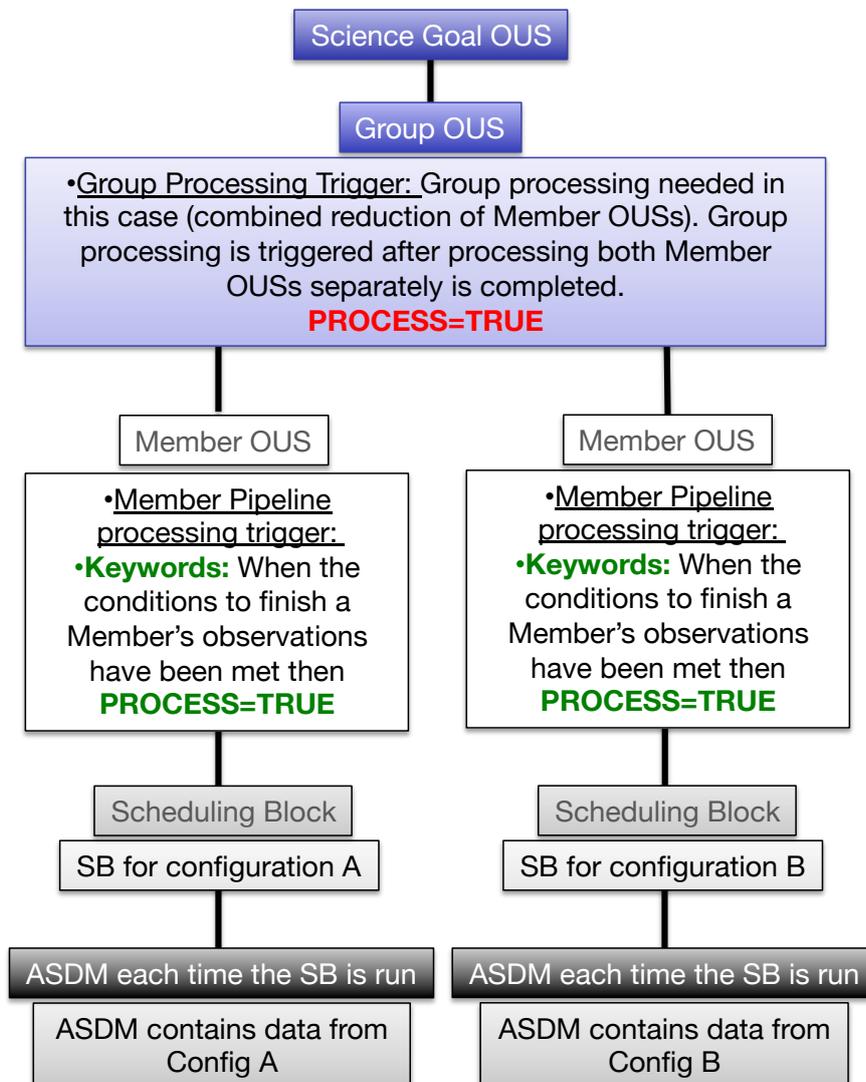


Figure 13.2: Pipeline processing of a Science Goal containing multiple ALMA configurations or arrays (e.g. 12-m, 7-m or TP). The data from the different configurations/arrays is processed independently in two separate Member OUSes, with science products created for each MOUS. Once the processing of the MOUSes is complete, then they will be processed together to form a combined science product from the two arrays (since they are part of the same Group OUS). Therefore this Science Goal would result in 3 Pipeline processings, with 3 sets of science products delivered to PI. (Note that the Pipeline is not yet commissioned for Group processing/array combination).

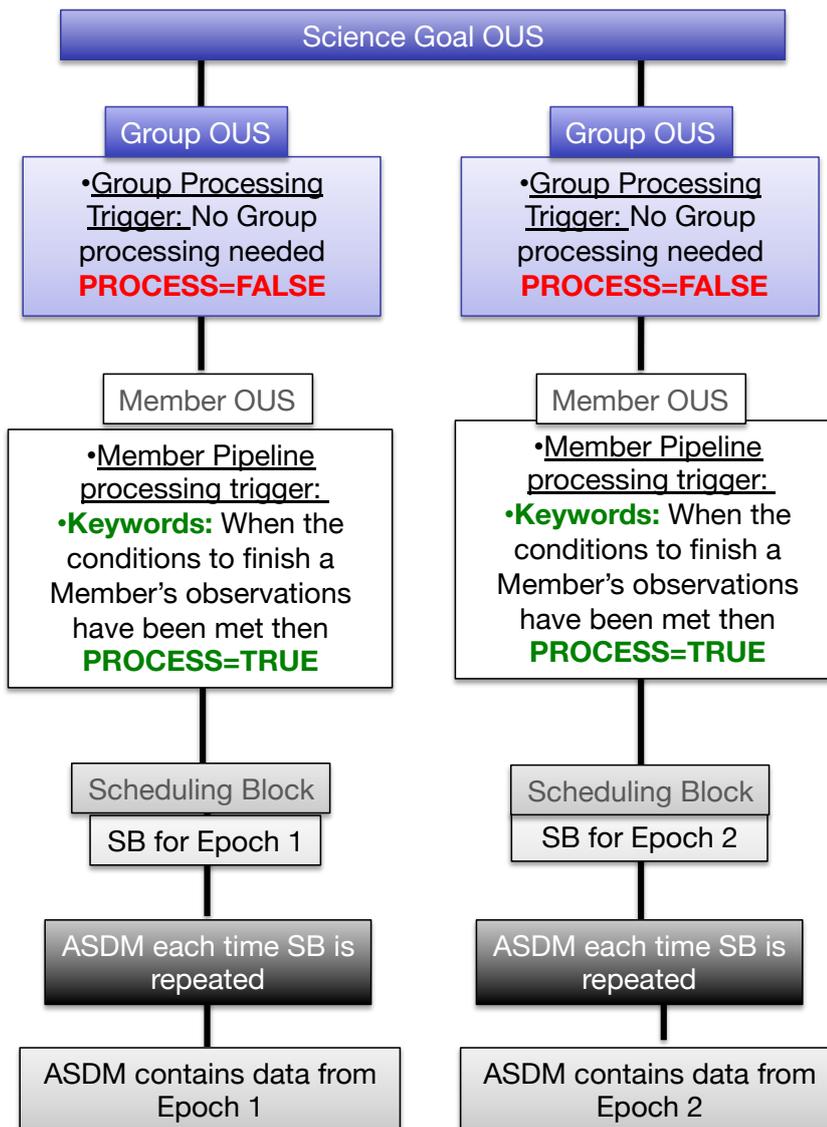


Figure 13.3: Pipeline processing of a Science Goal containing multi-epoch observations. Each epoch is processed independently by the Pipeline, with science products created for each epoch. Data from different epochs are not processed together to form a combined product (the Member OUS from each epoch are in different Group OUS). This Science Goal would result in 2 Pipeline processings, with 2 sets of science products delivered to PI.

13.2.5 Standard and non-standard Observing Modes

One of the definitions of whether an ALMA Observing Mode is labeled as a standard or a non-standard mode is whether the ALMA Pipeline has been commissioned to calibrate that Observing Mode in ALMA Operations or not. The standard and non-standard modes at Cycle 3 are listed in the Cycle 3 Call for Proposals.

13.3 Pipeline Processing of an ALMA OUS

13.3.1 Overview

An ALMA OUS may contain a number of different ALMA Execution Blocks. During data calibration, each execution block is calibrated independently unless the execution blocks form part of the same Observing Session. In this case only, one or more aspects of the calibration from one execution block may be applied to another. In the calibration processing, the Pipeline runs Execution Blocks sequentially through Pipeline tasks. For example, the Pipeline first imports the data from each Execution Block, it then performs deterministic flagging on each Execution Block, it then moves onto the next task etc. This proceeds until all the calibration steps have been performed for each Execution Block. At this point, all the calibration tables stored in the Context for each Execution Block are applied to each one using the Pipeline task `hif_applycal`. The calibration tables are only applied at the end point of the calibration, prior to that they have always been applied "on the fly" during Pipeline tasks. For imaging of the calibrators and/or science target(s) observed in an OUS, the data from all Execution Blocks is now imaged together for each target. During Cycle 3, some proportion of ALMA OUS will be calibrated using the Pipeline but imaged manually until Pipeline science target imaging is commissioned.

13.3.2 Pause after Data Import

In current practice, the ALMA Pipeline does not run straight through after triggering. The first time that the Pipeline is triggered, it will perform the import of each Execution Block of ALMA data only, with no calibration. The reason for this is multi-fold. Firstly, it is useful to check at this point the flux calibrator fluxes that have been obtained for any quasar flux calibrator. These values can be overridden if desired at this pause. Secondly, any fixes needed to the ALMA data prior to calibration can also be performed at this time. Thirdly, if the Pipeline has previously been run and it was noticed that the addition of extra manual flagging would improve the data reduction, the manual flagging can be added at this point.

13.3.3 Output of the Pipeline Processing

The main output of the ALMA Pipeline is calibrated visibility data and associated imaging products. Another important output is the Pipeline WebLog. Current lists of Pipeline products and a detailed description of Pipeline products at Cycle 3 are given in the document ALMA QA2 Products.

13.4 Pipeline Initialisation & Calibration Table Editing

The pipeline can be initialised in one of two ways: by creating a new pipeline state (`h_init`) or by loading a saved pipeline state (`h_resume`). `h_init` creates an empty pipeline context but does not load visibility data into the context (for e.g. interferometry datasets `hifa_importdata` can be used to load interferometry data). The Pipeline State is saved at any point with `h_save`. To view and edit the pipeline state/context then e.g. for interferometric processing `hif_show_calstate`, `hif_export_calstate` and `hif_import_calstate` can be used, see the Pipeline Reference Manual. It is these tasks that should be used in the case where it is wanted to replace a Pipeline-generated calibration table with a self-generated table.

13.5 General Interferometric ALMA Data Processing

13.5.1 Data Import & Calibrator Flux Retrieval

The Pipeline task `hifa_importdata` imports data into the interferometry pipeline employing the CASA task `importasdm`. It loads the specified visibility data into the Pipeline Context, unpacking and/or converting it as necessary. The visibility files may be ASDMs, tar files of ASDMs, measurement sets (MSs), or tar files of MSs, If ASDM files are specified, they will be converted to MS format.

The import can be set to apply binary data file flags. The data import stage also performs several checks on the data. One of these checks is which scan intents are present in the dataset. Scan intents (e.g. BANDPASS, PHASE) indicate to Pipeline throughout the processing the purpose of each scan observed during the OUS/session. Note that an incomplete set of calibration intents will not cause Pipeline to exit at this point. It will affect the QA score for this stage of the Pipeline however. See the Cycle 3 Pipeline Quickstart Guide for a full description of how the Cycle 3 QA scores are calculated.

Several files are created during the data import stage. For each ALMA execution block/ASDM a file of the flags provided by the online system will be created. The Pipeline will also create files to which additional manual flagging can be added to the Pipeline run. If the manual flag files are already present, then Pipeline will not overwrite them by default but will use the existing files in the directory. Own-named flag files can also be used, see the Pipeline Reference Manual information for `hifa_flagdata`. Whatever manual flagging needs to be added must be in place before `hifa_flagdata` is run. In ALMA Science Operations, this is run directly after the Pipeline "Pause After Import". In ALMA Operations, the Pipeline is run once without manual flagging added. The results of that run are examined by an ALMA Data Analyst, and if it is needed to add manual flagging, then the Pipeline is re-run with the manual flagging added before `hifa_flagdata` e.g. at the Pause after Import.

Note that, although the Pipeline can be started from MSes, it is always best practice to import data that is planned for Pipeline processing from ASDMs. This ensures that all the information needed by the Pipeline, such as which Session different ASDMs have been acquired from, has been correctly input into the Context.

