Doc 5.3, ver. 1.0 | March 21, 2017

ALMA Cycle 5 Technical Handbook

www.almascience.org

User Support

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

Revision History

Cycle	Version	Date	Editors	Remarks
1	1.0	2013-10-01	A. Hales, A. Lundgren, A. Biggs, J. Di Francesco, B. Vila Vilaro, S. Leon, D. Barkats, E. Fomalont, S. Asayama, S. Komugi	
2	1.1	2013-10-15	A. Lundgren, S. Leon, J. Hibbard, E. Müller, J. Di Francesco	
3	1.0	2015-03-20	A. Remijan, M. Adams, R. Warmels	Cycle 3
4	1.0	2016-03-22	J. Mangum and R. Warmels	Cycle 4
4	1.1	2016-09-07	J. Mangum and R. Warmels	Replace figs 7.9 and 7.10
5	1.1	2017-03-21	R. Warmels and A. Remijan	Cycle 5

Editors

Rein Warmels and Tony Remijan

Authors of this Edition

Itziar de Gregorio - Chapter 1; Antonio Hales - Chapter 2 and Appendix A; James Di Francesco - Chapter 3; Gianni Marconi and Bill Dent - Chapter 4; Seiji Kameno - Chapter 5; Bill Dent - Chapter 6; Eric Villard - Chapter 7; Baltasar Vila Vilaro - Chapters 8 ,11 and 12; Andy Biggs - Chapter 9; Ed Fomalont - Chapter 10; John Hibbard - Chapter 13; Felix Stoehr - Chapter 14; Shin'ichiro Asayama - Appendix B

Contributors to this Edition

Eiji Akiyama, Rodrigo Amestica, Yiping Ao, Yanett Contreras, Todd Hunter, Woojin Kwon, Mark Lacy, Harvey Liszt, Jeff Mangum, Brian Mason, Ralph Marson, Rie Miura, Arielle Moullet, Jennifer Donovan Meyer, Hiroshi Nagai, Kouichiro Nakanishi, Se-Heon Oh, Rosita Paladino, Anita Richards, Rachel Rosen, Kazy Rygi, Kamaljeet Saini, Richard Simon, Gie-Han Tan, Ken Tatematsu, Carmen Toribio, Lance Utley, Catherine Vlahakis, Al Wootten, Kijeong Yim, Llsang Yoon

In publications, please refer to this document as:
ALMA Partnership, 2017, S. Asayama, A. Biggs, I. de Gregorio, B. Dent, J. Di Francesco, E. Fomalont, A. Hales, J. Hibbard, G. Marconi, S. Kameno, B. Vila Vilaro, E. Villard, F. Stoehr
ISBN 978-3-923524-66-2

Contents

This page was intentionally left almost blank

Chapter 1

Introduction

The Atacama Large Millimeter/Submillimeter Array (ALMA) is an aperture synthesis telescope consisting of 66 antennas arranged in a series of different configurations. It operates over a broad range of observing frequencies in the millimeter and submillimeter regime. ALMA Early Science Operations started with Cycle 0 in September 2011 and the official inauguration took place in March 2013.

Cycle 5 operations will include standard and non-standard observing modes. Non-standard modes are observing modes for which the data cannot be processed by the pipeline and need to be reduced manually by ALMA staff. Up to 20% of the PI observing time in Cycle 5 will be allocated to proposals requesting non-standard modes. Users should refer to Appendix A (ALMA Capabilities) in the ALMA Cycle 5 Proposer's Guide, for the latest information and a description of standard and non-standard modes.

The Technical Handbook provides additional information on technical aspects and their limitations of the Cycle 5 setup for ALMA users. It should, however, not be necessary to use the Technical Handbook to prepare an ALMA proposal.

The handbook is divided into three main sections: The concepts of interferometry and the ALMA hardware components (Chapters 2–5), the observing concepts and software (Chapters 6–9), and finally the data quality and handling (Chapters 10–14). It also includes a number of appendices with expanded information on specific ALMA hardware components and calibration, and an acronym list.

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

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

Chapter 4 describes the details of the seven receivers offered in Cycle 5. The general technical specifications and a brief explanation on local oscillators and IF range is presented. Plots with the atmospheric transmission and the typical system temperatures per receiver are also included.

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

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

Chapter 7 describes several aspects of imaging to consider in ALMA observations. A short description on the different configurations proposed for Cycle 5 is included. Concepts of shadowing, beam shape, and spatial scale filtering are revisited. Mosaicing, and 12-m and 7-m Array data combination are also presented.

Chapter 8 describes the observing modes offered for Cycle 5 and how projects are observed. Single field interferometry, mosaics, single-dish observations, polarization, multiple region modes, solar, VLBI, and ephemeris observations are detailed in this section.

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

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

Chapter 11 describes the data quality assurance process. It also provides the criteria used for passing the quality assurance.

Chapter 12 describes the structure of the data.

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

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

The Technical Handbook concludes with two appendices, which contain supplemental material about the antenna design and transportation transporter (Appendix A), and the Local Oscillator (LO) and Intermediate Frequency (IF) system (Appendix B).

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

Chapter 2

Array Components

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

2.1 The ALMA Telescope

ALMA is composed of 66 high-precision antennas. Fifty of these antennas are 12 meter dishes in the 12-m Array, used for sensitive, high-resolution imaging. These fifty 12 m antennas are complemented by the Atacama Compact Array (ACA), also known as the Morita Array[1], composed of twelve closely spaced 7 m antennas (the 7-m Array), and four 12 m antennas for single-dish (or Total Power) observations (the TP Array), to enhance wide-field imaging of extended structures. In Cycle 5, ALMA will cover most of the wavelength range from 3.6 to 0.32 mm frequency), and when ALMA is completed the coverage will be from 10 to 0.32 mm (31-950 GHz).

The array is located on the Chajnantor plain of the Chilean Andes (lat.=–23.02917 deg., long.=–67.754649 deg.), a site that normally offers the exceptionally dry and clear sky conditions required to observe at millimeter and sub-millimeter wavelengths[2]. The ALMA antennas, weather stations, the two correlators and their computer interfaces, Local Oscillator generation hardware, timekeeping hardware, and the related array Real-Time Machine computer are all located at the 5000-meter altitude site referred to as the Array Operations Site (AOS). This site is connected via Gigabit fiber links to the Operation Support Facility (OSF), located at an altitude of 2900 meters, about 22 km from the AOS and 40 km from the town of San Pedro de Atacama. Science operations are conducted from the OSF and coordinated from the JAO Central office in Santiago. All three ALMA arrays are controlled via control software developed on the ALMA Common Software[3] (ACS).

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

The number of antennas in each array component (12-m, 7-m and TP Arrays), and the specific configurations available for an observing season (e.g. Cycle 5) will be published in the Capabilities section of the document *ALMA Proposer's Guide*. Complementary background information on ALMA and its capabilities for Cycle 5 can be found in the document *A Primer for Cycle 5*. Both documents can be found on the link http://www.almascience.org/documents-and-tools/.

