

ALMA Cycle 1 Technical Handbook

Andreas Lundgren (Editor)



www.almascience.org

ALMA, an international astronomy facility, is a partnership of Europe, North America and East Asia
in cooperation with the Republic of Chile.

User Support:

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

Revision History:

Version	Date	Editors
0.1	2011-10-24	B. Vila Vilaro, A. Lundgren,
0.2	2011-11-02	A. Lundgren
0.3	2011-12-02	A. Lundgren, T. Kamazaki, M. Saito
0.4	2012-02-14	A. Lundgren, Y. Kuroono
0.5	2012-02-28	A. Lundgren, B. Vila Vilaro
0.6	2012-03-16	A. Lundgren
0.7	2012-04-08	A. Lundgren, A. Biggs, R. Kneissl, S. Leon
0.8	2012-05-02	A. Lundgren, B. Dent, A. Biggs, S. Corder, S. Leon, B. Vila Vilaro, R. Kneissl
0.9	2012-05-15	A. Lundgren, A. Biggs, P. Andreani, B. Dent, L. Nyman, S. Leon, R. Kneissl, B. Vila Vilaro
1.0	2012-05-28	A. Lundgren

Contributors

Paola Andreani, Andy Biggs, Crystal Brogan, Stuartt Corder, Bill Dent, Darrel Emerson, Ed Fomalont, John Hibbard, Todd Hunter, Takeshi Kamazaki, Rüdiger Kneissl, Yasutaka Kuroono, Harvey Liszt, Stephane Leon, Carol Lonsdale, Nuria Marcelino, Jeff Mangum, Koh-Ichiro Morita, Lars-Ake Nyman, Sachiko Okumura, David Rabanus, Mark Rawlings, Masao Saito, Scott Schnee, Kimberly Scott, Kartik Sheth, Kenichi Tatematsu, Baltasar Vila-Vilaro.

In publications, please refer to this document as:

A. Lundgren, 2012, ALMA Cycle 1 Technical Handbook, Version 1.01, ALMA

Contents

1	Introduction	13
2	Receivers	15
2.1	Local Oscillators and IF ranges	17
2.2	The Cycle 1 Receivers	18
2.2.1	Band 3 receiver	18
2.2.2	Band 6 receiver	18
2.2.3	Band 7 receiver	19
2.2.4	Band 9 receiver	25
3	The Correlators	31
3.1	Digitizers	31
3.2	The 64-antenna correlator	32
3.2.1	FDM mode	32
3.2.2	TDM mode	32
3.2.3	Correlation and realtime processing	34
3.3	The 16-antenna ACA correlator	34
3.4	Correlator data processing	36
3.4.1	Integration time intervals, channel-average and spectral data	36
3.4.2	Online WVR correction	40
3.4.3	Final data product - the ASDM	40
3.5	Spectral resolution, smoothing and channel averaging	40
3.5.1	Channel averaging	41
3.6	Correlator speed and data rates	42
4	Spectral setups	43
4.1	Spectral setups for multiple lines	43
4.1.1	Calculating optimum continuum frequencies	44
4.2	Spectral setups for lines near the edge of the bands	45
4.3	Spectral setup in Band 9	45
4.4	Usable bandwidth	46
4.5	Spurious signals	46
4.6	Doppler setting	46
4.7	Limitations and rules for spectral setups in Cycle 1	47
5	Description of Atacama Compact Array in full operations and in Cycle 1	49
5.1	ACA during full operations	49
5.1.1	Concept and purpose of the ACA	49
5.1.2	System & operations description	49
5.1.3	7-m Array configurations in full operations	50
5.2	ACA 7-m and TP Arrays during Cycle 1	50
5.2.1	Capabilities of ACA in Cycle 1	50
5.2.2	Assessing the requirement for ACA data	51

5.2.3	Data processing	51
6	Cycle 1 Configurations	53
6.1	Introduction	53
6.2	Maximum Angular Scales (MAS)	53
6.3	The 12-m Array configurations	54
6.3.1	Beam Maps	55
6.4	The ACA and 12-m Array combination	55
7	Simulations of Cycle 1 observations	65
7.1	Introduction	65
7.1.1	The ALMA Simulation Tools	65
7.1.2	Basic steps of an ALMA simulation	66
7.2	Simulation Examples for Cycle 1 Configurations	68
7.2.1	Beam Maps	68
7.2.2	An M51-like Galaxy	69
7.2.3	A Toy Model Example	69
7.2.4	Combined ACA + 12m Observations	69
7.3	Sensitivity to Angular Scales	74
7.3.1	Filter Functions	74
7.3.2	Resolved Out Flux	74
7.4	Image Quality	74
7.4.1	Image Fidelity	74
7.4.2	Dynamic Range	74
8	ALMA Sensitivity Calculator	79
8.1	Calculating system temperature	79
8.1.1	Receiver temperatures	80
8.1.2	Sky temperature	81
8.1.3	Ambient temperature	82
8.2	The sensitivity calculation	82
8.2.1	12-m and 7-m Arrays	82
8.2.2	Total Power Array	83
8.3	User Interface	83
9	Observing Projects and Logical Data Structure	85
9.1	Observing Project Structure	85
9.2	Program Execution	87
9.3	Structure of an SB and associated scripts	87
9.3.1	Observing Groups	87
9.3.2	The Standard Interferometry Script	87
9.4	Data and Control Flow	89
9.5	The ALMA Science Data Model (ASDM)	91
10	Calibration and Calibration Strategies	93
10.1	Long-Term Effects	93
10.1.1	All-Sky Pointing	93
10.1.2	Focus Models	94
10.1.3	Baseline Vectors	94
10.1.4	Cable Delay	94
10.1.5	Antenna Characteristics	94
10.2	Short-Term Effects	95
10.2.1	Offset Pointing	95
10.2.2	Bandpass	95
10.2.3	WVR Corrections	96

10.2.4	Gain (Amplitude & Phase)	96
10.2.5	System and receiver temperature	97
10.2.6	Sideband Ratio	97
10.2.7	Flux	97
11	Quality Assurance	99
11.1	QA0	99
11.2	QA1	100
11.3	QA2	100
11.4	QA3	101
11.5	The Quality Assurance Report	101
11.6	Pass/Fail Criteria	101
12	Data Archiving	103
A	Appendix	107
A.1	Antennas	107
A.2	Antenna Foundations	110
A.3	Antenna Transportation	110
A.4	Cryostat	111
A.5	Amplitude calibration device	115
A.5.1	Atmospheric Calibration Procedure	116
A.6	Water Vapor Radiometers	117
A.7	The LO and IF System	120
A.7.1	Summary of operation	120
A.7.2	LO/IF components	122
A.7.3	Delay corrections, sideband suppression and interference rejection	129
B	Acronym List	131

List of Figures

2.1	ALMA Bands for Cycle 1 are shown in red superimposed on a zenith atmospheric transparency plot at the AOS for 0.5 mm of PWV.	16
2.2	IF ranges for the two sidebands in a heterodyne receiver.	17
2.3	Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the dashed line, and the Cassegrain focus is shown by the small circle to the upper right. A new version of this Figure is being prepared by the FE group.	19
2.4	Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges are constructed in Canada at NRC-HIA, Victoria.	20
2.5	Band 3 zenith transmission for 1, 5 and 15mm of pwv. Frequency is in GHz.	21
2.6	Typical system temperature (T_{sys}) at zenith for Band 3 with 1mm 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.)	21
2.7	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.	22
2.8	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 x3 multiplier inside the cryostat. The Band 6 cartridges are 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).	23
2.9	Band 6 zenith transmission for pwv=0.5, 1 and 5mm. Frequency is in GHz.	24
2.10	Typical T_{sys} at zenith for Band 6 with 1mm pwv, based on measured values of the receiver temperatures.	24
2.11	Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid. The band 9 cold optics arrangement is very similar this, although the actual numbers differ.	25
2.12	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.	26
2.13	Band 7 atmospheric zenith transmission for pwv=0.5, 1.0 and 5.0mm. 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.	27
2.14	Typical T_{sys} at zenith for Band 7 with pwv=1mm.	27
2.15	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.	28
2.16	Band 9 zenith transmission for pwv = 0.2, 0.5 and 1mm. Frequency is in GHz.	28

2.17	Typical T_{sys} at zenith for Band 9 with pwv=0.5mm. (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.)	29
3.1	Overview diagram of the ALMA digitizers, data transmission system and 64-antenna correlator. The digitised data from the individual basebands are fed to the TFBs and station cards (top right) for the parallelisation and processing required on each antenna datastream. The matrix of correlator boards (lower half) does the cross-correlation between antenna stations, and the resulting data is then integrated in the Long-term Accumulators (LTAs, lower right).	33
3.2	Correlator board for the 64-antenna correlator, with the 8x8 array of custom correlator chips, each with 4k lags. There are 128 such boards per correlator quadrant, adding up to a total of 32768 chips.	35
3.3	Overview diagram of the ACA correlator. The signals from the BackEnd Subsystem are processed in the DFP (Digital FFT Processor) and CIP (Correlation Integration Processor) modules and output to the Computing Subsystem (see text and Figure 3.4 for details of these). The modules are controlled by the MCI (Monitor and Control Interface), 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-antenna correlator.	35
3.4	Block diagram of ACA correlator. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. A Fast Fourier Transform (FFT) is performed using FPGA modules, resulting in signals in the spectral domain which can then be correlated. The auto and cross-correlation data are then accumulated in time, and the parallel streams averaged together in the CIP (Correlation and Integration Processor) module. The correlation spectra are then fed to the ACA CDP for further accumulation and processing.	37
3.5	Data capture sequence for the ACA correlator. A spectral window is available to select a frequency range from 512K spectra, whose total bandwidth is 2 GHz and frequency resolution is 3.815 kHz, in the Cycle 1 mode. Selected spectra are integrated and then sent to the output.	37
3.6	A quadrant of the ACA correlator. The left and middle racks compose a quadrant, which processes a baseband pair from sixteen antennas. The right rack is one of four computing racks, which have correlator control, post processing computers, network instruments and so on.	38
3.7	Basic data processing and accumulation steps in the correlator and archive. The timing intervals shown are described in the text.	39
4.1	Illustration of a frequency setup, based on the OT spectral display. Yellow areas are the IF ranges, red lines are the 2 GHz-wide basebands (bb0-3), and blue crosses represent four spectral windows (spw 0-3, one per baseband). The frequency of LO1 is shown by the central vertical line. The blue hashed area shows the possible tuning range of the frontend, and the curved line gives an indication of the atmospheric transmission. The molecules and transitions of interest are shown, and their frequencies are shown by the position of the vertical blue lines. The bandwidths of the spws are illustrated by the horizontal bars of the blue crosses.	44
4.2	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).	46
6.1	Recovered flux of a uniform disk, with a total flux of 1 Jy, observed at 100 GHz with three different 12-m Array configurations, as a function of the disk size.	56
6.2	Antenna positions for the 6 12-m Array Configurations with the corresponding pad names.	57
6.3	UV radial density distribution for the most compact configurations including the ACA	58
6.4	UV radial density distribution for the most extended 12-m Array configurations.	58
6.5	Spatial resolution in arcsec at different frequencies for the six 12-m Array configurations.	59

6.6	Shadowing fraction vs. Declination for the two most compact configurations of the 12-m Array and for the 7-m Array with a track duration of 2 hours (± 1 h HA).	59
6.7	uv-coverage for the two most compact (C32-1 and C32-2) and the two most extended (C32-5 and C32-6) 12-m Array configurations with a 2 hour integration (± 1 h HA) of a source at a declination of -30°	60
6.8	Synthesized beam radii vs. Declination for an observation of 2 hours (± 1 h HA) at 100 GHz. The 6 12-m Array configurations are shown with different colors starting with the most compact configuration (C32-1) in blue at the top, progressing through the more extended configurations. The width of the area gives an indication of the ellipticity of the beam (the wider the area, the more elliptical the beams are).	61
6.9	Maps of the synthesised beams for the six representative ALMA Cycle 1 configurations C32-1, C32-2, C32-3, C32-4, C32-5, C32-6 and the ACA 9 antenna configuration. Note: the axis (offset angle in arcseconds) vary between the different panels to show areas scaling roughly with the synthesized beam width. All maps are for Band 7 at 300 GHz, Dec= -35° , and an hour angle range of ± 1.5 hours	62
6.10	Antenna positions for the 7-m Array with the corresponding pad names.	63
7.1	CASA <code>sim_observe</code> and <code>sim_analyze</code> example showing in the top row 1) the antenna positions of the configuration, 2) the observed tracks in uv, 3) the input model image, and in the bottom row 4) the input model convolved with the Gaussian beam width of the configuration, 5) the observed and <code>cleaned</code> image, and 6) the difference between convolved input and observed image as outputs. See Section 7.1.2 for details on simulation input and parameters. The axis units are for 1)+2) in kilo-lambda (here for 1 mm observations equal to meters) and for 3)-6) in arcsec (offset angle). The brightness units are for 3) in Jy/pixel and for 4)-6) in Jy/beam.	67
7.2	Example of <code>clean</code> with the multi-scale option on the simulated dataset shown in Figure 7.1. The image quality is vastly improved over the reconstruction shown in panel 5). See Section 7.1.2 for details on simulation input and parameters.	68
7.3	Maps of the synthesized beams for two simulated observation examples in Band 7 at 300 GHz, Dec= -35° , and an hour angle range of ± 1.5 hours in the ACA 9 antenna configuration and in the C32-5 configuration. Note that the axis (offset angle in arcsec) are different in the two images.	68
7.4	M51 H α image as input for the simulations (left) and the Fourier transform of the image in uv space (right).	69
7.5	Top to bottom rows are for the ACA (7-m array + total flux density measurement from the ACA 12-m total power antennas), C32-1, C32-2 and C32-3 configurations. Right to left panels: Gaussian beam convolved input model; simulated and cleaned image; difference (convolved model - simulated). The simulations are noiseless to show primarily the rms resulting from dynamic range limitations and the angular scale filtering properties of the different configurations. All simulations are for Band 7 at 300 GHz, Dec= -35 deg, and an hour angle range of ± 1.5 hours. Brightness units are in Jy/beam.	70
7.6	Top to bottom rows are for the C32-4, C32-5 and C32-6 configurations. Right to left panels: Gaussian beam convolved input model; simulated and cleaned image; difference (convolved model - simulated). The simulations are noiseless to show primarily the rms resulting from dynamic range limitations and the angular scale filtering properties of the different configurations. All simulations are for Band 7 at 300 GHz, Dec= -35 deg, and an hour angle range of ± 1.5 hours. Brightness units are in Jy/beam.	71
7.7	The top panel represents the true sky brightness distribution: 4 Gaussians and 4 point sources (circled in black). The middle left and right panels show observations of the sky brightness distribution with the ACA and C32-1 arrays. The bottom left and right panels show observations of the sky brightness distribution with the C32-3 and C32-6 arrays. The simulation parameters are identical to the M51 simulation, i.e. Band 7 at 300 GHz, Dec= -35 deg, and an hour angle range of ± 1.5 hours. The brightness units are Jy/pixel for the input and Jy/beam for the output image. The top three panels show the same sky area, whereas the bottom two panels are zoomed into the center area.	72

7.8	M51 input image (left) compared to a simulated image synthesized with data from combining ACA 7-m and 12-m Array C32-4 configurations (right). The brightness units are Jy/pixel for the input (and 'inbright', the model peak brightness, had been set to 0.1 Jy/pixel) and Jy/beam for the output image.	73
7.9	C32-2 without the ACA 7-m array (left), Gaussian beam convolved input model (center), C32-2 with the ACA 7-m array (right).	73
7.10	Top: Sensitivity to angular scales (in relative units) for the different configurations C32-6 to C32-1 and ACA, from left to right, based on the uv tracks of the Dec = -35 deg, 3 hours (± 1.5 hours), observation. The 25% and 95% percentiles in the visibility (baseline) scale distribution are shown at the top of this panel. Bottom: Intensity of the structure models (M51 H α in red, toy model in blue) as a function of angular scale.	75
7.11	Gaussian functions as input models for the sky brightness distribution are observed with C32-1 at 100 GHz (DEC=-30deg, HA= ± 1 h). For a FWHM of 20 arcsec (left panel) in this noiseless simulation 96% of the integrated flux density can be recovered in this cleaned image, and for a FWHM of 40 arcsec (right panel) 22% are recovered.	76
7.12	Fidelity image for the simulated observation of M51 with C32-1	76
7.13	Crowded field with point sources and extended (Gaussian) structure in configuration C32-6. In the 3h observation on the left all four point sources, and the 0.3 and 2 arcsec FWHM Gaussians are visible. The 2 arcsec Gaussian is difficult to recover, particularly in the 5 min observation on the right, because there are few short baselines in that configuration. The weakest point source is not detected (and potential artificial sources are created) in the 5 min observation due to dynamic range limitations.	77
8.1	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 towards the bottom of the plot. The water line at 183 GHz (band 5) is particularly prominent. In general, higher frequencies require drier observing conditions.	81
8.2	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.	84
9.1	Block diagram of an Observing Project from the point of view of the observation preparation (top) and actual execution (bottom). Note that an ObsUnitSet (OUS) can be composed of other ObsUnitSets as explained in the text. Each time an Scheduling Block is executed, CONTROL creates a new ExecBlock structure (see text).	86
9.2	Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.	89
9.3	Data flow components.	91
9.4	ASDM Tables. Outlined set of tables are the core ones (i.e., present in all ASDMs).	92
9.5	Current list of ASDM associated tables generated by TelCal.	92
12.1	Archive design (front-end & Back-end) at the OSF to store metadata and raw and monitor data.	104
12.2	Location of the ALMA archives.	104
12.3	Data flow from the AOS down to the ARCs.	105
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).	109
A.2	Structure of an antenna pad (actual pad at the OSF) (left) and detail of antenna anchored to a pad (right).	110
A.3	The ALMA array with eight 12-m antennas (left), and an antenna being transported to the AOS (right).	111
A.4	Side view of ALMA frontend showing cryostat assembly, with room temperature unit below.	112
A.5	Bottom view of ALMA frontend, showing WCAs.	113

A.6 Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view). 114

A.7 Lateral view of the ACD on top of the ALMA frontends. 115

A.8 Top view of an ALMA front end showing the robotic arm of the ACD retracted during normal observations or on top of one of the frontend inserts for calibration. The current design has been improved by placing all the loads in a wheel. 116

A.9 WVR filters superimposed onto the 183 GHz water vapor emission line. 117

A.10 Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR. 118

A.11 Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer. 119

A.12 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 in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized tuning steps, e.g. LO2 (which has steps of 125 MHz, with a factor of “N”, with an offset of “fts”), TFB LO (steps of 30.5 kHz, with a factor of “L”) and the Bulk Delay Correction (steps of 250 ps, with a factor of “P”). See text for descriptions of each stage. 121

A.13 Summary block diagram of the LO distribution system. 123

A.14 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 (in this case, x3) is used to generate the LO1 frequency, at 282.9 – 365.1 GHz. The photonic LO signal (green) feeds a photomixer which creates a beat signal between the ML and SL frequencies. This is mixed with the LO and feeds the PLL. The FLOOG generates a small offset frequency which is different for each antenna. See text for details. 124

A.15 Block diagram of IF Processor. 125

A.16 The DTX and DRX Signal Digitization and Transmission system. 126

A.17 LO block diagram, showing the Central LO (CLOA) and the LO section in the WCA in each frontend. For description of acronyms, see text. 126

A.18 The ALMA reference signals. 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. 128

A.19 Block diagram of the Line Length Corrector system for ALMA. 128

This page was intentionally left almost blank

Chapter 1

Introduction

The Atacama Large Millimeter/Submillimeter Array (ALMA) is an aperture synthesis telescope that will consist of at least 66 antennas arranged in a series of different configurations. It will operate over a broad range of observing frequencies in the millimeter and submillimeter regime. During Cycle 1 only a limited number of antennas, frequency bands, array configurations and observing modes will be available. Users should refer to the Capabilities section on the ALMA Science Portal at <http://www.almascience.org/> for the latest information.

This Technical Handbook describes the Cycle 1 setup of the ALMA system. It is intended to provide additional technical information for ALMA users, to a deeper level than what is described in the ALMA Early Science Primer, and to provide more information on the limitations of the Cycle 1 setups. Although it contains chapters relevant for the preparation of proposals, it should not be necessary to read it in order to prepare a proposal. The appendix includes technical information about the antennas, cryostats, amplitude calibration devices, local oscillator (LO) and intermediate frequency (IF) systems as well as generation and distribution of reference signals.

Chapter 2

Receivers

The ALMA frontend can accommodate up to 10 receiver bands covering most of the wavelength range from 10 to 0.3 mm (30–950 GHz). In Cycle 1, Band 3, 6, 7 and 9 will be available (see available frequency and wavelength ranges for these bands in Table 2.1). 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 2.1 and the specifications are listed in Table 2.1. The receivers are described in more detail by the references listed in Table 2.2.

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 freq. (K) ⁶
3	84.0-116.0/ 2.59-3.57	92 - 108	2SB	4-8	7.5	<41 ⁷	<45 ⁷
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
9	602.0-720.0/ 0.42-0.50	610 - 712	DSB	4-12	7.5(15) ⁵	<175 (DSB)	<261 (DSB)

Table 2.1: 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. **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 2.2.2) **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 4.4). **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. **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 $T_{\text{rx}} < 41$ K at LO=104 GHz, and $T_{\text{rx}} < 45$ K for any other valid LO setting. Both values should be the average over all four IFs and 4 GHz bandwidth.

The ALMA receivers in each antenna are situated in a single frontend assembly (see Appendix, Section A.4). The frontend assembly consists of a large cryostat containing the receiver cold cartridge assemblies (including SIS

Topic	Author/Year	Meeting proceedings	ADS identifier
B3	Claude et al. 2008	SPIE 7020	2008SPIE.7020E..33C
B6	Kerr et al. 2004	15th Intl Symp Space Terahertz Tech	2004stt..conf...55K
B7	Mahieu et al. 2005	16th Intl Symp Space Terahertz Tech	2005stt..conf...99M
B9	Baryshev et al. 2007	18th Intl Symp Space Terahertz Tech	2007stt..conf..164B
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 2.2: Conference proceedings papers describing the receiver bands, optics and the water vapor radiometer.

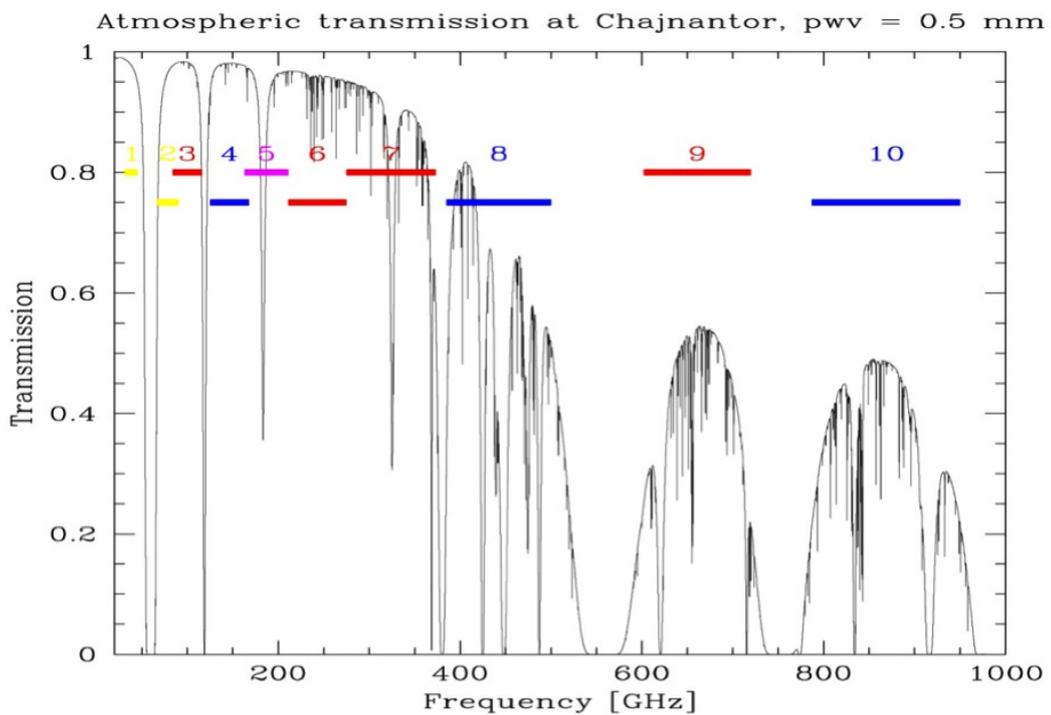


Figure 2.1: ALMA Bands for Cycle 1 are shown in red superimposed on a zenith atmospheric transparency plot at the AOS for 0.5 mm of PWV.

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. Each receiver cartridge contains two complete receiving systems sensitive to orthogonal linear polarizations. The designs of the mixers, optics, LO injection scheme, and polarization splitting vary from band to band, depending on the optimum technology available at the different frequencies; each receiver is described in more detail in the sections below. Table 2.1 summarizes the characteristics of the bands available in Cycle 1.

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 them 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.

2.1 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 IF systems are described in detail in the Appendix (Section 14.5).

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 and 7 have an output range of 4–8 GHz, Band 6 a range of 6–10 GHz and Band 9 a range of 4–12 GHz (Table 2.1). 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 2.1, and the IF bandwidth (per sideband) is $IF_{hi} - IF_{lo}$. This is illustrated in Figure 2.2. Note that the maximum IF bandwidth in Table 2.1 may be a few percent less than the IF range (see Section 4.4).

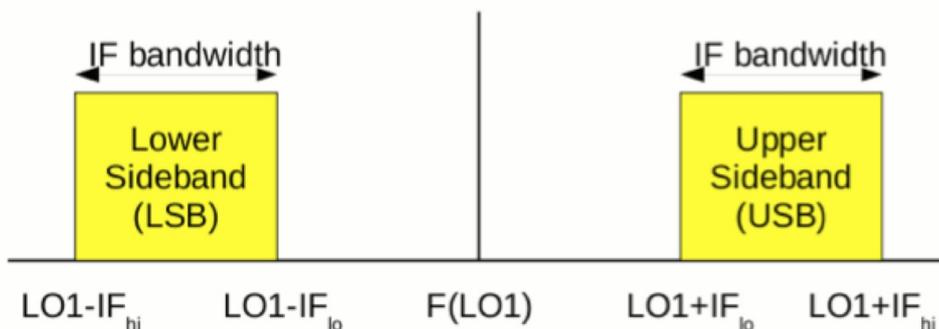


Figure 2.2: IF ranges for the two sidebands in a heterodyne receiver.

2.2 The Cycle 1 Receivers

The Band 3, 6 and 7 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 Section 5.4). 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 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, 6 and 7. In Cycle 1, only one sideband per spectral window will be 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 will be 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.

2.2.1 Band 3 receiver

Band 3 is the lowest frequency band available in Cycle 1, covering a frequency range of 84.0–116.0 GHz (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 2.3). 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 3 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 2.4. The Cold Cartridge Assembly (CCA) contains the cold optics, OMT, SIS mixers and the low-noise HEMT first IF amplifiers. At room temperature, the Warm Cartridge Assembly (WCA) includes further IF amplification and the Local Oscillator covering 92–108 GHz.

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 2.5) which means observations in Band 3 can, in principle, take place with 10 mm or more of pwv. The resulting system temperature (T_{sys}) shows the expected rise at the higher end, due to an atmospheric oxygen line (Figure 2.6).

2.2.2 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 2.7). 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 2.8.

The Band 6 IF frequency has been extended 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

¹Specifically, the 12CO/13CO/C18O 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

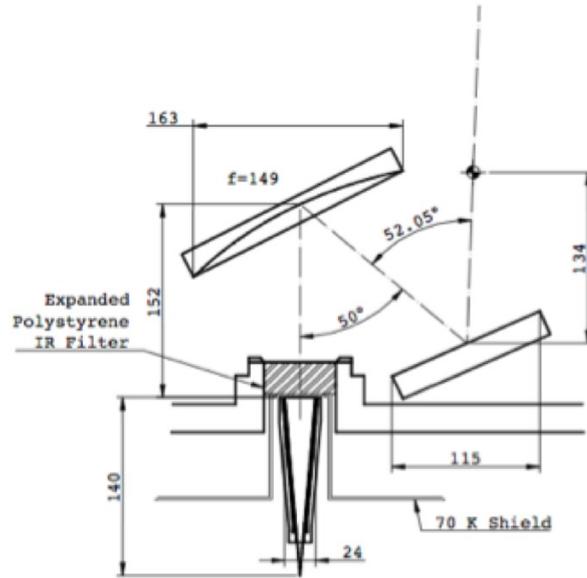


Figure 2.3: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the dashed line, and the Cassegrain focus is shown by the small circle to the upper right. A new version of this Figure is being prepared by the FE group.

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 width of the two basebands per polarization.

The atmospheric transmission in Band 6 is shown in Figure 2.9 for three typical pwv values. Most of the narrow absorption lines are from ozone. Most observations in Band 6 will be done with pwv <5mm.

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 1mm pwv are shown in Figure 2.10.

2.2.3 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, because the losses in the OMT are higher at these frequencies (Figure 2.11). A block diagram of the Band 7 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 2.12.

The atmospheric transmission in Band 7 is shown in Figure 2.13 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 1mm pwv are shown in Figure 2.14. Note that the atmospheric transmission (and hence T_{sys}) at frequencies below 300GHz is considerably better than that of the top half of band 7; in that respect the performance is closer to that of band 6.

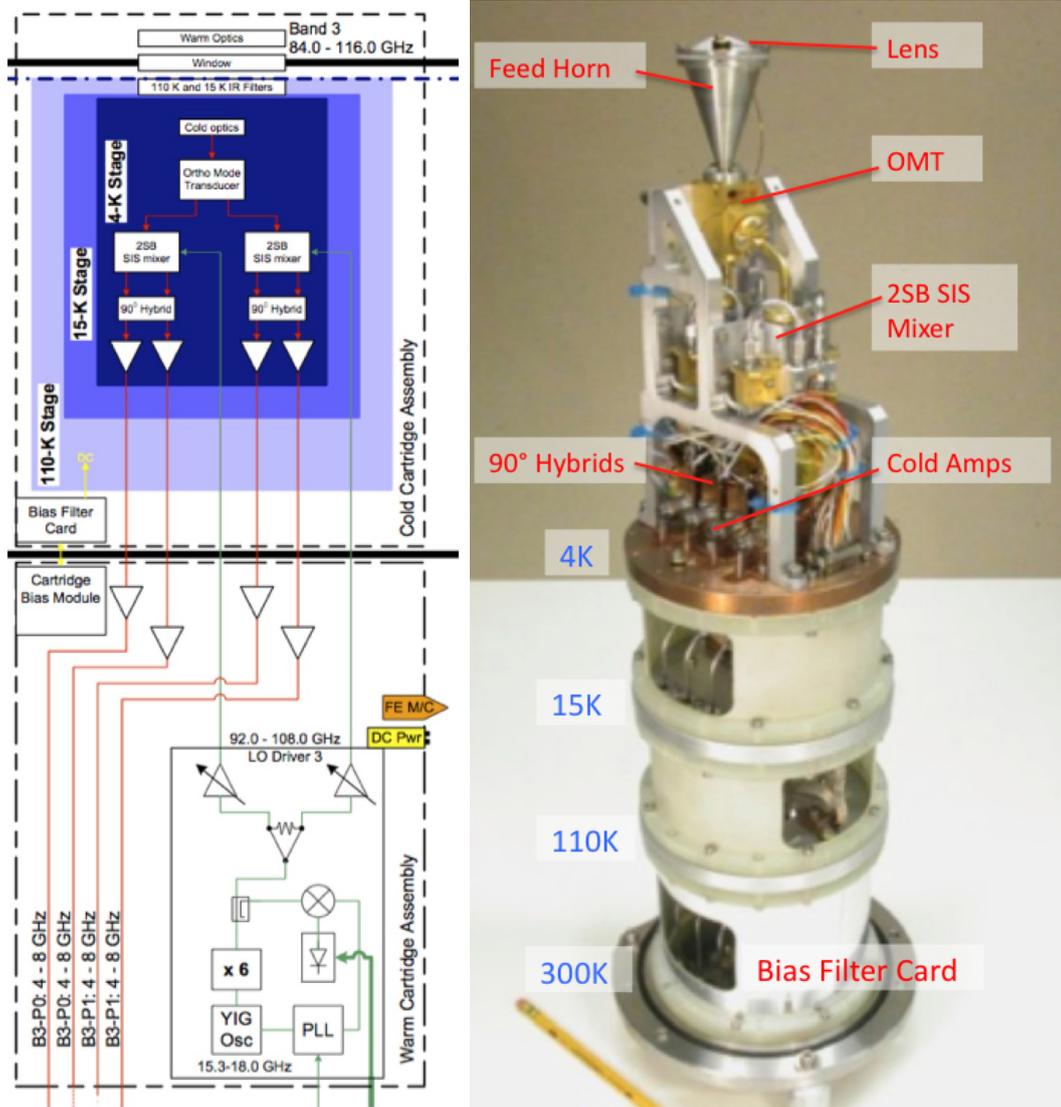


Figure 2.4: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges are constructed in Canada at NRC-HIA, Victoria.

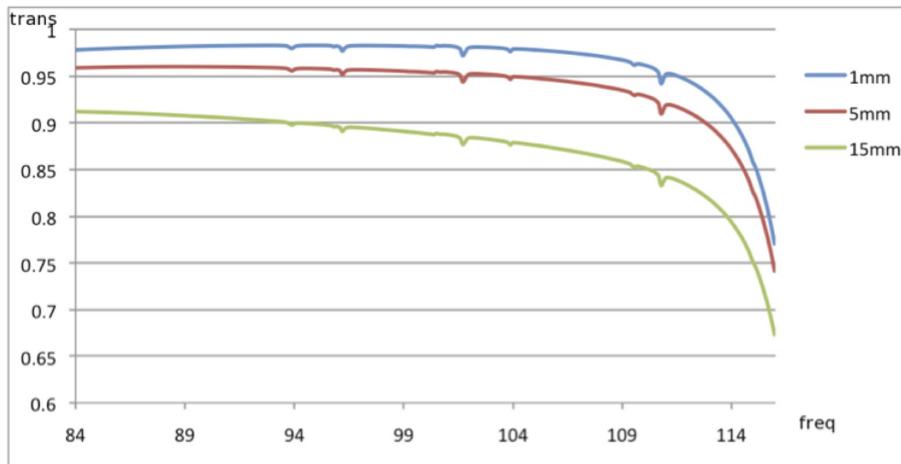


Figure 2.5: Band 3 zenith transmission for 1, 5 and 15mm of pwv. Frequency is in GHz.

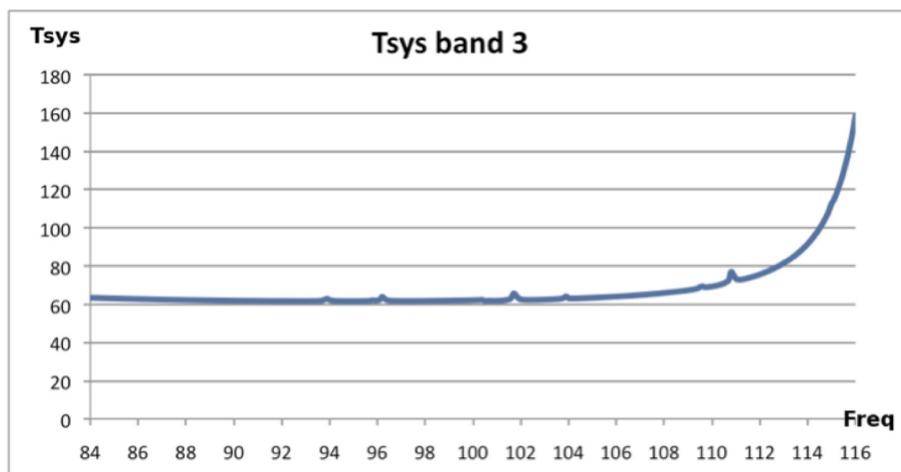


Figure 2.6: Typical system temperature (T_{sys}) at zenith for Band 3 with 1mm of pwv. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

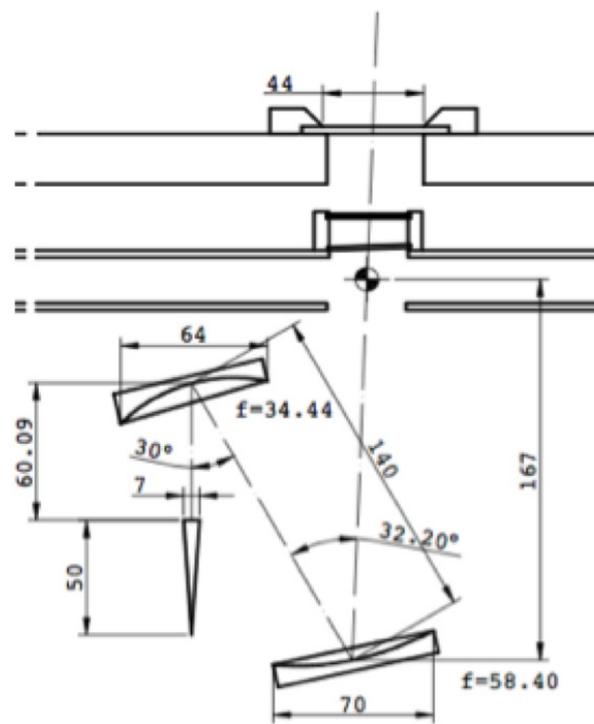


Figure 2.7: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocker. Note the slightly inclined inner window, designed to minimize standing waves.

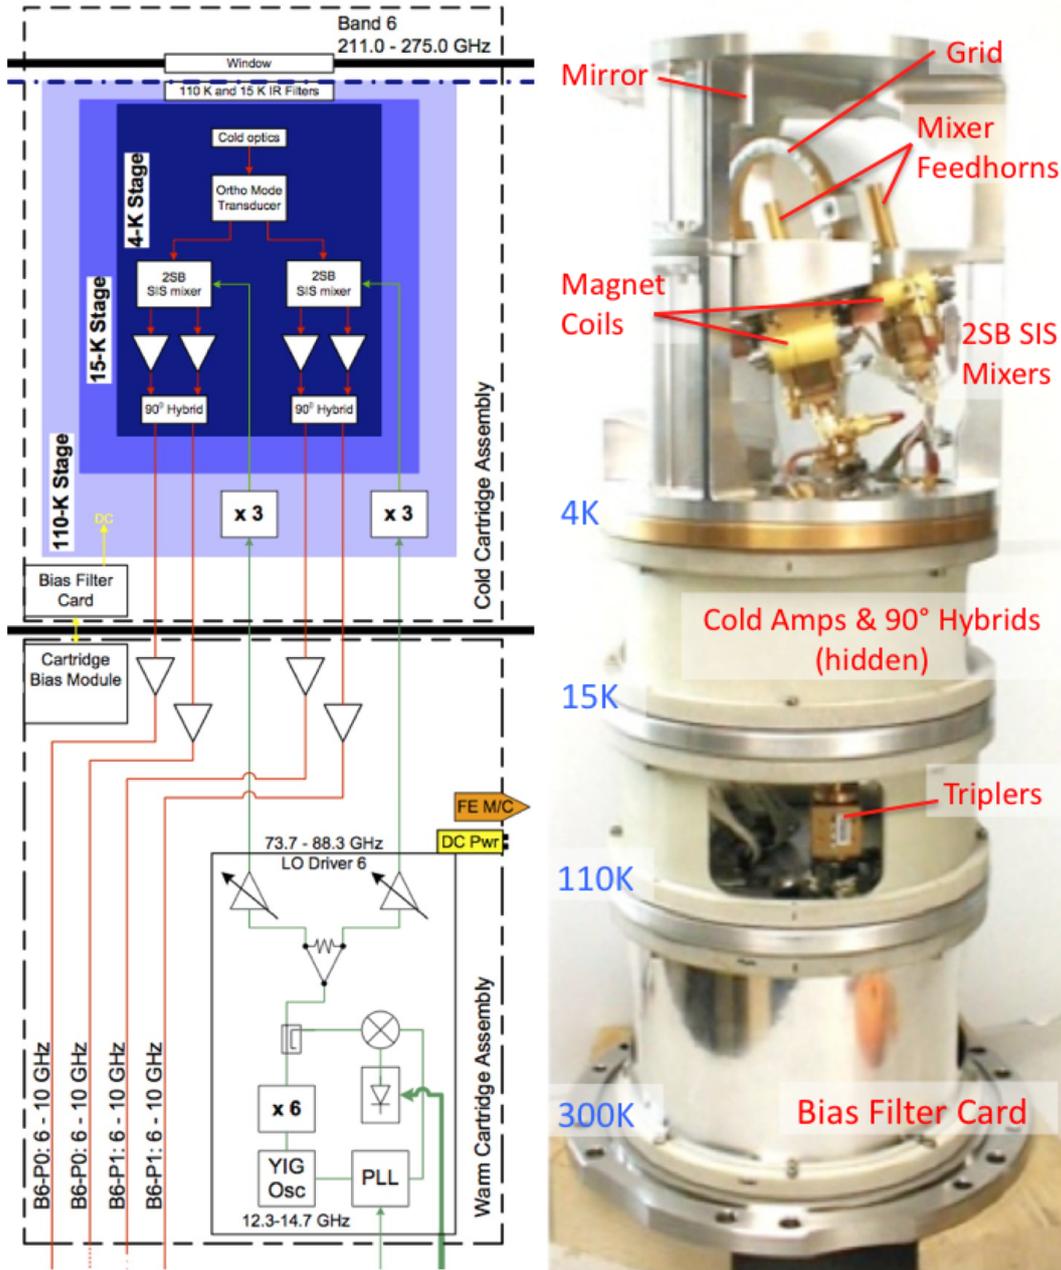


Figure 2.8: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra x3 multiplier inside the cryostat. The Band 6 cartridges are built at NRAO, Charlottesville. Note that the IF output range is actually 5-10 GHz. The range shown is the one recommended for continuum observations (see text).