13.5.2 Deterministic & Manual Flagging

The Pipeline applies several different types of flagging at this stage by default. These are:

- autocorrelations
- flags listed in the online flag files
- flags to data which have been affected by shadowing
- flagging to strip the dataset of data labelled with un-needed scan intents, such as POINTING, FOCUS, ATMOSPHERE, SIDEBAND RATIO
- flagging the edge channels of low-spectral resolution (TDM) data
- the manual flagging

13.5.3 Flagging lines in Solar System Object Flux Calibrators

The Pipeline contains a dictionary of the frequency ranges of lines that are commonly seen in Solar System Objects (SSO). It is possible to add lines to this dictionary via text file. If Pipeline detects these frequencies in an SSO flux calibrator, then it flags the lines unless the line frequency range exceeds a threshold fraction of the spectral window. If the threshold is exceeded, the Pipeline does not flag the line for technical reasons. Instead, it sets the parameter `refspwmap` such that this window will not be used in the later flux scale transfer to the bandpass and phase calibrator.

13.5.4 Reference Antenna Selection

The Pipeline creates a ranked list of reference antennas based on proximity to the centre of the array and degree of flagging.

13.5.5 Offline WVR Calibration for 12m Array Data

The Pipeline generates a gain table based on the Water Vapour Radiometer data in each ALMA Execution Block. It then applies the wvr calibration to the data specified by 'flag_intent', and calculates flagging views showing the ratio phase-rms with wvr/phase-rms without wvr. Thirdly, it searches the flagging views for antennas with anomalous high values. If any are found, then it recalculates the wvr calibration with the wvrflag parameter set to ignore their data and interpolate results from other antennas according to maxdist and minnumants. Fourth, if the overall QA score for the final wvr correction of a vis file is greater than the value in accept_threshold then it makes available the wvr calibration file for merging into the context and for use in the subsequent reduction.

13.5.6 Low gain flagging

Antennas with low or high relative amplitude gains are flagged by the Pipeline. The parameter threshold settings that determine which solutions are flagged are given in the Cycle 3 Pipeline Reference Manual. The CASA tasks employed during the heuristics for this step are bandpass and gaincal.

13.5.7 Bandpass calibration & flagging bandpass solutions

The Pipeline task hif_bandpass includes use of the CASA tasks gaincal and bandpass. Firstly a phase-only gaincal is run on a timescale shorter than the scan timescale (the scan timescale is the time spent on a source before switching to a different source). While it is common practice to perform this using a narrow frequency range of each spectral window only, in case of delays (slopes in phase vs frequency), at ALMA it is generally fine to include up to the whole frequency range in this gaincal. The solution table created here will only be applied on the fly when creating the bandpass solutions, other than that it is discarded. The reason to perform this “phase-up” of the bandpass calibrator is to fix any phase de-correlation. The CASA task bandpass is then run to create the bandpass solutions. Pre-averaging across several frequency channels can take place via the bandpass solint parameter; the amount of pre-averaging is calculated depending on the expected SNR of the bandpass calibrator (see low S/N heuristics section below). Flagging of the bandpass solutions may be performed.

13.5.8 Setting the fluxscale

To set the absolute fluxscale, either a solar system object or a regularly-monitored quasar will have been used. For a solar system object, the Butler-Horizons-JPL models are used, whereas the flux of a quasar flux calibrator will have been obtained already. The CASA task setjy is then used by the Pipeline to set the fluxes of the model visibilities for the flux calibrator.

13.5.9 Deriving the fluxes of the phase and bandpass calibrators

If the absolute flux calibrator is quite resolved, then only a subset of the ALMA array will be used to determine the fluxscale. This is because the model is more trusted for the shorter baselines. This is performed by using the CASA tasks gaincal and fluxscale.

13.5.10 Time Gain Calibration

The complex gains are derived from the data column (raw data) divided by the model column (usually set with `hif_setjy`). The gains are obtained for the specified solution intervals, SPW combination and field combination. One gain solution is computed for the science targets and one for the calibrator targets. Good candidate reference antennas can be determined using the `hif_refant` task. Previous calibrations that have been stored in the pipeline context are applied on the fly.

13.5.11 Applying the calibration

The Pipeline uses the CASA task `applycal` to apply the calibration tables to the dataset. A strict setting of `applycal` is used by default in the Pipeline (any data without full calibration available is flagged). This conservative approach is taken in order to prevent any non-properly calibrated data getting through the Pipeline.

13.5.12 Imaging

At minimum, the Pipeline images the bandpass and phase calibrators as a diagnostic check, producing an image per spectral window. Pipeline science target imaging may be commissioned during Cycle 3. Prior to this, Pipeline calibrated datasets will be manually imaged by ALMA staff.

13.5.13 Exporting the products

The `hif_exportdata` task exports the data defined in the pipeline context and exports it to the data products directory, converting and or packing it as necessary. The current version of the task exports the following products

- an XML file containing the pipeline processing request
- a tar file per ASDM / MS containing the final flags version
- a text file per ASDM / MS containing the final calibration apply list
- a FITS image for each selected calibrator source image
- a FITS image for each selected science target source image
- a tar file per session containing the caltables for that session
- a tar file containing the file web log
- a text file containing the final list of CASA commands
- a python script that can be used to calibrate the data
- a python script that can be used to restore raw data directly to calibrated measurement set

13.6 Low SNR Heuristics in Interferometric Pipeline Processing

If there is a low signal to noise ratio on one or more calibrators then the Pipeline can currently mitigate this in several ways.

13.6.1 Mapping the gain solutions from broad spectral windows to narrow ones

In datasets containing one or more narrow spectral windows and a weak phase calibrator, then the gain solutions derived for the narrow spectral window(s) may be significantly lower quality than those determined for any broader spectral windows in the dataset. The Pipeline can therefore apply the gain solutions determined for broad windows to narrower ones. However, care must be taken when performing this as there may be a phase offset between the solutions for different windows, particularly when these windows are not within the same ALMA baseband. Therefore, the observation of the usually stronger and higher SNR bandpass calibrator is used to determine the phase offset between the windows enabling correction for this.

13.6.2 Using the gain solutions obtained from the aggregate of all spectral windows

In some datasets, there is low SNR on the phase calibrator and all of the spectral windows used were narrow. For cases including this then the aggregate gain solution obtained from using all spectral windows to determine the solutions can be applied. Again the Pipeline ensures that phase offsets between the gain solutions for individual spectral windows and the aggregate solution is corrected for by determining this using the bandpass calibrator. This will become available in the Pipeline during Cycle 3.

13.6.3 Additional Bandpass Pre-averaging

For low SNR on the bandpass calibrator, the frequency pre-averaging can be performed using a larger frequency averaging interval.

13.7 Total Power Data Processing

Differently to interferometer operation, total power observations are designed in paired MOUS: one ‘science’ MOUS (i.e. the science target), and one calibration MOUS (used to derive the Kelvin-to-Jansky conversion factor). Both MOUS exist within the same GOUS and Pipeline processing is performed per MOUS at the beginning of deployment of Single Dish Pipeline.

Pipeline operation slightly differs between processing science targets and amplitude calibrators. Since Cycle 3 offers only spectral line emission observations for science targets, the processing for science targets always have line identification, baseline subtraction and subsequent flagging based on baseline rms after the Tsys and Sky calibration, followed by imaging. The processing for amplitude calibrators is generally similar in many ways, but aims to extract only the continuum signal; i.e. the same calibrations steps are followed as for the science data, although the imaging step immediately follows without line-identification, baseline subtraction, and baseline rms flagging.

13.7.1 Data Import

(Both Calibrator and Science data). The Pipeline first loads the data into the pipeline context unpacking and/or converting it as necessary.

Note that the imported files are in MeasurementSet (MSs) format, but the Single Dish Pipeline converts the MS-format files into Scantable format at the steps after generation of Tsys tables (see ‘Tsys & Sky Calibration Tables’).

13.7.2 Deterministic & Manual Flagging

(Both Calibrator and Science data). The Single Dish Pipeline applies online flagging of shadowing, identification of un-needed intents such as FOCUS, as well as some edge-channel flagging. It is also possible to insert manual flagging if an additional flagging than pipeline does is needed.

13.7.3 Tsys & Sky Calibration Tables

(Both Calibrator and Science data). The Pipeline creates Tsys calibration tables. Each spectral window of the science targets is automatically mapped to the corresponding, frequency-matched Tsys spectral window.

The Pipeline creates the sky calibration tables, linking each ‘ON’ (i.e. on-source) spectrum to the corresponding spectrum at ‘OFF’ (i.e. reference) positions. In contrast, the edges of the maps of the amplitude calibrators are treated as OFF points.

13.7.4 Applying the calibration

(Both Calibrator and Science data). The Pipeline uses the CASA task `sdcal2` to apply the Tsys and Sky calibration tables to the dataset. At the completion of this step, the resulting data are bandpass-corrected and have temperature units of K, in T_A^* .

13.7.5 Automated line identification and Baseline subtraction

(Science data only). The Single Dish Pipeline employs automated line identification algorithm which identifies candidate emission lines having a brightness above a certain threshold, and masks them. Then a low-order baseline is estimated and subtracted.

13.7.6 RMS Flagging

(Science data only). During the RMS flagging step, suspect data are flagged by outlier of RMS, by remarkably large RMS (relative to the expectation computed for the appropriate Tsys), and by rapid increase of RMS values.

13.7.7 Applying Kelvin-to-Jansky factors

(Science data only). The current strategy to obtain Kelvin-to-Jansky factors demands an observation of bright amplitude calibrator for each day the science targets are observed, and for each spectral window. After the calibrator source is imaged (as described in the following step), the measured brightness temperature (in K, T_A^*), can be compared to expected flux density (in Jy), and a Kelvin-to-Jansky factor conversion can be derived, relevant to each science observation, each day.

Since the factor is applied just before imaging of science targets, the Pipeline products (science targets) have a unit of Jy, while the pipeline processing product for the amplitude calibrators are always in Kelvin.

13.7.8 Imaging

(Both Calibrator and Science data). The next step produces a data cube (image fits file) for each spectral window per field. Aggregate continuum images will not be computed for the spectral line map.

Amplitude calibrator (continuum) images will be created inside the Pipeline only for the purpose of deriving the Kelvin-to-Jansky factor.

13.7.9 Exporting the products

(Both Calibrator and Science data). The final step is to export the image cube defined in the pipeline context to the data products directory, converting and or packing it as necessary. The current version of the task exports the following products:

- an XML file containing the pipeline processing request
- a tar file per ASDM / MS containing the final flags version
- a text file per ASDM / MS containing the final calibration apply list
- a FITS image for each selected science target source image
- a tar file per MS containing caltables for the MS
- a tar file containing the file web log, a text file containing the final list of CASA commands
- a python script that can be used to calibrate the data
- a text file containing the Kelvin-to-Jansky factors

No amplitude calibrator image shall be included in the data products. Instead, containing the Kelvin-to-Jansky factors will be delivered to the PIs.

Chapter 14

Data Archiving

14.1 Introduction

The ALMA Archive stores the observed raw data and metadata as well as the reduced data products and their metadata resulting of the QA2 processing. These data and metadata are made available to PIs and archival researchers for querying and download following ALMA's data access policy. In the future, VO services will be provided.

14.2 Data Flow and Archive

Data from the correlator, together with monitor and weather data, are sent via dedicated optical fiber links (1-10 Gbit/s) 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 capturing and data flow systems and will be imposed by the Observing Tool at the proposal validation stage.

The Pipeline processing system and the 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 200TB of data. The OT will issue a warning message when the data rate exceeds twice the average data rate. Currently proposals will not be rejected on grounds of large data rates and the justification is only required to ensure that a configuration yielding a less demanding data-rate and data-volume is scientifically not an option. In future cycles, however, it might be that data rate considerations play a larger role.