[1] dedicated to the honour of K.-I. Morita
[2] http://www.almascience.org/about-alma/weather
[3] http://www.eso.org/projects/alma/develop/acs/

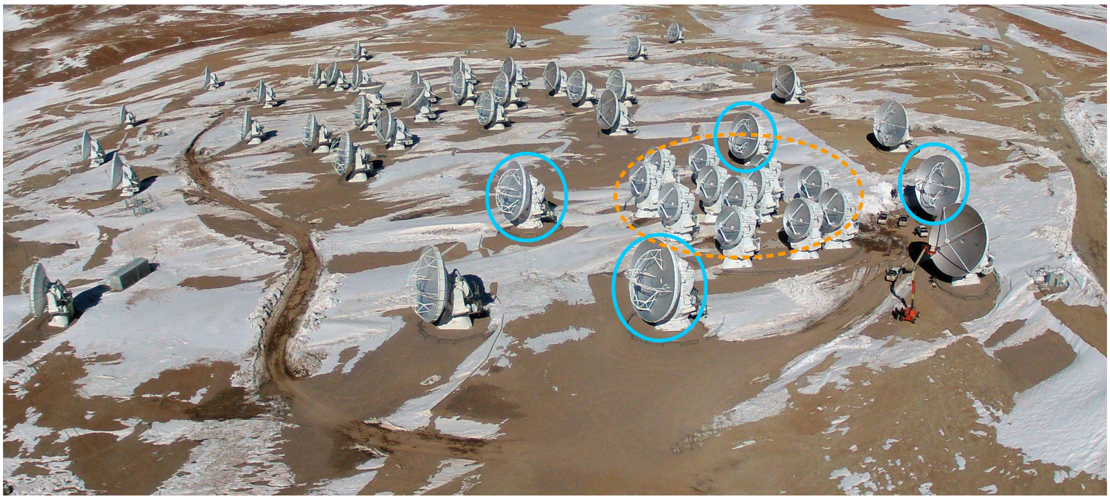

Figure 2.1: The ALMA 12-m Array in its compact configuration (left hand side of the image). The ACA with all 7 m antennas (dashed orange circles) and four single-dish 12 m antennas (blue circles) are distributed in the right hand side of the image. A few unoccupied stations can be seen, to which antennas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

2.2 The 12-m Array

The 12-m Array consists of fifty 12 m diameter antennas designed and built by the European and North American ALMA partners (each providing 25 units), according to the stringent ALMA Antenna Performance specifications (see Appendix A). Each antenna contains one front end, including a cryostat (see Appendix A.4), amplitude calibration device (ACD; A.5), water vapor radiometer (WVR; A.6), and back-end electronics (analog and digital racks). The WVRs are used to correct the phase fluctuations caused by water in the atmosphere along the line of sight of each 12-m Array element. The cryostat can contain up to ten cartridges, each covering one frequency band (see Chapter 4). Only one band observes at any time, but up to three can be switched on simultaneously, and rapid switching between them is possible. Each band receiver (Chapter 4) detects two orthogonal linear polarizations and down-converts the signals to an intermediate frequency with eight GHz of bandwidth per polarization. Bands 3-8 cartridges are dual sideband (2SB) and Bands 9 and 10 are double sideband (DSB).

The Local Oscillator (LO) signals (Appendix B) are transmitted to the antennas on optical fibers with a round trip measurement to correct for changes in the fiber length. There are four independent LO reference systems so that the 7-m Array, the TP array and two subsections of the 12-m Array (e.g., two ‘sub-arrays’) can conduct simultaneous independent observations. Please note that the sub-arrays feature of the 12-m Array is not yet a capability that is offered. The 8 GHz total IF bandwidth from the selected receiver is divided into four 2 GHz-wide basebands which are digitized at four Gsamples/s, with three-bit resolution, and transmitted on optical fibres. Total data rates are therefore 96 Gbits/s per antenna. With formatting, the bit rate is 120 Gbits/s.

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

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

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

2.3 The Atacama Compact Array

Using an interferometer to obtain images of extended or large-scale structures leads to the well known "zero spacing" problem. This problem arises from the constraint that, to avoid collisions, it is not possible to pack antennas closer than their diameter, leaving a hole in the distribution of baselines at short and zero baseline separations (corresponding to large angular structure). As a result, spatial information from baselines shorter than the closed-packing ratio is not recovered[4]. This problem has considerable impact on observations of extended objects, particularly those in which the emitted power is dominated by their large-scale structures.

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

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

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

2.3.1 The 7-m Array

The 7-m Array is composed of twelve 7 m diameter antennas designed and built by East Asia to the ALMA specifications (see Appendix A). Similar to 12-m Array antennas, each antenna contains one front end, including a cryostat, amplitude calibration devices, and one back end. Unlike the antennas of the 12-m Array, the 7 m antennas do not contain Water Vapor Radiometers (WVRs). The 7 m antenna cryostats are fitted with receivers nearly identical to those on the 12 m antennas, with small differences in the warm optics. The Local Oscillator signals transmitted to the 7 m antennas are originated identically to the ones sent to the 12-m Array, i.e., from inside the AOS building.

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

No baseline coverage from even the most compact configuration of the 12-m Array is obtained for spacings smaller than 15 m. The array configuration of the 7-m Array is designed to fill in missing spacings from about 9 m to $\sim$30 m (see Chapter 7). Each cartridge (see Chapter 4) receives two orthogonal polarizations.

[4]Strictly 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.

2.3.2 The TP Array

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

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

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

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

Chapter 3

Principles and Concepts of Interferometry

3.1 Introduction

Interferometry is the technique ALMA uses to obtain very high angular resolution observations of astronomical phenomena. In this Chapter, the principles and concepts behind interferometry are described, so that ALMA users can plan and understand their observations better. If more information is desired, the topic of interferometry is covered in more detail in the following seminal texts:

• *Interferometry and Aperture Synthesis in Radio Astronomy - Second Edition*, by Thompson, A. R., Moran, J. M., & Swenson, G. W. (Wiley-VCH)

• *Tools of Radio Astronomy - Fifth Edition*, by Wilson, T. L., Rohlfs, K., & Hüttemeister, S. (Springer)

• *Synthesis Imaging in Radio Astronomy II*, PASP Conference Series, Vol. 180, eds. G. B. Taylor, C. L. Carilli, & R. A. Perley (San Francisco - ASP)

• *Millimeter Interferometry: Proceedings of the IRAM Millimeter Interferometry School*, ed. A. Dutrey, available online at http://iram.fr/IRAMFR/IS/IS2002/archive.html.

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

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

The following expands upon these basic ideas. It begins by introducing the basic concepts of radio astronomy, and then moves to the principles of aperture synthesis.

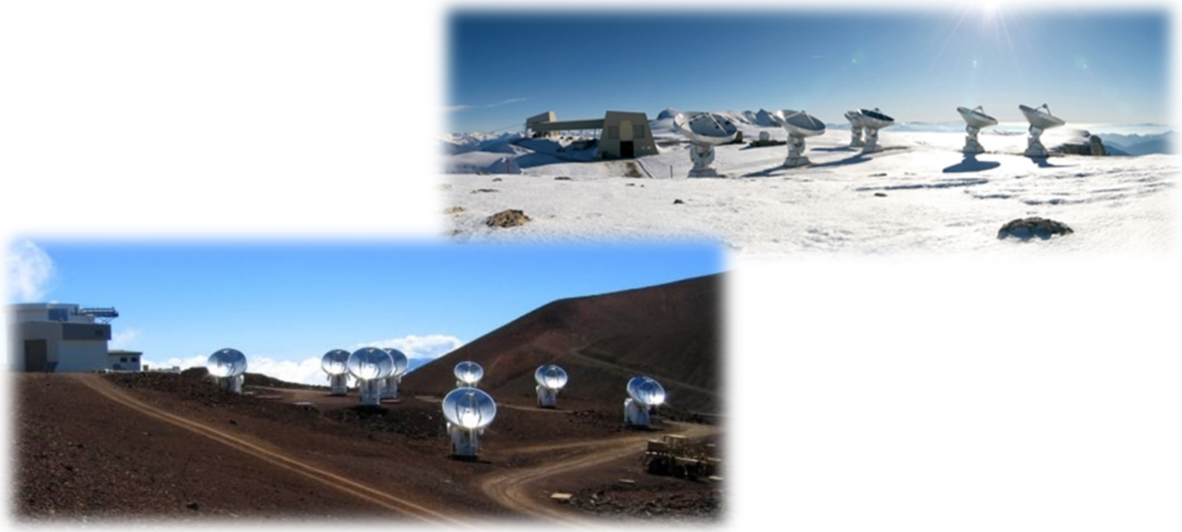

Figure 3.1: The Plateau de Bure Interferometer (top) and the Submillimeter Array (SMA) (bottom) are the precursors to the ALMA telescope; both are still in full operation and pioneered the science of millimeter wave interferometry.