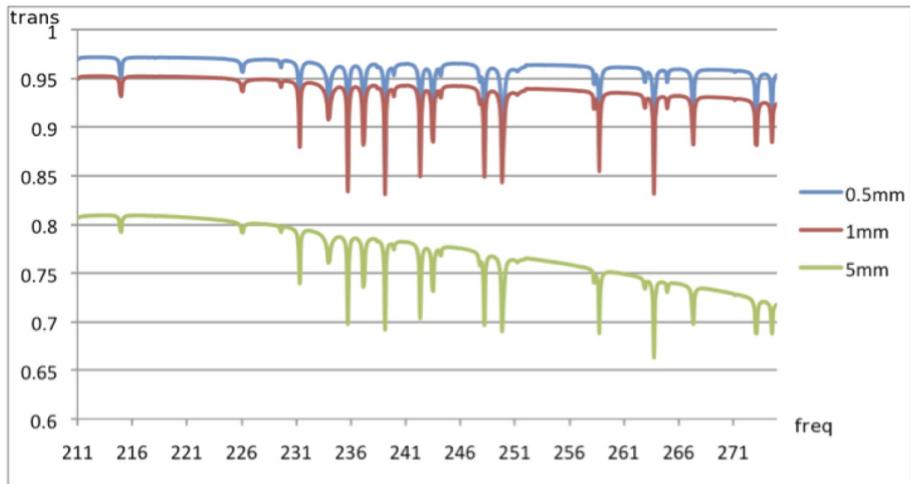


Figure 2.9: Band 6 zenith transmission for pwv=0.5, 1 and 5mm. Frequency is in GHz.

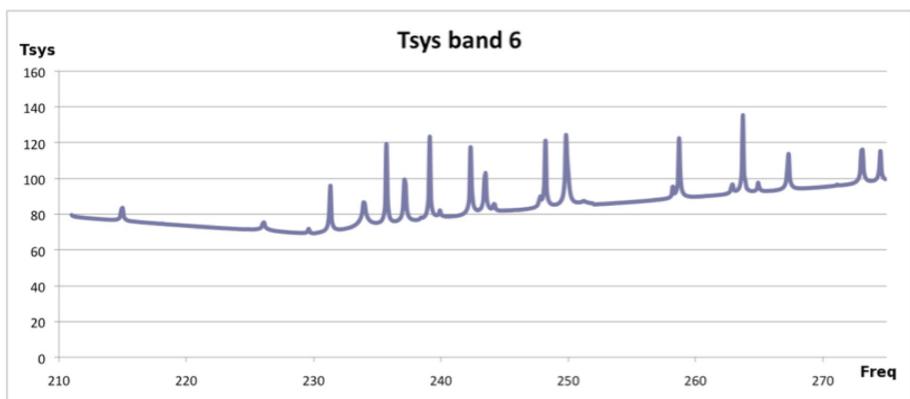


Figure 2.10: Typical T_{sys} at zenith for Band 6 with 1mm pwv, based on measured values of the receiver temperatures.

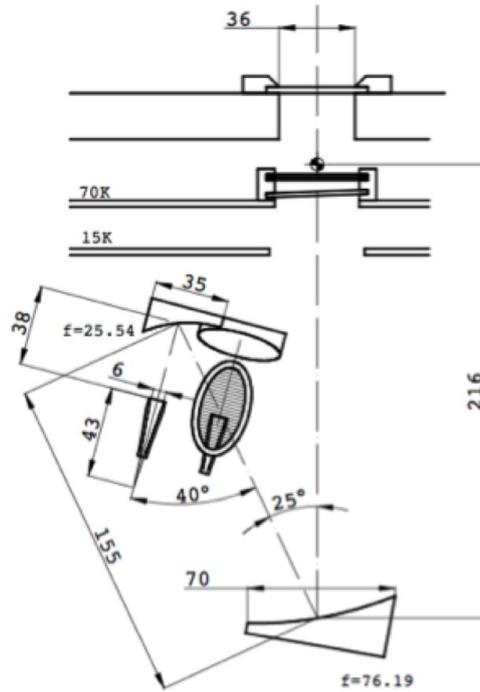


Figure 2.11: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid. The band 9 cold optics arrangement is very similar this, although the actual numbers differ.

2.2.4 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. 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 1, LO offsetting will be 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 which then washes out. The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 4.4), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 2.15.

The Band 9 atmospheric transmission is significantly dependent on the pwv, as illustrated in Figure 2.16 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 2.17 shows the expected T_{sys} for 0.5 mm of pwv, over most of the band given the expected receiver noise. 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.

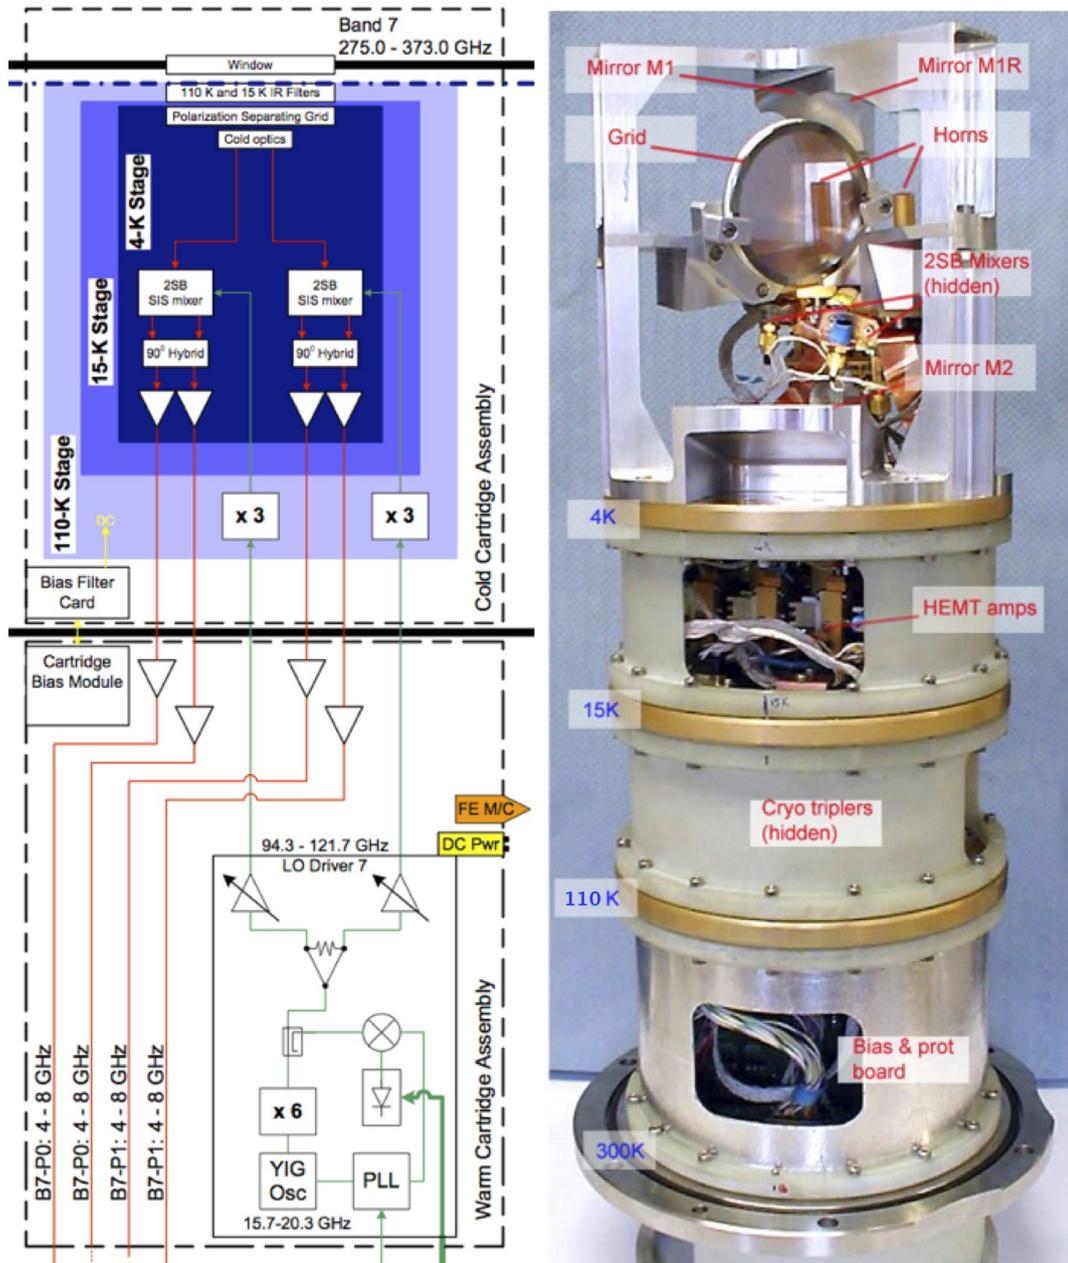


Figure 2.12: Band 7 frontend receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built at IRAM in France.

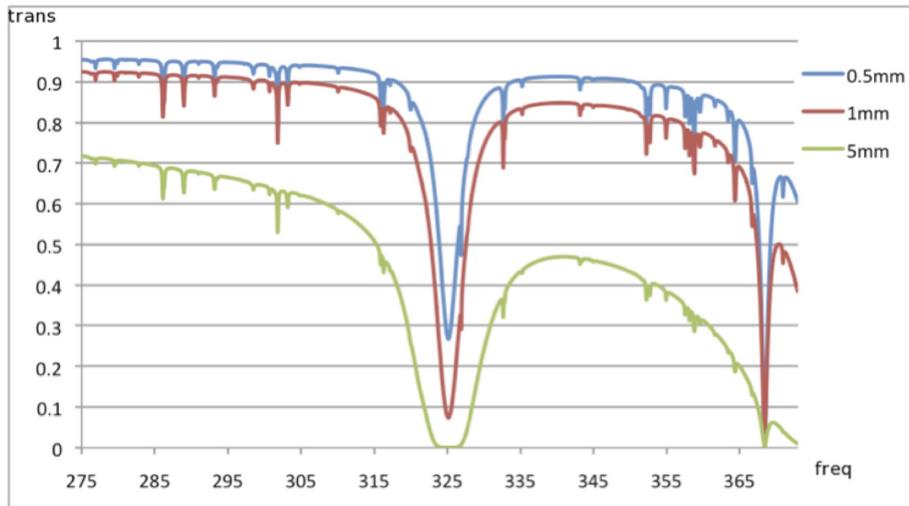


Figure 2.13: Band 7 atmospheric zenith transmission for pwv=0.5, 1.0 and 5.0mm. Frequency is in GHz. The deep atmospheric absorption at 325 GHz is due to water, and the less prominent absorption feature at 369 GHz is due to oxygen.

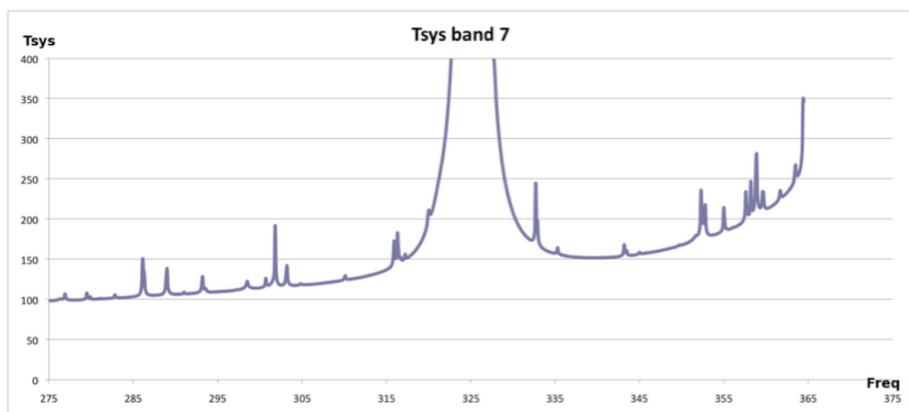


Figure 2.14: Typical T_{sys} at zenith for Band 7 with pwv=1mm.

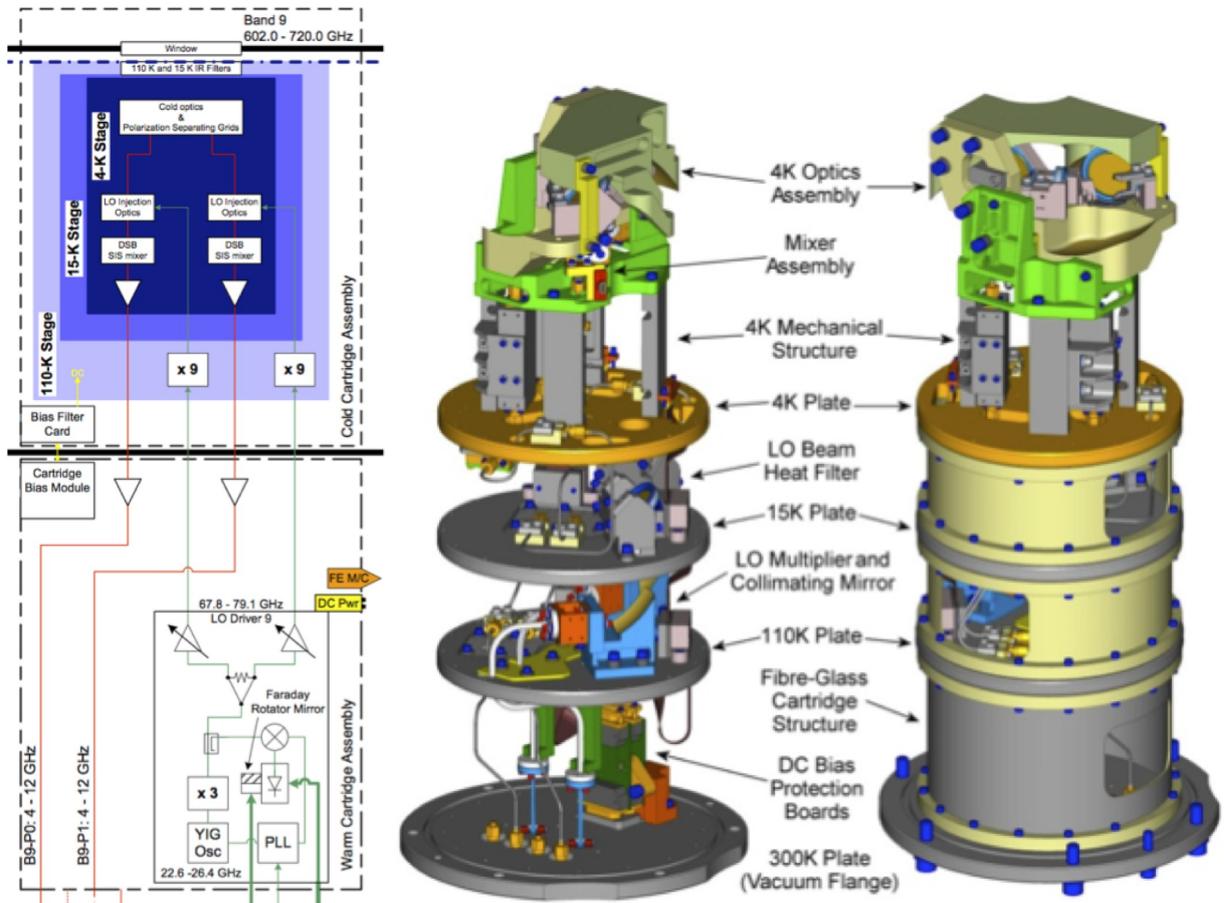


Figure 2.15: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 9 receiver was built at SRON in the Netherlands.

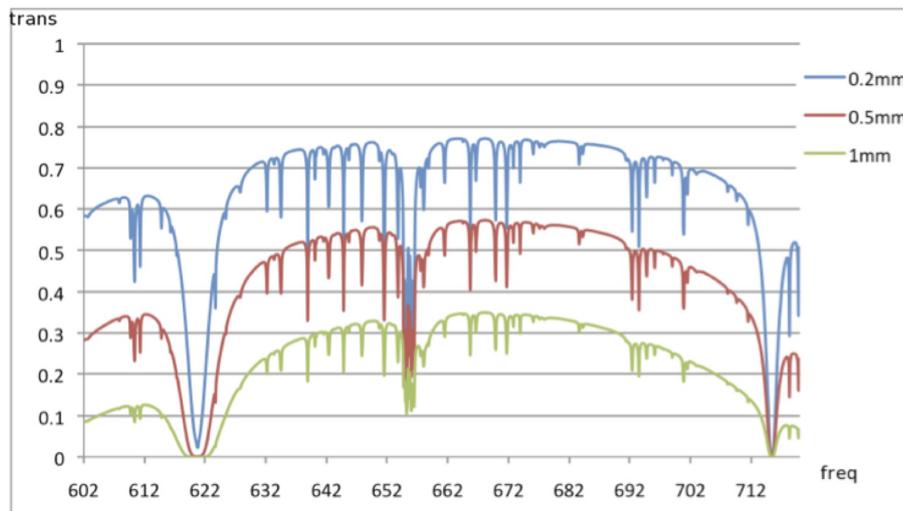


Figure 2.16: Band 9 zenith transmission for pwv = 0.2, 0.5 and 1mm. Frequency is in GHz.

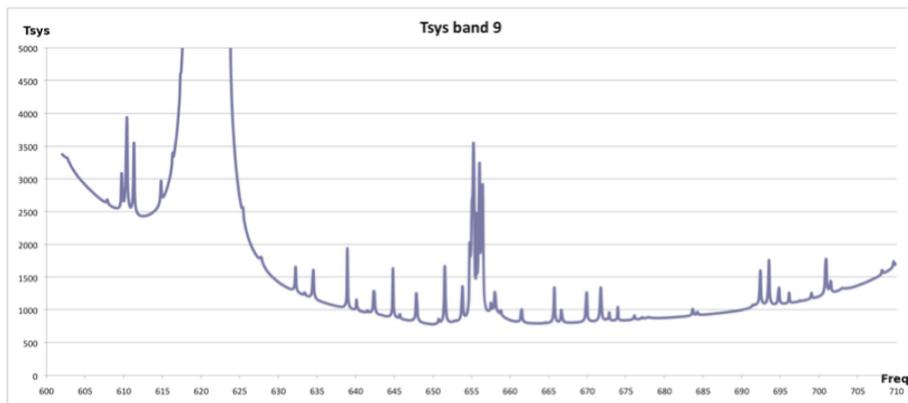


Figure 2.17: Typical T_{sys} at zenith for Band 9 with $\text{pwv}=0.5\text{mm}$. (T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

Chapter 3

The Correlators

The observed sky frequencies need to be down-converted to lower frequency bands in order to be digitally sampled and sent to the correlators. The frequency down-conversion in the IF system involves a set of Local Oscillators (LOs) and mixers, described in detail in Section A.7. The outputs from the IF system, known as "basebands", cover 2-4 GHz. These analog baseband signals are sampled and converted to a digital form (see subsection below) and fed over fibre-optics to the correlators. Two correlators will be available for Cycle 1: the 64-antenna correlator (also known as the 64 station correlator or the Baseline Correlator) and the 16-antenna ACA correlator. For Cycle 1, both have essentially identical characteristics as far as the 'end user' is concerned. They can handle up to 8 basebands per antenna simultaneously (4 basebands per polarization). The 64-antenna correlator is used primarily for the main 12-m Array, and the ACA correlator used primarily with the ACA 7-m Array and Total Power (TP, or single dish) Array. However, crossbar switching allows for some flexibility in this arrangement. Both correlators can be run simultaneously and independently; for example, this allows the 12-m Array to observe one object using the 64-antenna correlator, while the ACA correlator is used with the 7-m Array to observe either the same or a different object. It is also possible to take TP autocorrelation data simultaneously using a separate sub-array.

3.1 Digitizers

Correctly converting the analog baseband signals into digital form requires that the data be sampled at a rate equal to or more than the Nyquist frequency. The ALMA digitizers have a clock frequency of 4 GHz, so the 2-4 GHz basebands are effectively aliased down to 0-2 GHz. Samplers have a limited number of levels, specified by the number of bits in the resulting digital datastream. In the ALMA system, there are actually two sampling stages:

- The first stage of digitization is in the backend (BE) at the antenna. This is a 3-bit sampler (8 levels), and the digital signal is converted to a parallel datastream for transmission over fibre-optics to the correlators in the AOS technical building.
- The second stage of sampling is from the output of the station cards in the 64-antenna correlator feeding the correlator boards themselves. This can use 2,3 or 4 bit sampling. In the latter cases, the number of channels is reduced. In Cycle 1, only 2-bit sampling is available. Because of the limited number of bits used in sampling, there is a slight increase in the noise on the signal from the digitization (the quantization noise). This depends on the number of bits and the level into the samplers; at the optimum setting, the loss with 2-bit sampling is 12% (the mode available in Cycle 1), which decreases to 2% for a 4-bit sampler (to be available in the future).

Data is transferred in the ALMA digital format (Freund 2002, ALMA Memo No.420) through optical fibers from digital transmitters (DTX) in the antenna backends (BEs) to the data receivers (DRX) in the AOS

technical building.

3.2 The 64-antenna correlator

The 64-antenna correlator is of a digital hybrid design that enhances the performance of more traditional lag correlators (XF)¹, and is known as an FXF system. It is described in detail in Escoffier et al. (2007, A&A 462, 801). It operates in two basic modes, TDM (Time Division Mode) for low resolution wideband continuum observations, and FDM (Frequency Division Mode) for higher spectral resolutions. A simplified overview diagram of the 64-antenna correlator is shown in Figure 3.1. It consists of 4 quadrants, all of which will be available for Cycle 1. Each quadrant can handle a full baseband pair (defined as one of the basebands in each polarization) for anything up to 64 antennas for a total bandwidth per antenna of up to 16 GHz (that is, 8 GHz per polarization). The data is corrected for geometrical delays prior to processing, and for quantization errors during the processing.

3.2.1 FDM mode

In FDM mode each 2 GHz baseband is subdivided into as many as 32 sub-bands, with a full (nominal) bandwidth of 62.5 MHz each. This is done by digital filtering using Field Programmable Gate Arrays (FPGAs) in the Tunable Filterbanks (TFBs). To provide more than 62.5 MHz of bandwidth, multiple TFBs are programmed by the system to lie adjacent to one other in frequency. However, to avoid aliasing and edge effects, only 15/16 of the total bandwidth from each TFB is actually made available externally, giving a *usable* bandwidth per sub-band when operating in this mode of 58.5975 MHz (also the sub-band separation is 58.5975 MHz). The remaining edge channels are truncated within the correlator, and not visible in FDM data.

For Cycle 1 all TFBs in use in each baseband are stitched together in the correlator data processing to give a single Spectral Window (spw) per baseband². This provides nominal bandwidths from 62.5 to 2000.0 MHz in steps of 2, depending on the number of sub-bands used³. The full list of Cycle 1 correlator modes for dual-polarization operation is given in Table 3.1; this includes the nominal bandwidth internal to the system and the available bandwidth that the end-user sees in their spectrum. Normally the stitching is reliable and the individual sub-bands cannot be discerned in the final spectrum; however, one characteristic of correlator problems when in FDM mode is that the sub-band edges (or platforming) become visible.

At the output of the TFBs, the signals are re-quantized to 2, 3 or 4 bits (for Cycle 1, only 2-bits; see Section 3.1) for correlation. In this case, FDM modes provide up to 7680 channels per baseband pair in Cycle 1. For N polarization products, the number of channels is 7680/N, so for the most common operating mode (dual-polarization FDM) this gives 3840 channels (Table 3.1). FDM is used mostly (although not exclusively) for spectral line observing.

The center frequency of each Spectral Window can be tuned over the 2 GHz-wide baseband using a digitally-synthesised LO (the TFBLO, or LO4). However, there are limitations, for example the edges of the full bandwidth of the sub-bands cannot fall outside the 2 GHz baseband range, and the tuning is in steps of 30.5kHz (see LO section A.7 in the Appendix, and Section 4 below).

New for Cycle 1 is the ability to average spectral channels by multiples of 2, in order to reduce the data rate in FDM modes. So for example, in FDM 2-polarization mode, the number of channels per spw could be 3840 (the maximum), 1920, 960, 480, or 240. This is described in detail in Section 3.5.1.

3.2.2 TDM mode

In TDM mode, the TFBs are bypassed and the full 2 GHz baseband is fed through the correlator. The signal is distributed in time over several correlator chips, each of which performs the cross-correlation in parallel for some

¹X=correlation and F=Fourier Transform

²In future observing Cycles, the sub-bands will be independently tunable for full flexibility, allowing multiple spws per baseband

³But note the available bandwidth is 15/16 of the nominal bandwidth for reasons given previously

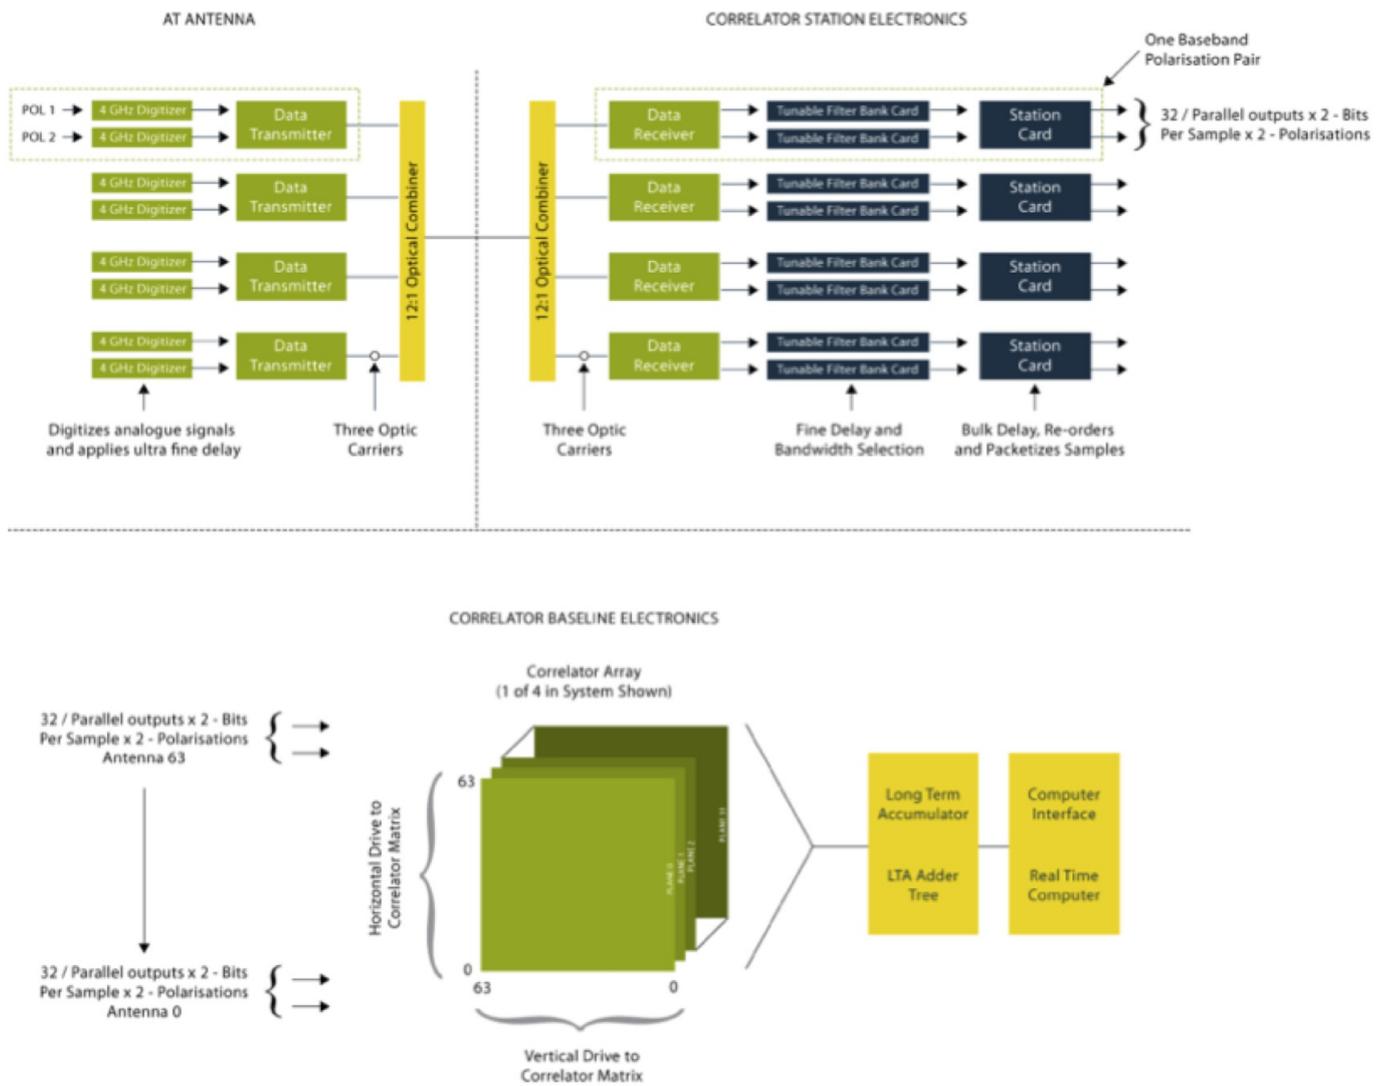


Figure 3.1: Overview diagram of the ALMA digitizers, data transmission system and 64-antenna correlator. The digitised data from the individual basebands are fed to the TFBs and station cards (top right) for the parallelisation and processing required on each antenna datastream. The matrix of correlator boards (lower half) does the cross-correlation between antenna stations, and the resulting data is then integrated in the Long-term Accumulators (LTAs, lower right).

Nominal Bandwidth (MHz)	Usable bandwidth (MHz)	Usable channels	Channel spacing (MHz)	Correlator mode
2000	~1875	~120	15.6	TDM
2000	1875	3840	0.488	FDM
1000	938	3840	0.244	FDM
500	469	3840	0.122	FDM
250	234	3840	0.0610	FDM
125	117	3840	0.0305	FDM
62.5	58.6	3840	0.0153	FDM

Table 3.1: Correlator modes (in dual polarization). Note that the channel spacing given here is *not* the same as the spectral resolution because of the applied weighting function; for calculating spectral resolution, see Section 3.5.

fraction of the time interval. The correlations are then re-combined to get the fully time-sampled data. The TDM mode provides for up to 256 channels per baseband (for N polarization products, the number of channels are reduced to $256/N$), and no edge channels are dropped in the correlator. As the full 2000 MHz baseband is covered, this requires some truncation of the edge channels in offline data processing – see Section 4.4. TDM is used mostly for continuum observing. It has the advantage of having a lower data rate than FDM, and is therefore used for pointing, focus, delay, system temperature, sideband ratio and other calibration observations where high-resolution is not required.

3.2.3 Correlation and realtime processing

In the cross-correlation modules, the multiply-and-add operations are performed at a clock rate of 125 MHz (4 GHz samples divided by 32). Each board (Figure 3.2) contains 64 4-k lag correlator chips, connected in a matrix to provide the required complex correlations. Cross-correlation data goes then to a time integrator (the Long Term Accumulator, LTA – Figure 3.1) that adds the data up to an integer multiple of 16 ms for cross-correlations and multiples of 1 ms for auto-correlations. The data can be further co-added over a dump period in the correlator Data Processor computers (CDP). Data for each sub-band are Fourier transformed and the final spectrum (when in FDM mode) is formed by stitching together the sub-bands.

In Cycle 1, each baseband pair is fed to one correlator quadrant, which can have a bandwidth mode given in 3.1. In Cycle 1 it will be possible to set the correlators to different modes in different basebands. So, for example, one baseband can be set to TDM for the continuum, and another baseband can be set to a high-resolution FDM mode to study an individual line. More details of spectral setups are given in 4.

Section 3.4 describes the further processing of the data from the correlators.

3.3 The 16-antenna ACA correlator

The 16-antenna (ACA) correlator is newly-available for Cycle 1. It is a 16-station correlator using the FX technique, whereby the incoming digitized data stream is first Fourier-Transformed (“F”) in realtime before the cross-correlations (“X”) are performed. All processing is done with high-speed FPGA chips, rather than custom chips as used in the 64-antenna correlator. An overview block diagram is shown in Figure 3.3, and a detailed description of the ACA correlator is given in Kamazaki et al., 2012 (PASJ, 64). Although the internal operation is different, the correlator is designed to mimic the performance capabilities of the 64-antenna correlator, for example in bandwidth/resolution modes etc. The ACA correlator does not use sub-bands, so there should not be any problems with sub-band stitching. It does allow for further flexibility of modes, but this will not be available in Cycle 1.

A detailed block diagram of the ACA correlator itself is shown in Figure 3.4. Like the 64-antenna correlator,

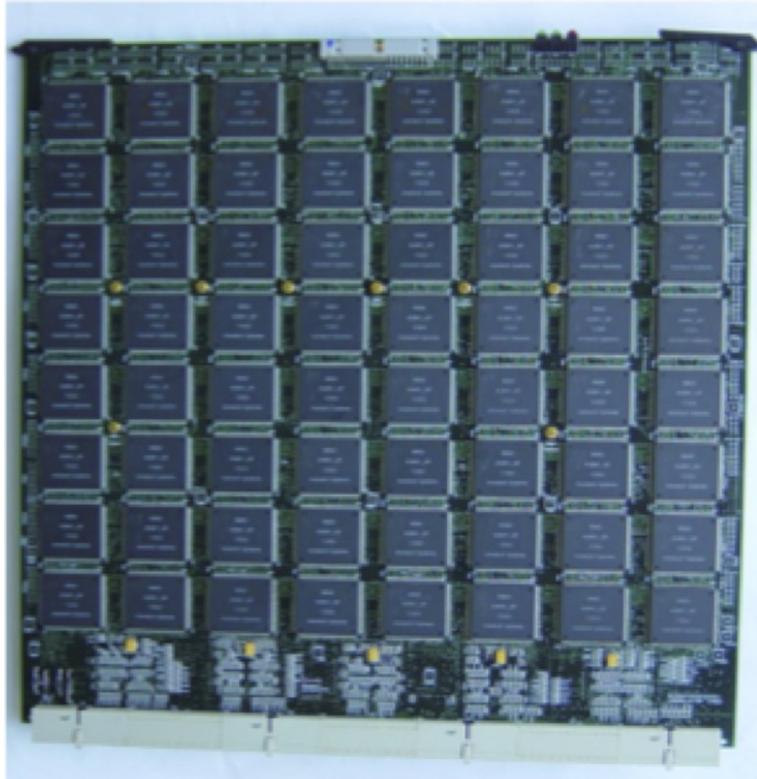


Figure 3.2: Correlator board for the 64-antenna correlator, with the 8x8 array of custom correlator chips, each with 4k lags. There are 128 such boards per correlator quadrant, adding up to a total of 32768 chips.

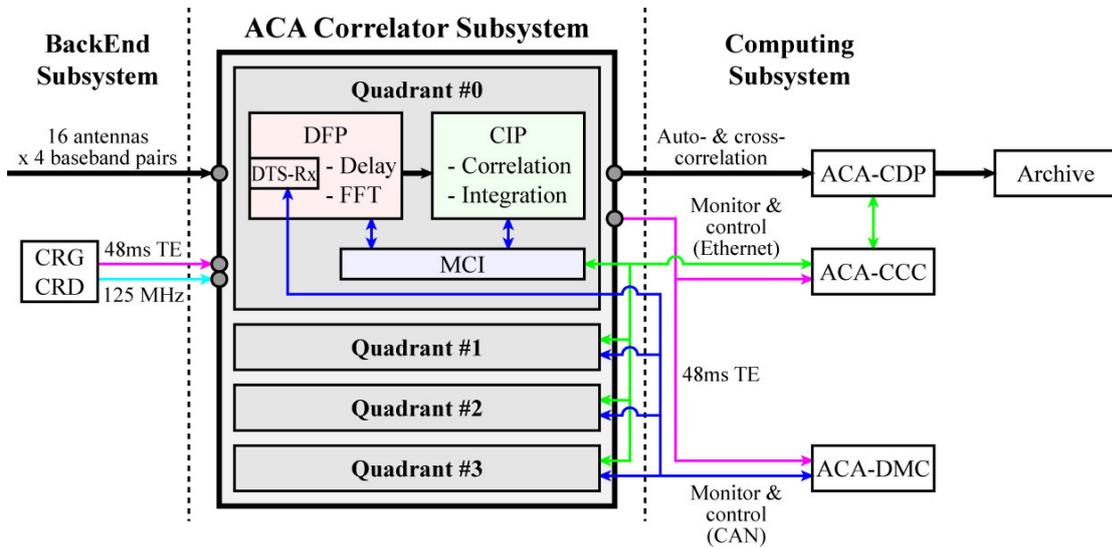


Figure 3.3: Overview diagram of the ACA correlator. The signals from the BackEnd Subsystem are processed in the DFP (Digital FFT Processor) and CIP (Correlation Integration Processor) modules and output to the Computing Subsystem (see text and Figure 3.4 for details of these). The modules are controlled by the MCI (Monitor and Control Interface), 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-antenna correlator.

each quadrant of the ACA correlator processes a baseband pair (2 GHz bandwidth, dual polarizations); so the complete system can process up to sixteen antennas independently of the 64-antenna correlator. After receiving the digitized baseband signals in the DRX (Digital Receiver, or DTS-Rx in the figure), the DFP (Digital FFT Processor) modules compensate for geometrical delays between antennas and performs the $1\text{M}(=2^{20})$ -point FFT (Fast Fourier Transform). This results in $512\text{K}(=2^{19})$ 16-bit complex values (hereafter, voltage spectra) every baseband per antenna, giving a frequency resolution of 3.815 kHz ($=2\text{ GHz} \div 512\text{K-point}$). All the voltage spectra are re-quantized to 4-bit complex from 16-bit complex and sent to the CIP (Correlation and Integration Processor) modules for the calculation of power spectra.

In the CIP modules, the voltage spectra in the frequency ranges specified by observers are multiplied with each other between antennas. In the case of different antennas or different polarizations of the same antenna, the cross-correlation spectra (cross-power spectra) are produced, while in the case of the same antenna and the same polarization, auto-correlation spectra (auto-power spectra) are obtained by the correlation multiplication of identical voltage signals. Then, the auto- and cross-correlation spectra are integrated up to 3.90625 MHz in the frequency domain according to the specifications by observers and output to the post-processing computers (ACA-CDP) in the Computing subsystem through optical fibers.

Once the data have been accumulated, the CDP processing proceeds in the same way as the 64-antenna correlator, and further spectral and temporal integration are performed, before data is sent to the archive. The maximum data rate from the ACA correlator is 3.6 MB/s, independent of the number of antennas. However, the average data rate to the archive will be lower during typical observing, because of overheads and the use of TDM modes during some observing. As the 7m antennas do not have WVRs, the WVR correction for the ACA is intended to use the TP antenna data, although more testing is required before this is fully operational.