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 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 ALMA Regional Centers (ARCs) in North America, Europe and Japan. 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.

The ARCs hold copies of the entire ALMA archive for backup, user support and data distribution to the ALMA PIs and archival researchers. The three ARCs provide a completely identical user experience to their communities and a user can download data from any ARC.

The ALMA archive consists of two parts - the ALMA frontend archive (hereafter AFA), which is used for storing all data and metadata of the observations and observing projects, and the ALMA Science Archive (hereafter ASA), which stores a subset of the science metadata in a relational database to enable fast science queries and VO access. The architecture is based on the Next Generation Archive System (NGAS) with Oracle

technology for replicating the metadata (see Figure 14.1).

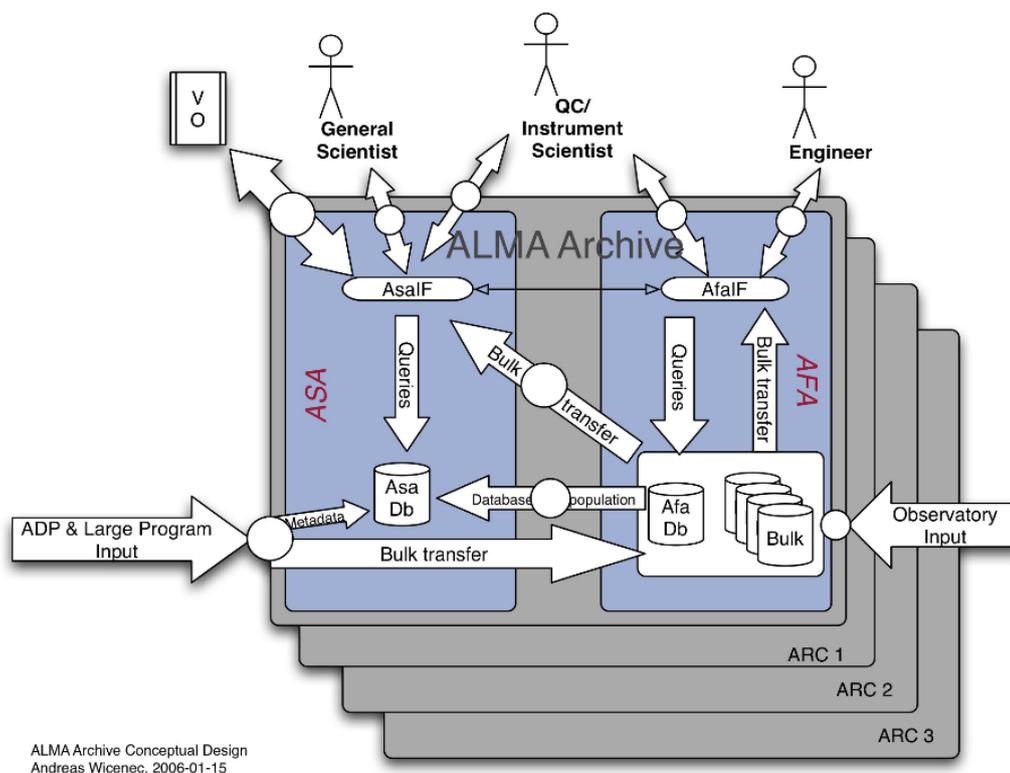


Figure 14.1: Archive design (frontend & backend) at the OSF to store metadata and raw and monitor data.

As soon as data is taken and has passed the QA0 quality control step (see Chapter 11.2), its 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. Metadata harvested includes 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.

Once the pipeline processed an ObsUnitSet (hereafter OUS, see 8 for details on the OUS structure) and the science products have passed QA2 they are ingested into the ASA.

14.3 PI Data Delivery and Data Delegation

The unit of data delivery to the PI is the OUS. As soon as QA2 on an OUS has passed and the data products have been replicated to the users' home ARC, PIs will receive a email notification containing a link to their data, and 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. 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. CoIs. 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.

14.4 Archive Query

At <http://almascience.org/alma-data/archive> users can query the holdings of the ALMA Science Archive (Figure 14.2). 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 14.2: The ALMA Science Archive Query Form

Queries can be made by physical quantities along the Position, Energy, Time and Polarization 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 will present users with the metadata of the public raw data, although they can choose to also see the metadata per project. On the results page, users can sort and subfilter the results and add or remove columns from the result table to narrow their search.

The results page also offers to download the results in VOTable, TSV (Tab-separated values) or CSV (Comma-separated values) format for further processing in tools like `topcat`¹. Modifications of the URL used for the exporting of the query results can also be modified to access the ALMA archive queries programmatically. `Astroquery`² can be used to encapsulate and simplify programmatic access to the ASA. Both tools mentioned are developed externally and are not part of the official ALMA software.

14.5 Request Handler

Once data of interest are defined, they can be selected via checkboxes on the results page and submitted to the ALMA Request Handler for download. The data are displayed in the full OUS hierarchy. This allows users to

¹<http://www.star.bris.ac.uk/~mbt/topcat/>

²<https://astroquery.readthedocs.org/en/latest/alma/alma.html>

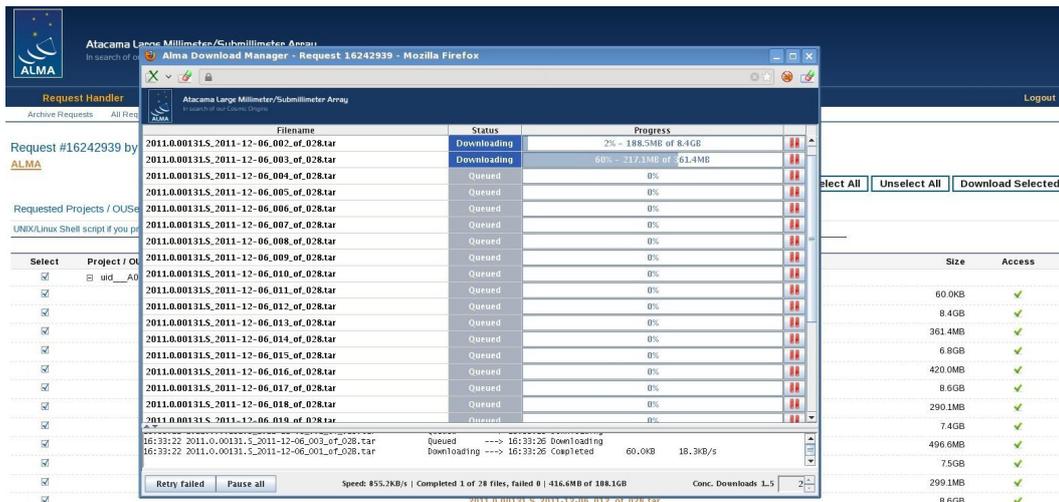


Figure 14.3: The ALMA Science Archive Request Handler

check immediately if there are data from other Member OUS or higher-level Group OUS products available. By default, only data products (i.e. FITS-format, and 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.

Four 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.
- Finally, 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".

The Request Handler also allows to select individual files for download. Only the data deliveries that users have permission to download can be selected. If they are not authorized 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 can have different policies regarding the details of the shipping of the hard-disks.

Appendix A

Antennas

A.1 Design and Properties

At the end of the construction period ALMA will have in total 66 antennas, 54 with a diameter of 12 m and 12 with a diameter of 7 m. Four of the 12 m antennas will be equipped with a nutating subreflector for total power observations. The four antennas used for total power observations and the twelve 7 m antennas will together form the Atacama Compact Array (ACA). The ALMA antennas are manufactured by three different contractors. These are VertexRSI (North America) which will provide 25 12 m antennas, Alcatel Alenia Space European Industrial Engineering MT Aerospace (AEM, Europe), which will provide 25 12 m antennas and Mitsubishi Electric Corporation (MELCO; East Asia), which will provide 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 during ALMA Cycle 3 have both 12 meter and 7 meter diameters, with the receivers mounted at the secondary (Cassegrain) focus. The 12 m dishes 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 dishes 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.

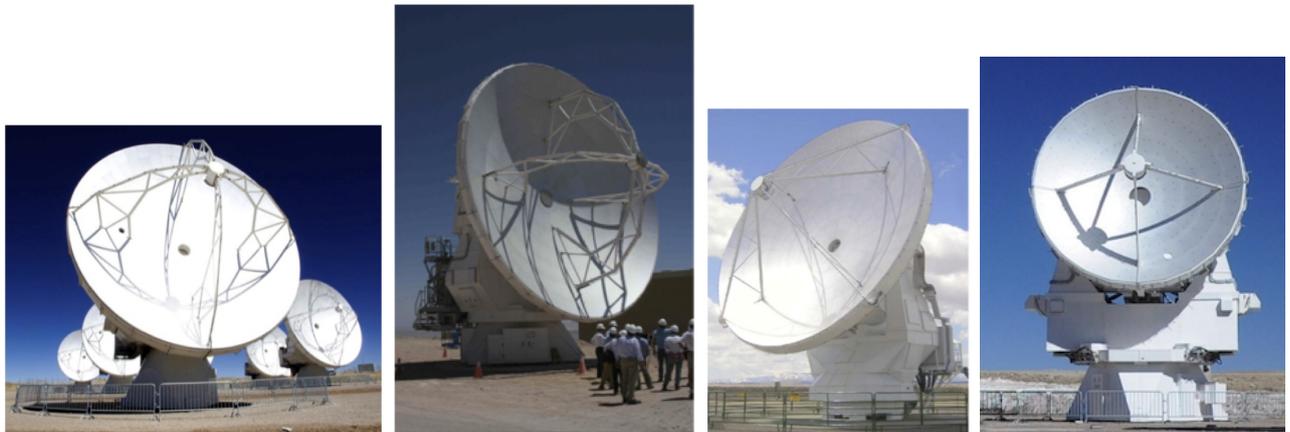


Figure A.1: The four different ALMA Antenna designs: Vertex 12 m, MELCO 12 m, AEM 12 m, and MELCO 7 m (from left to right).

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 aluminum 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).

	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 of 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.

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.

The different antennas use a combination of steel, aluminum, Carbon Fiber Reinforced Plastic (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 builtin 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.

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



Figure A.2: Structure of an antenna pad (actual pad at the OSF) (left) and detail of antenna anchored to a pad (right).

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 neither include the kinematic mount lower part nor it includes 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 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 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.

