

ALMA Cycle 4 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

Cycle	Version	Date	Editors	Remarks
1	1.0	2013-10-01	A. Hales, A. Lundgren, A. Biggs, J. Di Francesco, B. Vila Vilaro, S. Leon, D. Barkats, E. Fomalont, S. Asayama, S. Komugi	
2	1.1	2013-10-15	A. Lundgren, S. Leon, J. Hibbard, E. Müller, J. Di Francesco	
3	1.0	2015-03-20	A. Remijan, M. Adams, R. Warmels	
4	1.0	2016-03-22	J. Mangum and R. Warmels	
4	1.1	2016-09-07	J. Mangum and R. Warmels	Replace figs 7.9 and 7.10

Authors

Itziar de Gregorio - Chapter 1; Antonio Hales - Chapter 2; James Di Francesco - Chapter 3; Shin'ichiro Asayama - Chapter 4; Seiji Kameno - Chapter 5; Bill Dent - Chapter 6; Eric Villard - Chapter 7; Baltasar Vila Vilaro - Chapters 8, 10 and 12; Andy Biggs - Chapter 9; Ed Fomalont - Chapter 10; Erik Müller - Chapter 11; Liz Humphries - Chapter 13; Felix Stoehr - Chapter 14.

Contributors

Mark Adams, Eiji Akiyama, Paola Andreani, Yiping Ao, Denis Barkats, James Braatz, Crystal Brogan, James Chibueze, Stuartt Corder, Juan Cortes, Paulo Cortes, Carlos De Breuck, Jennifer Donovan Meyer, Darrel Emerson, Daniel Espada, Roberto Galvan-Madrid, Bunyo Hatsukade, Cinthya Herrera, John Hibbard, Akihiko Hirota, 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, 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, Kenichi Tatematsu, Ignacio Toledo, Rein Warmels, Al Wootten, Jorge F. García Yus, Martin Zwaan.

In publications, please refer to this document as:

ALMA Partnership, 2016, S. Asayama, A. Biggs, I. de Gregorio, B. Dent, J. Di Francesco, E. Fomalont, A. Hales, E. Humphries, S. Kameno, E. Müller, B. Vila Vilaro, E. Villard, F. Stoehr

ISBN 978-3-923524-66-2

Contents

1	Introduction	9
2	Array Components	11
2.1	The ALMA Telescope	11
2.2	The 12-m Array	12
2.3	The Atacama Compact Array	13
2.3.1	The 7-m Array	13
2.3.2	The TP Array	14
3	Principles and Concepts of Interferometry	15
3.1	Introduction	15
3.2	Single-dish Response	16
3.3	Visibilities and Aperture Synthesis	19
3.4	The Visibility or <i>uv</i> -Plane	22
3.5	Fields-of-view and Mosaics	23
3.6	Spatial Filtering	24
3.7	Multi-configuration Observations	26
3.8	Units and Conversions	28
4	Receivers	31
4.1	Local Oscillators and IF Ranges	33
4.2	The Cycle 4 Receivers	33
4.2.1	Band 3 Receiver	35
4.2.2	Band 4 Receiver	38
4.2.3	Band 6 Receiver	41
4.2.4	Band 7 Receiver	44
4.2.5	Band 8 Receiver	47
4.2.6	Band 9 Receiver	50
4.2.7	Band 10 Receiver	53
5	The Correlators	57
5.1	The 64-input Correlator	58
5.1.1	TDM mode	58
5.1.2	FDM mode	59
5.1.3	Correlation and realtime processing	60
5.2	The ACA Correlator	60
5.2.1	The 7-m Array	62
5.2.2	The Total Power Array	62
5.3	Digitizers	64
5.4	Online WVR Correction	64
5.5	Capabilities of the Correlators	64
5.5.1	Polarization	64
5.5.2	Spectral Setup	65

5.5.3	Time Resolution	66
5.6	Practical Performance	68
5.6.1	Sensitivity	68
5.6.2	Linearity	68
5.7	Final data product - the ASDM	69
6	Spectral Setups	71
6.1	Introduction	71
6.2	Frequency Definitions	73
6.3	Spectral Setups	73
6.3.1	Observing Frequencies for Continuum	74
6.3.2	Multiple Spectral Windows in the Same Baseband	76
6.3.3	Spectral Setups for Lines near the Edge of the Bands	76
6.3.4	Spectral Setup in Band 9 & 10 and DSB Considerations	77
6.4	Usable Bandwidth	78
6.5	Spectral Resolution	78
6.6	Spurious Signals	79
6.7	Doppler Setting and Velocity Reference Frames	79
6.8	Limitations and Rules for Spectral Setups in Cycle 4	80
7	Imaging with ALMA	81
7.1	Introduction	81
7.2	Cycle 4 Configurations	82
7.3	Shadowing	85
7.4	Resolution and Beam Shape	86
7.5	Response and Snapshot	90
7.6	Spatial Scale Filtering	92
7.7	Mosaicing	93
7.8	Multi-array and Multi-configuration Observations	94
7.9	Multi-array and Multi-configuration Imaging	97
8	Observing Modes	103
8.1	Observing Project Structure	103
8.2	Program Execution	105
8.3	The Observing Process	106
8.3.1	The source selection algorithm	106
8.4	Single Field Interferometry	107
8.5	Pointed Mosaic Observations	108
8.6	Single Dish (Total Power) Observations	108
8.7	Polarization	109
8.7.1	Sessions	112
8.8	Multiple Region Modes	113
8.9	Observation of Ephemeris Objects	113
8.10	Solar Observations	114
8.10.1	Solar Observing Modes	115
8.10.2	Array Observations	115
8.10.3	Single Dish Mapping	117
8.11	VLBI Observing Mode	117
8.11.1	General VLBI Considerations	118
8.11.2	ALMA Considerations for VLBI	119
8.12	High Frequency Observing	123

9	ALMA Sensitivity Calculator	125
9.1	Calculating the System Temperature	125
9.1.1	Sky temperature	125
9.1.2	CMB temperature	126
9.1.3	Receiver temperatures	126
9.1.4	Ambient temperature	127
9.1.5	DSB receivers	128
9.1.6	System temperature (OT version)	128
9.1.7	System Temperature (Applet and GUI version)	129
9.2	The Sensitivity Calculation	130
9.2.1	12-m and 7-m Arrays	130
9.2.2	Total Power Array	131
9.3	User Interface	131
9.4	Total Time Estimates	133
10	Calibration and Calibration Strategies	135
10.1	The Measured Visibility Function	135
10.1.1	The Data Format and the Correlator Model	135
10.1.2	The Visibility Function Dependence	135
10.2	Main Interferometric Calibration Strategies	136
10.2.1	Phase Referencing	136
10.2.2	Antenna-based Calibrations	136
10.2.3	Major Calibrations Affecting Interferometric Data	136
10.2.4	The Organization of the Calibrations	137
10.3	ALMA-Supplied a Priori Calibrations	138
10.4	Online Calibrations	140
10.5	Offline Calibrations	141
10.6	Special Calibrations	143
10.6.1	Phase Referencing, especially above 400 GHz	143
10.6.2	Phase Calibration with Narrow Bandwidths	144
10.6.3	Astrometry	145
10.6.4	Long Baseline Observations with Baselines > 2 km	147
10.6.5	Total-Power Calibrations	147
10.6.6	Acknowledgments	148
11	Quality Assurance	149
11.1	Cycle 4 "Best Efforts"	149
11.2	QA0	150
11.3	QA1	151
11.4	QA2	151
11.5	QA3	152
11.6	The Quality Assurance Report	152
11.7	QA Criteria	153
12	Data Flow and Logical Data Structure	155
12.1	Data and Control Flow	155
12.2	The ALMA Science Data Model	157
13	Pipeline Processing	159
13.1	Introduction	159
13.2	Pipeline In ALMA Operations	159
13.2.1	Pipeline Triggering	159
13.2.2	Pipeline Execution	159
13.2.3	Pipeline Runtime Database-querying	160
13.2.4	Pipeline and the ALMA Quality Assurance Process	160

13.2.5	Standard and Non-standard Observing Modes	164
13.3	Pipeline Processing of an ALMA OUS	164
13.3.1	Overview	164
13.3.2	Pause after Data Import	164
13.3.3	Output of the Pipeline Processing	164
13.4	Pipeline Initialisation & Calibration Table Editing	164
13.5	General Interferometric ALMA Data Processing	165
13.5.1	Data Import & Calibrator Flux Retrieval	165
13.5.2	Deterministic & Manual Flagging	165
13.5.3	Flagging lines in Solar System Object Flux Calibrators	165
13.5.4	Raw Data Flagging	166
13.5.5	Reference Antenna Selection	166
13.5.6	System Temperature Calibration & Flagging	166
13.5.7	Antenna Position Corrections	166
13.5.8	Offline WVR Calibration for 12m Array Data	166
13.5.9	Gain Flagging	166
13.5.10	Bandpass Calibration	166
13.5.11	Mapping Phase Gain Solutions	167
13.5.12	Setting the Fluxscale	167
13.5.13	Deriving the Fluxes of the Phase and Bandpass Calibrators	167
13.5.14	Time Gain Calibration	167
13.5.15	Applying the Calibration	167
13.5.16	Imaging	167
13.5.17	Exporting the Products	168
13.6	Low SNR Heuristics in Interferometric Pipeline Processing	168
13.6.1	Mapping the Gain Solutions from Broad Spectral Windows to Narrow Ones	168
13.6.2	Using the Gain Solutions Obtained from the Aggregate of all Spectral Windows	169
13.6.3	Additional Bandpass Pre-averaging	169
13.7	Total Power Data Processing	169
13.7.1	Data Import	169
13.7.2	Deterministic & Manual Flagging	169
13.7.3	T_{sys} and Sky Calibration Tables	170
13.7.4	Kelvin-to-Jansky Calibration Table	170
13.7.5	Applying the Calibration	170
13.7.6	Automated Line Identification and Baseline Subtraction	170
13.7.7	Baseline Flagging	171
13.7.8	Imaging	171
13.7.9	Exporting the products	171
14	Data Archiving	173
14.1	Introduction	173
14.2	Data Flow and Archive	173
14.3	PI Data and Data Delegation	174
14.4	Archive Query	175
14.5	Request Handler	175
A	Antennas	177
A.1	Design and Properties	177
A.2	Antenna Foundations	179
A.3	Antenna Transportation	181
A.4	Cryostat	181
A.5	Amplitude Calibration Device	184
A.5.1	Atmospheric Calibration Procedure	185
A.6	Water Vapor Radiometers	186

B	The LO and IF System	189
B.1	Functions of the LO & IF system	189
B.2	Summary of Operation	190
B.3	Frequency Generation and Distribution in ALMA	191
B.3.1	Reference and LO signal generation	193
B.3.2	LO1 signal generation and distribution	193
B.3.3	Line length corrections	194
B.3.4	The First Local Oscillator (LO1)	195
B.3.5	LO2 and the IF processor units and IF switch	196
B.3.6	Digitization and signal transmission	197
B.4	Other functions of the LO/IF	197
B.4.1	Delay corrections	198
B.4.2	Sideband suppression - LO offsetting	198
B.4.3	Sideband separation - 90 degree Walsh switching	199
B.4.4	Interference rejection - 180 degree phase switching	199
C	Calibration Sources	201
C.1	ALMA Catalog Description	201
C.2	Phase Calibrator Selection	202
C.3	Bandpass Calibrator Selection	202
C.4	Other Calibrator Types	204
C.4.1	Absolute and Relative Flux density Calibration	204
C.4.2	Polarization Calibrators	205
C.4.3	Astrometric Calibrators	205
C.4.4	Check Sources	205
D	Acronym Dictionary	207

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 4 operations will include standard and non-standard observing modes. Non-standard modes are observing modes that are less well characterized for which the data cannot be processed by the pipeline and need to be reduced manually by ALMA staff. Up to 20% of the PI observing time in Cycle 4 will be allocated to proposals requesting non-standard modes. Users should refer to Appendix A (Capabilities) in the ALMA Cycle 4 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 4 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 4. 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 4 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 4 and how projects are observed. Single field interferometry, mosaics, single-dish observations, polarization, multiple region modes, solar, VLBI, 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.

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-san 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 4, 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 (lat.=-23.02917 deg., long.=-67.754649 deg.), a site that normally offers the exceptionally dry and clear sky conditions required to observe at millimeter and sub-millimeter wavelengths². The ALMA antennas, weather stations, the two correlators and their computer interfaces, Local Oscillator generation hardware, timekeeping hardware, and the related array Real-Time Machine computer are all located at the 5000 meter altitude site referred to as the Array Operations Site (AOS). This site is connected via Gigabit fiber links to the Operation Support Facility (OSF), located at an altitude of 2900 meters, about 22 km from the AOS and 40 km from the town of San Pedro de Atacama. Science operations are conducted from the OSF and coordinated from the JAO Central office in Santiago. All three ALMA arrays are controlled via control software developed on the ALMA Common software³ (ACS).

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

The number of antennas in each array component (12-m, 7-m and TP Arrays), and the specific configurations available for an observing season (e.g. Cycle 4) will be published in the Capabilities section of the document *ALMA Proposer's Guide*. Complementary background information on ALMA and its capabilities for Early Science can be found in the document *A Primer for Early Science*. Both documents can be found on the link

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

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

³<http://www.eso.org/projects/almawdevelop/acs/>

<http://www.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 (dashed-orange circles) and 4 PM antennas (blue circles) are distributed in the right hand side of the image (highlighted). A few unoccupied stations can be seen, to which antennas of the 12-m Array can be moved by the transporter as the array is being reconfigured. At its most extended configuration, antennas in the 12-m Array will be about 16 km apart.

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 is possible. Each band receiver (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 cartridges are dual sideband (2SB) and Bands 9 and 10 are double sideband (DSB).

The Local Oscillator (LO) signals (Section B) are transmitted to the antennas on optical fibres 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 sub-sections of the 12-m Array (e.g., two ‘sub-arrays’) can simultaneous independent observations. Please 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 fibres (B). Total data rates are therefore 96 Gbits/s per antenna. With formatting, the bit rate is 120 Gb/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 5.2) is normally used for the 12-m Array, but it can also take inputs from the 7-m Array or TP antennas. It is an XF correlator (cross-correlates first, then Fourier transforms), but the correlator proper is preceded by Tunable Filter Banks, which makes it a digital hybrid XF correlator or FXF.

These can select sub-bands from the 2 GHz-wide basebands in a very flexible manner. From each of these, the correlator then generates 2016 cross-correlations and 64 auto-correlations (requiring 1.7×10^{16} operations per second). Either 2-, 3- or 4-bit resolution is used, and the sampling can be Nyquist or twice Nyquist. See Table 6.1 for an example table of 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.

2.3 The Atacama Compact Array

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

To achieve high-fidelity imaging of sources with emission on angular scales larger than those corresponding to the minimum spacing of the 12-m Array (the "Maximum Recoverable Scale" for that array - see Section 7.6), ALMA has been designed to include the Atacama Compact Array (ACA or Morita-san 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 12m 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 controlled via control software developed on the ALMA Common software (ACS) and is operated in a similar fashion to the 12-m Array. To achieve this unified operation, the ACA system is as compatible with the 12-m Array as possible at the level of hardware, interface, data, and observing modes. 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, or position switching. The raw time-series signals from the ACA antennas are processed in the ACA Correlator (see Section 5.2) to produce the cross-correlated and auto-correlated data.

2.3.1 The 7-m Array

The 7-m Array is composed of twelve 7-m diameter antennas designed and built by East Asia to the ALMA specifications (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 are originated identically to the ones sent to the 12-m Array, i.e., from inside the AOS building.

The ACA Correlator is normally used for the ACA, and can work with two sub-arrays. It is an FX correlator (Fourier transform first, then cross-correlate) with 3-bit input and 4 bits in the correlation. The correlator

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

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 9m to ~ 30 m (Chapter 7). Each cartridge (Chapter 4) receives two orthogonal polarizations.

2.3.2 The TP Array

The TP Array can fill in baseline coverage from 0m to about 12m, complementing the 7-m and 12-m Array's baseline coverage. It consists of four 12-m diameter antennas built by East Asia (A). The specifications of the TP antennas are almost identical to the ones for the 12-m Array. The TP antennas are located on stations surrounding the 7-m array. The TP antennas will be fitted with a nutating subreflector (also known as wobbler or chopper) in full operations, but is not yet employed. 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.

Due to the poorer point-source sensitivity of the 7-m Array, during Full Operations, the TP Array will be routinely used in the calibration observations of the 7-m Array, but this is not yet implemented. Since the 7-m Array is quite compact, atmospheric phase fluctuations will be smaller than for the 12-m Array.



Figure 2.2: The Morita Array - In remembrance of Professor Koh-Ichiro Morita. Koh-ichiro Morit, 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)
- *Millimeter Interferometry: Proceedings of the IRAM Millimeter Interferometry School*, ed. A. Dutrey, available online at <http://iram.fr/IRAMFR/IS/IS2002/archive.html>.

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

A range of discrete angular scales can be sampled by including many pairs of antennas in an array. Importantly, by tracking a source across the sky, the rotation of the Earth can be used to change the projected separations of the antenna pairs, allowing more angular scales to be sampled. An ensemble of the data, i.e., sinusoids of various amplitude and phase, can be then “summed” via the Fourier transform to produce an image of the sky brightness distribution. How well this image reflects the actual sky brightness distribution depends on how completely the relevant angular scales have been sampled. 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.



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.

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. (Note that ALMA receivers are constructed with two independent receptors so that both modes of polarization can be detected simultaneously.) 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 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

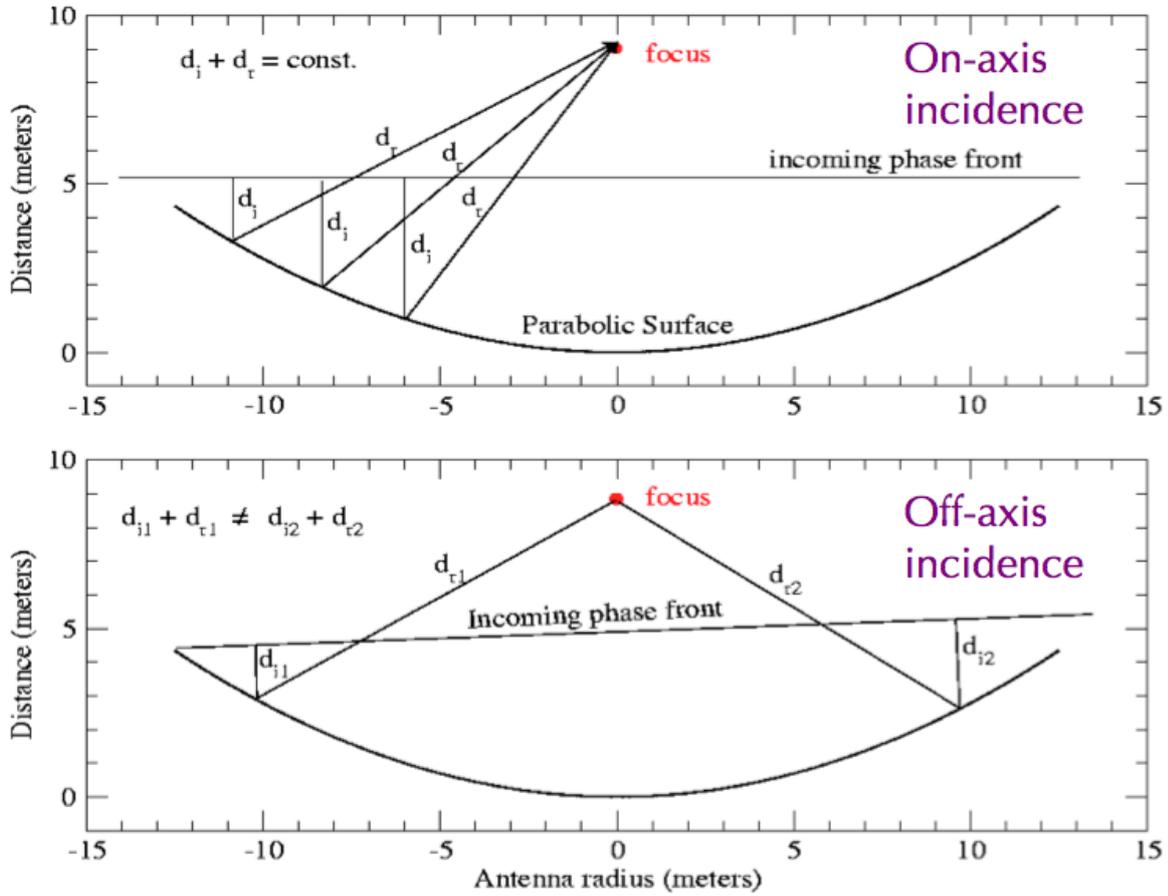


Figure 3.2: Schematic of an incoming plane-parallel wavefront reflecting off a antenna of diameter D and being brought to a focus. The top panel shows the case for a wavefront arriving on-axis. The bottom panel shows the case for a wavefront arriving off-axis. Note that the paths of incident EM power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

EM power is very distant, the EM power arrives at the primary surface essentially as plane-parallel wavefronts. Note that the antenna surface is parabolic in shape, so the path that each part of the front travels to the focus is constant. With zero path difference, the EM power arriving on-axis is coherently summed at the focus. This arrangement is only true, however, along the axis of the antenna. In the lower panel of Figure 3.2, the case for EM power arriving from an off-axis direction is shown. In this situation, the EM power does not add as constructively. In addition, the diameter of the antenna projected along the off-axis direction is less than the true diameter, decreasing the amount of power received from that direction. As a result, the antenna power response, i.e., its relative sensitivity, will be less than that found on-axis. In particular, at the off-axis angle of λ/D radians, 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:

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

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 Full Width Between Nulls (FWBN), and is given by:

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

Half the FWBN of the primary beam, $\sim 1.22 \lambda/D$, is considered the Rayleigh resolution of the antenna, i.e., its ability to distinguish objects on the sky separated by some angular distance. For convenience, the antenna power response is typically normalized to 1.0 along the axis. Figure 3.3 illustrates the antenna power response in one dimension (in log units); on the actual sky, the antenna power response is two-dimensional, and is obtained by rotating the function shown in Figure 3.3 about its central axis.

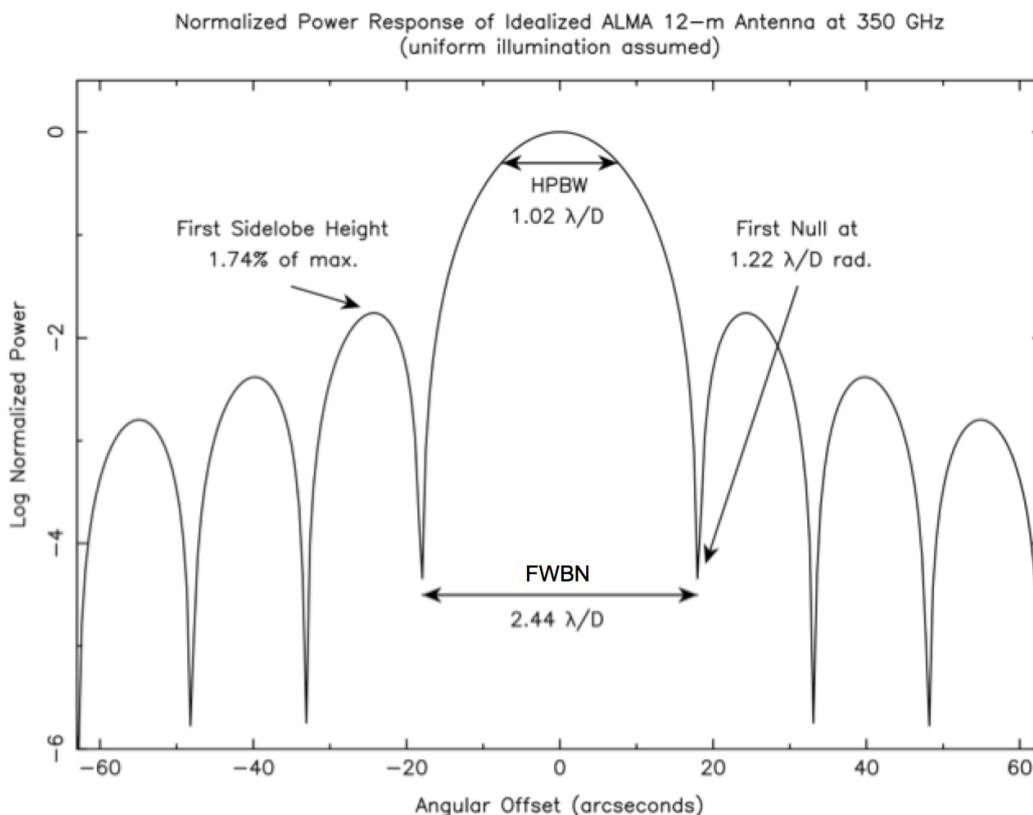


Figure 3.3: Normalized 1-D antenna power response for a 12 m antenna uniformly illuminated at 300 GHz. The power is in log units to emphasize the sidelobes. The HPBW of the primary beam is $\sim 1.02 \lambda/D$ and the FWBN is $\sim 2.44 \lambda/D$. The angle of the first null, i.e., the resolution, is $0.5 \times \text{FWBN} = \sim 1.22 \lambda/D$. Note that the HPBW measured from actual 12 m ALMA antennas is $\sim 1.13 \lambda/D$ because the illumination is not uniform.

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 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:

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

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 separations 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 parabolic surface with N elements has $N(N - 1)/2$ pairs of elements, and these could be observed sequentially to “synthesize” a measurement by a large filled-aperture telescope. Alternatively, numerous pairs of antennas, with each antenna considered an element, can be distributed to positions at distances much larger than it is possible to build a single filled-aperture telescope, and the signals received by these antennas can be combined in phase to approximate the resolving power of a single filled-aperture telescope.

The above situation only describes the emission received on-axis from antenna pairs. Of course, as noted above, emission 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., electronically) 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

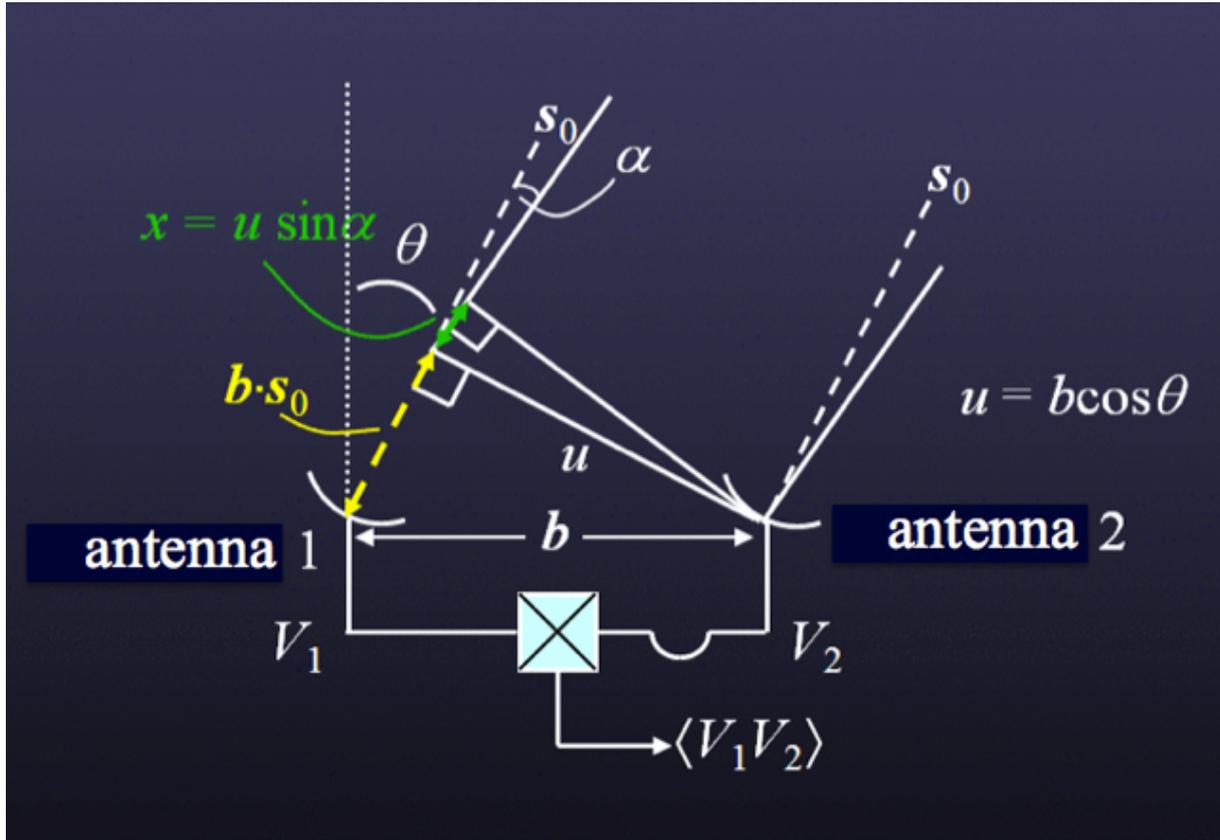


Figure 3.4: An ideal 1-D two-antenna interferometer consisting of two antennas, 1 and 2, separated by physical distance (i.e., a baseline) b . The antennas are both pointed towards a sky location given by s_o , which is at an angle θ from the meridian. The projected distance between the two antennas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

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:

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

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:

$$\langle V_1 V_2 \rangle = \left\langle \iint V_1(l, m) dl dm \iint V_2(l, m) dl dm \right\rangle \quad (3.13)$$

Under the assumption that signals emanating from different parts of the sky \ddot{A} 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:

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

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

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

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

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

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

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

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

3.4 The Visibility or 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:

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

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

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 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 samples of a complex-valued function with Hermitian symmetry, a single sampling gives two visibilities, one at (u, v) and its complex conjugate at $(-u, -v)$. To recover the true sky brightness distribution, however, knowledge of the distribution of visibilities across the uv -plane would be 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. Assuming the source emission is not variable, 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 earlier, such antennas have their own power response on the sky $P_N = \mathcal{A}(l, m)$ (e.g., the Airy function for an unobstructed, uniformly illuminated aperture). Indeed, the individual antenna response fundamentally limits the extent of an interferometric image made with a single pointing. In practice, the HPBW of the primary beam serves as the “field-of-view” of the single-pointing interferometric image. Moreover, $\mathcal{A}(l, m)$ 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 close to uniform sensitivity with a minimum number of pointings if the positions observed are arranged in a grid of equilateral triangles spaced by $\lambda/(\sqrt{3}D)$, where D is the diameter of the antenna. With this spacing, the fall-off of the primary beam response at one pointing is made up by the responses of the primary beams at adjacent pointings, except of course at the edge of the mosaic. Also, this grid allows Nyquist sampling to be achieved and all information can be retrieved. Mosaics can be made with adjacent positions that are spaced either closer or more distant, with non-uniform sensitivities. Mosaics provide increased areal coverage but 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 crucial 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 $\mathcal{A}(l, m)$) 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*. It is important to distinguish this beam from the single-dish response function $\mathcal{A}(l, m)$, which in the interferometry context is called the primary beam. The image resulting from the Fourier transform of a finite number of visibilities, $I^D(l, m)$, is sometimes referred to as the *dirty image*.

The measure of how similar an image is to the true sky distribution is sometimes referred to as *image fidelity*. Image fidelity depends on the specifics of coverage of the 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 (e.g., see Equation 3.25). Moreover, the complex visibilities of the point source have the same

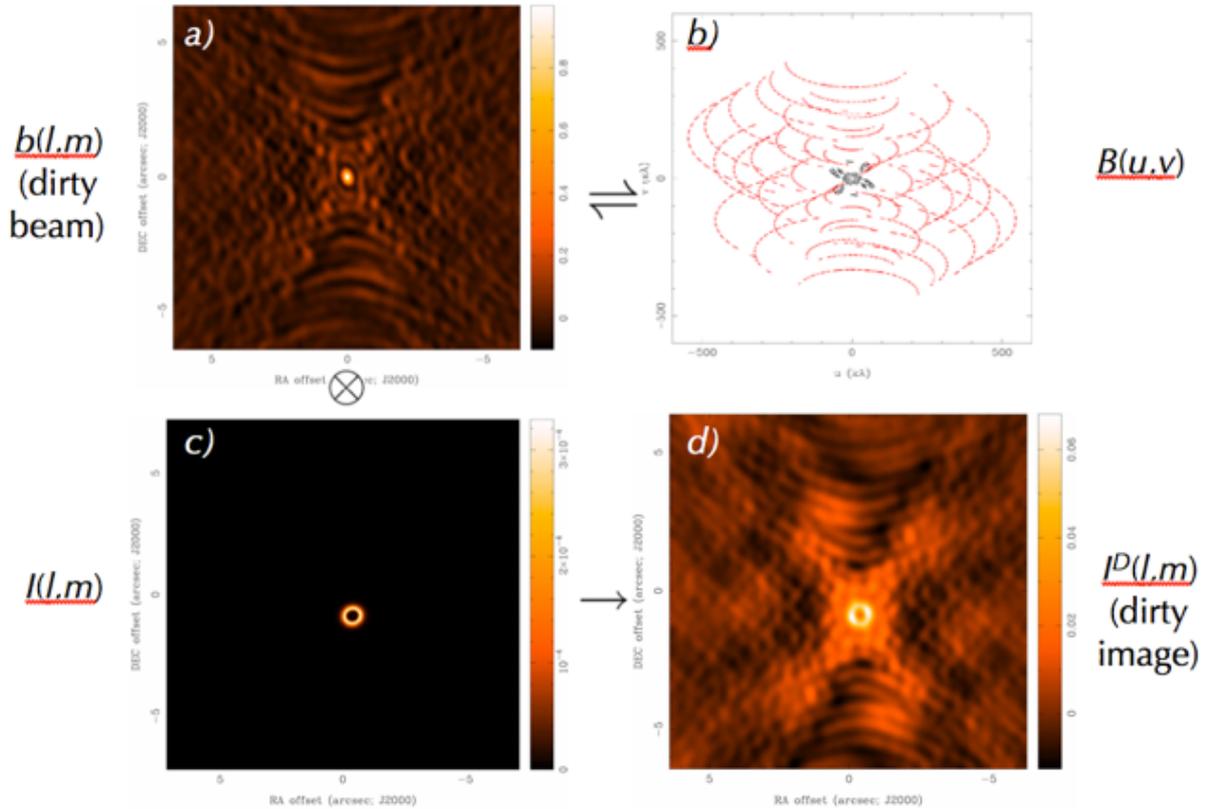


Figure 3.5: Imaging concepts. *Panel a (upper left)*: Example of a dirty beam, $b(l,m)$. *Panel b (upper right)*: The related ensemble of discrete points sampled in the uv -plane, $B(u,v)$. The black points were obtained from a compact configuration while the red ones were obtained from an extended configuration.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

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 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 7.6) and L_{max} is the longest baseline in the array.

Another important limitation of interferometric array observations is insensitivity to large angular scales. This insensitivity arises because interferometric arrays alone cannot sample spatial frequencies lower than those 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. The resulting visibilities are distorted, e.g., the antenna in the rear receives less power than the one in front. This situation is called *shadowing* and affected data are typically removed from the ensemble of observed visibilities.

Figure 3.6 illustrates the idea of spatial filtering using simulations of 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

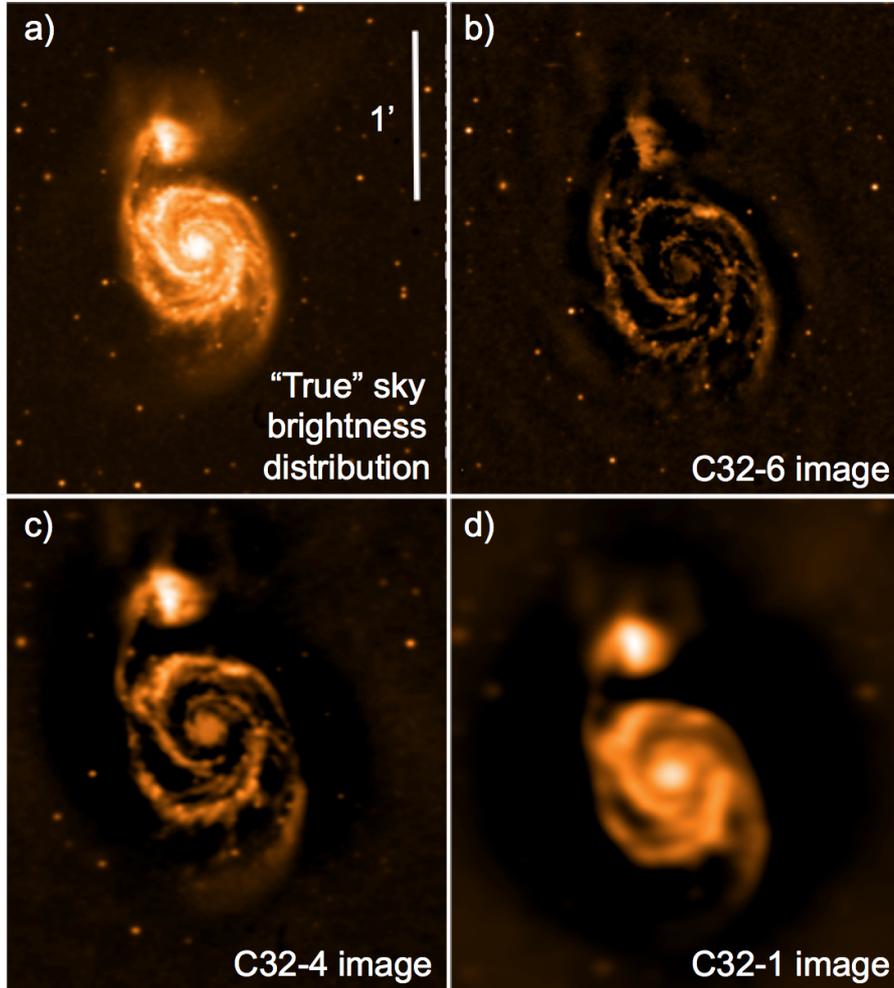


Figure 3.6: Examples of spatial filtering using the CASA task *simobserve* and actual ALMA configurations from Cycle 1. *Panel a (upper left)*: An optical image of the galaxy M51 used as a template for a true sky brightness distribution for the simulations. The frequency of the emission has been changed to 100 GHz, the image size has been scaled to $\sim 3' \times 3'$, and its declination has been changed to -40° to allow ALMA observations to be simulated. For the simulations, the galaxy was “observed” over a mosaic of 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 (Cycle 1) 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 $36.4''$.

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. 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. In addition, the accuracies of the astrometry and calibration of the different datasets are crucial. 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 4, 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.

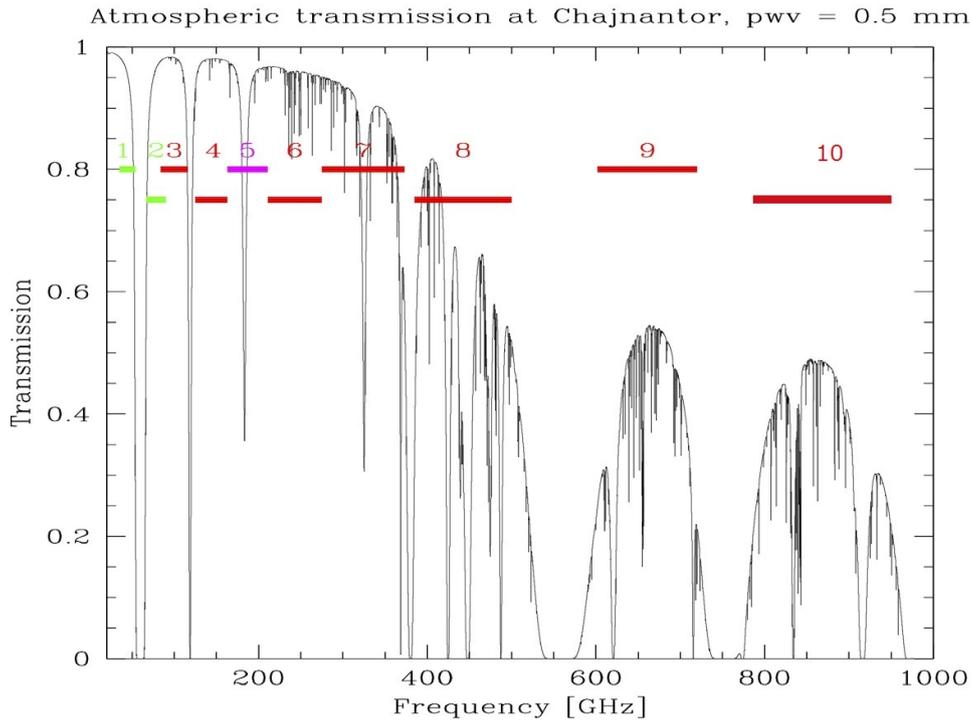


Figure 4.1: The ten ALMA receiver bands. Receiver bands for Cycle 4 are shown in red superimposed on a zenith atmospheric transparency plot at the AOS for 0.5 mm of PWV.

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 6.4) **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 6.4). **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 6.3.4 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 6.4).

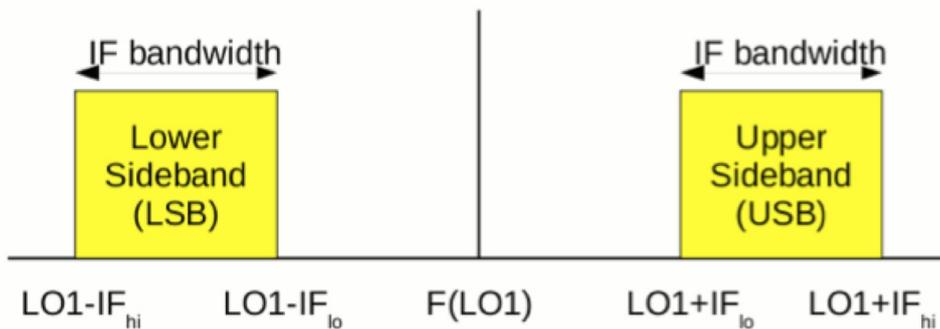


Figure 4.2: IF ranges for the two sidebands in a heterodyne receiver.

4.2 The Cycle 4 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 4, 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.

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

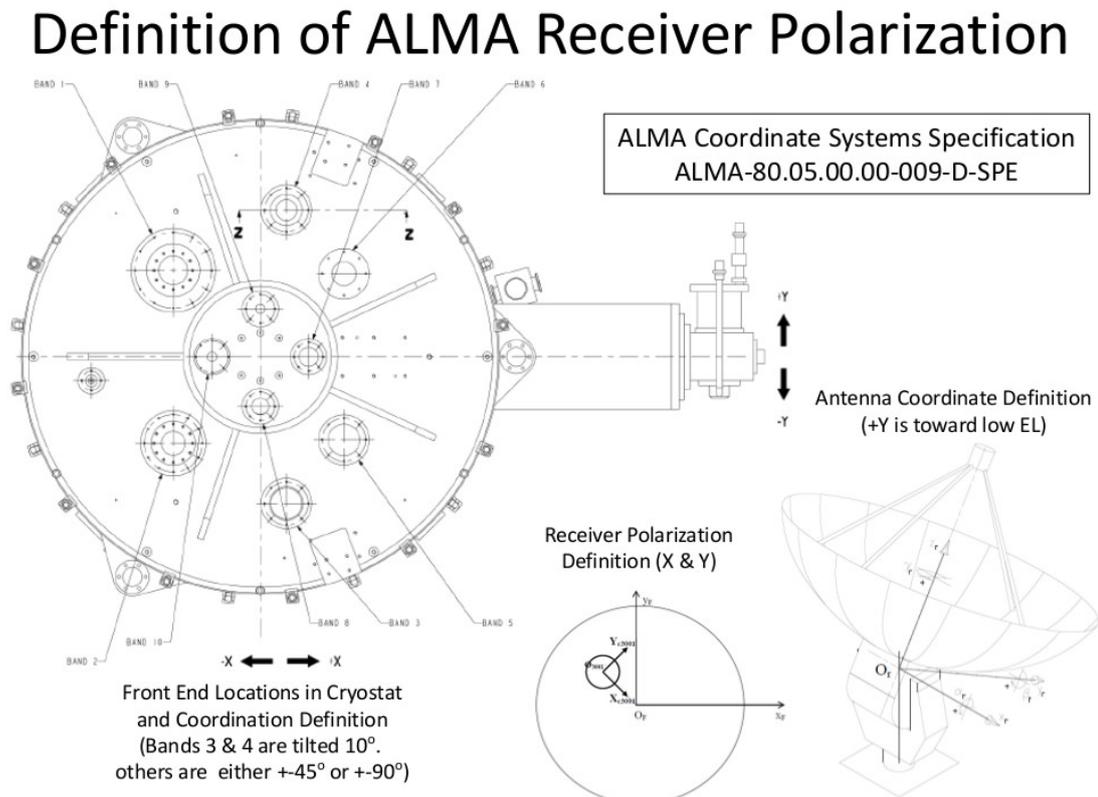


Figure 4.3: The layout of the ALMA receiver bands within the ALMA cryostat. Also shown is the orientation of the receiver polarization position angles.

4.2.1 Band 3 Receiver

Band 3 is the lowest frequency band available in Cycle 4, 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 4.4). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the SIS mixers.

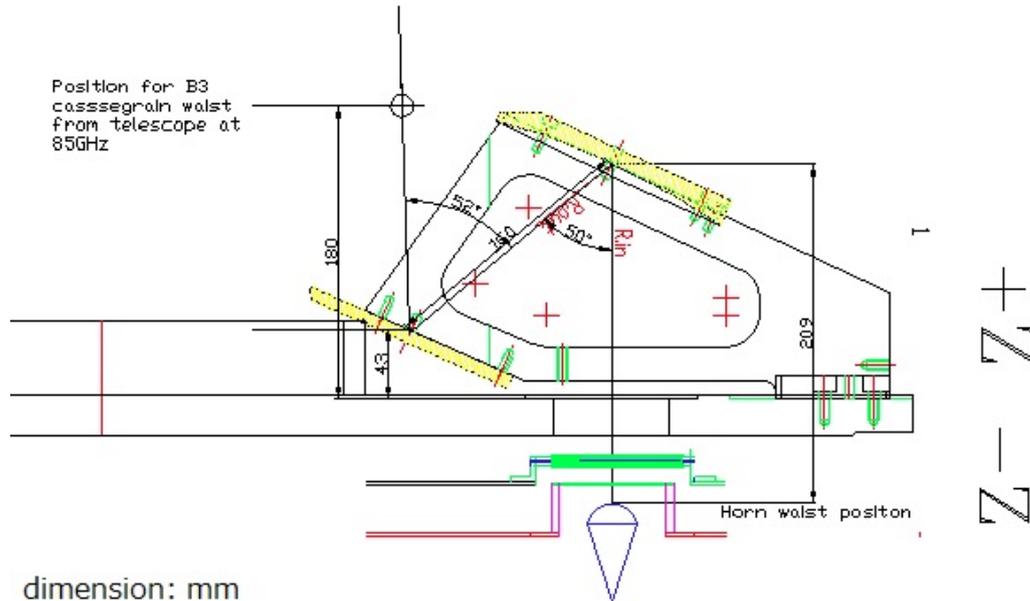


Figure 4.4: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the solid line, and the Cassegrain focus is shown by the small circle to the upper right.

A block diagram of the Band 3 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.5. The 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.

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 4.6) 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 4.7).

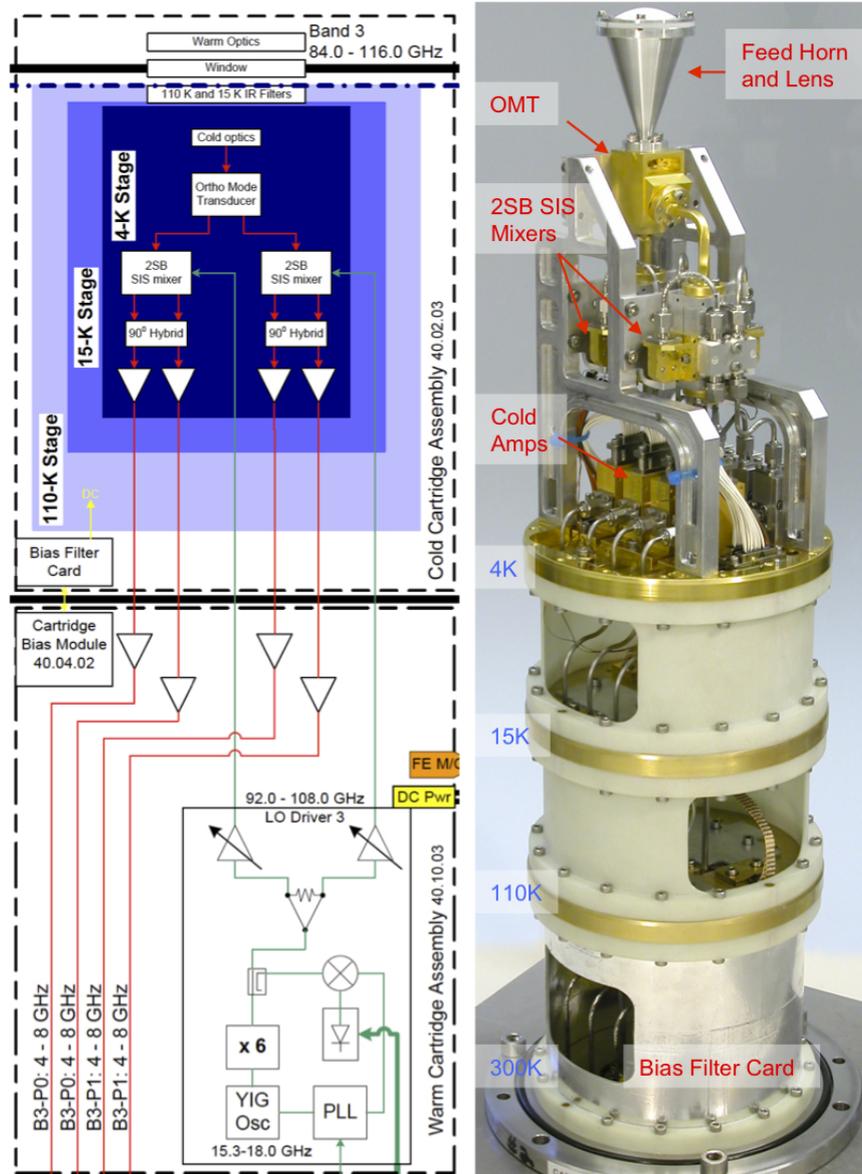


Figure 4.5: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges were constructed in Canada at NRC-HIA, Victoria.