The ACA correlator supports flexible spectral configurations, which are compatible with those of the 64-antenna correlator, in the Cycle 1 mode. However, it does not have tunable filter banks (TFBs) and realizes the FDM functions in another way. As previously mentioned, the ACA correlator adopts a conventional FX-architecture and always performs spectroscopy of inputs signals using FFT at the highest frequency resolution inside the correlator. Hence, it can easily output multiple spws with different frequency ranges and resolutions by selecting the highest resolution spectra and summing them up as shown in Figure 3.5. In the spectral integration, two types of integration methods are implemented in the ACA correlator. One is non-weighted spectral integration, which just adds up spectral data without weighting in the frequency domain. The other is frequency profile synthesis, which performs convolution with a weighting function in the frequency domain (Kamazaki et al. 2008, ALMA Memo No.580). The latter is necessary to make the frequency profile of the ACA correlator the same as the 64-antenna correlator. In Cycle 1, only the latter function is available, using the default Hanning smoothing.

3.4 Correlator data processing

Parameters of the data processing in the CDP and later are effectively identical for the 64-antenna and the ACA correlator, outlined below.

3.4.1 Integration time intervals, channel-average and spectral data

The correlators produces two datasets, the complete spectral data, and the channel-averaged data. These are normally written to the archive from the CDP (the correlator data processor) at different rates (see Figure 3.7). The channel-averaged data is a single complex number for each spw. In the case of TDM, the data is truncated in the CDP before averaging (only the central 15/16 are used). For FDM no truncation is applied. The primary purpose of the channel-averaged data (left of figure) is to provide a smaller dataset for TELCAL, the ALMA online processing software, which computes real-time pointing and focus corrections. It is stored in the ASDM (ALMA Science Data Model) dataset as a separate spectral window (see 3.4.3). Note that this should not be used as science continuum data in normal observing, but the continuum should be constructed offline in CASA (Common Astronomy Software Applications) using the appropriate portions of the spectral data. The spectral data (right of figure) contains the requested number of channels for each requested polarization product. The

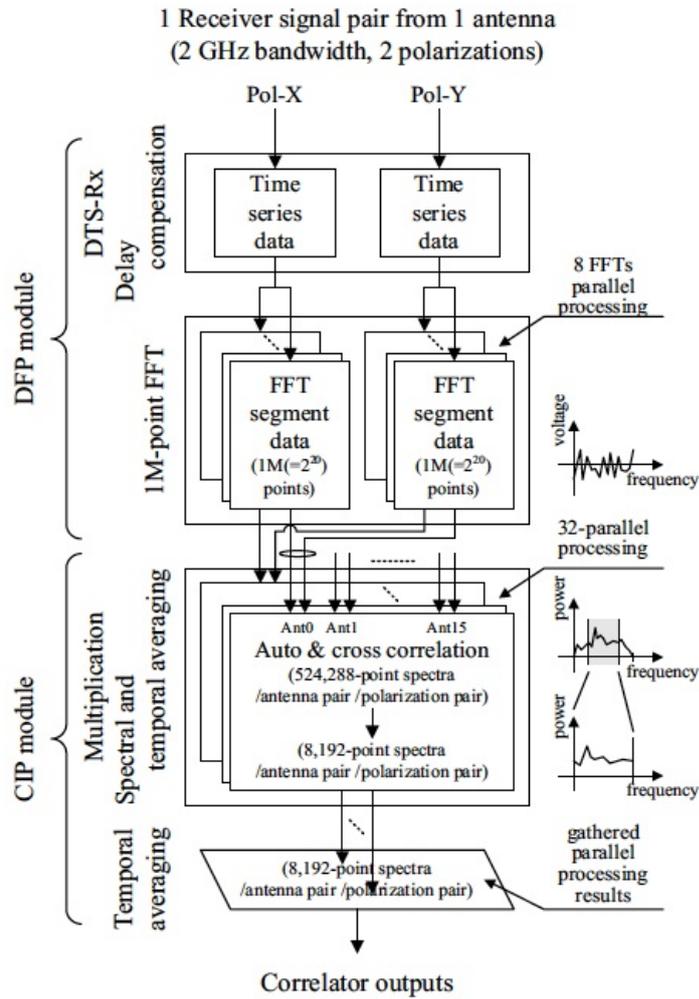


Figure 3.4: Block diagram of ACA correlator. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. A Fast Fourier Transform (FFT) is performed using FPGA modules, resulting in signals in the spectral domain which can then be correlated. The auto and cross-correlation data are then accumulated in time, and the parallel streams averaged together in the CIP (Correlation and Integration Processor) module. The correlation spectra are then fed to the ACA CDP for further accumulation and processing.

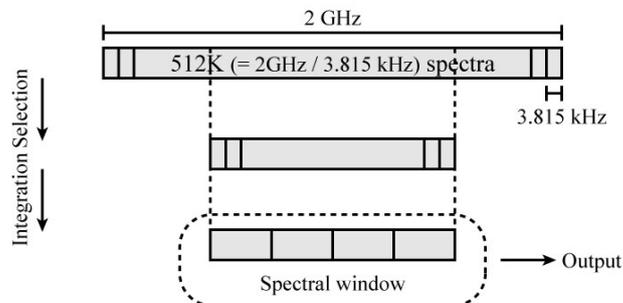


Figure 3.5: Data capture sequence for the ACA correlator. A spectral window is available to select a frequency range from 512K spectra, whose total bandwidth is 2 GHz and frequency resolution is 3.815 kHz, in the Cycle 1 mode. Selected spectra are integrated and then sent to the output.



Figure 3.6: A quadrant of the ACA correlator. The left and middle racks compose a quadrant, which processes a baseband pair from sixteen antennas. The right rack is one of four computing racks, which have correlator control, post processing computers, network instruments and so on.

time intervals involved at this stage of the observing and data acquisition are shown in Figure 3.7. In detail they are:

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

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

Spectral integration time: The spectral integration time (also known as the integration duration) is the shortest integration time before the full-resolution spectral data is written to the ASDM. So it is also the shortest time interval that can be selected (and potentially self-calibrated) in CASA. The spectral integration time must be a multiple of the dump duration. It is also usually set to be a multiple of the channel average time. For Cycle 1 using FDM observations, a Spectral Integration time of several seconds is used to avoid data rate problems (see 3.6). For TDM observations, which have fewer channels, a faster rate can be used.

Subscan duration: A further time interval, not shown in Figure 3.7, is the subscan duration or time. In an observation, this will effectively be the shortest time interval where no parameter is changed in the system. For example this could be a single integration on a source, one point of a mosaic, one scan of a raster, integration at one frequency, or on one of the loads, etc. The subscan duration must be a multiple of the spectral integration time. Typically a subscan might be 30 or 60 seconds; the ASDM might therefore contain 5-10 spectral integration times per subscan, resulting in 5-10 rows of data per baseline/polarisation/baseband product.

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

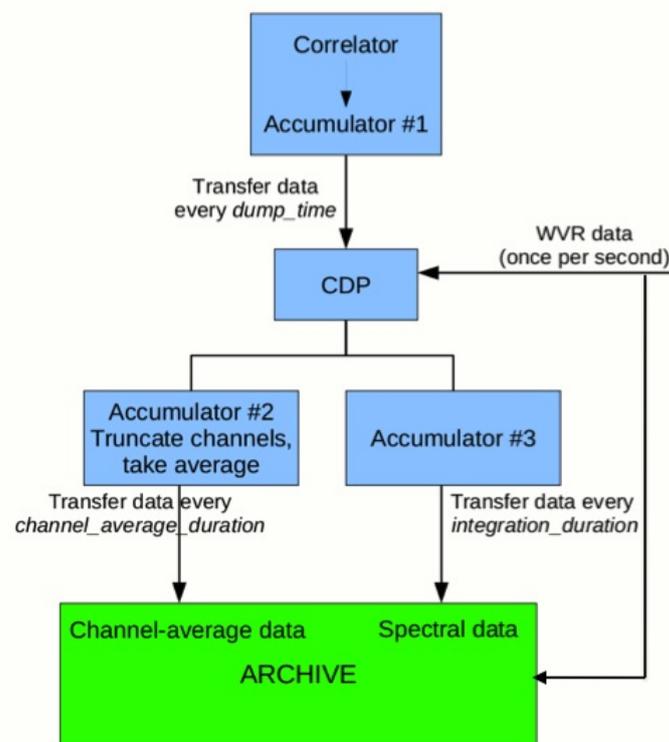


Figure 3.7: Basic data processing and accumulation steps in the correlator and archive. The timing intervals shown are described in the text.

3.4.2 Online WVR correction

New for Cycle 1 is the option to do semi-realtime correction of the pathlength fluctuations in the CDP using the WVR data from each antenna. The advantage of doing WVR correction in realtime is that it can be performed relatively rapidly (up to the largest value of the dump duration or the WVR chopperwheel rotation and readout period, approximately 1 second), thus enabling corrections to be made for faster variations in atmospheric transmission. At the same time, corrected data can be integrated in the CDP, allowing data to be transferred from the correlator to the archive at a slower rate. Note that once the data is combined, the online WVR correction is non-reversible. For some time during Cycle 1, it is expected that both the internally corrected and the uncorrected data will be written to ASDM once every Spectral Integration Time (although note that this will double the overall data rate; see 3.6). So, as in Cycle 0, it will be possible to do the WVR correction offline on the slower uncorrected data (or even not to do any WVR corrections). But it is expected that the realtime WVR corrections will track the sky more accurately, and this data will eventually be used directly. The optimal method of WVR correction, potentially including a combination of on-line and off-line corrections, is under investigation.

3.4.3 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 (eg uid://A002/X2fed6/X3f). The ASDM contains the meta-data (headers, descriptions of the observation setup, ancillary data etc), and the bulk data (the raw data itself), and is described in more detail in Section 9.5. The following describes the spectral data in the ASDM and how this maps to the correlator outputs.

In the ASDM, the bulk 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 upwards of 25 spectral windows. Luckily every scan/spw combination has an "intent" associated with it that indicates its purpose (pointing, system temperature, science, etc). This intent can be used in CASA to decode how to utilize each spectral window in the data reduction process.

3.5 Spectral resolution, smoothing and channel averaging

The channel spacing from the correlator is just that - the frequency spacing of the spectral channels in the data. The spectral resolution in general is not the same as the channel spacing. Moreover, the bandwidth used to calculate the sensitivity in the OT is not, in general, the same as the spectral resolution: typically it is larger (and it should not be less).

It is possible to select various weighting functions in the correlation (lag) dimension (see Table 3.2) which will affect the resolution. Without a weighting function (listed as 'Uniform' in the table), the spectral response will be a sinc function (the fourier transform of the unweighted top-hat function). This has a resolution of 1.2 times the channel spacing, but has undesirably-high spectral ripples or sidelobes, resulting in 'ringing' in the spectra when a narrow line, interference spike, or strong edge channels are present. Different levels of smoothing can be used to suppress these, but they also degrade the spectral resolution (for a full description,

see <http://mathworld.wolfram.com/ApodizationFunction.html>). The default weighting function is Hanning, which gives a spectral resolution $2.0 \times$ the channel spacing; in this case, the largest sidelobe response for an unresolved signal will be 2.6% (see Table 3.2). Note that this level of sidelobe will not normally be noticeable in astronomical spectra, as most lines will normally be resolved over several channels. If a different smoothing function is required in a specific science proposal, this should be justified in the technical case, and this can be selected in the phase II setup of the SB.

Mode	Resolution (FWHM) (channels)	Max spectral sidelobe
Uniform (none)	1.2	-0.22
Hanning (default)	2.0	-0.026
Hamming	1.81	-0.0069
Bartlett	1.77	+0.047
Blackmann	2.230	+0.0012
Welch	1.59	-0.086

Table 3.2: Spectral Resolution and sidelobe levels for different smoothing functions.

3.5.1 Channel averaging

New for Cycle 1 is the ability to bin or average spectral channels in the correlator. Channels can be averaged together in factors of $N = 2, 4, 8,$ or 16 . The main aim 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) and full FDM (with 3840 channels in the 2-polarization mode). It might be quite acceptable for those that need something with more resolution than TDM, but where the FDM channels at the full resolution are unnecessary. Table 3.3 shows the resolutions (in kHz) for different values of N , using Hanning smoothing, in the different bandwidth modes. In brackets are the channel spacings. $N=1$ is the default un-binned case, where the resolution is $2 \times$ the channel spacing.

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

Table 3.3: Spectral resolution and channel spacing (in brackets) in kHz for different correlator bandwidth modes (left column) and for different channel averaging factors (columns, $N=1$ to 16), using Hanning smoothing. The number of channels can be reduced from 3840 (for the un-averaged case, $N=1$) down to 240 (for $N=16$). Values are given for the 2-polarization case.

Note that the default Hanning spectral smoothing function (see above) gives a resolution $2x$ the maximum, 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 definitely required by the observations, then the number of channels be reduced when feasible. This is selected in Phase II of the SB creation. However, note that this is a non-reversible operation!

3.6 Correlator speed and data rates

The maximum data rate from the 64-antenna correlator in the Cycle 1 mode (using 32 antennas) is likely to be similar to Cycle 0, or about 17 Mbytes/sec. However, it is intended that this will be gradually increased with software upgrades to the design specification of 66 Mbytes/s. For the ACA correlator the maximum data rate is 3.6Mbytes/s. These rates set limitations on the shortest value of the spectral integration time in the correlator, particularly in FDM mode, where the number of channels without spectral averaging is 32 times larger than TDM mode. The FDM raw data rate (in Mbytes/sec) is approximately $N_A(N_A - 1)/2 \times 8192 \times N_{bb} \times 4/1e6/I_T$, where N_A is the number of antennas, N_{bb} is the number of basebands (normally 4) and I_T the spectral integration time. Note that this rate will double if both WVR online-corrected and uncorrected data is required from the correlator. As in cycle 0, the usual FDM spectral integration times set in the SBs are typically expected to be 6 seconds; this may be increased if the on-line WVR correction proves to give accurate and reliable path-length corrections. So under most circumstances the maximum data rate limitations should not have any impact on the science that can be done in Cycle 1, as it is expected that these high values will seldom be reached.

Note that the *average* data rate to the archive will be considerably lower than the maximum correlator data rate, because of observing overheads and the fact that some data taking within an SB will be in TDM rather than FDM mode (eg pointing and atmospheric calibration). The maximum average data rate to the archive is currently 6.6Mbyte/s (see Chapter 12).

Chapter 4

Spectral setups

Creating a spectral setup effectively consists of setting the local oscillators and correlator in the system such that the spectral windows cover the desired lines and/or continuum emission. The spectral setup is defined using the Observing Tool. For full details of the OT, see the user manual and reference manuals, available from the ALMA website (and also in the OT itself). In this section we describe how the OT information is used to set up the ALMA system.

During proposal preparation (Phase 1 and 2) using the OT, the users choose the frequencies (or transitions) to be observed, and the OT will look for the best solution for the settings of the LOs, correlator etc. There are effectively 4 different LOs in the system: LO1 which sets the frontend tuning frequency, LO2 which positions the basebands within the receiver IF output (there are 4 LO2's, one per baseband), LO3 which is the clock frequency of the digitizers (fixed at 4 GHz), and LO4 (also known as the TFB LO), which is a digital LO in the correlator allowing positioning of the spectral windows within each baseband. Their operation is described in more detail in Section A.7. The OT will choose the optimum values for the LO1, LO2 and LO4 frequencies based on the user's selections of frequencies to be observed. Compared with Cycle 0, Cycle 1 has fewer restrictions on the spectral setups (see Table 4.2 in section 4.7). The main limits are that the edges of the 2 GHz basebands cannot lie outside the receiver tuning range listed in Table 2.1, and the edges of the (untruncated) spectral windows cannot lie outside the 2 GHz-wide basebands. In Cycle 1 the settings of the spws in each baseband are independent, so the resolution and correlator mode as well as the value of LO4 can be different in each baseband. So for example it is possible to have a 62.5 MHz-wide spw centered on a particular line in spw1, and simultaneously use a broadband TDM mode for all the other spws. Note that in Cycle 1, only one spw per baseband is available.

Figure 4.1, which is adapted from the Spectral Editor of the Observing Tool (OT), illustrates a spectral setup. The blue hashed area represents the receiver tuning range (in this case, a section of Band 7), and the curved line the nominal atmospheric transmission for the chosen PWV. The LO1 is 335 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 red horizontal lines, can be moved around, but only within the two sidebands. The spectral windows, shown by the blue crosses, can cover either the whole baseband, or, for higher spectral resolution, a smaller fraction of the baseband. In this example, we have set up two wideband spectral windows to observe continuum in LSB (bb0) and USB (bb3). At the same time, we are observing two lines using high-resolution: the 13CO line in LSB (using bb1) and the CN line in USB (bb2). The basebands (and the spws) may partially overlap, as seen in this case with bb2 and bb3. The spws need not be in the center of the basebands (as seen in this case with the 13CO spw at the upper edge of bb1).

4.1 Spectral setups for multiple lines

The wide IF bandwidth and tuning ability allows for simultaneous imaging of multiple lines. Some examples (with the approximate line frequencies in GHz) are shown in Table 4.1 (for a source of redshift zero). Note

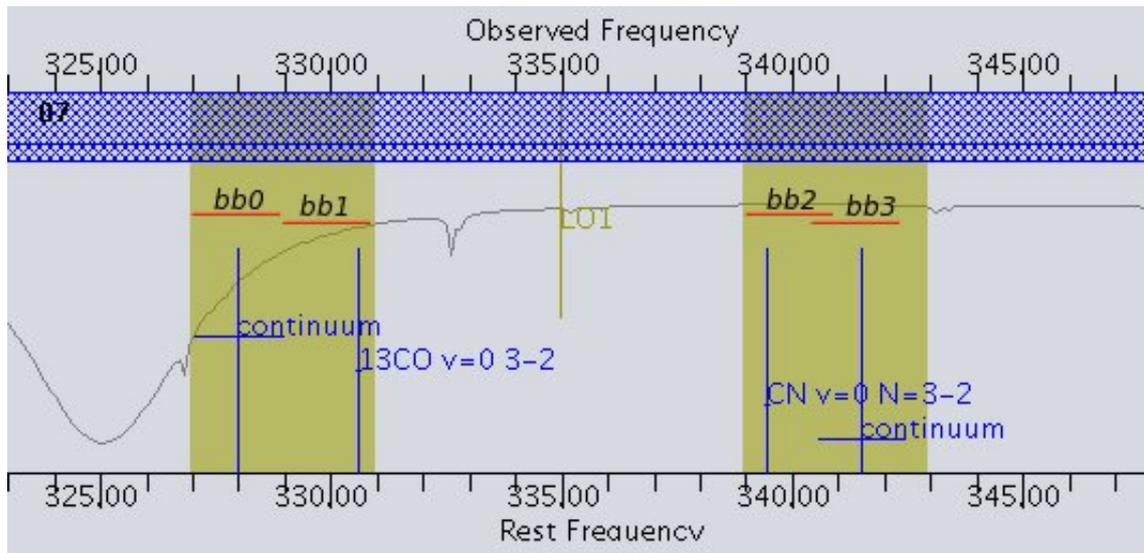


Figure 4.1: Illustration of a frequency setup, based on the OT spectral display. Yellow areas are the IF ranges, red lines are the 2 GHz-wide basebands (bb0-3), and blue crosses represent four spectral windows (spw 0-3, one per baseband). The frequency of LO1 is shown by the central vertical line. The blue hashed area shows the possible tuning range of the frontend, and the curved line gives an indication of the atmospheric transmission. The molecules and transitions of interest are shown, and their frequencies are shown by the position of the vertical blue lines. The bandwidths of the spws are illustrated by the horizontal bars of the blue crosses.

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 which are observable can be done using the OT spectral interface. In the case of continuum spws, in order to maximize the sensitivity, the widest bandwidth mode should be chosen in these spws (i.e. 1.875GHz in FDM, or 2 GHz in TDM - giving a usable bandwidth of ~ 1.9 GHz - see Section 4.4). 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 may eventually 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.

4.1.1 Calculating optimum continuum frequencies

If possible the continuum spws should be set to the sideband with the lowest system temperature (normally the best atmospheric transmission, as displayed on the OT - see Figure 4.1). If a single line in spw 1 is observed at frequency F , the continuum spws 2-4 can be set up at the following frequencies:

1. F (spw covering the primary line)
2. $[F - 2]$ GHz (for a continuum spw 2GHz below the primary line, in the same sideband)
3. $[(F - 1) + / - (2.F_{IF} + 1)]$ GHz (in the opposite sideband)
4. $[(F - 1) + / - (2.F_{IF} - 1)]$ GHz (in the opposite sideband)

In these, + means the opposite sideband is USB and - means the opposite sideband is LSB. F_{IF} is the central IF frequency (6.0 GHz for band 3 and 7, and normally 8.0 GHz for band 6 and 9). Note that this is an approximate rule - if possible, adjustment of the continuum spws (for example, choosing whether the opposite sideband is

LSB or USB) should be done to avoid deep atmospheric absorption features, particularly at band 9 and near the water lines around 325GHz in band 7 (see the left side of Figure 4.1).

Band	Species/transition	Frequency	Sideband	bandwidth	spw
3	HCO+ 1-0	89.188	LSB	62.5	1
	HCN 1-0	88.632	LSB	62.5	2
	CH ₃ OH/H ₂ CO	101.293/101.333	USB	125	3
	contin	101.3	USB	2 GHz/TDM	4
6	contin	231.6	USB	2 GHz/TDM	1
	¹² CO 2-1	230.538	USB	125	2
	C ¹⁸ O 2-1	219.560	LSB	250	3
	¹³ CO 2-1	220.399	LSB	125	4
7 (a)	contin	343.800	LSB	2 GHz/TDM	1
	¹² CO 3-2	345.796	LSB	250	2
	HCO+ 4-3	356.743	USB	500	3
	HCN 4-3	354.505	USB	500	4
7 (b)	¹² CO 3-2	345.796	USB	62.5	1
	contin	344.8	USB	2 GHz/TDM	2
	¹³ CO 3-2	330.588	LSB	62.5	3
	contin	331.6	LSB	2 GHz/TDM	4
9	¹² CO 6-5	691.472	USB	500	1
	CS 14-13	685.436	USB	500	2
	H ₂ S	687.303	USB	500	3
	C ¹⁷ O 6-5	674.009	LSB	500	4

Table 4.1: Some examples of multiple bright line/continuum configurations possible in Cycle 1.

4.2 Spectral setups for lines near the edge of the bands

The restriction that the selected bandwidth 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 ¹²CO(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 2.1). A setup using the wide bandwidth modes (TDM or FDM/1875 MHz) centered at the line frequency will not validate in the OT. There are two possible solutions: If the full bandwidth is absolutely required, the observing 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.

4.3 Spectral setup in Band 9

The spectral setup in band 9 is effectively the same as other bands. The only difference is that this is a double-sideband receiver (DSB), so the unwanted (or image) sideband needs to be suppressed using LO offsetting (see section A.7.3). The choice of which sideband a particular baseband is configured to observe depends only on the sign of the LO offset (which can be arbitrarily different for each baseband). So it is possible to set up two basebands at approximately the same LO1/LO2 frequency, but observe one line in the upper sideband in one baseband, and another line in the lower sideband using the next baseband, just by having different LO offset signs. If this sounds confusing, don't worry - just treat each baseband independently as far as sideband choice is concerned.

4.4 Usable bandwidth

The IF system contains an anti-aliasing filter which limits the bandwidth of the basebands. Nominally this filter has -1dB 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 bandwidth of 1.875 GHz, thus in FDM, the filters do not truncate the spectrum, and the full 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 4.2). 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 4.2.

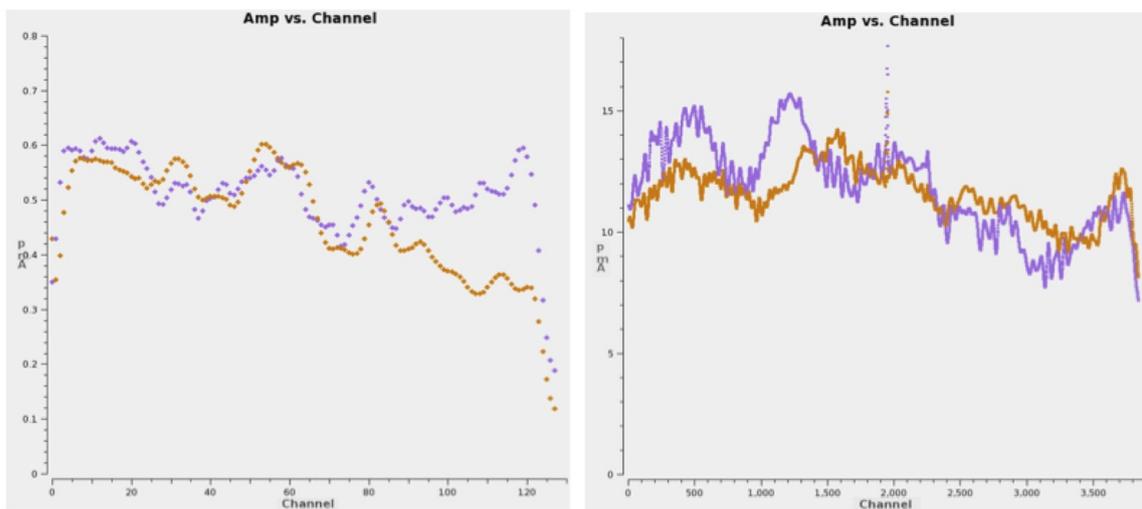


Figure 4.2: Comparison of TDM (left) and FDM (right) autocorrelation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in 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).

4.5 Spurious signals

Most spurious signals in the cross-correlation data are suppressed using Walsh-switching (see Section A.7.3). 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 Walsh switching show very few remaining spurious signals, although harmonics of the LO in the WVRs (ie multiples of 91.66 GHz) are seen. These are very narrow and can be flagged out during data reduction, but observing lines at these frequencies should be avoided if possible.

4.6 Doppler setting

New for Cycle 1 is the on-line correction for the science target velocity and Earth motion. If a target velocity and reference frame other than TOPO (topocentric) is selected in the OT, at the start of the SB at runtime, the velocity and reference frame of the science target and the velocity of the observatory relative to the chosen

reference frame are used to set the center frequencies of the science spws¹. The frequency setting for the science target will normally also be used in the SB for the bandpass and amplitude calibration.

4.7 Limitations and rules for spectral setups in Cycle 1

Although the Correlator will eventually allow complex combinations of spws in a single observation, these abilities are gradually being introduced and tested by Commissioning before release. Cycle 1 allows a broader subset of correlator capabilities compared with Cycle 0: new features are the ability to independently tune the spws around within the baseband (ie each spw can have a different TFB LO), and the ability to have different correlator modes for different spws (e.g. one baseband can be high-resolution FDM, and another one can use the TDM mode). There are still some rules (see Table 4.2) for spectral setups, although some of these are applicable to the current Cycle only.

¹Note that ALMA does not do Doppler *tracking*, where the frequency would be continuously corrected for Earth motion during the SB. In ALMA it is only set once per SB

Rule	Details	Note
1.	LO1 must lie within the LO tuning ranges given in Table 2.1.	general
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.)	general
3.	For 2SB receivers (Bands 3,6,7), the number of basebands in a sideband can only be 0, 1, 2 or 4. For DSB receivers (Band 9), there is no such restriction (the number can be 0, 1, 2, 3 or 4).	general
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. Centre Offset Frequency in the OT, phase II) must be $>(2000+B/2)$ and $<(4000-B/2)$ MHz. The Cycle 1 version of the OT forces this restriction. However there is a further restriction on this, as noted in the next rule.	general
5.	The spw usable bandwidth (ie 15/16 of the nominal spw bandwidth of multiple of 62.5 MHz) should be in an allowed region of the baseband. This is in addition to (4). In practice this means that the required region of the spw should be inside the range $\sim 2050 - 3950$ MHz to ensure that the edge of the anti-alias filter does not significantly affect the power.	general
8.	The line frequency must be in the center of the spw, otherwise the SB will not validate. This is mentioned in more detail in Section 4.2, for the case where a line is near the edge of the receiver band. It can be circumvented by entering an artificial line frequency.	Cycle 1
9.	Only 2-bit, Nyquist sampling is allowed	Cycle 1
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.	Cycle 1
11.	Only one spectral window is allowed per baseband. In the future, up to 8 spw will be allowed per baseband (or up to 32 total for all four basebands). Individual spw within a baseband may occupy all, 1/2, 1/4 or 1/8 of the resources available in the baseband and the sum of all the fractional resources used within one baseband must be ≤ 1 .	Future capability

Table 4.2: Rules for spectral setups. The note describes whether these are general restrictions, or restrictions in Cycle 1.

Chapter 5

Description of Atacama Compact Array in full operations and in Cycle 1

5.1 ACA during full operations

5.1.1 Concept and purpose of the ACA

Using an interferometer to obtain images of extended or large-scale structures suffers from the well-known “missing spacing” problem. This problem arises from the constraint that, to avoid collisions, it is not possible to pack antennas closer than 1.25 times their diameter, leaving a hole in the (u, v) sampling function at short and zero baseline separations, resulting in spatial information from baselines shorter than the closed-packing ratio not being recovered¹. This problem has considerable impact in 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 structures larger than the minimum spacing of the 12-m Array, ALMA has been designed to consist of a homogeneous array of 12-m antennas (12-m Array) and the Atacama Compact Array (ACA). In full operations the ACA will be composed of twelve 7-m antennas for interferometry (7-m Array) and four 12-m antennas for single dish observations (TP Array). The four single-dish provide uv-information equivalent to up to 12 m spacings as auto-correlations. The 7-m array samples baselines from 8.5 m to 43 m, bridging the (u,v) gap between the 12-m Array and the TP Array.

5.1.2 System & operations description

The ACA is controlled via 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. In full operations, 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 output from the ACA antennas are processed in the ACA Correlator (see Section 3.3) to produce the cross and auto-correlated data.

Due to the small point-source sensitivity of the 7-m Array, it is planned that for full operations, the TP Array will be routinely used in the calibration observations of the 7-m Array. Since the 7-m Array is quite compact, it is expected that the atmospheric phase fluctuations will be smaller than for the 12-m Array. The ACA has therefore been designed so that the WVRs installed on the TP Array antennas that surround the 7-m Array will be used for these corrections.

¹Strictly 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.

5.1.3 7-m Array configurations in full operations

No (u, v) coverage from even the most compact configuration of the 12-m Array is obtained for spacings < 12 m, and the limit is often 15 m to avoid shadowing conditions. The array configurations of the ACA 7-m Array are designed to fill missing spacings in (u, v) from about 9 m to 30 m. The 12-m TP Array further fills in the (u, v) spacings from near 0 m to about 10 m.

The ACA has two antenna configurations based on a slightly modified spiral design to obtain a homogeneous (u, v) sampling of the short spacing information. Due to its compactness, the most compact configuration (inner array) will cause significant shadowing for low elevation observations. The second configuration (NS array) that occupies pads further North and South makes it possible to map sources that are only visible at low elevations.

5.2 ACA 7-m and TP Arrays during Cycle 1

5.2.1 Capabilities of ACA in Cycle 1

The ACA can only be used together with the 12-m Array in one of its four most compact configurations. The capabilities offered for the ACA during Cycle 1 are the following:

1. Number of antennas and array configurations:
 - Two TP Array antennas
 - Nine of the 7-m Array antennas
 - One 7-m Array configuration (see Section 6.4. Configuration files are available from the "documents and tools" link of the ALMA Science Portal)
2. Stand alone operation:
 - No stand alone mode of ACA is offered
3. Receivers:
 - ALMA bands 3, 6, 7 and 9
4. Spectral setups:
 - Same as those in 64-antenna correlator, see Chapter 3
5. Observations:
 - ACA will only be offered to complement observations with the more compact configurations of the 12-m Array
 - The 7-m Array interferometric observation modes are single-field or mosaicing mode
 - For the TP Array, raster OTF mapping with a single off position is the default mode. Only spectral-line observations are offered (no continuum observations)
 - WVR calibration of the atmospheric phase fluctuations for the 7-m Array will not be available
6. Observing times:
 - A fixed ratio of 1:3 between the 12m-Array and ACA is adopted for the purposes of deriving the observing time required for Cycle 1 proposals

5.2.2 Assessing the requirement for ACA data

In principle interferometric observations at a wavelength of λ with a sampling coverage which has a minimum baseline of BL_{\min} have limited response for angular scales larger than $\sim 0.6\lambda/BL_{\min}$. If the target source extends beyond these angular scales for the 12-m Array, or the accuracy of flux measurements including these extended structures are crucial for a given projects, it is recommended to use the ACA in combination with the 12-m Array observations.

In addition, the *best practice* rule of thumb for mosaicing observations is that a mosaic should be carried out if the size of the sources is greater than one third of the primary beam. Therefore, a mosaic with the 7-m antennas will be needed, instead of a single pointing, for such situations.

5.2.3 Data processing

Combination of zero and short spacing array data with 12-m Array observations is supported by CASA during Cycle 1, although with limited scope compared with future cycles. Supported data combination processes in CASA include “Feather” and “CLEAN with sky model from low-resolution image as a model”. There is no commitment to support “joint deconvolution” during Cycle 1.

Chapter 6

Cycle 1 Configurations

6.1 Introduction

The antennas of the Cycle 1 12-m Array will be staged into six distinct configurations intended to smoothly transition from the most compact (with maximum baselines of ~ 160 m) up to the most extended configuration (maximum baselines of ~ 1 km). Users will not directly choose a specific 12-m Array configuration, but rather specify the desired spatial resolution and Maximum Angular Scale (MAS) (see next section) from which observatory staff will determine the optimal configuration for the proposed observations. While knowledge of the specific array configurations is not needed for project preparation, they are discussed in this Chapter to provide more details on their characteristics.

The Atacama Compact Array (ACA), which is described in Chapter 5, will be offered in Cycle 1. The ACA will be used to fill in gaps in uv coverage at short baselines not sampled by the 12-m Array, in order to achieve high fidelity imaging of objects with large scale structures. For Cycle 1, a single configuration of the 7-m Array consisting of nine antennas will be available, as well as two 12-m single-dish antennas in the Total Power Array (note that the TPA is only offered for spectral line observations).

6.2 Maximum Angular Scales (MAS)

An interferometer is not sensitive to angular scales larger than the equivalent shortest spacing between antennas in any configuration. The MAS is proportional to the inverse of the Nyquist frequency (twice the minimum baseline for a given configuration). Since this frequency would represent a very low sensitivity to the power of a source at the minimum baseline, it is usual to use a less stringent limit so that some power is detected at that uv distance. ALMA has adopted the 10% power sensitivity level, which corresponds to 1.2 times the inverse of the Nyquist frequency. The limit is only a guideline. This gives the following approximated formula (in arcsec):

$$\text{MAS (arcsec)} \simeq 0.6 \frac{\lambda}{BL_{min}} = \frac{37100.}{BL_{min}(\text{meter}) \times \nu(\text{GHz})} \quad (6.1)$$

where BL_{min} is the shortest baseline in meters and ν the observing frequency in GHz. A list of the minimum baseline lengths for the 12-m Array configurations is given in Table 6.1.

Because an interferometer is not sensitive to angular scales larger than the equivalent shortest spacing between antennas, only a fraction of the flux of a source with a given spatial scale can be recovered. Figure 6.1 shows the results of the filtering for observations at 100 GHz of a uniform disk of 1 Jy using three of the 12-m Array configurations.

6.3 The 12-m Array configurations

The Cycle 1 12-m Array will have a set of six configurations, from the most compact (C32-1) to the most extended (C32-6). The antenna positions of the 12-m Array configuration are shown on Fig.6.2 This will provide a uv coverage with minimum baseline lengths from ~ 15 m in the most compact configuration, to maximum baseline lengths up to ~ 1000 m in the most extended configuration (see Table 6.1). For any configuration the shortest baselines constrain the MAS (see Section 6.2) and the longest baselines constrain the highest spatial resolution which can be achieved. The uv radial density distribution (i.e., sensitivity in the uv plane) of the 12-m Array configurations are shown in Fig. 6.3 and Fig. 6.4. The spatial resolution θ of an array can be estimated with the following equation but may differ slightly from the final spatial resolution after the deconvolution of the uv plane ("cleaning"):

$$\theta \text{ (arcsec)} \simeq \frac{\lambda}{BL_{max}} = \frac{61800}{BL_{max}(\text{meter}) \times \nu(\text{GHz})} \quad (6.2)$$

where BL_{max} is the longest baseline in meters and ν the observing frequency in GHz.

In Table 6.2 the range of spatial resolutions, without any tapering, and the corresponding MAS are indicated at different frequencies for the six 12-m Array configurations. The Figure 6.5 shows the spatial resolution at different frequencies for the six 12-m Array configurations.

Configuration	C32-1	C32-2	C32-3	C32-4	C32-5	C32-6
Minimum baseline (meter)	14.7	15.0	21.0	21.0	25.9	43.
Maximum baseline (meter)	165.6	303.6	442.7	558.2	820.2	1091.0
RMS (meter)	80.5	127.7	188.5	243.1	341.6	503.0

Table 6.1: Basic parameters of the 12-m Array configurations during the Cycle 1.

Band	Freq	C32-1		C32-2		C32-3		C32-4		C32-5		C32-6	
		AR	MAS										
3	100	3.7''	25''	2.0''	25''	1.4''	17''	1.1''	17''	0.75''	14''	0.57''	8.6''
6	230	1.6''	11''	0.89''	11''	0.61''	7.6''	0.48''	7.6''	0.33''	6.2''	0.25''	3.7''
7	345	1.1''	7.1''	0.59''	7.1''	0.40''	5.0''	0.32''	5.0''	0.22''	4.1''	0.16''	2.5''
9	675	0.55''	3.6''	0.30''	3.6''	0.21''	2.6''	0.16''	2.6''	0.11''	2.1''	0.08''	1.3''

Table 6.2: Angular resolution (AR) and Maximum Angular Scale (MAS) for the six Cycle 1 12-m Array configurations. *Notes to Table:* **1.** Computation for source at zenith. For sources transiting at lower elevations, the North-Source angular measures will increase as the $\sin(\text{elevation})$. **2.** All angular measures scale inversely with observed sky frequency. **3.** "Maximum Angular Scale" is the largest angular scale that can be observed effectively. If the source contains smoothly varying structures that are larger than this in both directions, those components will be "resolved out". This is the well known "missing flux" problem intrinsic to interferometry.

For the compact configurations of the 12-m Array and for the 7-m Array, shadowing of the antennas will be an issue at low elevations. Figure 6.6 shows the shadowing fraction for the two most compact configurations (C32-1 and C32-2) of the 12-m Array during a 2-hour observation near transit (± 1 h HA) as a function of source declination. Since shadowing is dependent on the elevation of the source it becomes more prominent for longer observations around transit. The plot shows that shadowing will not be a major issue for the 12-m Array, except for sources that transit at very low elevations and with long observing times.

For reference, we show in Figure 6.7 a few examples of the uv coverage of the most compact and extended configurations of the 12-m Array for a source at a declination of -30° . For Cycle 1, the aim will be to obtain the requested resolution using observations from a single 12-m Array configuration only (See Appendix A in the Proposer's Guide). Each configuration results in a different beam-size and shape. In Fig. 6.8 the beam shape is displayed as function of source declination for an observation of two hours at 100 GHz (± 1 h HA).

6.3.1 Beam Maps

Beam maps can be produced as simulation output and inspected for each configuration, declination and length of observing track (here: $\text{DEC} = -35^\circ$ and $\text{HA} = \pm 1.5$ hours). Examples for 7-m and 12-m array configurations are given in Figure 6.9,

6.4 The ACA and 12-m Array combination

The ACA in Cycle 1 will have only one configuration, with minimum and maximum baselines of 8.9 m and 32.1 m, respectively (see Table 6.3 and Fig. 6.10).

	ACA
Minimum baseline (meter)	8.9
Maximum baseline (meter)	32.1
RMS (meter)	19.6

Table 6.3: Basic parameters of the ACA 7-m Array configuration during Cycle 1.

Given the short distance between the antennas of the ACA, a key issue is shadowing. Figure 6.6 shows that sources with declinations $< -75^\circ$ or $> 25^\circ$ are significantly affected by shadowing ($> 40\%$) within the range of ± 1 h HA, and should be observed within a more narrow hour angle range.