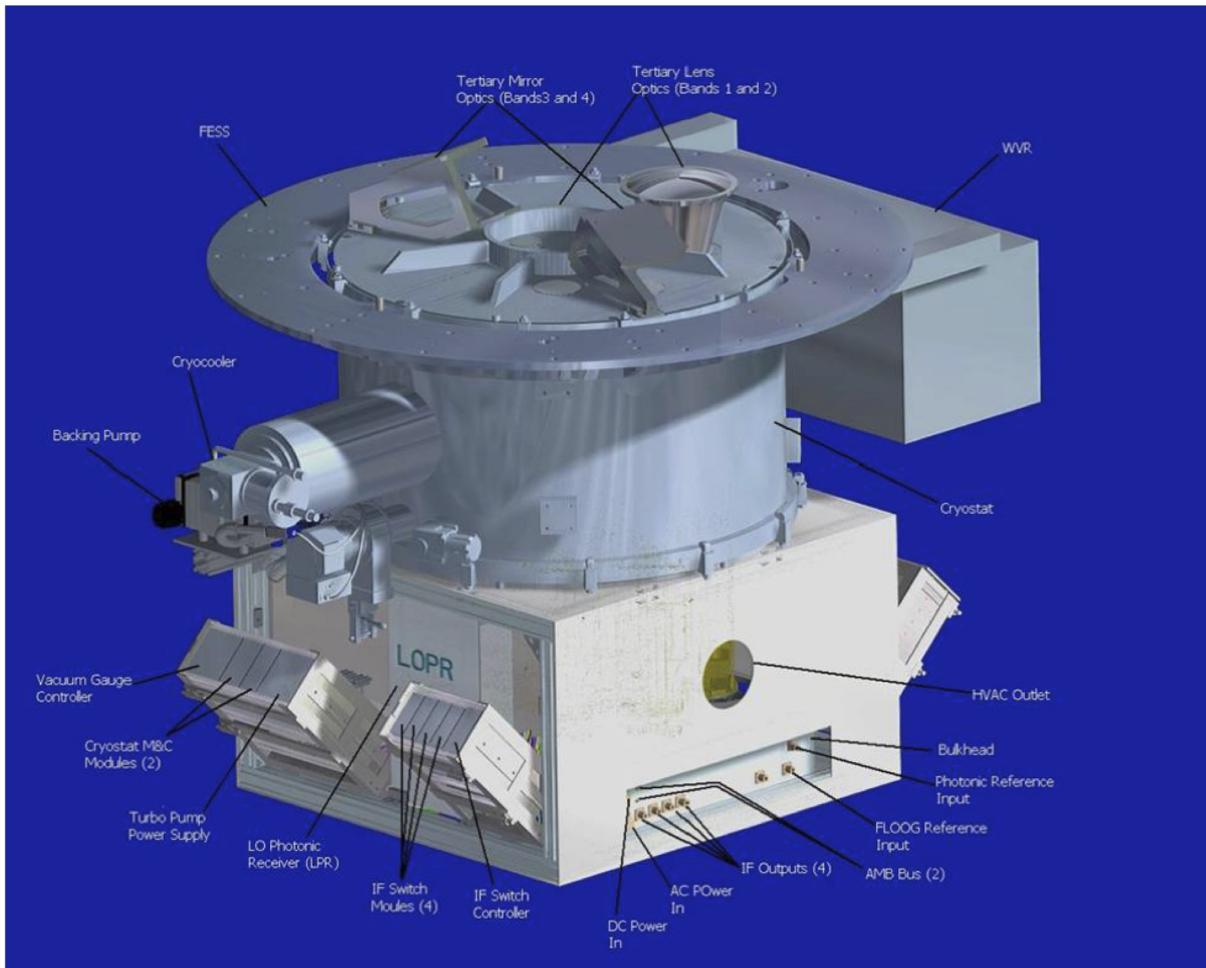


Figure A.4: Side view of ALMA frontend showing cryostat assembly, with room temperature unit below.

A.4 Cryostat

The ALMA frontend 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.

¹<http://www.almaobservatory.org/en/about-alma/how-does-alma-work/technology/transporters>

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 frontfrontend, and is described in Section A.5. Figure 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.

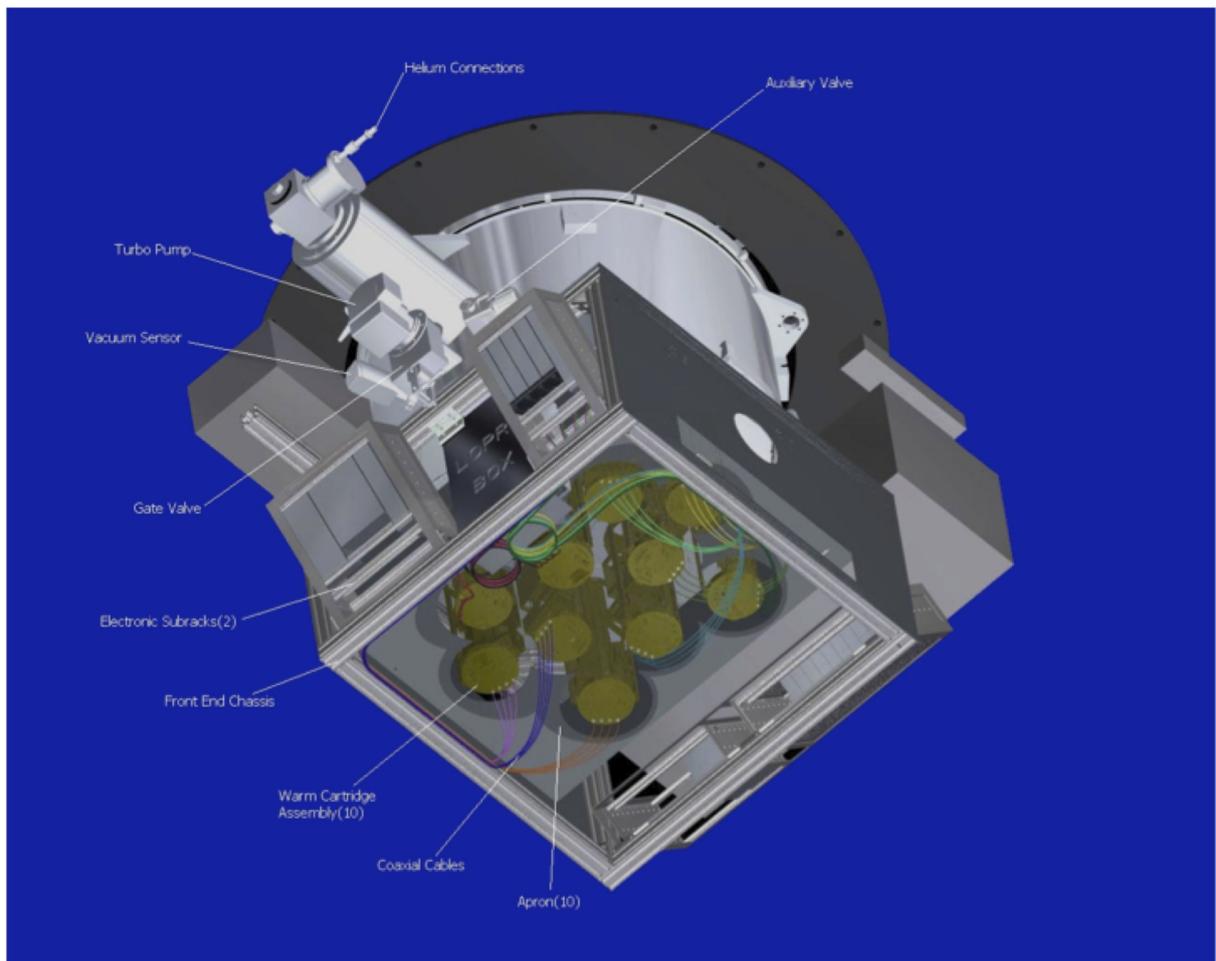


Figure A.5: Bottom view of ALMA frontend, showing WCAs.

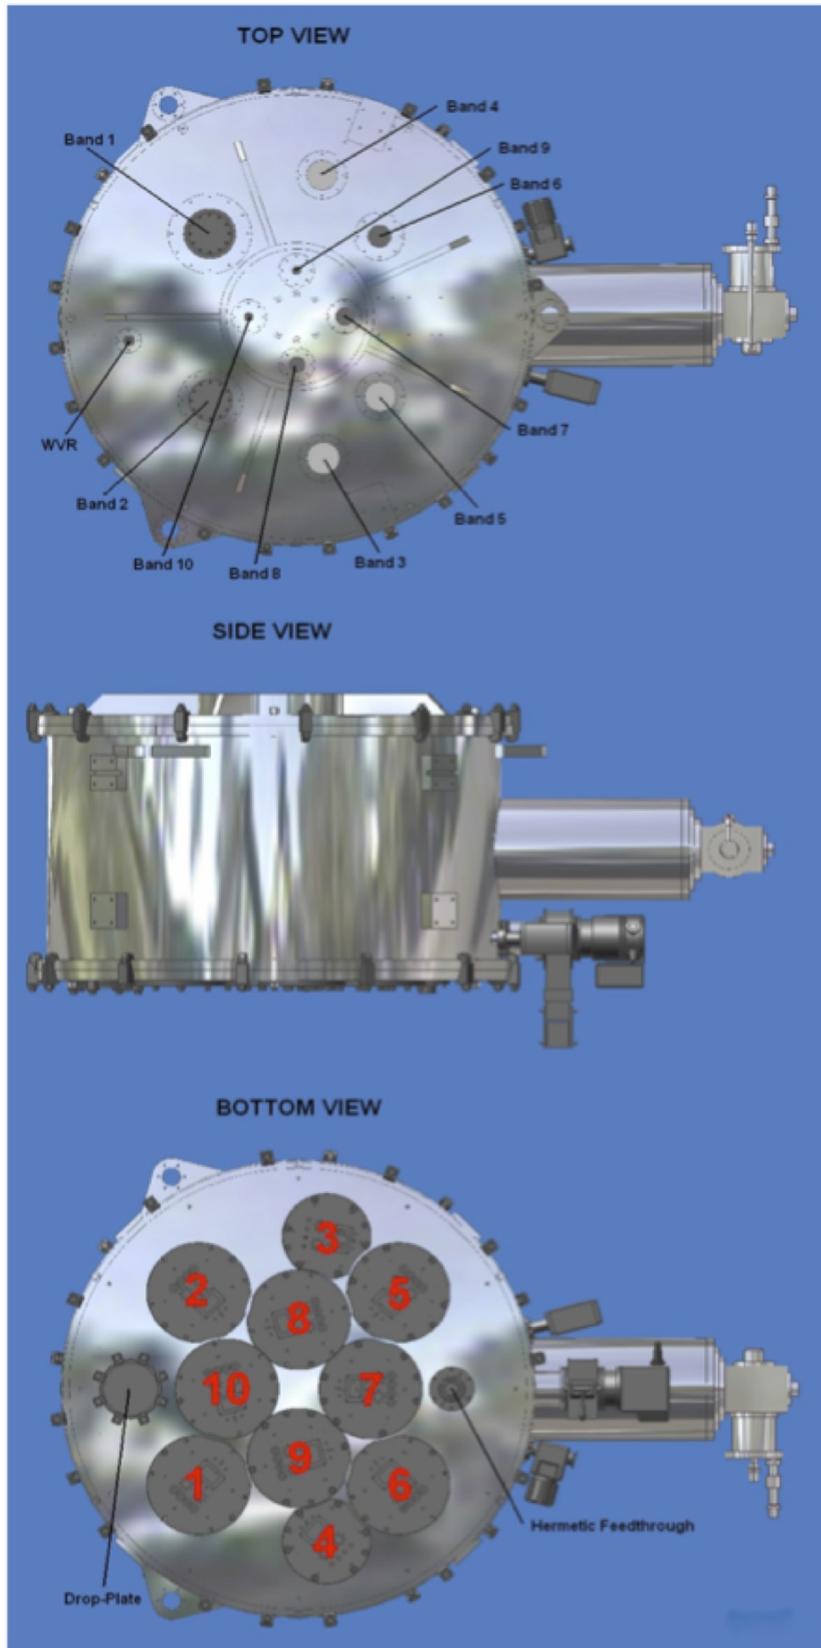


Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view).

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 frontend (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 (solar observations are not available during Cycle 3). 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.