3.2 Single-dish Response

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

$$\delta P = I_\nu \delta\Omega \delta A \delta\nu, \tag{3.1}$$

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

$$S_\nu = \int I_\nu \, d\Omega, \tag{3.2}$$

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

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

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

Antenna response is actually dependent on angle from the on-axis pointing direction of the antenna, due to diffraction. To demonstrate the angular dependence, Figure 3.2 shows in the top panel the case for EM power of wavelength λ arriving along the axis of an unobstructed antenna of diameter D. Since the source of the

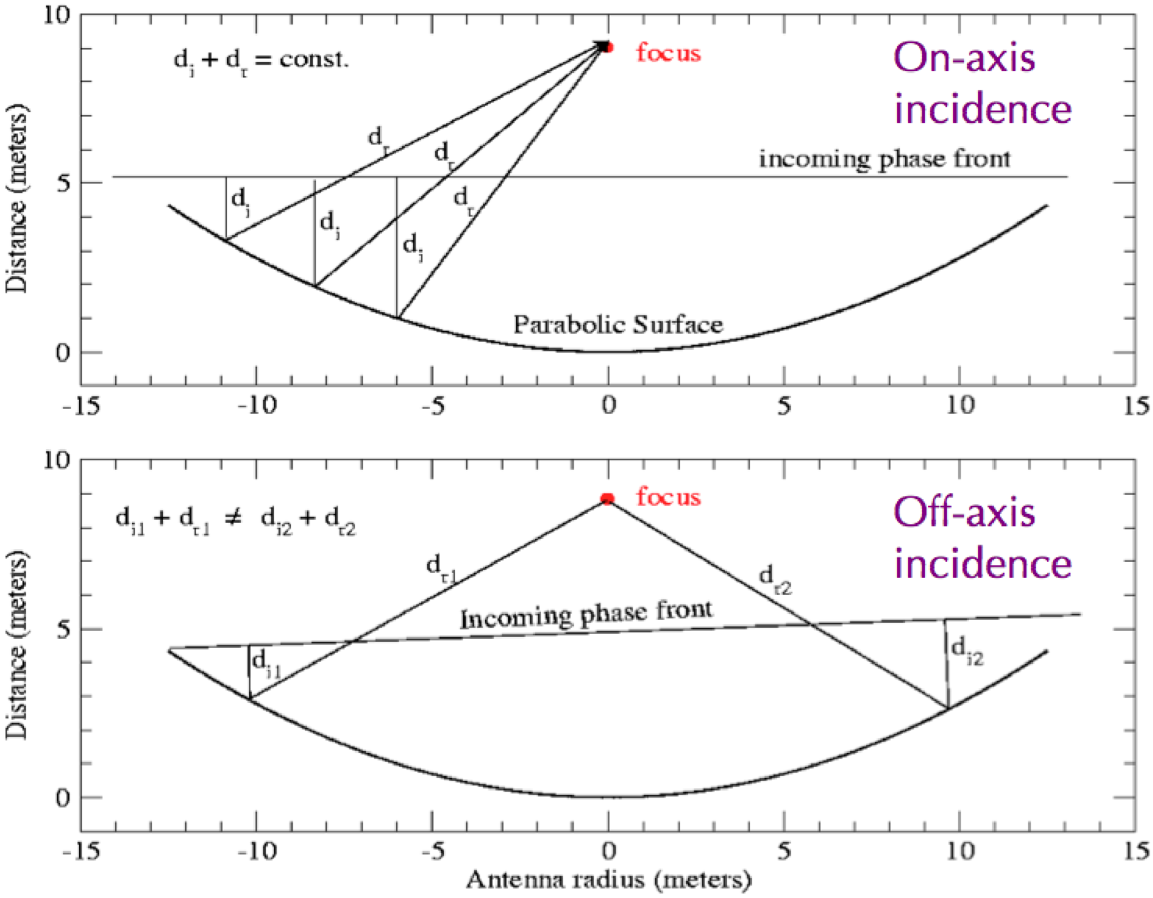

Figure 3.2: Schematic of an incoming plane-parallel wavefront reflecting off a antenna of diameter D and being brought to a focus. The top panel shows the case for a wavefront arriving on-axis. The bottom panel shows the case for a wavefront arriving off-axis. Note that the paths of incident EM power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

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

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

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

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

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

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

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

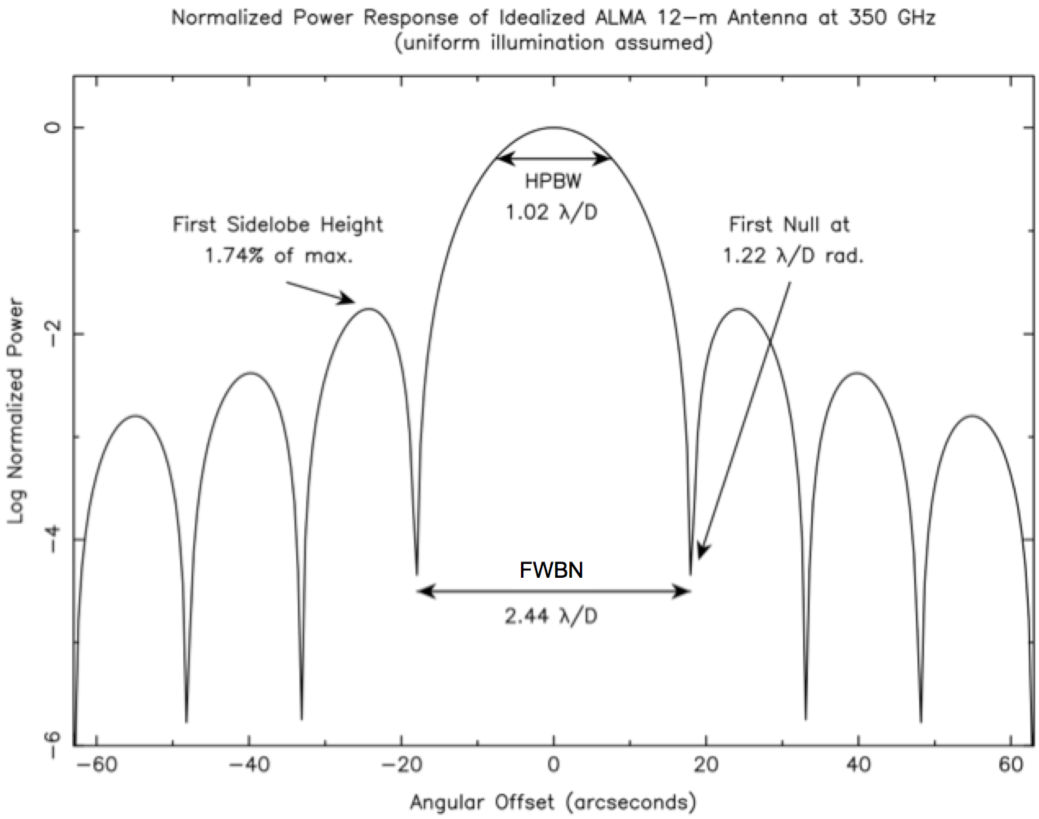

Figure 3.3: Normalized 1-D antenna power response for a 12 m antenna uniformly illuminated at 300 GHz. The power is in log units to emphasize the sidelobes. The HPBW of the primary beam is $\sim$1.02 λ/D and the FWBN is $\sim$2.44 λ/D. The angle of the first null, i.e., the resolution, is 0.5$\times$BWFN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