If a project needs to detect structures with spatial scales larger than the MAS for the 12-m Array, it will need to be observed in conjunction with the ACA to retrieve the short baselines. The integration time on both array will be different to combine properly both uv distribution and these time ratios are currently under investigation.

If the ACA is combined with the most compact 12-m Array (C32-1) the integration time ratio between both arrays should be ~ 3 . This allows the overlapping uv coverage to be matched in both arrays. For projects that only use the two most extended 12-m Array configurations, ACA observations will not be available because of the small uv plane overlap.

Table 6.4 summarizes the properties for the ACA+12-m Array observations for the four bands offered in Cycle 1.

Band	Freq.	C32-1 & ACA		C32-2 & ACA		C32-3 & ACA		C32-4 & ACA	
		AR	MAS	AR	MAS	AR	MAS	AR	MAS
3	100	3.7''	44''	2.0''	44''	1.4''	44''	1.1''	44''
6	230	1.6''	19''	0.89''	19''	0.61''	19''	0.48''	19''
7	345	1.1''	13''	0.59''	13''	0.40''	13''	0.32''	13''
9	675	0.55''	6.5''	0.30''	6.5''	0.21''	6.5''	0.16''	6.5''

Table 6.4: Angular resolution (AR) and Maximum Angular Scale (MAS) of the combined ACA 7-m Array and 12-m Array for the four bands offered in Cycle 1. The maximum baseline of the 12-m Array configuration sets the AR, while the minimum baseline of the ACA 7-m array sets the MAS.

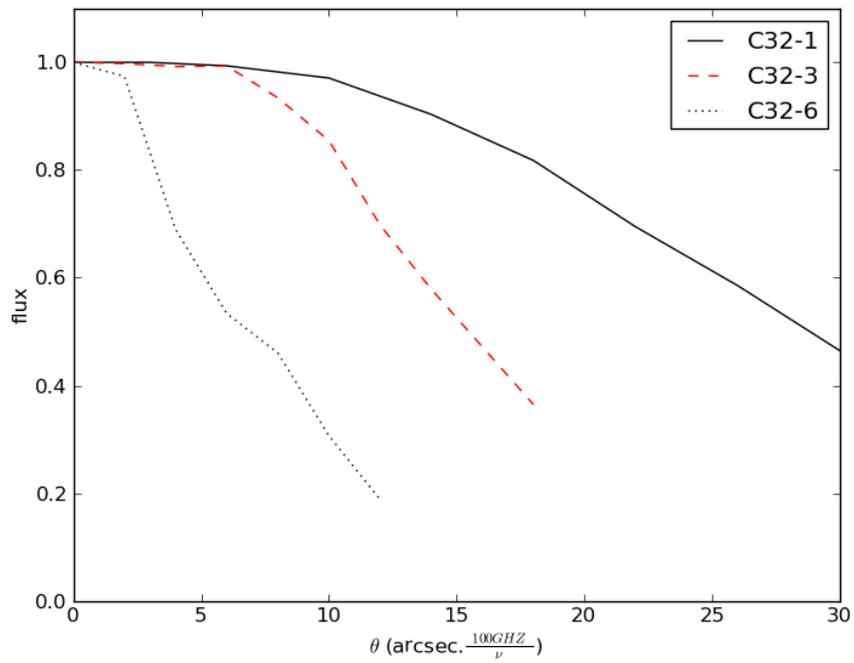


Figure 6.1: Recovered flux of a uniform disk, with a total flux of 1 Jy, observed at 100 GHz with three different 12-m Array configurations, as a function of the disk size.

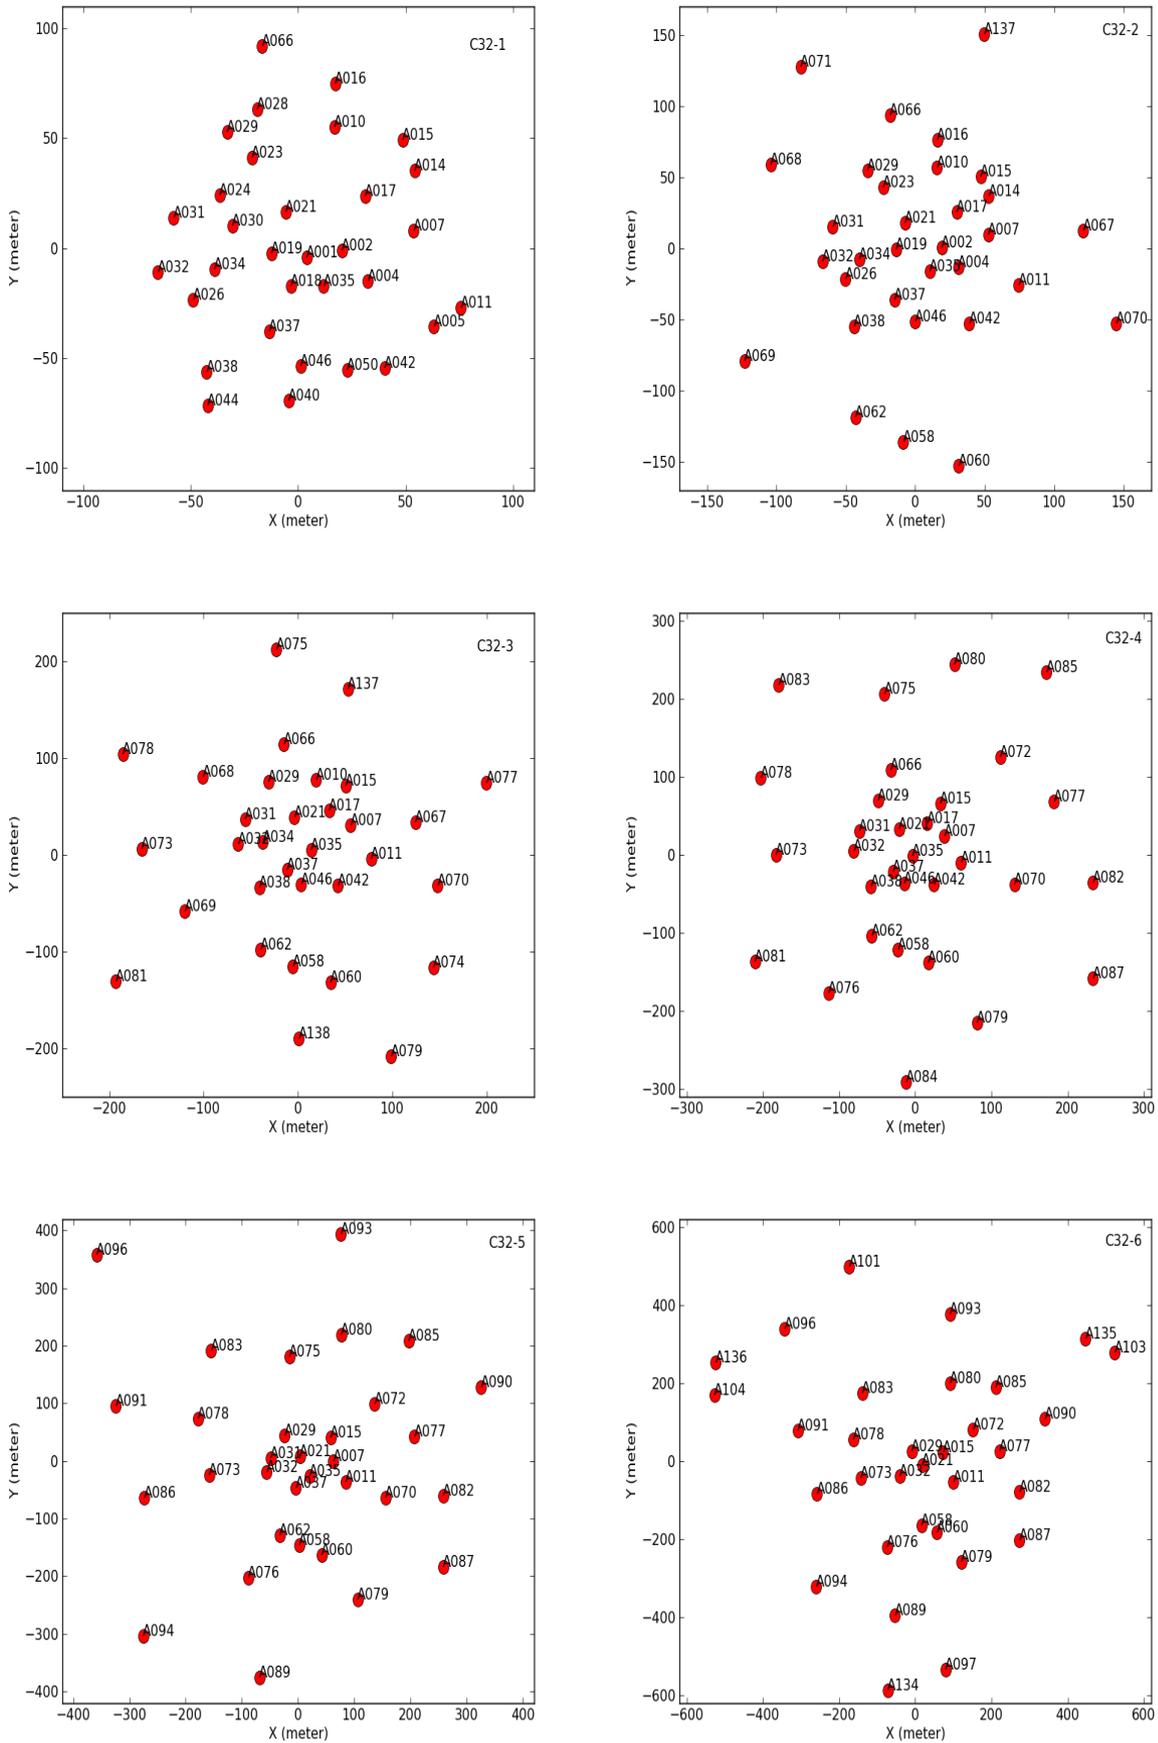


Figure 6.2: Antenna positions for the 6 12-m Array Configurations with the corresponding pad names.

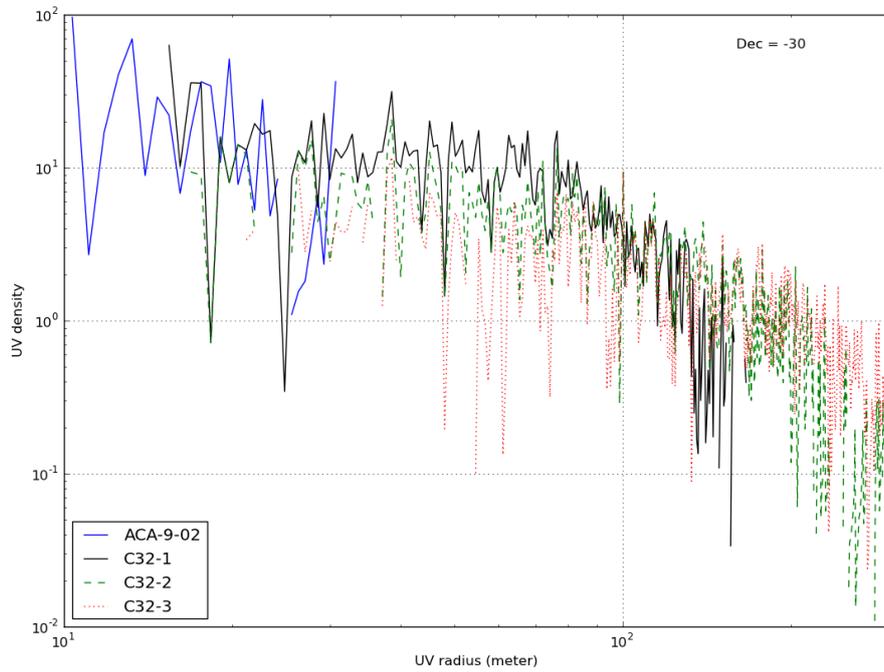


Figure 6.3: UV radial density distribution for the most compact configurations including the ACA

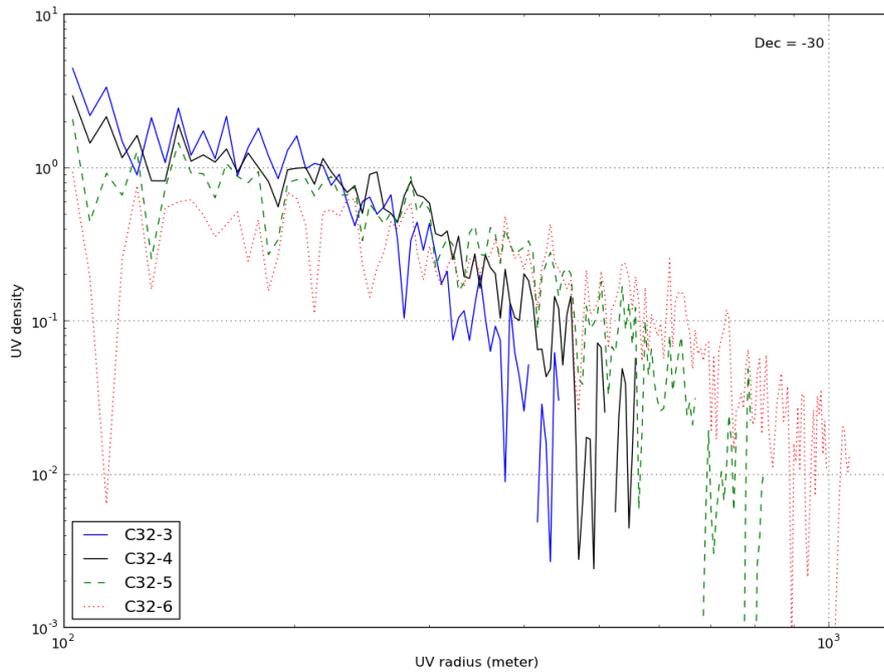


Figure 6.4: UV radial density distribution for the most extended 12-m Array configurations.

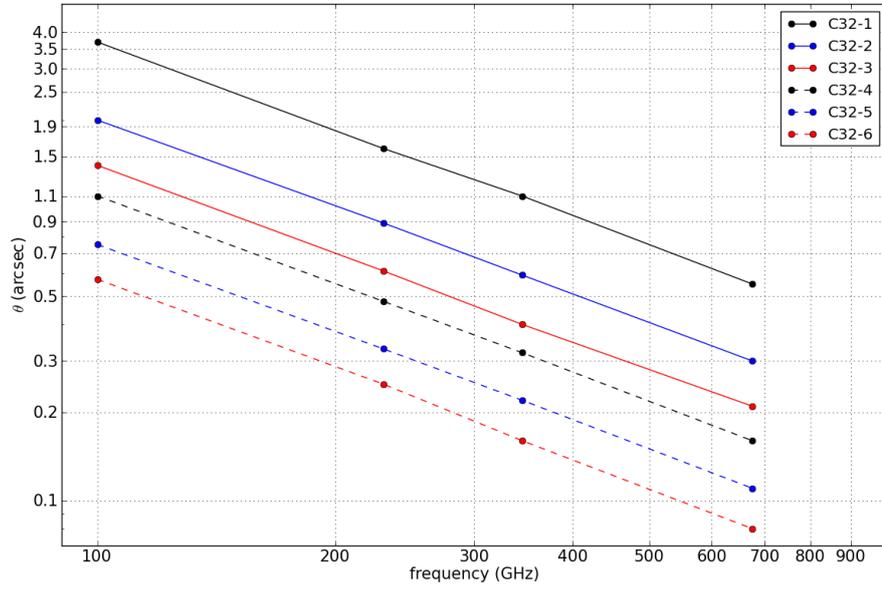


Figure 6.5: Spatial resolution in arcsec at different frequencies for the six 12-m Array configurations.

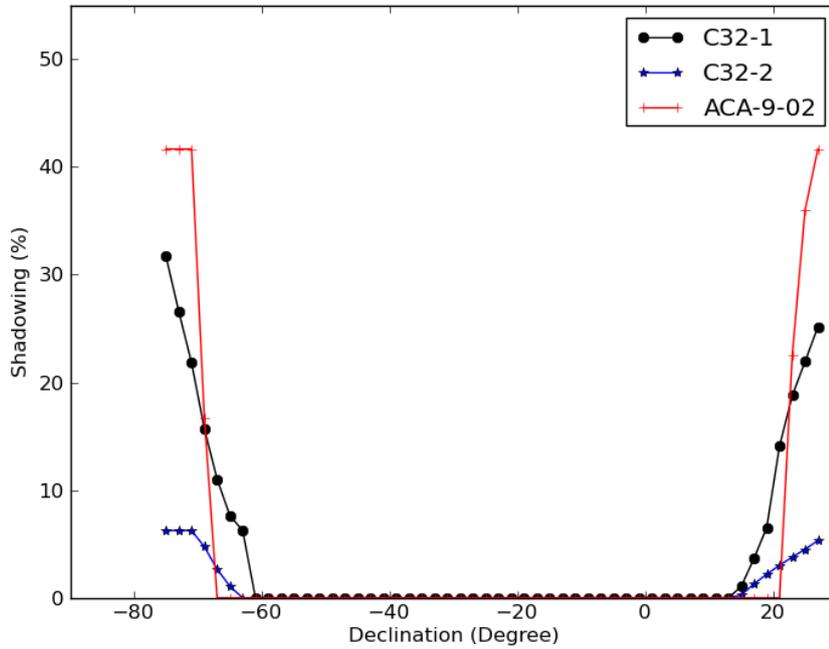


Figure 6.6: Shadowing fraction vs. Declination for the two most compact configurations of the 12-m Array and for the 7-m Array with a track duration of 2 hours (± 1 h HA).

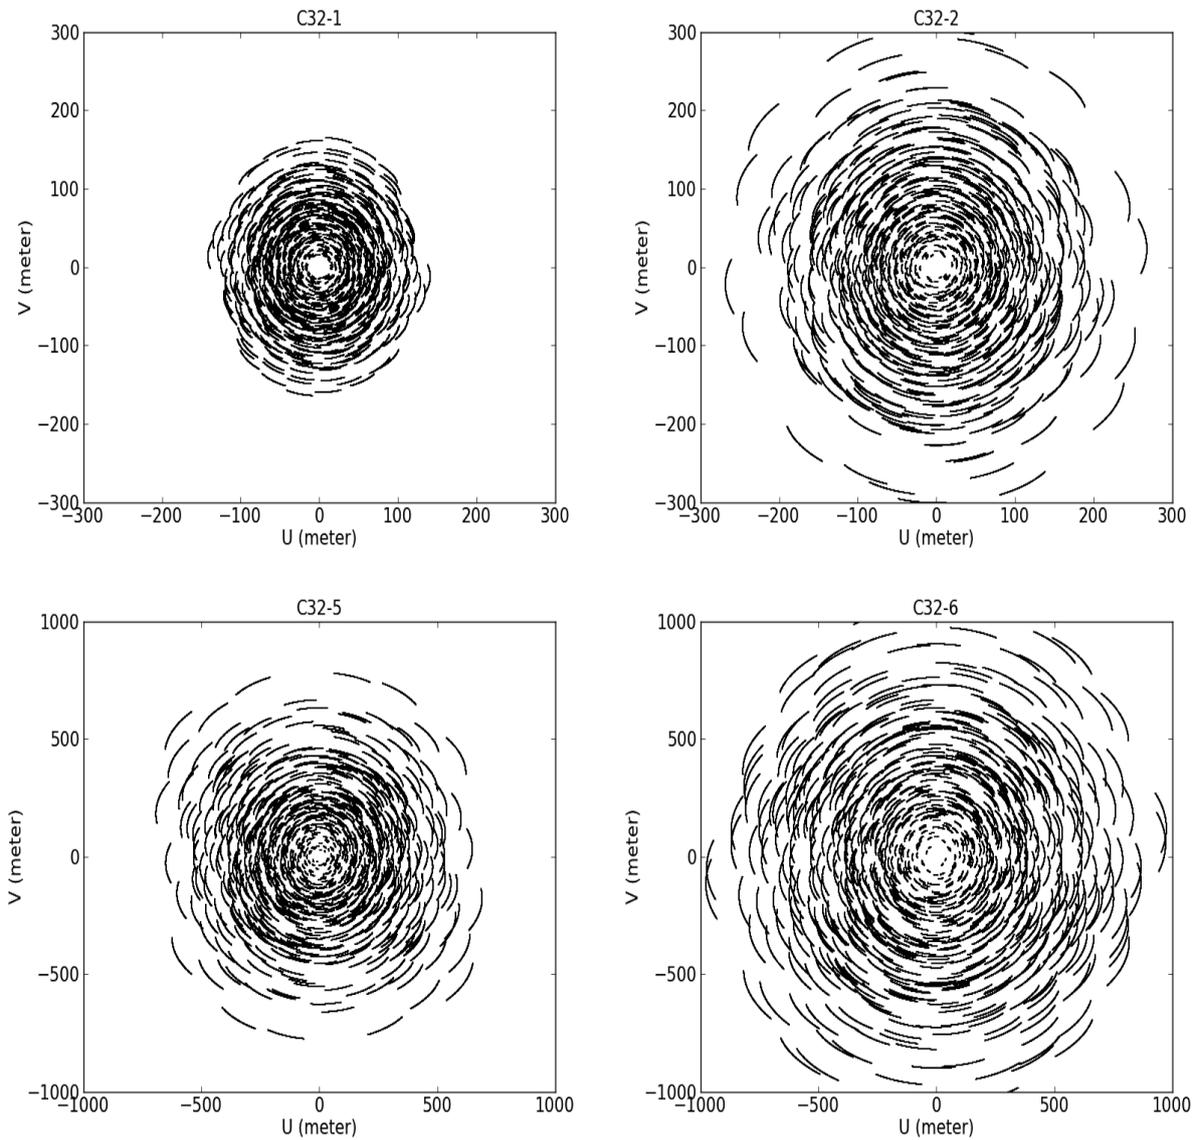


Figure 6.7: uv-coverage for the two most compact (C32-1 and C32-2) and the two most extended (C32-5 and C32-6) 12-m Array configurations with a 2 hour integration (± 1 h HA) of a source at a declination of -30° .

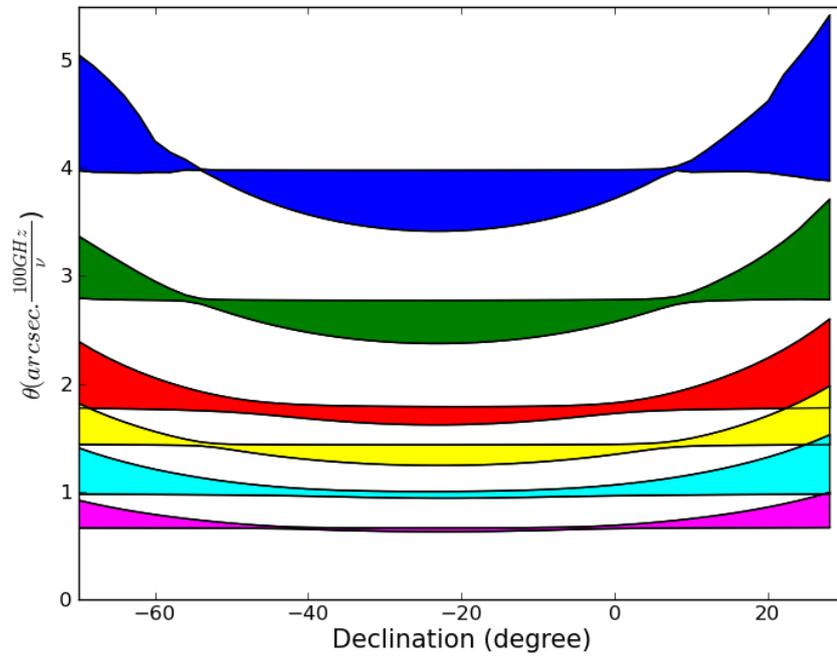


Figure 6.8: Synthesized beam radii vs. Declination for an observation of 2 hours (± 1 h HA) at 100 GHz. The 6 12-m Array configurations are shown with different colors .starting with the most compact configuration (C32-1) in blue at the top, progressing through the more extended configurations. The width of the area gives an indication of the ellipticity of the beam (the wide the area, the more elliptical the beams are).

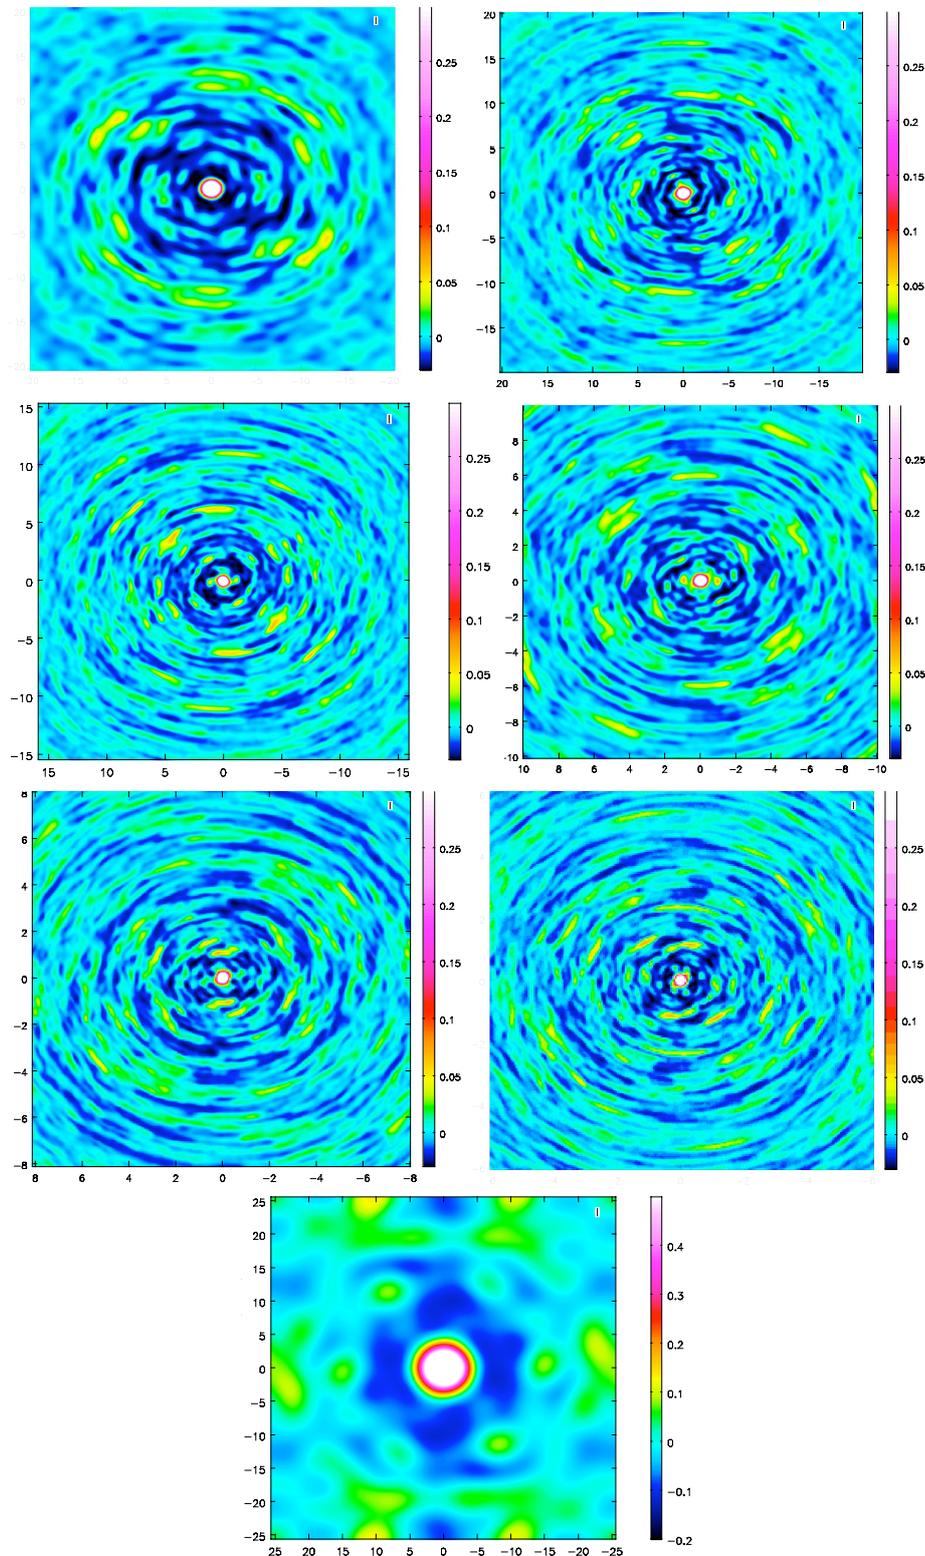


Figure 6.9: Maps of the synthesised beams for the six representative ALMA Cycle 1 configurations C32-1, C32-2, C32-3, C32-4, C32-5, C32-6 and the ACA 9 antenna configuration. Note: the axis (offset angle in arcseconds) vary between the different panels to show areas scaling roughly with the synthesized beam width. All maps are for Band 7 at 300 GHz, Dec= -35° , and an hour angle range of ± 1.5 hours

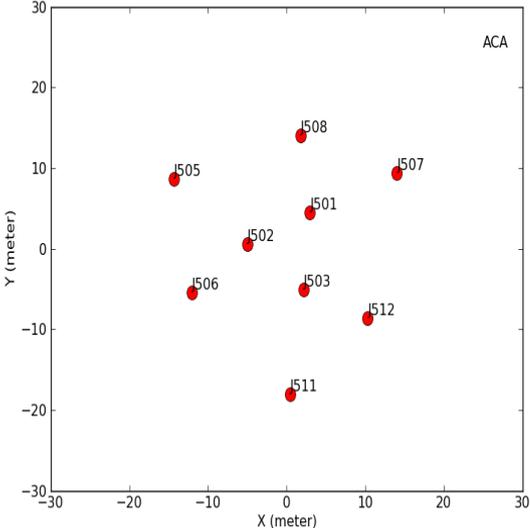


Figure 6.10: Antenna positions for the 7-m Array with the corresponding pad names.

Chapter 7

Simulations of Cycle 1 observations

7.1 Introduction

Simulations of ALMA data can help PIs to present a convincing technical case, which argues what the proposed ALMA observation in the chosen technical setup can achieve in terms of reaching the science goals. Also, simulations can be helpful in finding the optimal way to reduce the data and give valuable guidelines on how to interpret the results.

7.1.1 The ALMA Simulation Tools

The following two software tools are available to help users simulate images resulting from an ALMA observation:

- The CASA task `sim_observe`, which is described in the CASA Manual and from the website at http://casaguides.nrao.edu/index.php?title=Main_Page#Simulating_Observations
- The ALMA Observation Support Tool (OST), which is available at the website <http://almaost.jb.man.ac.uk/>

Both simulators use the `sm` toolkit package in CASA, and are available from links within the ALMA Science Portal (under Documentation and Tools). The CASA task `sim_observe` (in CASA version 3.3.0, and the `simdata` task in earlier versions) turn a model of the sky brightness (2 to 4 dimensions including frequency and Stokes) into the visibilities that would be measured by a given configuration of ALMA. The `sim_analyze` task (or `simdata` in CASA versions 3.3.0 and earlier) can also produce deconvolved (`cleaned`) images of the model visibilities, make a comparison with the input image convolved with the synthesized beam, and calculate a fidelity image¹.

In addition to these tasks, the CASA 3.4 release includes a representative set of ALMA configurations that span the advertised Cycle 1 capabilities (these configurations are also available on the ALMA Science Portal). This release also contains improvements to the algorithms for simulating mosaics, as well as better approximations for the ALMA primary beam shapes, and thus is to be preferred for making Cycle 1 mosaic simulations. Also, since CASA version 3.4.0 there is a configuration selector for ALMA Cycle 1, which allows to specify an angular resolution (example: `antennalist="alma_cycle1;0.5arcsec"`) and the corresponding configuration will be used to simulate the observation (Note that in CASA version 3.4.0 the simulation tasks are called `simobserve` and `simanalyze`).

With `sim_observe` one can add thermal noise (from receiver, atmosphere, and ground) and polarization leakage effects to the visibilities. It also uses an atmospheric model to calculate atmospheric signal corruption

¹ Image fidelity = $I(x) / |I(x) - T(x)|$, where $I(x)$ is the observed image intensity and $T(x)$ is the true image intensity given the sky model. See also ALMA Memo No. 272, available at <http://www.alma.cl/alمامemos/>.

effects (noise) as a function of frequency and site characteristics. The `sm` CASA toolkit can be used to simulate phase variations, gain fluctuations and drift, and bandpass errors. `sm` also has more flexibility in adding thermal noise than `sim_observe`. However, note that no general guidelines on suitable parameters for corrupting the simulated data, in particular for the longer baselines around 1 km available in Cycle 1, can be given yet.

The OST is a web-based interface to an ALMA simulator tool maintained by the EU ARC node in Manchester (UK) and at ESO. Like `sim_observe`, it is based on the CASA `sm` toolkit, but uses different wrapper scripts, and in particular, interpolates the atmospheric opacity as a function of frequency from a grid, whereas `sim_observe` calculates an atmospheric profile using J. Pardo's ATM library and average site conditions, altitude=5000 m, ground pressure=650 mbar, relative humidity=20%, and a water layer of the PWV provided by the user, `'user_pwv'`, at an altitude of 2 km above ground. Thus, slight differences in atmospheric noise between the tools can result.

Note that significant differences sometimes exist between the noise predicted by the ALMA Sensitivity Calculator and the measured RMS in the simulated images. These differences are primarily because the RMS measured in an image depends sensitively on the details of how the image is deconvolved. **The ALMA sensitivity calculator will be used for the technical assessment of ALMA proposals.** Therefore, before using the results from the simulators as a basis of a Technical Justification in an observing proposal, the RMS in the simulated image should be examined and the thermal noise part, i.e. not the imaging related part of the RMS, should be brought into agreement with the values given by the Sensitivity Calculator (see Chapter 8). After producing an initial noise simulation, in the thermal noise dominated regime, either `'simplenoise'` in flux density units or `'tsys-manual'` with a matched atmospheric brightness temperature can be used for this adjustment.

7.1.2 Basic steps of an ALMA simulation

After installing CASA or entering the OST website, a sky model (which may be any model or image of the target) can be imported and modified with a set of parameters, re-scaling spectral and spatial coordinates and the brightness. The mosaic pointings are calculated and can be saved in a text file, which can be altered or produced externally. Visibilities are calculated for a specific configuration of the ALMA array (input file) on a specific day. The measurements can be **corrupted** by adding thermal and phase noise, cross-polarization, etc. The same atmospheric conditions (as parametrized by the Precipitable Water Vapor PWV), as used by default in the sensitivity calculator should be used where noise in the map matters for image reconstruction (it should only be used to assess sensitivity qualitatively; proposed observing times should be based on the radiometer equation (see Chapter 8). Around the ALMA band centers the following rough PVW values can be adopted for computations: 2.7 mm (Band 3), 1.8 mm (Band 6), 1.3 mm (Band 7) and 0.7 mm (Band 9). In spectral areas with strong atmospheric attenuation, more stringent numbers should be used as listed in the 'Sensitivity Calculator'. A `cleaned` image can be produced from the visibilities, and `sim_analyzed` to give `clean` residual, input/output difference or fidelity images. An example output from the `sim_observe` / `sim_analyze` tasks is given in Figure 7.1 showing various plot options.

This example also demonstrates that the default `clean` within `sim_analyze` may not always be sufficient for a high quality reconstruction of the image. In Figure 7.2 the simulated measurement set is `cleaned` further with the multi-scale option, more iterations, suitable masking, and a smaller loop gain to yield a better reconstruction.

The output dataset and image will have the same channelization as the input image. Therefore, the number of spectral channels should be chosen carefully for line observations to avoid simulating a massive cube with only a few relevant planes. Furthermore, to be able to compute the image fidelity for continuum simulations, it may be necessary to make several data sets sampling the available bandwidth and concatenate and image the data separately.

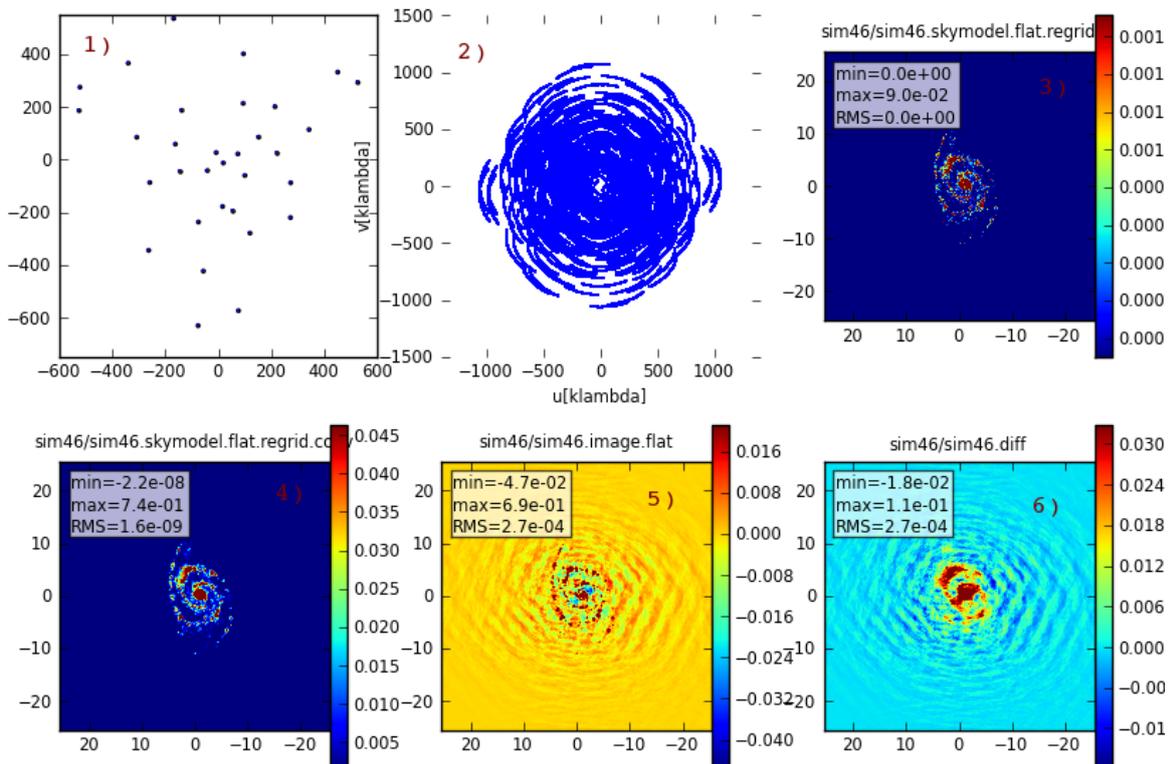


Figure 7.1: CASA `sim_observe` and `sim_analyze` example showing in the top row 1) the antenna positions of the configuration, 2) the observed tracks in uv , 3) the input model image, and in the bottom row 4) the input model convolved with the Gaussian beam width of the configuration, 5) the observed and `cleaned` image, and 6) the difference between convolved input and observed image as outputs. See Section 7.1.2 for details on simulation input and parameters. The axis units are for 1)+2) in kilo-lambda (here for 1 mm observations equal to meters) and for 3)-6) in arcsec (offset angle). The brightness units are for 3) in Jy/beam and for 4)-6) in Jy/beam.

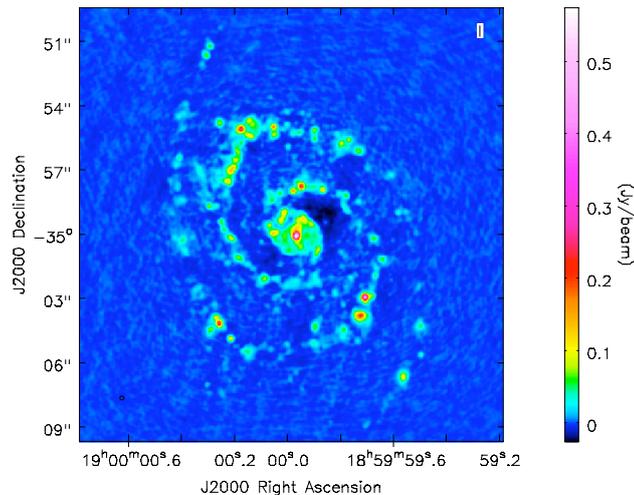


Figure 7.2: Example of `clean` with the multi-scale option on the simulated dataset shown in Figure 7.1. The image quality is vastly improved over the reconstruction shown in panel 5). See Section 7.1.2 for details on simulation input and parameters.