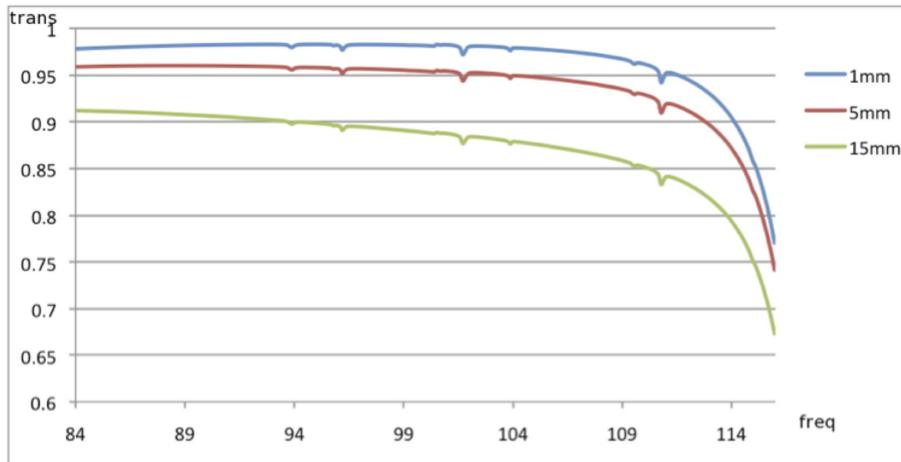


Figure 4.6: Band 3 zenith transmission for 1, 5 and 15 mm of PWV. Frequency is in GHz.

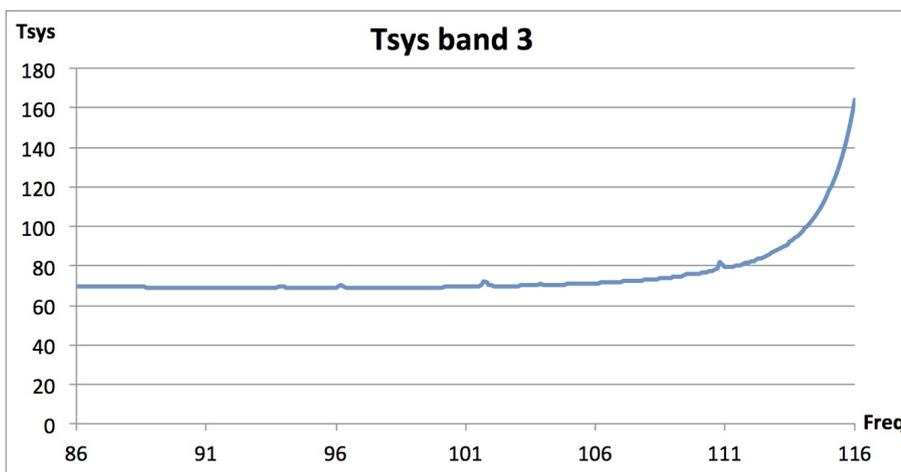


Figure 4.7: Typical system temperature (T_{sys}) at zenith for Band 3 with 1.262 mm of PWV. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 4.8). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the 2SB SIS mixers.

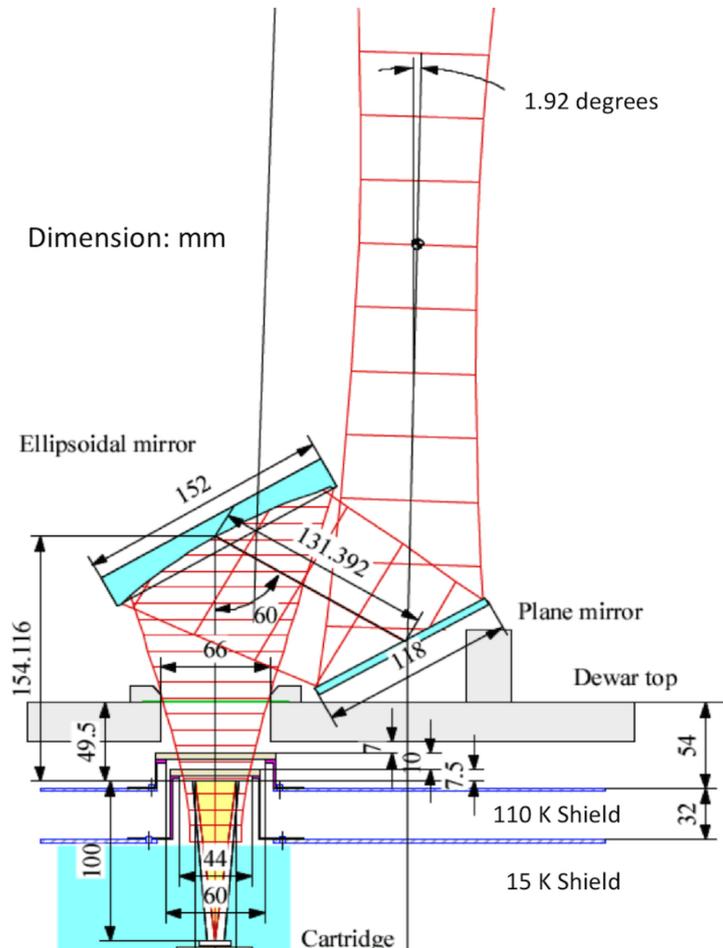


Figure 4.8: Optical layout of the Band 4. Red indicates Band 4 Gaussian beam of 5 times the beam width.

A block diagram of the Band 4 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.9. 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 4.10 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 4.11.

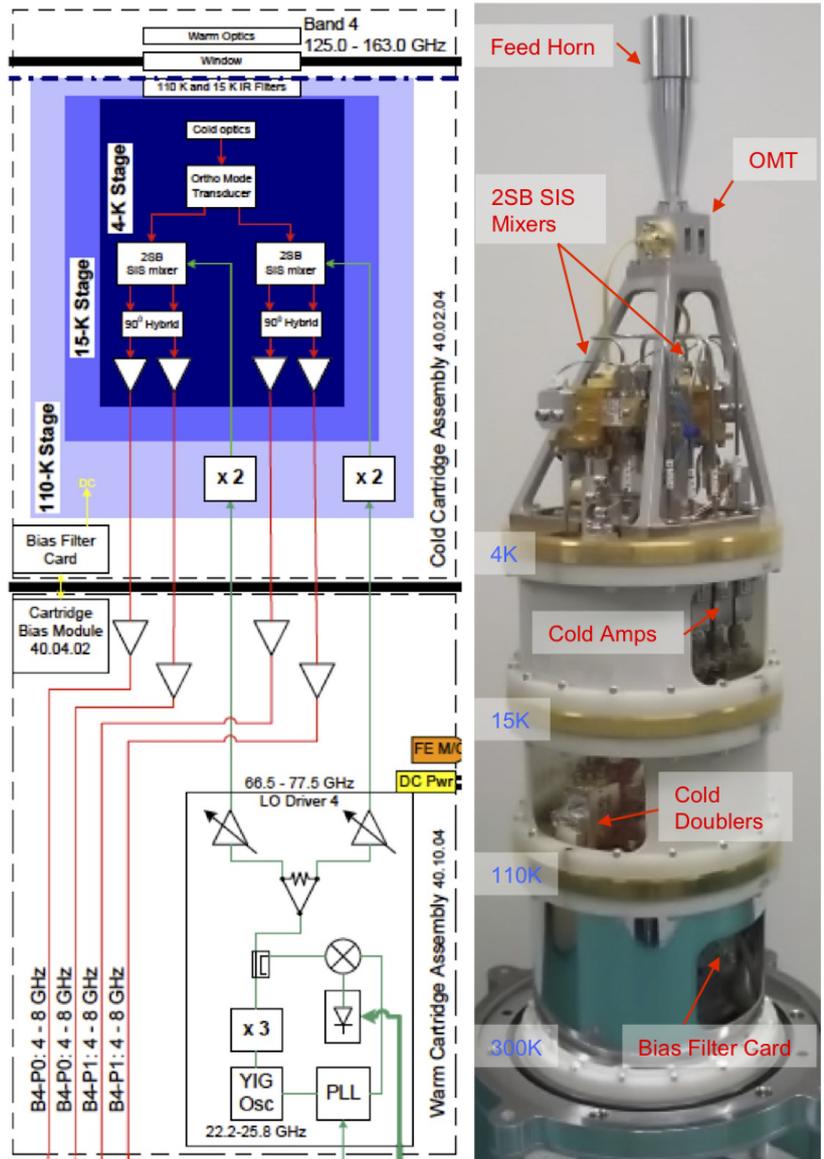


Figure 4.9: Block diagram of the Band 4 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 4 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 4 cartridges are constructed at the Advanced Technology Center (ATC).

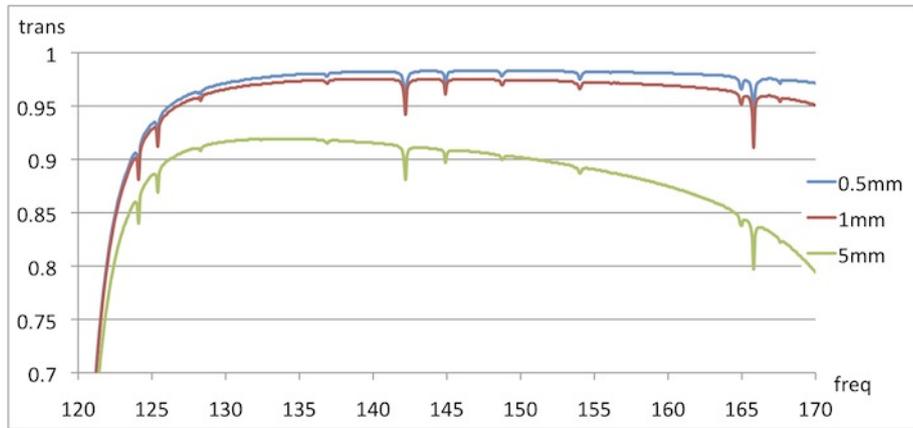


Figure 4.10: Band 4 zenith transmission for 0.5, 1 and 5 mm of PWV. Frequency is in GHz.

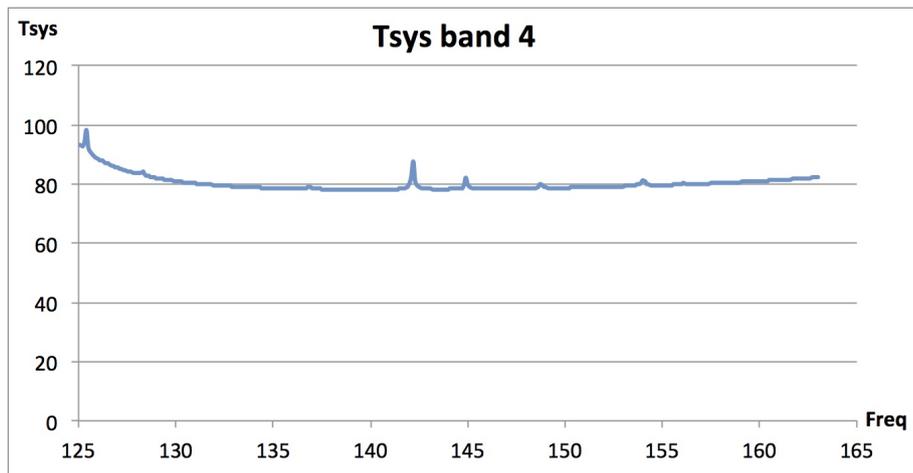


Figure 4.11: Typical system temperature (T_{sys}) at zenith for Band 4 with 1.262 mm of PWV. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 4.12). 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 4.13.

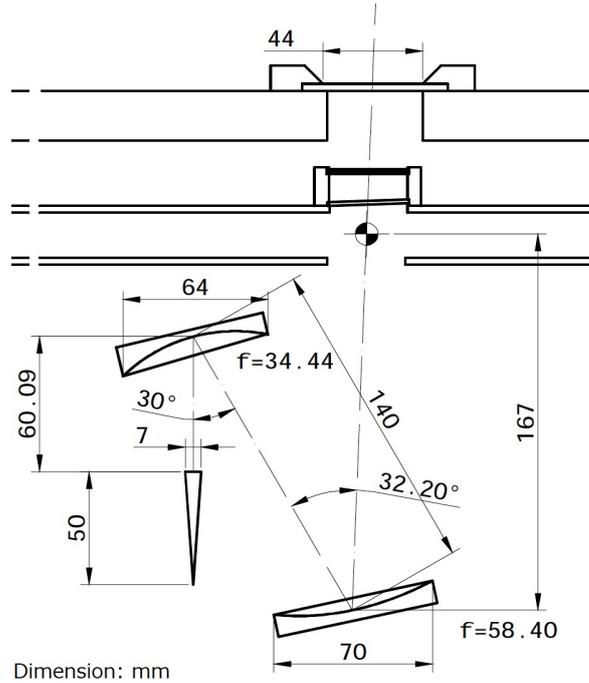


Figure 4.12: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocker. Note the slightly inclined 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 4.14 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 4.15.

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

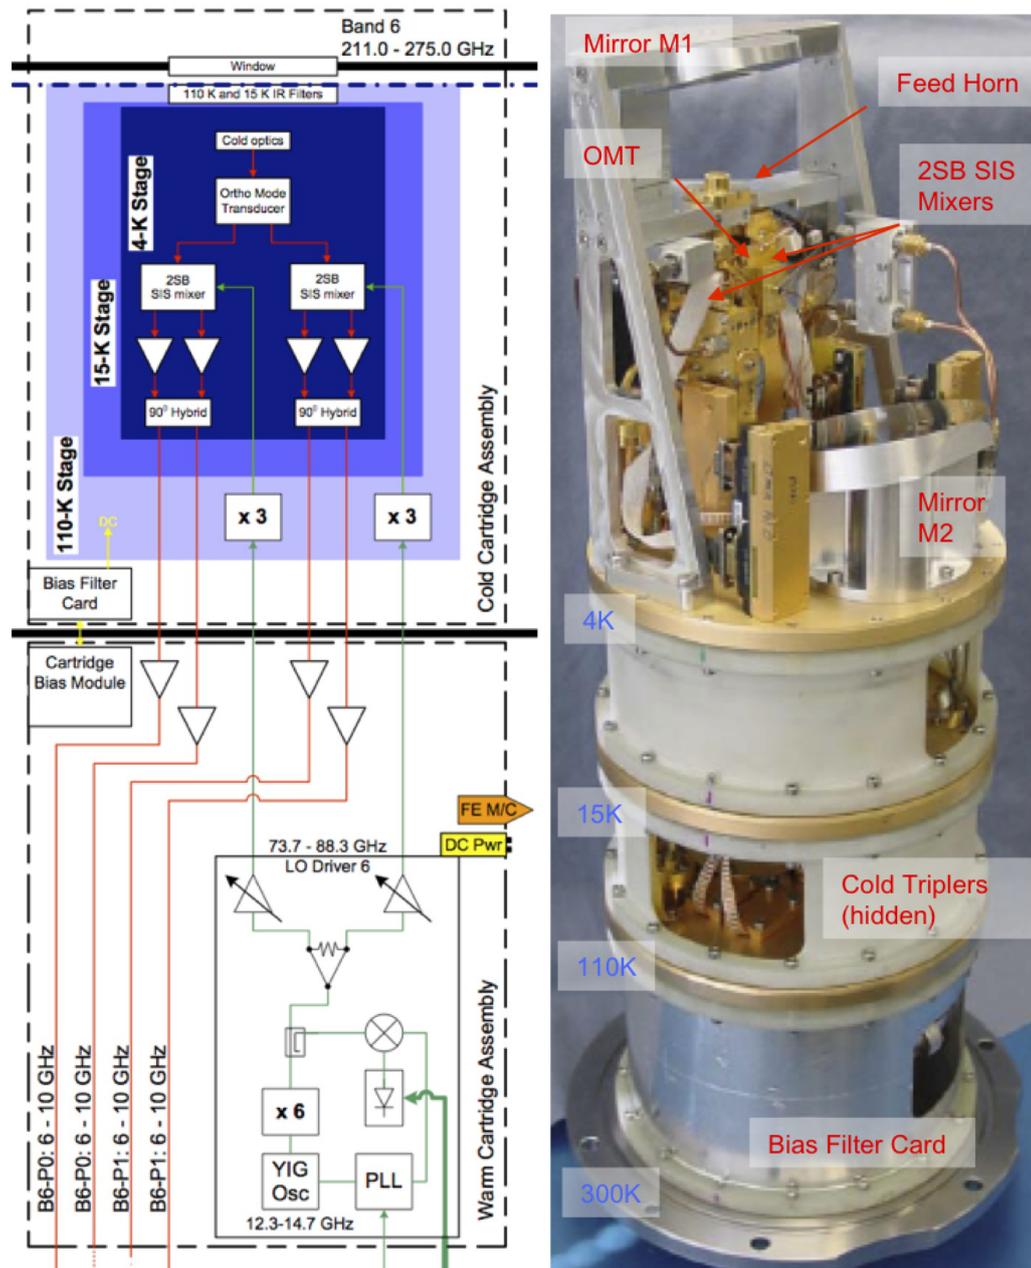


Figure 4.13: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra $\times 3$ multiplier inside the cryostat. The Band 6 cartridges were built at NRAO, Charlottesville. Note that the IF output range is actually 5-10 GHz. The range shown is the one recommended for continuum observations (see text).

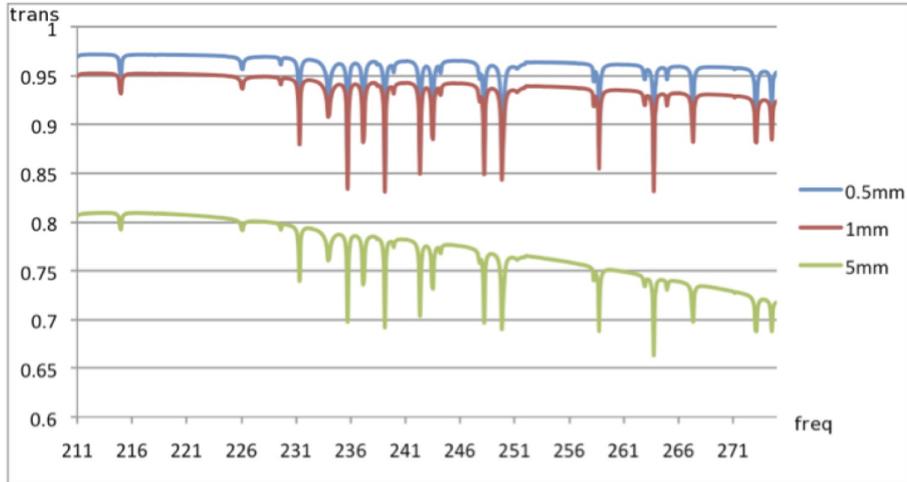


Figure 4.14: Band 6 zenith transmission for PWV=0.5, 1 and 5 mm. Frequency is in GHz. Most of the narrow absorption lines are from ozone.

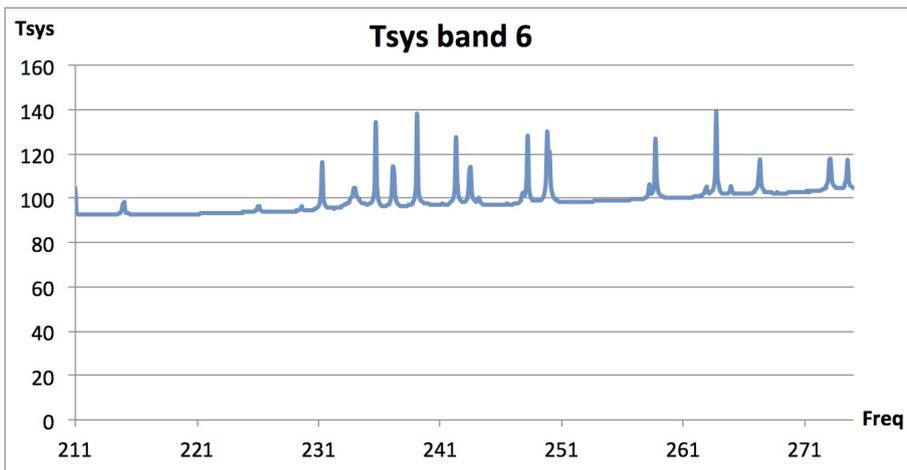


Figure 4.15: Typical T_{sys} at zenith for Band 6 with 1.262 mm PWV, based on measured values of the receiver temperatures. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature 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 4.16). A block diagram of the Band 7 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.17.

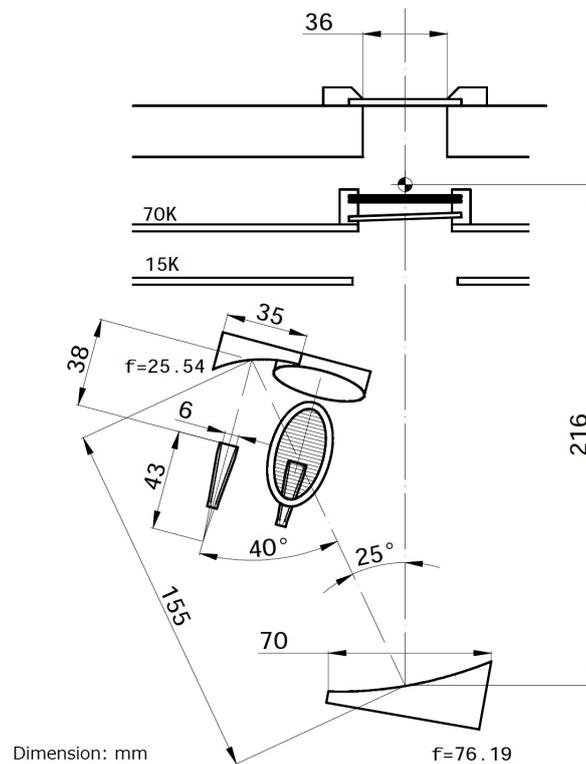


Figure 4.16: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid.

The atmospheric transmission in Band 7 is shown in Figure 4.18 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 4.19. 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.

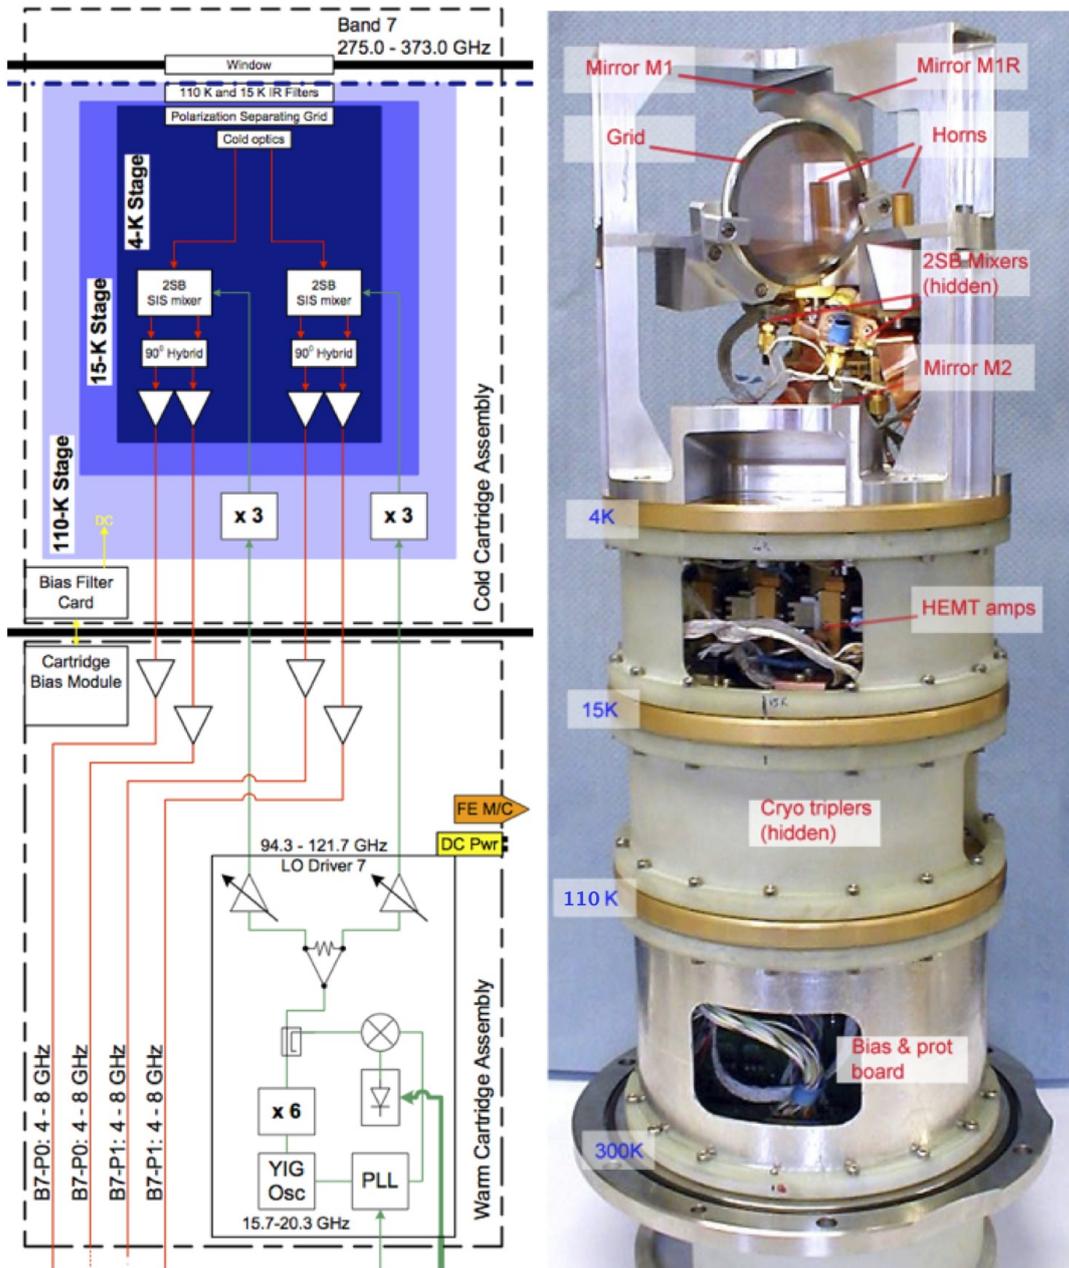


Figure 4.17: Band 7 frontend receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built at IRAM in France.

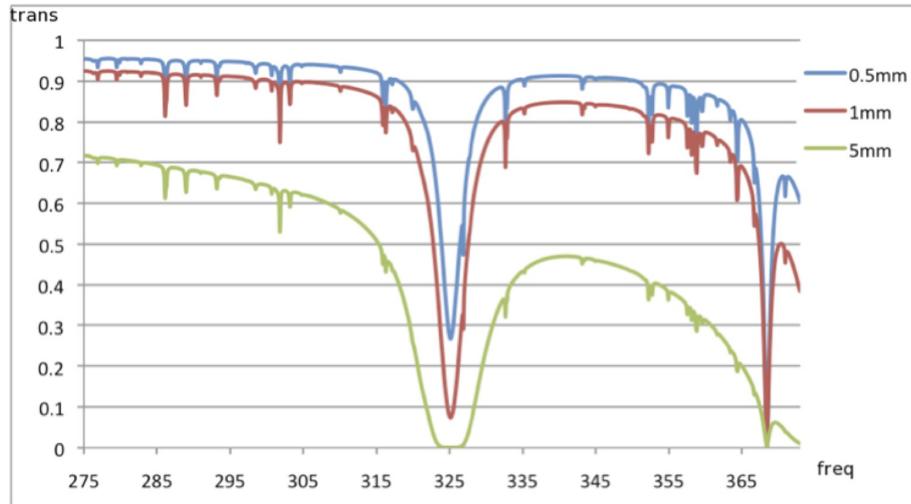


Figure 4.18: Band 7 atmospheric zenith transmission for PWV=0.5, 1.0 and 5.0 mm. Frequency is in GHz. The deep atmospheric absorption at 325 GHz is due to water, and the less prominent absorption feature at 369 GHz is due to oxygen.

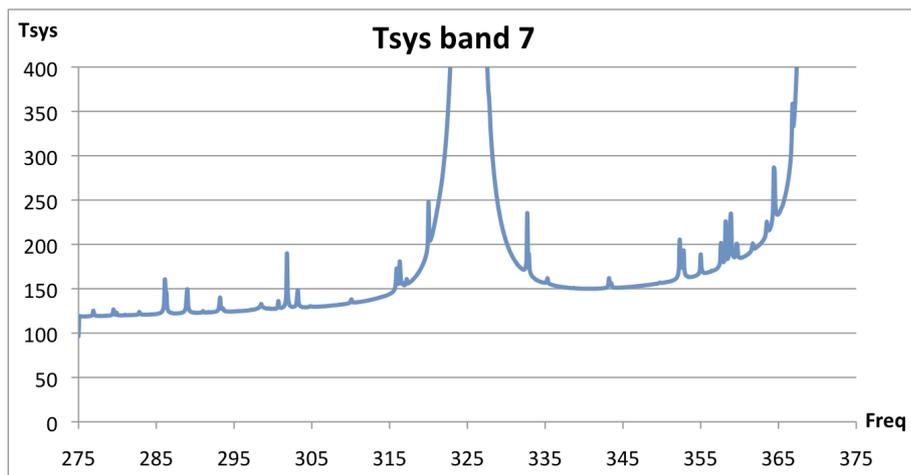


Figure 4.19: Typical T_{sys} at zenith for Band 7 with PWV=0.913 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 4.20).

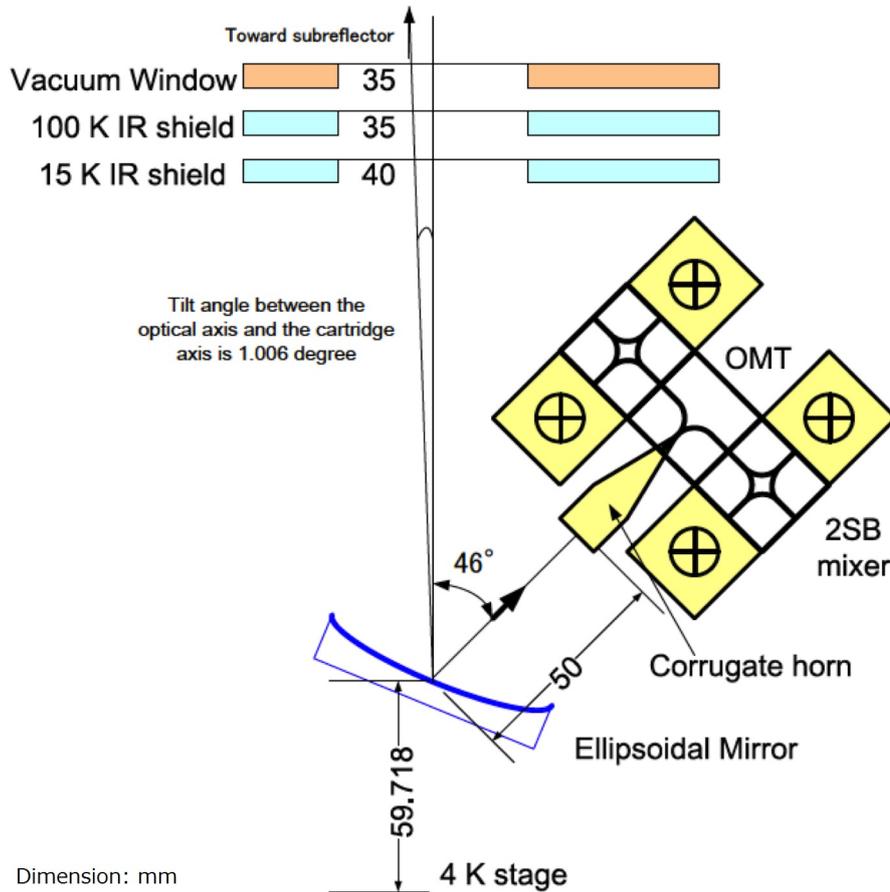


Figure 4.20: Optical layout of the Band 8.

A block diagram of the Band 8 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.21. 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 4.22 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 4.23.

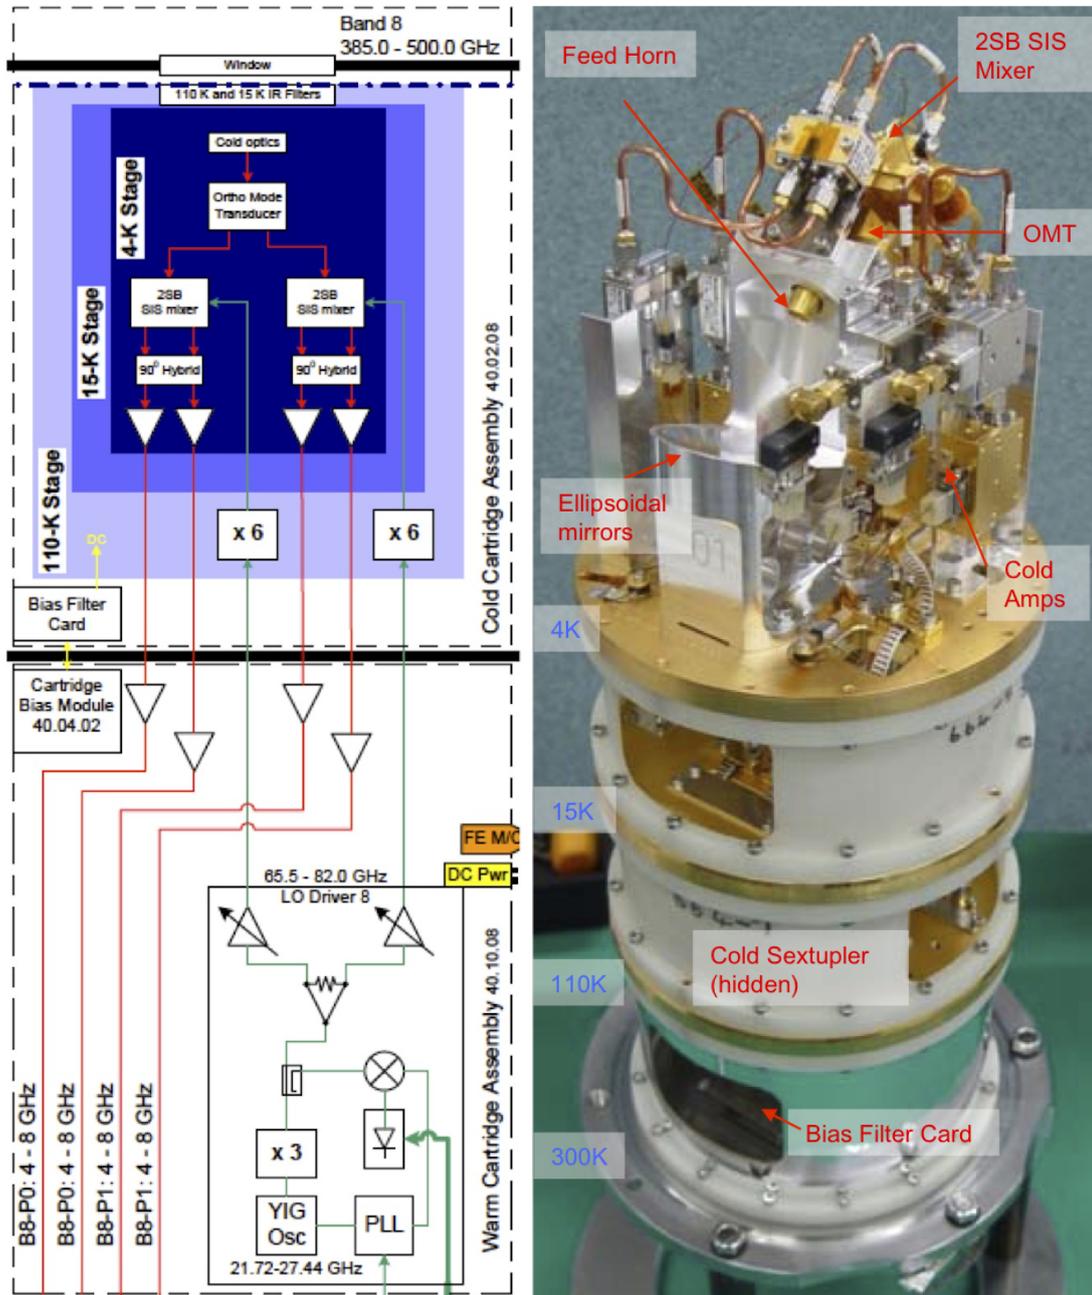


Figure 4.21: Block diagram of the Band 8 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 8 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 8 cartridges are constructed at the Advanced Technology Center (ATC).

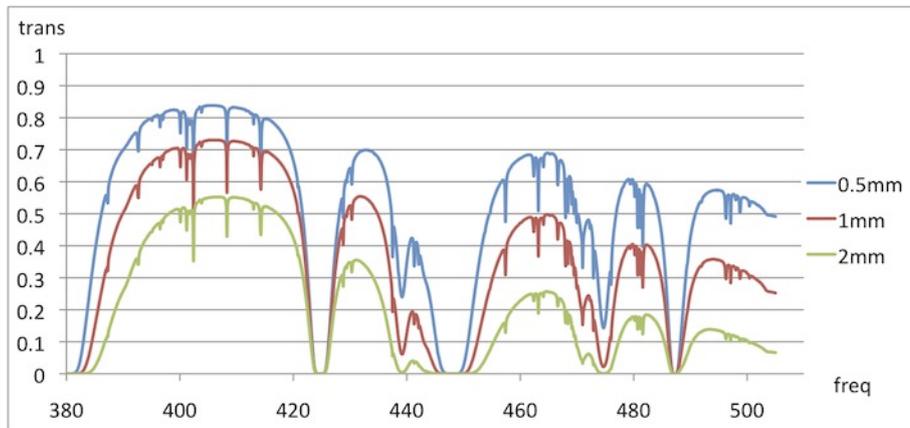


Figure 4.22: Band 8 atmospheric zenith transmission for PWV=0.5, 1.0 and 2.0 mm. Frequency is in GHz. The atmosphere in the Band 8 frequency range has some deep absorption by water and oxygen.

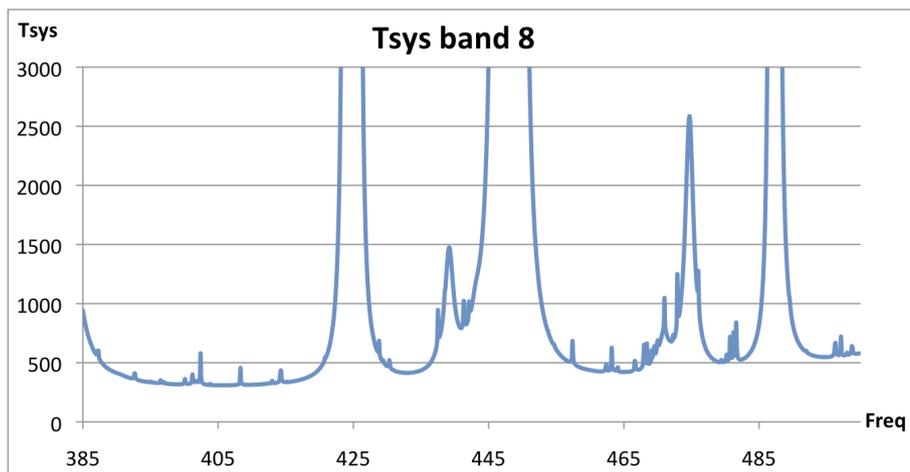


Figure 4.23: Typical T_{sys} at zenith for Band 8 with PWV=0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 4.24).

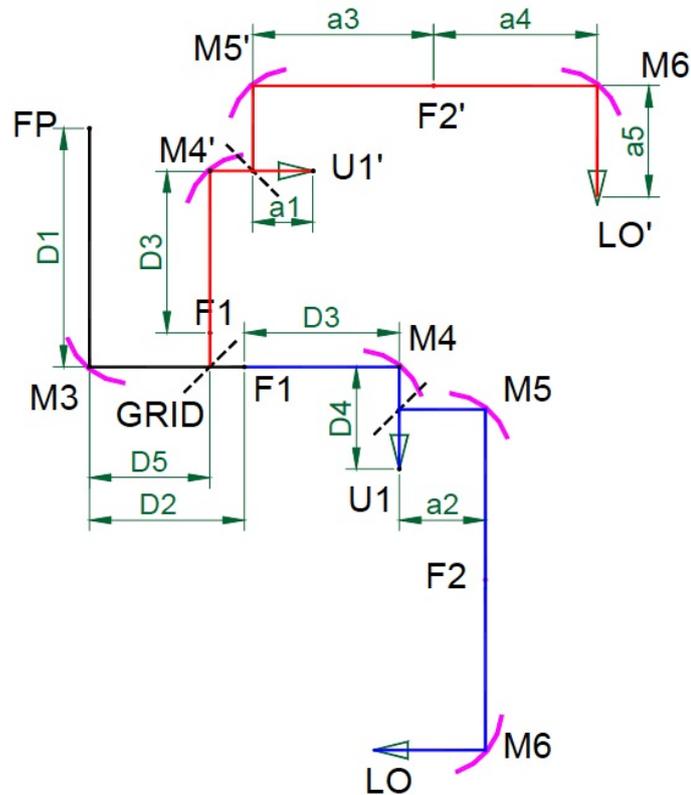


Figure 4.24: Basic Band 9 optics layout. The signal path is symmetrical for the vertical polarization (P0) and the horizontal polarization (P1). For P1 the path is as follows: a telescope focal point followed by mirror M3, grid, mirror M4, a beam splitter for LO insertion and finally the mixer horn U1. 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 4, 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 6.4), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.25.

The Band 9 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.26 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 4.27 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.

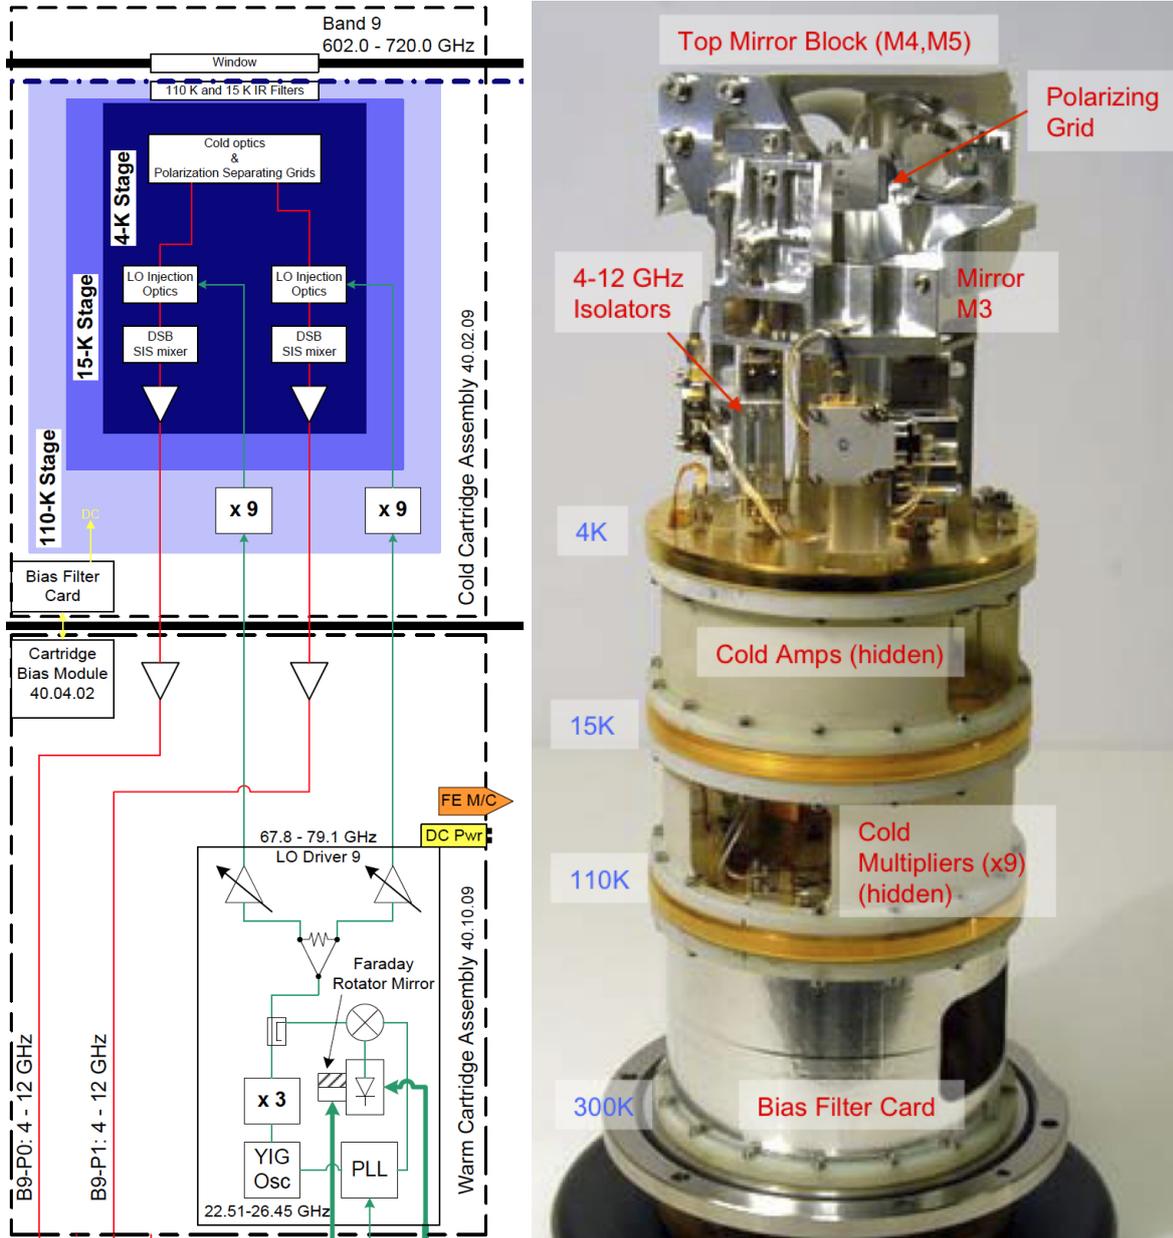


Figure 4.25: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 9 receiver was built at SRON in the Netherlands.

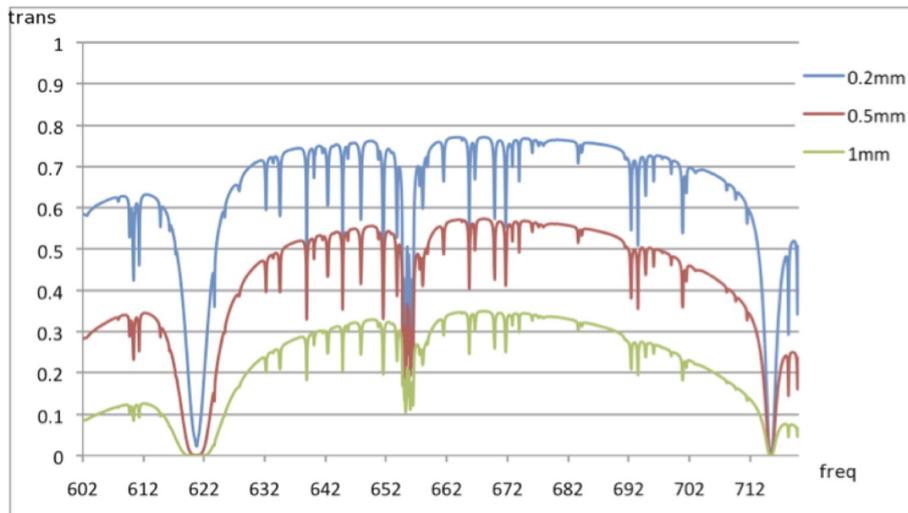


Figure 4.26: Band 9 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

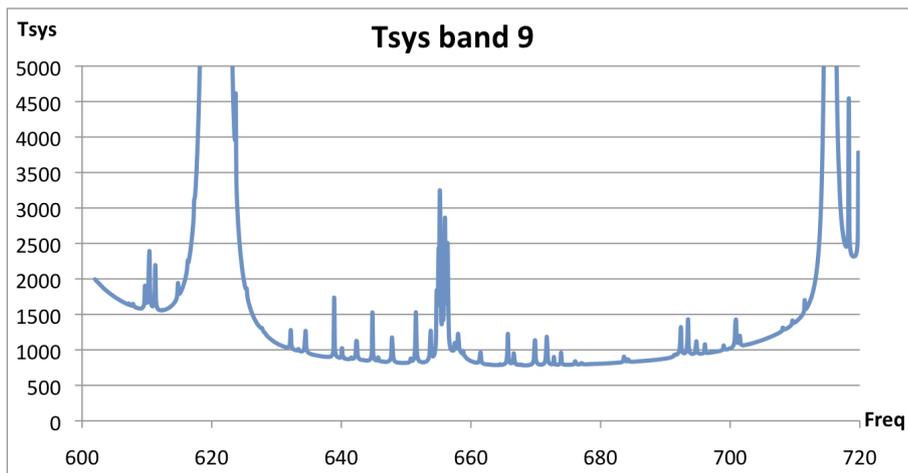


Figure 4.27: Typical T_{sys} at zenith for Band 9 with PWV = 0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 4.28). 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 6.4), covering 4-12 GHz. A block diagram of the Band 10 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.29.

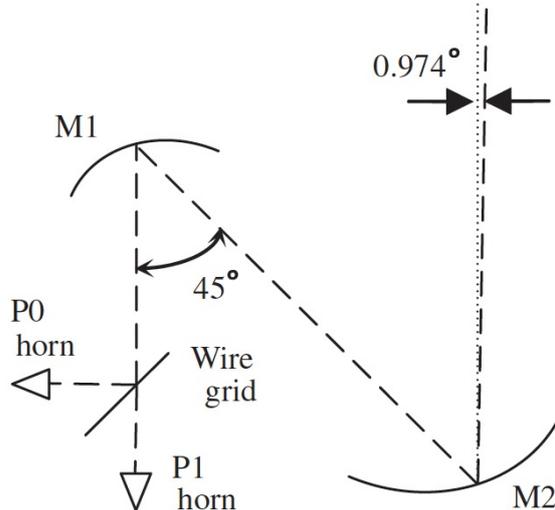


Figure 4.28: Schematic of ALMA Band 10 optics. ALMA Band 10 optics is composed of two elliptical mirrors, M1 and M2, a wire grid and two corrugated horns. The wire grid is used to separate the two linear polarizations, P0 and P1, and it is located after the two elliptical mirrors to minimize the number of optical components required.

The Band 10 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.30 for 3 low values of PWV. Figure 4.31 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.