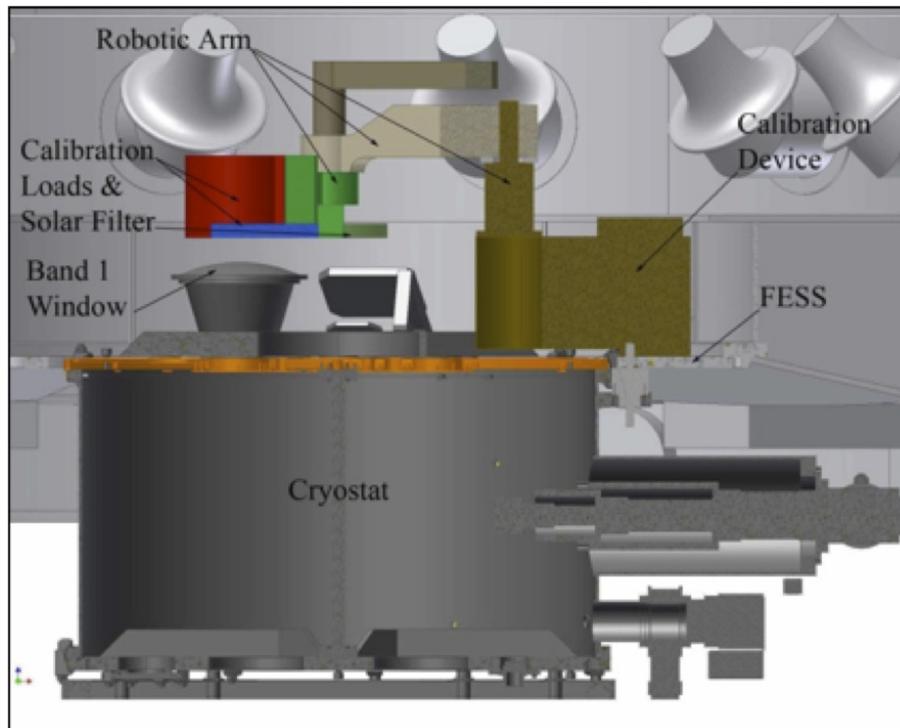


Figure A.7: Lateral view of the ACD on top of the ALMA frontends.

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 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.

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.).

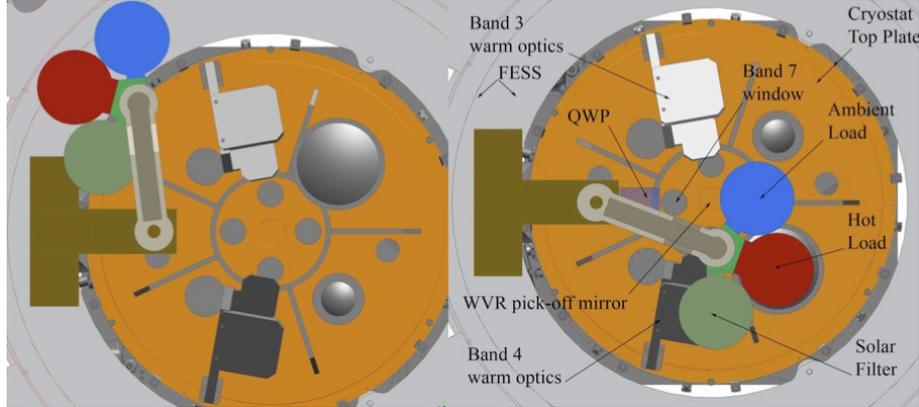


Figure A.8: Top view of an ALMA frontend showing the robotic arm of the ACD retracted during normal observations or on top of one of the frontend inserts for calibration. The current design has been improved by placing all the loads in a wheel.

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:

$$\alpha = \frac{\eta J_M + (1 - \eta) J_{SP} - J_{L_2}}{J_{L_1} - J_{L_2}} \quad (\text{A.2})$$

$$T_{Cal} = (J_{M_s} - J_{BG_s}) + g\eta^{\tau_s - \tau_i} (J_{M_i} - J_{BG_i}) \quad (\text{A.3})$$

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 that can be achieved for Cycle 3.

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 1Hz. On timescales shorter than this 1 Hz timescale, the watervapor 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.

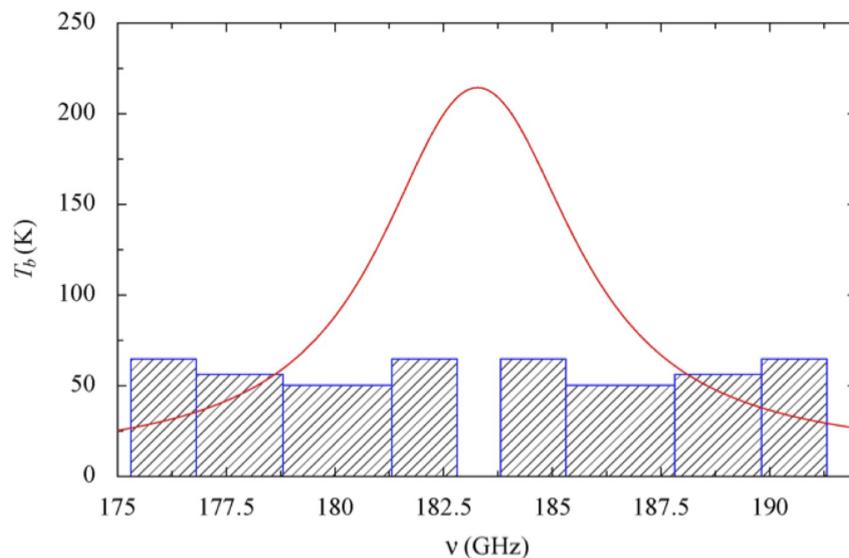


Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

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”

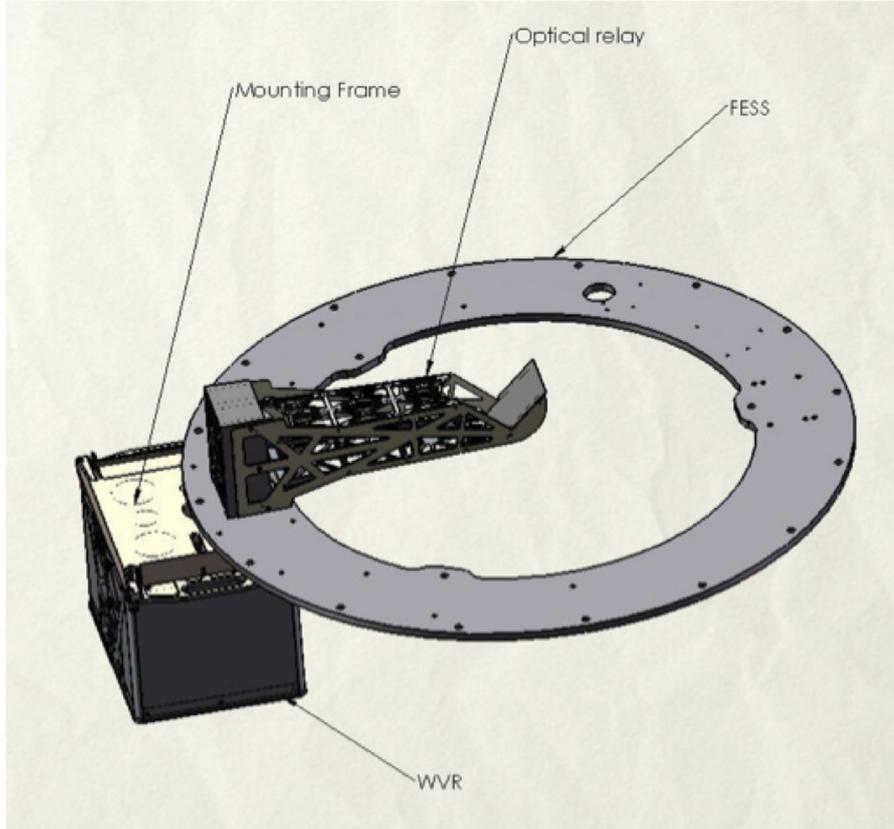


Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

atmosphere (i.e., PWV). 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):

$$\delta L_{corrected} \leq \left(1 + \frac{w}{1\text{mm}}\right) 10\mu\text{m} + 0.02\delta L_{raw}. \quad (\text{A.5})$$

where w is the total water vapor content 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 (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

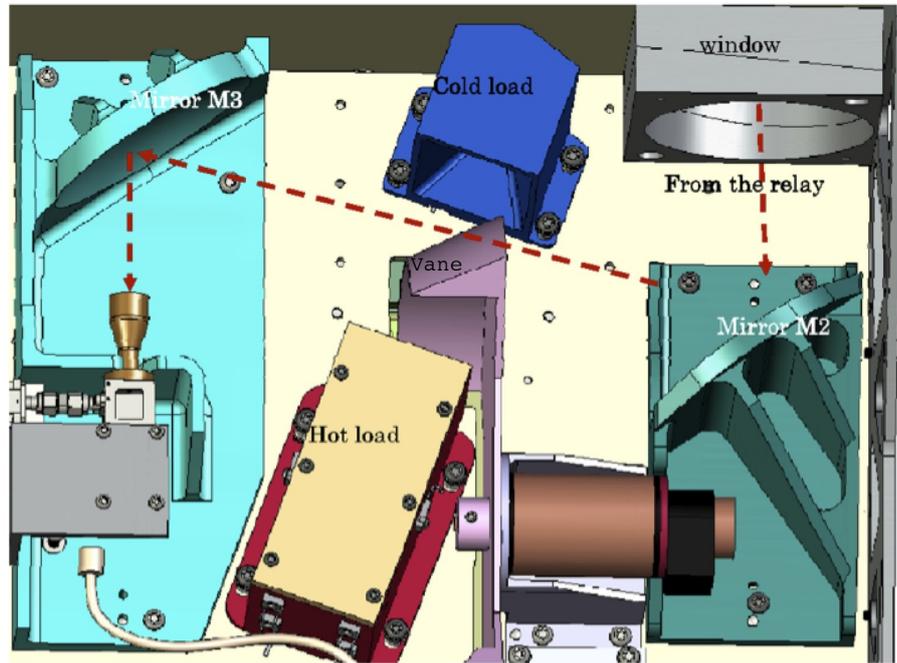


Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

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 product 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. It is envisioned that corrections at the scales of the sampling rates of the WVRs will be possible at the correlator and that refinements for longer timescales will be done offline in CASA using the `wvrgcal` tool.

Appendix B

The LO and IF System

In this Appendix we describe the signal path and LO chain used between the frontends and correlators, and how these are used to define spectral setups for the user. 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 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 jump to Chapter 6.

B.1 Functions of the LO & IF system

In the signal path from Frontend to correlator, ALMA uses three frequency conversions, and 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. Suppression of the image sideband or, 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 Correla-

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

tor². Section 6 shows how to setup the system to observe spectral lines (particularly multiple spectral lines) and continuum. We then discuss some other aspects of frequency setups, 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 B.2, and in Section B.3, we describe the hardware and how the LO frequencies are synthesised and distributed around ALMA.

B.2 Summary of Operation

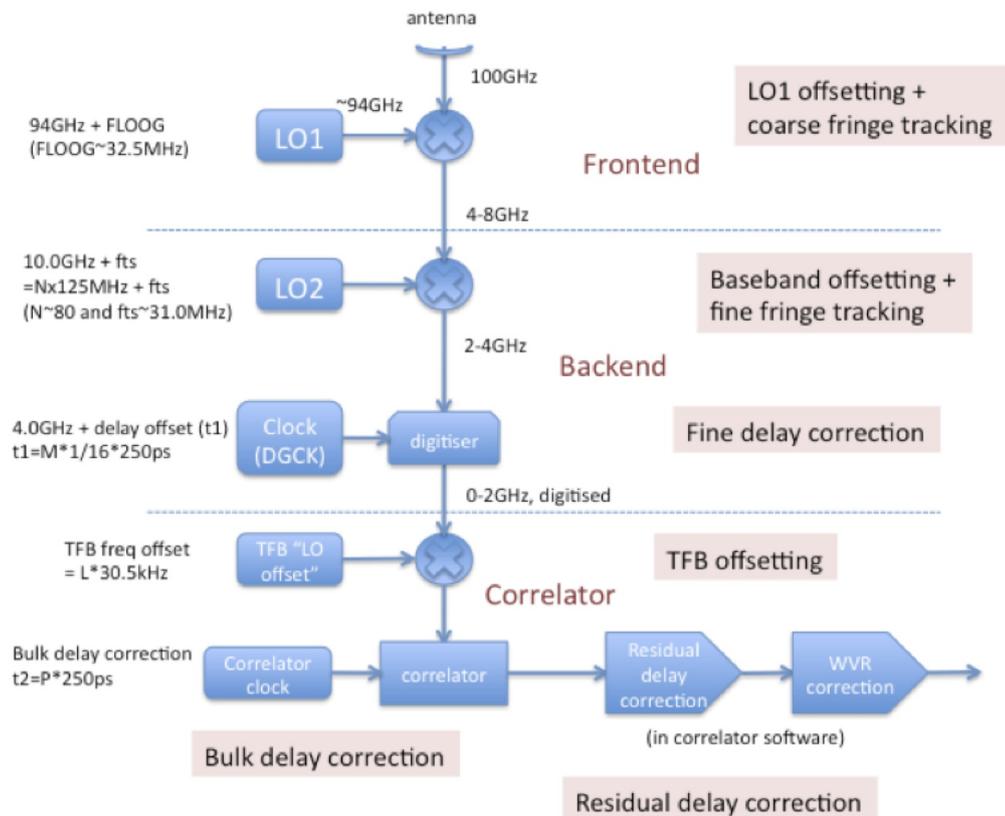


Figure B.1: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the Frontend, Backend, and Correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz seen in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized tuning steps, such as LO2 (which can be changed in a multiple (“N”) of 125 MHz plus an finely-adjustable offset of “fts”), 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 Band 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

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

lower sideband; although at the edges of the tuning band, only one sideband is possible. LO1 consists of a common 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 (by 180 or 90 degrees) in conjunction with a Walsh switching pattern on the antennas to remove DC systematic errors, for sideband suppression and, in the future, for the DSB receivers (band 9 & 10) for sideband separation.