7.2 Simulation Examples for Cycle 1 Configurations

The JAO provides representative configurations for simulation purposes to the community. These configurations give an indication of what resolution and image quality (based on uv sampling) can be expected from Cycle 1 observations. The real observations will not necessarily be taken in the exact same configurations.

7.2.1 Beam Maps

Beam maps can be produced as simulation output and inspected for each configuration, declination and length of observing track (here: $\text{DEC} = -35^\circ$ and $\text{HA} = \pm 1.5$ hours). Examples for ACA and C32-5 are given in Figure 7.3, and examples of the angular resolutions for all Cycle 1 configuration are given in Chapter 6.

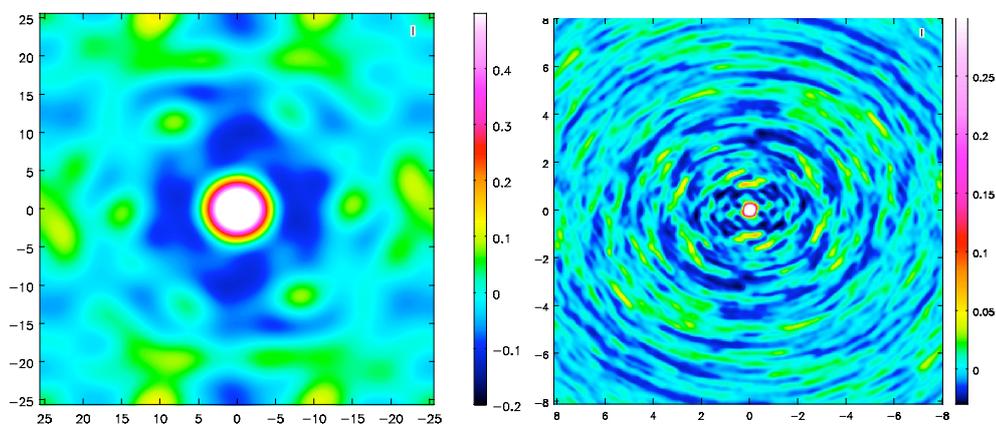


Figure 7.3: Maps of the synthesized beams for two simulated observation examples in Band 7 at 300 GHz, $\text{Dec} = -35^\circ$, and an hour angle range of ± 1.5 hours in the ACA 9 antenna configuration and in the C32-5 configuration. Note that the axis (offset angle in arcsec) are different in the two images.

7.2.2 An M51-like Galaxy

The following simulation examples, Figures 7.4, 7.5 and 7.6 (see also the previous Figures 7.1 and 7.2), use the ALMA Cycle 1 configurations described in Chapter 6 and a galaxy template: M51 ($H\alpha$) $4.3' \times 4.3'$ image down scaled in size by a factor of 10 (equivalent to placing it at $z \approx 0.015$), band 7 at 300 GHz (1mm), Dec = -35° and 3 hour (± 1.5 hours) observation.

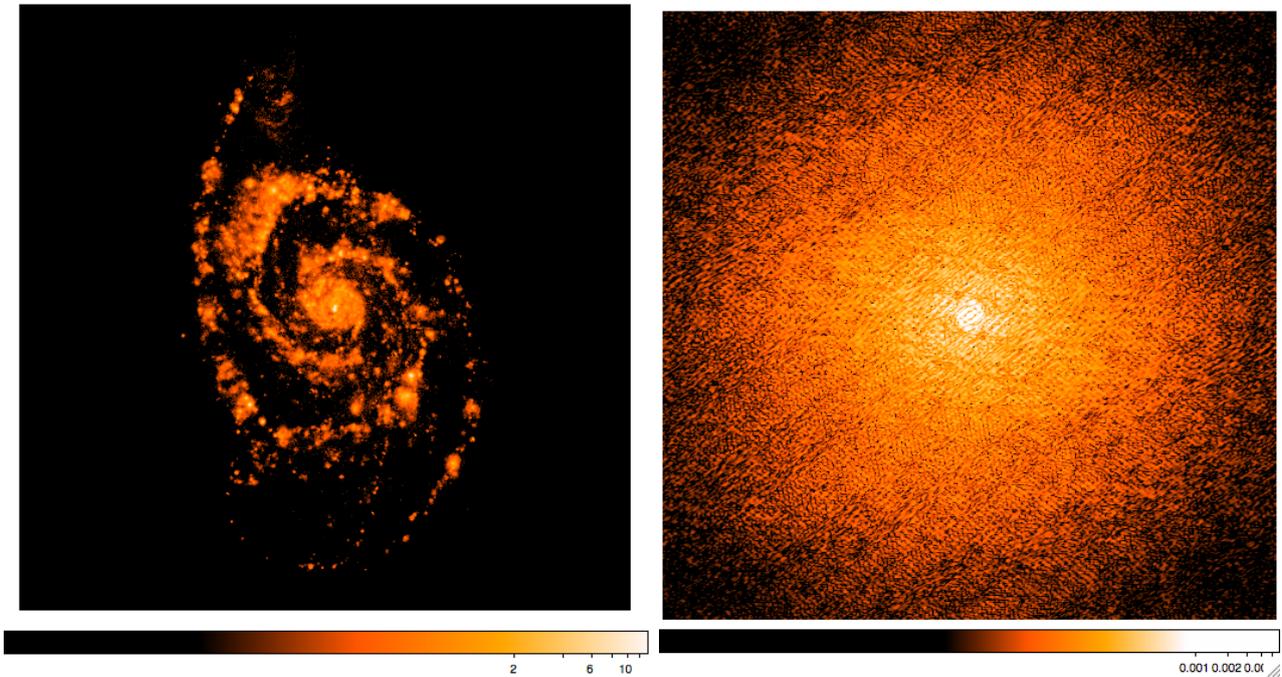


Figure 7.4: M51 $H\alpha$ image as input for the simulations (left) and the Fourier transform of the image in uv space (right).

7.2.3 A Toy Model Example

In some cases proposers may have neither physical maps nor numerical simulation outputs available to base their ALMA imaging expectations upon. In these cases even simple toy models can be informative in the pre-observation process.

The model used here for demonstration contains four point sources of flux densities 1, 0.5, 0.1 and 0.01 Jy, and four Gaussian functions with widths (FWHM) of 0.3, 2, 10 arcsec (peak flux density 0.01 Jy) and 40 arcsec (0.005 Jy), see Figure 7.7.

It can be seen that quite different structures are picked up in the different configurations. While the point sources are seen with the same flux density in all configurations, the extended Gaussian structures either dominate the surface brightness distribution for the large synthesized beams from the shorter baselines or are resolved out by the longer baselines with the small synthesized beams.

7.2.4 Combined ACA + 12m Observations

A comparison between the M51 input model image and a simulated combined ACA + 12-m Array observation is shown in Figure 7.8. The output image is synthesized from the complete range of angular scales available in Cycle 1.

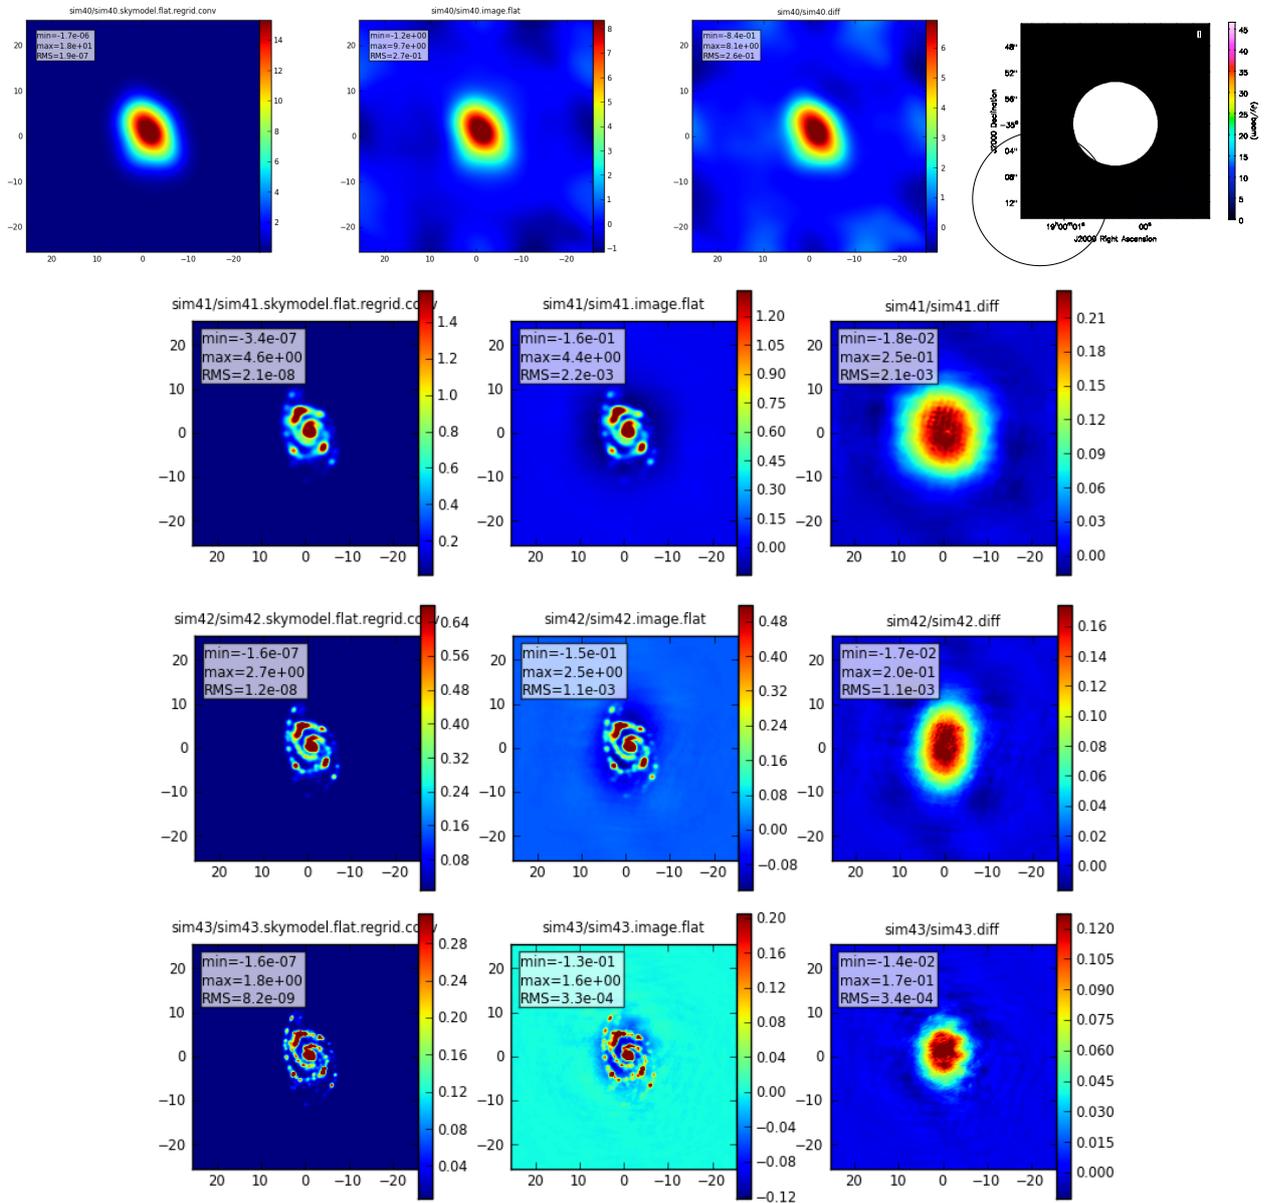


Figure 7.5: Top to bottom rows are for the ACA (7-m array + total flux density measurement from the ACA 12-m total power antennas), C32-1, C32-2 and C32-3 configurations. Right to left panels: Gaussian beam convolved input model; simulated and cleaned image; difference (convolved model - simulated). The simulations are noiseless to show primarily the rms resulting from dynamic range limitations and the angular scale filtering properties of the different configurations. All simulations are for Band 7 at 300 GHz, Dec=-35 deg, and an hour angle range of ± 1.5 hours. Brightness units are in Jy/beam.

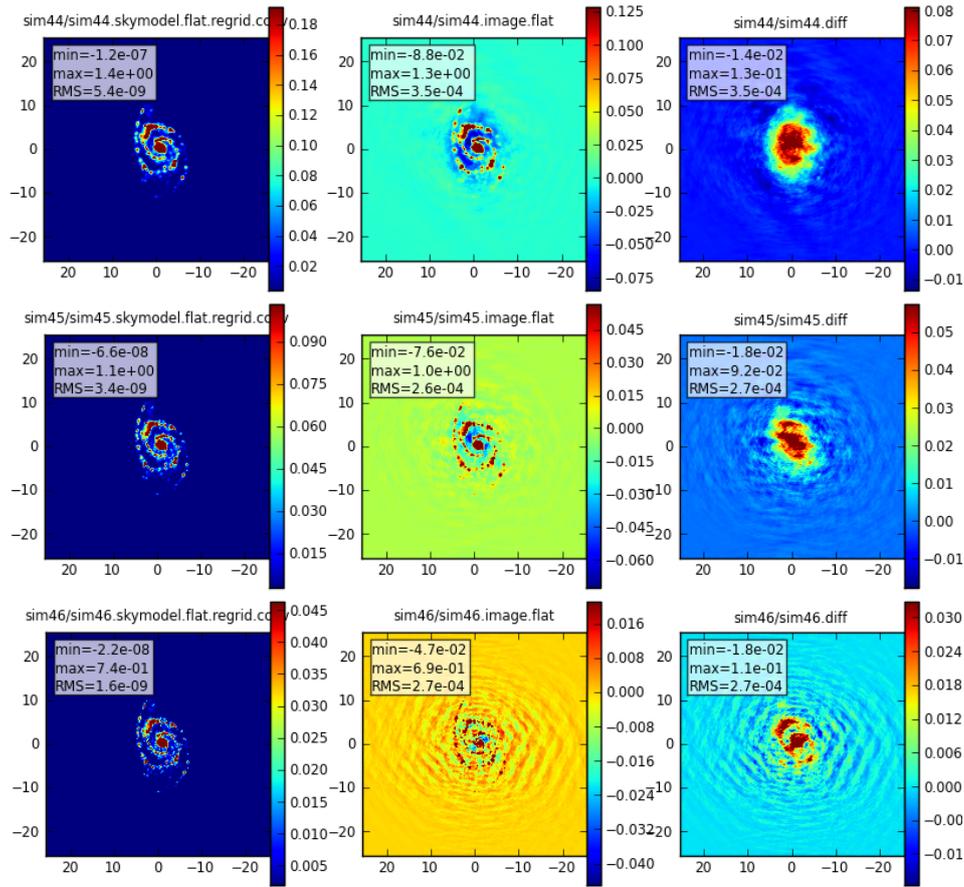


Figure 7.6: Top to bottom rows are for the C32-4, C32-5 and C32-6 configurations. Right to left panels: Gaussian beam convolved input model; simulated and cleaned image; difference (convolved model - simulated). The simulations are noiseless to show primarily the rms resulting from dynamic range limitations and the angular scale filtering properties of the different configurations. All simulations are for Band 7 at 300 GHz, Dec=-35 deg, and an hour angle range of ± 1.5 hours. Brightness units are in Jy/beam.

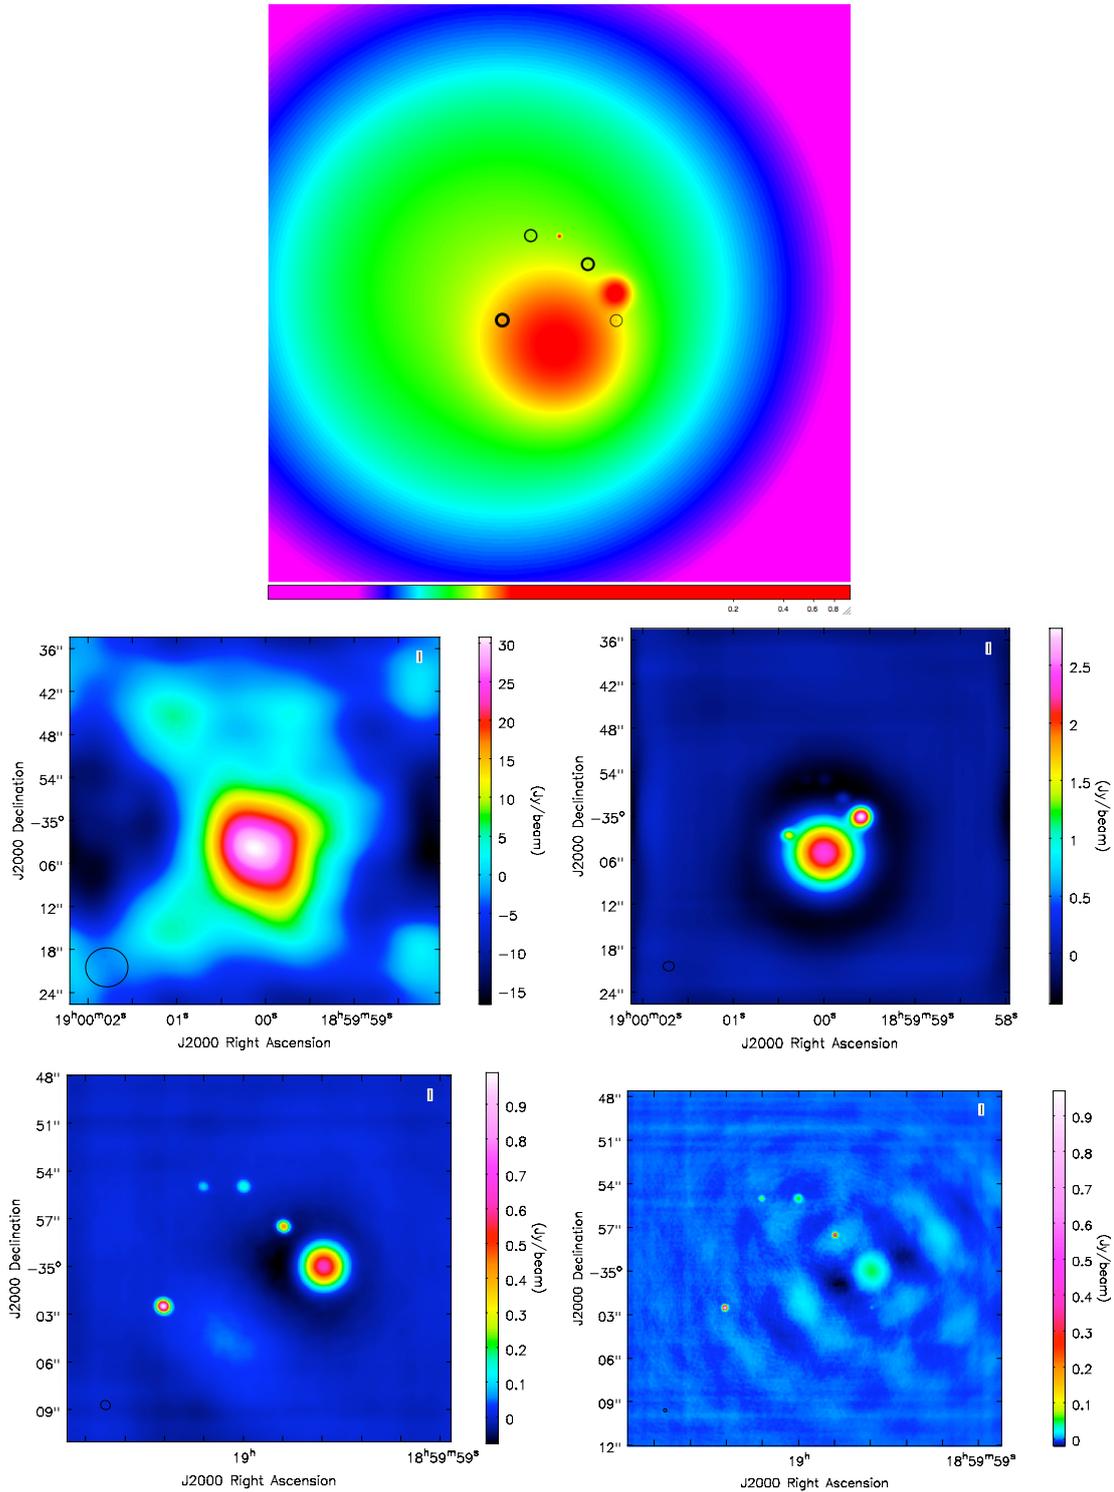


Figure 7.7: The top panel represents the true sky brightness distribution: 4 Gaussians and 4 point sources (circled in black). The middle left and right panels show observations of the sky brightness distribution with the ACA and C32-1 arrays. The bottom left and right panels show observations of the sky brightness distribution with the C32-3 and C32-6 arrays. The simulation parameters are identical to the M51 simulation, i.e. Band 7 at 300 GHz, Dec=-35 deg, and an hour angle range of ± 1.5 hours. The brightness units are Jy/pixel for the input and Jy/beam for the output image. The top three panels show the same sky area, whereas the bottom two panels are zoomed into the center area.

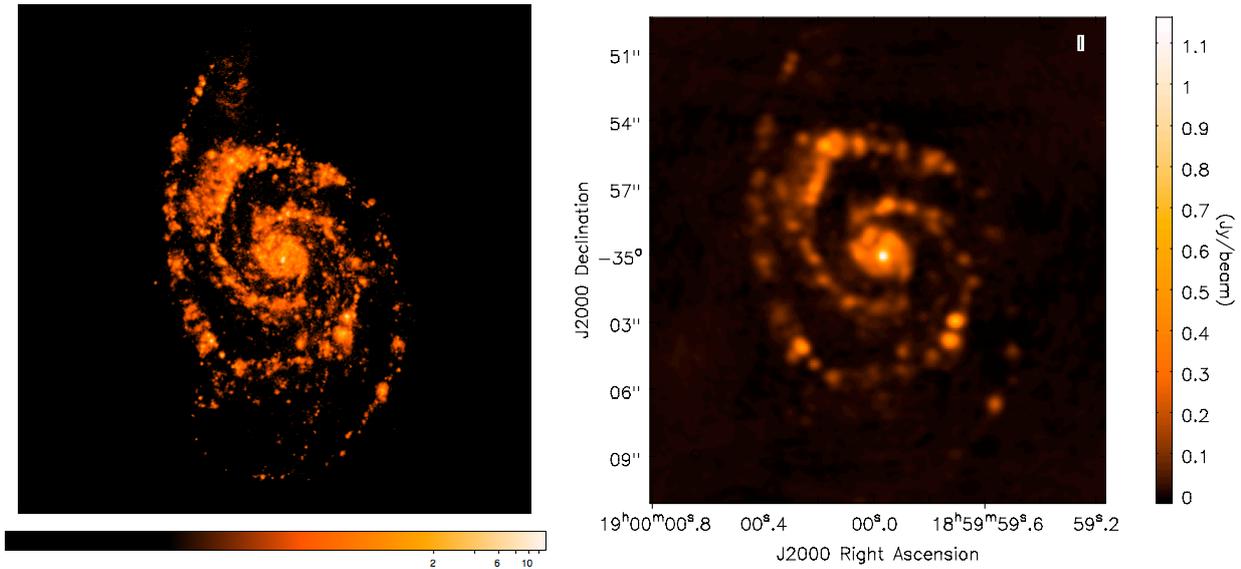


Figure 7.8: M51 input image (left) compared to a simulated image synthesized with data from combining ACA 7-m and 12-m Array C32-4 configurations (right). The brightness units are Jy/pixel for the input (and 'inbright', the model peak brightness, had been set to 0.1 Jy/pixel) and Jy/beam for the output image.

The addition of ACA observations with 12-m Array observations provides better imaging of complex sources with structures over a wide range of angular scales. The TP Array measures signals on angular scales not measured by the 12-m Array, while the 7-m array provides measurements on spatial scales intermediate between the two. The resulting reconstructed images can provide a more accurate representation of the source and a better measure of the total flux density; see Figure 7.9 for an illustrative example.

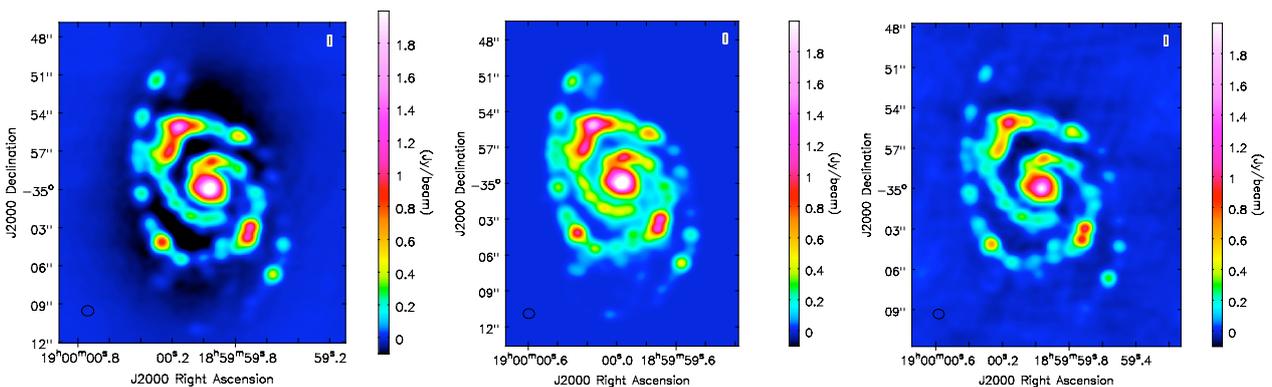


Figure 7.9: C32-2 without the ACA 7-m array (left), Gaussian beam convolved input model (center), C32-2 with the ACA 7-m array (right).

The CASA tasks available to combine different interferometric data sets and to add single dish data are `concat`, `clean`, and `feather`. The task `concat` makes a new measurement set (single file) from a list of given measurement sets. Also `clean` can take more than one measurement set as input. Concatenated measurement sets can be imaged with `clean`, or the task `feather` can be used to combine images with a specified weighting after a Fourier transform in uv space. The latter option can also be used to combine interferometric data with single dish images. Detailed recommendations on optimally combining Cycle 1 data taken with the 12-m Array, ACA and single dish will be updated on the CASA guides website.

7.3 Sensitivity to Angular Scales

7.3.1 Filter Functions

The relative response to different angular scales, also called the filter function in uv space (relevant here: radial distance from the origin), for the different Cycle 1 configurations can help to design the observation, see Figure 7.10. A comparison of the angular sensitivity with the intensity of the source as a function of angular scale, whenever available for the science target, is particularly informative. In the OT a range of angular scales can be given and the observatory will then carry out the observation based on these specifications.

7.3.2 Resolved Out Flux

Simulations can provide Informative results on the fraction of the flux that is resolved out (missing) by a given configuration from an extended structure, see Figure 7.11. The cause is that the observed structure has signal on angular scales to which the (too long) baselines of the configuration are not sensitive, see Figure 7.10.

7.4 Image Quality

In order to characterise the image quality required to achieve the science goal of a project, the fidelity and/or dynamic range can be used.

7.4.1 Image Fidelity

A fidelity image can be produced with CASA `sim_analyze` and various statistics (peak, average in a certain region, etc.) can be extracted and if relevant for the proposal be specified in the technical justification. An example of a fidelity image is given in Figure 7.12, which is taken from the M51 C32-1 simulation image reconstruction.

7.4.2 Dynamic Range

In Early Science Cycle 1 the dynamic range is not expected to be within the full ALMA requirements of > 1000 for all observing situations. If a high dynamic range (peak flux / sensitivity rms in the Observing Tool) is required in order to achieve the science goals, simulations can establish whether the observation is feasible. An example of a case where the dynamic range matters is given in Figure 7.13. While most of the structure accessible with the baselines of C32-6 is recovered even in a snapshot observation with ALMA, the detection of a faint point source near an extended Gaussian function of sky brightness is only possible with a longer observation, which reduces the sidelobe level and the resulting confusion in the de-convolved image.

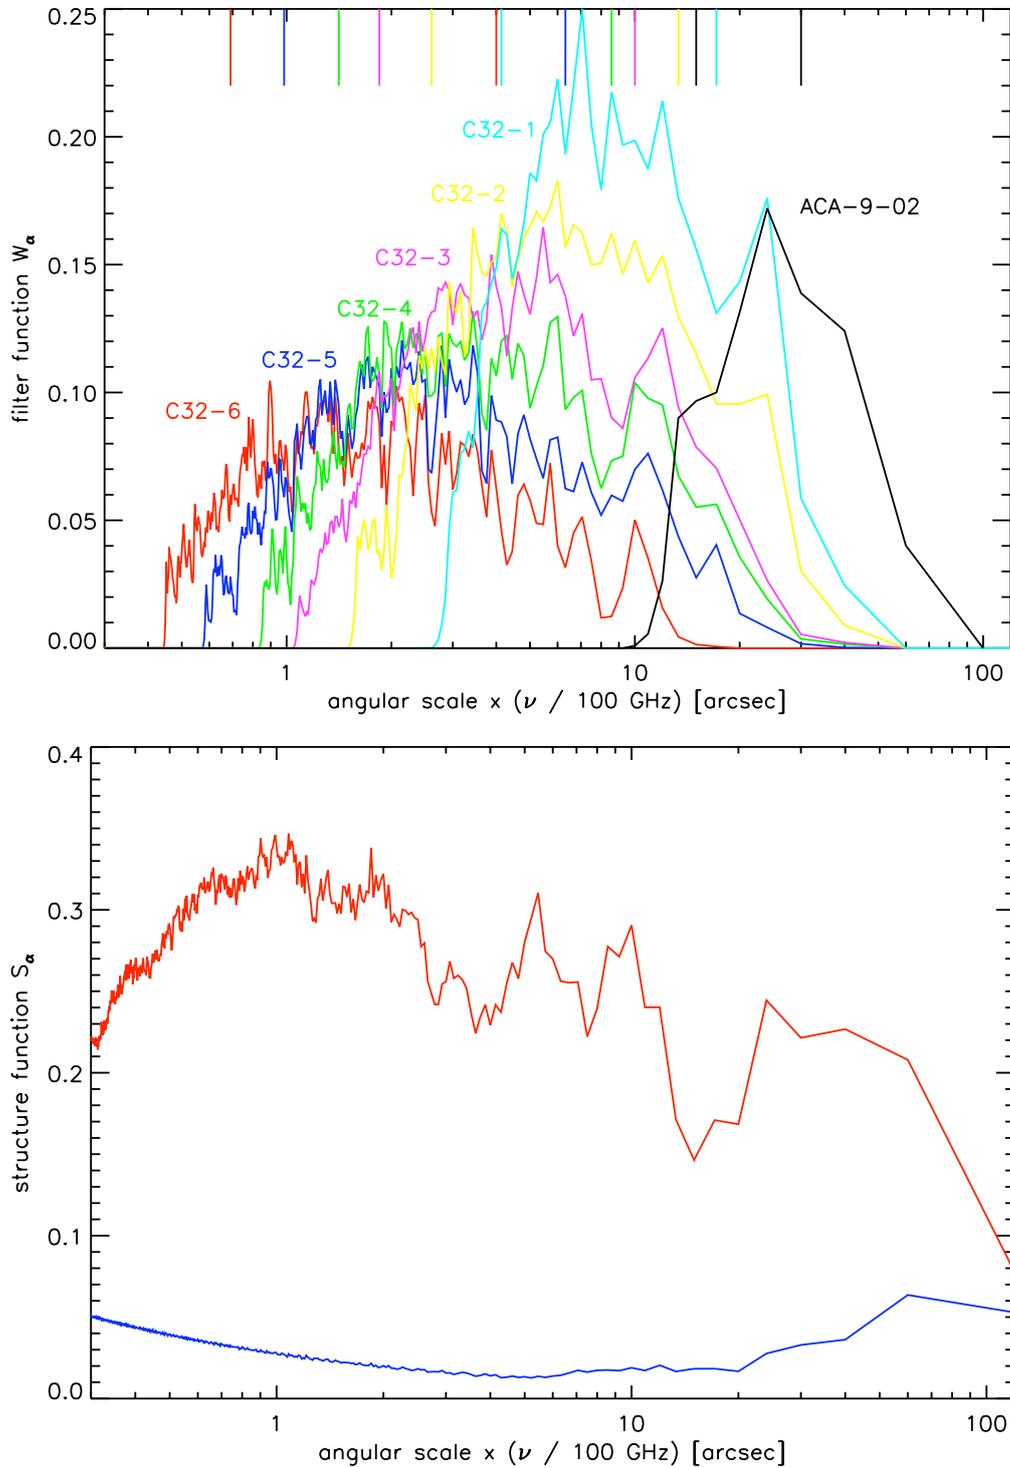


Figure 7.10: Top: Sensitivity to angular scales (in relative units) for the different configurations C32-6 to C32-1 and ACA, from left to right, based on the uv tracks of the Dec = -35 deg, 3 hours (± 1.5 hours), observation. The 25% and 95% percentiles in the visibility (baseline) scale distribution are shown at the top of this panel. Bottom: Intensity of the structure models (M51 H α in red, toy model in blue) as a function of angular scale.

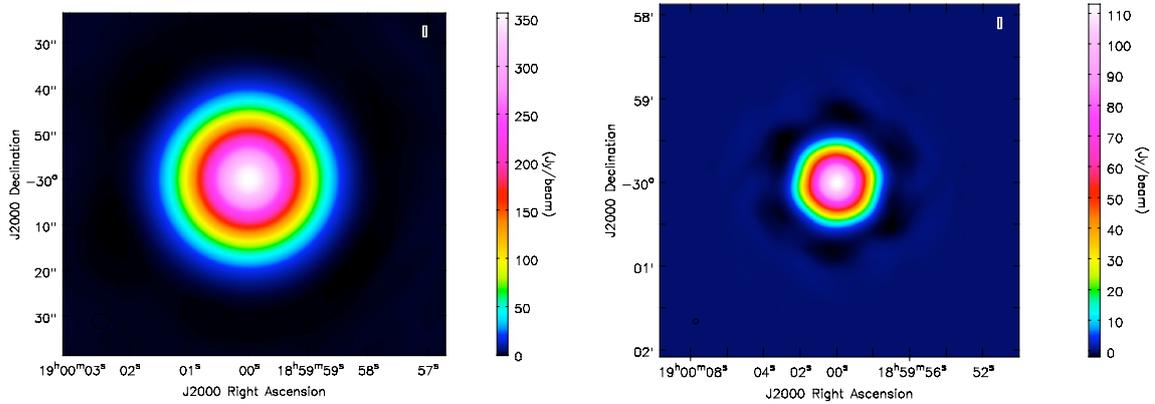


Figure 7.11: Gaussian functions as input models for the sky brightness distribution are observed with C32-1 at 100 GHz (DEC=-30deg, HA= \pm 1h). For a FWHM of 20 arcsec (left panel) in this noiseless simulation 96% of the integrated flux density can be recovered in this cleaned image, and for a FWHM of 40 arcsec (right panel) 22% are recovered.

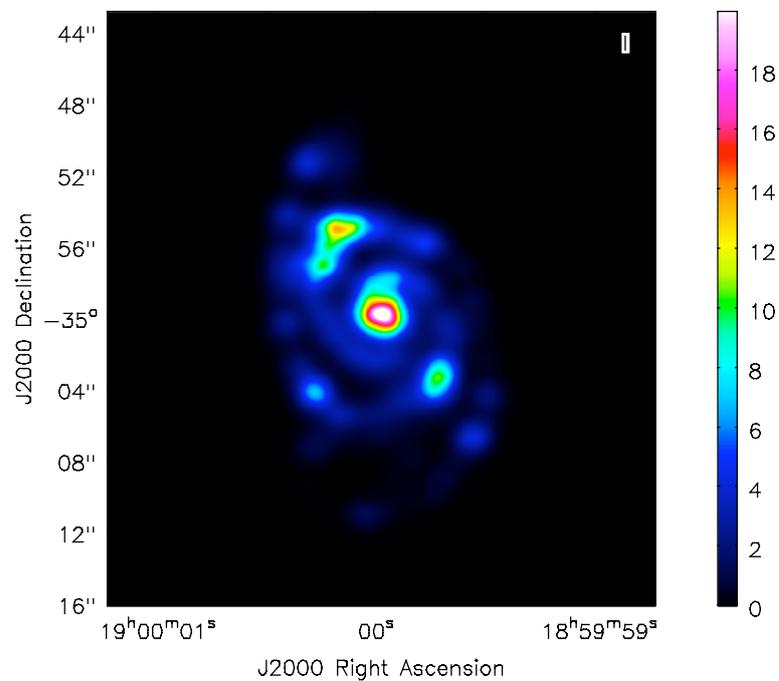


Figure 7.12: Fidelity image for the simulated observation of M51 with C32-1

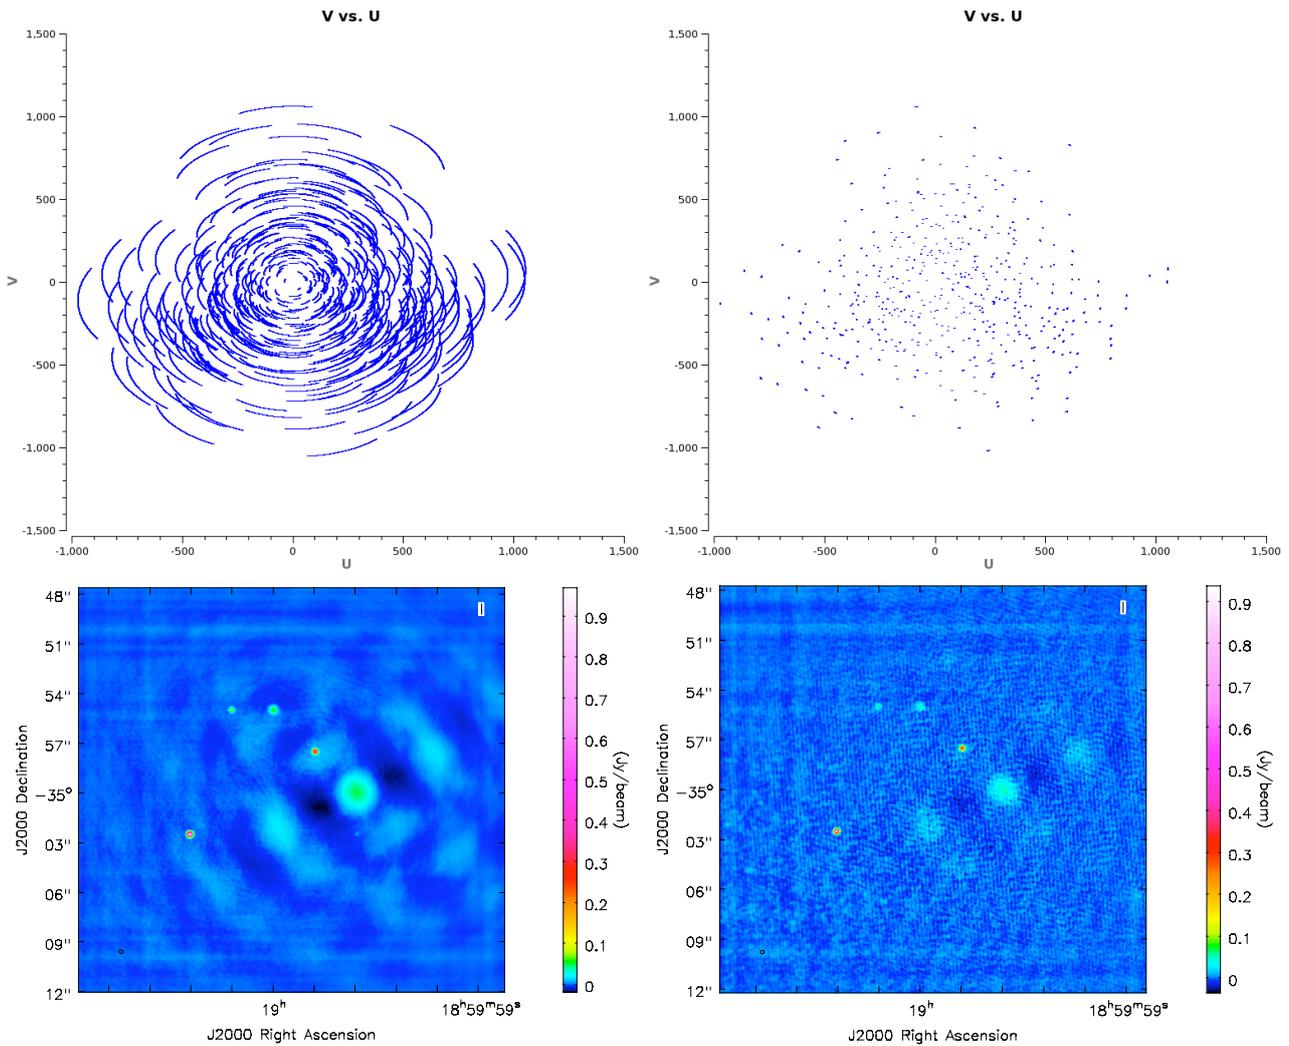


Figure 7.13: Crowded field with point sources and extended (Gaussian) structure in configuration C32-6. In the 3h observation on the left all four point sources, and the 0.3 and 2 arcsec FWHM Gaussians are visible. The 2 arcsec Gaussian is difficult to recover, particularly in the 5 min observation on the right, because there are few short baselines in that configuration. The weakest point source is not detected (and potential artificial sources are created) in the 5 min observation due to dynamic range limitations.