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

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

Defining θ and φ as orthogonal directional variables (e.g., sky coordinates), $I_\nu(\theta, \varphi)$ and $P_N(\theta, \varphi)$ can be defined as the directional functions of the sky brightness and the normalized antenna power, respectively. The total received power of an antenna at a given pointing is the integration over the sky of the product of the sky brightness distribution and the antenna power response:

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

In addition, the solid angle of the antenna power response $P_N(\theta, \varphi)$ can be found as:

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

3.3 Visibilities and Aperture Synthesis

For millimeter/radio astronomy, the angular resolution of a single-dish observation is very low compared to those found at optical wavelengths since λ is larger by many orders of magnitude. Though the D of millimeter/radio telescopes can also be much larger than those of optical wavelengths, the increase in D possible for single-dish telescopes is generally never enough to obtain the angular resolutions of ground-based optical telescopes, e.g., $1''$ or better. For example, the JCMT 15-m diameter antenna has an angular resolution at 850 μm of $\sim 14''$, and the Arecibo 300 m diameter antenna has an angular resolution at 21 cm of $\sim 3'$.

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

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

$$V(t) = \sum_i \Delta V_i(t) \tag{3.8}$$

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

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

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

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

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

From Equation 3.10, any measurement with a large filled-aperture telescope can be understood as being a sum in which each term depends on contributions from only two of the N elements. As long as the contributions from each element arrive at the focus in phase, *there is no need for the elements to be physically contiguous*. Generalizing, each cross-correlation term $\langle\Delta V_i \Delta V_k\rangle$ in Equation 3.10 can be measured with two smaller, physically separated antennas (at locations i and k) by measuring the average product of their output voltages with a correlating, i.e., multiplying, receiver. Moreover, if the source properties are unchanging, there is no need to measure all pairs at the same time. A given parabolic surface with N elements has $N(N-1)/2$ pairs of elements, and these could be observed sequentially to "synthesize" a measurement by a large filled-aperture telescope. Alternatively, numerous pairs of antennas, with each antenna considered an element, can be distributed to positions at distances much larger than it is possible to build a single filled-aperture telescope, and the signals received by these antennas can be combined in phase to approximate the resolving power of a single filled-aperture telescope.

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

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

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

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

Expanding to two dimensions, introduces β, a direction on the sky orthogonal to α. Also, $m = \sin\beta$ is defined as the small angle analog to l in this new direction, and $v = b\cos\varphi$ where φ is the angle of the position s_o on

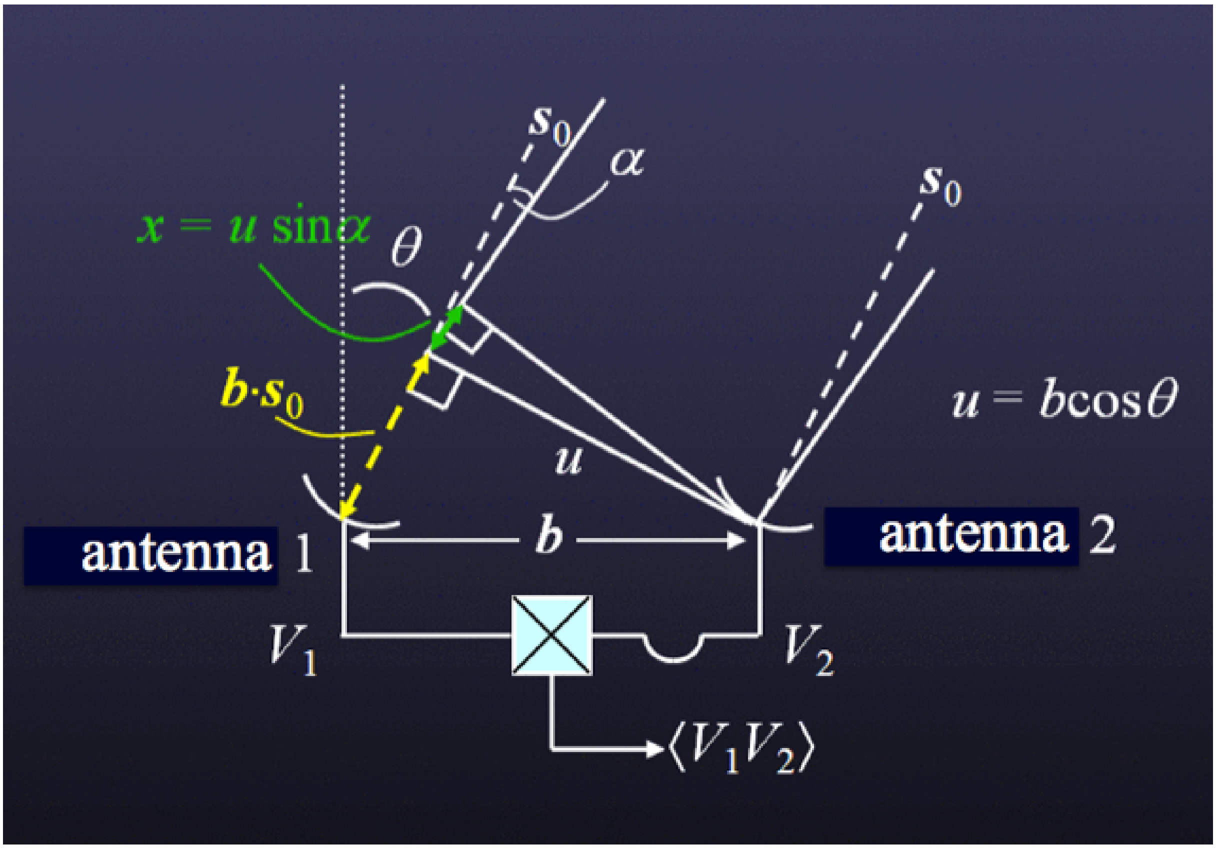

Figure 3.4: An ideal 1-D two-antenna interferometer consisting of two antennas, 1 and 2, separated by physical distance (i.e., a baseline) b. The antennas are both pointed towards a sky location given by s_o, which is at an angle θ from the meridian. The projected distance between the two antennas in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

the sky from the reference position orthogonal to θ. Finally, y is defined as the extra path length introduced in this new direction, in units of the wavelength of emission, i.e., $y = v\sin\beta = vm$. With these changes, the two-dimensional voltage response of antenna 2 is:

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

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

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

$$\langle V_1 V_2 \rangle = \langle \iint V_1(l,m) dl dm \iint V_2(l,m) dl dm \rangle \tag{3.13}$$

Under the assumption that signals emanating from different parts of the sky are incoherent, the time averages of the correlation of those signals will be zero. Thus, the product of the integrals in Equation 3.13 can be simplified to:

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

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

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

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

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

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

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

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

3.4 The Visibility or *uv*-Plane

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

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

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

By measuring the distribution of complex visibilities (in the visibility or *uv*-plane), in principle the sky brightness distribution can be recovered. In essence, an image is a "sum" (i.e., the Fourier transform) of the visibilities where each has an amplitude and phase representing the brightness and relative position of emission on a specific angular scale. The image and its Fourier transform are conjugates of each other, and each contains the same amount of information.