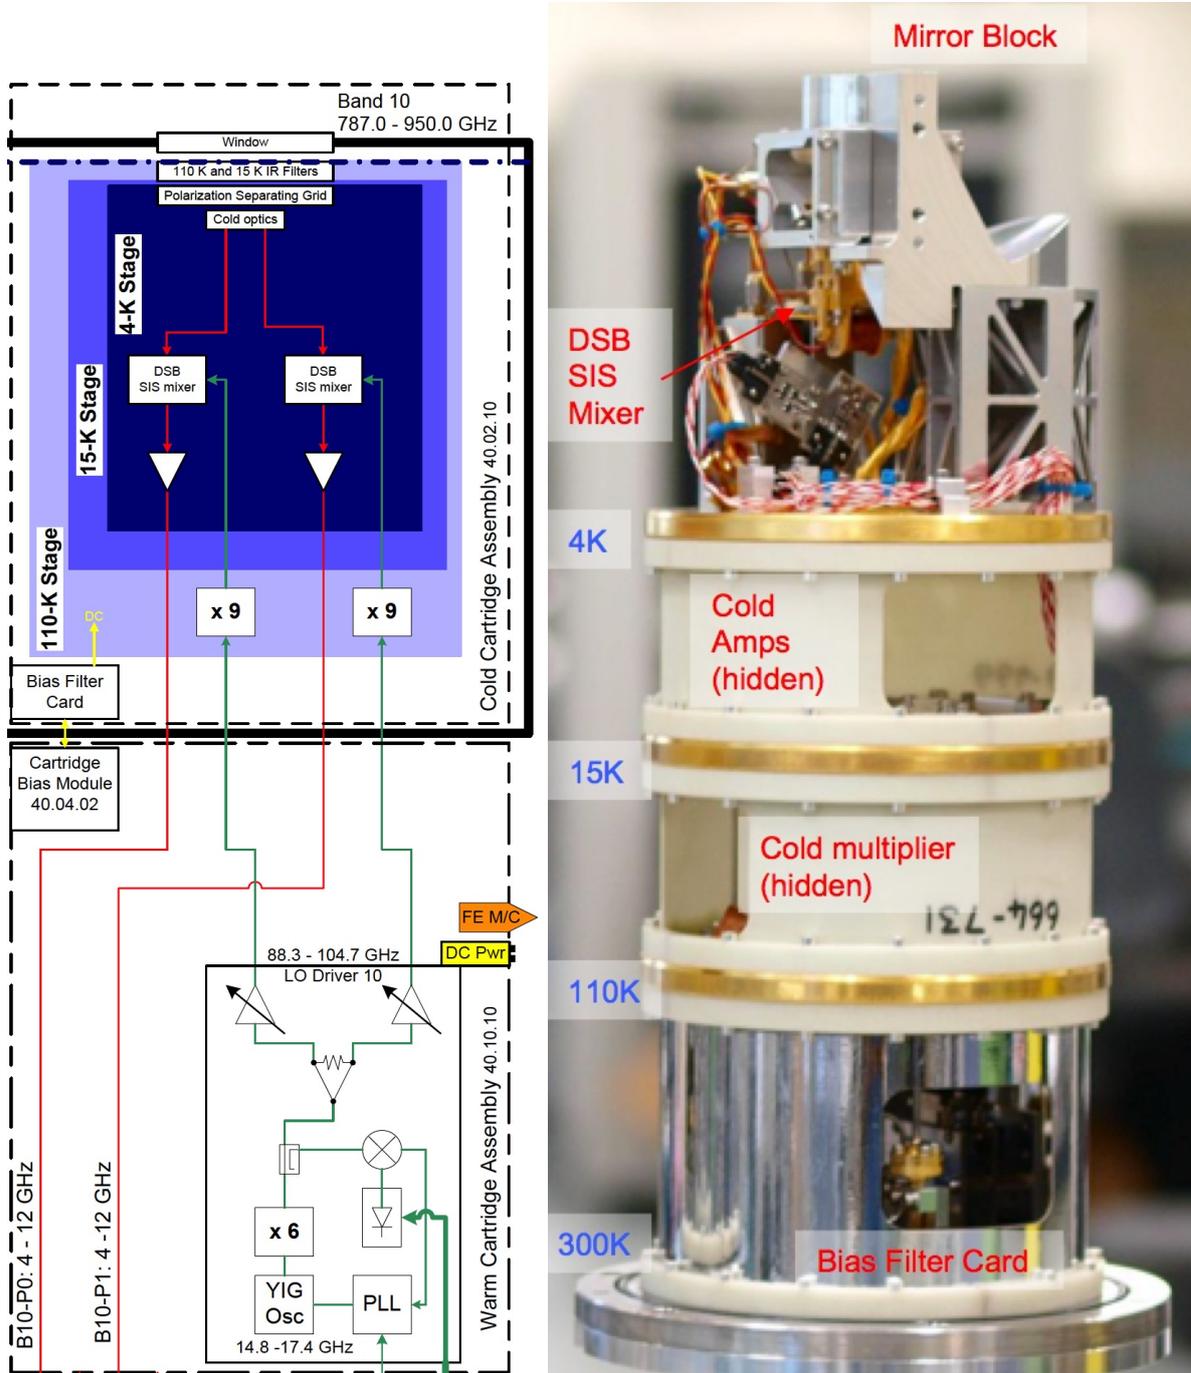


Figure 4.29: Block diagram of Band 10 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 10 cartridges are constructed at the Advanced Technology Center (ATC).

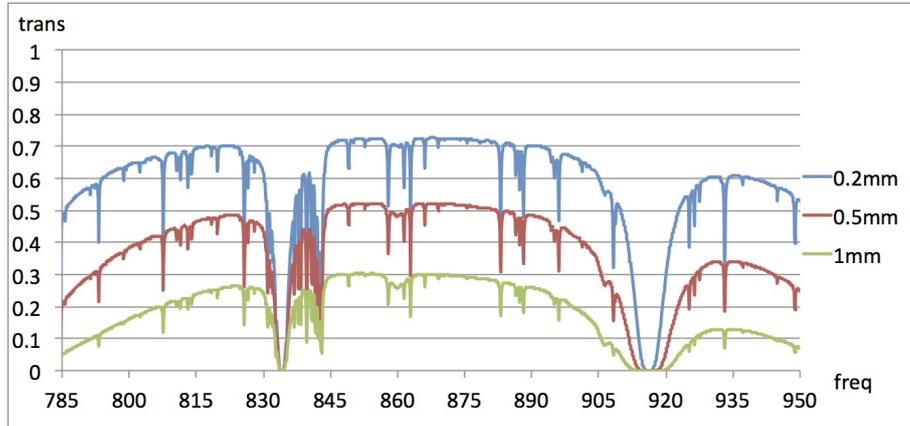


Figure 4.30: Band 10 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

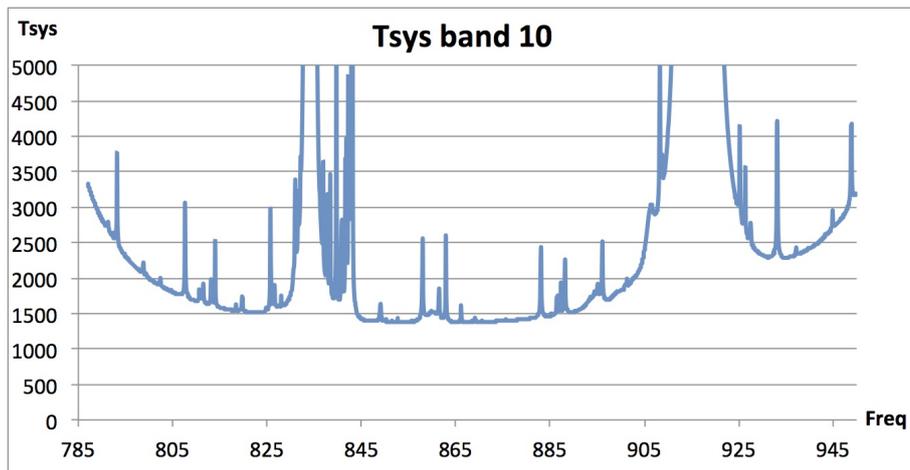


Figure 4.31: Typical T_{sys} at zenith for Band 10 with PWV = 0.472 mm. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

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 correlations and auto-correlations of them for each pair of antennas, and produces complex visibilities that users will reduce 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 antennas are processed in one of two correlators: the 64-input Correlator (also known as the Baseline Correlator) and the ACA Correlator. The 64-input Correlator is used primarily for the main 12-m Array, while the ACA Correlator is used for the Morita Array which forms the ACA 7-m Array and the Total Power Array (i.e. single-dish), 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 and/or the Total Power Array observing either the same or a different object.

Celestial signals received by the antennas are down converted to lower frequency bands using a set of Local Oscillators (LOs) and mixers as described in Appendix B. The outputs from the IF system form four 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 eight quantization levels (3 bits per sample) in digitizers, and then transferred via fiberoptic cable to one of the two correlators.

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

This chapter addresses capabilities of the correlators to be offered for Cycle 4 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 ³ system (Escoffier et al. 2007, A&A 462, 801), that brings a $32\times$ spectral resolution of the traditional lag (XF) correlators. 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 4. 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 with a dual polarization.

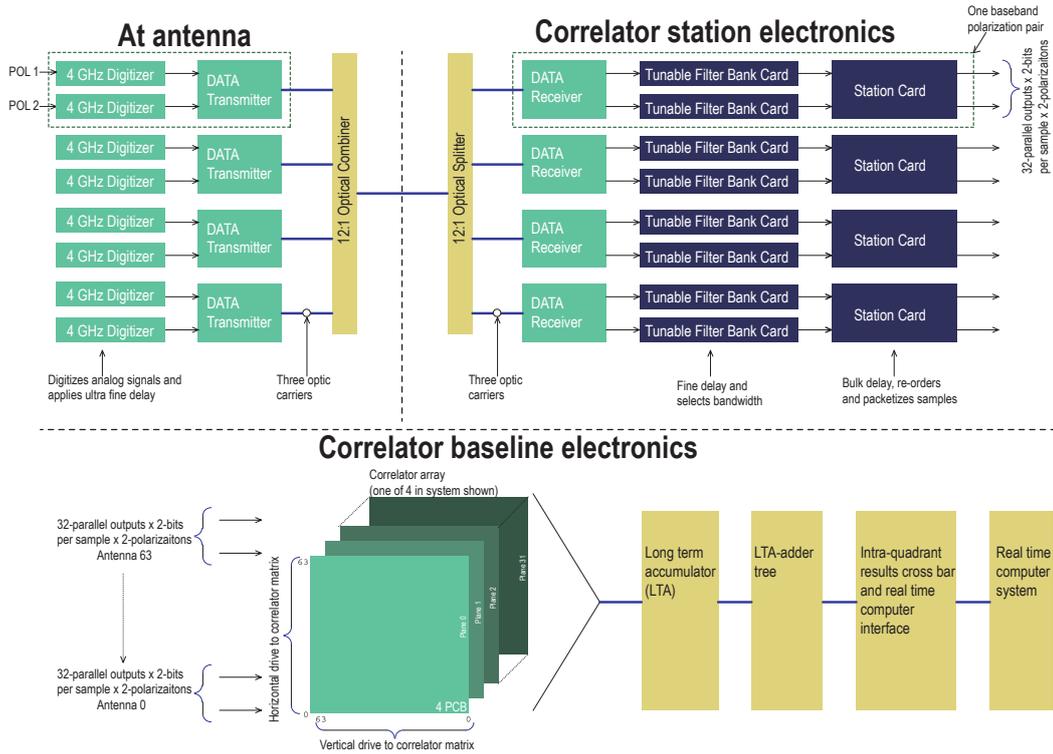


Figure 5.1: Overview diagram of the ALMA digitizers, data transmission system and 64-input Correlator (Escoffier et al. 2007, A&A 462, 801). The digitized data from the individual BBs are transferred to the Tunable Filter Banks (TFBs) and station cards (top right) for the parallelization and processing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

5.1.1 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

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 $N_{\text{ant}} = 64$.



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

Section 5.3).

The TDM mode provides a SPectral Window (SPW)⁵ that consists of up to $256/N_{\text{pol}}$ channels per BB, 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 6.4.

5.1.2 FDM mode

FDM is adequate for spectral-line observations that require a spectral resolution higher than that achieved with TDM. Each 2-GHz BB is split into 32 62.5-MHz sub-bands in a TFB card where digital filtering, digital mixer and LO (LO4) are implemented in Field Programmable Gate Arrays (FPGAs). The digital signal is re-quantized in the last stage of the TFB card which delivers 2-bit (4-level) samples⁷ at 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 together 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 valid sub-bands to form an

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

⁷Although the TFB card can deliver 2- or 4-bit samples after requantization and accepts 3-bit samples, the cross-correlation processor accepts only 2 bits in Cycle 4. See also Section 5.3.

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 SPWs in different BBs, all of SPWs in the same BB must have the same channel spacing.

The center frequency of each SPW can be tuned over the 2 GHz-wide BB using the digital LO (LO4) in the TFB card. However, the edges of the full bandwidth of the sub-bands cannot fall outside the 2 GHz BB range, and the frequency tuning is made in steps of 30.5-kHz set by the firmware in the TFB FPGAs. See also Section 5.5.2 for detail about spectral setting.

5.1.3 Correlation and realtime processing

The correlation cards perform the multiply-and-add operations to produce the correlation functions at a clock rate of 125 MHz (4 GHz samples demultiplexed by 32). Four quadrants handle four BBs. A quadrant of the correlator consists of 32 planes of 64×64 256-lag correlator circuits, and it yields auto-correlations and cross-correlations for 64 antennas⁸. There is a maximum of 8192 spectral points available with one BB being processed per quadrant thus providing high spectral resolution in FDM with narrower bandwidth than 2 GHz. 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 spectral setup details, see Chapter 6. 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 the correlator data processing.

5.2 The ACA Correlator

The ACA Correlator is dedicated to 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 antennas. The ACA Correlator is based on the FX⁹ design in which the incoming time-domain data stream is converted into frequency-domain spectrum via an FFT (Fast Fourier Transform) module before cross multiplication to form the power spectrum. The FFT part always accepts a full 2-GHz bandwidth and outputs 524288-ch spectra with a frequency resolution of 3.815 kHz. This design enables 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, a 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 the signal processing in the ACA Correlator itself is shown in Figure 5.3. Similar to the 64-input Correlator, each quadrant of the ACA Correlator processes a BB pair; so the complete system can process up to sixteen antennas independently from the 64-input Correlator.

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 spectrum¹⁰ (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 correspond to cross-correlations. Antenna-based power spectra, corresponding to auto-correlation, are also generated in the same way. Cross-polarization power spectra can be optionally produced. The (cross) power spectra are channel-averaged and time-integrated as designated before

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

⁹F and X stand for Fourier transform and correlation, respectively.

¹⁰ 2^{20} -point FFT produces a 2^{20} -point complex voltage spectrum of a double sideband including the image sideband that will be discarded.

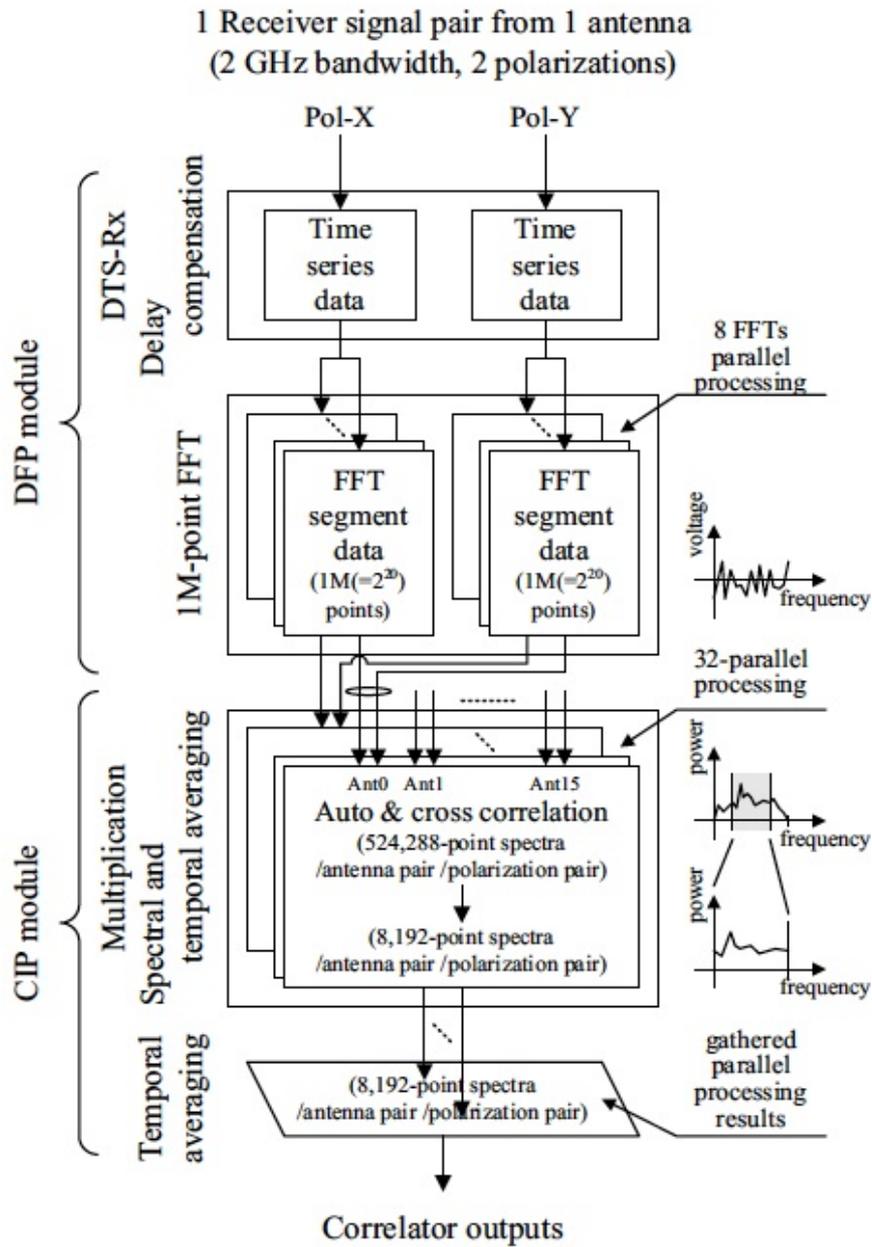


Figure 5.3: Block diagram of in the ACA Correlator. One of four quadrants is shown. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. FFT is performed in an FPGA delivering signals in the spectral domain which are then multiplied. The auto and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



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

they are sent to the ACA-CDP computers in the Computing subsystem through optical fibers.

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

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

5.2.1 The 7-m Array

The 7-m Array consists of twelve 7-m antennas whose correlation is performed in the ACA Correlator. The purpose of the 7-m Array is to cover low spatial frequencies in the (u, v) plane and to enhance the 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.

5.2.2 The 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 4. Single-dish dedicated observing modes such as position switching, raster scan, and OTF will be performed with the TP Array. Fast switching (~ 10 Hz) using the nutating subreflector will not be available during Cycle 4.

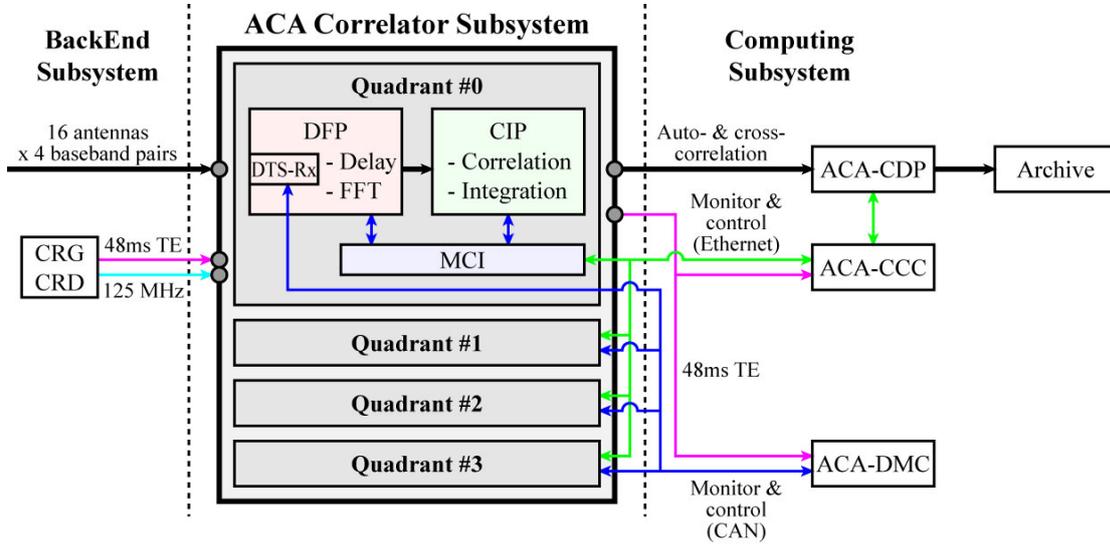


Figure 5.5: Overview diagram of the ACA Correlator. The signals from the BackEnd Subsystem are processed in the DFP and CIP modules and output to the Computing Subsystem (see text and Figure 5.3 for details of these). The modules are controlled by the Monitor and Control Interface (MCI), which communicates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

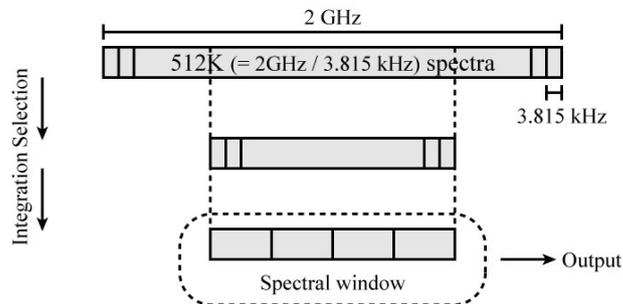


Figure 5.6: Data capture sequence of the ACA Correlator. The frequency range of an SPW is selectable from the 524288-point (512K) spectra across 2-GHz bandwidth. Additional channel-averaging within the selected frequency range before output.

5.3 Digitizers

A digitizer is a device converting at discrete times continuous voltage waveforms into quantized voltage levels encoded in a digital format. 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 $4 \text{ BBs} \times 2 \text{ polarizations}$. To sample and quantize the 2–4 GHz BB signals, the ALMA digitizer samples at the Nyquist frequency of 4 GHz and uses 8 quantization voltage levels (3 bit per sample). However, the 64-input Correlator reduces the quantization levels to 4 (2-bit). The ACA Correlator handles 3-bit digital form in FFT and the spectrum is requantized in 4-bit before cross multiplication. These quantization processes affect the sensitivity and linearity in the power spectra. The the 64-input correlator yield sensitivity loss of 12% (TDM) or 15.5% (FDM) due to the quantization, while the ACA Correlator does 4.8%, compared with an ideal analog system. These numbers suppose that the digitizer input power is in the expected range and that the digitizer adjustments are optimal. See also Section 5.6.1 about the quantization loss.

5.4 Online WVR Correction

The Water Vapor Radiometer (WVR; see Appendix A.6 for technical description) is a device mounted on each of the 12 m antennas to measure the amount of Precipitable Water Vapor (PWV) in order to correct for pathlength fluctuations in the troposphere. The WVR data are recorded together with the uncorrected visibilities, and in Cycle 4 (as in previous cycles), the WVR corrections are made offline in casa. However, the CDP of the 64-input Correlator can also apply a realtime correction of the pathlength fluctuation for each dump duration (online correction). This online corrected data is currently used for QA0 checks, but online WVR correction may eventually enable longer time-integration of the visibilities (thus slower output data rates) whilst maintaining coherence.

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

5.5 Capabilities of the Correlators

5.5.1 Polarization

Both correlators offer polarimetry capability by delivering all correlation products of $\langle XX^* \rangle$, $\langle XY^* \rangle$, $\langle YX^* \rangle$, and $\langle YY^* \rangle$ where X and Y stand for two linear orthogonal polarization components. These correlations relate to the Stokes visibilities of I , Q , U , and V as

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

where G_X and G_Y stand for complex antenna-based gains in X and Y polarizations and ψ is the parallactic angle. This formulation doesn't include cross-talk, also known as 'D-terms', whose calibration will be discussed in Chapter 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 4.

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)	Max. number of ch. for dual parallel pol. (XX and YY)	Ch. spacing (MHz)
TDM	2000*	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	2000*	4096	0.488
ACA	1000	4096	0.244
ACA	500	4096	0.122
ACA	250	4096	0.061
ACA	125	4096	0.0305
ACA	62.5	4096	0.0153

Table 5.1: Correlator modes and spectral performances per BB for dual parallel polarization (XX and YY). For full-polarization observations (XX , XY , YX , and YY), the maximum number of spectral channels per SPW becomes half and then the channel spacing is doubled. TDM and FDM stand for the Time Division Mode and Frequency Division Mode of the 64-input Correlator. The bandwidth and the number of channels in the ACA correlator became to remain unchanged through FPS (Frequency Profile Synthesis) since Cycle 4. NOTE: The channel spacing is *not* the same as the spectral resolution because of the applied weighting function. See table 5.2. * Usable bandwidth except band edges is ~ 1875 MHz.

The spectral profile which the correlators output is a convolution of the true spectrum with the spectral resolution function shown in figure 5.7. The spectral resolution function is given by 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. Note that the channel spacing table 5.2 is not the same as the spectral resolution, which depends on the weighting function in the lag domain.

Uniform weighting in the lag domain results in the sinc function with the FWHM of $1.21 \times$ the channel spacing and yields spectral sidelobes. This ‘ringing’ phenomenon affects the spectra when a narrow line, interference spike, or strong edge channels are present. Alternative weighting functions¹¹ 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 $2 \times$ 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 in the CDP. Channels can be averaged together; factors of $N = 2, 4, 8,$ or 16 are available. 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 who 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

¹¹For a full description, see <http://mathworld.wolfram.com/ApodizationFunction.html>.

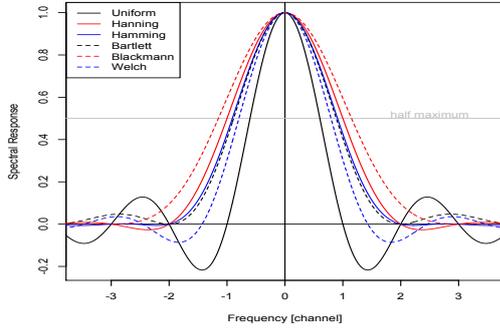


Figure 5.7: Spectral resolution functions corresponding to various weighting functions

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.

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
1875		3840	1920	960	480	240
		977 (488)	1129 (977)	1938 (1953)	3904 (3906)	7813 (7813)
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 parentheses) 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 $N = 16$). As N increases, the spectral resolution functions of adjoining channels are combined and then the spectral resolution approaches to the channel spacing. 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, to reduce the number of channels 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. FPS in the CDP can be applied to match the spectral resolution function to that of the 64-input Correlator. The bandwidth and the number of channels in the ACA correlator became to remain unchanged through FPS since Cycle 4.

5.5.3 Time Resolution

As shown in Figure 5.8 the correlated data are processed in the CDP and accumulators before being archived. The 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 the number of antennas, BBs, and spectral channels (sum of polarization products), respectively. N_{byte} (bytes/visibility) is the byte size of the real or imaginary part of a visibility. T_{int} is the integration time in the correlator. Typically, a configuration of $N_{\text{ant}} = 42$, $N_{\text{BB}} = 4$, $N_{\text{ch}} = 3840 \times 2$ polarization products (XX and YY), $N_{\text{byte}} = 2$ and $T_{\text{int}} = 6$ sec yields 18.1 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 optimal trade-off between time resolution, spectral resolution, polarization, number of BBs, and number of antennas.

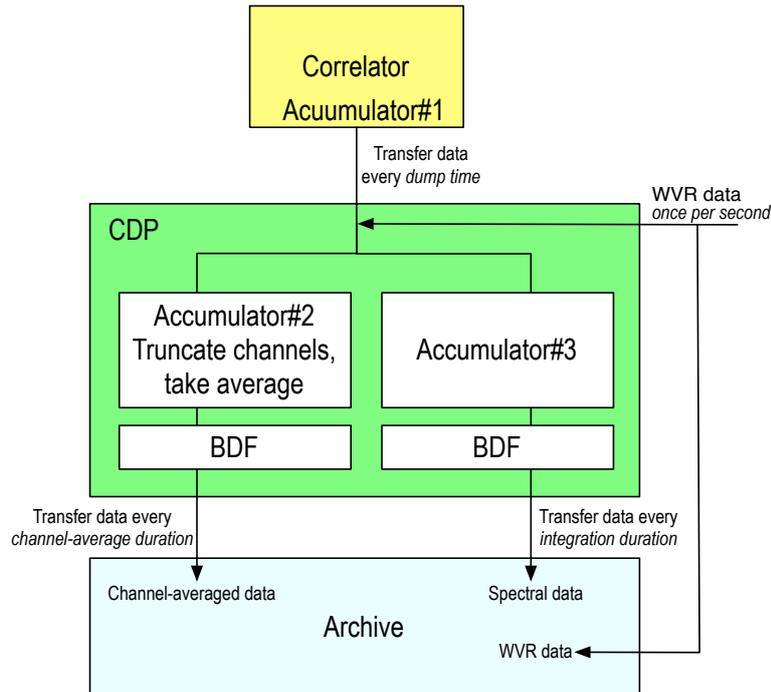


Figure 5.8: Basic data processing and accumulation steps between the correlator and archive. The CDP applies WVR correction before time integration of the spectral data and channel-averaged data. The set of WVR, channel averaged data, and spectral data are packed in Binary Data Format (BDF) and archived.

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, it is regulated at 40 Mbytes/s in Cycle 4. The 64-input Correlator employs $N_{\text{byte}} = 2$ where the real or imaginary part of the visibility is expressed with a 2-byte integer.

- Integration duration of spectral data with the ACA Correlator

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

- Integration duration of online WVR Correction

When on-line 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 4, 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 in a digitizer diminishes information contained in the input analog signals and causes a sensitivity loss. The ALMA digitizer employs 3-bit (8-level) quantization¹², and additional re-quantization processes are applied in the correlators.

In Cycle 4, the 64-input Correlator TDM mode makes use of 2-bit quantization that yields a quantization efficiency of $\eta_Q = 0.88$. In the FDM case, digital filtering is applied to 3-bit quantized signals followed by 2-bit re-quantization. The two stages of quantization would cause the net quantization efficiency of $\eta_Q = 0.96 \times 0.88 = 0.85$ if they were independent (see mode chart tables of Escoffier et al., ALMA Memo 556), nevertheless, the real efficiency can be slightly better than the multiplication (Iguchi et al. 2005, PASJ 57, 259)¹³.

The ACA Correlator feeds 3-bit quantized signals into the FFT processors where butterfly arithmetic is performed with 16-bit precision. The FFT output spectrum is re-quantized into 4-bit (16 levels) before cross multiplication. The combination of these quantizations results in 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 the signal level is adjusted for every scan, significant variation of power level (e.g. observation of strong sources, unstable weather, or frequency band near atmospheric absorption lines) can violate the conditions required to optimize the quantization efficiency.

5.6.2 Linearity

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

Figure 5.9 shows the Van Vleck relationships in 2-, 3-, and 4-bit quantizations. The correlation coefficients of quantized signals keep adequate linearity for small correlation coefficients¹⁴. 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 both correlators, the CDPs measure total power of digitized signals and determine the Van Vleck relationship to correct the (cross) power spectra.

¹²The 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).

¹³It is reported that the combination of 2-bit quantization, digital filtering, and 2-bit re-quantization yields $\eta_Q = 0.81$, better than $0.88 \times 0.88 = 0.77$.

¹⁴The departure from linear relation is $\delta\rho/\rho < 10^{-3}$ for $\rho < 0.2$

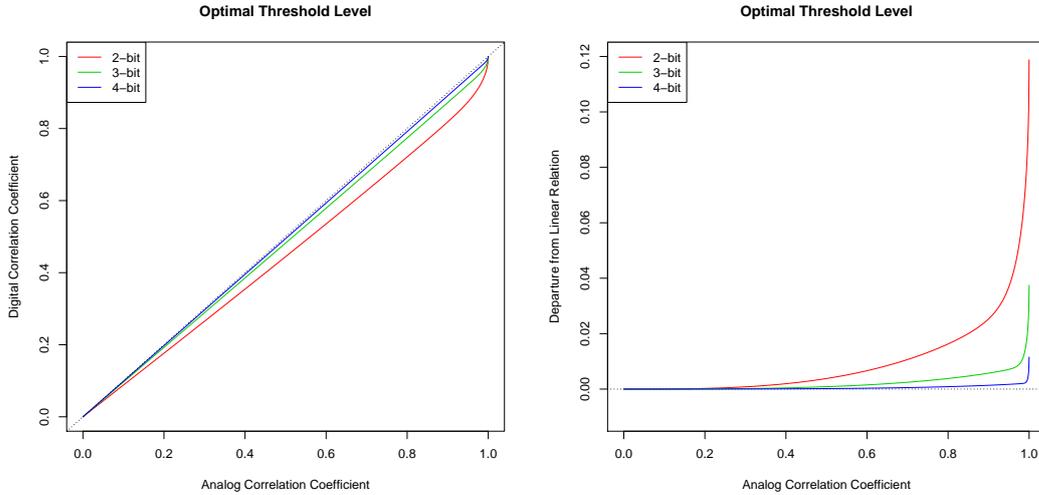


Figure 5.9: Van Vleck relationship of 2-, 3-, and 4-bit quantizations at the optimal signal power level and threshold voltages. (Left): Correlation coefficient of digitized signals as a function of analog correlation coefficients, ρ . (Right): Departures of the Van Vleck relationship from linear relation using the single factor of $\eta_4 = 0.881$, $\eta_8 = 0.963$, and $\eta_{16} = 0.988$ for 2-, 3-, and 4-bit quantizations, respectively.

5.7 Final data product - the ASDM

The final product from each observation in the archive is known as the ASDM (the ALMA Science Data Model), each of which has a unique hexadecimal name (e.g. `uid://A002/X2fed6/X3f`). The ASDM contains the 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 channel 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 more than 25 spectral windows. Every scan/SPW combination has an “intent” associated with it that indicates its purpose (pointing, system temperature, science, etc). The 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 describe the frequency setup of ALMA. In particular, we show how to define the spectral setups from the user viewpoint and how these set the ALMA hardware. We summarise the local oscillator (LO) and Intermediate Frequency (IF) systems and the various functions that they perform, as well as some of the other points pertaining to the frequency setup. For those interested in more details of the LO & IF components and how the hardware works, please jump to Appendix B.

To the ALMA system, a spectral setup includes the hardware LO, IF and correlator settings, such that each *Spectral Window* (SPW) covers the desired lines and/or continuum frequencies. To the end-user, however, the spectral setup is normally defined in the Observing Tool (OT) just in terms of the desired lines or observing frequencies, spectral window bandwidths and spectral resolutions: there is no need to worry about the details of each hardware setting. For full details of the OT and how to use it, see the user manual and reference manuals, available from the ALMA website¹ (and also in the OT itself).

6.1 Introduction

ALMA uses two stages of heterodyne conversion to shift the signals from the sky or observing frequency down to a range where electronics can be used to perform the digital sampling and cross-correlations. In general, the signals of frequency f_{sig} are mixed to 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 f_{sig} . The output signal may correspond to an input frequency above or below f_{LO} , known as the upper and lower sideband (USB or LSB).

Figure 6.1) illustrates the ALMA signal path and basic principle of operation. The first heterodyne stage uses the SIS mixer and LO1 to downconvert the astronomical signal to 4GHz- to 8GHz-wide IF bands (both USB and LSB). Up to four 2-GHz wide *Basebands* (or BBs) can be placed in the available IF range in either or both the USB or LSB IFs using the second stage of downconversion. Within each of these basebands it is possible to place, in principle, up to 32 spectral windows (although in practise it is unlikely that more than 4 will be used per baseband). Each Spectral Window forms a final contiguous spectrum, with bandwidths in FDM mode from 58.594 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 spectral lines) to be observed in each SPW and their required bandwidths or spectral resolutions. During the spectral tuning, there are effectively four different LOs used by the system (Figure 6.1):

- LO1 which sets the frontend tuning frequency
- LO2 which positions the basebands within the frontend receiver IF (each baseband uses a different LO2,

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

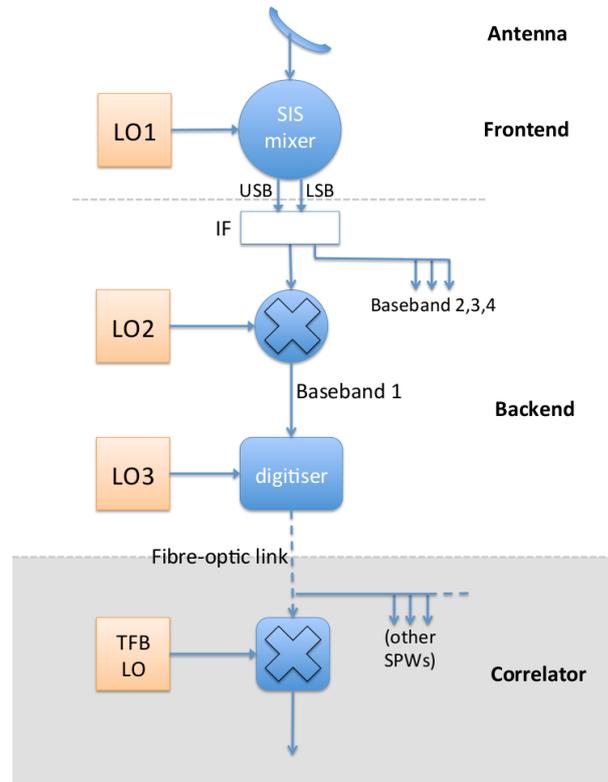


Figure 6.1: Summary block diagram of ALMA signal path and LO system for one spectral window, one baseband, one polarisation. The Backend IF system is located in each antenna, and the digitised signal for each baseband/polarisation is fed to the correlator in the central technical building through buried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

hence there are four LO2s)

- 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 LO1, 2 and 4 based on the requested observing frequencies. Because the LOs themselves are generated by combinations of frequency synthesizers, several different tuning solutions may be possible, and the algorithm picks the best one. If no solution is possible, the OT will notify the user. A more detailed description of the LO operation is given in Section B.2. Cycle 4 has similar restrictions on the spectral setups and tuning as Cycle 3 (see Table 6.2 in section 6.8). 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 each baseband in the correlator are independent, so the bandwidth, resolution and correlator mode as well as the value of LO4 can be different for each baseband. For example, it is possible to have a 58 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 4 also allows for multiple FDM spectral windows *within* each baseband, described in Section 6.3.2 below;

²Actually in Cycle 4 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.

however, currently the bandwidth of all the SPWs *within the same baseband* has to be equal. 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 the average motion of the Sun with respect to 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.

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 Setups

The wide IF bandwidth and tuning ability allows for routine 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 6.4). 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 are used in the OT Phase 1 when continuum is selected; these are noted in Table 6.1 and illustrated by the red in Figure 6.2.

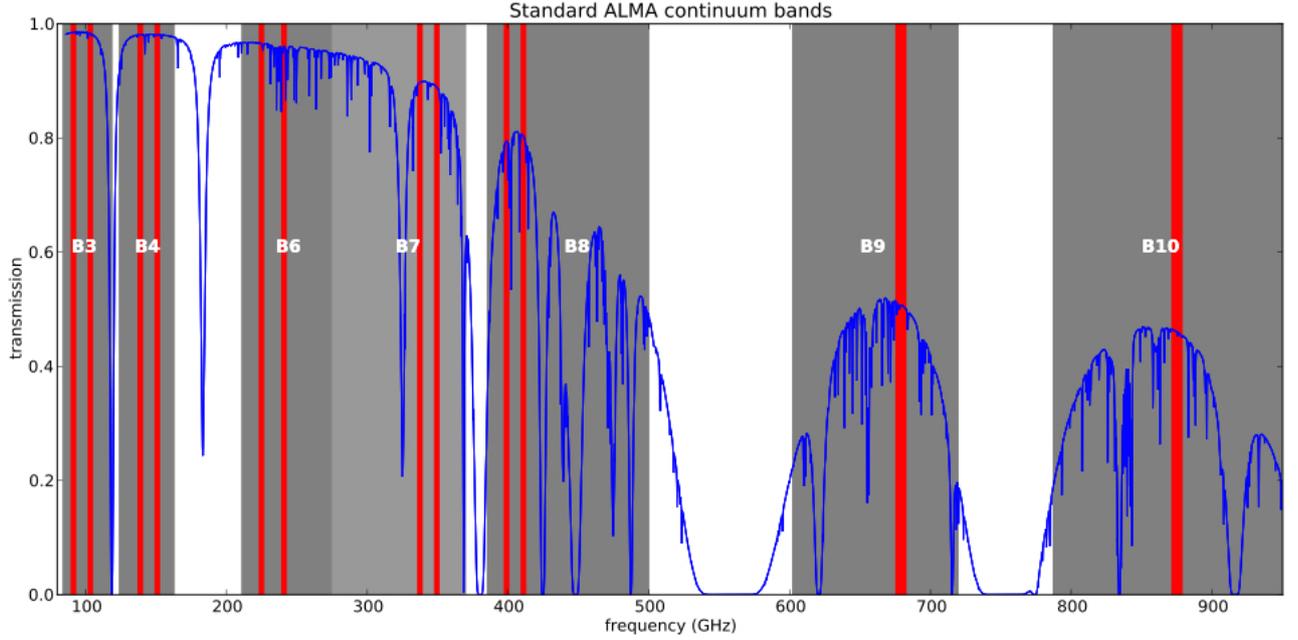


Figure 6.2: ALMA bands available in Cycle 4 (bands 3, 4, 6, 7, 8, 9, & 10), showing the frequency of the standard continuum settings as red shading. This gives the coverage of both USB and LSB, except for Band 9 & 10, which have 8 GHz of bandwidth using the USB only. The available frequency coverage of each band is shown by the grey shading, and atmospheric transmission for 0.6 mm of PWV is given by the blue line.

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

- BB1: f (BB covering the primary line)
- BB2: $f - 2.0$ (continuum BB in the same sideband, 2 GHz below the primary line)
- BB3: $f - (2 \cdot sb \cdot f_{IF})$ (continuum BB in the opposite sideband)
- BB4: $f - (2 \cdot sb \cdot f_{IF}) - 2.0$ (continuum BB in the opposite sideband)

where 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 BBs in the same sideband, so BB3 and BB4 would be at $f - 4.0$ GHz and $f - 6.0$ GHz. Note that this is an approximate rule - if possible, adjustment of the

³Note - this case assumes one SPW per BB

Band/setup	Species/transition	Freq. (GHz)	Sideb.	bandwidth	BB	SPW	Notes
3 a	Standard cont.	97.5 ¹	dual	TDM			LO1=97.5
3 b *	HCO ⁺ 1-0	89.188	LSB	58 MHz	1	1	HCO ⁺ /HCN/H ₂ CO
-	HCN 1-0	88.632	LSB	58 MHz	2	1	
-	CH ₃ OH/H ₂ CO	101.293/101.333	USB	125 MHz	3	1	
-	continuum SPW	101.3	USB	TDM	4	1	
3 c *	CO 1-0	115.271	USB	58 MHz	1	1	CO/CN/C ¹⁷ O
-	CN	113.499	USB	58 MHz	1	2	
-	C ¹⁷ O 1-0	112.359	USB	117 MHz	2	1	
-	continuum SPW	102.5	LSB	TDM	3	1	
-	continuum SPW	100.5	LSB	TDM	4	1	
4 a	Standard cont.	145.0 ¹	dual	TDM			LO1=145.0
4 b	CS 3-2	146.969	LSB	117 MHz	1	1	CS/DCO ⁺
-	DCO ⁺ 2-1	144.077	LSB	117 MHz	2	1	
-	SO ₂ 4(2,2)-4(1,3)	146.605	LSB	117 MHz	3	1	
-	H ₂ CO 2(0,2)-1(0,1)	145.603	LSB	117 MHz	4	1	
6 a	Standard cont.	233.0 ¹	dual	TDM			LO1=233.0
6 b *	¹² CO 2-1	230.538	USB	117 MHz	1	1	J=2-1 CO isotopes
-	continuum SPW	231.6	USB	TDM	2	1	
-	C ¹⁸ O 2-1	219.560	LSB	117 MHz	3	1	
-	¹³ CO 2-1	220.399	LSB	117 MHz	3	2	
-	continuum SPW	218.5	LSB	TDM	4	1	
7 a	Standard cont.	343.5 ¹	dual	TDM			LO1=343.5
7 b *	¹² CO 3-2	345.796	LSB	58 MHz	1	1	CO/HCO ⁺ /HCN
-	HC ¹⁵ N	344.200	LSB	58 MHz	1	2	
-	HC ¹⁵ N	345.340	LSB	58 MHz	1	3	
-	continuum SPW	343.200	LSB	TDM	2	1	
-	HCO ⁺ 4-3	356.734	USB	58 MHz	3	1	
-	H ¹⁵ CN 4-3	355.440	USB	58 MHz	3	2	
-	HCN 4-3	354.505	USB	117 MHz	4	1	
7 c	¹² CO 3-2	345.796	USB	58 MHz	1	1	J=3-2 CO/ ¹³ CO
-	continuum SPW	344.8	USB	TDM	2	1	
-	¹³ CO 3-2	330.588	LSB	58 MHz	3	1	
-	continuum SPW	331.6	LSB	TDM	4	1	
8 a	Standard cont.	405.0 ¹	dual	TDM			LO1=405.0
8 b	CI 3P1-3P0	492.160	USB	117 MHz	1	1	CI/ ¹³ CI
-	CS 10-9	489.751	USB	117 MHz	2	1	
-	continuum SPW	477.5	LSB	TDM	3	1	
-	continuum SPW	479.5	LSB	TDM	4	1	
9 a	Standard cont.	679.0	USB	TDM			LO1=671.0
9 b	¹² CO 6-5	691.472	USB	469 MHz	1	1	CO/CS
-	CS 14-13	685.436	USB	117 MHz	2	1	
-	H ₂ S	687.303	USB	234 MHz	3	1	
-	C ¹⁷ O 6-5	674.009	LSB	469 MHz	4	1	
10 a	Standard cont.	875.0	USB	TDM			LO1=867.0
10 b *	CO 7-6	806.652	LSB	469 MHz	1	1	CO/HCO ⁺
-	HCO ⁺ 9-8	802.458	LSB	117 MHz	2	1	
-	continuum SPW	805.5	USB	TDM	3	1	
-	continuum SPW	803.65	USB	TDM	4	1	

Table 6.1: Examples of spectral setups possible in Cycle 4. 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 6.2. * Template spectral setup released with Cycle 4 version of the OT.

continuum BBs (for example, choosing whether the opposite sideband is LSB or USB) should be done using the OT spectral display, to avoid deep atmospheric absorption features. This is particularly important at Band 8, 9 and 10, and near the water lines around 325 GHz in Band 7 (see Figure 6.2). Also, if contiguous spectral coverage of the continuum is desired, the BBs should be offset by 1.875 GHz rather than 2.0 GHz as suggested above.

6.3.2 Multiple Spectral Windows in the Same Baseband

Cycle 4 allows the capability to have up to four SPWs in each baseband, and each baseband can have different setups. This is useful for projects where several lines need to be observed simultaneously at high spectral resolution. However, in cases where four lines or fewer are observed in total, if the full bandwidth/resolution is needed for each line and they are well-separated, it may be preferable to place each line in a different baseband. Also, in cases where the lines are close to one another, or less spectral resolution/bandwidth is acceptable, it may be better to include multiple lines in the same SPW, and use the other SPWs in TDM mode to maximize the total (aggregate) bandwidth for phase calibration. Multiple SPWs per baseband does have some restrictions: each SPW within the same baseband must have the *same* bandwidth⁴, and the correlator resources used for each SPW in the same baseband (effectively the number of spectral channels, before channel-averaging) must be equal. However, it is possible to set up different channel-averaging for each SPW, so the final number of channels in each SPW may not be equal. An example of this is shown in Table 6.1 for spectral setup 6 b, where both $C^{18}O$ and ^{13}CO are observed in the same baseband (3.1 and 3.2). Multiple SPWs are set up in the Phase 1 OT simply by adding more spectral lines in the same baseband, and setting the SPWs to have a correlator fraction less than 1.0 (so for 4 SPWs, each has a correlator fraction of 1/4). There may be up to 4 SPWs per baseband in Cycle 4, however, the total number of spectral channels in each baseband is limited by the correlator resources. So doubling the number of SPWs in one baseband will result in half the number of channels per SPW, i.e. a lower resolution.

Figure 6.3, which is adapted from the Phase 2 Spectral Editor of the Observing Tool (OT), illustrates a more complex spectral setup with multiple SPWs and baseband which is possible in Cycle 4. The lines observed are given in Table 6.1 example 7 b. In this case, the frequency of LO1 is 350.0 GHz, and the upper and lower sidebands of the Band 7 receiver are shown as 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, labelled by the name of the primary spectral line or continuum. In this example, baseband 2 is set for a TDM SPW covering the whole 2 GHz baseband for a continuum measurement, whereas in some of the other basebands, multiple spectral windows are used. Baseband 1 has three 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.

6.3.3 Spectral Setups for Lines near the Edge of the Bands

The restriction that the baseband coverage cannot fall outside the maximum or minimum tuning range of the receiver is an issue for certain lines at the edge of the tuning range. One obvious example is the $^{12}CO(J=1-0)$ 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.

⁴although, as before, SPWs in *different* basebands can have different bandwidths

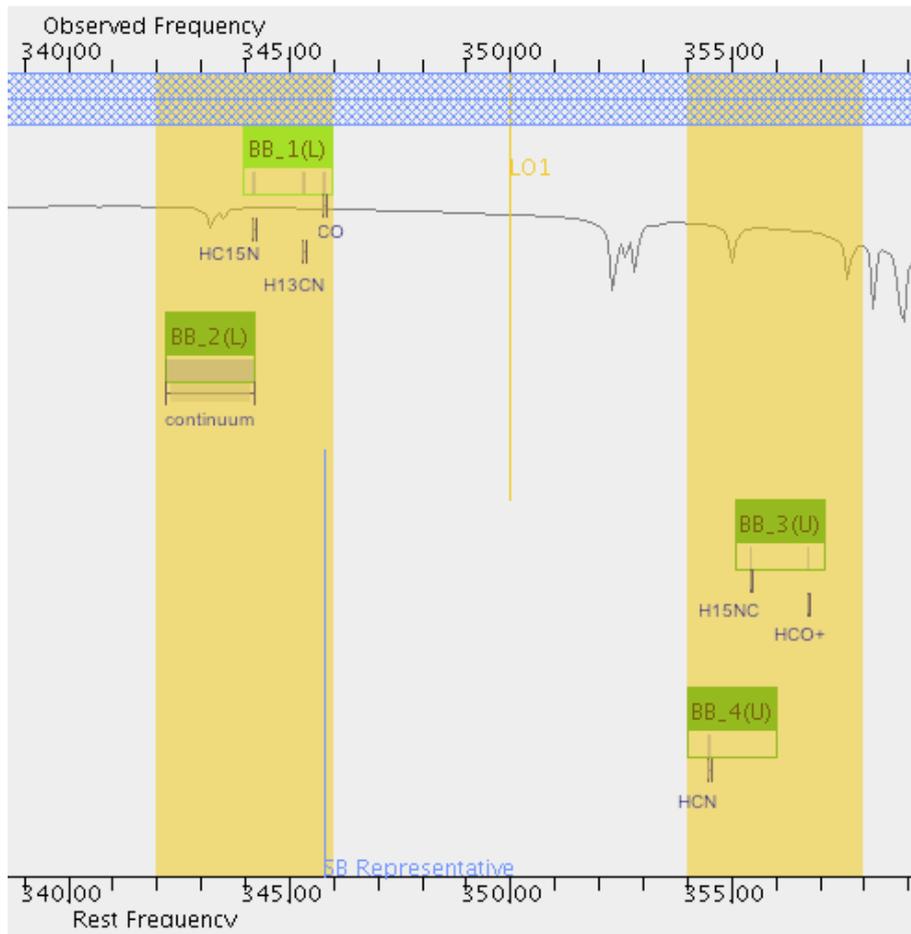


Figure 6.3: Illustration of a frequency setup, based on the OT spectral display. Yellow areas are the IF ranges, green bars are the 2 GHz-wide basebands (BB1-4), and smaller horizontal bars represent the spectral windows (1, 2 and 4 per baseband in this example). The frequency of LO1 is shown by the central vertical line. The blue hashed area represents the tuning range of the frontend (in this case a section of band 7), and the curved line represents the nominal atmospheric transmission for the chosen PWV. The bandwidths of the SPWs are illustrated by the widths of the horizontal bars. The lines observed are given in example 7b in Table 6.1.