Chapter 8

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 to experiment with various sensitivity options, and which is also used by the OT to calculate its time estimates based on the parameters entered into a project's Science Goals. The same tool 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, the final time estimate for a project cannot be influenced to the same extent. For instance, this will always assume a fixed number of antennas for the particular Cycle, and will always use a PWV octile that is appropriate to the frequency of observation.

8.1 Calculating system temperature

The system temperature (T_{sys}) is built up from a number of elements that contribute noise as

$$T_{sys} = \frac{1 + g}{\eta_{eff} e^{-\tau_0 \sec z}} \left(T_{rx} + \eta_{eff} T_{sky} + (1 - \eta_{eff}) T_{amb} \right) \quad (8.1)$$

The various terms are

- T_{rx} – receiver temperature
- T_{sky} – sky temperature
- T_{amb} – ambient temperature (ground spillover)
- g – sideband gain ratio. For bands 1 and 2 (Single Sideband; SSB) and 3-8 (Sideband Separating; 2SB), $g = 0$. For these bands there is no contribution to the system temperature as the image sideband is either filtered out (SSB) or separated in the receiver (2SB). Bands 9 and 10 are Double Sideband (DSB) receivers and the correlated signal includes noise from both sidebands; therefore $g = 1$
- 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 i.e. zenith angle (z) equals zero degrees.

¹<http://www.almascience.org/>

Table 8.1: Receiver temperatures assumed in the ASC. For most of the bands we are currently assuming the ALMA specifications. In the case of bands 3, 6, 7 and 9, however, we are using “typical temperatures measured in the laboratory”; these can be significantly better than the specifications.

ALMA Band	T_{rx} (K)	Source
1	17	Specification
2	30	Specification
3	45	Laboratory
4	51	Specification
5	65	Specification
6	55	Laboratory
7	75	Laboratory
8	196	Specification
9	110	Laboratory
10	230	Specification

T_{sky} and T_{amb} are 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.

$$T_{\text{n}} = T \times \left(\frac{h\nu/kT}{e^{h\nu/kT} - 1} \right) \quad (8.2)$$

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

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.

The temperature of the Cosmic Microwave Background is not explicitly included in Equation 8.1 as it is included in T_{sky} (see Section 8.1.2).

8.1.1 Receiver temperatures

For most of the ALMA bands, the calculator currently only uses the specifications for the receiver temperatures and not the actual measured values. However, for the Early Science bands (3, 6, 7 and 9), typical values measured in the laboratory are used as these are usually significantly better than the specifications. The values used in the ASC are given in Table 8.1. 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 T_{rx} , i.e. only a single value is used per band. The measured values are somewhat conservative and so are in between what we might expect at the middle and edges of the bands. Similarly, we only use one value for the specification, the one that applies “over 80% of the RF 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).

Table 8.2: Octiles of PWV measured at the ALMA site from years of monitoring data and used in the ASC.

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

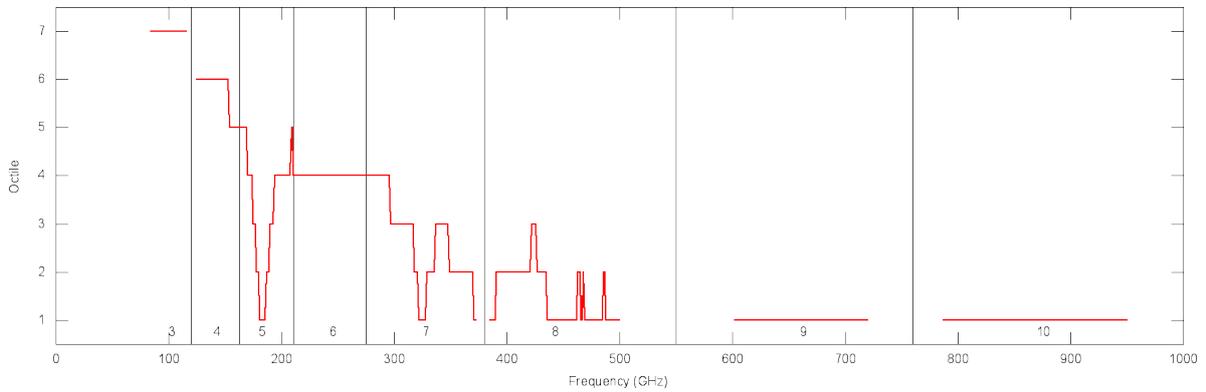


Figure 8.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown towards the bottom of the plot. The water line at 183 GHz (band 5) is particularly prominent. In general, higher frequencies require drier observing conditions.

8.1.2 Sky temperature

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

The ATM code assumes that the source is at the zenith and therefore the OT has to account for the greater atmospheric emission at lower elevations. The emission from the atmosphere is often approximated as

$$T_{\text{sky}} = T_{\text{atm}}(1 - e^{-\tau \sec z}) \quad (8.3)$$

where T_{atm} is the mean physical temperature of the atmosphere. This can be calculated using Equation 8.3 and from there the sky temperature as a function of zenith angle:

$$T_{\text{sky},z} = T_{\text{sky},z=0} \frac{(1 - e^{-\tau_0 \sec z})}{(1 - e^{-\tau_0})} \quad (8.4)$$

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

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.

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 8.2.

When estimating the time for a project, the OT will always select a PWV octile that is appropriate to the frequency being observed. It does this by calculating the time required for each octile and then choosing (and reporting) the highest (worst) octile for which the increase in time relative to the first is less than 50 per cent. A consequence of this definition is that the octile also depends on source declination i.e. sources at low elevations will require better weather conditions. The resulting curve of octile versus frequency is shown in Fig. 8.1, 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.

8.1.3 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). The value used by the ASC is corrected according to Equation 8.2.

8.2 The sensitivity calculation

Once T_{sys} has been determined it is possible to calculate the point-source sensitivity given a requested amount of observing time or vice versa. Note that the time calculated by the tool does not account for telescope overheads (calibration, etc.) and therefore the time is always assumed to be the true on-source time. There is also no accounting for the expected level of loss in sensitivity due to residual pointing and focus error.

8.2.1 12-m and 7-m Arrays

When dealing with the 12-m and 7-m Arrays, the standard equation used is

$$\sigma_{\text{S}} = \frac{2kT_{\text{sys}}}{\eta_{\text{q}}\eta_{\text{c}}A_{\text{eff}}\sqrt{N(N-1)}n_{\text{p}}\Delta\nu t_{\text{int}}}. \quad (8.5)$$

The various symbols are

- A_{eff} – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency (η_{ap}). The latter is given by the Ruze formula i.e. $\eta_{\text{ap}} = R_0 \exp(-16\pi^2\sigma^2/\lambda^2)$ where σ is the rms surface accuracy of the antenna – the specification of 25 μm and 20 μm for the 12-m and 7-m antennas respectively is currently used³. R_0 is equal to 0.72. See Table 8.3 for values of antenna efficiencies and effective areas in various ALMA bands
- η_{q} – quantization efficiency. A fundamental limit on the achievable sensitivity is set by the initial 3-bit digitization of the baseband signals. This is equal to 0.96
- η_{c} – correlator efficiency. Although this depends on the correlator (64-input or ACA) and correlator mode, this is currently fixed at 0.88. For the 12-m Array at Cycle 1, this is fine as all the available modes during Cycle 1 will have this efficiency. For the ACA, the efficiency *will* depend on the mode, but as the ASC is currently not mode-aware, the same factor of 0.88 has been assumed for the time being⁴
- N – number of antennas. This defaults to 32 for the 12-m and nine 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

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.

⁴The Observing Tool *is* mode-aware, although the ACA time required at Cycle 1 will be based on the 12-m Array time estimate.

Table 8.3: Aperture efficiencies at typical continuum frequencies observed in bands 3, 6, 7 and 9, for both the 12 and 7-m antennas. The effective area, A_{eff} , is equal to the efficiency multiplied by the physical area of the dish i.e. 113.1 m² and 38.5 m² for the 12 and 7-m antennas respectively.

Frequency (GHz)	$\eta_{\text{ap},12 \text{ m}}$ (%)	$\eta_{\text{ap},7 \text{ m}}$ (%)
100	71	71
230	68	69
345	63	66
690	43	52

- $\Delta\nu$ – 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. $n_p \Delta\nu$ is often referred to as the effective bandwidth
- t_{int} – integration time.

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

$$\sigma_{\text{T}} = \frac{\sigma_{\text{S}} \lambda^2}{2k \Omega} \quad (8.6)$$

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

$$\Omega = \frac{\pi \theta^2}{4 \ln 2}. \quad (8.7)$$

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

8.2.2 Total Power Array

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

$$\sigma_{\text{TP}} = \frac{2k T_{\text{sys}}}{\eta_{\text{q}} \eta_{\text{c}} A_{\text{eff}} \sqrt{N} n_{\text{p}} \Delta\nu t_{\text{int}}}. \quad (8.8)$$

This is the same as Equation 8.5, apart from there only being a factor of \sqrt{N} in the denominator, where N is the total number of total power antennas (two for Cycle 1).

8.3 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 Fig. 8.2.

- Source declination – this is used to calculate the maximum elevation of the observation and thus the minimum airmass i.e. the ASC assumes that the source is transiting.
- Observing frequency – this sets the receiver temperature, antenna efficiency and the atmospheric opacity.

Sensitivity Calculator

Common Parameters

Dec: -17:35:02.400

Polarization: Dual

Observing Frequency: 345.0 GHz

Bandwidth per Polarization: 7.5 GHz

Water Vapour: Automatic Choice Manual Choice

Column Density: 0.913mm (3rd Octile)

tau/Tsky: tau0=0.158, Tsky=36.784

Tsys: 151.779 K

Individual Parameters

	12m Array	7m Array	Total Power Array
Number of Antennas	32	9	2
Resolution	1.0 arcsec	5.974554 arcsec	17.923662 arcsec
Sensitivity(rms)	73.44625 uJy	0.76435 mJy	1.63573 mJy
(equivalent to)	0.75448 mK	0.21997 mK	0.05230 mK
Integration Time	10 min	10 min	10 min

Integration Time Unit Option: Automatic

Buttons: Calculate Integration Time, Calculate Sensitivity, Close

Figure 8.2: 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.

- Bandwidth per polarization – this otherwise straightforward parameter should be set to 7.5 GHz for continuum observations (see Section 8.2.1). 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 32 from the 12-m Array, nine from the 7-m Array and two from the TP Array.
- Angular resolution – this affects the time estimates when sensitivities are specified in temperature units. The calculator will not perform any calculations when Kelvins have been specified, unless a non-zero value for angular resolution has been entered. The calculator will also issue a warning if the angular resolution falls outside of the range corresponding to 125-m and 1-km baselines.

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

Chapter 9

Observing Projects and Logical Data Structure

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, which are needed for successful completion of a project.

A summary view of the constituents of an Observing Project is shown in Figure 9.1.

9.1 Observing Project Structure

From the point of view of the proposer, the fundamental unit of the Observing Project is the “Scheduling Block” (hereafter SB). An SB is the smallest sequence of observing instructions that can be scheduled. Scheduling Blocks are produced by the OT (including in Phase 1¹) and may be edited using the OT in Phase 2 after proposal acceptance and before/during proposal execution. Projects are broken down into a set of these fundamental units to take maximum advantage of the properties of the ALMA site and the continuously-varying status of the observatory as a whole, including the weather. Each SB has an execution time of typically 1hour. Within this time limit, the SB carries out the set-ups, calibrations and target observations to ensure that the acquired data can be properly calibrated and used in the production of the final data product. 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. An SB is atomic in the sense that it cannot be re-started in the middle. Therefore, an SB runs to completion, fails, or is terminated by the telescope operators.

Structurally, each SB contains two parts, a static and a dynamic part. The dynamic part is a script of observing commands (for standard observing modes it is actually a reference to an external library script) that direct the actual process of observations. The static part consists of fixed auxiliary data and meta-data that parameterize the dynamic part. This metadata is generated by the OT during the Phase 2 as part of the SBs. Observations are carried out by interpreting and executing the script supplemented by the relevant parameters contained in the SB.

Given the limited length of the SBs, it often will be necessary to observe several of them to obtain a final image. All the SBs required to reach a specific scientific goal within a project are thus gathered together under a larger structure called an Observing Unit Set (hereafter OUS). An OUS is composed either of a set of SBs or a set of OUS (needed for instance, in the case of the same set of observations for different array configurations). It also contains all the preconditions, performance and calibration requirements, and flow control that apply to that collection. The completion of the parent OUS in each of these triggers reduction of the relevant data by the automated Science Pipeline in Full Science Operations or semi-automated reduction by ALMA staff during

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

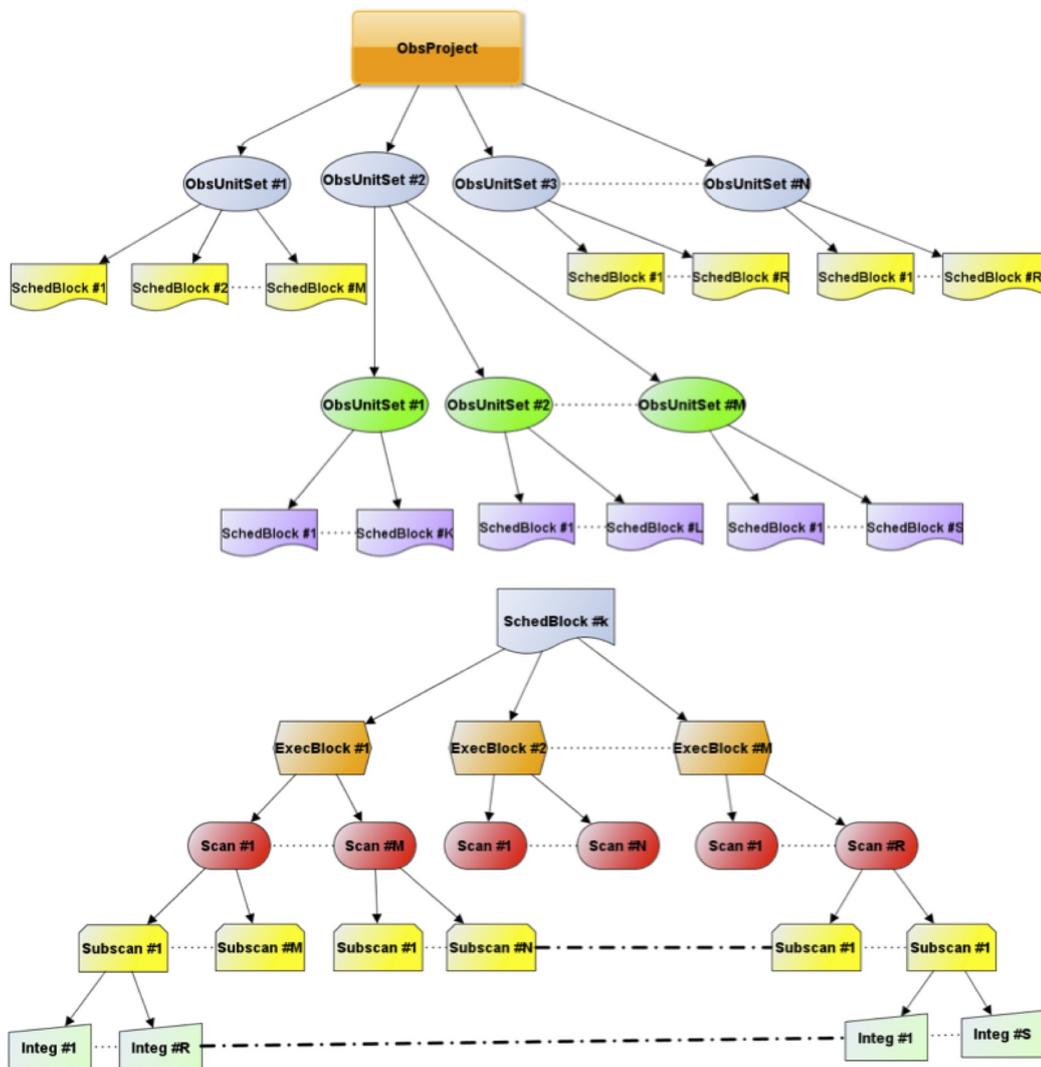


Figure 9.1: Block diagram of an Observing Project from the point of view of the observation preparation (top) and actual execution (bottom). Note that an **ObsUnitSet** (OUS) can be composed of other **ObsUnitSets** as explained in the text. Each time an **Scheduling Block** is executed, **CONTROL** creates a new **ExecBlock** structure (see text).

Early Science. The only other events that trigger data reduction are at the time of the completion of all the OUS that compose a given Observing Project, or at the end of an observing season, when all the projects are reduced, irrespective of their degree of completion.

9.2 Program Execution

Once a given SB has been selected for execution (by the Scheduler subsystem or by the Array Operator/AoD), it gets loaded into the system and the Control subsystem takes over. The Control subsystem creates an Execution Block structure that is attached to a given SB. Since there is the possibility of the same SB having to be executed several times, many Execution Blocks may exist for a given SB. Each Execution Block 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 observations is also shown in Figure 9.1. The Control subsystem constructs and executes a sequential series of scans. Each scan execution is in fact carried out by breaking it down into a series of subscans, each of which is itself broken into a series of integrations. Although commands are issued at the scan/subscan level, the correlator output corresponds to or is attached to a particular integration. Calibration results from the Telescope Calibration subsystem (TelCal) and QuickLook (QL) pipeline are usually attached to a subscan. However, some of these results may be accumulated over many subscans, and the resulting object will be attached to the scan to which the subscans belong. The Control subsystem is responsible for the creation of all the metadata needed downstream for data processing, for the Execution Block, the scans, subscans and integration objects.

9.3 Structure of an SB and associated scripts

9.3.1 Observing Groups

Within Phase II SBs for a given project, all observing Targets that are to be executed must be included in a “group” within the SB. Each SB can include multiple groups. As of the time of writing, the first group (group 1) of an SB is always the initial calibration group, with subsequent groups detailing the science observation(s). All Science targets and relevant calibrators² within a group are observed before the next group is started. Calibrators will not be observed during the execution of a group if they are not visible at the time of the observations. It is assumed that all the Science targets in a given group are sufficiently close as to share the same Phase and Pointing calibrators. A group, other than group 1, is considered complete when all Science targets in the group have been observed the requested time or have set below the elevation limit.

9.3.2 The Standard Interferometry Script

The Cycle 1 observing modes, single field interferometry, grouped source executions and pointed mosaics, use the same ALMA observing script, called the “standard interferometry” script. It is designed to provide well-calibrated observations of discrete sources. 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. This means that science targets and calibrators in a group may be observed repeatedly within an SB. A set of necessary calibration measurements (e.g. amplitude, bandpass, etc.) specified in group 1 are performed at the beginning of the observation sequence if appropriate sources are available, otherwise this information is collected at the end. Pointing is verified before the amplitude and bandpass calibrators are observed (unless they are close on the sky), and again before the main observations of the science target and phase calibrator cycle. Any additional (“secondary”) phase calibrators are observed as specified in the scheduling block. Atmospheric calibration to determine system temperature is performed during each iteration of primary phase calibration

²The user has several options to select optimum calibrators. He/she can let the OT set up adequate queries to the Calibrator Database in the ALMA Archive. He/she can also specify the calibrators or set-up the queries, but this has to be fully justified in the Technical Justification of the proposal.

(at high frequencies) and once every two primary phase calibrations at low frequencies. For a description of what each calibration entails, see Section 10.

Pointing and Atmospheric calibration is performed only in association with the relevant Science or Calibration target. This means that prior to an execution of a Amplitude, Bandpass or Phase Calibrator or Science target, associated Pointing and Atmospheric calibration targets are checked to determine if their cycle time has passed, allowing these calibrations to be done near in time to the relevant observation. Atmospheric calibration is associated by field name. Pointing calibration is associated by separation on the sky so that to each Amplitude/Bandpass/Phase Calibrator or Science target has associated with it the closest Pointing calibrator available. Ultimately, if the separation is more than 15 degrees on the sky, the association is discarded. Generally the Science target and Phase calibrator share an associated Pointing target so pointing calibration is done only for the pair.

The script is still being developed and may evolve somewhat, but the current version typically uses the following group structure and proceeds with the strategy indicated below:

Sources:

Initial calibration group:

- 1-2 Pointing calibrators (near and in association with the Amplitude and/or Bandpass calibrators)
- 1-2 Atmospheric calibrators (near and in association with the Amplitude and/or Bandpass calibrators)
- 1 Amplitude calibrator (Typically a Solar System object or a well-known quasar)
- 1 Bandpass calibrator

For each subsequent group:

- 1 (or more) Science targets(s)
- 1 (or more) Phase calibrator(s)
- 1 (or more) Pointing source(s) (associated with the Science and/or Phase calibrator)
- 1 (or more) Atmospheric calibrator(s) (associated with the Science target(s) and perhaps also with the Phase calibrator(s))

Strategy:

Execute Initial Calibration Group (as necessary):

1. Observe Focus calibrator(s) (optional and non-standard)
2. Observe Bandpass calibrator(s) (with associated Pointing, Atmospheric and implicit Sideband Ratio calibrators)
3. Observe Amplitude calibrator(s) (with associated Pointing and Atmospheric calibrators)

For all other groups:

1. Assume the primary Phase calibrator is the one in the group with the shortest cycle time, and use this to set the scan duration
2. Focus, bandpass and amplitude calibrations (optional and non-standard, with associated Atmospheric and Pointing calibrators)
3. Perform P Phase calibration observations; $P \geq 0$ (with associated Atmospheric and Pointing calibrators, as determined by the cycle time)

4. Observe the Science target list until complete, a Phase calibration is required, or SB execution limit is reached. Reaching the SB execution limit aborts the current group execution and bypasses all remaining groups (a single Science target is a list of one element)
5. Perform a primary Phase calibrator observation
6. Repeat above steps until group is complete or an SB execution time limit is reached

After all groups are complete or the SB execution limit is reached, the deferred Calibrators from group 1 and final Phase calibrator for the group which triggered the SB execution time limit are executed with associated Calibrators.

Each Science target in the Science target list is assigned a fraction of a sourceCycle time, a fraction determined by the time requested for that particular Science target divided by the total integration time requested for all Science targets. The sourceCycle time defaults to the uv-cell crossing time for the configuration in use and is effectively infinite for single element Science target lists. The scan time (see Section 3.4) for a given Science target is then determined by the minimum of the sourceCycle time fraction and the Phase calibrator cycle time with an absolute minimum of 30 seconds allowed for efficiency considerations. For a scan executed on a Science target which includes a pointed mosaic, each position in the mosaic is observed for one subscan duration (see Section 3.4) until the time for the scan expires. Both the Science target specific last mosaic pointing position and index within the Science target list are stored to ensure the next scan begins on the proper Science target and proper offset position.

9.4 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 9.2.

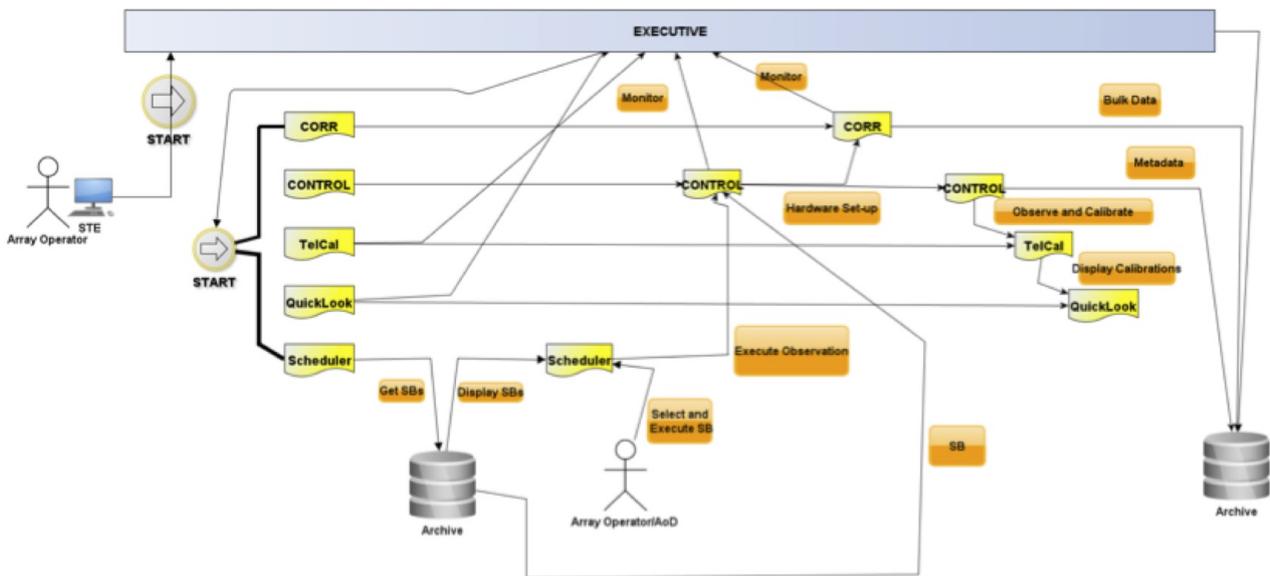


Figure 9.2: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

Each of the yellow boxes in Figure 9.2 represent an ALMA subsystem involved in the observations. Orange boxes include labels for the actions performed either by external agents 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 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 that it listens to, it is worth mentioning the publication 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. An SB is then selected from the list provided by the Scheduler subsystem, and the execution is started. All the SBs for a given observing season 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 a given observing cycle. The Telescope Operator/Astronomer on Duty have two options for using the Scheduler during Early Science. They can either let the Scheduler suggest possible SBs to execute, or they can carry out targeted searches of the local database. For Early Science, the Scheduler can produce a ranked list of optimum SBs to execute next based on weather conditions and forecast, hardware/configuration status, project completion status, representative source position on sky, proposal rank and score and Executive percentages. The Telescope Operator/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 it can run an automatic sequence of SBs without disagreeing with what an AoD would select.

Once the execution of the SB has been selected from the Scheduler, it is dispatched to the Control subsystem. Control executes the SB by commanding all relevant hardware and the correlator, resulting in raw data and metadata being made available to the Telescope Calibration (TelCal) and Quick-Look subsystems that report the results/progress/quality of the calibrations carried out during the observations. 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 correlator to within 48ms (TE). Control is also in charge of storage of data from all monitor points set in the hardware of the ALMA array.

Storage of data into the Archive follows two parallel paths and is handled by a part of the Control software called the “Data Capture” module. The best way to describe the Data Capture module is as an interface between the real-time domain of the data taking and the storage side. As an interface, it captures and stores into the Archive all the relevant meta-data information pertaining to a complete description of the data and their supporting calibration and monitor datasets and condenses that information into a set of XML tables. The contents of all these tables are defined in the ALMA Science Data Model (ASDM), which includes a set of 16 core tables (present for all datasets) and up to 23 additional tables (present as needed for a specific observing mode). Together with these tables, Data Capture also creates the relevant links of these metadata to the actual bulk data that is directly stored into the archive. Furthermore, it also provides calibration data to the TelCal and QuickLook subsystems for calibration reduction and display in semi-real time. Finally, when the SB is finished, the 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, since it has to handle Correlator/backend data, supporting (source information, spectral set-up, etc), and monitoring data needed for the reduction (weather, pointing, etc). All this information originates in different hardware/software elements, each of which can be sampling at different rates and with limited view of the behavior/state of the observing system.

A summary plot of the main elements involved in data flow is shown in Figure 9.3.

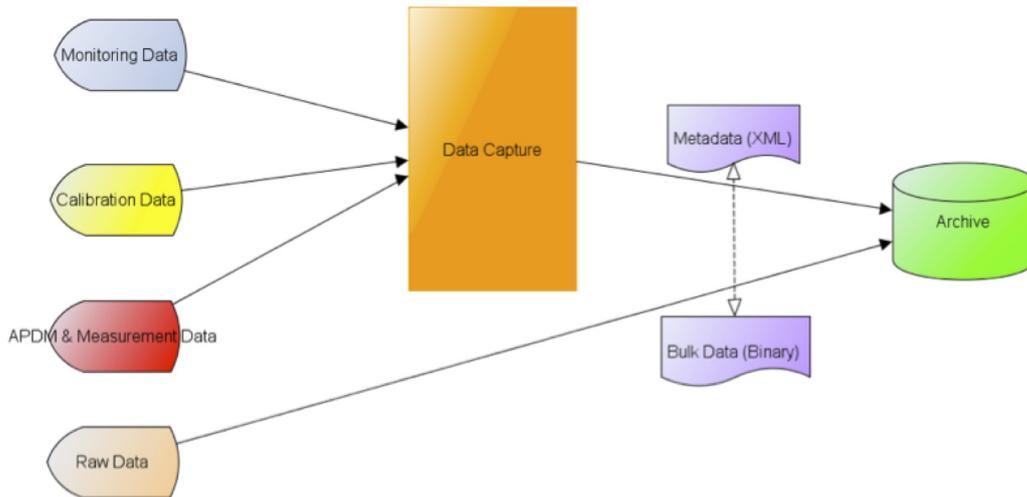


Figure 9.3: Data flow components.

9.5 The ALMA Science Data Model (ASDM)

The ASDM defines the collection of information recorded during an observation that is needed for scientific analysis. As described above it contains both bulk binary data and metadata (XML) organized in tables. The tables contain links to other XML tables and addresses pointing to the actual bulk data in the Archive. The ASDM contains 16 core tables that are common to all observing modes, and up to 23 additional tables that are only created for specific observations. On top of these, the online calibration system (TelCal) also creates associated tables whenever it processes any on the calibrations that can be done on-line. 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, the targets, auxiliary monitoring data, overall project and post-processing. A list of the core tables is shown in Figure 9.4.

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/calibrations become available (see Section 10). The current list is shown in Figure 9.5.

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

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

Figure 9.5: Current list of ASDM associated tables generated by TelCal.

Chapter 10

Calibration and Calibration Strategies

The calibrations needed by a given observation can be broadly divided into those that correct short-term effects (less than the duration of a typical SB, about 90 minutes, or that require being measured at least once per receiver tuning) and those that correct for longer term variations. Calibrations included in these two categories are:

- Long-Term: All-sky Pointing (including Band relative offsets), Baseline Vectors, Cable Delay, Focus Models, Antenna Characteristics
- Short-Term: Offset Pointing, Bandpass, Phase fluctuations (WVR), Gain (Amplitude & Phase), Flux, Receiver Temperature, System Temperature, Sideband Ratio and Absolute Flux Calibration

Long-term effects do not need to be calibrated at the time of a given science observation. ALMA staff will carry out periodic measurements of these long-term effects and apply the required corrections to the ALMA system so that they are shared by all observational projects. Any small residual errors that remain can be corrected offline. Short-term effects will have to be measured during the science observations and the time taken by these calibrations is added to the total time required to reach the SNR and imaging goals on the science targets. A brief description of the objectives and strategies of each of the calibration follows.

10.1 Long-Term Effects

10.1.1 All-Sky Pointing

These observations are used to correct the overall mechanical and orientation imperfections of each antenna & pad assembly. These imperfections include tilts of the AZ axis, zero-points of the encoders, encoder run-outs and tilts in AZ and EL, non-perpendicularity of the AZ and EL axes, sagging of the sub-reflector structure, differential sagging of the telescope structure in AZ between two successive antenna supports, plate offsets among ALMA Bands, etc. The preferred method of measurement is a series of interferometric observations of 50-100 pointing sources (mostly unresolved continuum sources) covering a wide range in AZ and EL. A 14-18 parameter pointing model (with well-known angular dependencies) is fitted to the data and used to derive the required corrections to the nominal pointing positions to correct for all of these structural effects. It is customary to define one of the ALMA bands as pointing reference (currently Band 3 for all antennas) and define the other Band pointing models as offsets from the reference. This all-sky pointing model has a specification RMS error of $2''$ (for any point on the visible sky and after application of the model) for all ALMA antennas.

10.1.2 Focus Models

Any homologous antenna design causes the optimum focus position to vary as a function of elevation. For a perfect antenna, the only offset that would need to be corrected for would be deflections relative to the axis of the parabolic dish. However, because ALMA deals with off-axis receivers and due to the presence of small non-uniform sagging of the sub-reflector mount/support, corrections in the plane perpendicular to the optical axis are also needed. Models (the same for all antenna types) for the offsets of the different ALMA Bands are derived by multiple interferometric/single-dish observations of bright unresolved sources. With these measurements the positions of the maximum power transmission to the receivers as a function of position in three orthogonal directions ($X=AZ$, $Y=EL$, Z =optical axis) over a wide range of telescope elevations are derived. The subreflectors are also tilted to maximize the power transfer and avoid reflections of signals originating in the receiver cabin. The reflections are avoided by molding the sub-reflectors with an additional inner conic section. As in the case of the All-Sky pointing model, one of the ALMA receiver bands is defined as reference (Band 7). The ALMA antenna subreflectors can be positioned with a very high accuracy (i.e., $5 \mu\text{m}$). The corrections of the focus models are applied during observations in a way that allows for correction of pathlength changes.

10.1.3 Baseline Vectors

To accurately compute appropriate signal delays, uv sampling, etc, the relative telescope position (i.e., the baseline vector) for all antenna pairs must be known to within a fraction of a wavelength. Baseline vectors will be measured regularly by the ALMA staff (whenever antennas are moved and periodically). Baseline determination in ALMA follows a two step approach. An initial estimate of the pad positions is obtained from GPS measurements and/or cartographic maps. Interferometric observations of unresolved continuum sources covering a wide range in hour angle and declination are then used to refine the baseline solutions. It is expected that ALMA will be able to determine baseline vector lengths to an accuracy of $65 \mu\text{m}$ for the main array and $33 \mu\text{m}$ for the ACA. The more stringent requirement for the ACA comes from its lower sensitivity, which requires, on average, to look for calibrators twice as far away on the sky than for the main array. Baselines for the Total Power antennas will also be measured, because it is expected to support the calibration observations of the ACA 7m antennas in interferometric mode. The baseline vector files pertinent to a given dataset will be delivered to the PIs together with the data.

10.1.4 Cable Delay

Errors in the signal delay from the receiver to the correlator, associated with each antenna/spw/pol, produce a baseband phase slope as a function of frequency. The majority of this phase slope must be removed before observations can proceed. The approximate value and temporal change of this delay is monitored by appropriate ALMA test signals, but a more accurate delay is obtained from short observations of bright sources that are done periodically by the ALMA staff. Any further small residual delay (less than 180 degree phase shifts across a given spectral window) can be removed as part of the bandpass calibration (see below). Because of the role that the Total Power antennas play in calibration of the ACA 7m interferometer, the delays to these antennas will also be accurately measured.

10.1.5 Antenna Characteristics

Surface Measurements/Adjustments

The aperture efficiency of an antenna is a very steep function of the ratio of the surface errors (compared to an ideal aperture) to the wavelength of the observations. The practical shortest wavelength an antenna can operate at is given by the surface errors as of about $1/13$ of the wavelength. Therefore, the surface accuracy of all ALMA 12m antennas should be less than $25 \mu\text{m}$ (under all primary operational conditions) and less than $20 \mu\text{m}$ for the ACA 7m antennas (nighttime conditions) for optimal scientific results, especially at the higher frequencies. The difference between an ideal surface shape and the one of a given antenna is measured with

a strong CW (Continuous Wave) signal from a near-field beacon and interferometrically using bright celestial sources. The interpretation of these data results in adjustments of the antenna surfaces to obtain the nominal aperture efficiency as a function of source elevation, and the angular sensitivity distribution of each antenna (beam pattern, see below).

Beam Patterns

Accurate knowledge of the primary beam patterns, for all ALMA bands, polarizations and ALMA antenna types/footnoteThere are three different designs of the ALMA 12m antennas, which results in different beam sidelobe properties, is required for high-quality imaging and for assessing telescope performance. Interferometric or SD (Single Dish) observations of bright unresolved sources, obtained by scanning the antenna to be measured in square patterns across the source, are usually used. These allow for the derivation of high SNR maps of the antenna total power beams. Currently the model beam patterns are used in the offline reduction instead of the measured results, but CASA has the capability to import them if necessary. For observation planning, the current best estimate of the ALMA primary beam profile is a Gaussian with a $\text{FWHM} = 1.17\lambda/D$, where $D=12\text{m}$ (or 7m for the ACA). The coefficient of 1.17 is based on the optics specification of a -12dB Gaussian illumination taper for the receiver feeds (see ALMA Memo 456). Detailed measurements of the individual beam profiles at various frequencies in each band for each antenna will not be available for some time.

10.2 Short-Term Effects

10.2.1 Offset Pointing

To be able to center sources within the primary beam field of view at the highest frequencies offered by ALMA, a pointing accuracy of about $0.6''$ is needed This is clearly not possible with the All-Sky Pointing model, and must be measured separately. This is achieved with interferometric observations of nearby pointing calibrators (within 4 degrees of intended target). Since some of the calibrators needed within an SB can be widely separated from the science target, i.e. bandpass and amplitude, it is customary to carry out offset pointing observations for those calibrators too. Once the science target observations are started, usually one pointing calibration is enough for the typical duration of an SB. These offset pointing measurements will only update the two collimation offsets (AZ, EL) of the antenna pointing model. Due to the reduced number of suitable pointing sources at the highest frequencies (particularly ALMA Bands 7 and 9), pointing calibrators will usually be observed in a lower frequency band (Band 3/6) and these observations along with relative focal plane offsets between receivers will be used to update the offset pointing model.

10.2.2 Bandpass

The spectral response (amplitude and phase) of the combined atmosphere and receiving system is generally not flat for a real interferometric array. For single dish observations, a similar situation occurs for the spectral response within each spectral window. ALMA data coming out of the correlator may appear flatter than in other instruments because the cross-correlations are normalized by the auto-correlations at the correlator. In fact, the ALMA specification for spectral window "flatness" is a variation of response of less than 3dB across the band in autocorrelation mode. It is important to measure the spectral response accurately to the level required by a given observation, as this will be the limiting factor for the accuracy of measurements requiring high spectral dynamical range (especially for high spectral resolution modes). At least one bandpass observation should be included per tuning (i.e., SB execution, because each re-execution entails a tuning optimization) for interferometric observations that require high spectral accuracy. This calibration is done by observing a bright continuum source with a well-known spectral energy distribution (flat-spectrum sources are the easiest to use). A single bandpass is calculated per antenna, spectral window and polarization. The observation should be long enough to obtain a signal-to-noise on the bandpass calibrator that is at least as high as the desired spectral dynamic range of the target observation at the desired spectral resolution of the observations. In the case of

single dish observations, the bandpass characteristics are corrected using the spectral curves derived from the T_{sys}/T_{rx} calibrations (see below). Several are usually executed during the duration of an SB. A simple formula for the time required for a bandpass observation can be derived assuming that the required spectral SNR is that of the target source:

$$\tau_{BP} = \tau_T \frac{\Delta\nu_T}{\Delta\nu_{BP}} \left(\frac{S}{B}\right)^2 \frac{T_{sys}^{BP}}{T_{sys}^T} \quad (10.1)$$

Where τ is the required on-source time, T, BP are the target and bandpass observations, $\Delta\nu$ is the spectral resolution (can be different for the target and bandpass observations to improve the SNR of the bandpass observations), S and B the fluxes of the target and bandpass calibrator and T_{sys} the system temperatures.