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

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

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

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

3.5 Fields-of-view and Mosaics

Note that an interferometer with idealized antennas was described. Each antenna of an actual interferometer, however, has finite diameter. As described earlier, such antennas have their own power response on the sky $P_N = \mathcal{A}(l, m)$ (e.g., the Airy function for an unobstructed, uniformly illuminated aperture). Indeed, the individual antenna response fundamentally limits the extent of an interferometric image made with a single pointing. In practice, the HPBW of the primary beam serves as the "field-of-view" of the single-pointing interferometric image. Moreover, $\mathcal{A}(l, m)$ is actually included formally in the correlator output:

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

Hence, an interferometer actually measures the Fourier transform of the sky brightness distribution multiplied by the antenna power response. To recover $I(l, m)$, the image resulting from the Fourier transform of the complex visibilities must be divided by $\mathcal{A}(l, m)$ as the last step of image processing. This so-called *primary beam correction*, however, is only necessary if the resulting image contains extended emission or emission located far from the phase center.

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

3.6 Spatial Filtering

Though the principles described in Sections 3.3-3.5 should enable the true sky brightness distribution to be recovered, it is impossible in practice to sample completely the *uv*-plane and obtain all visibilities. The incomplete *uv*-plane sampling effectively provides a fundamental limit to the level of detail discernible in the sky brightness distribution, i.e., down to a minimum scale defined as the resolution. In addition, incomplete sampling results in *spatial filtering* of the true sky brightness distribution, i.e., the resulting images do not contain information on angular scales unobserved by the interferometer. In particular, the lack of coverage at the shortest baselines (i.e., lower than those sampled by the smallest baselines) results in an intrinsic lack of sensitivity to large-scale emission. It is crucial for ALMA users to understand these limitations. Further examples can be found in Chapter 7 ("Imaging"), Sections 7.6 and 7.7.

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

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

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

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

Using the convolution theorem, Equation 3.24 can be rewritten as:

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

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

The measure of how similar an image is to the true sky distribution is sometimes referred to as *image fidelity*. Image fidelity depends on the specifics of coverage of the *uv*-plane sampled by the interferometer. Since the numbers of samples are necessarily finite and discrete, there are invariably gaps in any practical sampling of the *uv*-plane. These gaps mean that no information is obtained about the true sky brightness distribution on those specific angular scales. Note that visibilities corresponding to those unobserved scales can have any value. With no information, however, it is typically assumed that $\mathcal{V}(u,v) = 0$ at unsampled locations in the *uv*-plane. Including these visibility domain gaps through the Fourier transform produces aliased features in the resulting image, the magnitude of which depends on the extents and locations of gaps in the *uv*-plane and the brightness of emission on sampled scales. If the *uv*-plane has been reasonably well sampled, the synthesized beam will consist of a compact positive feature surrounded by positive and negative features of lower relative amplitude. These latter features, also called *sidelobes*, can complicate the image since brightness is distributed via the point spread function throughout the image. The resulting image can have significant artifacts depending on the sky brightness distribution and the sampling of the *uv*-plane. A dirty image, however, can be improved through deconvolution techniques to minimize the effect of incomplete spatial frequency sampling (e.g., CLEAN and its variants; see Chapter 7).

Though speaking generally about true sky brightness distributions so far, a special note should be made for the case of point sources. Obviously, a point source is a distribution of emission that is not extended relative to the resolution of the observation. In this case, the morphology of the dirty image will equal that of the dirty

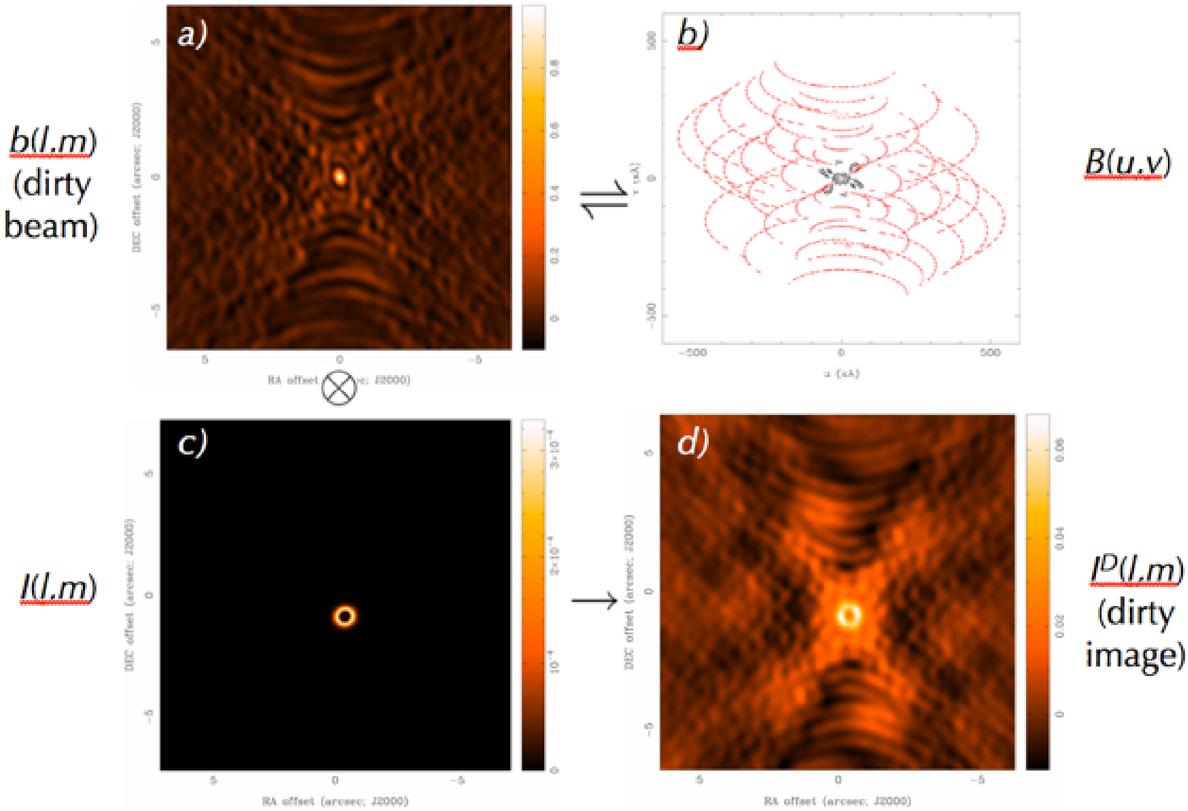

Figure 3.5: Imaging concepts. *Panel a (upper left)*: Example of a dirty beam, $b(l, m)$. *Panel b (upper right)*: The related ensemble of discrete points sampled in the uv-plane, $B(u, v)$. The black points were obtained from a compact configuration while the red ones were obtained from an extended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

beam (e.g., see Equation 3.25). Moreover, the complex visibilities of the point source have the same amplitudes on all observed angular scales. Of course, sources may appear point-like at low resolutions but may appear extended in higher-resolution observations.