6.3.4 Spectral Setup in Band 9 & 10 and DSB Considerations

The spectral setup in Band 9 and 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.

6.4 Usable Bandwidth

The IF system contains an anti-aliasing filter which limits the bandwidth of the basebands. Nominally this filter has -1 dB points at 2.10 and 3.90 GHz, giving a maximum bandwidth of 1.8 GHz. However, the IF response is such that the usable bandwidth is slightly wider – i.e. closer to 1.9 GHz. In FDM mode, the correlator outputs a maximum bandwidth of 15/16 of the nominal bandwidth (and reduces the number of channels by the same factors, from 4096 to 3840 - 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 6.4). It is recommended that 4 (in double-polarization) or 8 (single-polarization) channels are removed or flagged manually offline. This results in approximately the same usable bandwidth in both TDM and FDM modes and is illustrated in Figure 6.4.

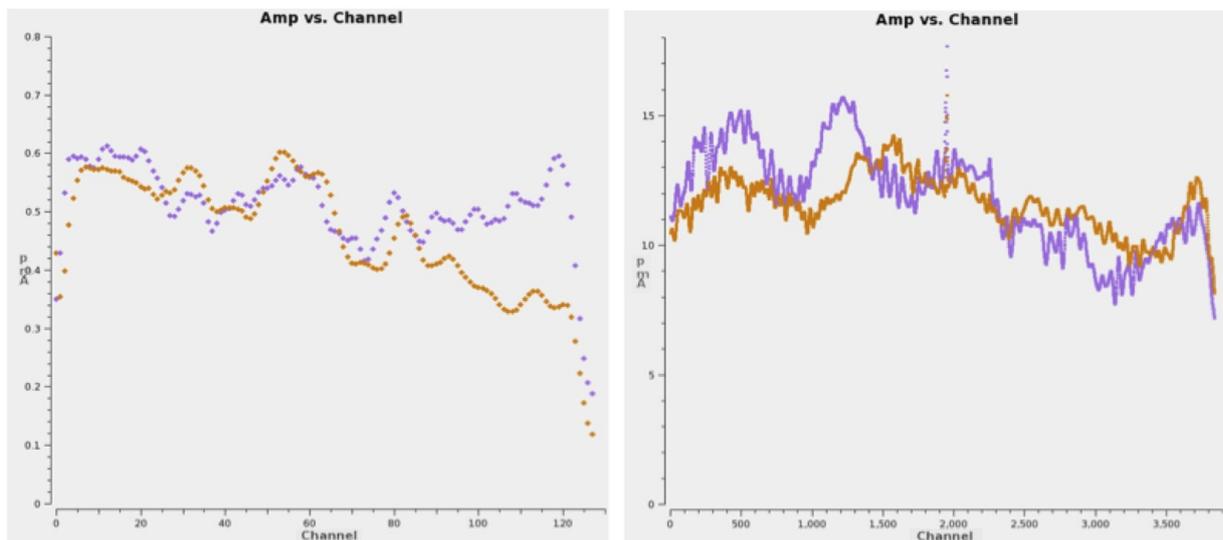


Figure 6.4: Comparison of TDM (left) and FDM (right) autocorrelation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in 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 Spectral Resolution

The spectral resolution of each SPW is set by a combination of the inherent resolution/bandwidth setup in the correlator, the correlator weighting function and the channel-averaging factor (which is carried out at a late stage of processing in the correlator). This is described in the correlator section 5.5.2. The default setup uses Hanning smoothing with no channel-averaging, which results in a spectral resolution of $2 \times$ the channel spacing. Selecting a channel-averaging factor of 2 therefore results in only a small loss in the final resolution (of $\sim 15\%$), but has the advantage of halving the data rate. Unless the highest spectral resolution is needed, it is recommended to apply some channel-averaging, to reduce the data rate and size of the data product; this is more critical where fast dump times are needed, e.g. for long baseline observing.

6.6 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.*

6.7 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 recommended for use in ALMA are:

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

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

LSRK Velocity with respect to the Kinematic local standard of rest, at 20.0 km s^{-1} in the direction $18^h, +30^\circ$ [B1900.0]. This frame is based on the mean velocity of the stars in the Solar neighbourhood.

Other velocity reference frames are also possible, including Heliocentric and LSRD (Dynamical LSR, based on the Solar velocity with respect to a rotating frame around the Galactic Center).

If a target velocity and reference frame other than topocentric is selected in the OT, at the start of the observation, the velocity of the science target in the chosen reference frame *and* the velocity of the observatory relative to the chosen reference frame are used to set the center frequencies of the science SPWs⁵. It is also possible to just make the correction for the motion of the observatory relative to the chosen reference frame, ignoring the source velocity in that frame. The frequency setting for the science target will normally also be used in the same execution for both the bandpass and amplitude calibration. For sources with an external ephemeris file, the rate of change of distance between the target and the observatory taken from the ephemeris is used to compute the source velocity at the start of the execution. This velocity is used throughout the observation (like non-ephemeris observations, the velocity from the ephemeris is not updated). Combining executions with different velocities must be done in CASA offline (normally in the Clean process).

Note that the velocity and redshift can be defined in three different ways, resulting in different conversions from velocity (or redshift) to frequency. These differences only become significant at high velocities. The velocity formulae are:

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

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

Relativistic $z=(f_o-f)/f$. The user specifies the redshift rather than velocity.

In these equations, v and z are the source velocity and redshift, and f and f_o are the observed and the rest frequencies of the line. For further details, see, for example, <http://www.gb.nrao.edu/~ghigo/gbtdoc/doppler.html>

⁵Note that ALMA does *not* do Doppler *tracking*, where the frequency would be continuously updated for Earth motion during the observation. Doppler corrections are only *set* once at the start of each execution.

6.8 Limitations and Rules for Spectral Setups in Cycle 4

The correlators already allow for a broad flexibility of SPWs in a single observation, however, the full capabilities are gradually being introduced and tested by the observatory before release. The current rules for spectral setups are given in Table 6.2.

Rule	Description
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 range 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 is possible to extend some part of the SPW to <50 MHz of the IF edges; in particular this is necessary for the Band 6 setup 6b in Table 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.	The number of SPWs per baseband can be 1, 2, 3 or 4. For 3 SPWs, the correlator resources per SPW should be set to 1/4, so only 3/4 of the available resources are used in this case.
12.	All the SPWs <i>within the same baseband</i> are required to have the same bandwidth and use the same correlator resources. An individual SPW within a baseband may occupy 1, 1/2 or 1/4 of the resources available in the baseband and the sum of all the fractional resources used within one baseband must be ≤ 1 . The correlator resources are essentially the inherent number of spectral channels. However, the channel-averaging can be different for different SPWs.

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} :

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

An ALMA observation can be represented by N discrete points in the uv -plane $\mathcal{V}_k(u_k, v_k)$ where $k = 1, \dots, N$ 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 is a noise-variance weight derived from the sensitivity of the k^{th} visibility measurement. It accounts for variables such as the integration time, the system temperature, and the bandwidth. These factors are determined by the instrument, technical set-up, and observing conditions, and are not controlled by the imaging process. T_k is a “taper” weight—by convention usually a multiplicative, Gaussian factor in the uv plane—that can be used to down-weight the longest baselines, suppressing small-scale sidelobes, and increasing the synthesized beam width and surface brightness sensitivity. Finally, D_k is a (uv) “density” weight that can be used to offset the high concentration of measurements near the center of the uv -plane, typically increasing resolution and reducing the sidelobes caused by gaps in the uv -coverage.

There are two commonly used limiting forms of density weighting (D_k):

- $D_k = 1$ 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.

As discussed in § 7.4, ALMA imaging typically uses a tunable, intermediate form of density weighting called Robust or Briggs weighting.

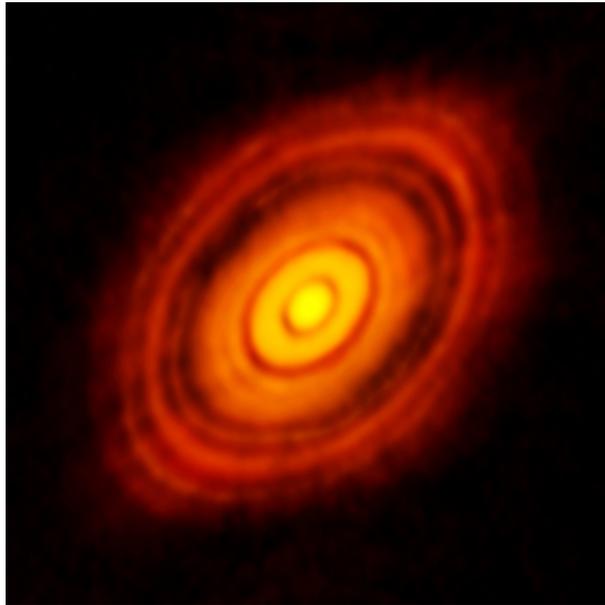


Figure 7.1: This is the sharpest image ever taken by ALMA. It shows the protoplanetary disc surrounding the young star HL Tauri. These new ALMA observations reveal substructures within the disc that have never been seen before and even show the possible positions of planets forming in the dark patches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

7.2 Cycle 4 Configurations

Cycle 4, 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 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 4, the 12-m Array, with (at least) 40 antennas, will be arranged in 9 different configurations. The maximum baselines achievable will be 12.6 km for Bands 3-6, 6.8 km for Band 7 and 3.7 km for Bands 8-10. The 7-m Array, with (at least) 10 antennas, will be available in only one configuration. Starting in Cycle 4, the three most extended configurations, C40-7, C40-8 and C40-9, will be offered in combination with other configurations. 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, 7.3 and 7.4 show antenna locations for the configurations of the 12- and 7-m Arrays. The three 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 user’s 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.

	Band	3	4	6	7	8	9	10
	Frequency (GHz)	100	150	230	345	460	650	870
Configuration								
7-m	θ_{res} (arcsec)	12.5	8.4	5.4	3.6	2.7	1.9	1.4
	θ_{MRS} (arcsec)	66.7	44.5	29.0	19.3	14.5	10.3	7.7
C40-1	θ_{res} (arcsec)	3.7	2.5	1.6	1.1	0.80	0.57	0.42
	θ_{MRS} (arcsec)	29.0	19.4	12.6	8.4	6.3	4.5	3.3
C40-2	θ_{res} (arcsec)	2.4	1.6	1.0	0.69	0.52	0.37	0.27
	θ_{MRS} (arcsec)	22.1	14.8	9.6	6.4	4.8	3.4	2.5
C40-3	θ_{res} (arcsec)	1.5	0.97	0.63	0.42	0.32	0.22	0.17
	θ_{MRS} (arcsec)	13.7	9.1	5.9	4.0	3.0	2.1	1.6
C40-4	θ_{res} (arcsec)	0.93	0.62	0.40	0.27	0.20	0.14	0.11
	θ_{MRS} (arcsec)	8.9	5.9	3.9	2.6	1.9	1.4	1.0
C40-5	θ_{res} (arcsec)	0.54	0.36	0.23	0.16	0.12	0.083	0.062
	θ_{MRS} (arcsec)	6.0	4.0	2.6	1.7	1.3	0.93	0.69
C40-6	θ_{res} (arcsec)	0.35	0.23	0.15	0.10	0.076	0.054	0.040
	θ_{MRS} (arcsec)	3.1	2.1	1.3	0.90	0.67	0.48	0.36
C40-7	θ_{res} (arcsec)	0.21	0.14	0.090	0.060	0.045	0.032	0.024
	θ_{MRS} (arcsec)	1.8	1.2	0.77	0.52	0.39	0.27	0.20
C40-8	θ_{res} (arcsec)	0.12	0.079	0.052	0.034	-	-	-
	θ_{MRS} (arcsec)	1.3	0.87	0.57	0.38	-	-	-
C40-9	θ_{res} (arcsec)	0.066	0.044	0.029	-	-	-	-
	θ_{MRS} (arcsec)	0.78	0.52	0.34	-	-	-	-

Table 7.1: Resolution (θ_{res}) and maximum recoverable scale (θ_{MRS}) for the 7-m Array and 12-m Array configurations available during Cycle 4 as a function of a representative frequency in a band. The value of θ_{MRS} is computed using L05 from Table 7.2 and equation 7.7; 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.) The computations were done for a source at zenith; for sources transiting at lower elevations, the North-South angular measures will increase proportional to $1/\sin(\text{ELEVATION})$.

Configuration	7-m	C40-1	C40-2	C40-3	C40-4	C40-5	C40-6	C40-7	C40-8	C40-9
Minimum baseline (m)	8.9	15.1	15.1	15.1	15.1	16.7	15.3	81.3	168.4	270.7
5th percentile or L05 (m)	9.2	21.2	27.7	44.8	69.0	101.2	198.2	379.0	501.8	707.0
80th percentile or L80 (m)	30.8	97.3	141.9	233.5	369.7	632.4	999.9	1748.2	2989.6	5446.8
Maximum baseline (m)	45.0	155.6	272.6	460.0	704.1	1124.3	1813.1	3696.9	6855.1	12644.8

Table 7.2: Basic parameters of the 7-m Array configuration and the nine 12-m Array configurations offered during Cycle 4. The baselines are projected for a transiting source ($HA = \pm 1h$) at a declination of -23° . Note that C40-8 will not be available for Bands 8-10, and C40-9 will not be available for Bands 7-10.

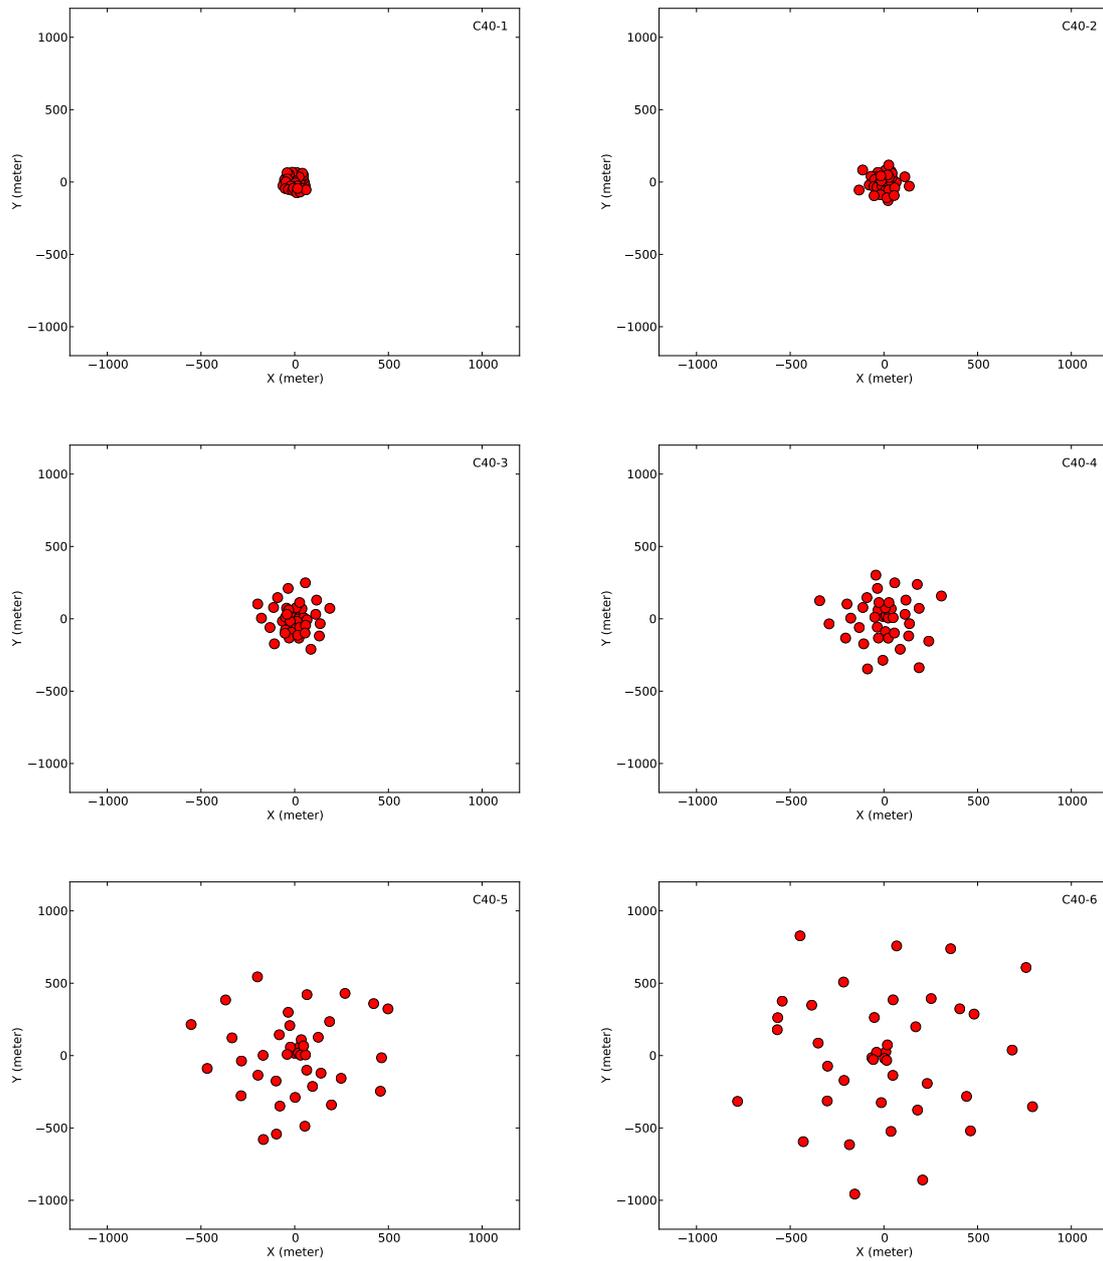


Figure 7.2: Representative 12-m Array compact configurations for Cycle 4.

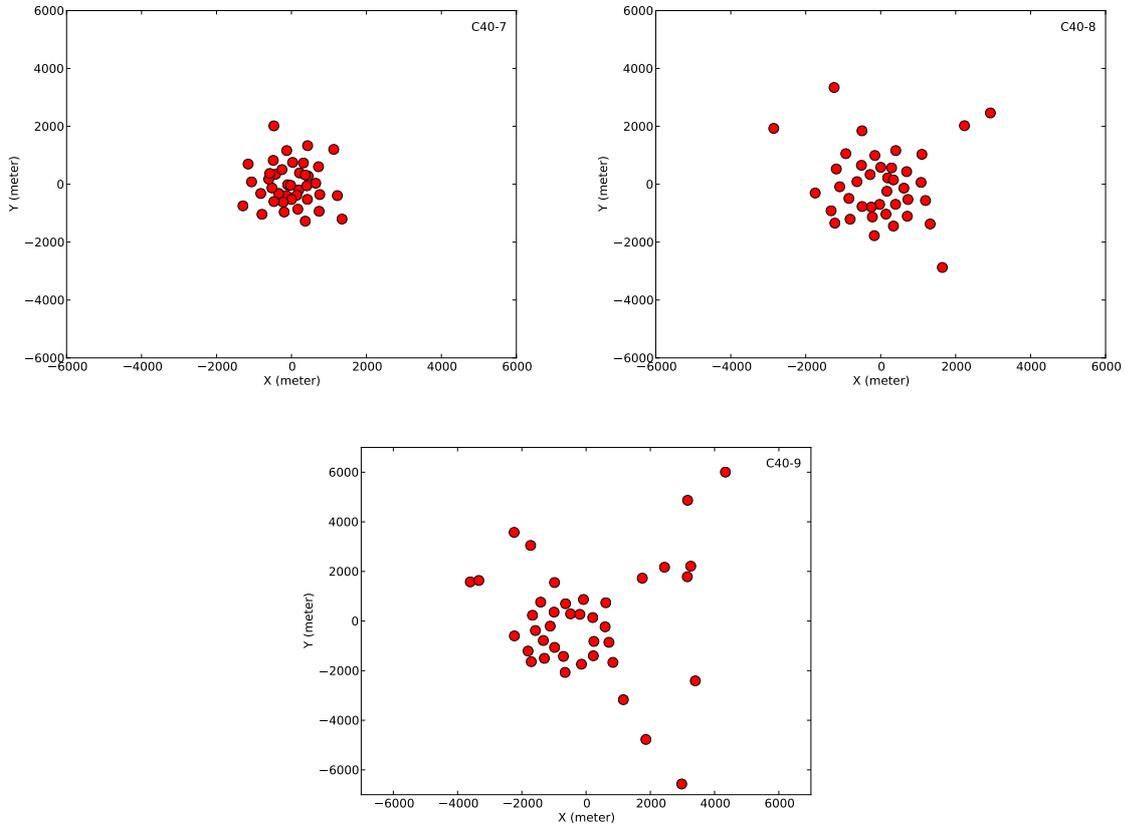


Figure 7.3: Representative 12-m Array extended configurations for Cycle 4.

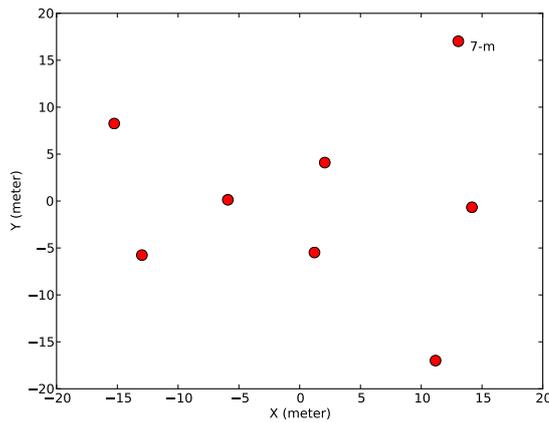


Figure 7.4: Representative 7-m Array configuration for Cycle 4.

7.3 Shadowing

ALMA is located at latitude= -23.02917° , longitude= -67.754649° . During observations, particularly in very compact configurations like those of the 7-m Array or C40-1, 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

and can be quite severe for sources that are observed at very low elevation. Targets as far North as declination $+47^\circ$, corresponding to a maximum source elevation at Chajnantor of $\sim 20^\circ$, can in principle be observed from the ALMA site, but shadowing by adjacent antennas becomes an increasing problem at low elevations. For example, Figure 7.5 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 (C40-1). 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.

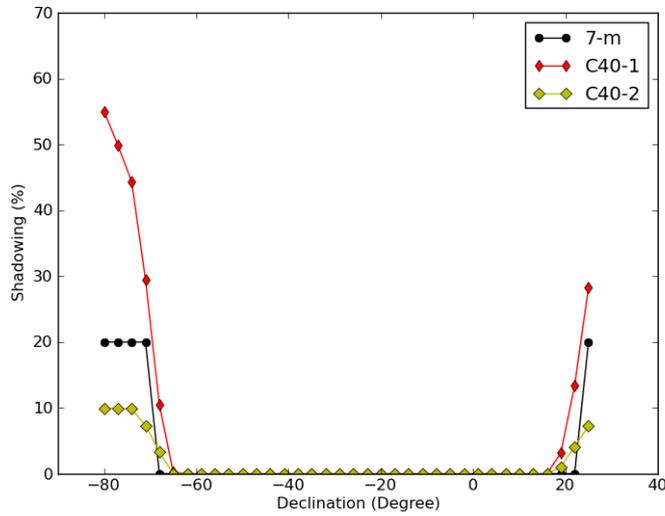


Figure 7.5: Shadowing fraction (%) for the most compact configurations as a function of declination.

7.4 Resolution and Beam Shape

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

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

where k is a factor that depends on the uv -plane weighting function, λ is the observing wavelength in meters, and L_{max} is the longest baseline in meters. For the ALMA configurations, the following equation was determined, which uses the 80th percentile of the uv -distance as a more robust proxy to the longest baselines.

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

where λ is the observing wavelength in meters, and L_{80} is the 80th percentile of the uv -distance in meters.

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

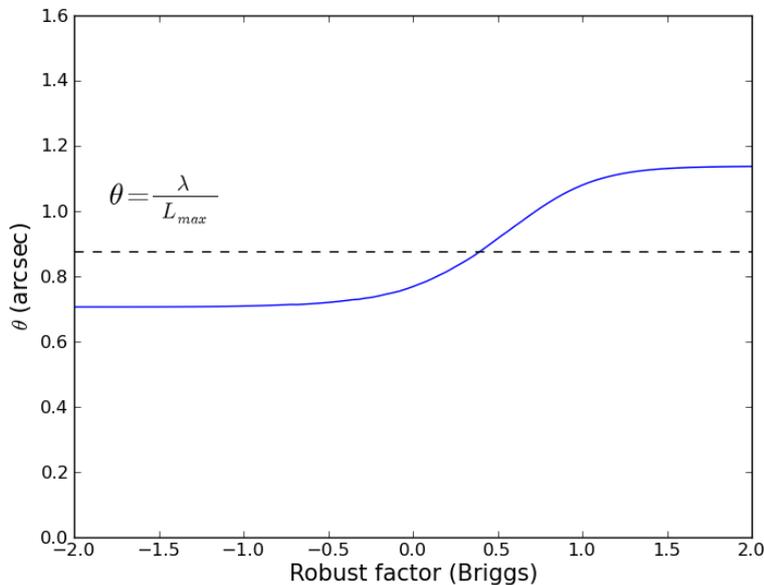


Figure 7.6: Angular resolution achieved using different values of the CASA *robust* parameter for a 1-hour observation at 100 GHz and a declination of -23 deg in the C40-4 configuration. Note that *robust* = -2 is close to uniform weighting and *robust* = 2 is close to natural weighting. The dotted line corresponds to $\frac{\lambda}{L_{max}}$.

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 7.6 shows the angular resolution achieved for an observation with the C40-3 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 7.7 shows the different beam shapes for sources observed at declinations of -70° (similar to the SMC and LMC) and -30° . Also, Figure 7.8 shows the uv -coverage for these sources, revealing the elongation of the uv -plane sampling distribution, together 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 7.9 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 7.10 shows the geometrical mean of the major and minor axes of θ_{res} at the same frequency. The beam width scales with λ , but bear in mind that not all configurations can be used with the higher-frequency bands e.g., C40-8 is not available at bands 7-10.

Note that in the OT, the angular resolution and sensitivity are now computed consistently, assuming Briggs weighting with *robust* = 0.5. The most appropriate value of *robust* to use when imaging a dataset must be

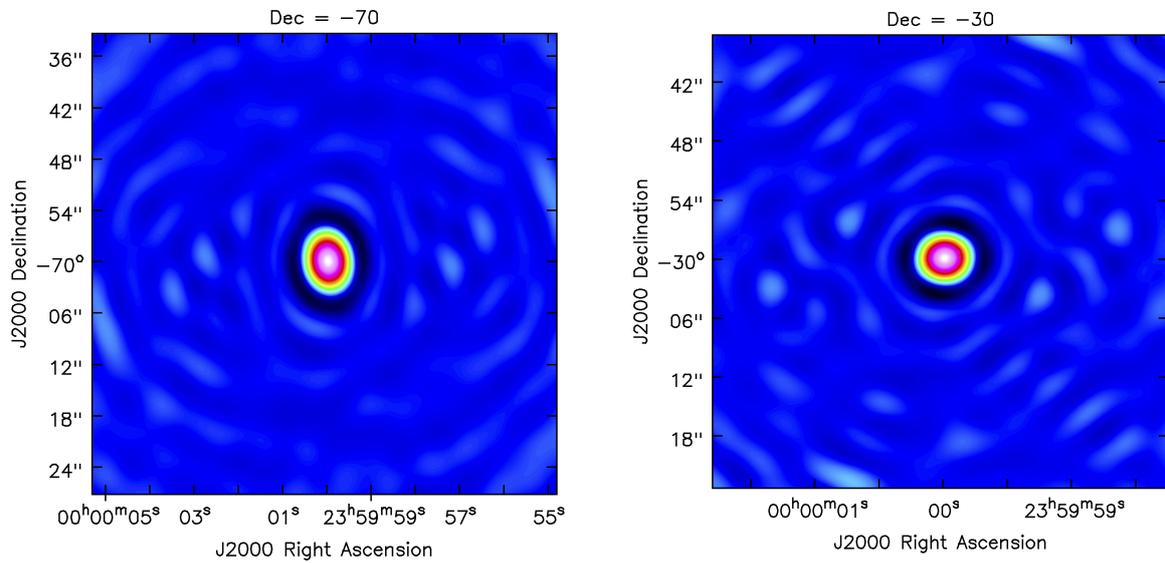


Figure 7.7: Beam shape for configuration C40-1 with a 2-hour observation of a transiting source at a declination of either -70° (left) or -30° (right).

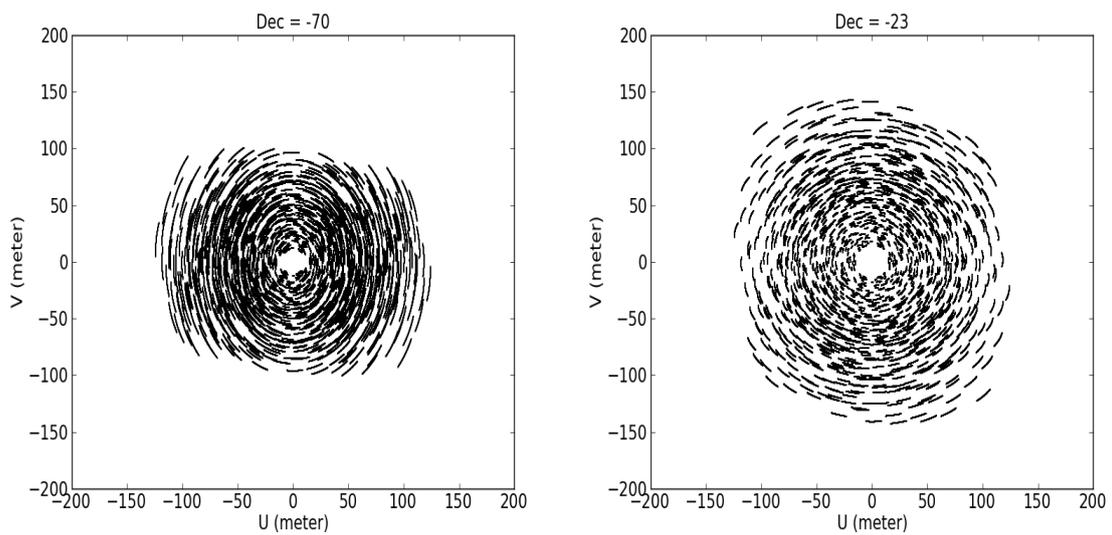


Figure 7.8: uv -plane coverage for configuration C40-1 with a 1-hour observation of a transiting source at a declination of either -70° (left) or -23° (right). For the source with a declination of -70° , the shadowing fraction represents 24.9% of the total.

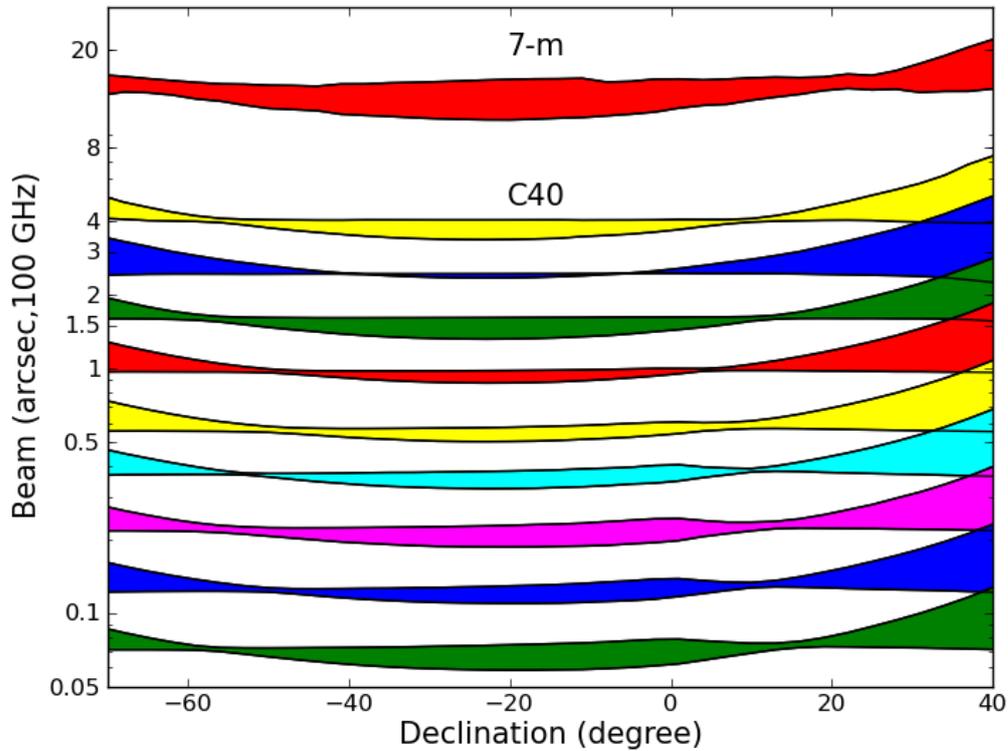


Figure 7.9: Cycle 4 angular resolutions as a function of source declination. Each color shows, for a particular configuration, the range of the major and minor axes of the synthesized beams expected from a 2-hour observation at 100 GHz. 12-m configurations are arranged with C40-1 at the top and C40-9 at the bottom.

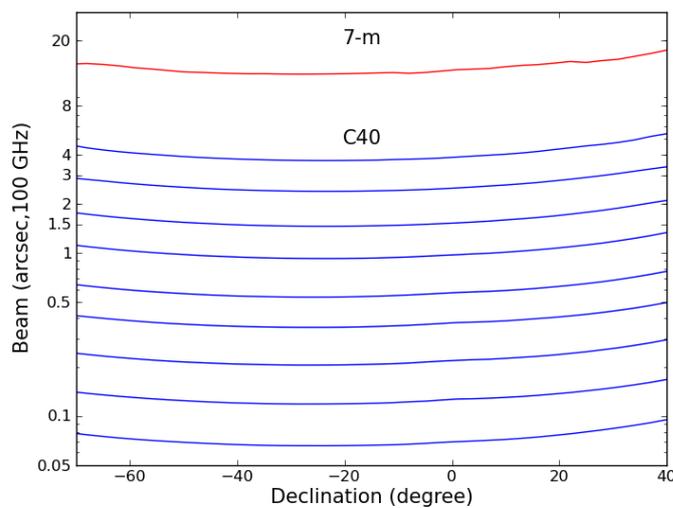


Figure 7.10: Geometrical mean of the major and minor axes of the synthesized beams as a function of source declination. These correspond to a 2-hour observation at 100 GHz. 12-m configurations are arranged with C40-1 at the top and C40-9 at the bottom. The 7-m Array is also shown.

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.)

As discussed in § 7.9, combining data from compact configurations with those from more extended configurations primarily has the effect of filling in negative sidelobes in the synthesized beam and thereby improving the interferometer’s response to larger angular structures on the sky. However it does also slightly degrade the resolution. This degradation can be mitigated by moving towards uniform density weighting.

7.5 Response and Snapshot

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

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

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 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 deg.

Configuration	7-m	C40-1	C40-2	C40-3	C40-4	C40-5	C40-6	C40-7	C40-8	C40-9
Natural	43.6%	11.8%	7.5%	10.6%	7.8%	11.9%	11.6%	13.5%	12.0%	11.7%
Briggs ($R = 0.5$)	42.6%	8.8%	6.8%	7.0%	8.1%	8.8%	9.3%	10.5%	7.9%	10.6%
Uniform	43.8%	9.1%	10.9%	11.3%	11.0%	9.5%	10.6%	11.4%	8.2%	11.3%

Table 7.3: Sidelobe levels for a 1-hour observation of an unresolved 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 7.11 and 7.12 show examples of a uv -plane sampling distribution and cleaned image for a snapshot of 1-minute duration and a longer integration of 1-hour duration expected with the configuration C40-1. As shown in Figure 7.11, the uv -coverage in 1 minute is quite uniformly sampled, but much less dense than that of the 1-hour integration. In addition, the cleaned images of the snapshot and the longer integration (see Figure 7.12) are similar in terms of angular resolution and the apparent differences are quite small. The main difference between the two (besides sensitivity) is in the dirty beam: the sidelobe level for the snapshot will be much higher than for the 1-hour integration. The snapshot image will then require more careful cleaning in order to avoid introducing spurious sources from the strong sidelobes. Condon et al. (1998) gave a very comprehensive description of how this issue impacted the 20-cm NRAO VLA Sky Survey (NVSS).

A useful practice to disentangle sidelobe effects from real point sources, especially with relatively strong point sources, is to perform CASA simulations using a component list of the strongest sources together with

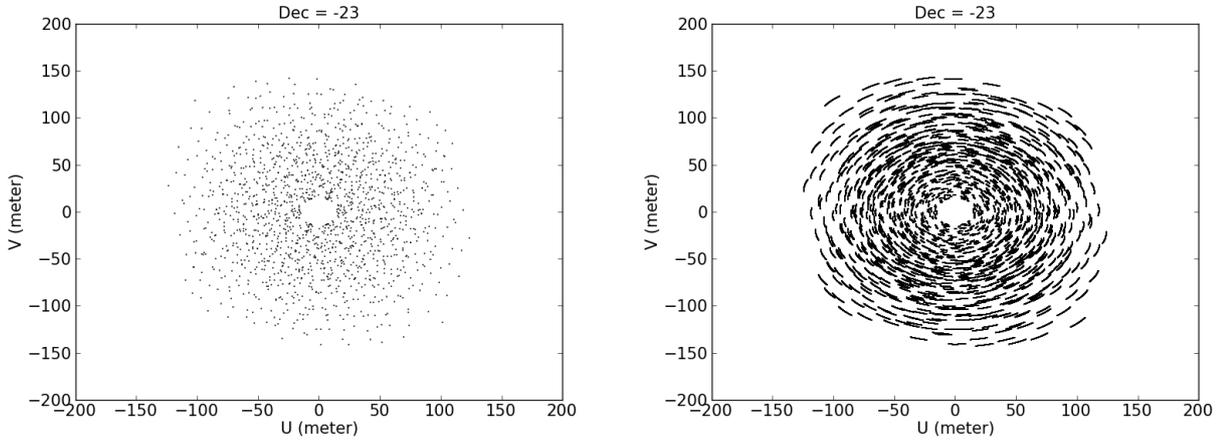


Figure 7.11: uv -plane sampling distributions of a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the C40-1 configuration to observe a source at a declination of -23° .

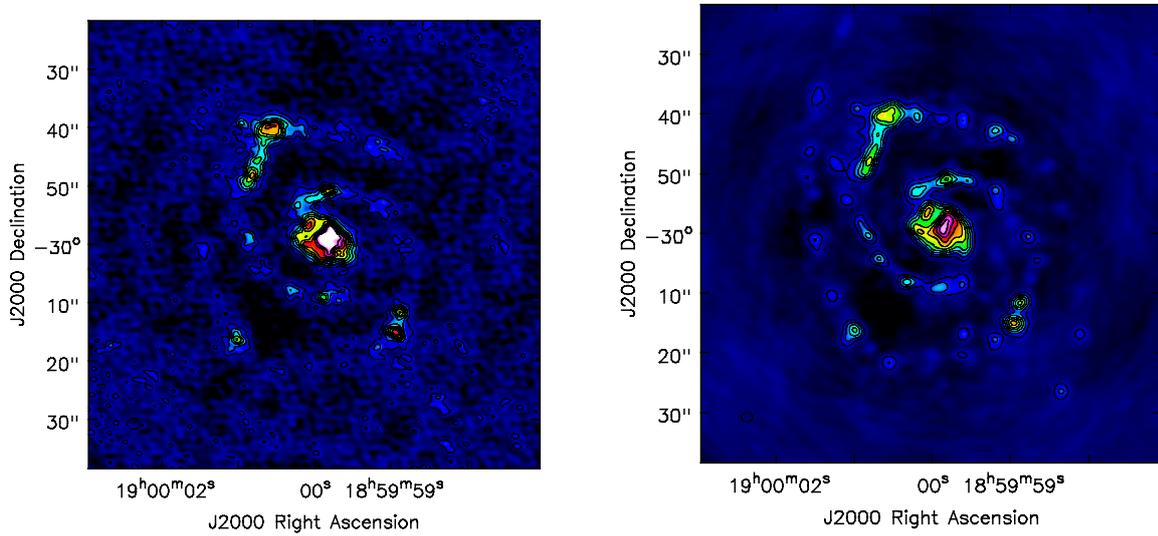


Figure 7.12: Images obtained from a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the C40-1 configuration to observe a source at declination of -23° . Black contours at 100, 200, 300, 500, 700 and 900 mJy beam^{-1} are overlaid.

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 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 7.13 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:

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

where λ is the observing wavelength in meters, and L_{min} is the shortest baseline in meters. For the ALMA configurations, the following equation was determined, which uses the 5th percentile of the uv -distance as a more robust proxy to the shortest baselines.

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

where λ is the observing wavelength in meters, and L_5 is the 5th percentile of the uv -distance in meters.

Table 7.1 lists the θ_{MRS} for the various Cycle 4 array configurations. 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

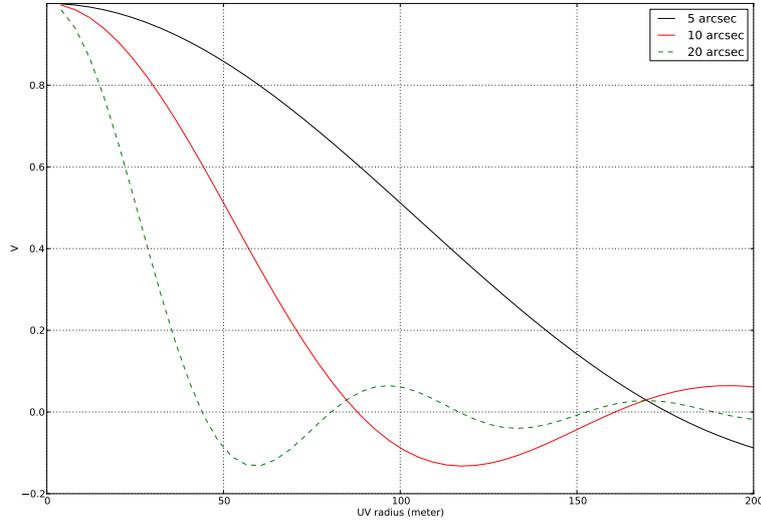


Figure 7.13: Expected visibilities of three model uniform disks (annular averages) at 100 GHz as a function of uv -distance.

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 ~ 23 Jy, whereas that detected by the configuration C40-6 ($k_{uv} \approx 200k\lambda$) 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.

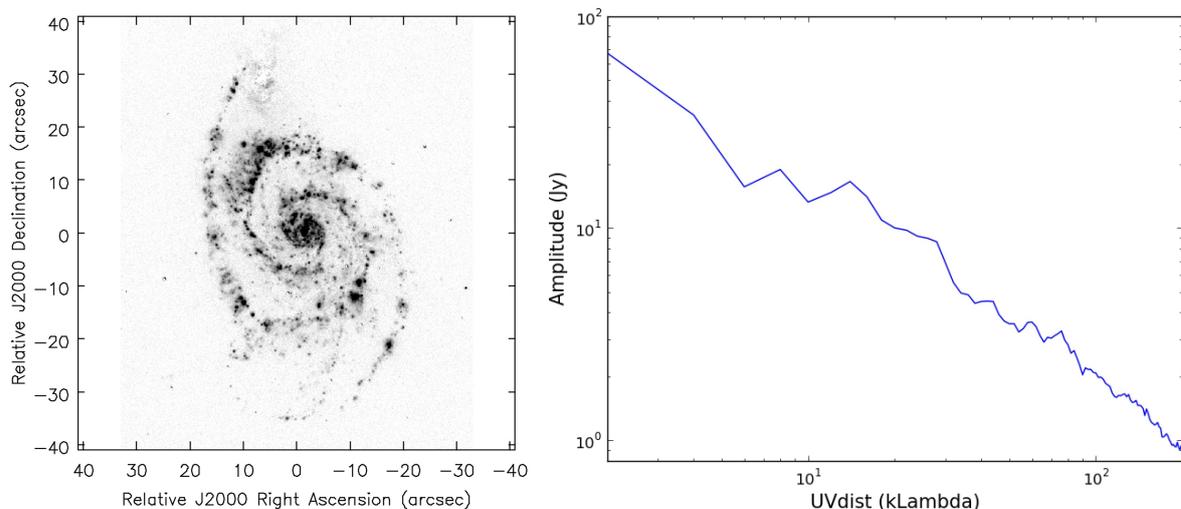


Figure 7.14: Image of $H\alpha$ emission from M51 used as a model of emission at 100 GHz (*left*) and the expected visibility amplitudes with uv -distance (*right*).

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:

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

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 ($HPBW = 1.13 \frac{\lambda}{D}$ 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 $HPBW/\sqrt{2}$ (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 set up a mosaic of adjacent pointings with a user-defined spacing, though it is not recommended to exceed spacings greater than the constant noise pattern ($HPBW/\sqrt{2}$) 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} , which are a function of the longest baselines of the array. Similarly, the maximum recoverable scales, θ_{MRS} , of each array depends on the shortest baselines present. In order to achieve the requirements entered by the PI in the

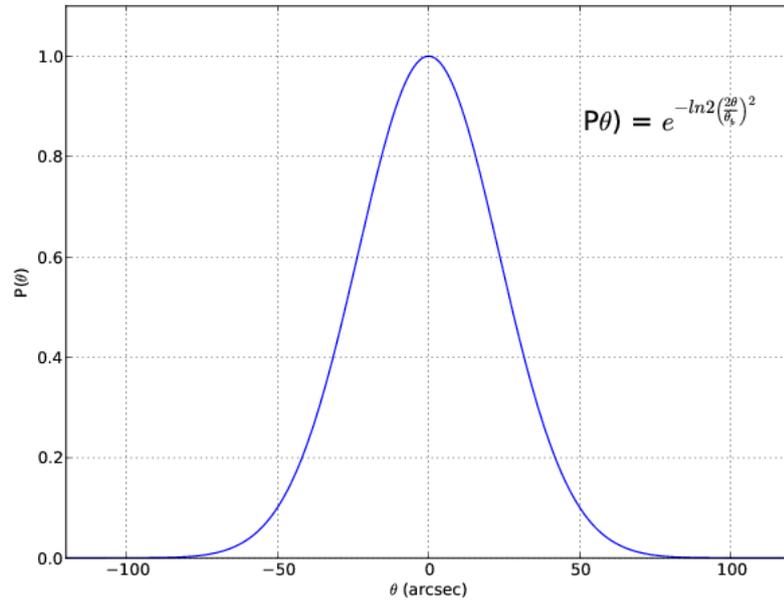


Figure 7.15: Approximation of the primary beam profile of an ALMA 12-m antenna at 112 GHz with HPBW = 52".

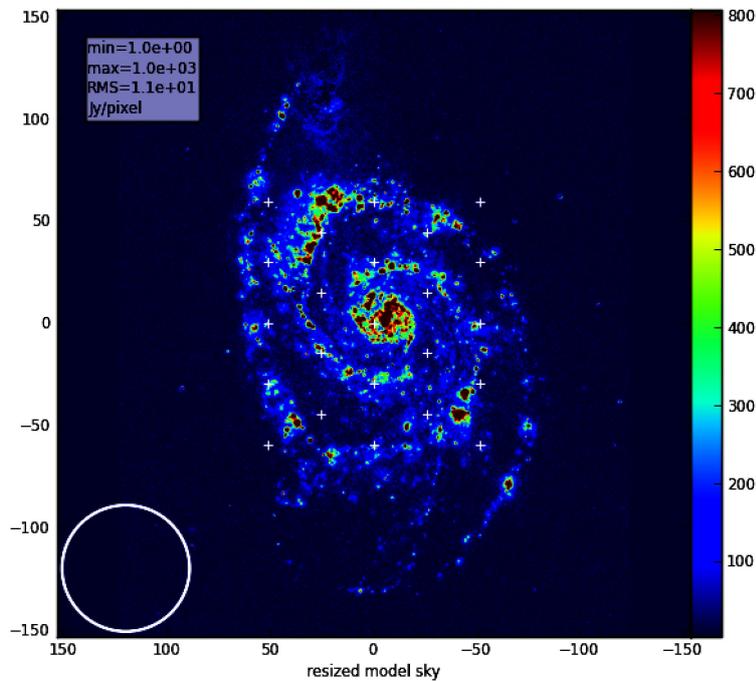


Figure 7.16: An example of mosaicing with a field of 2 arcmin at 100 GHz using an hexagonal pattern with Nyquist sampling (white crosses).

OT, θ_{res} and θ_{LAS} , multiple configurations may be needed. Note that θ_{LAS} is a property of the science target, while θ_{MRS} is a property of an array configuration. In particular, the smallest configuration for a project must ensure that θ_{MRS} is larger than or equal to the θ_{LAS} of the science target.

Based on the user-entered values of θ_{res} and θ_{LAS} , the ALMA Observing Tool will automatically attempt to choose configurations that will produce a final image with the requested scientific goals. Note however that there are significant restrictions when doing this. For example, at most two 12-m configurations are allowed and the three longest-baseline configurations (C40-7, -8 and -9) cannot be combined with any smaller configurations. If such restrictions result in the requested θ_{LAS} not being achievable, the OT will return a validation error and θ_{LAS} must be modified. In contrast, a new possibility for Cycle 4 is “stand-alone ACA”, i.e. observations that use the 7-m array, with the TP array added if required. It is not possible to request observations that only require the TP array.

In detail, the process of selecting the various arrays required to satisfy the entered constraints is the following:

1. An interferometric array (12- or 7-m) is selected that comes closest to achieving the entered value of θ_{res} .
2. If the θ_{MRS} of this array is less than the requested θ_{LAS} , the largest structures in the source will not be well imaged and smaller arrays must be added, if possible.
3. If the initial array was a medium-sized 12-m configuration, a smaller 12-m array can be added. If this array also has $\theta_{MRS} < \theta_{LAS}$, the 7-m array will then be added. If the initial array was one of the smaller 12-m configurations, it will only be possible to add the 7-m array.
4. If, after adding additional interferometric arrays the requested θ_{LAS} has still been achieved, the TP array will be added. Note that this is not possible for single-continuum observations and for Bands 9 and 10 so there is a fundamental limit to how large a source that can be reliably imaged in these situations.

This procedure is carried out automatically by the OT and is entirely dependent on the values of θ_{res} and θ_{LAS} that are entered by the user. A detailed list of the allowed combinations is given in Table 7.4.

Table 7.4 also shows the relative integration times that will be used by each array. Broadly speaking, because of their smaller collecting areas, the TP and 7-m Arrays require more on-source integration time than the 12-m Array. The needed time ratios have been calculated by requiring each array to have similar sensitivity in the region of the uv plane where they provide overlapping baselines; details of this calculation are in Mason & Brogan (2013). The actual time ratios adopted for ALMA observations are guided by the calculations outlined in this memo, but are subject to further operational constraints. Cycle-4 implements for the first time *configuration-dependent* time ratios. That is, the time multiplier for e.g. the 7-m array depends on which 12-m configuration it is being combined with. A trend can be seen whereby the smaller the 12-m configuration, the more 7-m array time is required. This is because increasingly compact 12-m configurations contain more sensitivity on baselines similar to those provided by the 7-m array, which must therefore observe for longer in order to contribute significantly to the final image. The TP array’s time multiplier is fixed at 1.7 times the 7-m array time.