2. In the "backend" (BE), the IF processor (IFP) splits the IF into basebands, each with frequency range of 2-4 GHz, via a set of filters and tunable second LOs (LO2) (see IFS/IFP section in 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 a second fine-tuned synthesiser (fts) which provides an offset in the range 20.0-42.5MHz (marked as "fts" in Figure B.1). The limited fts range and the 125MHz quantization means that LO2 setting is not fully contiguous; consequently there can be up to ~30MHz 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 30MHz. 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 2-4GHz 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 spectral windows 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 fibre-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, we describe some of these components.

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

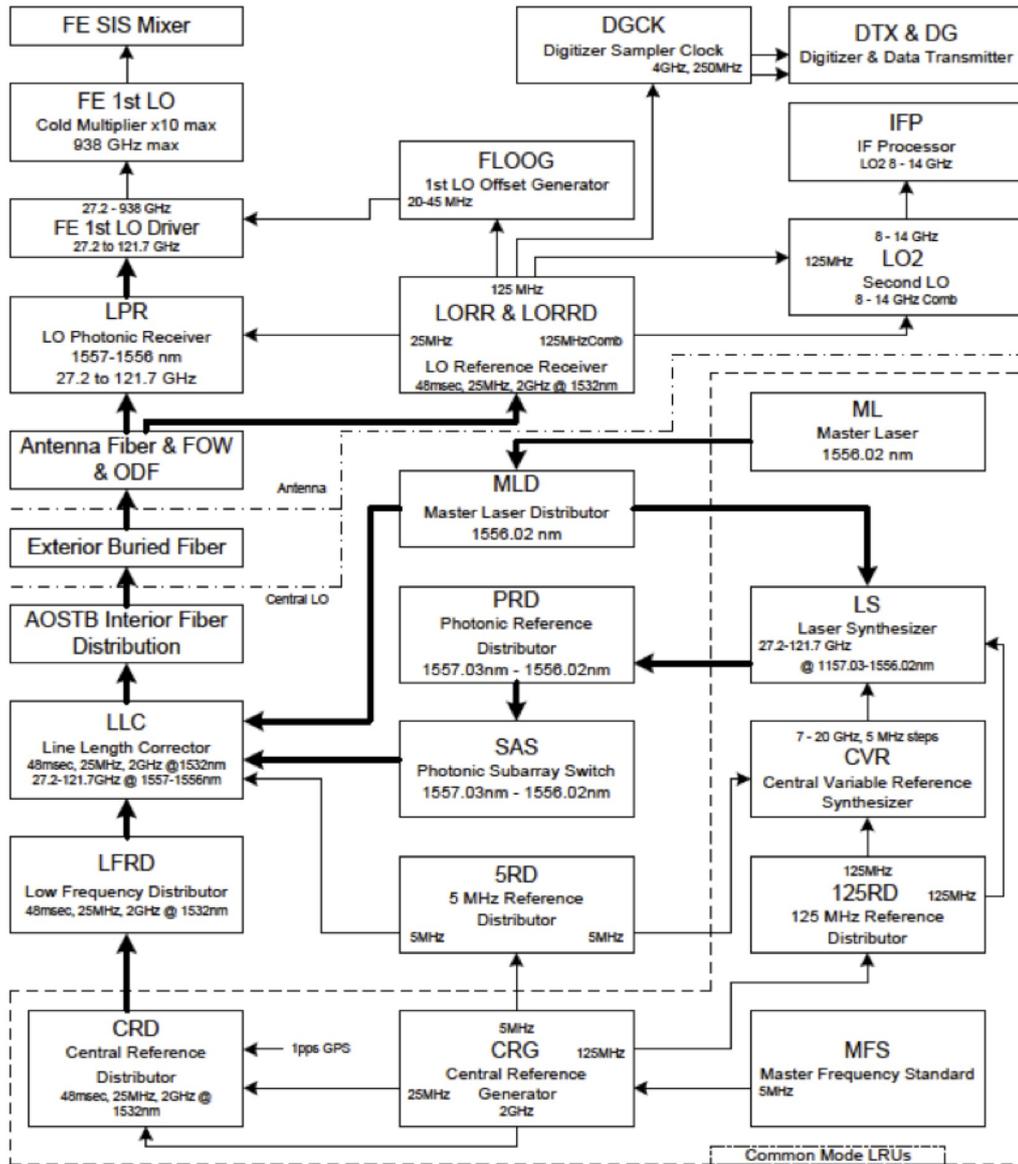


Figure B.2: Summary block diagram of the LO distribution system. The components in the lower section of the diagram (below the the dashed-dotted line) form the Central LO (CLO), located in the AOS Technical Building. The components above are located in each Antenna (only one antenna is depicted in this 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.

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 (CRG) 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) 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 fibre-optic distribution system (see Section B.3.3).

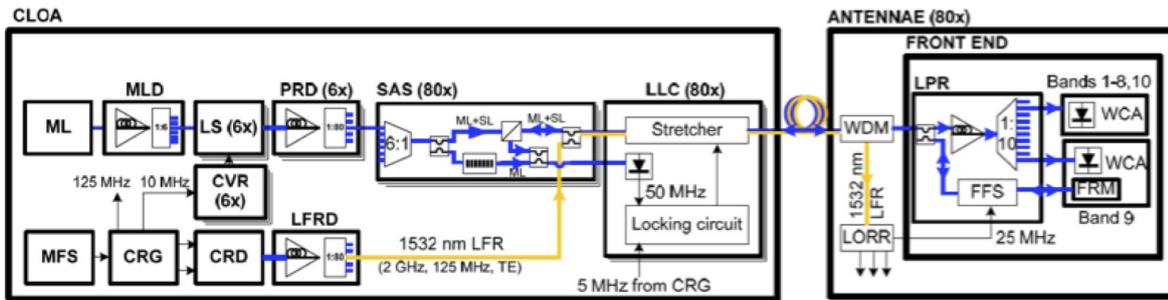


Figure B.3: LO block diagram, showing the Central LO (CLO) and the LO section in the WCA in each front end. Yellow lines represent common synchronization and reference signals fed to all antennas, blue lines are LO signals which are individual to each antenna. For description of acronyms, see text.

B.3.2 LO1 signal generation and distribution

LO1 is distributed through fibre-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 6 Slave Lasers (Laser Synthesizers, LSs) that can produce 6 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 (eg band-to-band transfer or spectral scans).

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 reference signals that are fed into the 6 Laser Synthesizers (LS). The LSs controls the frequency of the the Slave Lasers producing a frequency offset of the SL signal of 27-120 GHz with respect to the ML signal. The SL signals are added to the ML signal. The

offsets between the ML and SL signals provides the beat note which is used to generate the LO1 frequency in the photomixers (LPRs) 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 6 LSs it is possible to generate 6 separate LO1 frequencies.

- The Photonic Reference Distribution (PRD) feeds the optical signals to the Sub Array Switch (SAS) which can distribute the signals to the different sub-arrays.

Figure B.4 shows the three laser signals after combination in the Sub Array Switches (SAS). The Master and Slave laser signals have wavelengths of about 1556 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 fibres are distributed in buried trenches, and fed into the Cassegrain cabin on each antenna through Az and El fibre wraps. All are fed through Line Length Correctors (LLCs), which are used to correct for changes in the optical fibres. The LLCs are described in B.3.3 below.

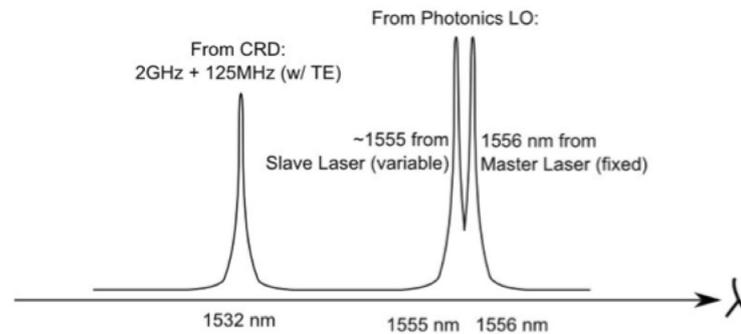


Figure B.4: The ALMA fibre signals. The 1532nm carrier contains the frequency reference signals, and the 1556/1555nm carriers are used to remotely generate LO1. Within each antenna, the optical fibers are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing 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 (PDM) 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 and mixed at the SubArray Switch (SAS) 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 at 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 receives thus twice a frequency shift of 25 MHz, thus 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

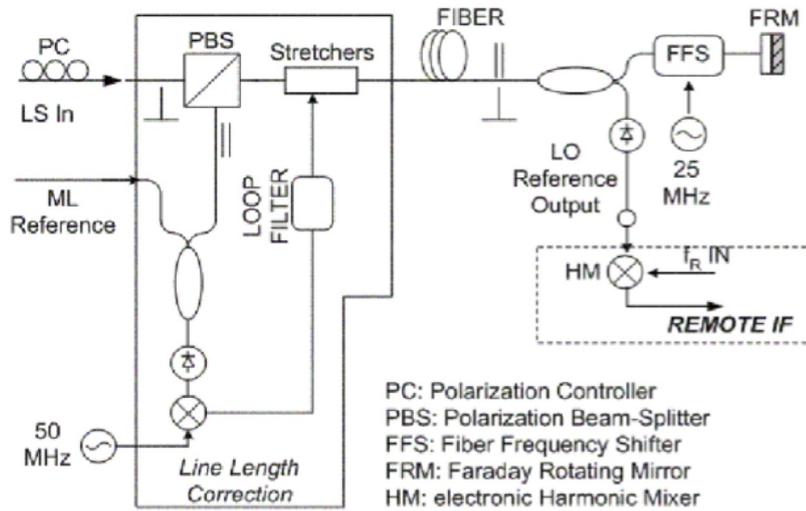


Figure B.5: Block diagram of the Line Length Corrector system for ALMA. The FRM (Faraday Rotation Mirror), shown upper right, is located in the band 9 cartridge of every front-end. It rotates the polarisation of the incoming light, and the resulting reflected signal is fed back through the buried optical 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.

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 1kHz) copes with the small range variations (about 0.1 mm). The LLCs are reset to 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 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 – 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 10–30GHz (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 a IF Switch unit (IFS) situated in the frontend, 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 polarisation channel is shown). The Band 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 polarisations. Each output has a IF bandwidth of up to 4 GHz. For Band 9 & 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.