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

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

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

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

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

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

where L_{min} is the minimum baseline in the array configuration. (Strictly speaking, for an input Gaussian visibility distribution of FWHM θ_{MRS}, the ratio of the brightness at source center of an image made by an array with L_{min} to the same made with no central hole in visibility sampling is $1/e$; see Wilner & Welch 1994.) The smallest baseline possible in an array occurs when two antennas are adjacent to each other. Of course, the antennas cannot be moved physically closer together than their diameters. Note that in projection antennas can appear closer than their diameters. In those cases, found typically when observing low elevation sources in compact configurations, the antenna in front blocks the reception by the antenna in the rear. The resulting visibilities are distorted, e.g., the antenna in the rear receives less power than the one in front. This situation is called *shadowing* and affected data are typically removed from the ensemble of observed visibilities.

Figure 3.6 illustrates the idea of spatial filtering using simulations of actual ALMA configurations. Here, panel *a* shows an optical image of the galaxy M51 used as an example of a true sky brightness distribution, after changing the frequency of the image to 100 GHz and placing the image center at δ = -40°, a declination easily observable with ALMA. Panels *b*, *c*, and *d* show the recovered images obtained by simulating 32-antenna observations of the galaxy using the CASA task *simobserve* only in very extended, moderately extended, and compact configurations, respectively, and CLEANing. Angular resolutions of 0.55″, 1.1″, and 3.6″ are obtained, respectively, and the corresponding maximum recoverable scales are 6.6″, 14.4″, and 36.4″, respectively. Larger scale emission from the galaxy has been filtered out in the very extended configuration observations (panel *b*), leaving only the compact structures of its arms. On the other hand, these compact structures are not very discernible in the compact configuration observations (panel *d*). A reasonable compromise is found in the moderately extended configuration observations (panel *c*), yet some small-scale detail and larger-scale emission remain missing. Note that though these images are missing angular scales, good science can still be obtained with them, as long as their limitations are properly understood. Combining data obtained from multiple configurations, or having more antennas, would increase the fidelity of recovered images.

3.7 Multi-configuration Observations

As previously discussed, a given configuration with baselines ranging from L_{min} to L_{max} is sensitive to angular scales from $\sim\theta_{MRS}$ to θ_{res}. Sensitivity to a broader range of angular scales is possible by combining data

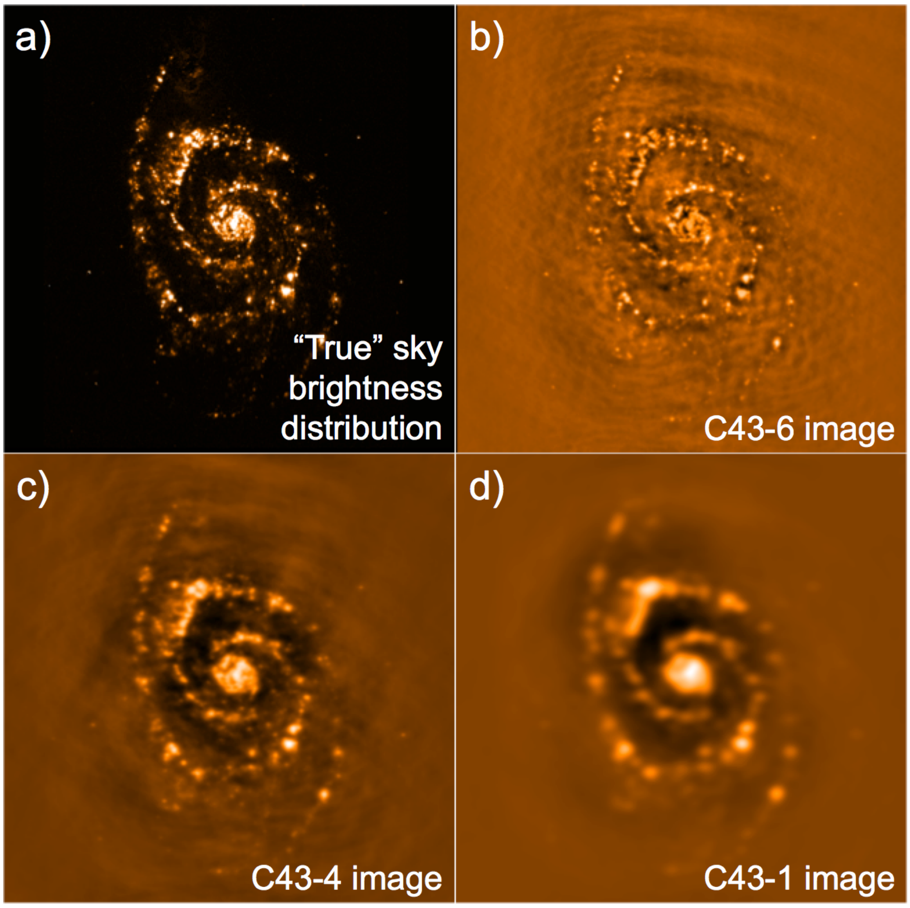

Figure 3.6: Examples of spatial filtering using the CASA task *simobserve* and actual ALMA configurations from Cycle 1. *Panel a (upper left)*: An optical image of the galaxy M51 used as a template for a true sky brightness distribution for the simulations. The frequency of the emission has been changed to 100 GHz, the image size has been scaled to $\sim3' \times 3'$, and its declination has been changed to -40° to allow ALMA observations to be simulated. For the simulations, the galaxy was “observed" over a mosaic of 39 pointings, for ∼10 hours in total. The resulting dirty images were CLEANed. *Panel b (upper right)*: The high-resolution image of the galaxy obtained when observed in an ALMA (Cycle 1) configuration with minimum and maximum baselines of 40.6 m and 1091.0 m, respectively (C32-6). The resulting synthesized beam is $\sim0.55''$ and the maximum recoverable scale is $6.6''$. *Panel c (lower left)*: Medium-resolution image of the galaxy when observed in an ALMA configuration with minimum and maximum baselines of 20.6 m and 558.2 m, respectively (C32-4). The resulting synthesized beam is $\sim1.1''$ and the maximum recoverable scale is $14.4''$. *Panel d (lower right)*: Low-resolution image of the galaxy when observed in an ALMA configuration with minimum and maximum baselines of 14.2 m and 81.4 m, respectively (C32-1). The resulting synthesized beam is $\sim3.8''$ and the maximum recoverable scale is $36.4''$.

obtained in multiple configurations, where more extensive coverage of the uv-plane is attained. For example, the same source can be observed in different configurations of the 12-m Array, with more extended configurations providing higher angular resolutions. For sensitivity to extended structures, the more compact 12-m Array configurations or the more tightly clustered 7-m Array can be used. Finally, the individual Total Power (TP) Array antennas of the ACA can be used to map the largest angular scales and address the zero-spacing problem. During proposal preparation, ALMA users should take note of the maximum recoverable scale needed to ensure that the proposed observations will be able to recover the scales needed to address the science in question. The ALMA Observing Tool will determine which combination of configurations will yield the desired angular resolution and maximum recoverable scale.