As the compact arrays do require extra observing time, it is important to give consideration to the requested value of θ_{LAS} and to justify the final choice. 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 4, it is recommended that at least one hour of unique 7-m Array hour-angle coverage is needed to provide sufficient uv -coverage for good image combination with the 12-m Array, and snapshot observations are therefore strongly discouraged. A comprehensive plot about the array combination is shown in Fig. 15 of the ALMA Primer for Cycle 4 (Moriarty-Schieven et al., 2015).

The process whereby the OT chooses 12-m configurations is done assuming the exact antenna positions that make up the nine configurations that have been selected for Cycle 4. In reality, the configuration available to the scheduling software at the time of SB execution rarely conforms to this ideal, due to the fact that the array is often being reconfigured i.e. antennas are being moved from one pad to another, a process that takes some time. Therefore, it is usually the case that the available configuration’s properties, θ_{MRS} and θ_{res} , do not exactly conform to any of those assumed by the OT. However, providing that the available configuration provides a similar range of spatial scales to that which was assumed, it can be considered for SB execution providing that

θ_{res} (arcsec)	θ_{LAS} (arcsec)	Array combination	Time ratios	Total Time
0.066	< 0.78	C40-9	1	$1.0 \times \Delta_{extended}$
0.066	0.78-3.1	C40-9 + C40-6	1 : 0.3	$1.3 \times \Delta_{extended}$
0.066	> 3.1	-	-	-
0.12	< 1.3	C40-8	1	$1.0 \times \Delta_{extended}$
0.12	1.3-6.0	C40-8 + C40-5	1 : 0.3	$1.3 \times \Delta_{extended}$
0.12	> 6.0	-	-	-
0.21	< 1.8	C40-7	1	$1.0 \times \Delta_{extended}$
0.21	1.8-8.9	C40-7 + C40-4	1 : 0.3	$1.3 \times \Delta_{extended}$
0.21	> 8.9	-	-	-
0.35	< 3.1	C40-6	1	$1.0 \times \Delta_{extended}$
0.35	3.1-13.7	C40-6 + C40-3	1 : 0.3	$1.3 \times \Delta_{extended}$
0.35	13.7-66.7	C40-6 + C40-3 + 7-m	1 : 0.3 : 0.4	$1.7 \times \Delta_{extended}$
0.35	> 66.7	C40-6 + C40-3 + 7-m + TP	1 : 0.3 : 0.4 : 0.68	$1.98 \times \Delta_{extended}$
0.54	< 6.0	C40-5	1	$1.0 \times \Delta_{extended}$
0.54	6.0-22.1	C40-5 + C40-2	1 : 0.3	$1.3 \times \Delta_{extended}$
0.54	22.1-66.7	C40-5 + C40-2 + 7-m	1 : 0.3 : 1.4	$2.7 \times \Delta_{extended}$
0.54	> 66.7	C40-5 + C40-2 + 7-m + TP	1 : 0.3 : 1.4 : 2.38	$3.68 \times \Delta_{extended}$
0.93	< 8.9	C40-4	1	$1.0 \times \Delta_{extended}$
0.93	8.9-29.0	C40-4 + C40-1	1 : 0.3	$1.3 \times \Delta_{extended}$
0.93	29.0-66.7	C40-4 + C40-1 + 7-m	1 : 0.3 : 3	$4.3 \times \Delta_{extended}$
0.93	> 66.7	C40-4 + C40-1 + 7-m + TP	1 : 0.3 : 3 : 5.1	$6.4 \times \Delta_{extended}$
1.5	< 13.7	C40-3	1	$1.0 \times \Delta_{extended}$
1.5	13.7-66.7	C40-3 + 7-m	1 : 1.4	$2.4 \times \Delta_{extended}$
1.5	> 66.7	C40-3 + 7-m + TP	1 : 1.4 : 2.38	$3.38 \times \Delta_{extended}$
2.4	< 22.1	C40-2	1	$1.0 \times \Delta_{extended}$
2.4	22.1-66.7	C40-2 + 7-m	1 : 5	$6.0 \times \Delta_{extended}$
2.4	> 66.7	C40-2 + 7-m + TP	1 : 5 : 8.5	$9.5 \times \Delta_{extended}$
3.7	< 29.0	C40-1	1	$1.0 \times \Delta_{extended}$
3.7	29.0-66.7	C40-1 + 7-m	1 : 5	$6.0 \times \Delta_{extended}$
3.7	> 66.7	C40-1 + 7-m + TP	1 : 5 : 8.5	$9.5 \times \Delta_{extended}$
12.5	< 66.7	7-m	1	$1.0 \times \Delta_{extended}$
12.5	> 66.7	7-m + TP	1 : 1.7	$1.7 \times \Delta_{extended}$

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 and 7-m Array observations are run in parallel; b) for intermediate values of θ_{res} , please see text.

the other parameters (weather, project rank, LST range and executive balance) are also appropriate. This flexibility in scheduling is vital to the efficient execution of SBs and is built in to the second stage of ALMA's Quality Assurance process (QA2) such that all observed projects should meet the user's imaging requirements, within some threshold. Note that some additional flexibility is provided by the assumption of Briggs' weighting during the imaging stage, Fig. 7.6, a scheme which allows a trade off between angular resolution and sensitivity.

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. When all observations are complete, we will also— on a best-efforts basis— combine them to make and deliver a final image, again if it passes quality assurance.

The combination procedure comprises either one or two steps, depending on whether total power data are collected. First, the interferometric data will be imaged together in a single, “joint” deconvolution. This can be done in CASA by passing all interferometric measurement sets— 7-m and one or more 12-m arrays— directly to the CLEAN task¹. All 7-m and 12-m data delivered during Cycle 4 will be directly combinable. *Manually* calibrated data deliveries from earlier ALMA cycles that were calibrated using CASA versions earlier than 4.3 require an additional step to put the 7-m and 12-m data weights on an equal basis; however *all* ALMA data calibrated by the ALMA pipeline have correct weights, irrespective of cycle and CASA version. The steps to correct data weight issues for manually calibrated data from earlier cycles are described in the “Data Weights and Combination” CASA guide². Using the STATWT method is usually the easiest and most reliable approach assuming sufficient line-free spectrum can be identified.

If total power data are part of your project then there is a second step to the combination procedure, which is to combine the deconvolved interferometric (7-m+12-m) image cube with the total power image cube. In Cycle 4 this will be done by “feathering” them together. Feathering is a commonly used technique in radio imaging of combining two images together by forming a weighted sum of their Fourier transforms. The steps are as follows:

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

The resulting cube provides high angular resolution due to the interferometric data, but will have more accurate large-scale features and total flux densities because the single dish data are present as well. An excellent discussion of feathering and comparison to other techniques for combining single dish and interferometer data is given in Stanimirovic (2002).

A 1-hour observation with the C40-1 configuration and 5 hours with the 7-m Array is shown in Figure 7.17 (uv -coverage) and Figure 7.18 (uv -plane sampling distribution). The Cycle 4 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 C40-compact configuration (C40-3), a compact and extended configuration (C40-3 + C40-6), 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 C40-6 configuration at 100 GHz provides an angular resolution of $0.27''$. Combined with the C40-3 configuration, the angular resolution is lowered to $0.31''$. Finally, the angular resolution with the full combination (C40-3 + C40-6 + 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.

These simulations illustrate two general features of multi-configuration/array observations. First, they are effective in retrieving a wider range of spatial scales than are accessible to a single configuration or array. Second, the addition of more compact configurations or arrays does tend to broaden the synthesized beam (PSF) slightly. This broadening can be mitigated by changing the uv weighting, *e.g.* to uniform, at the expense of sensitivity.

¹Or TCLEAN, an improved implementation of CLEAN. In a near future CASA release CLEAN and TCLEAN will be renamed and the old implementation deprecated.

²<https://casaguides.nrao.edu/index.php/DataWeightsAndCombination>

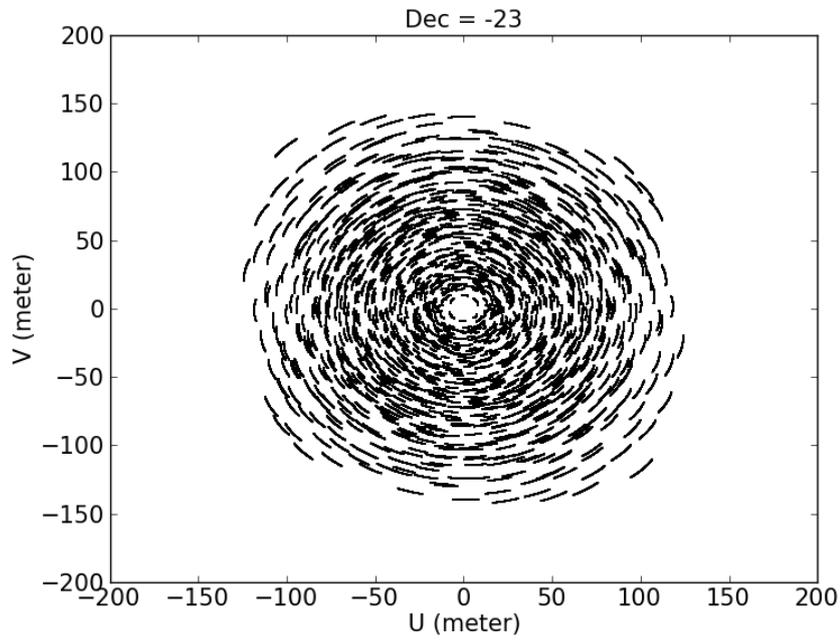


Figure 7.17: Expected uv -coverage for C40-1 (1 hour) and 7-m observations (5 hours)

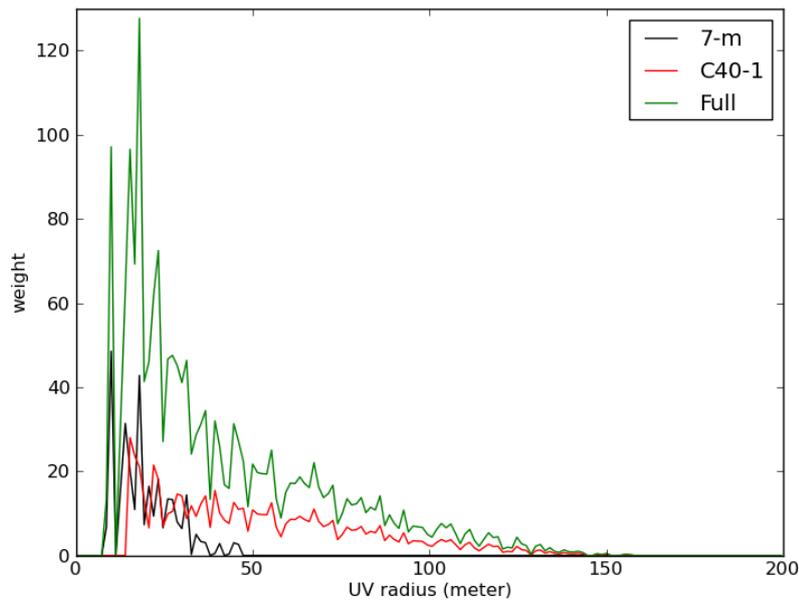


Figure 7.18: Radial density of uv -coverage with uv -distance using C40-1 (1 hour) and 7-m observations (5 hours)

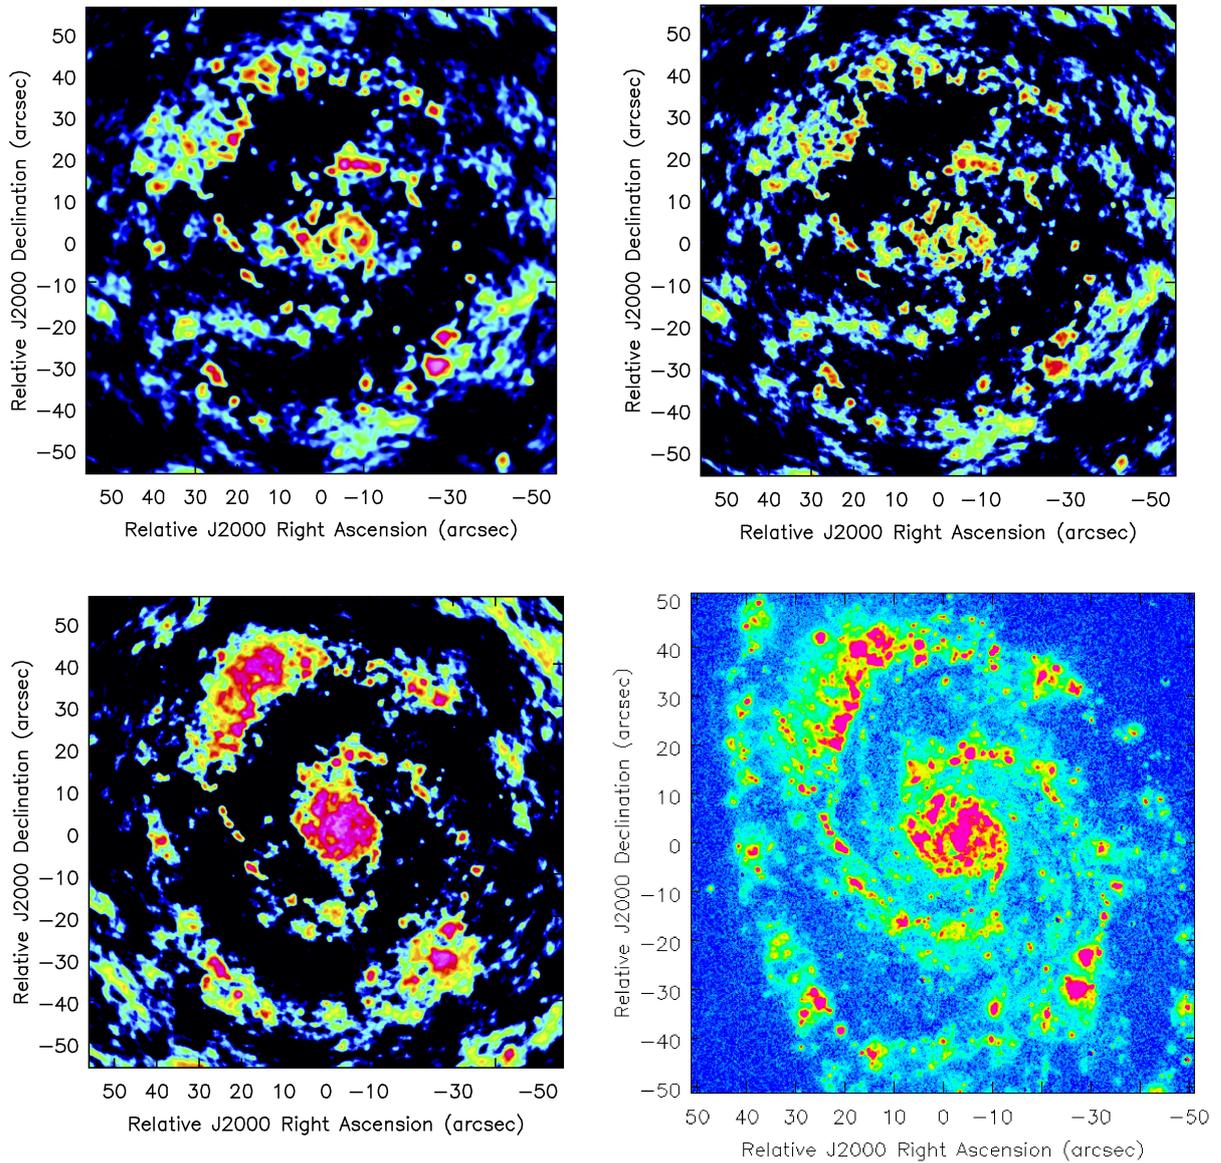


Figure 7.19: Images obtained using C40-3 (*top left*; 1 hour), C40-3 + C40-6 (*top right*; 1 + 3.3 hours) and C40-3 + C40-6 + 7-m (*bottom left*; 1 + 3.3 + 1.3 hours) array combinations, and the image model itself (*bottom right*).

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] Mason, B. & Brogan, C., 2013, Relative Integration Times for the ALMA Cycle 1 12-m, 7-m, and Total Power Arrays, ALMA memo 598
- [5] Moriarty-Schieven, G. ed., 2013, *Observing with ALMA: A Primer for Early Science*, ALMA Doc. 2.1, ver. 1
- [6] Stanimirovic, S. 2002, in *Single-Dish Radio Astronomy: Techniques and Applications*, edited by S. Stanimirovic, D. Altschuler, P. Goldsmith, & C. Salter, vol. 278 of *Astronomical Society of the Pacific Conference Series*, 375. astro-ph/0205329

Chapter 8

Observing Modes

The ALMA Observing modes are the set of capabilities that ALMA offers each cycle to its base community. In order to execute an approved ALMA project, the particular observing mode needs to be implemented in the ALMA online software package and verified using the required hardware. Here we described the software implementation along with a description of particular observing modes.

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

8.1 Observing Project Structure

The organization of each science project is subdivided into a well defined structure with clear hierarchical levels. At the bottom of this structure are the *Scheduling Blocks* (hereafter SB). They are the minimum set of instructions that describe an ALMA observation. The SBs are produced by the ALMA Observing Tool (hereafter OT) which is the tool used during the Phase 1¹ stage of the project life cycle. The SB may be edited using the OT in Phase 2 proposal acceptance and before/during proposal execution². Projects are broken down into a set of these fundamental units for flexibility, given the properties of the ALMA site and the continuously-varying status of the Observatory as a whole (including the weather), and to encapsulate the scientific objectives of the proposal into single entities.

An SB contains a large amount of information about what it should be observed, 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 type of observation required e.g., single-dish, single field interferometry, etc. The SB has relatively little influence over the order in which the various sources are observed and does not describe all of the calibrations that should be performed (a prime example being the measurements of the system temperatures).

Each SB typically executes set-ups, calibrations, and target observations within 1 hour. The end of an SB 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 SB runs to completion, fails, or is terminated by the *Astronomer on Duty* (or AoD). Given the limited length 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

¹For a description of Phase 1 and Phase 2 see the Cycle 4 Proposer's Guide

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



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

been chosen based on statistical measurements about 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 two 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.

All SBs that have to be processed together to produce calibrated science 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. For polarization observations, all the SBs that belong to the same session should be grouped together for data processing. Unfortunately, this is not explicitly described in the SB.

If the calibrated science observations of a Member OUS have to be combined with science observations of another Member OUS, then these are grouped together into one Group OUS. Otherwise, they will be placed into different Group OUSes. A typical case of Member OUSes that belong to the same Group OUS are observations of the 12-m Array, the 7-m Array and the TP array. According to the method above, these observations would be in three separate 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 (see Chapter 13 for more details). It is expected that in future cycles, processing will also happen 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, irrespective of their degree of completion. As polarization observations do not currently fit this scheme, the data reduction will be done manually by the ALMA staff.

The SB/Member OUS/Group OUS are the smallest structures that hold science observations that need to be observed/processed/combined together. These data hierarchy therefore maximizes observing flexibility and ensures that data gets processed as early and as fast as possible. In the OT, observations are divided into different Science Goals (see the ALMA Proposer's Guide). In order to follow this structure, each 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 Astronomer on Duty, i.e., AoD), it is read into the online CONTROL software, particularly by the Science Software Requirements subsystem (hereafter SSR). The SSR subsystem commands the lower level CONTROL software to create an Execution Block (EB) structure that is attached to the SB. As 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 the SB was executed along with references to the acquired data. The internal hierarchical structure of the EB is also shown in Figure 8.1. The SSR subsystem constructs a sequential series of scans for each of the required calibrations and command 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 (the correlator software only understand subscan sequences). 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 five 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. 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 SBs construction, particularly 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.

8.3 The Observing Process

For a given project, a set of targets are specified in the SB. These targets can correspond to either calibration or science executions. Targets are organized in observing groups where the first group (group 1) is always the initial calibration group, with subsequent groups detailing the science observations³. All Science targets and relevant calibrators within a group are observed before the next group is started. An SB can have multiple groups.

All groups other than group 1 are considered complete when all Science targets in the group have been observed for the requested time or have set below the elevation limit. After all groups are completed or the SB execution 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, grouped source executions, pointed mosaics, and polarization use the same ALMA observing script, called the *standard interferometry* script. The set of necessary calibration measurements (e.g. amplitude, bandpass, etc) usually specified in group 1 are performed at the beginning of the observation sequence if appropriate 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.

Pointing, Atmospheric, and Sideband-Ratio calibrations are associated to the main (Bandpass, Amplitude, Polarization, and Phase) calibrations on a as-needed basis which is determined by the SSR at run-time. However 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 turn until observation of the primary phase calibrator (the calibrator with shortest cycle time in the group) is required. A typical cycle time for the phase calibrator may be 7-10 minutes. 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 alternatives 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 calibration for a given scheduling block is done using astronomical sources selected at runtime. To correctly select the appropriate calibration sources, selection criteria were implemented in the SSR using the following steps:

- A list of sources is retrieved from the ALMA catalog based on the query center and the search radius defined in the SB.
- For each source in the list we check whether there are flux measurements at the SB representative frequency. If no flux measurement is present, we extrapolate the flux using a spectral index of -0.7 to the SB representative frequency.

³However, this is not strictly correct. Additional calibrations can be added into subsequent groups and they will be interleaved according their cycle times. One example is the polarization calibration

- The flux is weighted based on proximity to the representative frequency, target and time since the last measurement present in the source catalog.
- The returned list of sources that passed the signal to noise criteria are ranked based on source strength, separation respect to the query center and flux error.
- The final list is sorted based on ranking.

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

- Bandpass calibrator. A signal to noise of 50 given 15 minutes integration is used in the flux estimation. The smallest spectral window (bandwidth) is used for this calculation. The default search radius is defined to be 45^{deg} from the query center, but the user can overwrite this default as needed. A minimum elevation is calculated to avoid antenna shadowing.
- Phase calibrator. A signal to noise of 15 given 2 minutes integration is used in the flux estimation. The widest spectral window bandwidth (usually 2 GHz) is used for this calculation. A 20 degrees from the query center search radius is used as default (smaller cycle times and radii are used for the highest frequencies and longest baselines, as described in Chapter 10). However, value is taken from the SB and the closest source to the target, from the selected list, is returned.
- Flux calibrator. A solar system object is searched first, then a grid source. First observable grid source is picked up. A detailed list of the Solar System objects and grid sources used for flux calibration is available in Appendix C.

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 flux scale 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 flux scale calibrator 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.

The interferometer is a device for measuring the spatial coherence function (Clark 1999), the phase information of the sources must be preserved and discerned from phase variations at the individual antenna elements. This is achieved with the phase calibration during the course of an observation. Since the phase is expected to change much more rapidly in time than the amplitude, 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 fluctuates rapidly (specially at higher frequencies), radiometric observations of atmospheric water lines are done to correct for these additional phase variations. This corrections can be applied online and offline (see section 10.4.2 for further details).

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 execution. In pointed mosaic observations, each of the fields will be assigned different field IDs in the data, but the same source ID. 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 by a number of subscans corresponding to the individual rows in the mosaic (see Section 5.5). Both the Science target specific last mosaic pointing position and index within the Science target list are handled by the SSR to ensure the next scan begins on the proper Science target and proper offset position. The mosaic observations can be set up by specifying a *Rectangular Field* in the OT. If the user wishes to execute several small mosaics within an SB, the *custom-mosaic* under *Multiple Pointings* should be selected.

8.6 Single Dish (Total Power) Observations

The purpose of adding data from single-dish observations using autocorrelations 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* (hereafter TP), although in practice they are taken in autocorrelation mode rather than using a total power (square law) detector. Four 12-m antennas connected to the ACA Correlator are available for this purpose. For this cycle, only spectral-line observations are offered in this mode.

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 (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 autocorrelation data will be written to the ASDM to minimize data rate and size, whereas calibrations use cross-correlation information for analysis.

Most of the calibration scans are executed in group 2 due to limitations in the observing script in cycle 4. Group 2 can contain pointing observations (if a calibrator is attached to that group and none have been carried out in group 1), and the atmospheric calibrations at regular intervals. another round of pointing calibration, then a sideband ratio measurement. For the higher frequency bands, a focus calibration is executed next. Then another atmospheric calibration is taken at the reference position of the science target, then the actual OTF map of the science target.

The OTF map is observed as a series of raster rows, scanning in the coordinate system specified in the SB. In cycle 4, 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 turn, to minimize scanning effects. The reference positions, assumed to be positions which are free of emission, are specified either in absolute coordinates or as offsets from 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, with a maximum length of a scan (consisting of some number of subscans) of 600 seconds. The reference position is observed as specified by its cycle time during the science target scans. Pointing is calibrated on a bright calibrator near the Science target with a frequency indicated in the SB, and atmospheric calibrations are taken every 15 minutes at the reference position to measure the system temperature.

The calibrated TP map of the science target will be in units of Kelvin, on the antenna temperature (T_a^*) scale. Since all TP observations are expected to be combined with 12-m and 7-m Array data which come in calibrated units of Jy/beam, the TP data must also be converted into these units. This conversion from Kelvin

to Jy/beam requires knowledge of the main beam efficiency η_{mb} and the beam size θ . These are measured separately for each project by obtaining a continuum map of a bright quasar or a planet with known flux⁴. These “amplitude calibration” maps are reduced in the same way as the science observations, then the emission will be compared with their model or observed flux (taken from the most recent calibrator survey measurements in the case of quasars) to calculate the Kelvin to Jansky/beam 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 splitter⁵. In this way, the incoming radiation is separated into two orthogonal components (X and Y) which are down-converted and digitized independently. For each baseline, the digital signals are cross-correlated at the correlator where the outputs are the four cross-correlations XX, YY, XY, and YX (or V_{xx} , V_{yy} , V_{xy} , and V_{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), and 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 the frequency and the 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 parallactic angle (ψ). By design, the instrumental polarization is small (a few percent), but not negligible.

⁴It is expected that during Cycle 4, these measurements will only be needed for the high-frequency ALMA bands, while for bands up to band 7, the values will be automatically derived from fits to the aperture efficiency of each antenna as a function of time, elevation, outside temperature, surface accuracy, etc., which will be regularly monitored by the observatory.

⁵ For the ALMA receiver bands offered in polarisation mode in Cycle 4, bands 3 and 6 the polarisation splitter is an *Ortho-Mode-Transducer* (OMT), and for band 7 it is a polarized grid

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. It has also been observed that the usage of the different photonic references also introduces an offset in the XY-Phase.
3. The usual off-line calibration procedure arbitrarily sets the phase of a reference antenna to zero; this results in the inverse of its X-Y offset being imprinted on the cross-hand phases of the other antennas.
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 4 (see below for a detailed explanation.)

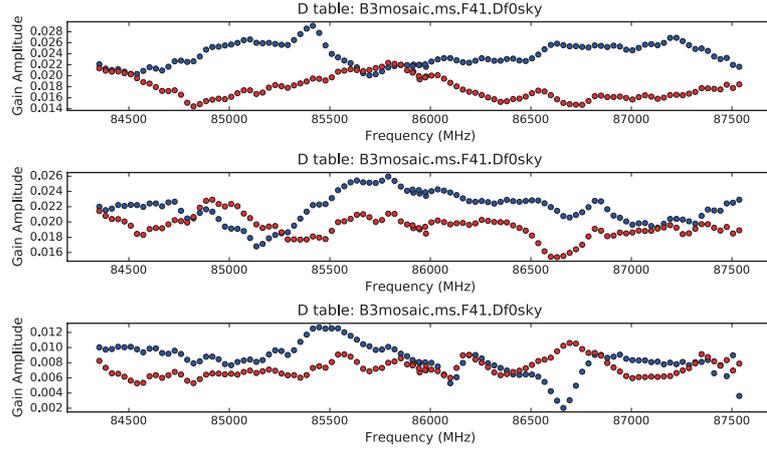


Figure 8.2: The D – term plots of DA42 (top), DV03 (middle), and PM01 (bottom). The vertical axis is the fraction of the input signal voltage in one polarization that leaks into the output of the other polarization in voltage units and the horizontal axis is the frequency in MHz. The blue and red symbols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

By taking into consideration the instrumental polarization and the parallactic angle dependence, 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 8.2 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 frequency⁶. 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.

ALMA therefore normally uses an unresolved quasar as a polarisation calibration source. These are monitored to ensure they are bright enough, and the exact polarization is determined during data reduction provided that the source is observed over a wide enough range ($> 60^\circ$) of parallactic angles to separate the effects of the

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

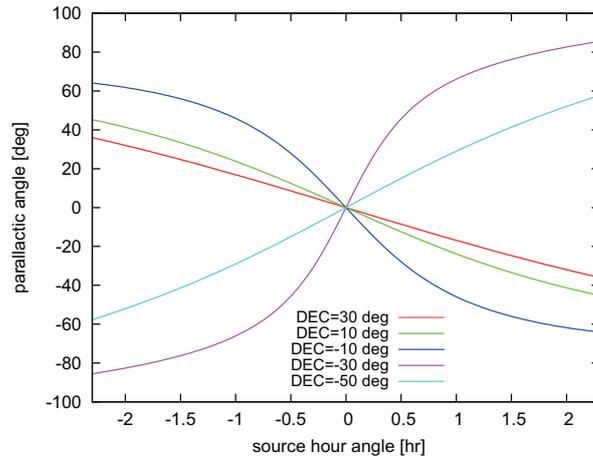


Figure 8.3: Parallax angle (ψ) plot as a function of hour angle for the declination of 30, 10, -10, -30, and -50 degrees.

source polarisation, the D – terms and the cross-hands phase spectra and delay. Fig 8.3 shows that the rate of change of parallax angle is fastest near transit and slowest for low elevation sources. The D – terms may have a slight elevation dependence, and therefore it is favourable 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.

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 (based on initial polarisation mode verification results), with a contribution of 2 degrees coming from the specification on the orientation of the receiver feeds, and an additional scaling of the error with the number of antennas in final image. Of course, these levels of accuracy will only be reached if there is sufficient signal-to-noise in the polarised emission. 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 4, the off-axis instrumental polarization calibration will not be employed. In Cycle 4, polarisation imaging with accuracy better than 0.3% can be achieved within the inner 1/3 of the primary beam FWHM and this determines the largest acceptable angular size of sources which can be observed in full polarisation. Users must justify that their science goals can be met by sources exceeding these lower limits (of 0.1% polarisation for a point source in the centre of the field or 0.3% within up to the inner 1/3 of the FWHM).

For this cycle, polarization observations have been extended to FDM correlation mode at arbitrary frequencies and frequency resolutions within the offered Bands (3, 6, and 7). Figure 8.4 shows the D – term solutions of a representative antenna (DV22) obtained during commissioning. Figure 8.5 shows a zoomed-in view in frequency of the D – term solutions shown in Figure 8.4. These figures show that the spectral shape of the D – term solutions is smooth and no spur like structure is seen even in highest frequency resolution mode (30.5 kHz at the right bottom panel of Figure 8.5). Figures 8.4 and 8.5 clearly demonstrate that the instrumental polarization is mostly spectral resolution independent for these datasets. No mosaic, ACA, circular polarization and TP Array observations are offered. Further, a minimum execution time of 3 hours will be imposed to ensure sufficient parallax angle coverage for calibration.

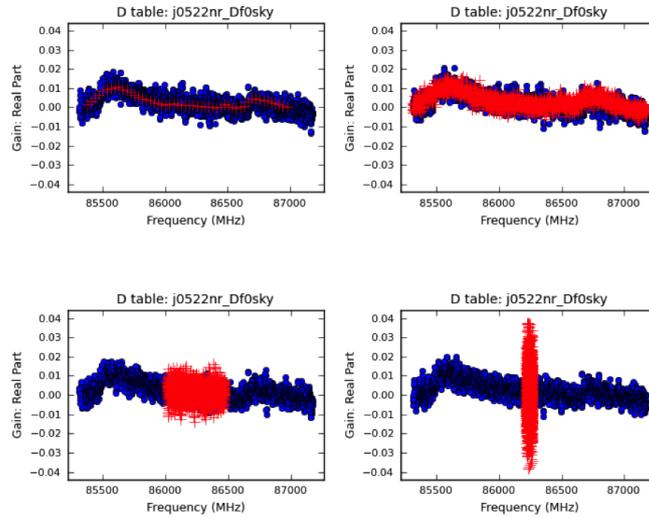


Figure 8.4: A frequency vs real part the *D-terms* plot for an antenna (DV22) obtained towards J0522-364. Blue circles and red crosses are showing the *D-term* obtained using a setup of FDM with 2GHz bandwidth and using a setup with higher spectral resolutions (FDM 500 and 62.5 MHz bandwidth respectively).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

8.7.1 Sessions

Though not directly related to polarization itself, sessions are the observing scheme designed to observe 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 calibration⁷ science target loop will be interrupted (preserving phase calibration) when an additional calibration is-needed (e.g. polarization)⁷. The current mode for sessions in ALMA is the observation of polarization projects. Because for an Alt/Az antenna the frame of the sky rotates, the calibration of the 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 interleave the calibration when needed. In general, we can achieve the parallactic angle coverage by running the polarization SB 2 or 3 times, which will give between 3 to 4 hours of observation, which is a perfect example of the session scheme. Other cases in which 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 4. 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.

⁷The cycle time parameters are user controlled and can be explicitly specified in the OT target parameters section

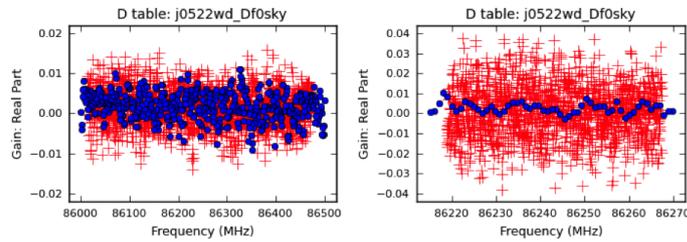


Figure 8.5: A close-up view in frequency space of the real part of D-term shown in Figure 8.4. Blue circles and red crosses are showing the *D-terms* obtained from setup #1 (FDM, 2 GHz bandwidth) and setup #2 (FDM, 500 and 62.5 MHz bandwidth respectively).

8.8 Multiple Region Modes

In the frequency division modes (FDM) of the correlator (see section 6.3.2), 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 disjoint 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 4, a maximum of 4 spectral windows per baseband will be offered.

8.9 Observation of Ephemeris Objects

Observation of solar system objects (with the exception of the Sun) is supported as in Cycle 3. Several well known solar system objects including planets, satellites and asteroids can be selected from a pulldown menu in the Observing Tool. For other sources including non-sidereal objects, an external ephemeris file can be supplied as an input. The ephemeris file must be in JPL Horizons format. A typical ephemeris file may consist of the date (time), Right Ascension, Declination, range and range rate, for example:

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

```
2012-Jun-26 13:05    06 22 58.14 +23 16 10.74 1.01653122901917 -0.2541942
2012-Jun-26 13:06    06 22 58.30 +23 16 10.66 1.01653122901917 -0.2531491
```

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

8.10 Solar Observations

Solar observing was part of the original science case for ALMA but it has not been implemented until now pending extensive tests to determine appropriate observing modes and calibration strategies. In Cycle 4, solar observing with ALMA will be supported for the first time, albeit with limitations. In particular, the following conditions apply to solar observing in Cycle 4:

- Only Band 3 and Band 6 continuum observations using the following default frequencies are offered

```
*****
                                LSB                                USB
Band      LO Freq.           BB1           BB2           BB3           BB4
*****
3          100 GHz           92-94 GHz    94-96 GHz    104-106 GHz  106-108 GHz
6          239 GHz           229-231 GHz  231-233 GHz  245-247 GHz  247-249 GHz
```

- Interferometric observations must be done using the TDM mode (see section 5.1.1); frequencies are fixed to 2 GHz-wide spectral windows centered on the frequencies shown in the table above. The spectral-line FDM observing mode is not offered for solar observations (section 5.1.2).
- Simultaneous observations with Bands 3 and 6 are not offered: each execution block can only include one band.
- Observations will be carried out while the array is in a compact configuration corresponding to C40-1, 2, or 3 (see section 5.3.2 of this Proposers Guide).
- Because the WVR receivers are saturated when dishes point at the Sun, on-line WVR phase correction will not be applied and off-line WVR correction for on-source (solar) data is not possible.
- Interferometric observations will use a combined array comprising both 12m and 7m dishes and will be processed with the baseline correlator (section 5.1).
- To minimize shadowing of 7m antennas, observations will be carried out between 13:00 UT and 20:00 UT.
- Mosaicking of larger FOVs can be carried out with up to 150 different pointing offsets relative to the phase center specified in the ephemeris, typically in a rectangular pattern.
- One target (one ephemeris file) is allowed for each science goal.
- A single observation (an execution block) cannot exceed 2 hours, which will include the time overheads for bandpass and flux calibration. These calibration overheads amount to about 25 mins.
- The integration time of interferometric observations is fixed to 2 seconds.
- Observations may be performed using dual linear polarization (XX, YY) or single polarization (XX) correlations; full polarization measurements are not currently offered for solar observations.
- Total-power full-Sun fast-scanning single-dish observations are offered as context and to recover the largest angular scales for interferometric observations; proposals requesting only total-power single-dish observations will not be accepted in Cycle 4.

- Single-dish total-power mapping will be “monomode” (double-circle scan pattern) with a 2400arcsec-diameter circular FOV centered on the solar disk (see below).
- Single-dish full-Sun will not be executed when the Sun is at elevations above 70° because the required fast-scan azimuth slew speeds are too high.
- The time cadence of full-sun images obtained from total power observations is fixed to about 7 minutes for Band 3 and 10 minutes for Band 6.

8.10.1 Solar Observing Modes

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

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

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

Two so-called “mixer de-tuned” or “mixer de-biased” settings have been adopted for solar observations in Band 3 and Band 6 for Cycle 4. These are referred to as solar observing modes MD1 and MD2.

- **Band 3:** MD1 mode uses a bias voltage that sets the SIS mixer to the 2^{nd} photonic step below the voltage gap whereas MD2 mode employs a bias voltage corresponding to the 2^{nd} photonic step above the gap. For MD1 mode, the Band 3 receiver temperature suffers a modest increase, to ~ 50 K, and receiver compression is limited to $\sim 10\%$. For MD2 mode, however, the Band 3 receiver temperature increases significantly, to 800 K, but receiver performance is believed to be essentially linear.
- **Band 6:** MD1 mode settings are nominal for ALMA; i.e., the bias voltage is that used under normal observing conditions. Receiver compression is again of order $\sim 10\%$. MD2 mode employs a bias voltage corresponding to the 1^{st} photonic step above the voltage gap. Again, the receiver temperature is significantly higher, 800 K, but the receiver is essentially linear.

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

8.10.2 Array Observations

The Sun is an extremely large source compared with the primary beam of the either the 7m or 12m antennas. The primary beam of both antennas is filled with complex emission when pointing at the Sun, as are the beam sidelobes. The ALMA array ultimately measures the brightness temperature contrast relative to the background

Sun, which is resolved out by the array. As noted earlier, in order to recover the absolute brightness temperature of solar targets, it is necessary to include not only interferometric observations (by the 7m and 12m antennas), but also total power measurements made with a single dish. Single dish fast-scan mapping of the Sun in total power mode is addressed in the next subsection.

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

Another consideration, again regardless of whether SFs or the MD modes are used, is the system IF attenuator settings. The input power changes significantly as the antennas move from the (solar) source to a calibrator and back. The IF chain has two variable attenuators (in steps of 0.5 dB) to ensure that signal levels remain within nominal limits: one in the IF Switch and one in the IF Processor. A concern is whether the variable attenuators themselves introduce unacceptable (differential) phase variation between source and calibrator settings, thereby corrupting phase calibration referenced against suitable sidereal calibrators; and whether there are differences between the spectral window bandpass response between source and calibrator scans as a result of attenuator settings. Careful testing has shown that this should not be a significant concern. While these tests show that phase shifts caused by the attenuation level changes do in practice difference out, verification that this is the case cannot be checked from observing data obtained using the solar scheduling block. As a check, the observatory will carry out a test observation of a calibrator source using normal and MD attenuation levels before solar observations begin on a given day or at least once before a campaign program.

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

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

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

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

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

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

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

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

8.10.3 Single Dish Mapping

Fast scanning observations are ideal for recovering the flux or brightness distribution on angular scales ranging from the ALMA primary beam width to the scale of the target in question (typically a few arcminutes), or up to the full disk of the Sun. Briefly, fast-scan mapping entails making total power (and more recently, autocorrelation measurements) as the telescope pointing is driven continuously and smoothly through a sampling pattern on the target that avoids sudden acceleration or deceleration of the antenna drive motors. A major advantage of fast scanning is that it minimizes the impact of atmospheric variation, and the full solar disk can be mapped in as little as 7 minutes.

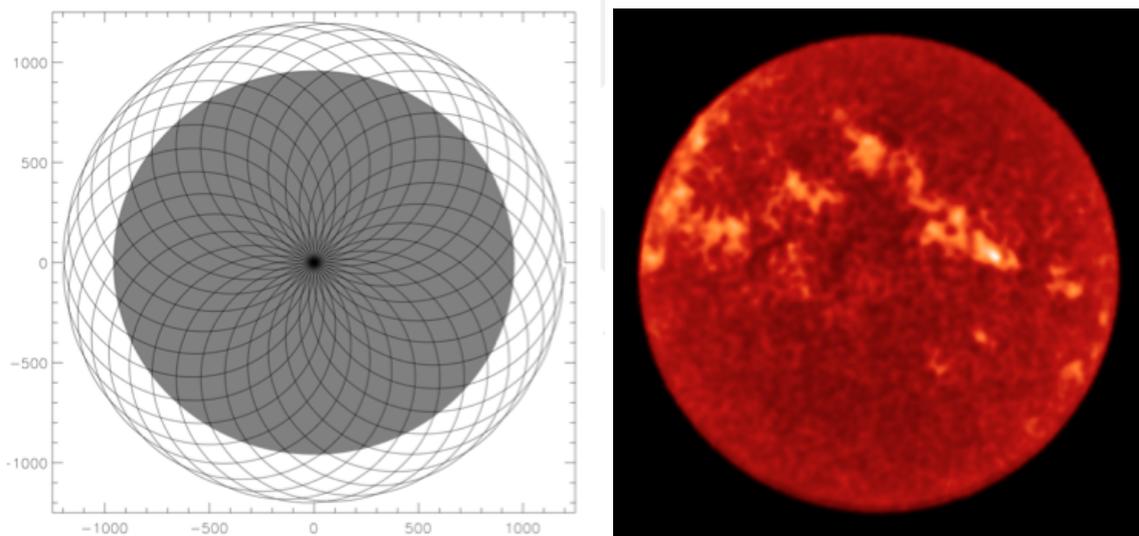


Figure 8.6: Left: Schematic illustration of a double-circle scan pattern used to map the full disk of the Sun. Actual patterns sample the source more densely to ensure that it is no less dilute than Nyquist sampling. Right: Example of a fast-scan map of the Sun obtained in Band 6.

While various types of scan patterns have been developed and tested for ALMA dishes, Cycle 4 supports the use of a “double circle” pattern, which map a circular region on the sky. The double-circle pattern is particularly well suited for full-disk mapping because its coverage matches the shape of the solar disk and it repeatedly revisits the region of the center of the disk, allowing atmospheric opacity variations to be corrected. Standard observing procedures include focus and pointing checks on suitable sources prior to the fast-scan mapping.

8.11 VLBI Observing Mode

The Very Long Baseline Interferometry (VLBI) Observing Mode (VOM) is a variant of the standard interferometry mode with some additional capabilities to allow ALMA to participate in global VLBI networks operating at

millimeter and submillimeter wavelengths. In Cycle 4, ALMA VLBI mode observing will be offered in Band 3 in conjunction with the the Global Millimeter VLBI Array (GMVA)⁸ and in Band 6 in conjunction with the Event Horizon Telescope (EHT) network⁹. For details about proposing for these opportunities, see the ALMA Cycle 4 Call for Proposals See <https://science.nrao.edu/observing/call-for-proposals/1mm-vlbi-cycle4/>.

For all VOM observations in Cycle 4, ALMA will be operated as a phased array of 12-m antennas (most likely containing between 31 and 37 phased antennas). For example, with 37 12-m antennas, this is an equivalent diameter of 73 m. From Section 9.2.1, Table 9.3 the aperture efficiencies in Band 3 and Band 6 are 0.71 and 0.68, and T_{sys} is 70 K and 100 K, respectively. Since $\text{Gain} \equiv A_{\text{eff}}/(2k_B)$ and $\text{SEFD} \equiv T_{\text{sys}}/\text{Gain}$, we have

$$\begin{aligned} \text{Gain} &= 0.71\pi(36.5^2) \times 10^{-26} = 1.08 \text{ K/Jy} \\ \text{SEFD} &= 65 \text{ Jy} && \text{(Band 3)} \\ \text{Gain} &= 0.68\pi(36.5^2) \times 10^{-26} = 1.03 \text{ K/Jy} \\ \text{SEFD} &= 97 \text{ Jy} && \text{(Band 6)} \end{aligned}$$

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

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

8.11.1 General VLBI Considerations

The supporting VLBI networks (i.e., the GMVA for Band 3 and the EHT for Band 6), together with ALMA, place certain restrictions on the observing process; these are largely driven by the complexity of orchestrating the simultaneous, reliable execution of a common VLBI schedule, as well as seeing to the calibration requirements of the individual telescopes that comprise the global array. In practice, for the ALMA VLBI observer there will be very little difference between planning an observation in the two observing bands, aside from a few minor details (i.e. differing bandwidths and frequencies). Successful proposals will be passed to a network Scheduler who builds the common VLBI schedules (encoded in a so-called VLBI EXperiment, or VEX file built with SCHED¹⁰) for a particular observing campaign. In these schedules, some number of contiguous hours will be devoted to each set of science targets and necessary VLBI calibrators.

The schedule will be worked out between the Scheduler and the PI well in advance of the observation. Most sites participating in the VLBI network also have a “friend” of VLBI who assists the Scheduler in working out details specific to that site. Schedules from previous campaigns are available from the aforementioned GMVA website and it is expected that schedules for Cycle 4 with ALMA with either the GMVA or the EHTC will be similar.

In general, the required time on the VLBI science target(s) and VLBI calibrators will be broken into a number of VLBI scans of several minutes duration, separated by gaps of order several minutes to allow time for local calibrations (pointing, system temperature measurements, etc.) and antenna slew time sufficient to meet the needs of each observatory. VLBI observations that include ALMA will require that the gaps between VLBI scans are of sufficient length to allow for the performance of the calibrations normally required for standard interferometric observing, including flux, gain, bandpass, and polarimetric calibrators (see below). Since all VLBI science targets for Cycle 4 are required to be sufficiently bright to allow phase self-calibration (≥ 500 mJy), observations of a gain calibrator will be required only every 20-30 minutes to allow calibration of the amplitudes.

It is important to recognize that most VLBI stations currently record *circular* polarization (either left, right or dual), in contrast to ALMA which records dual linear polarizations. Consequently, *every VLBI mode*

⁸See <http://www3.mpifr-bonn.mpg.de/div/vlbi/globalmm/>

⁹See <http://www.eventhorizontelescope.org/>

¹⁰<http://www.aoc.nrao.edu/software/sched>

observation with ALMA must be treated as a polarization observation in order that a transformation from linear to circular polarization can be made correctly during post-correlation processing. (A special tool, `PolConvert` has been written to be used with the standard VLBI correlation software, `DiFX`¹¹, which converts ALMA data to a circular basis, greatly simplifying the analysis.) In particular, this requires a minimum session duration of 3 hours (session length, not time on science target) in order to achieve adequate parallactic angle coverage on the polarization calibrator sources for the computation of polarimetric “leakage” or D terms.

Aside from differences in polarization, other VLBI observatories all have receiver capabilities that differ substantially from ALMA, especially with regard to bandwidth. In particular, none of the GMVA and EHT sites is currently able to match ALMA’s full recording bandwidth. An additional complication is that the sample rate used at ALMA is non-traditional for VLBI, so some care must be taken in the experiment set-up to ensure that the VLBI data can be correlated with the other stations. The VOM spectral set-ups chosen for each band in Cycle 4 was selected to address these constraints and are therefore fixed. For Band 3 (at 3mm) there are 4 1.875 GHz bands centered at 86.124, 88.124, 98.124 and 100.124 GHz, respectively. The GMVA sites will be tuned to a 256 MHz band centered on 86.124, so only the first band is available for VLBI. In Band 6 (at 1mm) the bands are centered at 213.1, 215.1, 227.1 and 229.1 GHz, respectively and the number of recorded bands will depend on the capabilities of the EHTC—for Cycle 4, use of the two upper bands is expected at some sites, and the 227.1 GHz band at all sites.

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

In a VLBI observation, the observatories are usually not connected by any network, so recordings (onto disk modules) of the antenna baseband signals are forwarded to common correlation facilities (managed by the appropriate network, *i.e.* MPIfR Bonn and MIT Haystack Observatory under arrangements with ALMA) which take responsibility for the correlation of the VLBI data. Correlated VLBI data products will then be made available to the observer, as discussed in Section 8.11.2

8.11.2 ALMA Considerations for VLBI

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

In order to participate in a VLBI observation, there are two special requirements at ALMA. The first is that the ALMA control system must tune receivers appropriately and point to the targets specified in the VEX schedule at the appropriate times. The second is that the signals from the ALMA antennas must be coherently summed, decimated, and recorded on media suitable for delivery to the common VLBI correlator. In order to form this coherent sum, the (arbitrary) phases of the signals at the antennas need to be adjusted during the observation so that they can be added “in phase”. The necessary signal processing for this takes place within special cards and circuitry within the ALMA correlator as well as with special software within the ALMA control system which were developed by the ALMA Phasing Project (APP)¹² We present some of the details of this in the following sections, with particular emphasis on those aspects which are essential either for the proposer or data analyst to be aware of. We begin with a discussion of the VOM Scheduling block in the next section. Then Section 8.11.2 discusses the scan sequences used by the VOM. That is followed by a section with some important details about how the ALMA Phasing System (APS) works (Section 8.11.2). Finally Section 8.11.2 makes a few comments about the analysis of VLBI data (a full discussion would be considerably outside the scope of this document).

¹¹<http://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/start>

¹²An NSF MRI/ALMA development fund project.

VLBI Observing Mode Schedule Block and Execution

As with other modes, VOM observations are executed with a scheduling block which provides observation-specific details to an SSR observing script. In the case of the VOM, the script is called `StandardVLBI.py`, although, as noted elsewhere, VLBI is a non-standard observing mode in Cycle 4.