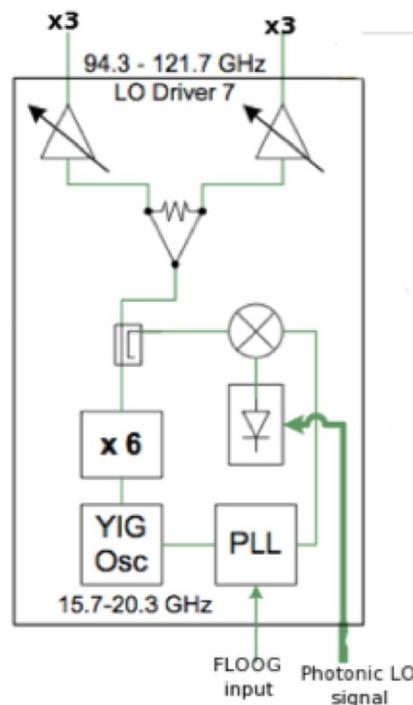


Figure B.6: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier, not shown here, is used to generate the LO1 frequency at 282.9 – 365.1 GHz (in this case, a x3 multiplier). 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 8-14GHz. LO2 is generated from a harmonic of 125MHz, plus a fine-tuned synthesiser (fts) of range 20-42.5MHz, 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).

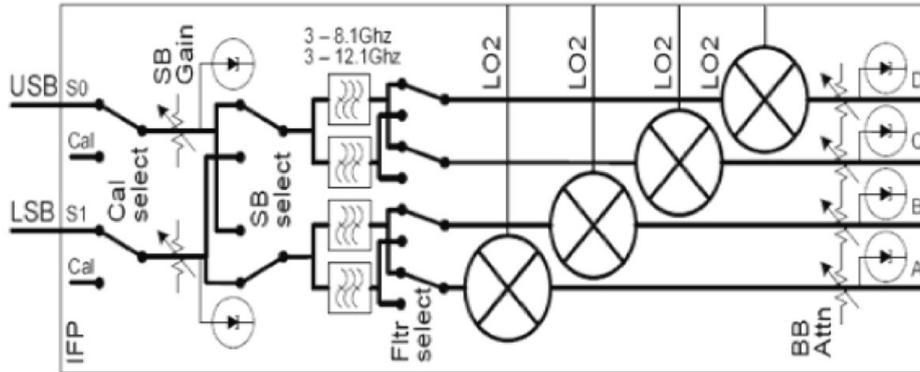


Figure B.7: Block diagram of one polarisation channel of the IF Processor. This has two IF inputs (at left), and feeds 4 IF outputs to the digitisers (to the right of the diagram).

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 6.5.1). These filters cause a decrease in the effective IF range to approximately 1.875 GHz.

B.3.6 Digitization and signal transmission

The outputs of the IF Processor units are fed into the Data Transmission System modules (DTS), that include digitizers and formatters to convert the signals to optical wavelengths for transmission via optical fibers. There are four DTS 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 DTS 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 DTS 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 DTS 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 Receiver units), 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 DTS module.

B.4 Other functions of the LO/IF

In addition to frequency downconversion, the LO/IF performs several other tasks, detailed below.

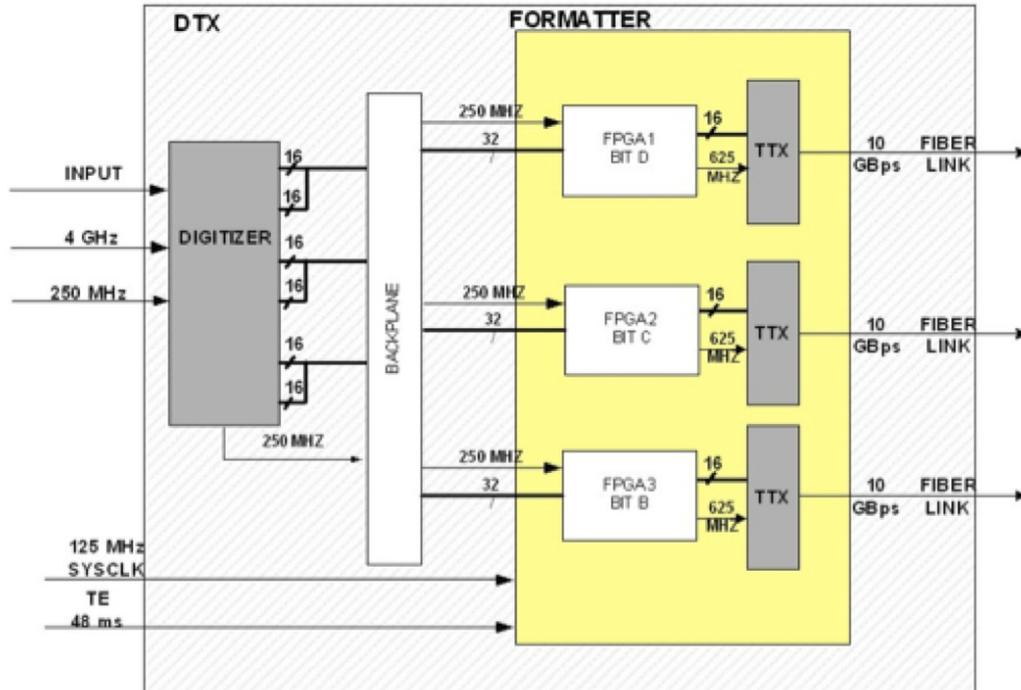


Figure B.8: The DTX Signal Digitization and Transmission system, as used in each antenna.

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 these is 1/16 of the FLOOGs (i.e., 1/16th of 250 ps). 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 ps/16) by applying a linear phase gradient across the passband after correlation. Also, the correlator applies relative 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. Band 3, 4, 6, 7, & 8) are inherently single sideband (SSB), either through having a mixer or quasioptic design which rejects the unwanted sideband⁵. 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 (Band 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

⁵Although they have mixers to allow both sidebands to be observed separately and simultaneously

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 Sideband separation - 90 degree Walsh switching

For future Cycles, it will be 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 Band 9 & 10 (DSB) it will effectively double the bandwidth from 8 GHz to 16 GHz per polarization from the DSB receivers. This is under development and will not be available for Cycle 3.

B.4.4 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 pattern cycle time of 16 ms. The Walsh 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

Appendix C

Calibration Source Selection

C.1 Antenna-based Sensitivity Determinations

The proper selection of two calibrator types is critical to the success of the scientific experiment: the **bandpass calibrator** to determine the relative frequency dependence of amplitude and phase during the experiment, and the **phase calibrator** to determine the temporal variations of the amplitude and phase over the experiment. An experiment will be defined here as one scheduling block or a series of consecutive scheduling blocks that are tied together using scan sequencing. Some scientific experiments may span several configurations and each should be treated as separate experiments in determining appropriate calibrators, although the choice may overlap.

The main criterion of the calibrator selection is that their correlated visibility amplitudes over all baselines is sufficiently strong in order to obtain the specified signal-to-noise for the to determine of the frequency spectrum and temporal variations during the experiment. Other calibration types, like determining the flux scale of the observations and the polarization leakage, are described elsewhere.

C.1.1 ALMA Sensitivities

The present ALMA sensitivity must be known in order to choose the calibrators. These sensitivities are presently determined from an ASCII table that contains the nominal ALMA sensitivity as a function of frequency. It was obtained from the ALMA Sensitivity Calculator by running it many times. A partial list of the table is shown in Table C.1.

For example, if the experiment frequency is 345 GHz, then the ALMA sensitivity in terms of an image RMS made from 60-sec of integration over a 7.5-GHz bandwidth (all four TDM spectral windows) is 0.28 mJy, the average of the 340 and 350 GHz rows. The nominal T_{sys} is used and an assumption of a typical Water Vapor column in the line of sight. To be conservative, the RMS value should be taken from the least sensitive of the SPW frequencies in the experiment.

New observing band data will be added as needed, and in the future these data will be downloaded from a system file; but for the present this sensitivity file is made by hand.

C.1.2 Calibration Parameters and Derived Antenna-based RMS

The calibration parameters, adopted in 2011 but subject to change, are as follows. The relevant SNR given are for an **antenna determination**.

- Phase calibration minimum SNR (over entire BW): 15.0

SENSITIVITIES FOR 32 12-m antennas
DUAL POLARIZATION, TOTAL BW = 7.5 GHz
INTEGRATION TIME 60.0 sec
RMS in mJy

FQ	RMS	TSYS	WVR	TSKY
280.0	0.194	133.0	1.27	26.3
300.0	0.212	144.0	1.27	32.2
310.0	0.234	157.0	1.27	39.7
320.0	0.480	319.3	1.26	98.5
325.0	16.57	9999.9	1.26	249.0
330.0	0.404	267.4	1.26	84.9
340.0	0.273	179.5	1.26	50.7
350.0	0.291	188.6	1.26	55.3
360.0	0.362	234.3	1.26	73.2
367.0	0.719	462.0	1.26	128.0
372.0	1.19	760.0	1.26	166.0

Table C.1: Table showing ALMA examples of sensitivities as function of frequency.

- Maximum phase calibration scan length: 2.0 minutes
- Maximum phase calibration separation from target: 15.0°
- Bandpass Cal/chan minimum SNR: 50.0
- Maximum bandpass calibration scan length: 10.0 minutes
- Maximum bandpass calibration separation from target 50.0°

The phase calibration SNR is equivalent to an antenna solution rms accuracy over the scan of 4% in phase and 2% in amplitude. The maximum scan length is chosen so that the phase calibrator observing time is significantly less than that for the calibrator with normal phase referencing. Because of short-term pseudo-random tropospheric phase fluctuations and longer term systematic phases differences between the target and calibrator, a calibrator measurement with high SNR is not useful.

The bandpass calibration SNR is equivalent to a channel rms accuracy of 1° in phase and 0.5% in amplitude. Any variation of the normalized bandpass with time can only be determined using multiple bandpass calibration scans. At present, one bandpass calibrator is observed for each experiment (observing block) which is rarely longer than 90 min, hence the limit of a 10 minutes scan seems reasonable. The effect of temporal tropospheric phase fluctuations can be removed from the bandpass determination.

C.2 Choosing the Calibrators

The image sensitivity for the experimental parameters can be obtained from the above sensitivity table, scaled by the square-root of the appropriate bandwidth and integration time. The conversion of the image sensitivity S_{image} to *antenna-based* sensitivities S_{ant} also depends on the number of antenna, N_{ant} , and is given by Equation C.1

$$S_{ant} = S_{image}(N_{ant} - 3)^{0.5} \quad (C.1)$$

C.2.1 The Phase Calibrator

From the experiment frequency, ALMA sensitivities, the number of antennas, the total bandwidth, and the maximum scan length, the minimum flux density for a phase calibrator can be determined. Two minimum flux densities are determined: (1) The entire observing bandwidth and both polarizations can often be combined

BAND	Frequency (GHz)	Assumed Tsys (K)	Assumed Pwv (mm)	Flux Density (mJy)
3	100	80	5.2	5.5
4	145	89	2.7	6.3
6	230	95	1.26	6.9
7	345	151	0.91	11.9
8	405	280	0.66	23.0
9	690	877	0.47	101.7

Table C.2: The minimum phase calibrator flux density in mJy need for an antenna-based phase determination (see Eq. C.1) with 4 deg rms (15 SNR) with a 2 min integration time, assuming 37 antennas. The total bandwidth in both polarizations of 7.5 GHz is assumed.

coherently using the bandpass calibrator scan, and one amplitude and phase solution per antenna per scan can be obtained. (2) An amplitude and phase solution is determined for each SPW and polarization per scan. With four SPWs and two polarizations, this minimum flux density is $\sqrt{8}$ larger than the full bandwidth solution. The full bandwidth solutions are often made at the highest frequencies in order to find calibrators that are relatively close to the target.