Data combination appears to work best when the signal-to-noise ratios (SNR) of the datasets are similar. Otherwise, information on scales covered by lower SNR data is relatively less reliable, making interpretation of the images difficult. In addition, the accuracies of the astrometry and calibration of the different datasets are crucial. Assuming the SNR of the individual datasets is high, combination is best done in the visibility domain rather than the image domain to minimize the effect of artifacts produced by aliasing, i.e., incomplete uv-coverage, in either dataset. For example, interferometer data obtained from different configurations should be combined in the visibility domain and then the new ensemble should be Fourier transformed to produce a new dirty image.

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

3.8 Units and Conversions

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

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

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

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

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

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

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

Note again that flux densities observed by ALMA are typically in units of Janskys, where 1 Jy = 10^{-26} W m^{-2} Hz^{-1} = 10^{-23} erg s^{-1} cm^{-2} Hz^{-1}. The ALMA Observing Tool converts between temperatures and flux densities using these formulae.

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

References

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

Chapter 4

Receivers

The ALMA front end can accommodate up to 10 receiver bands covering most of the wavelength range from 10 to 0.3 mm (30–950 GHz). Each receiver band is designed to cover a tuning range which is approximately tailored to the atmospheric transmission windows. These windows and the tuning ranges are outlined in Figure 4.1. In Cycle 5, Bands 3, 4, 5, 6, 7, 8, 9, and 10 are available (see available frequency and wavelength ranges for these bands in Table 4.1). The receivers are described in more detail in the following sections as well as in the references listed in Table 4.2.

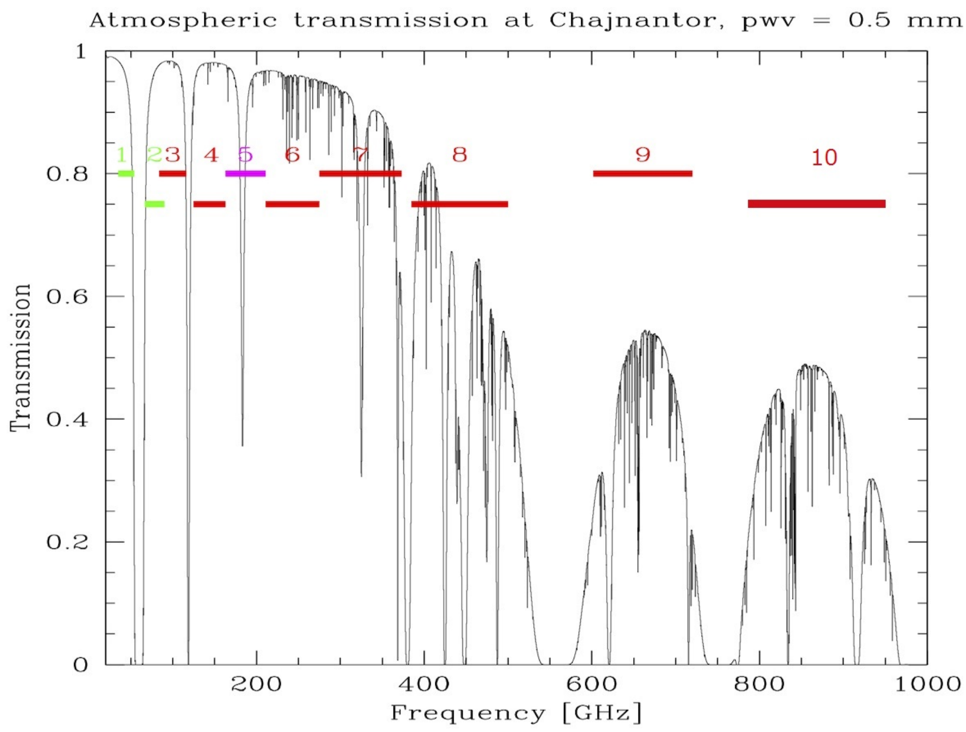

Figure 4.1: The ten ALMA receiver bands. Receiver bands for Cycle 5 are shown in red or purple superimposed on a zenith atmospheric transparency plot at the Array Operation Site (AOS) for 0.5 mm of Precipitable Water Vapor (PWV). Band 5 (shown in purple) is newly available in Cycle 5.

The ALMA receivers in each antenna are situated in a single front-end assembly (see Appendix A, Section A.4). The front-end assembly consists of a large cryostat containing the receiver cold cartridge assemblies (including Superconductor-Insulator-Superconductor (SIS) mixers and Local Oscillator (LO) injections) and the Intermediate Frequency (IF) and LO room-temperature electronics of each band (the warm cartridge assembly,

Band	Frequency/ Wavelength range (GHz)[1]/(mm)	LO range (GHz)	Sideband mode[2]	IF range (GHz)	Inst. IF bandw. (GHz)[4]	T_{rx} over 80% of band (K)[6]	T_{rx} at any frq. (K)[6]
3	84.0-116.0/ 2.59-3.57	92 - 108	2SB	4–8	7.5	<41[7]	<45[7]
4	125.0-163.0/ 1.84 - 2.40	133 - 155	2SB	4–8	7.5	<51	<82
5	163.0-211.0/ 1.84 - 2.40	171 - 203	2SB	4–8	7.5	<65	<105
6	211.0-275.0/ 1.09-1.42	221 - 265	2SB	5–10[3]	7.5	<83	<136
7	275.0-373.0/ 0.80-1.09	283 - 365	2SB	4–8	7.5	<147	<219
8	385.0-500.0/ 0.60-0.78	393 - 492	2SB	4–8	7.5	<196	<292
9	602.0-720.0/ 0.42-0.50	610 - 712	DSB	4–12	7.5(15)[5]	<175 (DSB)	<261 (DSB)
10	787.0-950.0/ 0.32-0.38	795 - 942	DSB	4–12	7.5(15)[5]	<230[8] (DSB)	<344 (DSB)

Table 4.1: Receiver Specifications. *Notes to Table:* **1.** Frequency range is the maximum available, at the extreme upper and lower limits of the IF passband. For FDM mode, the coverage is a bit smaller (See Section 6.4) **2.** Sideband modes: SSB means single sideband receiver, 2SB means dual sideband receiver where the two sidebands are available simultaneously, DSB means double sideband receiver. See text for details. **3.** Usable IF range is extended to allow simultaneous observations of multiple lines. However, the autocorrelation noise performance is degraded by a factor of up to about 1.5 below 5.5 GHz (Section 4.2.4) **4.** Maximum instantaneous IF bandwidth, limited by the back-end filters and spectrometers. As both upper and lower sidebands both pass through the same IF bandwidth but are subsequently separated, the effective signal bandwidth given in this column for 2SB receivers is twice the actual IF filter bandwidth. In addition, this is per polarization, so the total effective bandwidth for each receiver is then another factor of 2 higher. Note that the effects of the anti-aliasing filters have been included (see Section 6.4). **5.** In Cycle 5, the maximum bandwidth will be approximately double in cross-correlation mode, because both sidebands can be separated and correlated using 90-degree phase switching (see Section 6.3.4 and Section B.4.4). **6.** 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 typical values measured, which in many cases are better than specifications. **7.** The specification for Band 3 receivers is T_{rx}<41 K at LO=104 GHz, and T_{rx}<45 K for any other valid LO setting. Both values should be the average over all four IFs and 4 GHz bandwidth. **8.** The specification for Band 10 receivers is T_{rx}<230 K within a selected 80 % portion of that band (787-950 GHz).