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

Like the `StandardInterferometry.py` script, the `StandardVLBI.py` ensure that the necessary calibrations as specified by the OT are performed with the desired repetition rate. However, unlike the `StandardInterferometry.py` script, it must give precedence to the scheduled VLBI scans which need to occur at the appointed times. As mentioned above in Section 8.11.1, the VEX file will have provided gaps of various durations which are in fact created to allow the observing script to make these ALMA-specific calibration observations. Thus the script organizes its work by first noting when the next VLBI scan will occur and then noting which calibrations might fit into the time until the VLBI scan and finally executing those that fit. The VLBI scans in the VEX file include a start time specified in so-called VEX time format, `YYYYyDOYdHHhMMmSSs`, so the next result from an execution perspective is something similar to:

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

VLBI Scan Sequence

The VEX file specifies the observations in terms of VLBI scans which begin and end at some appointed times. The VOM translates this request into a sequence of (ALMA) correlator scans which typically start prior to the specified time in order to allow the phasing system to stabilize on a good solution. Each of these correlator scans appears in the ALMA meta-data, and is processed by the telescope calibration system (TelCal) according to the specified intents. The relevant ones for this discussion are intents for phasing and for WVR correction. The phasing intents direct TelCal to calculate the phases of the signals (*i.e.* the signals of each polarization in each band of each antenna) relative to those of a designated “reference” antenna. These phases are calculated from “channel average” data products calculated within the correlator. (Typically there are several channel averages calculated during one of these correlator subscans; see below). With those phases known, commands are issued to the tunable filter banks (TFBs) within the station cards in the correlator to adjust the antenna phases by exactly the computed values so as to bring all signals into phase with those of the reference antenna. Once these adjustments have been made, a “residual” phase correction typically needs to be made on the following correlator scan, and indeed, such corrections continue to be made every scan until the end of the recording. The phasing loop is closed in the sense that small errors in the phase adjustments may be corrected on subsequent scans. This is the so-called “slow” loop which is illustrated schematically in Figure 8.7.

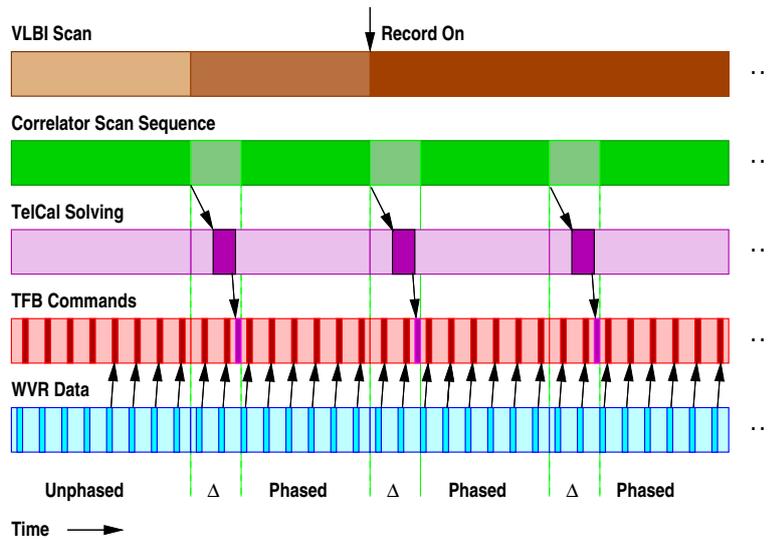


Figure 8.7: VLBI Scans in the Phasing System. Each VLBI scan is partitioned into “subscans” for correlation and “slow” timescale processing (seconds) in TelCal. WVR adjustments are made in the CDP on a “fast” timescale (every second). See text for full discussion and explanation.

While the correlator scans are shown as time-contiguous in Figure 8.7 (green bar), there are in fact small gaps to allow the data to be dumped out for subsequent correlator processing. Even so, the solutions from TelCal (purple) do not arrive at the TFBs (red) until after the correlator scan has started—the time of application is noted and the early part of the scan is excluded from the subsequent phase calculation. Usually one phase application is needed to get an acceptable solution, so the recording should start no sooner than the start of the third correlator subscan. These parameters are all programmable, and judicious choices for them are made when the schedule block is created. For Cycle 4, the subscans are likely to be 16-s long with 4-s (channel average) integrations; there are then about 30s of observations prior to the start of the recording. Note that the ALMA correlator does its processing by subscans—each subscan must be dumped and processed by the correlator and ultimately archived—with a gap of several seconds between subscans. On the other hand, the signals from the antennas flow continuously to the summation logic and recorders, so there are no comparable gaps in the VLBI data.

In addition to the slow loop, there is also a “fast” loop that may be enabled and which uses the WVR data from each antenna to correct the antenna phases in response to the wet component of the atmosphere. These corrections are performed at a frequency of ~ 1 Hz (blue in Figure 8.7). This loop is open (corrections are continuously made, but feedback is only achieved through the slow loop). Note also that all WVR corrections are made online in the ALMA correlator, but obviously after the WVR data was collected, so the phasing corrections are at their best in stable atmospheres, and deteriorate with increasing atmospheric instability. There is no limit to the duration on the operation of these loops (other than general ALMA ones with regard to total data collected).

VLBI Phasing System

Moving beyond the timing of the phasing system, there are several other aspects of the system that an ALMA VLBI observer should be familiar with. All ALMA observations operate with an array of antennas (*i.e.* those available to a scheduled observation). Of these, a portion are controlled by the phasing system to form the coherent sum signal which is used for the VLBI recordings. The previously mentioned reference antenna is merely a designated member of this “phased array”. At least two antennas from the active array (designated “comparison antennas”) are held outside the phased array. These comparison antennas are thus available for diagnosing the performance and efficiency of the phased array system. Finally, after construction the summed signal is decimated to two bits per sample (as is the case with all of the ALMA signals in the VOM), so it

is possible to have the ALMA correlator correlate the sum signal during online processing. This is done by co-opting some antenna that is not used by the VOM and placing the sum signals into the correlator as if it had come from this “sum antenna”. The sum signal therefore appears as antenna “APP001” in any ASDM file containing ALMA data acquired with the VOM. There are a few things to note about the data and metadata for this antenna:

- it has no useful metadata, as it is not a physical antenna
- it has no WVR data, for the same reason
- it is **highly** correlated with other antennas in the phased array since the non-physical receiver noise is present in both

The selections of antennas for the phased array and comparison array are made when the scheduling block is created, although these may be adjusted if necessary during the observation. Normally, the antennas to be phased are all chosen to lie within a certain radius from the array center. (The efficiency of the phasing system decreases on longer baselines, so in practice, maximum baseline lengths in the phased array will typically be $\lesssim 1$ km). The reference antenna is in general selected to be centrally located, and typically at least one of the comparison antennas is chosen to be within the radius of the phased array. An additional detail is that because of the decimation to two-bits, the number of phased antennas must always be odd.

A final important note concerns delays between antennas. The ALMA control system estimates the total delays between all antennas and adjusts the signals through a variety of techniques (in hardware and software) so that the visibilities found in the data set have essentially zero delay on each baseline. (*I.e.* there is no slope in a plot of visibility phase as a function of with frequency on any baseline.) Unfortunately, a (significant) component of the delay is removed with a frequency dependent phase rotation in the online correlation processing, not in the hardware, so there is some significant delay present in the signals at the TFBs where the phase corrections are made and correspondingly in the logic which creates the “sum antenna” signal. Thus in order for the phasing system to work properly, it is necessary to turn off this component of software delay correction and to instead correct it in the phasing calculations. To do this, the observing band is subdivided into a relatively large number of channel averages (at least 8) and the normal phase-solving procedure produces different, independently phases that can be applied to the TFBs corresponding to each channel average. This does remove most of the delay, but it does leave a small residual of phase-slope within each channel average. The only restriction is that the source needs to be bright enough in each of the channels so that a usable phase solution can be found. The flux limit of 0.5 Jy in Cycle 4 ensures that this is the case.

The activities of the phasing loop are fully documented in an ASDM table as described in Section 8.11.2.

VLBI Observing Mode Analysis

As pointed out in Section 8.11, any VOM observation actually involves two concurrent interferometric observations: one on the ALMA scale that results in correlated data products between all participating ALMA antennas and one on a global scale, where ALMA serves as a station in a global VLBI array. A description of the reduction and analysis of ALMA VLBI data is largely beyond the scope of this document¹³. This process is, however, informed by the analysis of the ALMA-scale observation with regard to observatory performance (SEFD, system temperature, *etc.*), which for the most part may be obtained in the usual way. However, since the VLBI signal that is recorded is that of the “sum” antenna whose properties in turn depend on the phasing performance, we should point out a few details. The most significant point is that the effective area of the “ALMA dish” is roughly proportional to the square root of the number of phased antennas. (In Cycle 4 only 12-m dishes are to be used, and they have similar effective areas. If the 7-m dishes were to be included the relation is more complicated.) In practice, a number of effects conspire to lower performance from this ideal relation¹⁴ to about 90% of what would be expected in a perfect system. A second point here is that the sum antenna signal is decimated to two bits from the individual antenna signals which were also decimated to two bits. Thus there is an additional reduction in sensitivity by 0.88 that must be taken into account in the VLBI

¹³A VLBI analysis handbook is in preparation.

¹⁴ALMA-05.11.63.03-0001-A-REP.pdf

analysis. In any case, the net effective area of the ALMA “sum” dish may be calculated from the ASDM on a per-scan basis.

The performance of the phasing system is captured in the metadata in ASDM `CalAppPhase` tables. There is one table entry per scan per baseband.

In each table, there are temporal bounds on validity (*i.e.* the range of time within the scan of stable phase) and entries to list the disposition of the antennas amongst the various categories (*i.e.* phased antennas, reference antenna, and comparison antennas) and whether this represents a change from the previous scan. The table primarily reports a number of phase values, N_v , which is dependent on the “packing mode”. In Cycle 4, there is one value per channel average per polarization per antenna. Additionally, the `CalAppPhase` table reports lower-level details of the phasing system, including online per-antenna assessments from TelCal of the antennas drawn from the quality of the phasing solution. This table may be used in conjunction with the ALMA data for a particular scan to calculate the phasing efficiency and thus to make an appropriate correction for the absolute VLBI correlation amplitudes.

As indicated in Section 8.11.1, the correlation products are delivered to the PI for analysis. Traditionally for the GMVA this has been in the form of FITS-IDI files. A complication for work with ALMA data is that the polarization conversion is performed with a tool (`PolConvert`) that requires input from the CASA analysis of the ALMA dataset. This conversion can be performed on the FITS-IDI file or at the correlator using the raw correlator output if the the CASA analysis is available. Arrangements for doing this will be worked out between the PI of each project and the supporting observatory collaborations (GMVA and EHTC) following the observation.

8.12 High Frequency Observing

High-Frequency Observations (specifically Bands 8, 9, and 10) require some special observing techniques, mostly because of the combined ill-effects of lower atmospheric transmission and lack of calibrators. These techniques, which are just addenda to standard observing modes, will be progressively implemented in the ALMA observations for Cycles 4 and beyond, after some extensive testing is carried out. It is expected that all of them will improve not only the quality of the final data products, but also the overall observing efficiency at high frequencies. For reference, we add a brief description of the nomenclature that would be used by Contact Scientists and the ALMA project in general whenever they are discussed, reported or communicated to PIs:

- **High Frequency Cone Searches and Calibrator Surveys:** Observations with the high frequency ALMA bands require that the availability of calibrators near the target source are checked near the time of the planned observations. Searches for phase calibrators and check sources (see Chapter 10 for details) within 10 degrees of the science targets in each SB will be carried out by ALMA staff within 90 days of the high frequency observations. The results of those measurements will be immediately ingested into the ALMA Calibrator Catalogue, so that they are available for queries at the time of the actual observations.
- **Band-to-Band Transfer (B2B):** Whenever suitable calibrators at the observing frequency are unavailable, the B2B technique will be used. B2B consists in phase calibrator observations at a lower frequency, where many more calibrators are available, straddling the science target observations. In order to properly scale the phases between the low-frequency band and the high-frequency band observations, a separate source is observed in both at some point during the SB execution. Such source does not have the requirement of being close to the target, and therefore the probability of finding one at high frequencies is significantly higher. For more details, please read Section 10.6 of this Handbook.
- **Bandwidth Switching (BWSW):** For high frequency observations (Bands 8,9,10) with spectral set-ups including only narrow band spectral windows, the PI will be asked to add an additional wide band spectral window. With this, it will be possible to transfer the phase solutions of the wide-band spectral window to the rest. Whenever adding an spectral window is not possible (*i.e.*, projects that have occupied all the spectral windows with narrow-band set-ups), there is also the option to use the Differential Gain Calibrator approach, which will add some observations with a different spectral set-up (*i.e.*, wider spectral windows) in the normal cycle of observations. For more details, please read Chapter 10.

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

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

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

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

¹<http://almascience.org/>

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

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

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

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

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

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

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

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

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

The octiles characterize the amount of PWV that can be expected at the ALMA site i.e. a value of PWV *at least as good* as the first octile value can be expected 12.5 per cent of the time, a value at least as good as the second octile 25 per cent of the time, and so on. The octiles corresponding to the ALMA site (determined from many years of monitoring) are shown in Table 9.1.1.

When estimating the time for a project, the OT will always select a PWV octile that is appropriate to the frequency being observed. It does this by calculating the time required for each octile and then choosing (and reporting) the highest (worst) octile for which the increase in time relative to the first is less than 100 per cent. A consequence of this definition is that the octile also depends on source declination i.e. sources at low elevations will require better weather conditions. The resulting curve of octile versus frequency is shown in Figure 9.1, for a source declination of zero degrees. A user can override this choice in the GUI 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 ALMA Bands 4, 5 and 10, the calculator currently only uses the specifications for the receiver temperatures (required over 80% of the receiver bandwidth) and not the actual measured values. For ALMA Bands 3, 6, 7, 8 and 9, however, typical values measured in the laboratory are used as these are usually significantly better

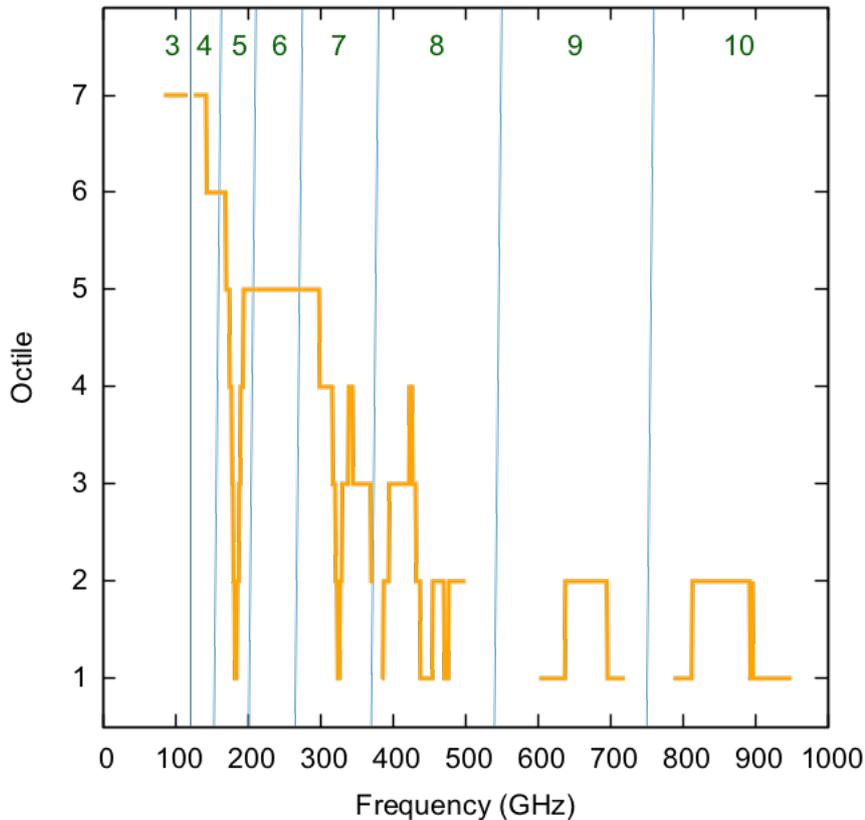


Figure 9.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown at the top of the plot. The water line at 183 GHz (Band 5) is particularly prominent. In general, higher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

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.

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 9.4.

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

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).

to Equation 9.4 and thus its total contribution is frequency dependent and can vary between the different sidebands.

9.1.5 DSB receivers

Due to the way that the astronomical signals are down-converted to an intermediate frequency, every heterodyne radio receiver simultaneously detects radiation from two sidebands. This means that, if nothing were done to prevent it, a 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 tuning software will always place spectral windows 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. Where possible, this situation can be avoided by defining dummy spectral windows such that the line of interest is forced into the other sideband (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:

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

where

- T_{rx} – receiver temperature

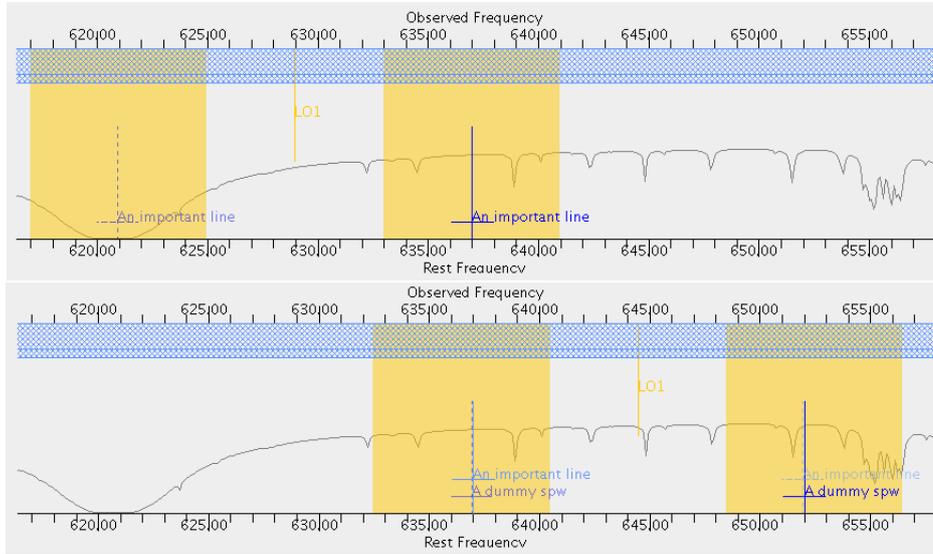


Figure 9.2: A Band-9 spectral setup as displayed in the spectral visual editor in the OT. The top figure shows an example of how the image equivalent of a single spectral window can fall into an area of poor atmospheric transmission, leading to much higher T_{sys} than necessary. Placing a dummy spectral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

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

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

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

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

9.1.7 System Temperature (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 its contribution to the system temperature. In this case, the same equation is used for

all receiver types and the DSB noise contribution is simply double the single sideband case. This is controlled via the sideband gain ratio, g :

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

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

9.2 The Sensitivity Calculation

Once T_{sys} has been determined it is possible to calculate the point-source sensitivity given a requested amount of on-source observing time or vice versa. At no point is any account made for the expected level of loss in sensitivity due to 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:

$$\sigma_{\text{S}} = \frac{w_r 2 k T_{\text{sys}}}{\eta_{\text{q}} \eta_{\text{c}} A_{\text{eff}} (1 - f_s) \sqrt{N(N-1)} n_{\text{p}} \Delta\nu t_{\text{int}}}. \quad (9.8)$$

The various parameters are

- w_r – robust weighting factor. Pipeline imaging and subsequent QA2 assessment is performed assuming that the visibilities are weighted using robust weighting, specifically a Briggs robustness factor of 0.5. Simulations have shown that this factor is equal to 1.1
- A_{eff} – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency where $\eta_{\text{ap}} = R_0 \exp(-16 \pi^2 \sigma^2 / \lambda^2)$ and σ is the rms surface accuracy of the antenna. The latter is set to the design goal of 25 μm and 20 μm for the 12- and 7-m antennas respectively³. R_0 is the product of a number of different efficiencies and is equal to 0.72. See Table 9.3 for values of antenna efficiencies and effective areas in various ALMA bands
- f_s – shadowing fraction. For the more compact 12-m configurations and the ACA 7-m array, antennas can block the field-of-view of other antennas in the array and thus reduce the total collecting area. The shadowing fraction is a function of source declination as shown in Fig. 7.5.
- η_{q} – quantization efficiency. A fundamental limit on the achievable sensitivity is set by the initial 3-bit digitization of the baseband signals. This is equal to 0.96
- η_{c} – correlator efficiency. This depends on the correlator (64-input or ACA) and correlator mode, although the efficiency of all 64-input correlator modes is equal to 0.88. The ACA efficiencies *do* depend on the mode, but this is only taken into account in the OT, not by the Java applet (see below) which therefore also assumes a value of 0.88
- N – number of antennas. This defaults to 40 for the 12-m and ten 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

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 η_{ap} 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.

- t_{int} – total on-source integration time.

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

$$\sigma_{\text{T}} = \frac{\sigma_{\text{S}} \lambda^2}{2k \Omega} \quad (9.9)$$

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

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

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

9.2.2 Total Power Array

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

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

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

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

9.3 User Interface

The main way that a user interacts with the Calculator is through a GUI in the OT or via 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. It is also used to calculate the amount of shadowing that is likely to affect the ACA 7-m and the smaller 12-m configurations.

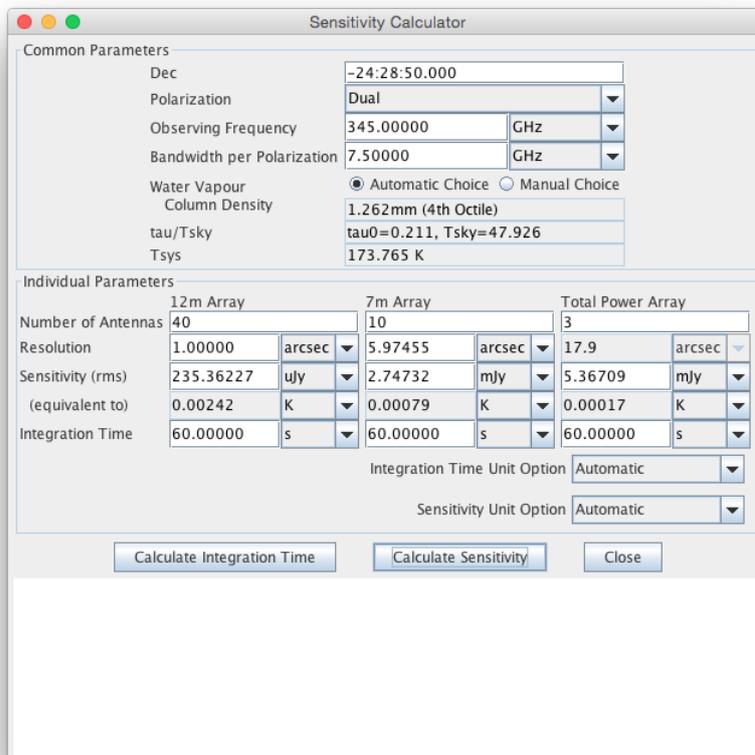


Figure 9.3: Screenshot of the GUI version of the ALMA Sensitivity Calculator as implemented in the ALMA Observing Tool. The white area at the bottom is for displaying error messages i.e. parameters out of bounds. The example here shows the achievable sensitivity for all three arrays for an on-source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

- 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.
- Number of antennas – the ASC currently defaults to the values for Early Science, namely 40 from the 12-m Array, ten from the 7-m Array and three from the TP Array.
- Angular resolution – this affects the time estimates when sensitivities are specified in temperature units. The calculator will not perform any calculations when Kelvins have been specified, unless a non-zero value for angular resolution has been entered. The calculator will also issue a warning if the angular resolution falls outside of the range corresponding to 125-m and 1-km baselines.

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 phase difference 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). and are in the form of the ALMA Science Data Model (ASDM). Details of the flow of the astronomical data from the antennas through the correlator to the archive are given in Chapter 14.

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*, associated with an assumed position, σ in the sky, is included in correlator processing in order to remove the approximate a priori signal delays that are a function of baseline length, observing frequency (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 this correlator model removed from the correlation of the input signals, the visibility function can then be averaged coherently over many seconds and a large bandwidth.

10.1.2 The Visibility Function Dependence

The fundamental observational segment is the *scan*. It is an observation with one field name, associated with a phase center σ (which may be moving), and one frequency band, ν_b . The observed visibility function, \mathcal{V}^o , associated with a scan/band is

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

where:

i, j is baseline between antenna i and j .

t_k is each time-stamp in the scan.

$\nu_{s,c}$ is the frequency sampling within the band.

where s is a spectral window (SPW),

and c is the channelization of each SPW (from 1 to 4096).

p_q represents up to a maximum of four polarizations products.

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 *autocorrelation* data, where $i = j$ is also contained in the data set. 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 Phase Referencing

Nearly all ALMA interferometric observations use *phase referencing* as the major calibration method. This method interleaves scans between the science target and a nearby *calibrator* source that is usually a quasar of known position, spectrum and flux density.

10.2.2 Antenna-based Calibrations

Most of the changes to the true visibility function are associated with the signal flow for a given antenna from the troposphere above that antenna, through all of the electronic processing along the signal path, up to the correlator. Hence, the calibration equation that describes the changes can be written as:

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

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 antenna-based calibration equation. The second line contains the effect of noise and any additional baseline-calibration, and these will be ignored for the most part. The polarization dependence and more complicated frequency separation are not included, but discussed below.

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

$$G_i(t_k, \nu) = a_i(t_k, \nu_c) e^{i\psi(t_k, \nu_c)} \quad (10.3)$$

where $a_i(t_k, \nu_c)$ 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.

10.2.3 Major Calibrations Affecting Interferometric Data

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 is as follows: ¹

I = Ionosphere: Opposite rotation of linear polarization position angle caused by propagation through the ionosphere—not significant at the current 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 (pointing) position, caused by focus and pointing offset.

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

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

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

d = Polarization leakage: The non-orthogonality and small mis-alignments of the two orthogonal polarized feeds causing slight mixing of polarizations.

T = System Temperature: A measure of the noise fluctuations from the entire system, dominated by the receiver system at low frequencies, and the atmosphere noise at high frequencies

M = Delay model used in the correlator to compensate for the best-estimate temporal delay changes at each antenna.

D = Instrumental delay offset of the signals between antennas.

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

G = Temporal gain: The relative amplitude and phase of an antenna response as a function of time.

F = Flagging: Editing of data found during the calibration process.

Do not be intimidated by the above list. The main ALMA calibrations of interest to most users deal with the temporal gain G and the bandpass response B . The calibration process is significantly simplified because the temporal dependence (GT) and frequency dependence (BF) for nearly all instrumental changes are only *lightly coupled*, so their variations can be determined independently, or at least interactively.

$$G_i(t_k, \nu_c) \approx GT(t_k) * BF(\nu_c) \quad (10.4)$$

10.2.4 The Organization of the Calibrations

There are three major 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, and antenna positions. They are measured periodically by the ALMA staff using methodologies that best suit the calibration. These calibrations are briefly outlined below.

Online parameters and 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 thermal antenna pointing offsets; measurement of $T_{\text{sys}}/T_{\text{RX}}$ using the Amplitude Calibration Devices (ACDs); measurement of the water vapor emission of each antenna toward the source in order to determine the wet path delay variation along the line-of-sight; measurement of the variable delay using a Line-Length Correction system (LLC).

Offline calibrations. The offline calibrations are those determined after the completion of the experiment. These calibrations are often associated with 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 . Another important calibration is the conversion of the correlated visibility amplitude to Jansky flux density units.

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 pointing direction is called *primary beam pattern*. This sensitivity correction is needed for large angular source emission, and for 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 over an area that extends well outside of the main primary beam area. The measurements are processed in special-purpose software that determine the accurate antenna surface topology, 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 software packages and applied during the image-make step (see Chapters 3 and 7 for more details). There are four different antenna-types (12-m European, Japanese, American, 7-m Japanese) that have significantly different primary beam shapes.

Pointing Models (P). All-sky 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, most of which are at the arc-second level. The pointing differences among the receiver bands require only an elevation and azimuth collimation offset with respect to that of a reference frequency (usually Band 3 or 6).

The complete pointing models are checked and adjusted on weekly timescales, primarily by making "all-sky pointing runs", which are observations of about 50-100 interferometric pointing measurements on a well-distributed set of bright quasars of accurately known position. The total observing time is about one hour and each pointing measurement scan consists of a measurement of the quasar amplitude at five subscans at different antenna pointing offsets (depending on frequency and antenna diameter) from the assumed pointing direction. From these data, the pointing offset from the assumed position for each antenna are determined to an accuracy of around $0.1''$.

The pointing models are stored in the Telescope Monitor and Configuration Data Base (TMCDB) which includes a history of all changes. The results of additional pointing measurements made over many weeks during spot checks, are also used to update the pointing models when appropriate. Currently only night time pointing measurements are considered in these *static* models. During daytime observations, the pointing changes dynamically and the modeling of them is continuing.

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 by fitting total power data recorded while the antenna performs azimuth strokes across a planet or bright quasar. After these major adjustments, the pointing model is refined as described in the previous paragraph.

In all science executions additional short pointing measurements are made as needed using a bright quasar near each of the science targets and calibrators. The pointing observations are made at the target observing frequency up to Band 8, and in Band 8 for higher frequencies. The choice of pointing tuning will be made by the observatory in order to maximize available array elements by avoiding bands or frequencies affected by hardware problems, and minimize pointing error. This 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 to determine a six parameter model, plus offsets for each receiver band relative to the reference band (Band 7) at elevation 50° and temperature 0°C. The models contain an offset for each axis which describes the optimum subreflector position for Band 7 ("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, and are impacted slightly to hardware changes. Focus measurements are traditionally made in a method similar to that for interferometric pointing measurements, except the subreflector is offset (in one axis per scan) relative to the nominal model position. The subreflector position along the axis is changed and that giving the maximum cross-correlation signal is determined for each antenna. During Cycle-4 a faster technique in which the three-axis focus offsets are derived from holographic beam maps (which do not require the subreflector to move), may be used. 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.

For frequencies less than 500 GHz during the night, the optimum focus position across the focal plane, X and Y positions, is described accurately by the temperature-dependent antenna models and are rarely checked with observations. Generally when switching to a new receiver band it is a useful cross-check to make a focus calibration in the Z axis (in-out direction), and can be done during the pointing and delay offset measurements as by-products. For Bands 9 and 10 where focus accuracy is more critical, the observations require use of the brightest quasar, and for the atmospheric conditions to be excellent. At night, errors below 20 microns are achievable.

During the daytime the focus can deviate from the models by several hundred microns, so focus calibration observations are generally executed between every science execution above 240 GHz in order to track these variations, but the poor atmospheric stability produces focus measurement uncertainties. Generally, focus errors can be larger than 100 microns in daytime, so that observation during the day at Bands 8, 9 and 10 need frequent focal checking. Further daytime modeling and observing strategies are in progress. The focus model used during an observation is stored in a `FocusModel` table in the ASDM datasets. As for pointing, modeling and correction of focus during daytime is continuing in progress.

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 antenna/SPW/pol signal path. These are made with short interferometric observations of a strong quasar at the relevant tunings. If the delay offset is more than about 50 ps, producing a phase change over a 2 GHz bandwidth of 36°, the delay offset in the TMCDB entry will be adjusted. Delay measurement offsets are repeatable to better than 10 ps, but they can vary by a few tens of ps from a variety of effects. Hence, during calibration and science observations, the online delay errors are usually below 50 ps, so the use channel-average data over the 2-GHz bandwidth will have little decorrelation due to the delay errors, and they can be used for continuum analysis, both in PI observations and for or pointing, focus and beam measurement calibrations without the need for a bandpass calibration.

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, with a goal of about 50 micron accuracy. The antenna positions are crucial for the determination of the accurate correlator delay model, M , and the derived phase of an observation.

The baseline runs (as they are called) are executed by the ALMA staff about once per week during good weather conditions, but also when one or more antennas have changed pads. The antenna positions are determined from about 50 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 coordinate offsets of the antenna position from that referenced in the TCMDB that is used in the correlator model.

The typical a priori position offset for an antenna that has been placed on a pad is well-above the antenna location specification of 0.05 mm, so baseline observation 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 measured and updated in the TCMBD for up to a week. Thus, the correlator delay model will be incorrect and the visibility phase for data associated with this antenna will have significant phase variations with time and between sources. There are offline procedures in the data reduction scripts (see below) to correct for the difference between the antenna position used in the TCMBD and the more accurate values measured after an accurate baseline run. For the longest ALMA baselines, it is difficult to measure the antenna position accuracy better than 1 mm because of the effect of the variable atmospheric refraction over the array.

10.4 Online Calibrations

Water Vapor Radiometer Measurements (M). Fluctuations in the line-of-sight precipitable water vapor (pww) 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 from 5 m to many km. All ALMA 12-m antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures the emission of the atmospheric water line along the same path of the other ALMA bands, in four spectral frequencies near 183 GHz with a sampling of about 1 Hz. These radiometers are described in Section 5.4. The 7-m (CM) antennas do not have wvr receivers. There have been tests to use the wvr signals from the four PM antennas that surround the CM array, and interpolating the effective wvr correction for each CM antenna, but are believed to be unnecessary for this short-baseline data.

The conversion of the water vapor emission into appropriate delay changes (M) is calculated using two methods and will be discussed in the following section.

System and Receiver Temperature (T). At millimeter and submillimeter wavelengths, the additional random noise associated with a measurement has two major origins: the noise associated with the receiving system, T_{RX} , and the noise associated with the atmospheric transmission, T_{sky} which acts like a black body emitter, as well as attenuating the astronomical signals. These noise contributions, P , are usually expressed as a temperature, T , where $P = kT$ where k is the Boltzmann's constant. This sky noise 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 (Band 3 receiver characteristic).

To measure these effective temperatures, the ALMA front ends are equipped with an Amplitude Calibration Device (ACD) which can move a hot and ambient blackbody load in front of each receiver feed. The description of the system temperature electronics is given in Chapter 4. The noise measurement from these two loads give the total system temperature, T_{sys} and the receiver temperature, T_{RX} for each spectral channel. The combination of these measured temperatures to determine their effect on the observation sensitivity is given in Chapter 9, equation 1.7. The accuracy of these measurements depend on the 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 polarization products) of the Baseline correlator²³. The ACA correlator uses the same spectral configuration (resolution) for temperature calibration⁴.

²In TDM mode the autocorrelation zero lags can be used to fully compensate for quantization effects on the autocorrelations. However, in FDM mode which adds a filter bank before the correlations are produced an external measure of the total digitized 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 excessively 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-4 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.

The temperature measurement include the receiver-weighted contributions from both sidebands, whereas in the cross-correlation data, LO switching sequences are used to highly suppress the image sideband cross-correlation. The measured T_{sys} and T_{RX} must be scaled by a sideband gain factor to produce single sideband temperatures which can be directly applied to the visibility data to put them in temperature units. For cycle 3 scans of strong sources were included in every science execution to measure baseband-averaged sideband gains to correct the system temperatures, but these measurements are not accurate except under very ideal observing conditions. For cycle 4 the nominal value of 0.99 for single sideband frequency bands and 0.5 for double side-band frequency bands will be used to convert the system temperature.

In principle, every scan could have a system temperature measurement. At frequencies below about 400 GHz where the system temperatures are relatively constant over 10 min or over 10° in the sky, measurements can be limited in time and among sources. At higher frequencies, each of the sources should have a measurement to remove systematic amplitude offsets between them. The specifications for T_{sys} calibrations are to reach 1% correction accuracy.

The online measurements of T_{sys} and T_{RX} are stored in the `CalAtmosphere` table in the ASDM datasets, and are displayed on-line to show measurement problems immediately. The system temperature measurements can also be re-derived offline from the original hot and cold load data if there were on-line conversion problems or if a modified processing algorithm is used. The T_{sys} spectra, suitably interpolated between time and sources, 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 discussed in Pointing Models, corrections are made during observations in order to reach ALMA's $0.6''$ rms pointing accuracy specification. To some extent these additional 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 automatically included in the observation schedules, with no PI input. At present a pointing calibration is triggered when moving to a target more than 20 degrees from the previous pointing calibration.

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 on line. 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 first step in the calibration 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). An inspection of the incoming data at the ALMA site, called *QA0*, is made as the data are collected to check the data quality, and is discussed in Chapter 11. Each experiment is then calibrated using the ALMA Pipeline for *standard* experiments (Chapter 13), or more flexible processing is available for *non-standard* experiments.

On-Line 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 archive data base in two ways, xml and binary flags. One of the initial offline calibration steps is to apply these flags.

Calibration-based Flagging (F). During the off-line calibrations, described below, further times and antenna of data flagging are often found. Examples are, noisy edge channels in the spectral windows, low relative signal in one antenna compared with the others, phase jumps. In many cases, the characteristics

of these bad data can only be seen in the calibrator observations; hence, most flags associated with a calibration scan should be appropriately extended to other sources scans, SPW and baselines.

Bandpass Calibration (B). The response across each SPW 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, this residual 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 SPW's. After removing any significant phase variations over the bandpass scan, 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 in Appendix C. The bandpass with time and with angle is extremely stable⁵ 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.

Gain Calibration (G). Virtually all ALMA observations use the phase referencing technique to calibrate the temporal antenna-based variations in amplitude and phase in order to calibrate the target source visibility data. The larger temporal changes 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 observations can be used to align the phases of all data streams that can then be summed before determining a solution. The combined SPW/pol (often 8 independent data streams) antenna solutions for the phase calibrator is then more sensitive than that for each SPW separately.

The plots of the antenna-based amplitudes and phases 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: for example, is the T_{sys} correction anomalous, are there comments in the observing logs? (2) The phase difference between adjacent calibration scans should be relatively small. For example, if the antenna-based phase change between scans exceeds about 60° , interpolation to the intervening target scan is uncertain. This occurs at longer baselines and high frequencies, and flagging of the target data during these intervals may be useful to improve the image quality. The execution of the *Go/Nogo* observations (short observations of a strong source made periodically through the observation period) to measure the short-term phase variations are used to avoid periods of poor stability.

Absolute Flux Density Calibration(T). The Calibration of the data using the T_{sys} measurements converts the raw correlation amplitude into Kelvin. The relationship of the correlation units in Kelvin to Jy is called the *system equivalent flux density (SEFD)*, and depends on the antenna diameter and efficiency that depends on surface accuracy and elevation dependencies. The SEFD is very stable with time and similar among the antennas of the same type. For example, the 12-m antennas at Bands 3, 6 and 7 have an SEFD of about 35 K/Jy. This a priori conversion from Kelvin to Jy is accurate 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 usually included in the observations. Perhaps, the most reliable Kelvin to Jy conversion is obtained observing solar system objects of relatively small angular size that have accurate visibility models versus frequency and baseline length (see Appendix C for the current list used). Also about 30 bright quasars (grid sources) at Bands 3 and 7 are monitored by ALMA every 10 days, 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). The polarization calibration of the XY and YX correlations is made on a bright quasar that is polarized by at least 3%. The SB observing length needed for the calibration is about three hours during which the calibrator rotates at least 60° in parallactic angle. Two calibration steps are needed. First, after the XX and YY data are calibrated (bandpass and gaincal), the XY and

⁵http://almascience.org/documents-and-tools/alma-technical-notes/ALMATechnicalNotes15_FINAL.pdf

YX frequency dependence (bandpass and phase offset) for each SPW will not be zero, but reflect the X- and 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. After these array phase-frequency terms are then removed, the second calibration step determines the *leakage* between the X and Y polarizations (called D-terms) for each antenna and SPW. This is a part of the calibration that requires the parallactic angle change of at least 60° during an experiment. There are two leakage terms: one related to the non-orthogonality between the X and Y feeds and one term related to the ellipticity (X leakage into Y, and vice-versa) of each feed. Both of the above calibration results are determined from the polarization calibrator observations alone, but then applied to the entire data set in order to calibrate the polarization properties for all of the other sources. Another polarization term is related to the average orientation of the feeds (which is known), but can be measured using a strongly polarized source of known polarization position.

Investigations are in progress to decrease the three-hour time needed for polarization calibration, regardless of the time needed for the target integration. One method is to use a quasar of known polarization (many sources must be monitored) in which case only a short observation is needed to determine the D-terms. Another method is to use the D-term results from previous experiments at the same frequencies since they do not vary significantly with time. But, these methods will not be available in Cycle 4.

Because the XY, YX response over the primary beam is considerably more complicated than that for the XX and YY polarizations, only sources with a limited angular extent can be accurately imaged. It is expected that in cycle 4, the linear polarized imaging will be sufficiently accurate to the the 30% of the power pattern. For later cycles, full primary beam and mosaicing may be offered.

Circular polarization data is included in the four Stokes parameters that are measured. However, the calibration accuracy for circular polarization is still uncertain, and the variation of the circular polarization properties beyond the 80% sensitivity region of the primary beam is significant at the few percent level. Hence, circular polarization results are not supported. Please contact the ALMA help desk or your local ARC for more information about circular polarization results from your data.

10.6 Special Calibrations

10.6.1 Phase Referencing, especially above 400 GHz

The quality of a phase-referenced image improves the closer the calibrator is to the target because the systematic short- and long-term phase differences between the target and calibrator depend roughly on their separation. At frequencies lower than 400 GHz, calibrators can often be found within 5° of any target, and images of dynamic ranges (peak to rms), if the source is strong enough, can be obtained above 200 with simple phase referencing. At higher frequencies, calibrators within 10° of a random target can be found infrequently for a 12-m array observations and seldom for a 7-m Array observations, and phase-referenced images at these frequencies are of poorer quality. In Appendix C, the phase calibrator minimum flux densities needed for detection are given in Table C.2. The probability of finding a calibrator within a specific distance from a target is given in Figure C.2 for different frequencies and arrays.

The ALMA catalog (see Appendix C) contains over ten thousand sources, mostly found in cm-wave radio surveys. In order to determine which of these calibrators are sufficiently strong at ALMA bands, special-purpose ALMA observations, called *cone searches*, measure their flux density in bands at 100 GHz and 300 GHz. Often the observing list contains candidate calibrators that are within about ten degrees from upcoming science targets. Most calibrators are quasars that have a radio spectral index, α , in the mm-wave region between -0.5 and -0.8; hence, their extrapolation to 900 GHz from the measurements at 100 and 300 are accurate to about 30%. Also, quasars are variable. If the most recent measurement is more than 90 days from the approximate experiment date, then a current flux density will be remeasured in an upcoming cone search experiment. The minimum flux density of a calibrator that can be detected depends on the frequency and total bandwidth of the experiment.

Two methods are available to use calibrators that are as close as possible to a target:

SPW/Pol averaging. The temporal amplitude and phase variations over a period of several hours amongst the spectral windows and two polarization are small (1% and 2°) for all ALMA bands. This amplitude scale and phase offset between streams can be accurately determined from the observations of the bandpass calibrator scan, and is then applied to the entire experiment of a few hours. The Gain Calibration (G), described in the previous section, can be then be measured for this averaged data stream, with a typical increase in the antenna-solution SNR per time by a factor 2.8 (from 8 original data streams). This allows finding detectable calibrators closer to a target. This technique is called *low SNR calibration*.

Band-to-Band Calibration. For phase referencing calibration at the higher frequencies (> 400 GHz), finding a calibrator closer than 5° to the target is relatively rare, whereas a closer calibrator at lower frequencies will be available, and potentially improve the phase referencing accuracy. Thus, the technique of band-to-band phase transfer using observations of a calibrator in a lower ALMA Band (3, 6 or 7), and then scaling the antenna-based phase solutions (G) to the higher frequency band (8, 9, 10) of the target has been tested.

This inter-frequency calibration is possible because of the excellent instrumental phase stability of ALMA. Once the phase difference is measured between two ALMA data streams (different SPW, pol and bands), nearly all remaining phase change are non-dispersive (delay-like) and scale with frequency. The major delay changes are those in the troposphere and changes in the path length in the antenna system, including those associated with antenna position errors.

The band-to-band phase scaling requires three short scans to be made at any time during the experiment, t_0 : An observation of a bright source, called the diff-gain-source (DGC), at the low frequency, an observation of the DGC at the high frequencies, plus an observation of the low frequency phase calibrator (LF) at the low frequency. The phase difference between the DGC observations gives its instrumental phase difference at t_0 . The phase difference between the DGC low frequency and LF low frequency gives the correction to this phase difference for the LF. After removal of this LF phase difference, any further phase changes over time will scale precisely with frequency.

Since ALMA can observe only one frequency and source at a time, the needed phase differences are estimated by alternating observations of the DGC at the two frequencies and the LF at the low frequencies. The temporal phase variation in the troposphere thus introduce errors in the phase difference which add significant errors in the band-to-band phase scaling from low to high frequencies.

Experimentation with the Band-to-band technique is continuing. However, with the success of the cone searches in finding high frequency calibrators, usually with 10° of a target, and the phase offset uncertainty of the band-to-band calibrations, band-to-band will not be offered in Cycle 4.

10.6.2 Phase Calibration with Narrow Bandwidths

When observing a target within narrow SPW bandwidths, a suitable phase calibrator, especially at the higher frequencies, may not be detectable with a 2-min calibrator scan. Two methods of calibrating these narrow bandwidths are available:

One wide SPW: Low-SNR calibration. If the aggregate bandwidth in the combined SPW's in the experiment has a bandwidth > 1 GHz (usually because one of the SPW is wide-band), then there will be sufficient sensitivity to detect in a 2-minute scan to obtain antenna-based solution of a phase calibrator that is reasonably close to the target. In this case, only one spectral set-up is needed for the entire experiment.

The calibration method is identical to that described in the SPW/Pol averaging section where: 1) The amplitude ratio and phase offsets between the SPW's are determined from the observation of the bandpass calibrator scan; 2) The phase calibrator phase is determined from the one wideband SPW (or the aggregate) and then transferred to the narrow band SPW's. An alternative observing strategy, where all of the target spectral windows are narrow-band, is described next. But, this requires constant switching between the narrow-band and wide band configurations over the experiment.

Band-Width-Switching. If all of the SPW's are narrow bandwidth, it is likely that a suitably nearby phase calibrator cannot be detected within a 2-min scan. The strategy for calibrating this set-up is to observe the phase calibrator with wide-band spectral windows and the target in the desired narrow bandwidths. This requires the switching of the spectral setup between scans.

The amplitude ratio and phase difference between the wide band SPW and its associated narrow band SPW is measured by observing a strong calibrator (DGC) in three consecutive scans: wide band, narrow band, wide band. The calibrations of two wide band scans can be interpolated to the time of the narrow band scan, removing most of the temporal variations. The amplitude and phase difference between the wide and narrow bands can then be determined.

The advantage of this strategy is that all of the target SPW's can be used for narrow-band spectroscopy, whereas the previous strategy must use one wide band SPW for calibration purposes, and is simpler to schedule and easier to calibrate.

10.6.3 Astrometry

The absolute positional accuracy of an ALMA image made with normal phase referencing depends on the resolution and the quality of the phase calibration. In general, the closer the calibrator is to the target, the better will be the image registration in the sky. With reasonable phase stability, the registration accuracy of an ALMA image will be about 1/20 of the resolution, with a minimum of about 5 mas before other instrumental and atmospheric problems become dominant. This accuracy assumes that all calibrator and target positions are in the ICRF system and the position accuracy of the calibrator (see the ALMA on-line catalog) is less than about 3 mas.

For an experiment in which the target absolute position (or target motion between several experiments) is the main concern, more precise knowledge of the expected errors, and a suggested observation strategy is suggested. The simplest calculation is the *theoretical* astrometric accuracy, limited only by the image by the signal-to-noise (SNR), which is the ratio of radio source peak on the image. This is given by

$$\Delta p = 60 \text{ mas} * (100\text{GHz}/\text{FREQ}) * (10\text{km}/\text{BSL})/\text{SNR} \quad (10.5)$$

where

Δp = Theoretical accuracy limit by SNR only.

BSL = Maximum configuration baseline in km.

FREQ = The observing frequency in GHz.

SNR = The signal to noise at the source peak.

For example, at Band 7 (350 GHz), for the 1 km configuration (C40-5), for $\text{SNR} = 20$ and $\text{BSL} = 1$ km, the theoretical astrometric accuracy is 8.5 mas. However, this SNR calculation alone rarely indicates the true astrometric accuracy because: (1) beyond a SNR of 100, further accuracy is not likely because of imaging and calibration uncertainties at the level of 0.01 of the resolution; (2) short-term and medium-term atmospheric phase variations degrade and distort the image and the astrometric accuracy; and (3) large-scale atmospheric structure over the array and array parameter errors (antenna position offsets) produce astrometric errors that are > 5 mas or even for the largest ALMA configuration.

Most ALMA long baseline observations and all high-frequency observations will include observations of a *check source*. It is a quasar of known position, about the same distance from the calibrator as the target, but observed less frequently. After the same data reduction as that for the target, the offset of the check source image position (peak or centroid) from its a priori position is a measure of the astrometric accuracy of the experiment. The check source offset cannot be used directly as a correction to the target position (more complicated scheme given below), but provides an estimate of the astrometric error expected for the target. Also, the check source image coherence (peak / total flux density) is also a good indication of the quality of the calibration. If the check source image is badly distorted, then more serious problems are occurring with the data and reductions.

Virtually all astrometric phase errors are delay-like; that is, scale with frequency. Hence, simply changing an observation from a low frequency to a higher frequency, other things being equal, will not produce higher astrometric accuracy with the higher resolution since the phase errors will also increase with the same frequency ratio.

The scaling of the delay errors with baseline length b , however, is more complicated and consists of three major components: (1) short-term (30-sec) variations, mostly from the water vapor cloud-lets; (2) longer-term (3-30 min) variations, from a combination of larger-scale (km-size) water vapor and dry air component over the array; (3) systematic errors (antenna position errors) and 10-40 km sized atmospheric features over the ALMA site. Component (1) delay errors scale as $b^{0.3}$ beyond about 300 km baselines. Component (2) scales roughly as $b^{0.5}$. Component (3) scales as $b^{0.8}$. Hence, the gain in astrometric accuracy at the longer baselines is only as $b^{0.2}$.

Recommendations for astrometric accuracy, based on the above phase properties of the ALMA environment/array are:

1. Choose the closest calibrator with an accurate position. The use of the low SNR calibration by combining polarization and the SPW's (continuum source) to detect a close calibrator should be considered
2. Make sure a check source is included in the experiment.
3. Choose a frequency for which the SNR at the peak of the target is maximum. Thus, for a thermal object, Band 7 is a good compromise between flux density strength and system sensitivity. For non-thermal objects, a lower frequency may be more optimum.
4. Choose as high a resolution as needed to separate any blending of the desired component with other nearby emission. However, do not choose a resolution for which the target is too resolved, so that the peak flux density decreases faster than the increased resolution.
5. Request better than average weather (low pwv and low phase fluctuations) since the astrometric accuracy is directly dependent on the phase fluctuation level. For example, the normal *Go/noGO* criterion of $< 30^\circ$ phase fluctuations for observations should be reduced to 15° . Also request observations at as high an elevation as possible. Astrometric accuracy degrades very quickly for elevations lower than about 50° .

For astrometric accuracy better than about 10 mas rms, multi-calibrator observations should be considered. But, before asking for a more complicated observing strategy, check the following: First determine the optimum frequency and baseline-length (resolution) suitable for the target. Secondly, make sure that the observations are sufficiently long so that the estimated target image SNR is > 15 ; otherwise, multi-calibrator observations will not be warranted, although a check source should always be included. Finally, make a strong argument for observations in very good weather, perhaps that is required for the next higher observing band.

The astrometric accuracy can be increased by a factor of two to four by observing the target with two to four quasars with accurate positions that are near and surround the target. The main phase reference calibrator should be the quasar that is closest to the target. The other (secondary) calibrators should surround the target within a 10° radius. These other calibrators need only be observed once every three or four phase calibrator/target cycles. The analysis essentially determines the positions of the secondary calibrator images during the experiment (they will change and be offset from their known position by many mas), fits a linear function of the secondary position offsets versus separation from the phase cal (deriving a linear refraction wedge over the array), and then corrects the target position for the effects of this refraction wedge.