10.2.3 WVR Corrections

Fluctuations in the line-of-sight Precipitable Water Vapor (PWV) of two antennas in a given baseline can cause significant decorrelation, especially at the high frequency bands of ALMA. Fluctuations in PWV levels are driven by wind, and can have rapid and strong variations over a large range of spatial scales. All ALMA 12m antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures at a rate of 1 Hz the emissivity of the atmospheric water line within four spectral bands near 183 GHz. From these measurements, atmospheric models are used to derive the amount of PWV in the line-of-sight and from the difference between the values for the two antennas in a baseline, the phase fluctuations at the observing frequency. The PWV measurements are also used to derive sky opacities at the frequencies of the science observations when combined with the ATM atmospheric model. It is expected that the phase fluctuation corrections will be applied automatically by the correlator by the start of Cycle 1. Since the ACA is very compact, the effect of the atmospheric PWV fluctuations is quite smaller than that for the 12m Array. For Cycle 1, the PWV fluctuations will not be measured for these antennas, and therefore, small phase fluctuations will remain in the data. For future cycles it is expected to correct the fluctuations using a method involving measurements with the TP array antennas surrounding the ACA and extrapolation to the positions of the 7m dishes. A more detailed description of the WVR hardware, specifications and operation can be found in Section A.6.

10.2.4 Gain (Amplitude & Phase)

The amplitude and phase transfer properties of the atmosphere and receiving system vary more slowly than the PWV fluctuations but vary on timescales of about 5-10 minutes and are still rapid compared to a typical SB. They are corrected using nearby calibrators (within 10° of target). Usually, observations of the science targets are done in cycles straddled by observations of the gain calibrators. Since the WVR on the ALMA 12m antennas help to correct the phase fluctuations at short timescales (i.e., 1 sec to 2 minutes), the gain calibrations can be done more sparsely (every one to a few minutes). Faster cycle times between the gain calibrators and the target sources will only be needed for special observations at long baselines and under very adverse weather conditions. Ideally the calibrators should be unresolved for simplicity of reduction, but CASA can also cope with resolved calibrators if necessary (using models). Also, band to band phase transfer techniques are possible for example using gain observations at Bands 3 or 6 to calibrated science target data at Bands 7 or 9. This type of calibration requires observation of at least one calibrator at both the high and low bands to calibrate the band to band phase difference. Note that band to band phase transfer techniques are only needed when there are no sufficiently bright/nearby calibrators available for the high frequency bands. For Cycle 1, the fact that full polarization observations will not be offered implies that there might be some small calibration errors in the observations in those cases where the calibrators used are polarized. The observatory will strive to minimize these errors by using calibrators with small polarization as much as possible. An estimate of the observing time required to reach a given phase RMS is given by:

$$\tau(sec) = 300 \left(\frac{T_{sys}}{100 K}\right)^2 \frac{31}{N_{ant} - 1} \left(\frac{12 m}{D_{ant}}\right)^4 \frac{2 GHz}{\Delta\nu} \left(\frac{0.1 rad}{\sigma_\phi}\right)^2 \left(\frac{1 mJy}{S_\nu}\right)^2 \exp\left(\frac{0.1\epsilon}{25 \mu m} \frac{\nu}{100 GHz}\right) \quad (10.2)$$

where τ is the required time in sec, T_{sys} the system temperature, N_{ant} the number of antennas, D_{ant} the diameter of one of the antennas in the array (12m for the 12m array and 7m for the 7m array), $\Delta\nu$ the total bandwidth used in determining the phase (it is usually the whole bandwidth of a given spectral window), S_ν the flux of the calibrator at the frequency of the observations, ϵ the surface RMS of an antenna (25 μm for the 12m array and 20 μm for the 7m array), ν the frequency of the observations and σ_ϕ the desired phase RMS.

10.2.5 System and receiver temperature

At millimeter and submillimeter wavelengths, the atmosphere both attenuates astronomical signals and acts as a black body emitter that adds additional noise to any measurements. This effect is a strong function of frequency, elevation, the column of wet and dry constituents of the atmosphere, and the temperature of the atmosphere. To measure this effect, ALMA front-ends are equipped with an Amplitude Calibration Device (ACD) which consists of a robotic arm with two "loads" of known emissivity, one at the antenna cabin temperature and the other heated to 370 K. Consecutive observations of these two loads and an observation of the sky are used by ALMA to measure the T_{sys} (see description of the measurements in Section A.5.1). Because of its dependence on the possibly changing ambient temperature and elevation, T_{sys} needs to be tracked every 5-15 minutes depending on frequency. The specifications for T_{sys} calibrations are to reach 1% repeatability¹ for ALMA bands up to 7 and 3% repeatability at higher frequencies.

10.2.6 Sideband Ratio

ALMA receivers for Cycle 1 include types 2SB and DSB. Both types can have significant leakage of signals from one of the sidebands to the other, affecting the flux calibrations of the results (for 2SB receivers the specification is 10dB). Some techniques (i.e., LO offsetting) will be used during the interferometric observations to significantly reduce the leakage, but the residual has to be measured. No technique is available for the SD observations in Cycle 1. Since the ALMA receivers have very stable tuning properties, the Sideband Ratio (SBR) will usually only have to be measured once at the start of an SB using sources with well-known spectral properties. The results are automatically applied to all the T_{sys} calibrations during the observations.

10.2.7 Flux

Calibration of the data using the T_{sys} method above does not yield an absolute calibration of the actual absolute flux density of a given source. It just allows the data to have a high repeatability irrespective of weather conditions, elevation of the observations, etc. In order to derive an accurate absolute flux calibration scale, a measurement of a source with known flux density and structure is required for each spectral setup, each time an SB is executed. For observations ALMA Solar System bodies in conjunction with models for their size and brightness temperature are most often used for absolute flux calibration, though frequently monitored quasars and the photosphere of stars are also options. A measurement of a flux calibrator is needed once per tuning (or per SB if high accuracy is desired). The absolute calibration accuracy for Cycle 1 is 5% for Band 3, 10% for Bands 6 and 7, and 15% for Band 9.

¹Repeatability is defined as the ability to observe the same non-variable target under different observing conditions (day, time, weather, receivers, etc) and obtain the same intensities. The percentages given in the text represent the expected departures from a perfect match.

Chapter 11

Quality Assurance

The goal of ALMA Quality Assurance (QA) is to deliver to the PI a reliable final data product that has reached the desired control parameters outlined in the science goals, that is calibrated to the desired accuracy and free of calibration or imaging artifacts. For Cycle 1 there are several restrictions that make it a bit more difficult to achieve this goal. It has therefore been decided that QA will also be done on a “best effort” basis, covering all the major issues of the data. The QA process analysis will be based on a calibration plan that specifies which observations must be acquired and at which intervals in order to monitor system performance and environmental time evolution. Furthermore, it will also tackle issues related to the merging of data for each science goal taken with different configurations, the inclusion of single-dish data, and the ultimate image quality. Errors introduced by user-supplied parameters, such as incorrect source coordinates, inadequate frequency setting (e.g. an incorrect redshift), inadequate sensitivity limits (leading to an inadequate integration time or inadequate uv plane coverage) are outside the scope 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: At the time of data acquisition

QA1: Observatory-Task Quality Assurance

QA2: Data Reduction

QA3: Post Data Reduction

The QA0, 1 and 2 stages will be handled by the Program Management Group (PMG) and the Data Management Group (DMG) (with some contribution from ARC personnel) using the AQUA Tool (see Section 11.5). The QA3 stage will be handled by the ARCs via JIRA¹ tickets created by the ARC personnel. The final output of the ALMA QA0-QA2 process is a “QA Report” per project (or ObsUnitSet) that summarizes all the relevant QA information for each of the different QA stages up to the Data Reduction. This report will be included in the data package delivered to the PI. QA3 will be handled separately, as discussed below. In Chile, responsibility for data quality assurance rests with the Data Manager within the Department of Science Operations, drawing upon the resources of the Program Management Group and the Data Management Group. A more detailed description of the different stages of QA is as follows:

11.1 QA0

QA0 is a near-real-time verification of data quality. It deals with rapidly-varying performance parameters (at scales of an SB execution length or shorter) and thus has to be performed at the time of data taking. Assessment

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

is performed by AoDs (Astronomers on Duty) at the OSF, based on semi-real time output of the calibrations (TelCal) 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 will be complemented with Monitor and Control display tools to monitor specific parameters not directly tracked by the calibrations (e.g., total power level variations, weather parameters, etc). QA0 metrics/parameters have been selected to check the health of the whole signal path from the atmosphere down to the back-ends. The parameters can be grouped into the following categories:

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

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

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

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

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

The tolerances for these parameters that have been adopted by ALMA are listed below.

11.2 QA1

QA1 tackles slowly varying (timescales longer than a week) array performance parameters. They will all be measured by AoDs executing standard calibration SBs created as specified by the Calibration Plan. The QA1-related parameters will, in general, be measured at predefined periods during the month as “Observatory Tasks” and, in cases of detection of significant deterioration of performance, during operations. Currently, the different tasks to measure these parameters are done by different software packages. This situation will be changed in the near future by including some of the packages within TelCal and/or CASA. Reduction of the “Observatory Tasks” will be a joint effort of the AoDs and System Astronomers (DMG). The output of the reduction is a set of parameters with errors that are ingested into the TMCDB and can then be used during observations. The most critical QA1 parameters will also be displayed in real time on the Operations Monitor and Control (OMC) for reference by the AoDs. 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

The tolerances for these parameters that have been adopted by ALMA are listed in Appendix A.

11.3 QA2

QA2 deals with QA at the level of the data reduction by the Science Pipeline or semi-interactively by the ALMA Data Reducers Team. It is only at the stage of data reduction and analysis of data products that some of the science goals set by the PI can be compared with the results (i.e., SNR, dynamic range, etc). During Cycle 1, data reduction will still be partly carried out by ALMA staff in a semi-automated fashion using similar scripts to those in the automated Pipeline. However, it is expected that for the basic ALMA standard observing modes the automated Pipeline will be used.

The current list of QA2 parameters can be classified, following the data reduction flow, as:

Calibration Issues: Relative and Absolute Calibration Quality Among Datasets, Bandpass Calibration (flatness and dynamical range achieved), Overall Gain Calibration (granularity, Phase Transfer Quality, WVR Calibration Improvements)

Reduction Process: Data flagging (amount, cause, etc), Cleaning Convergence

Final Data products: SNR Achieved, Sidelobe Levels, Dynamic Range, Contamination by Bright Sources Outside FOV/aliasing, Resolution (spectral & spatial), Comparison of Deconvolution Algorithms, Residual Structures

11.4 QA3

QA3 is post-reduction evaluation of the data products delivered to the PIs. It will be triggered by PIs (or ARC personnel) using the helpdesk to report any problems with the delivered data products to the ARCs. The ARC receiving the helpdesk report will retrieve the data from the archive and evaluate the nature of the problem. The assessment by the ARC should include BOTH an assessment on whether the problem is present only in a particular dataset or 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 process, or the QA process, the ARC will communicate their findings to the DSO, which in return and in collaboration with the ARCs, will work on solving the problem. The result will be communicated back to the reporting investigator.

If the problem is of limited impact (i.e., the specific dataset), the dataset would be fully re-reduced using the SCO Science Pipeline and replicated to the ARC Archives after QA. If the problem affects a significant number of datasets so that a re-reduction might significantly slow down the SCO pipeline operations, the ARCs pipelines could be used once a solution has been implemented. The data would then be re-ingested into the SCO Archive and replicated.

11.5 The Quality Assurance Report

The Quality Assurance Reports that will be delivered to the PIs will be generated using the ALMA Quality Assurance (AQUA) package developed for this purpose. 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 DMG at the end of the data reduction process that also has to be cleared for the data products to be deliverable to the PI. The final report per ObsUnitSet delivered to the PI will be a concatenation of all the relevant QA0,QA1 reports per execution and the QA2 report. Comments on each stage of the QA process (with supporting images, if required) will be added to the Report.

11.6 Pass/Fail Criteria

The QA0 pass/fail criteria that have been adopted by ALMA during Cycle 1 are based on the following:

- **Antennas:** Less than 27 antennas available in the 12-m Array and/or less than 7 antennas in the 7-m Array
- **Receivers:** Less than 27 receivers available at the observing ALMA band in the 12-m Array and/or less than 7 in the 7-m Array

- **Pointing:** Large scatter in pointing for more than 5 antennas ($> 1/10$ HPBW) even after pointing calibration
- **Focus:** More than 3 antennas with offsets from focus model $> 1/5\lambda$ in Z and $> \lambda$ in X/Y
- **Bandpass:** Bandpass too weak, with amplitude wiggles > 3 dB on the autocorrelations (not due to atmospheric features) or strong CW signals
- **System Temperatures:** More than 5 antennas with T_{sys} values $> 50\%$ of the others, with differences between polarizations $> 50\%$ for a given antenna or with $T_{\text{sys}} > 2000$ K
- **Gain:** Phase RMS > 0.5 rad for a significant fraction of the baselines
- **Execution:** All datasets whose execution failed at some point and contain less than 20% of the planned observations
- **Storage:** Data that could not be read from the Archive

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

For QA2, the main criteria are the achievement of the requested noise RMS in the images (it must be within 15% of the goal), the resulting synthesized beam shape and the calibration quality (phase RMS).

Chapter 12

Data Archiving

Data from the correlator, together with additional monitor and weather data, are sent via dedicated optical fiber links (1-10 Gbit/s) to the OSF, where they are archived. The system has been designed to cope with 6.6 MB/s average data rates and 66.6 MB/s peak rates (for short periods of time). The archive at the OSF is designed to provide up to a year of temporary storage for the instrumental data (in the form of “ALMA Science Data Model”, or ASDM, files) and the monitoring data. The instrumental data are then transferred to the main archive at the SCO, where the pipeline is run and from where the data and pipeline products are distributed to the three ALMA Regional Centers (ARCs) in North America, Europe and Japan (Figure 12.2). The process of copying the data to any of the archives involves a replication of the metadata (support data) and a mirroring of the bulk data (ASDM and FITS files), see Figure 12.3. These two processes do not need to be simultaneous and they can use different routes (i.e., media delivery and internet). It is expected to use Oracle stream technology for the copies of the metadata. The ALMA archive consists of two parts - a front-end archive, which is used for storing details of observational scheduling blocks and related data needed to execute observations, and the back-end (science) archive, which stores the instrumental and processed data. The architecture is based on the NGAS system (New Generation Archive System) with Oracle technology for the metadata, as shown in Figure 12.1.

From Cycle 1 on, an automated reduction Pipeline, located at the SCO will be used to provide an initial processing of the data. Copies of the Pipeline will also be available at the ARCs, to support reduction of projects with large data throughput and/or specific user-requested re-reductions of limited sets of data. For Cycle 1, reduction will be carried out by ALMA staff based on the Science Pipeline scripts. Successful proposers, or “Principal Investigators” (PIs), will access their data via the Science Portal after authentication. The data will be directly downloadable by PIs via the Internet or shipped to them on physical media from the nearest ARC if the datasets are large (i.e., tens of GBs, TBD). The data products delivered to the PI will include ALL the relevant information for him/her to repeat the reduction process, final fully reduced and calibrated FITS (Flexible Image Transport System) cubes, Flagging and Calibration tables, Quality Assurance reports and Observing logs. Raw ASDM files will also be made available upon request. The PI will be notified by the corresponding ARC whenever data products are available from a completed pipeline run on a “ObsUnitSet” (the amount of an observing project that is observed before running the data through the science pipeline; this is set up during the project planning “phase 2” process - see the Proposer’s Guide). The proprietary period for an ObsUnitSet is twelve months from the time the data is available to the PIs (i.e., at the ARCs Archives), after which data will be available from the Science Portal without the need to authenticate (unless shipment on storage media is requested).

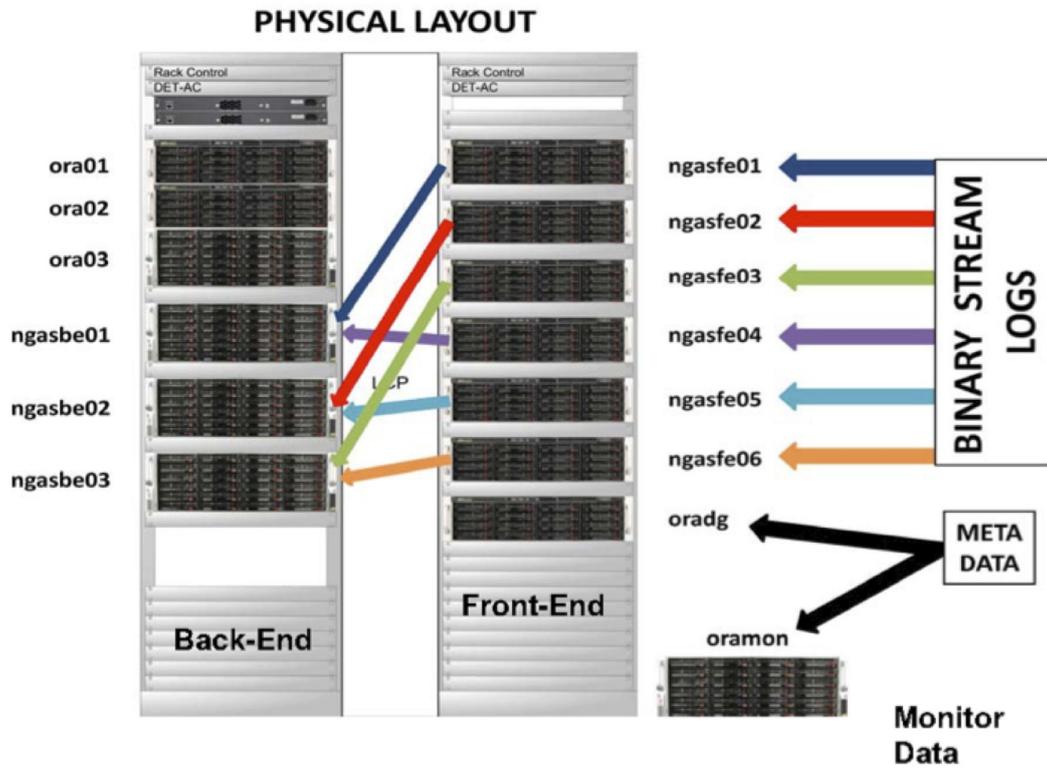


Figure 12.1: Archive design (front-end & Back-end) at the OSF to store metadata and raw and monitor data.

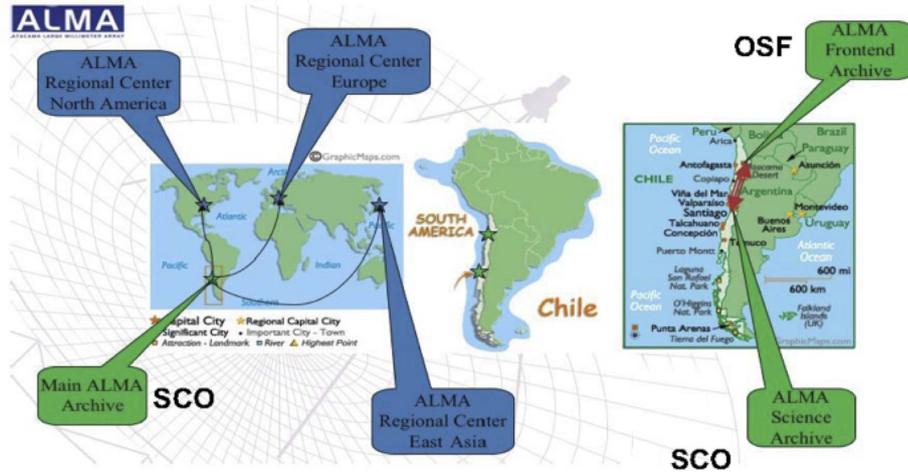


Figure 12.2: Location of the ALMA archives.

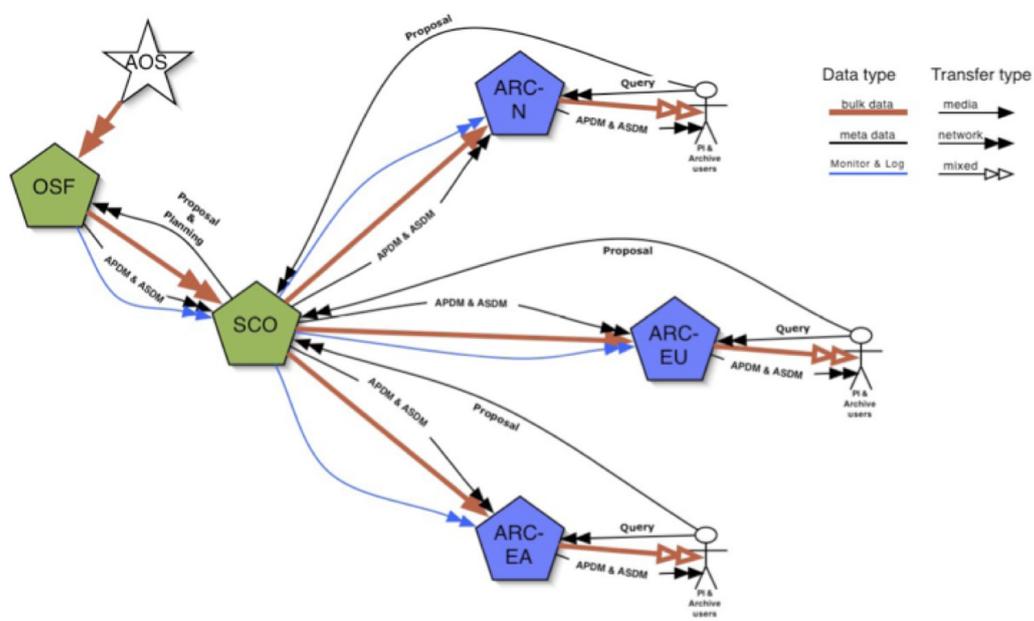


Figure 12.3: Data flow from the AOS down to the ARCs.

Appendix A

Appendix

A.1 Antennas

At the end of the construction period ALMA will have in total 66 antennas, 54 with a diameter of 12 m and 12 with a diameter of 7 m. Four of the 12 m antennas will be equipped with a nutating subreflector for total power observations. The four antennas used for total power observations and the twelve 7 m antennas will together form the Atacama Compact Array (ACA). The ALMA antennas are manufactured by three different contractors. These are VertexRSI (North America) which will provide 25 12 m antennas, Alcatel Alenia Space European Industrial Engineering MT Aerospace (AEM, Europe), which will provide 25 12 m antennas and Mitsubishi Electric Corporation (MELCO; East Asia), which will provide the four 12 m total power antennas and the twelve 7 m antennas (Figure A.1).

All antennas have been designed to meet very stringent ALMA performance criteria, and to successfully operate under the extreme environmental conditions at the Array Operation Site (AOS), i.e. strong winds, large temperature ranges and gradients, solar irradiation and snow. The primary operating conditions used for the design of the antennas 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 1 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.

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 nearinfrared solar radiation.

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

All antennas have a Cassegrain cabin that is kept at a constant temperature of 20 degrees Centigrade and contains the receivers, the amplitude calibration device and associated electronics.

A shutter protects the inside of the Cassegrain cabin when the antenna is not operating. A membrane transparent to the frequencies that can be observed with ALMA is located below the shutter to prevent airflow from the cabin to the outside when the shutter is open. The current design uses a 0.5 mm thick Goretex membrane.

The different antennas use a combination of steel, aluminum, Carbon Fiber Reinforced Plastic (CFRP) and Invar in order 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.

	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 12m	CFRP	7/205	Al	+	Steel	Direct	Reference Frame metrology
Melco 7m	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.

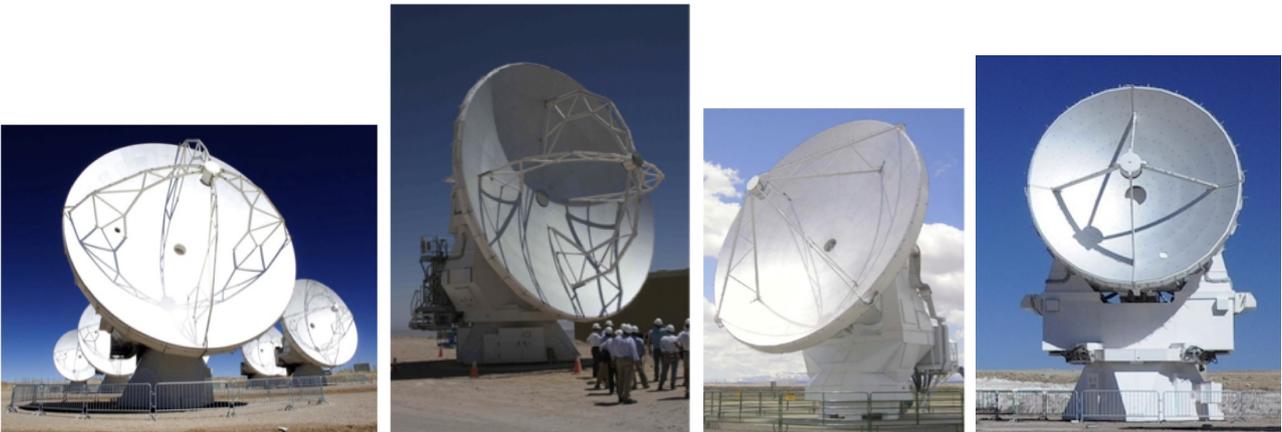


Figure A.1: The four different ALMA Antenna designs: Vertex 12 m, MELCO 12 m, AEM 12 m, and MELCO 7 m (from left to right).

A.2 Antenna Foundations



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

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

The minimum stiffness which the foundation must exhibit at each insert is:

- Vertical stiffness (Z) > 13 x 10⁹ N/m
- In X/Y plane > 9 x 10⁹ N/m

This 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.

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) in order to achieve the required antenna positioning accuracy. The antennas can be positioned to within a few millimeters, ensuring accurate placement on the antenna foundation pads. More information on the transporters can be found on the ALMA EPO pages¹.

¹<http://www.almaobservatory.org/en/technology/transporters>



Figure A.3: The ALMA array with eight 12-m antennas (left), and an antenna being transported to the AOS (right).

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 roomtemperature electronics for the IF and LO for each band (the warm cartridge assembly, WCA) and fore-optics and entrance windows for each band. The water vapor radiometer (WVR) is mounted to one side of the cryostat using a pickoff mirror to direct the antenna beam into the WVR. The Amplitude Calibration Device (ACD) is mounted above the frontend, and is described in Section A.5. Figure A.4 and A.5 show overviews of the frontend unit, with the cylindrical cryostat on top and the roomtemperature 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 in order 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.

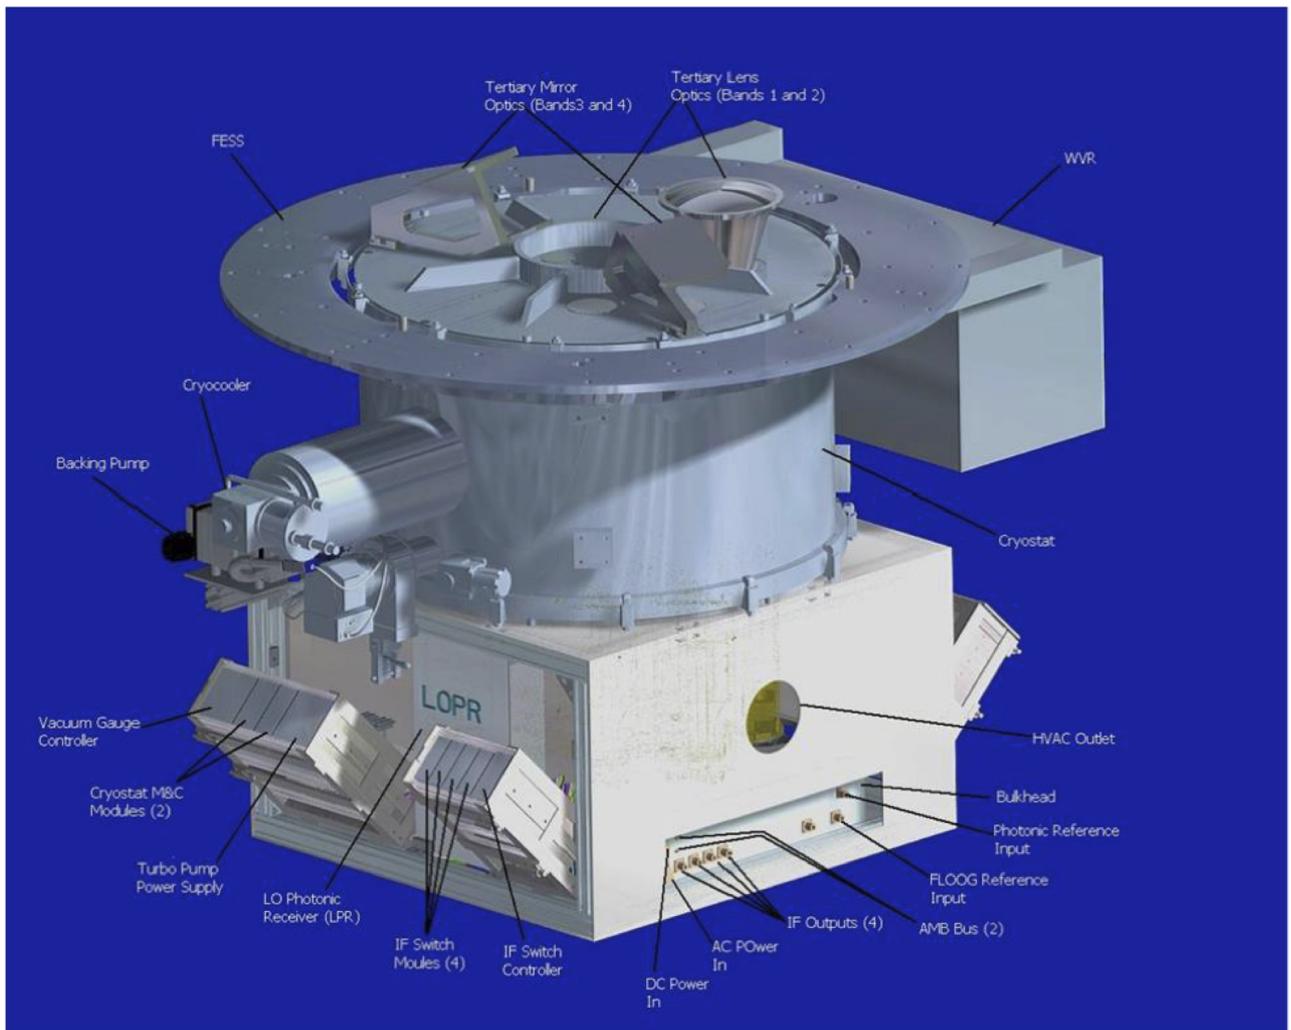


Figure A.4: Side view of ALMA frontend showing cryostat assembly, with room temperature unit below.

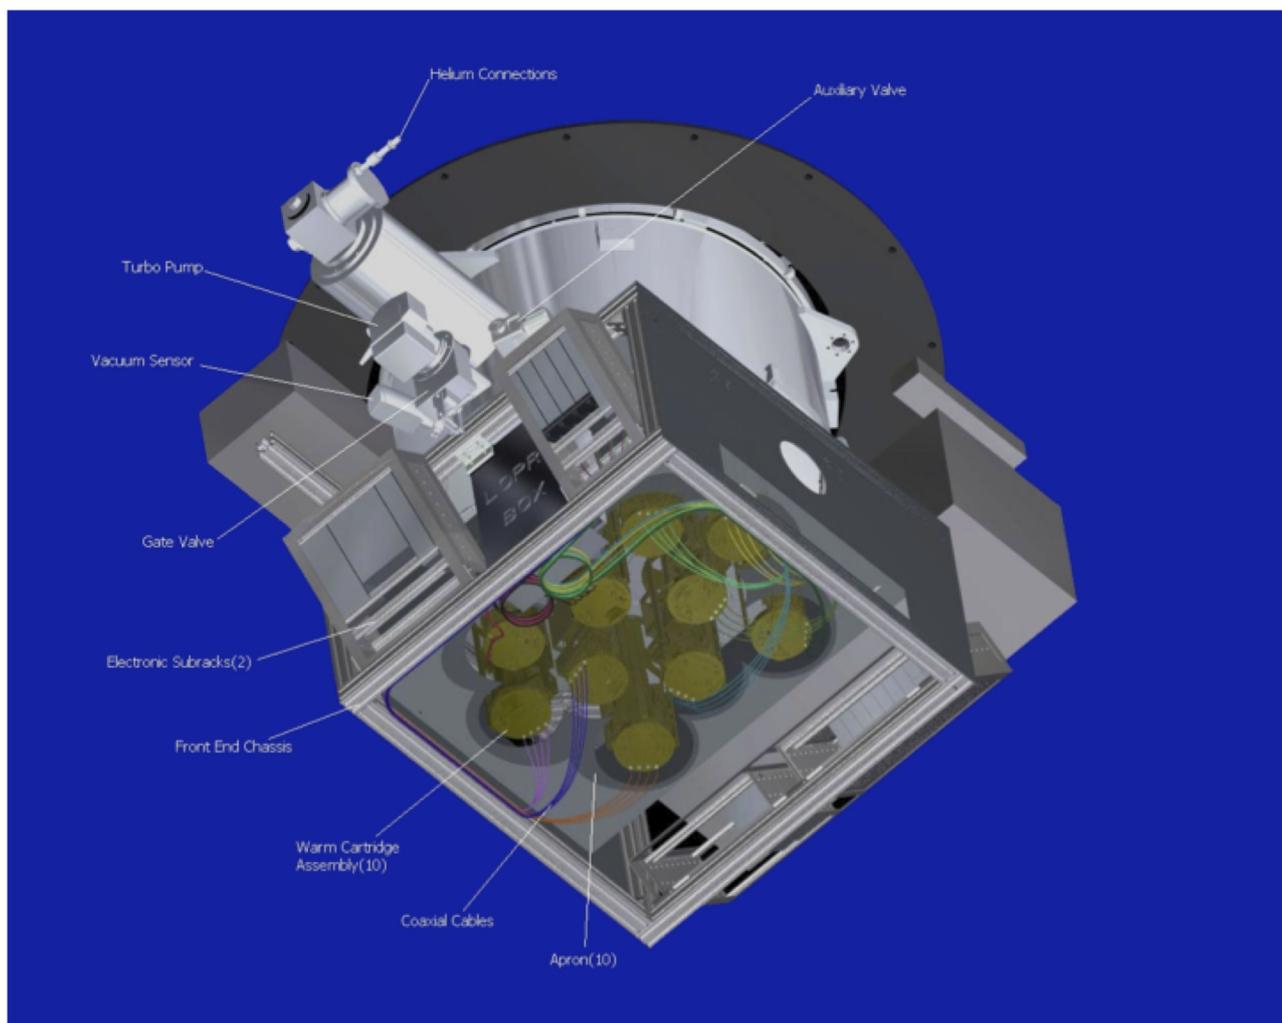


Figure A.5: Bottom view of ALMA frontend, showing WCAs.

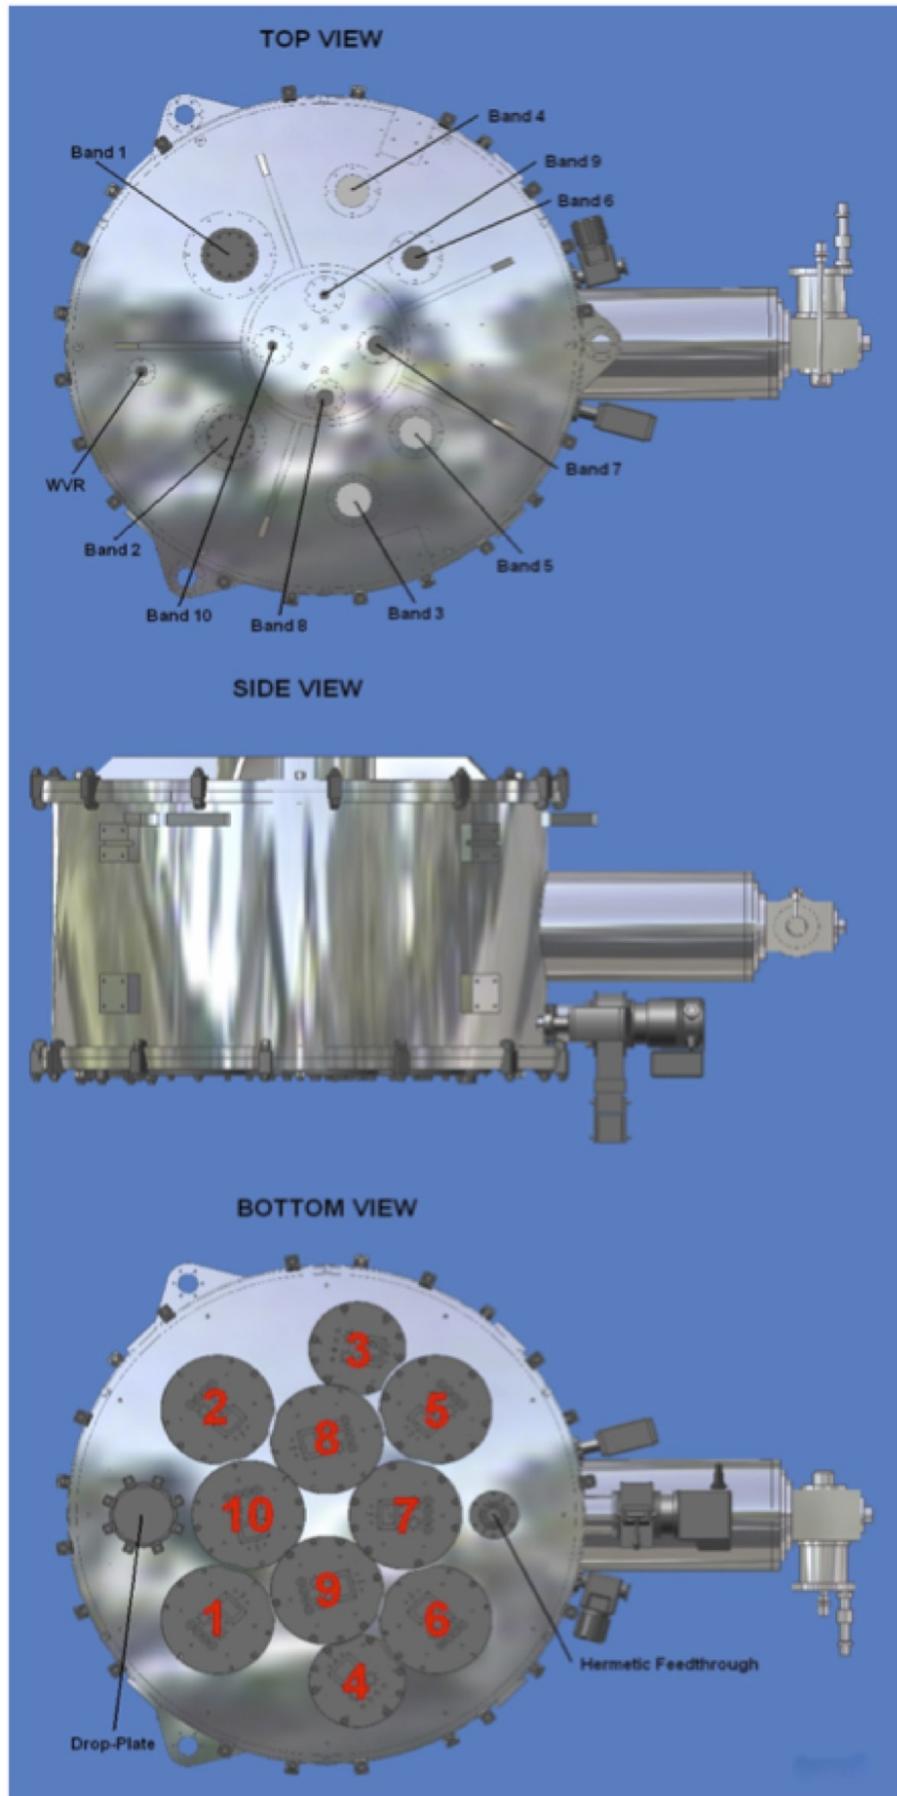


Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCs for each band (lower view).

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 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 1). 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.