The ALMA cataloged is then searched for all unresolved or nearly unresolved calibrators that are less than 15° from the target or the centroid position for a group of targets that are above the minimum flux density of the separate SPW case and the full bandwidth coherence case. The relative goodness of the calibrators is determined from: (1) the angular separation of the calibrator from the target; (2) the accuracy of the estimated flux density at the experiment frequency that is derived from date, frequency and error estimate of the ALMA catalog entries. **The most important consideration is that the estimated flux density, with all realistic uncertainties, must not go much below the threshold needed for successful phase calibration.** The algorithms now being used are undergoing review.

Examples of minimum phase calibration flux densities for some ALMA bands are given in the following table. They assume an integration time of 2 minutes and a bandwidth of 2 GHz, one polarization. The flux limit decreases by a factor 2.8 if 8 GHz in both polarizations are first coherently combined. These values assume the use of 32 12-m antennas.

The calibrator-target separation is the most important consideration in the choice of phase calibrator. The use of a strong calibrator that is significantly further from the target than a weak calibrator that does satisfy the minimum SNR consideration, is not recommended. However, weak phase calibrators near the minimum threshold should be avoided if there estimated flux density errors are large. For this reason, the more conservative minimum flux density (needed for solutions for each SPW and polarization) is often used, but with the expectation of combining coherently all of the bandwidth if needed. Another reason for using the more conservative limit, is that the phase calibrator will be sufficiently strong for accurate pointing observations during the experiment.

The paucity of recent observations at a frequency at or above 300 GHz is a major concern. Since very few calibrator observations over these frequencies have been measured, their estimated flux density at higher frequencies can only be estimated by extrapolating a lower frequency blues using a radio spectrum with a spectrum of the form $\nu^{-0.7}$. This can introduce a further error of up to 50%.

With all of the above considerations, the best calibrator is chosen for the experiment. The calibration queries also list the next four calibrator possibilities. For strong calibrators, the maximum phase calibration scan length can be significantly less than 2 min; however it is always more than 30 sec.

C.2.2 The Bandpass Calibrator

Although the same strategy in picking the phase calibrator was originally used to determine the bandpass calibrator, a more simpler scheme is now used. Forty-four bright, compact quasars are now monitored periodically and are tied to solar system objects with accurately known absolute flux density models, with ALMA at Bands 3 and 7. All are sufficiently bright to be used as bandpass calibrators (with a caveat given below). In addition,

many have sufficiently stable flux densities than can be followed to about 5 to 10% between frequencies 90 GHz to 400 GHz, and can be used as secondary flux density standards for an ALMA experiment. Nevertheless, some sources can vary by a factor of three in a month or less (egs. 3c454.3 in 2012), but the monitoring will be sufficiently frequent to remove such variables from consideration as soon as possible.

Source name	R.A.	Dec.	B3 flux (Jy)	B6 flux (Jy)	B7 flux (Jy)
J0106-405	01:07:24.60	-40:29:44.3	0.66	0.30	0.22
J0132-169	01:33:25.19	-16:50:24.3	1.29	0.89	0.78
J0237+288	02:38:43.07	+28:51:41.4	1.83	1.59	0.78
J0238+166	02:39:26.44	+16:40:34.2	0.91	0.96	0.63
3c84	03:20:44.96	+41:33:32.8	15.50	10.21	6.57
J0334-401	03:34:44.61	-40:05:28.7	0.97	0.98	0.34
J0423-013	04:23:58.75	-01:18:36.3	3.46	3.60	1.44
J0510+180	05:10:51.43	+18:01:37.3	3.63	2.09	2.18
J0519-454	05:20:13.94	-45:45:42.7	1.75	1.18	1.01
J0927+390	09:27:53.63	+38:58:36.4	4.74	3.31	1.73
J1037-295	10:37:54.09	-29:38:16.5	0.69	0.83	0.54
J1058+015	10:59:11.37	+01:29:35.7	1.68	1.58	0.80
J1107-448	11:07:46.03	-44:53:32.7	1.47	0.75	0.65
J1130-148	11:30:48.11	-14:53:56.9	1.30	0.78	0.49
J1146+399	11:47:40.03	+39:54:00.9	0.99	0.57	0.39
J1147-6753	11:48:11.56	-67:58:17.6	1.04	0.58	0.70
J1159+292	12:00:13.03	+29:10:11.8	0.73	1.20	0.45
3C273	12:29:48.23	+01:58:40.6	7.14	2.95	2.22
3C279	12:56:53.33	-05:51:43.3	18.54	13.07	6.30
J1337-129	13:38:23.37	-13:01:30.2	5.93	4.14	2.54
J1426+364	14:27:10.27	+36:21:43.0	0.34	0.17	0.15
J1427-421	14:28:48.32	-42:10:00.2	7.49	3.93	4.10
J1517-243	15:18:30.37	-24:25:14.5	1.68	1.29	0.82
J1550+054	15:51:15.97	+05:24:56.4	1.05	0.59	0.38
J1613-586	16:18:27.85	-58:50:12.9	1.40	1.17	0.54
3C345	16:43:26.41	+39:47:28.9	3.00	1.68	1.16
J1733-130	17:33:49.77	-13:05:11.8	2.39	1.66	0.91
J1751+096	17:52:12.51	+09:39:08.5	2.73	2.25	1.28
J1800+388	18:00:52.70	+38:48:56.1	0.18	0.06	0.11
J1924-292	19:25:44.10	-29:12:43.9	4.83	3.97	1.98
J2025+337	20:25:44.79	+33:46:06.1	2.16	1.05	0.93
J2056-472	20:57:14.66	-47:11:32.4	0.96	0.99	0.41
J2148+069	21:48:48.21	+07:01:46.0	2.39	1.01	0.63
J2157-694	21:58:16.56	-69:37:25.6	1.29	0.88	—
J2202+422	22:03:19.72	+42:20:58.8	8.91	2.72	3.15
J2232+117	22:33:19.12	+11:48:22.1	2.03	4.89	0.89
3C454.3	22:54:40.47	+16:13:32.5	3.21	2.47	1.72
J2258-279	22:58:52.60	-27:53:45.1	1.48	3.25	0.51
J2357-5311	23:58:37.84	-53:06:29.1	1.10	0.78	0.43

Table C.3: ALMA bandpass calibrator sources and their representative fluxes. Data downloaded on Sept 19, 2013.

The present list of bandpass sources is shown in Table C.3. Some changes will be made for extremely variable sources or by the addition of sources that have become strong at Bands 7 and 9.

The density of these sources in the sky is sufficiently large so that at least one will be within 40° of any target, and generally the brightest one that is not below elevation 30° during the experiment will be chosen. [Note: the elevation limit at the beginning of Cycle 1 will be 57° to avoid shadowing with the most compared

ALMA and Morita arrays.]

The channel to channel bandpass sensitivity for FDM mode with 3840 channels in a spectral window, especially at frequencies above 300 GHz, cannot reach the 50:1 SNR per channel unless one of the few very bright quasars is near the target. Since the instrumental frequency variations between consecutive channels are less than about 1 is no need to determine the bandpass calibration for each channel. However, the number of channels that can be averaged and still meet the 50:1 SNR specification is under investigation. Presently, 64 channels are averaged in order to find the bandpass calibrator that can meet the 50:1 SNR. If none of the above calibrators meets this level, the best one is chosen.

C.2.3 Other Calibrator Types

Absolute Flux density Calibration: An experiment will contain either a solar system object and/or a bandpass calibrator with a flux density, extrapolated from ALMA monitoring data, to the experiment date. Inclusion of both types are recommended if one of the eight solar system objects is available during the experiment. If extreme absolute amplitude accuracy is needed ($< 3\%$), additional amplitude-type calibrators should be included in the experiment.

Polarization Calibrators: The compilation of sufficiently strong calibrators that are also polarized more than about 4% is underway. These are needed to calibrator the leakage signal between each antenna/SPW X and Y polarization signals. Since only about 25% of the bandpass calibrators are sufficiently polarized, additional strong sources are being checked. More information is given in the polarization chapter of the technical handbook.

Astrometric Calibrators: The positional expectation for a typical ALMA observation is about 10% of the resolution, if there is sufficient SNR in the target. If higher precision is need, perhaps reaching 2% of the resolution, additional secondary phase calibrators are needed. More details will be forthcoming when ALMA opens up to astrometric-quality proposals.

Check Sources: An accurate assessment of the quality of Band 9 observations is difficult to make. ALMA staff can in cases like this add a check source in the experiment, a unresolved quasar with an accurate position. This source should be at about the same angular separation from the phase calibrator as the target, although it can be significantly weaker than the minimum phase calibrator flux density. After data editing and calibration, an image of the check source will provide information about the calibration quality and the coherence of the observations.

Appendix D

Acronym Dictionary

ACA	Atacama Compact Array
ACD	Amplitude Calibration Device
ACS	ALMA Common Software
ALMA	Atacama Large Millimeter/Submillimeter Array
AoD	Astronomer on Duty
AOS	Array Operation Site
APDM	ALMA Project Data Model
AQUA	ALMA Quality Assurance software
ARC	ALMA Regional Center
ASA	ALMA Science Archive
ASC	ALMA Sensitivity Calculator
ASDM	ALMA Science Data Model
AZ	Azimuth
BB	Baseband
BE	Backend
BL	Baseline
BLC	BaseLine Correlator
BWFN	Beam Width between First Nulls
CASA	Common Astronomy Software Applications package
CCA	Cold Cartridge Assemblies
CCC	Correlator Control Computer
CDP	Correlator Data Processor
CFRP	Carbon Fiber Reinforced Plastic
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
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

DTX	Data Transmitter module
EB	Execution Block
EL	Elevation
EPO	Education and Public Outreach
ES	Early Science
ESO	European Southern Observatory
FDM	Frequency Division Mode
FE	Frontend
FITS	Flexible Image Transport System
FLOOG	First LO Offset Generator
FOM	Fiber Optic Multiplexer
FOV	Field of View
FPGA	Field-Programmable Gate Array
FT	Fourier Transform
FWHM	Full Width Half Maximum
FWHP	Full Width to Half Power
GPS	Global Positioning System
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
JPL	Jet Propulsion Laboratory
LFRD	Low Frequency Reference Distributor
LLC	Line Length Corrector
LO	Local Oscillator
LO1	First LO
LO2	Second LO
LO3	Digitizer Clock Third LO
LO4	Tunable Filterbank LO
LORR	LO Reference Receiver
LS	Laser Synthesizer
LSB	Lower Sideband
LTA	Long Term Accumulator
MFS	Master Frequency Standard
ML	Master Laser
MLD	Master Laser Distributor
NGAS	New Generation Archive System
NRAO	National Radio Astronomy Observatory
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 Distributor
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
SED	Spectral Energy Distribution
SIS	Superconductor-Insulator-Superconductor Mixer
SL	Slave Laser
SNR	Signal-to-Noise Ratio
SPW	SPectral Window
SRON	Stichting Ruimte Onderzoek Nederland (Netherlands Institute for Space Research)
SSB	Single Sideband
2SB	Sideband separating Mixer
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable Filterbanks
TFB LO	Local Oscillator at the Tunable Filterbanks
TMADB	Telescope Monitor and Configuration DataBase
TP	Total Power
Tsys	System Temperature
T_{rx}	Receiver Temperature
USB	Upper Sideband
VLA	Very Large Array
VO	Virtual Observatory
WCA	Warm Cartridge Assembly
WVR	Water Vapor Radiometer
XF	Correlation-Fourier Transform Type Correlator
YIG	Yttrium-Iron Garnet Oscillator



The Atacama Large Millimeter/submillimeter Array (ALMA), an international astronomy facility, is a partnership of the European Organization 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 National Science Council of Taiwan (NSC) 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.