The normal calibration of these data should be standard in Cycle 4 since the observations and calibrations are no different than the standard interferometric experiment, except for the use of several check sources. Further multi-source astrometric analysis of the resulting images or antenna-based gains are the responsibility of the PI.

Please discuss the several observational options and solicit further advice from the help desk at your ARC. Since there is some restrictions in the cycling of source observations that can be scheduled, the observation strategy should be discussed.

10.6.4 Long Baseline Observations with Baselines > 2 km

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

1. Does the source have sufficiently small-scale structure for long baseline observations? For example, at Band 7 with a 5-km array, the rms image noise in a 600 sec observation over a bandwidth of 7.6 GHz is 0.1 mJy. For this resolution of $0.06''$, the rms brightness temperature will be 0.3K. For line emission within a 10 km/s velocity, the brightness sensitivity rms limit is about 50K. This requires an object with more than 500K temperature in order to obtain sufficient SNR in the line image.
2. If the target is known to have significant emission at the longer spacings, then self-calibration of the target may be possible. This is more important for long baselines, because the phase variations will be larger, and 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 spacings must be comparable to that needed to detect a phase calibrator, given in phase sensitivity table in Appendix C, Table 1. For example at Band 7, with 7.5-GHz bandwidth, the correlated flux density of the target must be > 125 mJy at the longer baselines. Such precise knowledge about the target may not be known, but the PI should have some indication of small-scale structure, otherwise the observational results may be disappointing.
3. 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, cone-search observations or the weak calibrator survey program may be used to find a more optimum calibrator. Also, fast-switching Band 3 to Band 6 and 7 observations are now being tested for the long baselines.
4. If the source is extended more than $2''$ at Band 3 (scaling inversely with frequency), then the long baseline configuration image will not be able to include all of the emission. For Cycle 4, an additional more compact configuration must be proposed as a second science goal in the proposal.

10.6.5 Total-Power Calibrations

For sources that are significantly more extended than the 7-m primary beam size, twelve-meter antennas (often the PM antenna) are used to make on-the-fly raster observations that cover the specified source area. The observing modes are discussed in Chapter 8 and the imaging in Chapter 7.

1. *Tsys* Calibration: The data are calibrated into brightness temperature in units of Kelvin using the equation $T_{a*} = T_{sysON} - T_{sysOFF}/T_{sysOFF}$. where T_{sys} is the system temperature and ON and OFF are, respectively, the data on-source (i.e. during the raster scanning) and off-source (observations towards a defined reference position which should contain only background emission). For very extended sources. it is important that the reference position is free from target emission and for this reason the off position may need to be up to several degrees away.
2. Baseline Subtraction: After apply the *Tsys*ON-OFF calibration, the off-source regions within the on-the-fly area should be zero for all spectral channels. However, due to the changing and turbulent wet atmosphere in the primary beam, low-order spectral changes will be present across the spectrum. These are removed by determining the average spectrum for the off-source positions, and removing them from all raster positions. Be aware, that this technique will remove large-scale continuum emission. Further baseline correction can be made after imaging
3. Flux Calibration: In order to calibrate the single-dish data from units of Kelvin to units of Jy/beam, a conversion factor is applied. This Jy/K conversion factor is obtained from periodic calibration measurements (beam maps) of strong sources of known continuum flux density (planets or quasars). The beam size is also measured.

10.6.6 Acknowledgments

I thank Satoko Takahashi for contributions to the high frequency sections, Neil Phillips for contributions to the on-line a priori calibration section, and Catherine Vlahakis contributions to the total power section and by improving the readability and accuracy of the text.

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. 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.

11.1 Cycle 4 "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, ALMA cannot guarantee that all the scheduled Cycle 4 projects are completed (specially those with grades B and C), or that specific criteria like sensitivity

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

or angular resolution are precisely met, give the restrictions of the array configurations, system performance at the time of the observations, etc. The observatory will therefore attempt to meet the sensitivity and resolution stated by PIs in their Science Goals, within certain tolerances as described later in this chapter. In principle, each project component ("Schedule Block" or SB) is scheduled for the number of executions that are expected to reach these goals (based on nominal Array performance values). Additional executions may be needed during the observing cycle if specific executions do not pass QA0 (see below).

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). It is expected that during Cycle 4, an upgrade of the AQUA tool will be carried out, allowing the complete computation of the contents of the Calibration Summary files within the AQUA environment using the information available in the ASDM files for each Executionblock.

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 (unavailable for Cycle 4).

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 data products, 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 4 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 4 (i.e., forty 12-m antennas on the main array, ten 7-m antennas on the Morita

Array, and three TP 12-m antennas), 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. The number of executions may be changed (i.e., increased/decreased) based on the total execution fraction of a given member OUS at a given point in time.

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 are, 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 that the science goals set by the PI can be compared with the actual value in the data products (i.e., rms, angular resolution, SNR, dynamic range, etc).

During Cycle 4, 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 imaging. The Pipeline provides a Weblog page with detailed information on the quality of the calibration, which will provide the basis for Pipeline QA2. This information will also be made available to the AQUA tool during the second half of Cycle 4, so that it can be included into the overall quality assurance reports. For manually-reduced datasets, the QA2 will be carried out with special purpose scripts whose outputs will be included in the data packages delivered to PIs. 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 EcBlocks 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: 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 the specifications listed in section 11.7. 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 scheduled for additional observations. If these data fail to obtain enough additional observations to pass QA2 by the end of an observing season and they are not grade A proposals, 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 4 User Policies document.

11.6 The Quality Assurance Report

During Cycle 4 the Quality Assurance Reports of QA0, QA1 basic QA2 will be accessible to the PIs using Snoopi. These reports will be generated using the AQUA tool. For the additional QA2 reports, for Pipeline calibrated datasets, the PI will receive, as part of the data delivery, the Weblogs. 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 4, a complete 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 (including the Weblog information). 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 do not fulfill the QA2 pass criteria. 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 good quality, they will be released to the PI at the end of the cycle.

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 (see below). Furthermore, the execution has to be "significant" in terms of the fractional execution (i.e., more than 0.2) or in terms of sensitivity (i.e., more than 50% of expected Cycle 4 sensitivity per execution of a given SB). If it does not meet these criteria, but has some useful data (calibrator or source), it is classified as QA0 SemiPass. Such data are available through the ALMA archive, but they 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 4 are based on the following:

- **Antennas:** QA0_SEMIPASS for fewer than 19 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. Situations 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 the bandpass calibrator signal is 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 19 antennas in the 12-m Array, or more than 3 antennas in the 7-m Array, or more than 1 Total Power antenna have T_{sys} values $> 50\%$ of the others, have differences between polarizations $> 50\%$, or have $T_{\text{sys}} > 2000$ K.
- **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
- **Time-Critical Observations:** QA0_SEMIPASS if observations have not been carried out at the requested dates and times.

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 at the GOUS (Group OUS) level (or at the Member OUS, MOUS, level if no data combination is required), 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 9/10, in flux units per beam, respectively), the synthesized beam shape (i.e., for some of the ALMA cycle 4 configurations², the differences with the requested angular resolution could be as high as 30% in linear dimensions or 69% in area of the synthesised beam on the sky) and the calibration quality (phase RMS of 5 degrees after calibration, and absolute flux scales within the ALMA specifications). For the individual MOUS that must be combined at the GOUS level, the RMS values quoted above should be scaled by the relative sensitivities of the configurations, using the most compact 12-m Array configuration as reference.

²During Cycle 4, SBs shall be executed following the recommendations given by the OT, which will select the configuration that produces the closest values to the requested angular resolution.

Chapter 12

Data Flow and Logical Data Structure

This chapter describes 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 selected from the list provided by the Scheduler subsystem, and the execution is started. An SB is the smallest, calibratable element of a project and they are kept in the proposal along with other ancillary information, in a series of XML tables (XML is the “Extensible Markup Language” format, common in database systems). 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) 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.

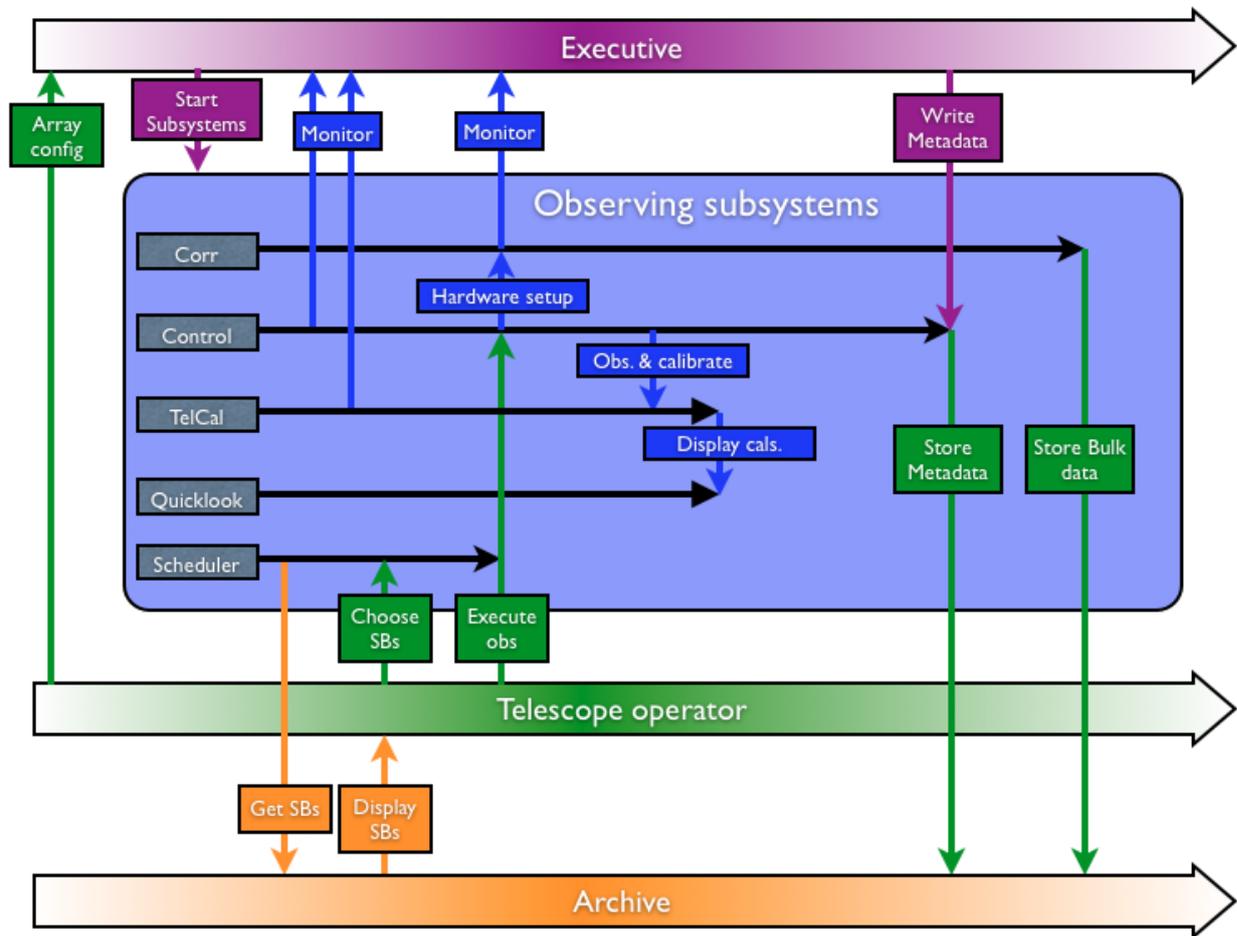


Figure 12.1: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

Once the execution of the SB has been selected from the Scheduler, it is dispatched to the Control subsystem. Control parses the SB XML document into python objects (with the help of the SSR scripts¹ for a given observing mode, which interpret the XML metadata in the SB), 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

¹SSR stands for “Software Science Requirements”. It is a collection of Python scripts that read relevant parameters from the XML metadata of SBs, and pass all the relevant information to the CONTROL software to execute the needed observations. SSR scripts also define the cadence of *all* the calibration observations during the execution of the SB.

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 (see, for instance, the 2006 ADASS contribution by Hafok, Caillat and McMullin). 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. In addition, it also monitors datasets and condenses all 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 displaying 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.

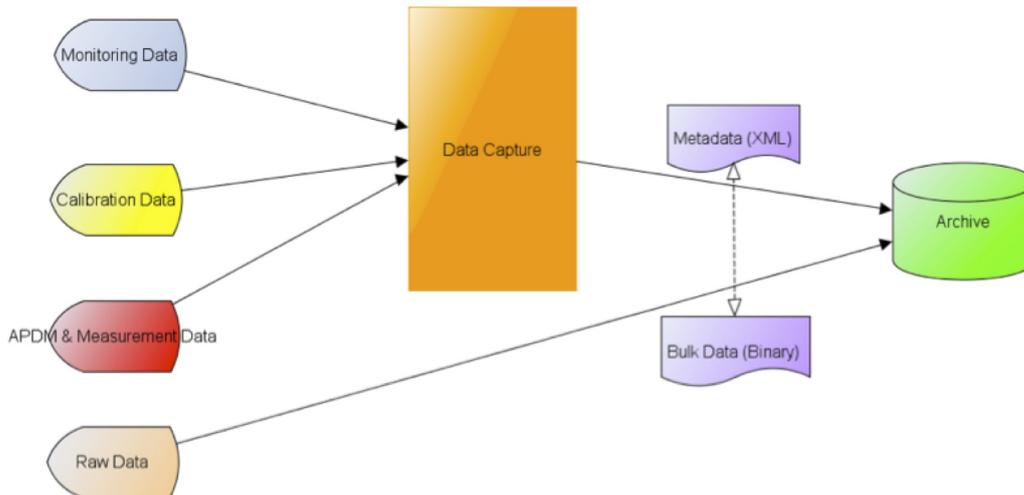


Figure 12.2: Data flow components.

12.2 The ALMA Science Data Model

The ALMA Science Data Model² (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

²In order to be able to have a general schema for any ALMA project, the information on any Observing Project is stored in the ALMA Archive following the APDM, the ALMA Project Data Model. The APDM defines a series of generic structures that are shared by all Observing Projects, including the project status, the proposal information, the outcome of the proposal review process, and the scheduling blocks with their interdependencies. All these structures are filled with the relevant information (or with links to it in the Archive) at the time of the creation of the project, and updated accordingly during its complete lifecycle.

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 commissioned for use on Cycle 3 data onwards. 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 needed to perform ALMA Quality Assurance 2 (QA2), i.e. it calibrates the data and performs imaging. This results in data products that can be delivered to the 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 Pipeline Quickstart Guide. The Pipeline QA scores, and the algorithms used to determine them, will 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.

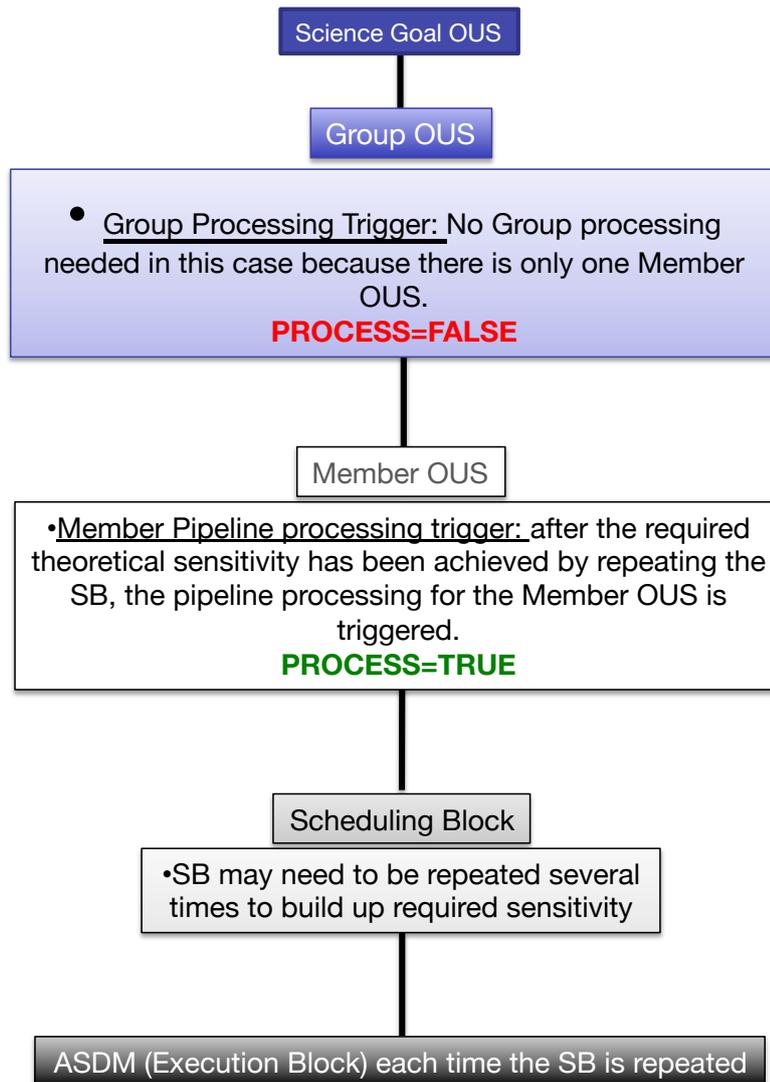


Figure 13.1: Pipeline processing of a Science Goal for single epoch, single array data. The same SB may need repeating N times to build up the requested sensitivity. This Science Goal would result in 1 Pipeline processing, with 1 set of science products delivered to PI.

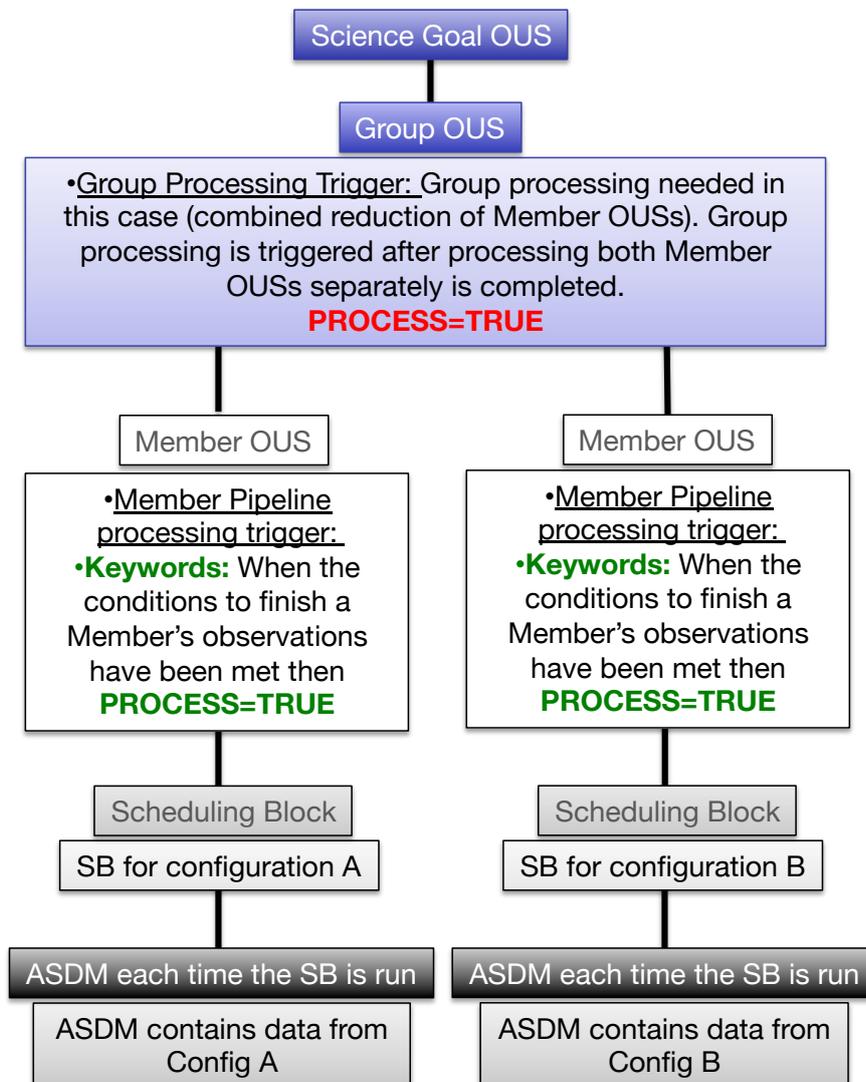


Figure 13.2: Pipeline processing of a Science Goal containing multiple ALMA configurations or arrays (e.g. 12-m, 7-m or TP). The data from the different configurations/arrays is processed independently in two separate Member OUSes, with science products created for each MOUS. Once the processing of the 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).

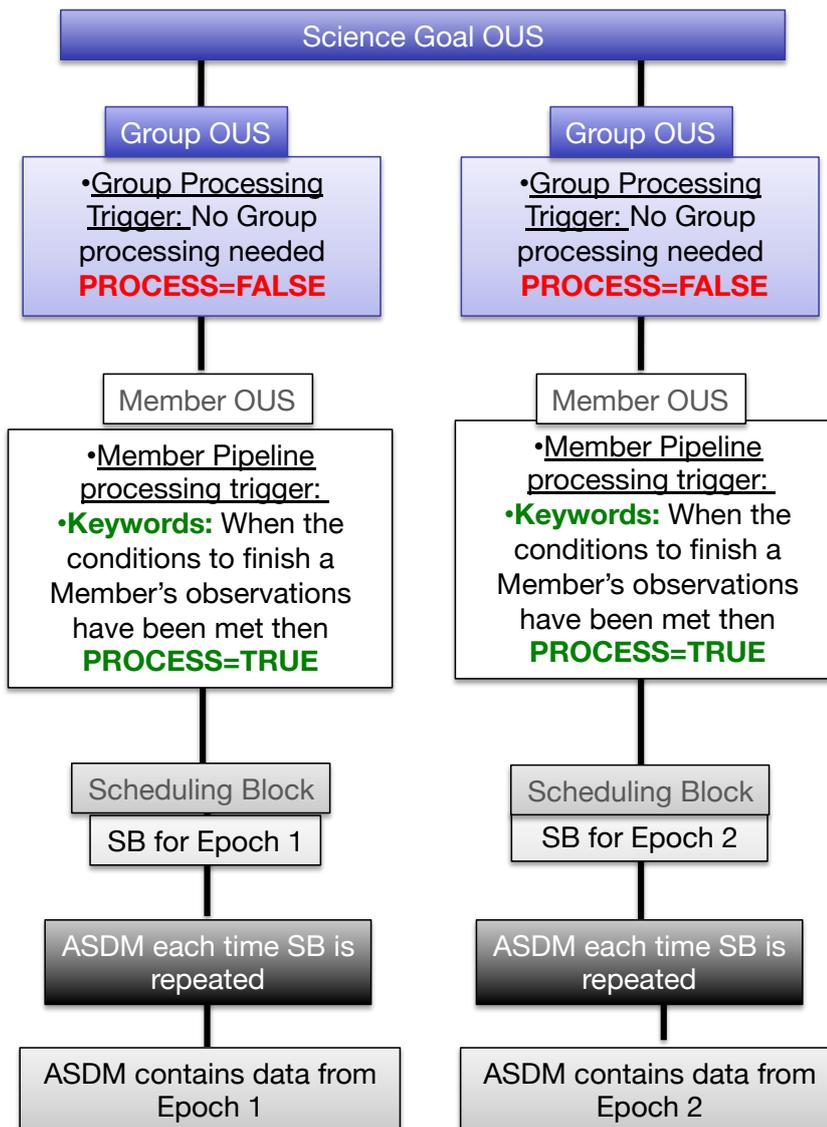


Figure 13.3: Pipeline processing of a Science Goal containing multi-epoch observations. Each epoch is processed independently by the Pipeline, with science products created for each epoch. Data from different epochs are not processed together to form a combined product (the Member OUS from each epoch 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 4 are listed in the Cycle 4 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 science target(s) observed in an OUS, the data from all Execution Blocks is now imaged together for each target. During Cycle 4, some proportion of ALMA OUSes may be calibrated using the Pipeline but imaged manually. In other cases the pipeline will be used to perform both calibration and science target imaging. It is also possible that manually-calibrated data will be imaged using the Pipeline.

13.3.2 Pause after Data Import

In current practice, the ALMA Pipeline does not necessarily 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 4 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 is 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 in the dataset. 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 4 Pipeline Quickstart Guide for a description of how the Cycle 4 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 Raw Data Flagging

The task *hif_rawflagchans* task flags deviant baseline/channels in the raw data. Data from the bandpass, phase and flux calibrators are examined and, where bad frequency ranges are determined on a per baseline basis, these data are flagged throughout the dataset. For more information on the parameter settings for this task, see the Pipeline Reference Manual.

13.5.5 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.

13.5.6 System Temperature Calibration & Flagging

The system temperature (T_{sys}) calibration table is generated by the Pipeline task *hifa_tsyscal*, which employs the CASA task *genical*. This table is then flagged by the task *hifa_tsysflag*. Flagging of T_{sys} spectra takes place because of high median values, high median derivative (ringing), edge channels, field shape (T_{sys} spectra whose shape differs from those associated with measurements made for the bandpass calibrator), and on the basis of spikes/birdies.

13.5.7 Antenna Position Corrections

The Pipeline task *hifa_antpos* can correct antenna positions via the *antennapos.csv* file.

13.5.8 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 visibility 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.

13.5.9 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 4 Pipeline Reference Manual. The CASA tasks employed during the heuristics for this step are *bandpass* and *gaincal*. The amplitude gain solutions are also flagged on the basis of deviant rms.

13.5.10 Bandpass Calibration

The Pipeline task *hifa_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). The solution interval is determined by the expected signal-to-noise ratio (SNR) on the bandpass calibrator. 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).

13.5.11 Mapping Phase Gain Solutions

Phase gain solutions derived from broad spectral windows can be applied to the data in narrow spectral windows by the Pipeline, in order to improve calibration results. The *hifa_spwphaseup* task computes the mapping needed and also creates a calibration table that removes bulk phase offsets between the different spectral windows using the bandpass calibrator (i.e. it “phases them up”).

13.5.12 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.

13.5.13 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. In addition a limited uvrange is used. This is because the model is more trusted for the shorter baselines. This is performed by using the CASA tasks *gaincal* and *fluxscale*.

13.5.14 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.

13.5.15 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. For phase solutions, the mapping determined by *hifa_spwphaseup* is used.

13.5.16 Imaging

At minimum, the Pipeline images the bandpass and phase calibrators as a diagnostic check, producing an image per spectral window. The Pipeline can also perform imaging of the science target(s), either in the same Pipeline run as the data calibration or starting from measurement sets that have already been calibrated. Data that were not calibrated by the Pipeline can therefore be run through the imaging stage only. In some cases, Pipeline calibrated datasets may be manually imaged by ALMA staff however.

Both for calibrator and science target imaging, the Pipeline uses the task *hif_makeimlist* to determine which sources need to be imaged. This task also determines, the phase centre, cell size and image size (*imsize*) to be used.

The imaging is then performed using the Pipeline task *hif_makeimages*. This task employs the CASA task *tclean* to perform imaging, and the *tclean* commands used, including parameter settings, can be found in the Pipeline *casa_commands.log*. The task *hif_findcont* identifies line-free frequency ranges of the dataset with which to form the Pipeline continuum products. The Pipeline additionally creates images of the continuum-subtracted dataset. Note that the Pipeline image cubes always use a frequency, rather than velocity, axis.

13.5.17 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 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(s)

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.

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. The solution interval in the pre-bandpass *gaincal* can be lengthened.

13.7 Total Power Data Processing

In Cycle 4, there will be a maximum of one single dish (Total Power) Science SB per Member OUS. The Science SB will have one or multiple science targets, each with a corresponding rectangular observing field. Once all of the required number of Execution Blocks (EBs) for the Science SB are completed and the status changes to “FullyObserved”, the Single Dish Pipeline is triggered.

In Cycle 4 spectral line observations in Bands 3 to 7 are offered as a standard observation mode¹, the Pipeline always performs T_{sys} and blank-sky calibration, Kelvin-to-Jansky conversion, baseline subtraction, subsequent flagging, and imaging, to process data for science targets. The Kelvin-to-Jansky factors, which are obtained by the Observatory, are applied per antenna per spectral window during calibration inside the Pipeline processing.

The Observatory has compiled amplitude calibrator (planets and bright quasars) observations with 12m antennas in the TP Array to derive the trends of the Gain (Jy/K) conversion factor for each of the antennas as a function of frequency, ambient temperature, elevation, and day-night dependence. The Kelvin-to-Jansky conversion factor is then derived from these trends. The Observatory stores the derived factors in an internal database containing the relevant parameters which include observational information such as temperature and elevation of the calibrators when observed. Note that no amplitude calibrator image will be included in the data products. Instead, a list of the applied Kelvin-to-Jansky factors will be added to the delivered package for the PIs.

13.7.1 Data Import

The Pipeline first loads the ASDM into the pipeline context, unpacking and/or converting to measurement set (MS), the same as for the data import in Interferometric Pipeline processing. At the same time as importing the data, the binary data file flags are applied to the imported MSes, and a flag file containing online flagging information and a flag file template for manual flagging per MS, are generated.

13.7.2 Deterministic & Manual Flagging

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

- flags listed in the online flag files
- flags of data which have been affected by shadowing
- flags of un-needed intents such as POINTING, FOCUS, SIDEBAND, and ATMOSPHERE.

¹The standard and non-standard modes in Cycle 4 are listed in the Cycle 4 Call for Proposals.

- flags of edge of channels of low-spectral resolution (TDM) data

It is also possible to insert manual flagging if further flagging is needed for the Pipeline-processed data.

13.7.3 T_{sys} and Sky Calibration Tables

The Pipeline creates T_{sys} calibration tables after heuristic flagging is applied. Each of spectral windows for the science targets is automatically mapped to the corresponding frequency-matched T_{sys} spectral window.

The Pipeline also generates the blank-sky calibration tables (so-called sky calibration tables) that contain a collection of spectra of the ‘OFF’ (i.e. reference) positions, linking to each ‘ON’ (i.e. on-source) spectrum.

13.7.4 Kelvin-to-Jansky Calibration Table

The Kelvin-to-Jansky factors are derived by the observatory (i.e., independently from the Pipeline Operation) and inserted in a file which will serve as input to the Single Dish Pipeline. The Pipeline generates Kelvin-to-Jansky calibration tables using that input file. The users can notice that the WebLog gives a list of the Kelvin-to-Jansky factors applied to the data in the the summary page.

13.7.5 Applying the Calibration

The Single Dish Pipeline applies the T_{sys} , sky and Kelvin-to-Jansky calibration tables to the dataset: i.e., the Pipeline calculates $(ON - OFF)/OFF * T_{\text{sys}} * \alpha_{\text{Jy/K}}$, where ON and OFF are spectra of ON and OFF positions, respectively. T_{sys} is system temperature, and $\alpha_{\text{Jy/K}}$ is a Kelvin-to-Jansky factor. At the completion of this step, the resulting data have intensity units of Jy. Note that in case that no Kelvin-to-Jansky factor is provided to the Pipeline, the Pipeline just interprets $\alpha_{\text{Jy/K}} = 1.0$, and the units appearing in the WebLog and in the Pipeline products will be in Jy.

13.7.6 Automated Line Identification and Baseline Subtraction

The Single Dish Pipeline generates baseline fitting tables and subtracts low-order baseline from the spectra.

The Pipeline employs an automated line-identification algorithm using clustering analysis. Line identification is performed on the integrated data of all EBs (per spectral window). The algorithm searches for emission (or absorption) line candidates having a brightness above a certain threshold (based on the median absolute deviation) on the integrated spectrum per space-domain grid cell per time-domain group. The candidate emission (absorption) lines are first screened by whether those lines are identified also in neighbouring grids and different time-domain group. The clustering analysis is performed on the central frequency and line-width for all candidate emission, to determine the emission lines to be masked. Then the central frequency and line-width for each emission lines are estimated by taking into account the spatial distribution of the two parameters. The corresponding channel ranges are masked in the spectrum. The Pipeline analyzes each emission-masked spectrum through Fast Fourier Transform to obtain the power spectra, which is used to determine the number of segments for cubic spline fitting. Finally Pipeline performs baseline fitting using a cubic spline. See “ALMA Science Pipeline Reference Manual” for more detailed explanation on the line identification algorithm. This step (line finding and baseline subtraction) and the subsequent flagging step are executed twice in the standard procedure to improve accuracy of the line finding algorithm.

The derived baseline information such as parameters for line finding, fitting function, and fitted baseline are filed in the baseline calibration tables.

13.7.7 Baseline Flagging

The Pipeline examines spectra before and after baseline subtraction based on baseline rms and T_{sys} , and generates flag tables containing a collection of flagged data to apply them just before the imaging step. The baseline rms is calculated using emission-free channels (excluding masked channel of emission line by line finding algorithm). The flagging is done per spectrum (per dump), according to the following criteria:

- abnormally large T_{sys}
- remarkably large baseline rms compared to other spectrum
- rapid increase of baseline rms over time, based on deviation from the moving average of baseline rms
- significantly large baseline rms relative to the expectation computed for the observed T_{sys}

The threshold for each criteria are determined by Pipeline heuristic. See “ALMA Science Pipeline Reference Manual” for more detailed explanation on the criteria.

13.7.8 Imaging

At the imaging step, for the science target, the Pipeline produces a data cube with the native frequency resolution for each spectral window per field. Note that all imaging parameters are set to a cell size of 1/9 of theoretical beam size, and a grid function of Spheroidal Function.

13.7.9 Exporting the products

The final step of the Pipeline processing is to export the image cube defined in the pipeline context to the data products directory, converting and or packing it if necessary. The current version of the task exports the following products:

- an XML file containing the pipeline processing request
- a tar file for each ASDM / MS containing the final version of flags
- a text file for each ASDM / MS containing commands to apply calibration tables
- a FITS image for each science target source for each spectral window
- a tar file containing all calibration (T_{sys} , sky, and baseline) tables for each MS to be applied
- a tar file containing the WebLog files
- a text file containing the final list of CASA commands
- a python script containing the final list of Pipeline tasks can be used to calibrate the data
- a text file containing the Kelvin-to-Jansky factors

Chapter 14

Data Archiving

14.1 Introduction

The ALMA archive is at the centre of the ALMA data flow (Fig. 14.1). It is a combined database and binary data storage system that is accessed by the different software subsystems through the same software layer. The ALMA Archive stores all metadata and data of ALMA from the user accounts, over the proposals, the observatory and antenna configuration and monitoring to the raw and reduced science data. The science data and metadata are made available to PIs and archival researchers for querying and download following ALMA's data access policy.

The ALMA archive is divided into two parts. The ALMA Frontend archive (hereafter AFA), which provides the core persistence functionality and into the ALMA Science Archive (hereafter ASA). The latter holds a small subset of metadata of the AFA in a relational database and provides access to external interfaces like the Archive Query interface and in the future Virtual Observatory (VO) tools. The storage architecture is based on the Next Generation Archive System (NGAS) with Oracle technology for replicating the metadata.

Each of the three ALMA Regional Centers (ARCs) in North America, Europe and East Asia holds a copy of the entire ALMA archive for backup, user support and data distribution to the ALMA PIs and archival researchers. The ARCs provide a completely identical user experience to their communities and a user can download data from any ARC.

An ALMA Science Archive Manual is available from the ALMA Science Portal (<http://almascience.org/documents-and-tools>).

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 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 six times the average data rate. In such cases, PIs may get contacted in Phase 2 in order to see whether a reduction of the data-rate is possible.

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 ARCs. The process of copying the data to any of the archives involves a replication of the metadata (support data) and of the bulk data (ASDM and FITS files). All data transfer is done over the network. Metadata replication from SCO to the ARCs happens within a few seconds, transfer of bulk data can take longer (up to several hours), depending on the amount of data to transfer to the individual ARC.



Figure 14.1: The ALMA archive is in the centre of the ALMA data flow.

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 include all the information needed to describe the observations, including date and time, source coordinates, frequency settings for each spectral window, and spectral and spatial resolution.

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 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 an 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>

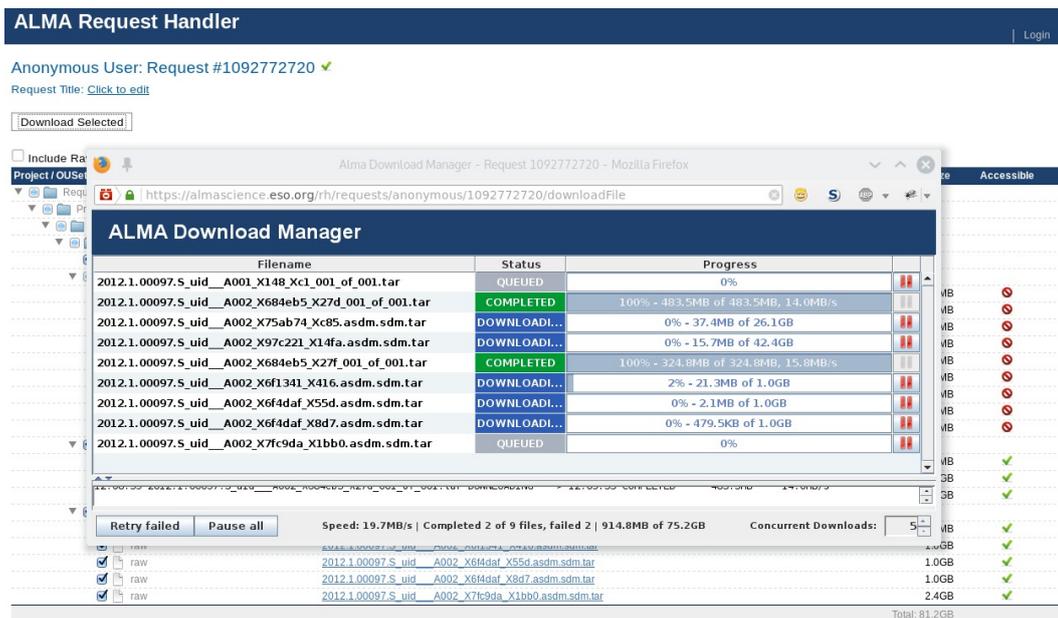


Figure 14.3: The ALMA Science Archive Request Handler

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 may 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 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: $-20^{\circ}\text{C} \leq T_{amb} \leq +20^{\circ}\text{C}$
- Gradient of temperature: $\Delta(T_{amb}) \leq 0.6/1.8^{\circ}\text{C}$ in 10/30 minutes
- Wind Velocities $\leq 6/9$ 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 4 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.

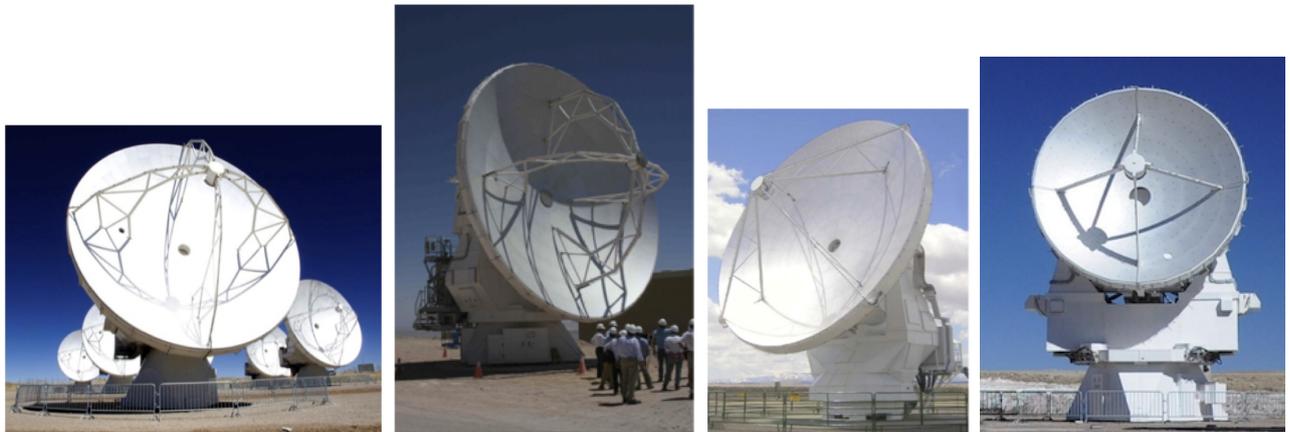


Figure A.1: The four different ALMA Antenna designs: Vertex 12 m, MELCO 12 m, AEM 12 m, and MELCO 7 m (from left to right).

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 aluminium subreflector. Subreflector adjustment is used to maximize the transfer of power into the receivers by compensating for changes in the focus position due to gravitational- and temperature-induced deformations. The backplane of the subreflector is attached to a hexapod that controls its position and orientation. The hexapod has six degrees of freedom, displacement and tilt around the three axes, horizontal (X), vertical (Y) and along the optical axis (Z).

	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, aluminium, 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.

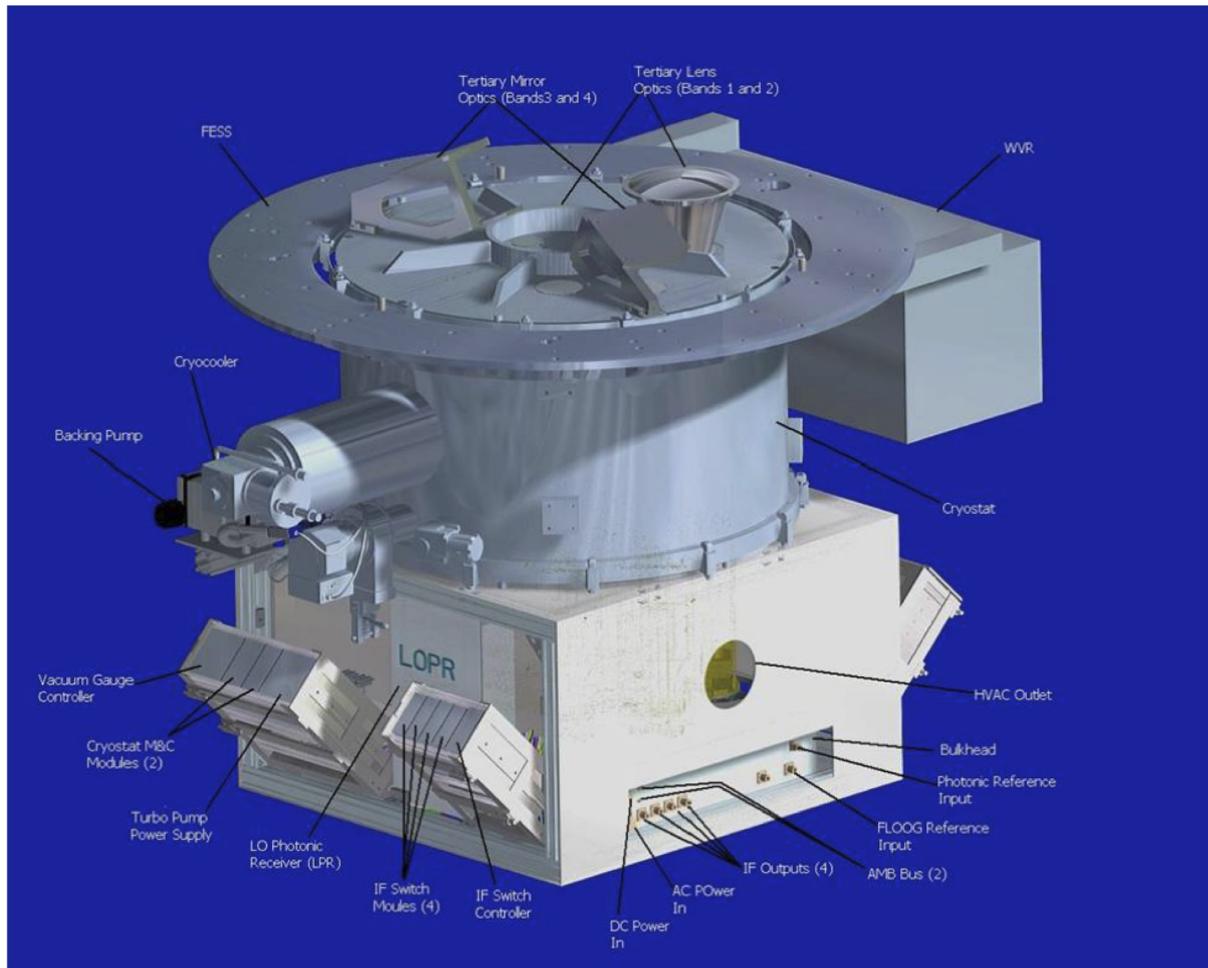


Figure A.4: Side view of ALMA frontend showing cryostat assembly, with room temperature unit below.

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 frontend, 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.

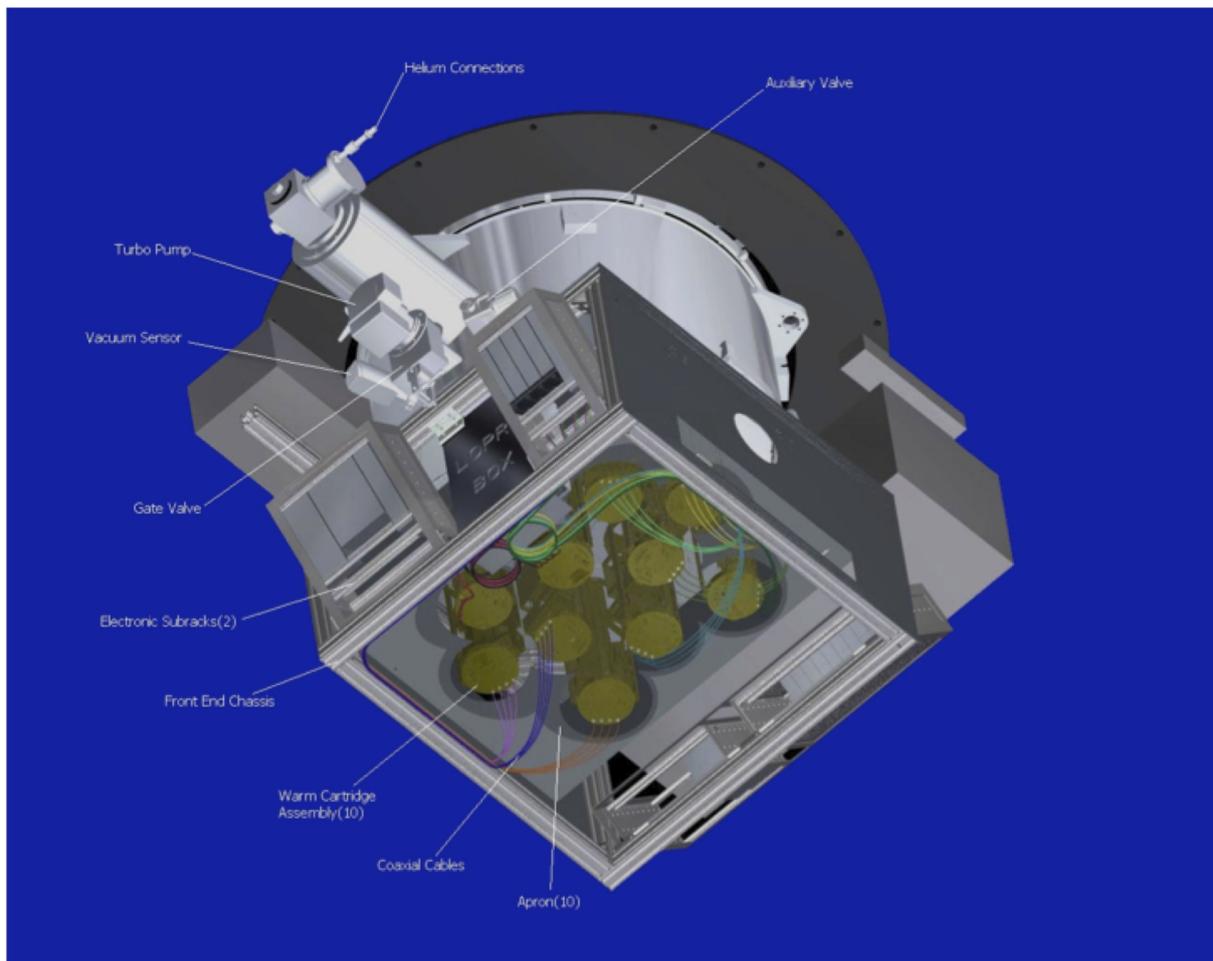


Figure A.5: Bottom view of ALMA frontend, showing WCAs.

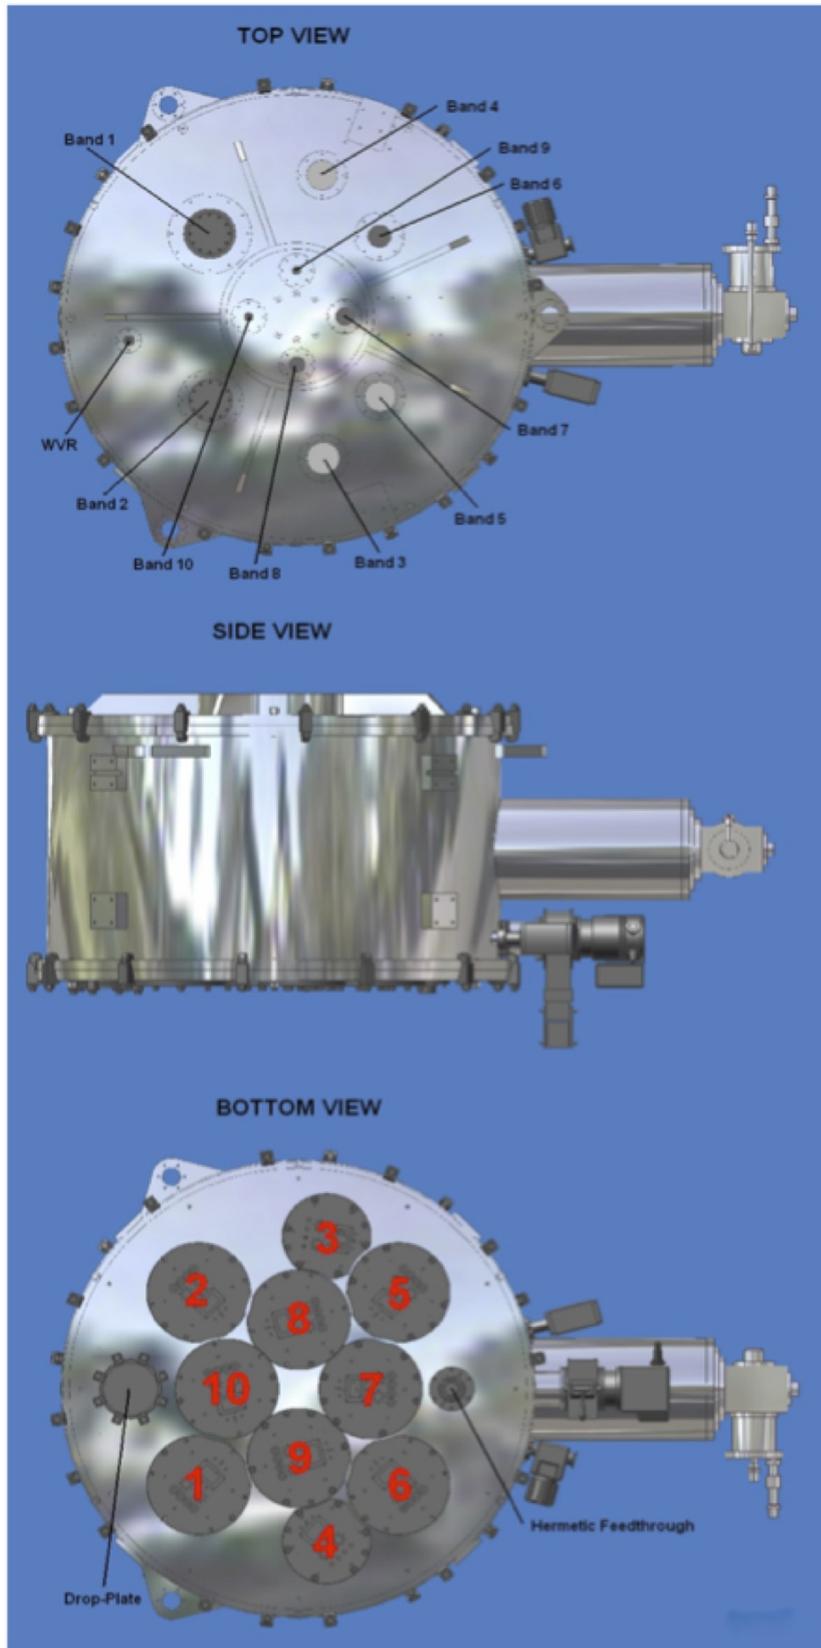


Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view).