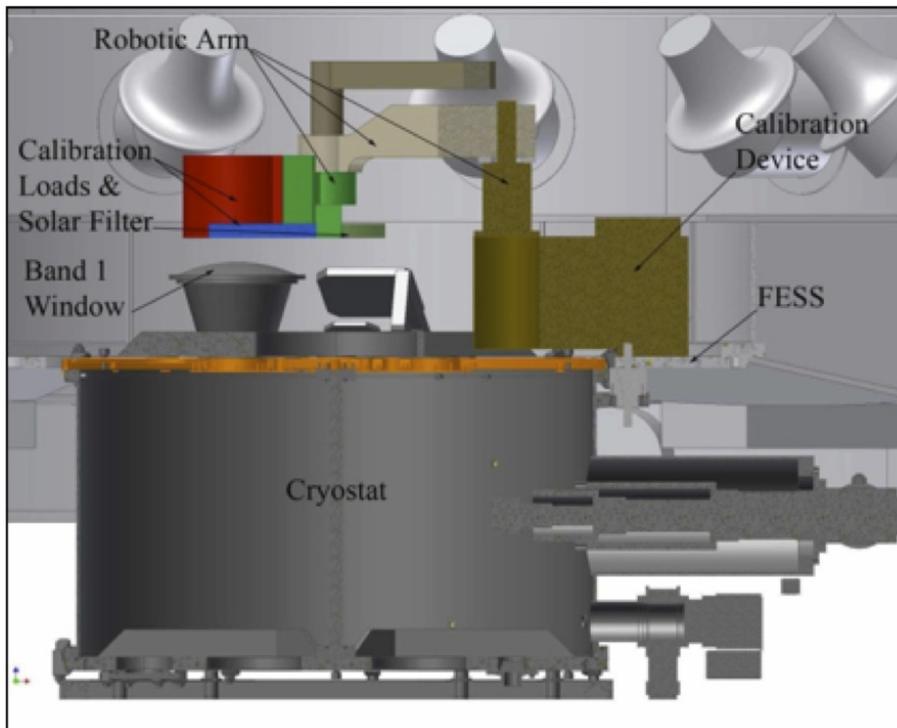


Figure A.7: Lateral view of the ACD on top of the ALMA frontends.

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, frontend status, etc.).

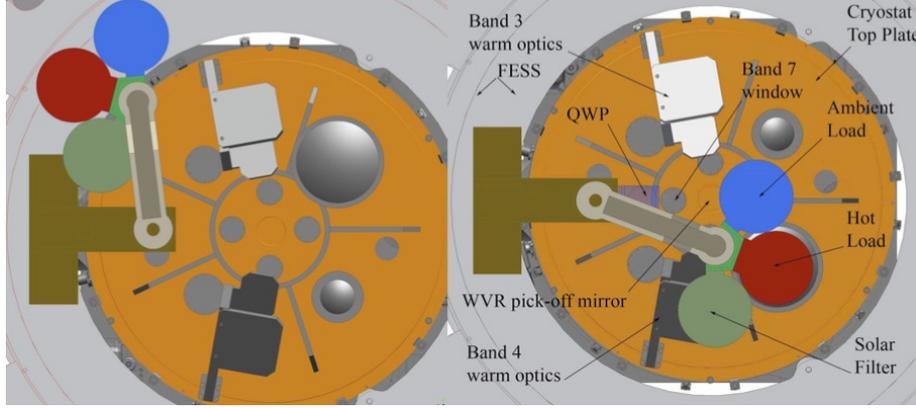


Figure A.8: Top view of an ALMA front end showing the robotic arm of the ACD retracted during normal observations or on top of one of the frontend inserts for calibration. The current design has been improved by placing all the loads in a wheel.

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 1.

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 12m antenna (The ACA antennas do not have WVRs). For the WVRs to be effective, the measurements have to be taken with a cadence that is fast enough to map the actual variations in the atmosphere. The relevant shortest timescale is the antenna diameter divided by the wind speed as the path delay is averaged over the whole antenna beam and cannot therefore be corrected at any finer time resolution than that. The effective diameter is about 10 m for the ALMA antennas and the relevant windspeed is usually 10 m/s or a bit less so the fastest necessary sampling speed is 1Hz. On timescales shorter than this 1 Hz timescale, the watervapor path fluctuations are expected to lead to small apparent pointing fluctuations which are analogous to the seeing effects in single-aperture optical telescopes. ALMA selected the 183 GHz line because it is quite bright and allows a more compact design than would the 22 GHz water line. It was decided to measure the temperature of the 183 GHz line in four regions offset from the center using filters of different bandwidths. The positions of the filters are indicated as blue boxes superimposed on the profile of the water vapor line in Figure A.9. The sensitivity specification for the WVRs is 0.08–0.1 K per channel RMS.

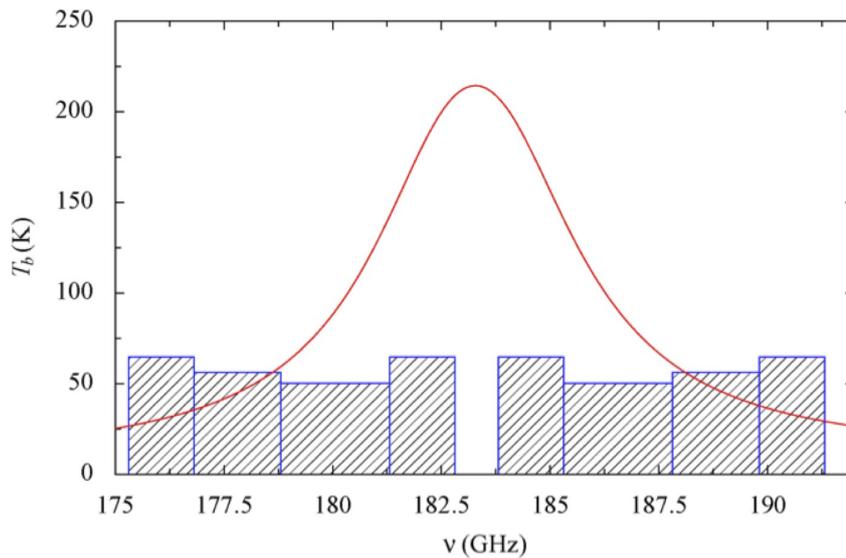


Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

It is very important that the WVR illuminates the same area of the sky as the ALMA band receivers in the near-field region. This is because the origin of the water vapor fluctuations is usually located in the lower troposphere (i.e., near the observatory), with one to several layers of water vapor clumps encompassing a wide range of sizes. Since the ALMA frontends 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”

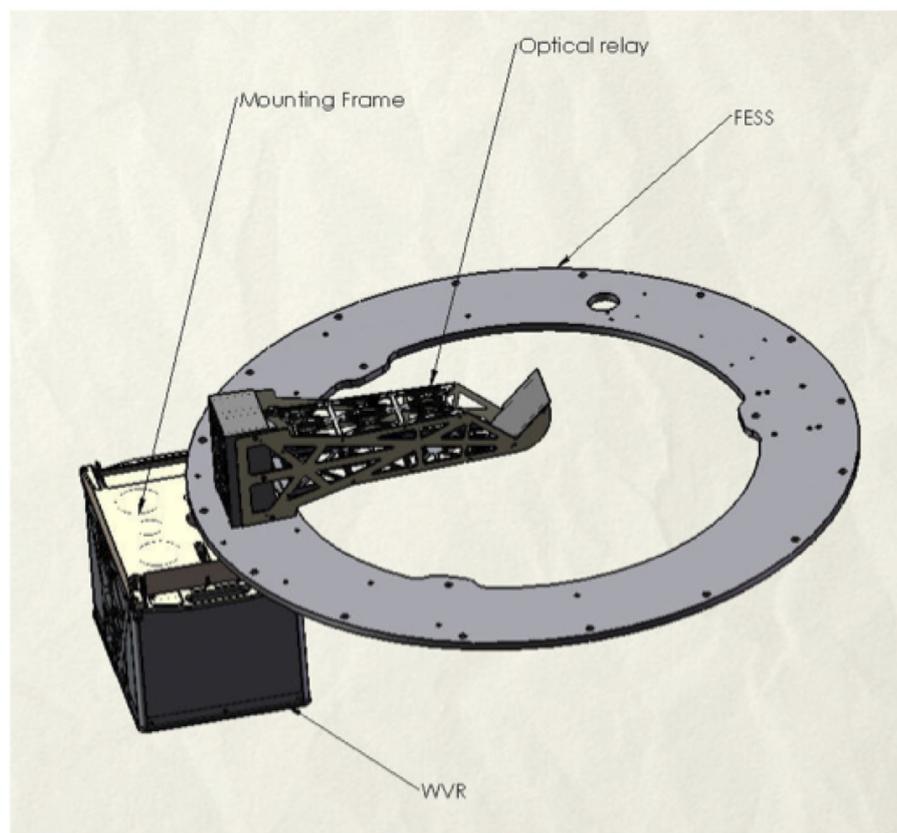


Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

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 1mm pwv, the residual term in the formula would be $20\mu\text{m}$. 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.

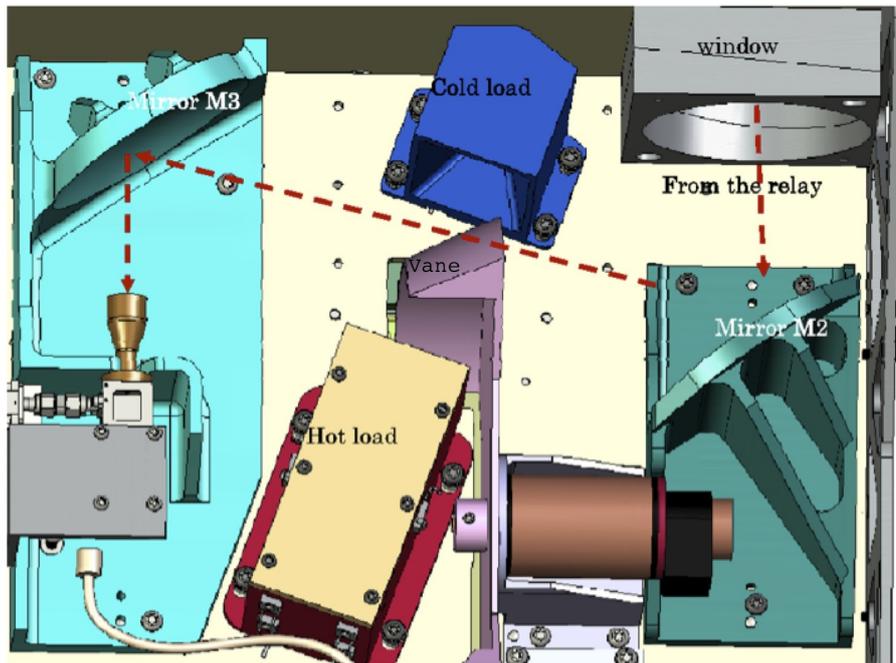


Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

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 (max error). The mixer system is an un-cooled DSB Schottky diode pumped by an LO at 15 GHz that undergoes 2 stages of multiplication. The receiver noise temperature is about 1000 K. After amplification, the IF signal is split into four complete chains (one per filter) and a bandpass filter is applied to select the four desired sampling regions in the profile of the water vapor emission line. In each IF chain, the signal is detected with diodes and after a Voltage-to-Frequency conversion, sent to the Control section for accumulation and control. There is a possibility of LO leakage out of the WVRs that could affect the

ALMA receivers in the same antenna and others nearby. To avoid coherence, all the WVRs are tuned to a frequency slightly different (offsets by consecutive integer multiples of 10 kHz up to the total number of WVRs available). The final 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.

A.7 The LO and IF System

The Local Oscillator (LO) and Intermediate Frequency (IF) subsystems lie on the signal path between the Frontends (FE) and the correlators; they perform numerous functions, including:

1. Downconversion of the observed sky frequencies to a band in the range 2–4 GHz, which aliases down to 0–2 GHz for digitization.
2. Setting the correct power levels into the digitizers
3. Setting the center frequencies of the spws in the Correlator FDM modes (this is actually done in the correlator, but is effectively part 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. Providing 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.
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.

Frequency downconversion takes place in several stages and involves 2 hardware Local Oscillators (LO1 and LO2), a 4 GHz sampler and an LO synthesised in the TFBs in the Correlator. An overview of the operation is given in A.7.1. More detailed descriptions of the individual subsystems are given in Section A.7.2.

A.7.1 Summary of operation

Figure A.12 shows a simplified block diagram of the operation of the LO/IF system, including example frequencies for an observing frequency centered on 100 GHz. Referring to this diagram, the system operates in the following way:

1. The frontend 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 all the ALMA bands, since the mixers for Bands 3 and 7 have an output IF of 4-8 GHz, Band 6 a range of 5-10 GHz and Band 9 a range of 4-12 GHz. Over most of the frontend tuning range, LO1 and the frontend mixer can be used in upper or lower sideband; although at the edges of the tuning band, only one sideband is possible. LO1 consists of a common component for all antennas, plus a small offset component generated in the FLOOG (First LO Offset Generator) which is different for each antenna (see LO1 Section A.7.2). The FLOOG is used to perform coarse fringe tracking (i.e. rough correction for the small offsets in 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, and will also be used for sideband suppression and, for the DSB receivers, for sideband separation.

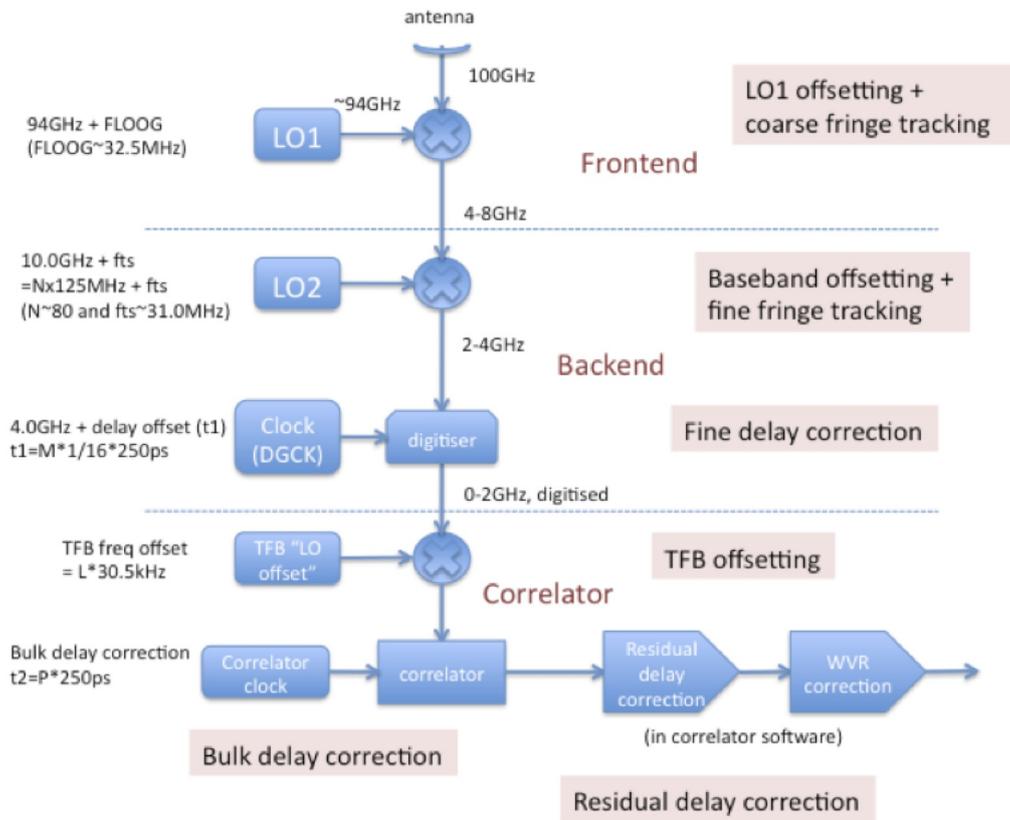


Figure A.12: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the Frontend, Backend, and Correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized tuning steps, e.g. LO2 (which has steps of 125 MHz, with a factor of “N”, with an offset of “fts”), TFB LO (steps of 30.5 kHz, with a factor of “L”) and the Bulk Delay Correction (steps of 250 ps, with a factor of “P”). See text for descriptions of each stage.

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 A.7.2). LO2 is used to offset the individual baseband frequencies within the IF range. The LO2 and second mixer only operates in LSB, with an LO2 tuning range of 8-14 GHz which can only be set in steps of 125 MHz plus a fine-tuned offset ("fts" in Figure A.12). The fts range and the 125MHz quantization means that LO2 setting is not fully contiguous, and there may be up to ~30MHz difference between the desired and the set value. An algorithm used by the OT and the realtime system generates the best LO "solution" which minimizes these offsets. However, without additional correction, with multiple basebands the requested lines may be offset from the spw center by up to this amount; this becomes significant in high-resolution narrowband correlator modes. The remaining difference is compensated for by applying an opposite offset to the TFBLO (LO4 - see item 4 below.)³. The finely-tunable fts is used for fine fringe tracking. LO2 can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband.
3. The analog IF signal from the second mixer is 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, digital filters (known as TFBs, or "Tunable Filterbanks") are applied to the signal, each of which can be individually offset in frequency (the TFB offsetting). This is effectively applying a digital LO (the TFB LO, 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 II of the Obsprep, the TFB is centered on the baseband if the TFB "offset" is set to the default of 3000.0 MHz; it can be moved +/-900 MHz from that frequency, the range depending on the spw bandwidth. The TFB outputs are resampled and sent to the correlator. The TFB LO can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband. The correlator is used to perform the finest level of residual delay correction.

A.7.2 LO/IF components

Figure A.13 shows the components involved in the LO generation and distribution. Outputs from this are LO1, the LO in the IFP (LO2) and the digitizer clock (DGCK). All of these are required in the antenna (in the upper section of Figure A.13). A fibre-optics system is used to distribute these signals from the Central LO (CLO) in the AOS Technical Building out to the antennas. Path length correction using Line Length Correctors (LLCs) is necessary in the fibre distribution. In the following subsections, we describe some of these components.

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) at a fixed wavelength of 1556 nm and the tunable Slave Laser (SL) which is offset from the Master Laser signal, and generated in the CLO (see Section A.7.2). The offset frequency can be anywhere in the range 27 – 122 GHz. The beat note from the two lasers constitutes the Photonic LO Reference, and the LO1 reference signal is produced from this by photomixers located in the Warm Cartridge Assembly (WCA) of each receiver. The reference signal is used to drive a YIG (Yttrium Iron Garnet) oscillator via a Phase Locked Loop (PLL) circuit. This produces the LO1 signal via a set of multipliers (see example Figure A.14 for Band 7). The same photonic reference signal is distributed to all antennas in the same subarray. However, to correct for different delay rates required in different antennas, the First LO Offset Generator (FLOOG) in each antenna generates a variable offset frequency of 20-45 MHz which is also fed into each PLL.

³This is done automatically when the OT generates a spectral setup in an SB from a proposal. However, it is repeated at runtime. With Doppler setting in Cycle 1 (Section 4.6), the LO solutions at SB generation and runtime may be different, requiring a different LO4 offset from the SB. This issue will be fixed before Cycle 1 observing starts.

⁴The OT "adjust" button quantizes the value entered by the user to this unit

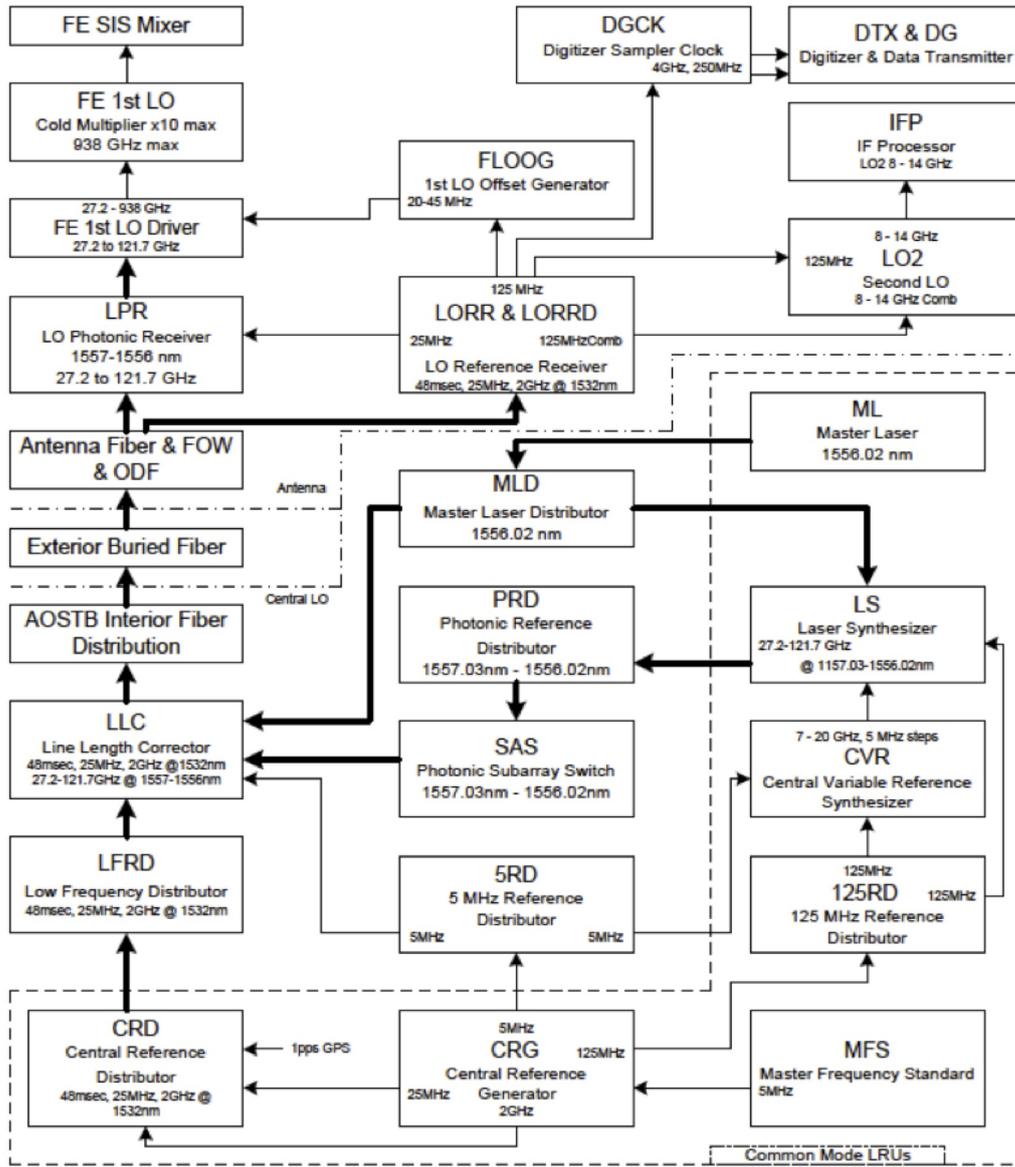


Figure A.13: Summary block diagram of the LO distribution system.

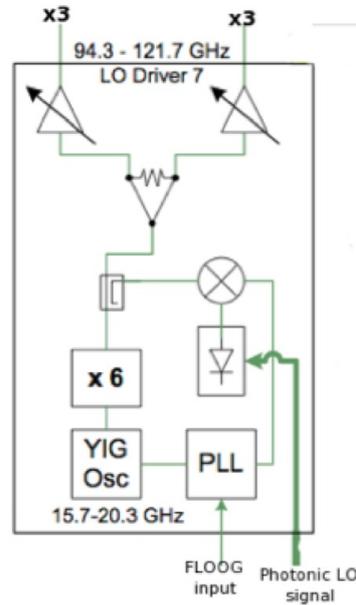


Figure A.14: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier (in this case, x3) is used to generate the LO1 frequency, at 282.9 – 365.1 GHz. The photonic LO signal (green) feeds a photomixer which creates a beat signal between the ML and SL frequencies. This is mixed with the LO and feeds the PLL. The FLOOR generates a small offset frequency which is different for each antenna. See text for details.

The IF switch and IF Processor units

The Band 3, 6 and 7 receivers are dual-sideband (2SB), where both the upper and lower sidebands signals are provided separately and simultaneously. So there are four outputs from each receiver cartridge, two per polarization carrying one of the sidebands each with a signal IF bandwidth of 4 GHz each (or 8 GHz per polarization). For Band 9, the receivers are double-sideband (DSB), where the mixer produces a downconverted output from signals in both USB and LSB. It has only two outputs, one per polarization. However, the signal IF bandwidth of Band 9 is 8 GHz per polarization.

The output of each frontend 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 A.15). The IF processors divide the incoming 4-8 GHz 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 output bandwidth of the receiver (see Section 4 and Table 4.2 for limitations). The LO2s are common to both mixer polarizations which means that both polarizations will have the same spectral setups.

The IF processor unit has Total Power detectors for tuning/optimization of the IF power levels into the digital samplers, although they are not used for observations. Figure A.15 shows a basic block diagram of the IF Processor. It is important to note that the switch network 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 LO2s can be used for sideband separation when combined with the first LO (Section A.7).

The IF processor also has anti-aliasing filters, which define the 2 GHz baseband width and remove out-of-band signals (Section A.7.2). This results in the higher noise levels on the upper and lower 50-100 MHz of

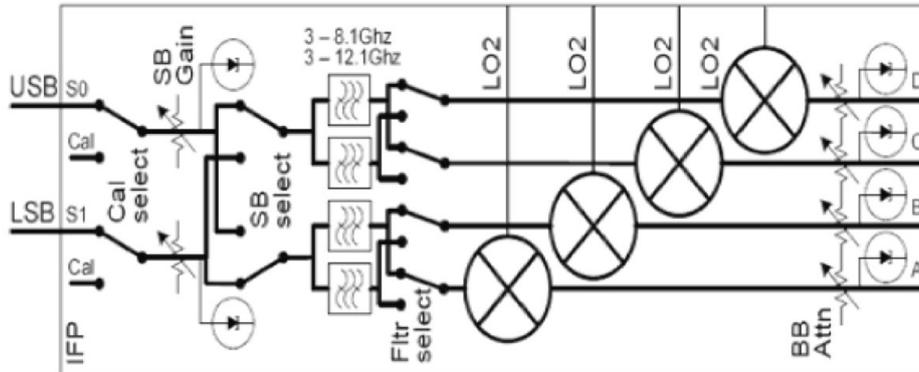


Figure A.15: Block diagram of IF Processor.

channels in the TDM correlator mode (see Section 4.4). These filters cause a decrease in the effective IF range to approximately 1.875 GHz.

Digitization and 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 A.16 shows a block diagram of a single DTS module.

Reference and LO Signal Generation and Distribution

The Central Local Oscillator (CLO) system generates and distributes the reference, timing and LO signals to all ALMA components in order to ensure that antenna movement, electronics, and data acquisition are synchronized. The signals are distributed to the antennas through optical fiber using the light of three infrared lasers. Figure A.17 shows a block diagram of the CLO.

Reference Signal generation

The ALMA frequency and phase standard is a Rubidium atomic clock, the Master Frequency Standard (MFS), which produces a signal at a frequency of 5 MHz. This signal is fed into the Central Reference Generator (CRG) module, which produces several signals as multiples of the 5 MHz signal. The 5 MHz signal is fed into the Line Length Corrector (LLC, See Section A.7.2). The 125 MHz signal becomes one of the standards used by many components in the ALMA system. At the AOS Technical building it is used by the Slave Lasers in the Laser

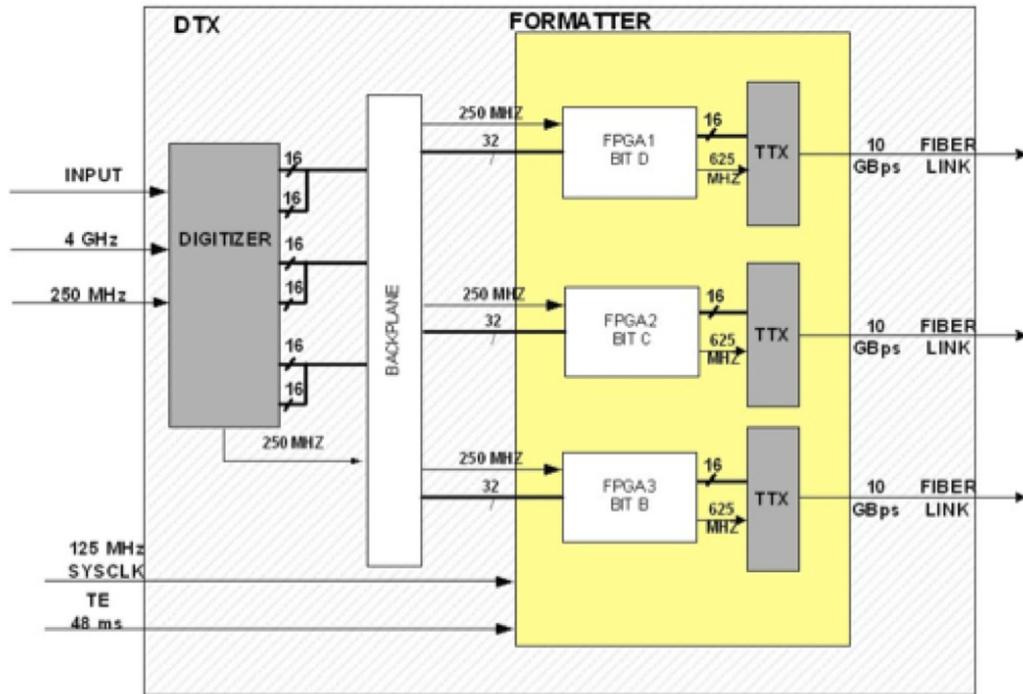


Figure A.16: The DTX and DRX Signal Digitization and Transmission system.

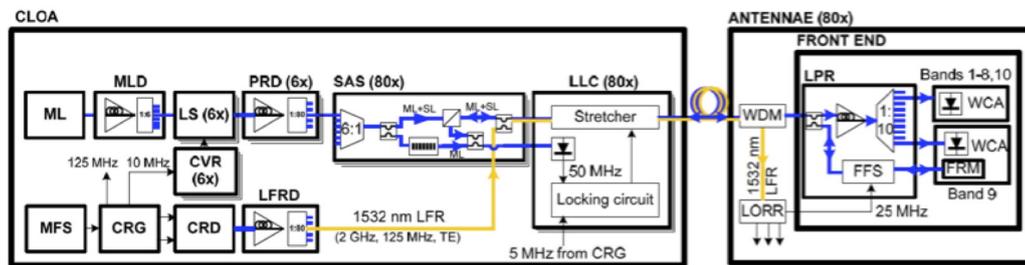


Figure A.17: LO block diagram, showing the Central LO (CLOA) and the LO section in the WCA in each frontend. For description of acronyms, see text.

Synthesizer modules (Section A.7.2). At the antennas, it is fed into the FLOOG, the Digital Clock (DGCK, see Section A.7.2) and the LO2 Synthesizers. The 2 GHz signal goes into the DTS cards at the antennas. All these reference signals are modulated into a 1532 nm IR laser in the Central Reference Distributor (CRD) module. The CRD has an internal 48ms (TE) clock that is also modulated into the same signal. The modulated 1532 nm signal is sent to an optical distributor (with 80 outputs), the Low Frequency Reference Distributor (LFRD), that feeds it into the Sub Array Switch (SAS) modules, where it is merged with the signals from the Master and Slave lasers (see Section A.7.2).

LO signal generation

The 1st LO is generated photonicly in each antenna frontend by mixing the two infrared laser carriers from the Master (ML) and Slave Lasers (SL) to produce a fixed frequency for all the antennas. There are 6 Slave Laser systems (Laser Synthesizers, LSs) that produce 6 different LO1 frequencies which allow simultaneous observations at different frequencies with different subsets of the array (subarrays).

The laser frequencies are generated in the CLO in the following way:

- The Master Laser (ML) generates a 1556 nm fixed optical reference signal, which feeds the Master Laser Distributor (MLD) – essentially a 6-way splitter.
- The Central Reference Generator (CRG) produces reference signals that are fed into the 6 Laser Synthesizers (LS). The LSs controls the frequency of the the Slave Lasers producing a frequency offset of the SL signal of 27-120 GHz with respect to the ML signal. The SL signals are added to the ML signal. The offsets between the ML and SL signals provides the beat note which is used to generate the LO1 frequency in the photomixers in the WCAs in the frontend. It is used to set up the LO1 frontend observing frequency, and is set by the software. With 6 LSs it is possible to generate 6 separate LO1 frequencies for the different subarrays.
- The Photonic Reference Distribution (PRD) feeds the optical signals to the Sub Array Switch (SAS) which can distribute the signals to the different subarrays,

Both the reference signals and the LO signals are fed through the Line Length Corrector (LLC), which is used to correct for changes in the optical fibres. The LLCs are described below.

Optical Signal Distribution

Figure A.18 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 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.

LO Path Length Corrections

The LO Reference signals are generated at the AOS Technical Building and distributed via optical fibers to all the antennas. To guarantee that the phase of the LO signals is stable during the observations for fibers of up to 15km in length, compensation for the 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. A block diagram of the Line Length Correction system implemented by ALMA is shown in Figure A.19.

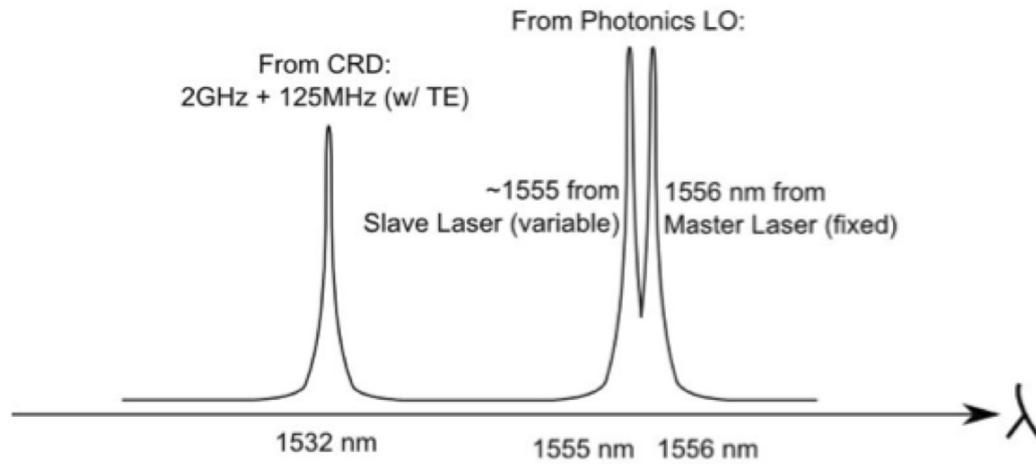


Figure A.18: The ALMA reference signals. Within each antenna, the optical fibers are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing signals, and the LO Photonic Receiver for the LO Reference signals.

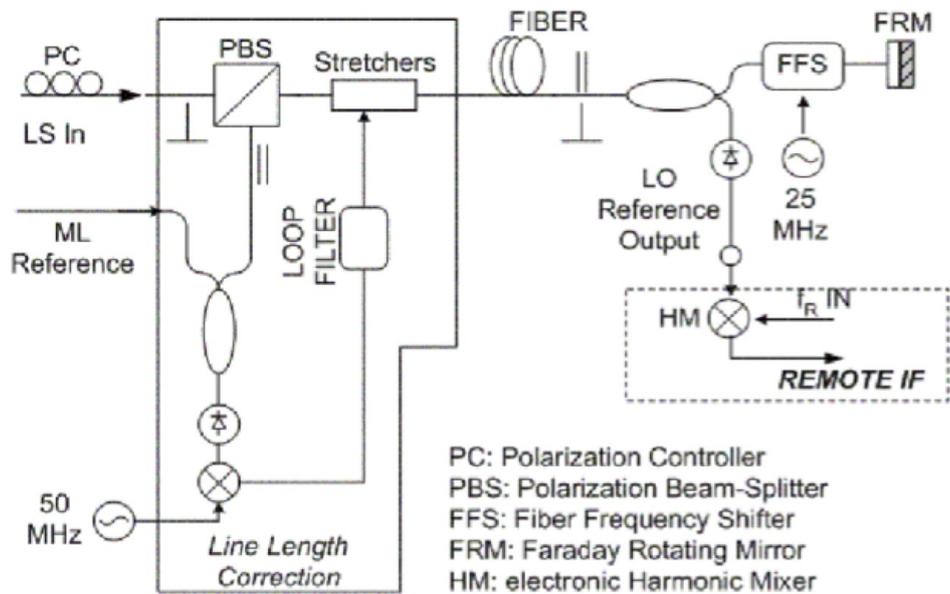


Figure A.19: Block diagram of the Line Length Corrector system for ALMA.

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 the Band 9 cartridge in each frontend. 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 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 5mm in two modes. A “slow” mode (about 10Hz) copes with the large deformations (about 3mm, allowing for some headroom at the ends of the ranges) and a “fast” response mode (about 1kHz) copes with the small range variations (about 0.1mm).

A.7.3 Delay corrections, sideband suppression and interference rejection

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

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 A.12). 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.

Sideband suppression - LO offsetting

Some of the ALMA receivers (e.g. band 3, 6 & 7) are inherently single sideband (SSB), either through having a mixer or quasi-optic design which rejects the unwanted sideband. But their intrinsic sideband rejection is typically only about 10-15 dB - adequate for rejection of the unwanted sky noise, but not enough to remove strong lines from the other sideband. Others receivers (e.g. band 9) are double sideband (DSB), and the relative response of the two sidebands may not be equal, significantly affecting calibration. Accordingly, additional schemes are necessary for more effective removal of the unwanted sideband (known as sideband suppression), and for correlation of both sidebands independently (or sideband separation - see next section). Sideband suppression in ALMA is done using the FLOOG, and either LO2 (2LO offsetting) or a combination of LO2 and LO4 (3LO offsetting). A small frequency offset F_o is added to LO1 and subtracted from the other LOs, so that while the signal sideband remains at the same frequency, the image sideband is shifted $2F_o$ away from it's

nominal value. A different value of F_o is applied at every antenna, 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.

LO offsetting cannot suppress the image sideband in autocorrelation signals, and other techniques such as sweeping the LO2 offset need to be used; these are under investigation for Cycle 1.

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 (DSB) it will effectively double the bandwidth from 8 GHz to 16 GHz per polarization. This is under development and will not be available for Cycle 1.

Interference rejection - 180 degree Walsh 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 16ms. 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 B

Acronym List

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
ASC	ALMA Sensitivity Calculator
ASDM	ALMA Science Data Model
AZ	Azimuth
BB	Baseband
BE	Backend
BL	Baseline
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
EL	Elevation
EPO	Education and Public Outreach
ES	Early Science
ESO	European Southern Observatory
FDM	Frequency Division Mode
FE	Frontend
FITS	Flexible Image Transport System
FLOOG	First LO Offset Generator
FOM	Fiber Optic Multiplexer
FOV	Field of View
FPGA	Field-Programmable Gate Array
FT	Fourier Transform
FWHM	Full Width Half Maximum
GPS	Global Positioning System

HA	Hour Angle
HEMT	High Electron Mobility Transistor
IF	Intermediate Frequency
IFP	Intermediate Frequency Processor
IRAM	Institut de Radioastronomie Millimetrique
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
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	Netherlands Institute for Space Research
SSB	Single Sideband
2SB	Sideband separating Mixer
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable Filterbanks
TFB LO	Local Oscillator at the Tunable Filterbanks
Tsys	System Temperature
T_{rx}	Receiver Temperature
USB	Upper Sideband
VLA	Very Large Array
WCA	Warm Cartridge Assembly
WVR	Water Vapor Radiometer
XF	Correlation-Fourier Transform Type Correlator
YIG	Yttrium-Iron Garnet Oscillator



The Atacama Large Millimeter/submillimeter Array (ALMA), an international astronomy facility, is a partnership of Europe, North America and East Asia in cooperation with the Republic of Chile. ALMA is funded in Europe by the European Organization for Astronomical Research in the Southern Hemisphere (ESO), in North America by the U.S. National Science Foundation (NSF) in cooperation with the National Research Council of Canada (NRC) and the National Science Council of Taiwan (NSC) and in East Asia by the National Institutes of Natural Sciences (NINS) of Japan in cooperation with the Academia Sinica (AS) in Taiwan. ALMA construction and operations are led on behalf of Europe by ESO, on behalf of North America by the National Radio Astronomy Observatory (NRAO), which is managed by Associated Universities, Inc. (AUI) and on behalf of East Asia by the National Astronomical Observatory of Japan (NAOJ). The Joint ALMA Observatory (JAO) provides the unified leadership and management of the construction, commissioning and operation of ALMA.