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 4). 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.

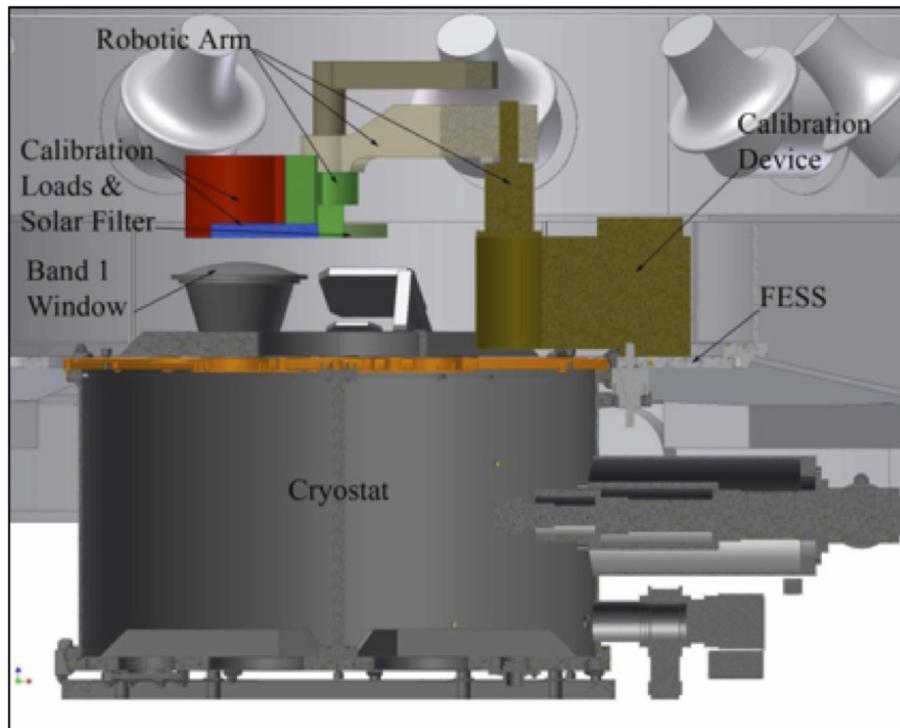


Figure A.7: Lateral view of the ACD on top of the ALMA frontends.

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.).

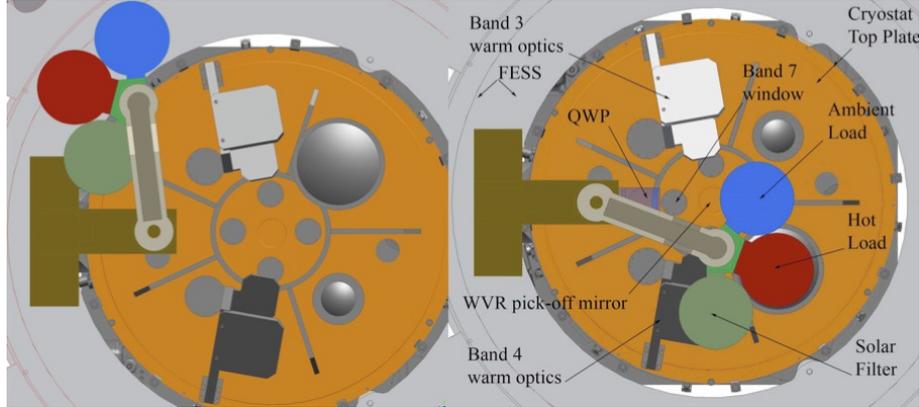


Figure A.8: Top view of an ALMA frontend showing the robotic arm of the ACD retracted during normal observations or on top of one of the frontend inserts for calibration. The current design has been improved by placing all the loads in a wheel.

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 4.

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 water vapor path fluctuations are expected to lead to small apparent pointing fluctuations which are analogous to the seeing effects in single-aperture optical telescopes. ALMA selected the 183 GHz line because it is quite bright and allows a more compact design than would the 22 GHz water line. It was decided to measure the temperature of the 183 GHz line in four regions offset from the center using filters of different bandwidths. The positions of the filters are indicated as blue boxes superimposed on the profile of the water vapor line in Figure A.9. The sensitivity specification for the WVRs is 0.08–0.1 K per channel RMS.

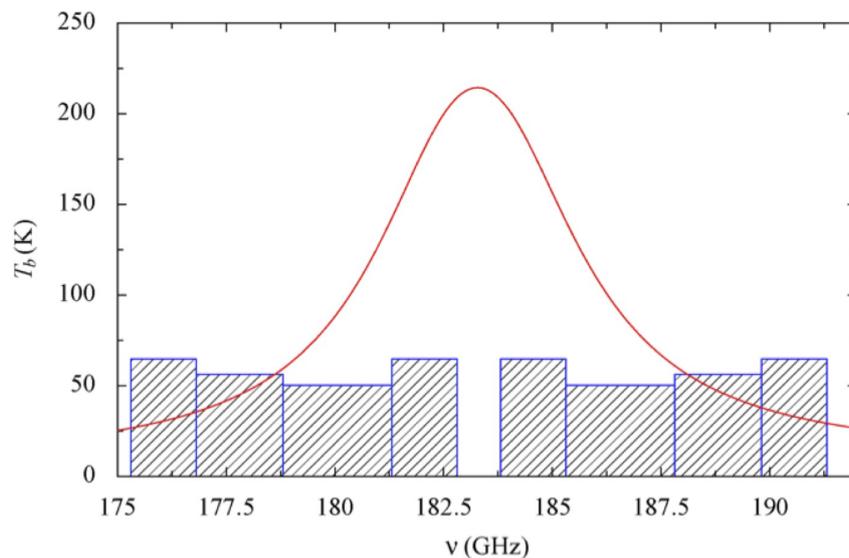


Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

It is very important that the WVR illuminates the same area of the sky as the ALMA band receivers in the near-field region. This is because the origin of the water vapor fluctuations is usually located in the lower troposphere (i.e., near the observatory), with one to several layers of water vapor clumps encompassing a wide range of sizes. Since the ALMA backends are located at the Cassegrain focus, an offsetting optical system (see Figure A.10) had to be designed to allow the WVR to measure along the optical axis of the antennas.

The WVRs are only able to detect the variations in atmospheric brightness temperatures due to the “wet”

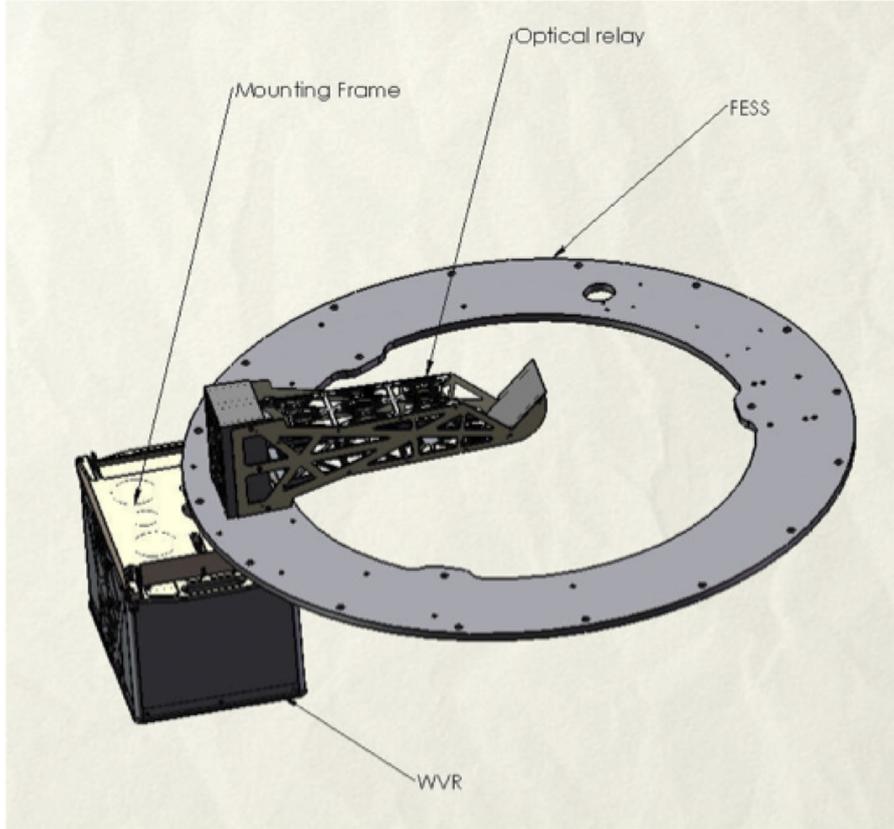


Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

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

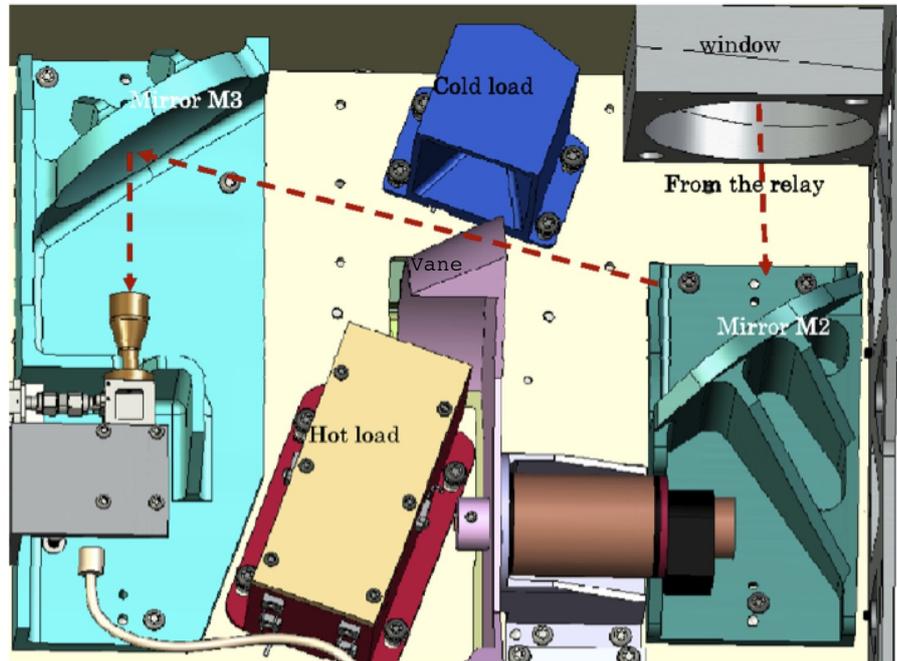


Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

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

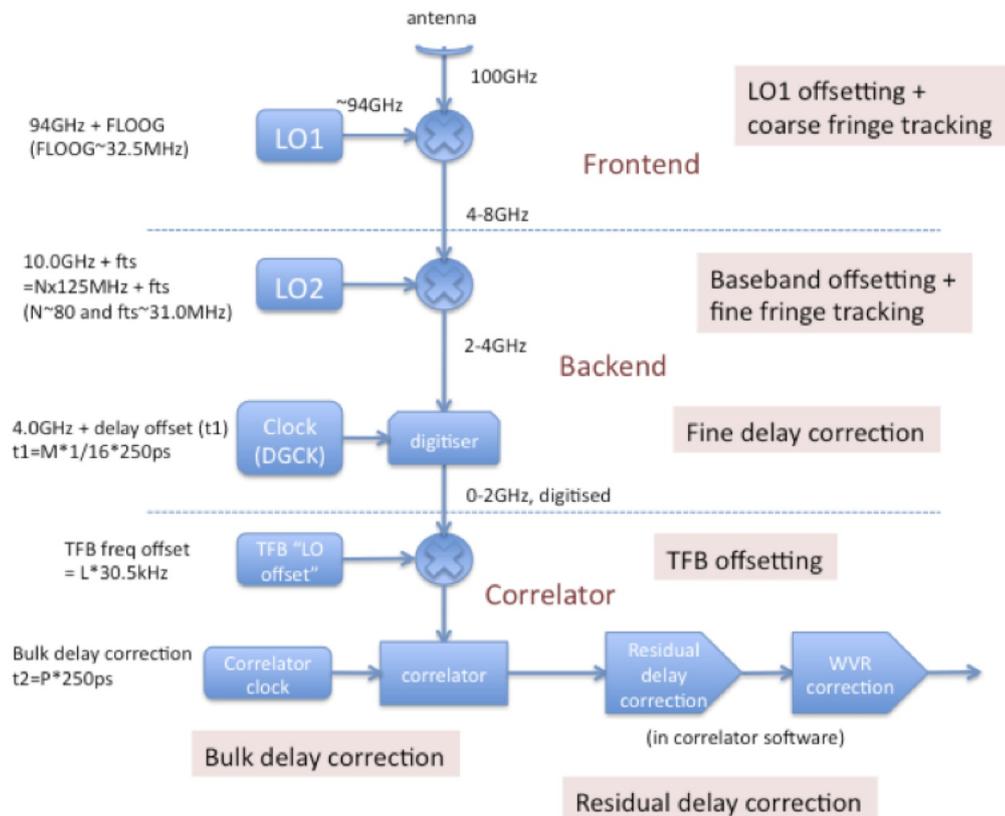


Figure B.1: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the Frontend, Backend, and Correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz seen in the USB. Some LOs (e.g. LO1) are continuously tunable; others have 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

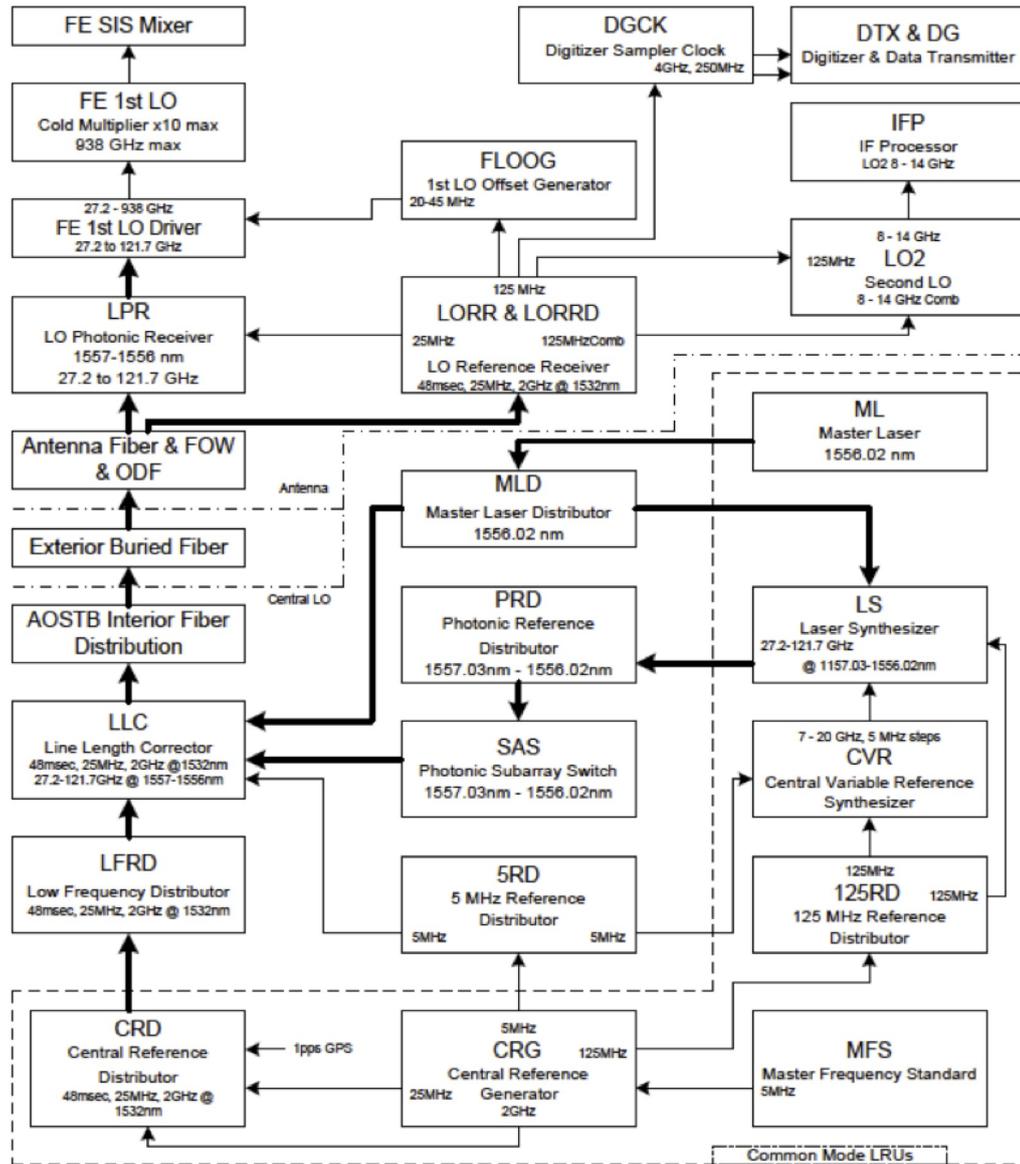


Figure B.2: Summary block diagram of the LO distribution system. The components in the lower section of the diagram (below the the dashed-dotted line) form the Central LO (CLO), located in the AOS Technical Building. The components above are located in each Antenna (only one antenna is depicted in this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

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.

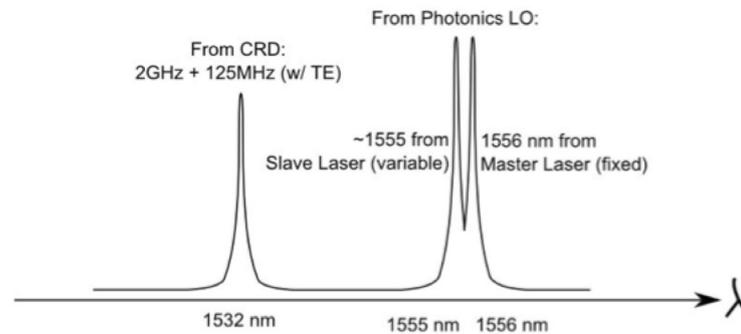


Figure B.4: The ALMA fibre signals. The 1532nm carrier contains the frequency reference signals, and the 1556/1555nm carriers are used to remotely generate LO1. Within each antenna, the optical fibers are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing 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

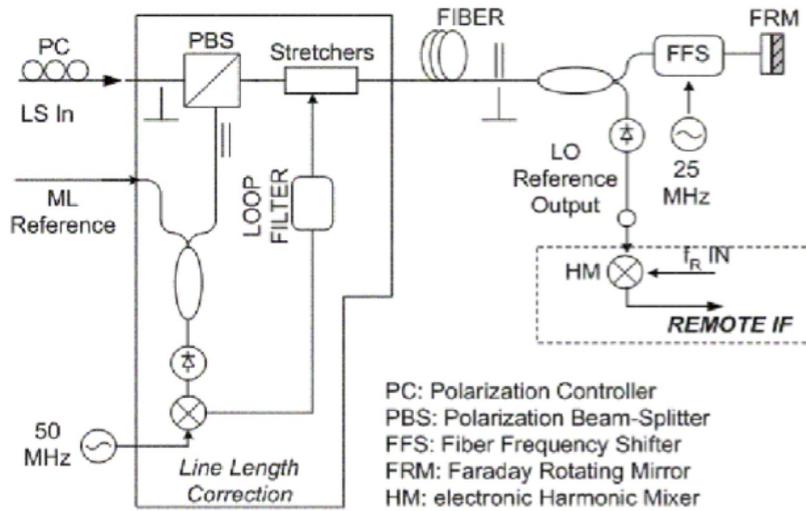


Figure B.5: Block diagram of the Line Length Corrector system for ALMA. The FRM (Faraday Rotation Mirror), shown upper right, is located in the band 9 cartridge of every front-end. It rotates the polarisation of the incoming light, and the resulting reflected signal is fed back through the buried optical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

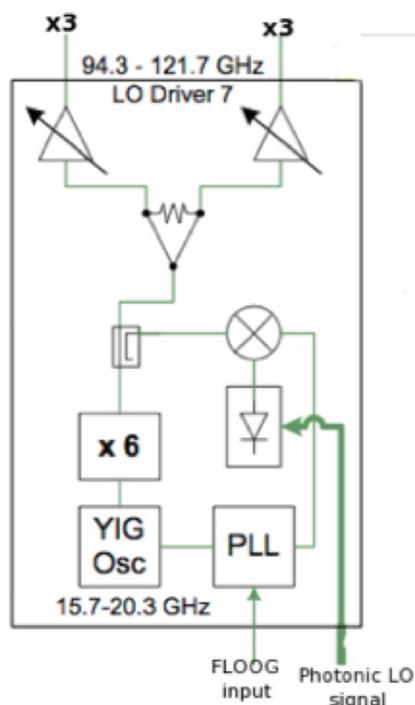


Figure B.6: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier, not shown here, is used to generate the LO1 frequency at 282.9 – 365.1 GHz (in this case, a x3 multiplier). 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).

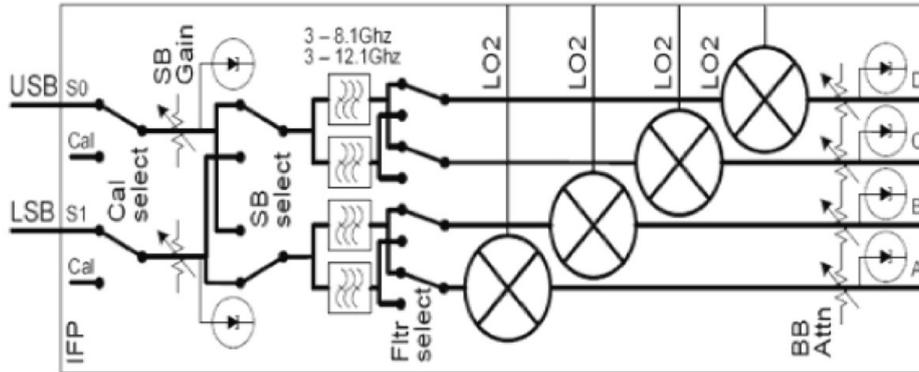


Figure B.7: Block diagram of one polarisation channel of the IF Processor. This has two IF inputs (at left), and feeds 4 IF outputs to the digitisers (to the right of the diagram).

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 6.4). 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.

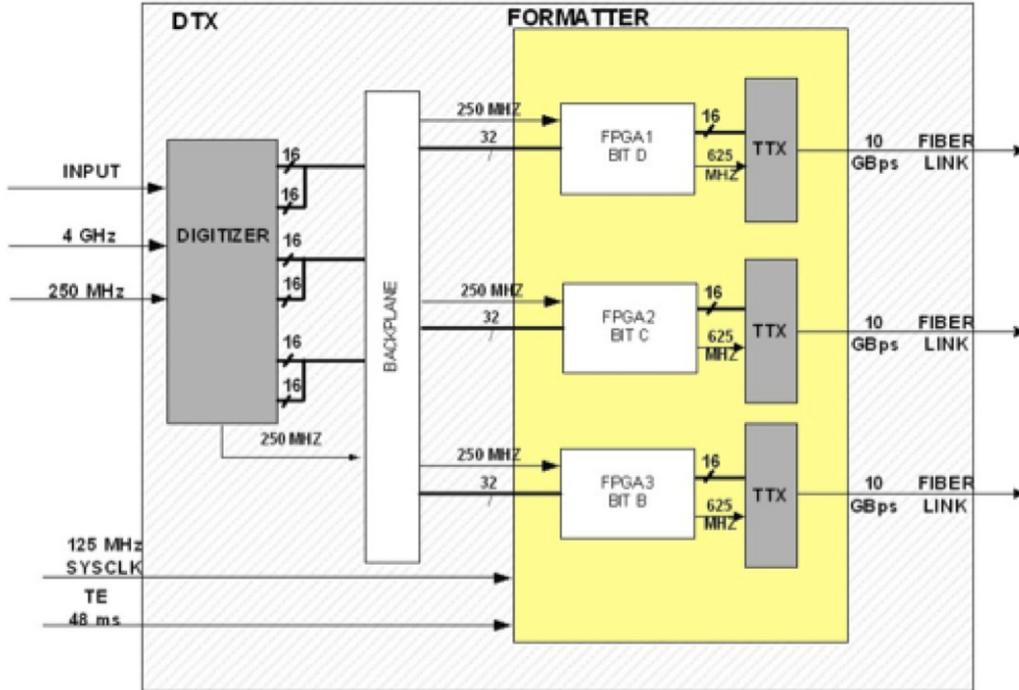


Figure B.8: The DTX Signal Digitization and Transmission system, as used in each antenna.

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 4.

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 Sources

C.1 ALMA Catalog Description

The ALMA catalog of sources contains information for over 10,000 calibrators south of $\delta = 45^\circ$. The indicated reference¹ brings up a query page from which 1) information on any particular calibrator or 2) a cone search for calibrators near any position, can be listed. An example of a cone search is shown in Figure C.1. The left-hand side shows the typical query inputs needed. The right hand side side shows the results of the query. Plots and a hard copy of the source lists can also be obtained. Additional information, such as polarization, can be obtained with the listing.



Figure C.1: ALMA catalog example; Left: Search form for calibrators within 7° of target at 18:24:00 -29:40:00. Right: List of calibrators near target.

The ALMA catalog is a compilation from many major catalogs (mostly at lower frequency) that have been generated over the last 20 years. The catalogs include Crates 8.4 GHz², VLA calibrator catalog³ at20g catalog⁴, vlbi (Petrov) catalog⁵, and SMA catalog⁶. No blind ALMA surveys have been made to find new sources since the above catalogs go sufficiently deep to include all sources that would be useful as ALMA calibrators. Since most of the catalogs are based on observations at lower frequencies, the flux density of the sources must be observed with ALMA in order to determine if their flux densities between Bands 3 and 10 are sufficiently strong (see below), and if they have extended structure. Many Observatory calibrations been developed to measure the flux density of thousands of sources in the catalog (weak calibrator survey) and to check the flux density of calibrators that are near targets (cone-search) that will be observed in the near future. These efforts will

¹<https://www.almascience.org/sc>

²ApJS, 171, 61, 2007

³<http://www.aoc.nrao.edu/~gtaylor/csourc.html>

⁴Murphy et al, 2010, MNRAS, 402, 2403

⁵http://astrogeo.org/vlbi/solutions/rfc_2016a/rfc_2016a_cat.txt

⁶<http://sma1.sma.hawaii.edu/callist/callist.html>

continue into Cycle 4.

About 2000 sources have been observed with VLBI techniques (Petrov catalog) and have positions that are accurate to at least $0.002''$. These are the sources that should be used as phase calibrators, and they are used for all ALMA system calibrations, for example, the baseline measurements. All calibrator positions are measured on the ICRF/ICRS astrometric system and it is strongly recommended that all targets that are observed with ALMA have positions on this quasar reference system. Most quasars in the ALMA frequency range have spectral index, α , ($S \propto \nu^\alpha$) that are between -0.5 and -0.8; hence with a measurement at Band 3 and 7, the flux density at Band 9 can be inferred to an accuracy of about 25%. On the average, the separation between a target and a good quality calibrator in the catalog is about 3° at Band 3, and about 8° at Band 9.

The flux density for a set of strong quasars that cover the sky is monitored every 10 days with ALMA observations at Bands 3 and also at Band 6 or 7, depending on the phase conditions. These grid sources are used as secondary amplitude calibrators and tied to the solar system objects with known brightness models. The flux density of quasars vary but with their frequent monitoring their flux densities at any time are accurate to about 5% at Bands 3, 6, 7, and 10% at Bands 8, 9 and 10. Most quasars are stable to 20% over months, but there are a few highly varying quasars that periodically change by a factor of three in a month or less.

C.2 Phase Calibrator Selection

The proper selection of a satisfactory phase calibrator is critical for the image quality of the target source. The major criterion is that its flux density is sufficiently strong in order to obtain an SNR greater than 15 for an antenna-based solution for each scan. The secondary consideration is that the closer the calibrator is to the target, subject to the sensitivity limit, the more effective it will be in reducing the phase variations in the target. For two potential calibrators that are well above the sensitivity limit, the one closer to the target should be selected.

With the addition of many recent ALMA observations of the flux density of many sources in many cases the optimum calibrator can be determined using an automatic query system as the SB is being prepared. However, cone-searches around targets will still be made for long baseline and high frequency projects where the closeness of the calibrator to the target is more important. If a selected calibrator does not have a measured flux density more recently than 90 days, and is less than about two times the detection limit needed for an experiment, a short ALMA observation (perhaps part of a cone search) will be made well to determine its current flux density. An observation at any band is sufficient since the quasar flux density spectral indices lie in a narrow range.

The lower flux density limits for a phase calibrator for all of the observing bands are given in Table 10.1 below. The flux limits can be scaled by $(N_{ant}/40)^{-0.5} (T_{sys}) [(t_{scan}/120) (\nu_{BW}/7.5GHz)]^{-0.5}$. Figure C.2 shows the probability of finding a detectable calibrator (40 antennas, 7.5 GHz bandwidth) in the ALMA catalog within a specified angular separation. For example at 492 GHz, there is an 80% chance of finding a calibrator within 4.0° of a random target.

C.3 Bandpass Calibrator Selection

The bandpass calibrator must be stronger than the phase calibrator since accurate calibrations will be obtained for relatively narrow frequency ranges. However, the ALMA bandpass shape changes by less than 1% over tens of degrees in the sky and several hours of time so the bandpass calibrator need not be close to the target (within $< 45^\circ$) and a sufficiently bright source should be available for most experiments. Also, a concern is obtaining sufficient SNR for FDM (narrow-band) channels. Testing has shown that for bandwidths < 128 MHz, the bandpass is flat to 1%; so calibration over narrower bandwidths are not needed⁷.

The bandpass calibration limit of 50 SNR for each antenna/SPW/polarization gain determination, giving a 1% rms error in amplitude and 2° in phase for each of the SPW/polarization channels. The flux density low limits for suitable bandpass sources with an observation of 15 min and a bandwidth of 128 MHz are given in the

⁷https://wikis.alma.cl/twiki/pub/AIV/EOCMemos/EOC_Memo27_Dec17_FINAL.pdf

Band	Frequency (GHz)	Assumed T_{sys} (K)	Assumed PWV (mm)	Flux Density (mJy)	Cal-target separation
3	86.243	83	5.186	4.4	1.1
3	115.271	173	5.186	9.3	1.5
4	146.969	86	1.796	4.7	1.3
6	230.538	102	1.262	5.7	1.5
7	345.796	151	0.658	9.0	1.8
8	461.041	431	0.472	28.5	2.1
8	492.161	638	0.472	43.7	2.5
9	624.208	4302	0.472	346.0	9.0
9	658.007	1721	0.472	145.2	5.0
9	691.473	1231	0.472	109.5	4.7
10	806.652	2405	0.472	257.7	7.8

Table C.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 40 antennas and elevation of 50 degrees. The total bandwidth in both polarizations of 7.5 GHz is assumed.

Table 2, below. The flux limits should be scaled by $(N_{ant}/40)^{-0.5} (T_{sys}) [(t_{scan}/900) (\nu_{BW}/128MHz)]^{-0.5}$. The above antenna-based specifications for 40 antennas produce an image channel to image channel rms variations of about 0.2%. For nearly all projects, the target SNR may give much larger variations in the spectral properties of the source than the bandpass calibration inaccuracy.

However, for the longer baselines in TDM mode, changes in delay caused by the atmosphere and baseline errors can produce phase slopes over the 2 GHz bandpass of up to 10 deg between the bandpass calibrator and the phase cal/target. If the phase cal is sufficiently strong, then its residual delay can be removed from itself and the target.

Band	Frequency (GHz)	Assumed T_{sys} (K)	Assumed PWV (mm)	Flux Density (mJy)
3	86.243	66	5.186	58
3	115.271	173	5.186	121
4	146.969	86	1.796	61
6	230.538	102	1.262	74
7	345.796	151	0.658	119
8	461.041	431	0.472	375
8	492.161	638	0.472	574
9	624.208	4302	0.472	4551
9	658.007	1721	0.472	1909
9	691.473	1231	0.472	1435
10	806.652	2405	0.472	3390
10	850.000	1817	0.472	2771

Table C.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, 40 antennas, and elevation of 50 degrees. The sensitivity assumes a bandwidth of 128 MHz, regardless of the number of spectral channels within this width.

The present list of bandpass sources that are frequently monitored is shown in Table C.3. Some changes will be made for extremely variable sources or with the addition of sources that have become strong at Bands 7 and 9. Each source has several aliases so that J1229+0203 is also recognized as 3C273 or 3c273.

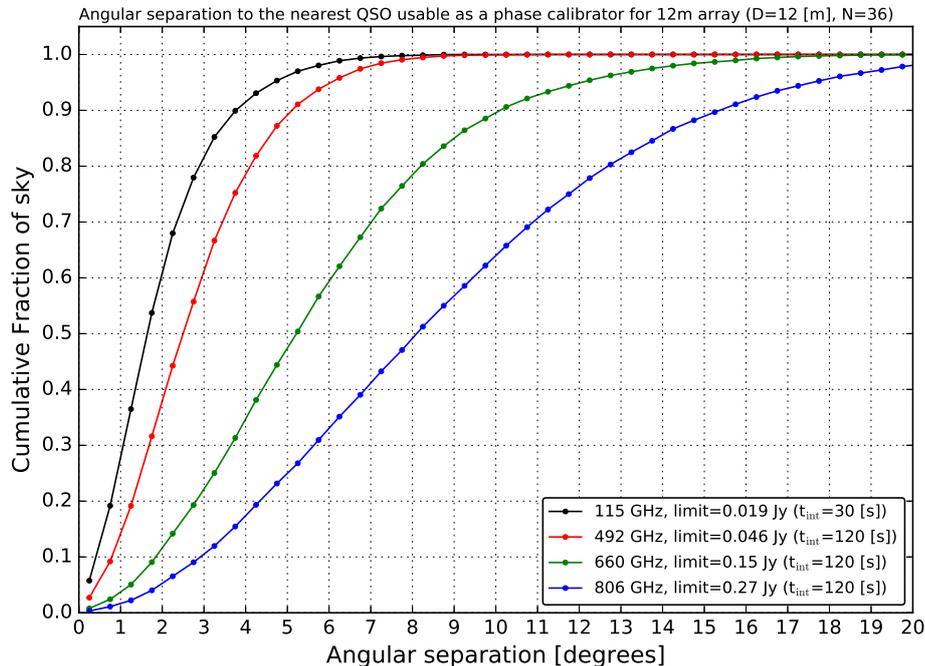


Figure C.2: The probability of finding a phase calibrator within an angular separation of a random target. The y-axis is the probability, the x-axis is the target-calibrator angular separation. The plot is given for four different ALMA frequencies using the 12-m array with 40-antennas and 7.5 GHz bandwidth.

C.4 Other Calibrator Types

C.4.1 Absolute and Relative Flux density Calibration

In order to determine the *absolute* flux density scale of an experiment, the experiment should contain a solar system object with a known brightness model; Mars, Uranus, Neptune, Ceres, Pallas, Juno, Vesta, Ganymede, Callisto, Titan. Their accuracy varies with resolution and frequency, from 3% to 10%. If a solar system object is not available, the bandpass sources chosen is usually a grid sources (Table C.3. These sources are monitored with ALMA every two weeks at Band 3 and Band 7, so that their estimated flux density at any time and frequency can be estimated to about 5% accuracy below 400 GHz and about 10% to 15% accuracy above this frequency.

J0106–4034	J0132–1654	J0237+2848	J0238+1636	J0319+4130
J0334–4008	J0423–0120	J0510+1800	J0519–4546	J0522–3627
J0538–4405	J0635–7516	J0750+1231	J0854+2006	J0927+3902
J1037–2934	J1058+0133	J1107–4449	J1130–1449	J1146+3958
J1147–6753	J1159+2914	J1229+0203	J1256–0547	J1337–1256
J1426+3625	J1427–4206	J1517–2422	J1550+0527	J1617–5848
J1642+3948	J1733–1304	J1751+0939	J1800+3884	J1924–2914
J2025+3343	J2056–4714	J2148+0657	J2202+4216	J2232+1143
J2353+1608	J2258–2758	J2357–5311		

Table C.3: ALMA bandpass calibrator (grid) sources

The ALMA system antenna noise properties are stable and well-known so that the a priori absolute flux density scale can be determined from the measured antenna SEFD (Source Equivalent Flux Density), and vary by less than 10% over the array elements. For example, the SEFD at Band 3 for most antennas is about 35. After multiplying the raw visibility amplitude (units of correlation with a typical value of 10^{-4}) and correcting by the system temperature, the visibility units are converted to antenna temperature in Kelvin. Then multiplication by the SEFD will produce visibility units in Jy. For example, with a system temperature of 100K, a 10^{-4} uncalibrated correlation amplitude would correspond to a correlated flux density of 0.35 Jy.

Thus, obtaining absolute amplitude accuracy better than 3% is challenging even at the lower frequencies, and 10% at the higher frequencies. Specialized ALMA observations of all of the solar system objects and quasars are done periodically in order to check the relative values of the brightness models, in collaboration with planetary model builders.

The *relative* flux density of a target source with respect to the assumed phase reference flux density is obtained by the normal phase referencing technique to an accuracy of about 1% over an experiment of one hour. Since most quasars vary less than a few percent over a few days, target variations over hours to days can be accurately measured as long as the same calibrator is used. For longer periods, an accurate absolute flux density of the calibrator must also be determined.

C.4.2 Polarization Calibrators

The measurement of the linear polarization of sufficiently strong calibrators (many are grid sources) are listed in the ALMA catalog. Additional ALMA observations are made in order to monitor the strongest, polarized, sources. The calibration methods are discussed in Section ().

C.4.3 Astrometric Calibrators

The discussion of the astrometric methods and observations for ALMA are discussed in Section 10.6.3. A good astrometric calibrator should have a position error of less than about 2 mas in order not to introduce astrometric and phase errors into the calibration measurements. This precision is only available from VLBI observations at cm wavelengths, but are currently known for over 2000 quasars. The use of a calibrator with less accurate radio positions should be avoided under most circumstances.

C.4.4 Check Sources

The quality of ALMA observations can be checked if a point source of known position and flux density is included in the SB, and observed and reduced in a similar manner as the target. This source is called a *check source* and is described in Section 10.6.3. A check source will be included when the observation resolution is $< 0.25''$ (longer spacings at Bands 3, 4, 6, and 7) and for all observations at Bands 8 and above. It will be selected from the ALMA catalog and should have the same properties as that of a suitable phase calibrator, but can be appreciably fainter.

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
EM	Electromagnetic
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
FWBN	Full Width Between the Nulls
FXF	Filtering, Correlation, and Fourier transform type correlator
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
JCMT	James Clerk Maxwell Telescope
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
SSR	Science Software Requirements
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable 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

Index

- uv*-coverage, 86
- uv*-plane, 22
- uv*-taper, 90
- 12-m Array, 9, 12
- 64-input Correlator, 9, 12, 57, 58
- 7-m Array, 9, 11, 13, 62

- ACA Correlator, 9, 60
- Airy function, 19
- Alcatel Alenia Space European Industrial Engineering
MT Aerospace (AEM), 177
- Aliasing, 28, 59
- ALMA
 - Archive, 173
 - Common Software (ACS), 155
 - Frontend Archive (AFA), 173
 - location, 11
 - Science Archive (ASA), 173
 - Science Data Model (ASDM), 157
- ALMA Common Software (ACS), 11
- ALMA Project Data Model (APDM), 103
- ALMA Regional Centre (ARC), 173
- ALMA Science Data Model, 135
- ALMA subsystem, 155
 - Control, 155
 - Executive, 155
 - QuickLook, 155
 - Scheduler, 155
 - Telescope Calibration, 155, 156
- Amplitude, 15, 22
- Amplitude Calibration Device, 140
- Amplitude Calibration Device, 137
- Amplitude Calibration Device (ACD), 31, 184
- Angular resolution, 82, 86
- Antenna
 - delay offset, 139
 - position, 139
 - power response, 19
- Antenna beam, 138
- Antenna(s)
 - beam size, 17
 - CM, 12
 - design and properties, 177
 - foundations, 11, 179
 - pads, 179
 - panels, 178
 - performance specifications, 12, 177
 - PM, 12
 - power response, 19
 - response, 16
 - specifications, 177
 - surface, 177
 - transportation, 181
 - transporters, 11
 - types, 178
- AOS Technical Building, 191
- Aperture, 17
 - synthesis, 19
- Application Specific Integrated Circuit, 62
- Array
 - performance, 152
- Array Operations Site (AOS), 11
- Astronomer on Duty, 103, 150
- Atacama Compact Array, 9, 11
- Atacama Large Millimeter/Submillimeter Array, 9
- Atmosphere
 - opacity, 125
- Atmospheric calibration (ATM), 185
- Atmospheric transmission, 31
- Auto-correlation, 20
- Autocorrelation
 - data, 136

- Barlett weighting, 66
- Baseband(s), 12, 14, 57, 71
 - desired center, 73
- Baseline, 20
- Baseline Correlator, 57
- Beam shape, 86
- Bessel function, 19
- Binary Data Format, 67
- Brightness, 16
 - temperature, 28
- Brightness temperature, 28
- Bulk binary data, 157

- Calibration
 - amplitude, 109
 - antenna pointing, 141
 - atmospheric, 108
 - bandpass, 107
 - baseband, 142

- flux, 107
- flux density, 142
- gain, 142
- phase, 107
- plan, 151
- polarization, 142
- receiver temperature, 140
- summary, 150
- system temperature, 140
- system-defined, 106
- user-defined, 106
- water vapor radiometer, 140
- Calibration observation, 105
- Calibrations
 - a priori, 137
 - offline, 137
 - online, 137
- Cartridge assembly, 31
- Cassegrain focus, 178
- Chajnantor, 11
- Chajnantor Plateau, 10
- Check source, 145
- Cittert-Zernike theorem, 22
- Clock rate, 60
- Complex visibility
 - function, 135
- Complex visibility, 22
 - distribution, 22
- Configurations, 82
- Contact scientist, 152
- Correlator
 - 64-input, 12
 - ACA, 12, 57
 - capabilities, 64
 - dump time, 67
 - FX, 13, 60
 - XF, 12
- Correlator and Integration Processor, 60
- Correlator Data Processor, 60
- Correlator delay model, 135
- Correlator mode
 - FDM, 191
- Cosmic Microwave Background, 126
- Cross-correlation, 20
- Cryostat, 181
- Cycle 0, 9
- Data
 - delegation, 175
 - delivery, 174
 - download, 176
 - replication, 174
- Data and Control Flow, 155
- Data Capture module, 157
- Data Management Group, 149
- Data rate, 12
 - average, 173
 - peak, 173
- Data receivers, 64
- Data Reducers Team, 151
- Deconvolution techniques, 24
- Delay
 - artificial, 20
 - geometrical, 20
- Digital Transmission System Receiver, 60
- Digital Transmitter, 64
- Digitizers, 64
- Director's Discretionary Time, 174
- Dirty
 - beam, 24
 - image, 24
- Dirty beam, 82
- Doppler correction, 72
- Doppler shift, 189
- Doppler tracking, 79
- DTS and FFT processor, 60
- Dump rates, 13
- Early Science
 - Operations, 9
 - Primer, 9
- Earth rotation, 23
- Electromagnetic power, 16
- Ephemeris file, 79, 113
- Execution Block (EB), 105
- Feathering, 98
- Field of view, 23, 93
- Field programmable gate array, 59
- FITS format, 176
- Flux density, 16, 28
- Focus model, 139
- Fourier transform, 15
- Fraunhofer approximation, 19
- Frequency
 - center offset, 73
 - definitions, 73
 - generation, 191
 - rest, 73
 - sky, 73
 - synthesizer, 72
- Frequency and phase standard, 193
- Frequency Division Mode (FDM), 58, 65
- Frequency Profile Synthesis, 60
- Fringe, 22
 - rotation, 189
 - spacing, 23
 - tracking, 191
- Frontend, 31, 181

- Full Width Between Nulls, 18
- Full Width Between the Nulls, 18
- FXF, 58
- Gibbs phenomenon, 90
- Grid source, 107
- Group Observing Unit Set (GOUS), 105
- Half Power Beam Width, 17
- Hanning function, 65
- Helpdesk, 152
- Image fidelity, 24
- Imaging
 - simultaneous, 73
- Integration time, 96
- Interferometer
 - resolution, 26
 - two-antenna, 20
- Interferometry
 - single field, 107
- Intermediate frequency, 33, 71
- Jansky, 16
- Joint ALMA Observatory (JAO), 11
- JPL Horizon format, 113
- Koh-ichiro Morita, 14
- Lasers, 193
- LO and IF systems, 189
 - operations, 190
- Local Oscillator, 33, 57, 71
 - offsetting, 34, 77
- Local Oscillator (LO), 12
- Long Term Accumulator, 60
- Lower sideband, 71
- Maximum data rate, 67
- Maximum recoverable scale, 13, 26, 92, 94
- Measurement Set, 141
- Member Observing Unit Set (MOUS), 105
- Metadata, 157
- Metrology, 179
- Mitsubishi Electric Corporation (MELCO), 177
- Morita Array, 9
- Mosaicing, 23, 93, 108
- Multiple pointings, 108
- Multiple Region Modes, 113
- NED, 175
- Next Generation Archive System, 173
- Nutator, 14
- Observations
 - on-the-Fly Mapping, 108
 - total power, 108
- Observing
 - project, 103
 - proposal, 103
 - script, 103, 106
 - session, 105
- Observing mode(s), 9
- Observing session, 155
- Observing Unit Set, 153
- Observing Unit Set (OUS), 105
- On-the-fly scan (OTF), 13
- Operation Support Facility (OSF), 11
- Otho mode transducer, 35
- Output radiance, 125
- Pampa la Bola, 11
- Parallactic angle, 112
- Phase, 15, 22
- Phase 1, 71, 103
- Phase 2, 71, 103
- Phase calibration, 76
- Phase center, 21
- Phase switching, 79
- Pipeline, 10
- Plateau de Bure, 16
- Pointing model, 138
- Polarization, 16, 78, 109
 - instrumental, 109
 - leakage, 143
 - off-axis, 110
- Position switching, 13
- Precipitable water vapor, 64
 - octiles, 125
- Precipitable water vapour
 - octiles, 126
- Primary beam, 17, 93
 - correction, 23
- Primary focus, 178
- Primer for Early Science, 11
- Program Management Group, 149
- Proposer's Guide, 9, 105
- Proprietary period, 174
- Quality Assurance, 149
 - categories, 150
 - classification, 149
 - report, 152
- Quantization, 57, 68
 - efficiency, 68
- Radiometer
 - Dicky-switching, 187
- Rayleigh resolution, 18
- Rayleigh-Jeans limit, 28
- Receiver(s)

- 2 sideband, 128
 - double sideband, 128
 - heterodyne, 128
- Rectangular field, 108
- Reflector
 - nutating, 177
 - secondary, 19
- Robustness, 87

- San Pedro de Atacama, 11
- Scan, 105, 135, 156
- Scheduling Block, 150
- Scheduling Block (SB), 103
- Science
 - goal, 103, 108
 - parameter(s), 103
- Science Archive Content Manager, 150
- Science Pipeline, 151
- Science Software Requirements, 105
- Sensitivity, 74, 125
 - point-source, 130
 - surface brightness, 131
- Shadowing, 26, 85
- Shutter, 179
- Sideband
 - double, 32, 77
 - double (DSB), 12
 - dual, 12, 32
 - single, 32
- Sidelobes, 18, 24, 81, 90
- Simbad, 175
- Sky brightness distribution, 15, 22
- Sky frequency, 71
- Snapshot, 90
- Solar observations, 115
- Solar system objects, 113
- Spatial
 - filtering, 24, 92
 - frequency, 15, 21
- Spatial coherence function, 107
- Specific intensity, 16, 28
- Spectral setup, 71, 73
- Spectral window (SPW), 59, 71
- Standard observing modes, 151
- Standard Test Environment, 155
- Stokes parameters, 109
- Sub-arrays, 12
- Submillimeter Array (SMA), 16
- Subreflector, 14, 178
- Subscan, 156
- Subsystems
 - QuickLook, 105
 - Control, 106
 - Scheduler, 105
 - Telescope Calibration, 105
- Surface characteristics, 138
- Synthesized beam, 24

- Technical Handbook, 9
- Telescope
 - filled-aperture, 20
- Telescope Monitor and Configuration Data Base, 138
- Telescope operator, 155
- Temperature
 - sky, 125
 - system, 125
- Tiltmeters, 179
- Time Division Mode (TDM), 58, 65
- time resolution, 66
- Tools
 - AQUA, 150, 152
 - JIRA, 149
 - Monitor and Control, 150
 - Project Tracker, 150
 - QuickLook, 150
 - TelCal, 150
- Total Power Array (TP Array), 9, 11, 13, 14, 62
- Transporter, 181
- Tunable Filter Bank (TFB), 12, 58, 189
- Tuning range, 76

- Upper sideband, 71
- User Policies, 152

- Van Vleck relationship, 68
- Velocity definition
 - Optical, 79
 - Radio, 79
 - Relativistic, 79
- Velocity reference, 79
 - Barycentric, 79
 - LSRK, 79
 - Topocentric, 79
- VertexRS, 177
- Virtual Observatory, 175
- Visibility, 15
- Visibility function
 - weighted, 81
- Vizier, 175
- Voltage response, 19
- VOTables, 175

- Walsh-switching, 57, 79
- Water Vapour Radiometer (WVR), 12, 64, 186
- Weblog page, 151
- Weighting
 - Briggs, 87
 - natural, 81, 86
 - uniform, 81, 87

WVR Correction, 64

WVR correction, 57

XML tables, 157

Zero spacing, 108

Zero-point fluctuations, 127

Zero-spacing problem, 13, 26



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.