Topic	Author/Year	Technical papers or Meeting proceedings	ADS identifier
B3	Claude et al. 2008	SPIE 7020	2008SPIE.7020E..33C
B4	Asayama et al. 2014	PASJ, 66 (3), 57(1-13)	2014PASJ...66...57A
B5	Bhushan et al. 2010	21th Intl Symp Space Terahertz Tech	2010stt..conf..137B
B6	G. A. Ediss, et. al 2004	15th Intl Symp Space Terahertz Tech	2004stt..conf..181E
B7	Mahieu et al. 2012	Trans. THz Sci. and Tech., 2(1) 29-39	2012ITTST...2...29M
B8	Sekimoto et al. 2008	19th Intl Symp Space Terahertz Tech	2008stt..conf..253S
B9	Baryshev et al. 2007	19th Intl Symp Space Terahertz Tech	2008stt..conf..258B
B10	Uzawa et al. 2009	20th Intl Symp Space Terahertz Tech	2009stt..conf...12U
Optics	Rudolf et al. 2007	IEEE Trans. on Antennas & Propagation	2007ITAP...55.2966R
WVR	Emrich et al. 2009	20th Intl Symp Space Terahertz Tech	2009stt..conf..174E

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

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 front end (see Appendix A.5). Each receiver cartridge contains two complete receiving systems sensitive to orthogonal linear polarizations. The designs of the mixers, optics, LO injection scheme, and polarization splitting vary from band to band, depending on the optimum technology available at the different frequencies; each receiver is described in more detail in the sections below.

Up to three bands can be switched on at a time (more would risk overloading the cryostat cooler). From a hardware point of view, it takes only about 1-2 seconds to switch between two bands which are powered up (limited by the LO locking and antenna movement). For bands that are not switched on, the time to fully thermally stabilize a receiver from an off state is up to ∼15 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. It means that only one receiver can be used at a given time.

4.1 Local Oscillators and IF Ranges

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

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

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

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

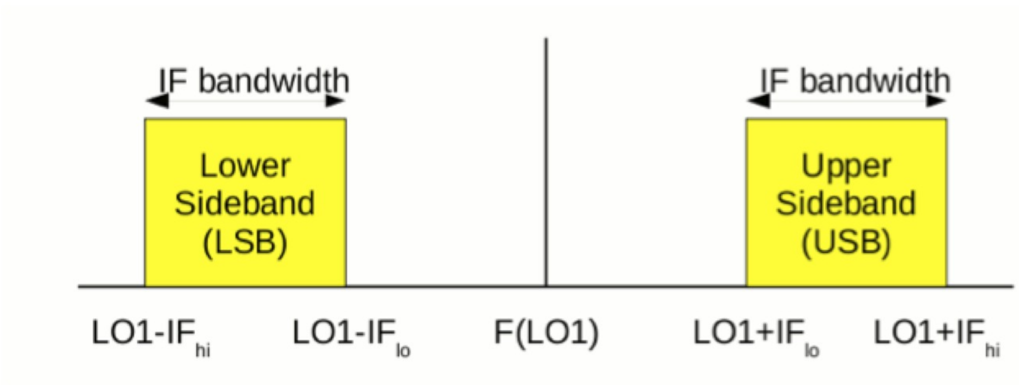

Figure 4.2: IF ranges for the two sidebands in a heterodyne receiver.

4.2 The Cycle 5 Receivers

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

Bands 9 and 10 use double sideband (DSB) receivers, where the IF contains noise and signals from both sidebands. They only have two outputs, one per polarization. The IF effective bandwidth is 7.5 GHz per polarization (after passing through the IF processing units), so the total instantaneous receiver bandwidth is the same as Bands 3 - 8. In Cycle 5, however, both sidebands can be simultaneously correlated using 90-degree phase switching (see Chapter 6), although this does not remove the noise from the opposite sideband. The effective system bandwidth available in Bands 9 and 10 is therefore double that of the 2SB receivers.

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

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

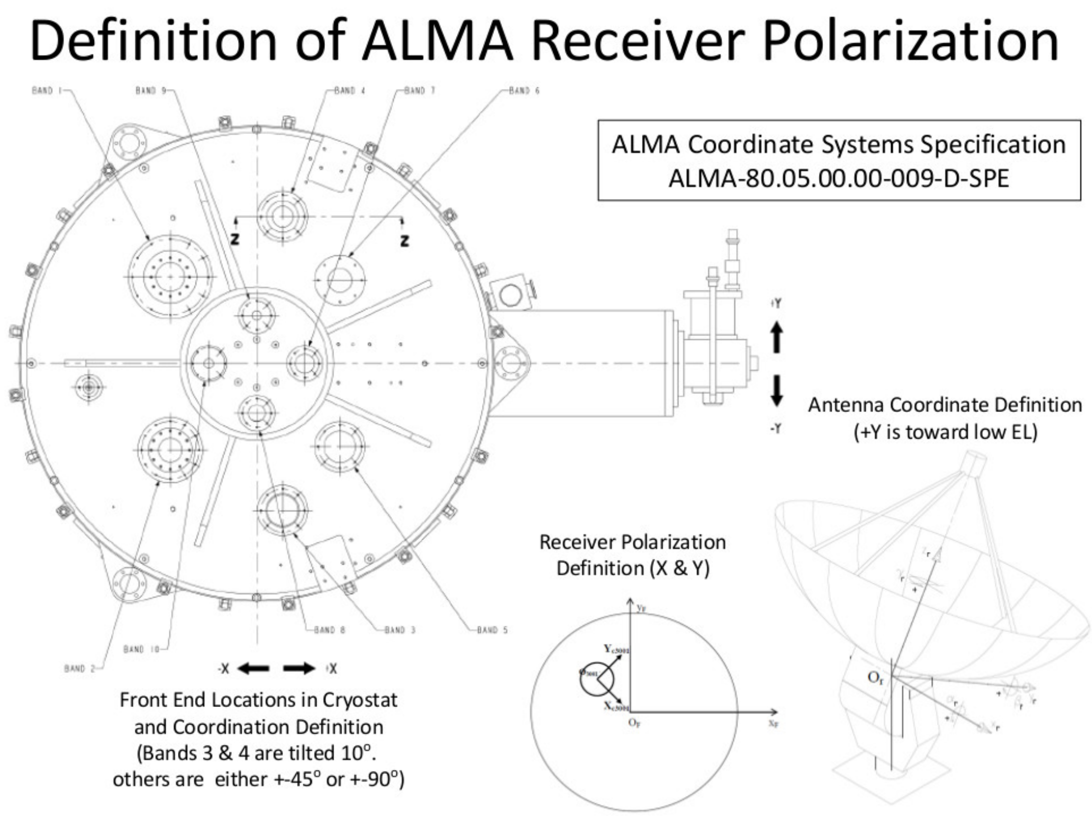

Figure 4.3: The layout of the ALMA receiver bands within the ALMA cryostat. Also shown is the orientation of the receiver polarization position angles.

4.2.1 Band 3 Receiver

Band 3 is the lowest frequency band available in Cycle 5, covering a frequency range of 84.0–116.0 GHz (in the 3 mm atmospheric window). The cartridge is fed by a "periscope" pair of ellipsoidal pickoff mirrors located outside the cryostat, which refocus the beam through the cryostat window, allowing for a smaller window diameter (Figure 4.4). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the incoming signal into two linear orthogonal polarizations and feeds the SIS mixers.

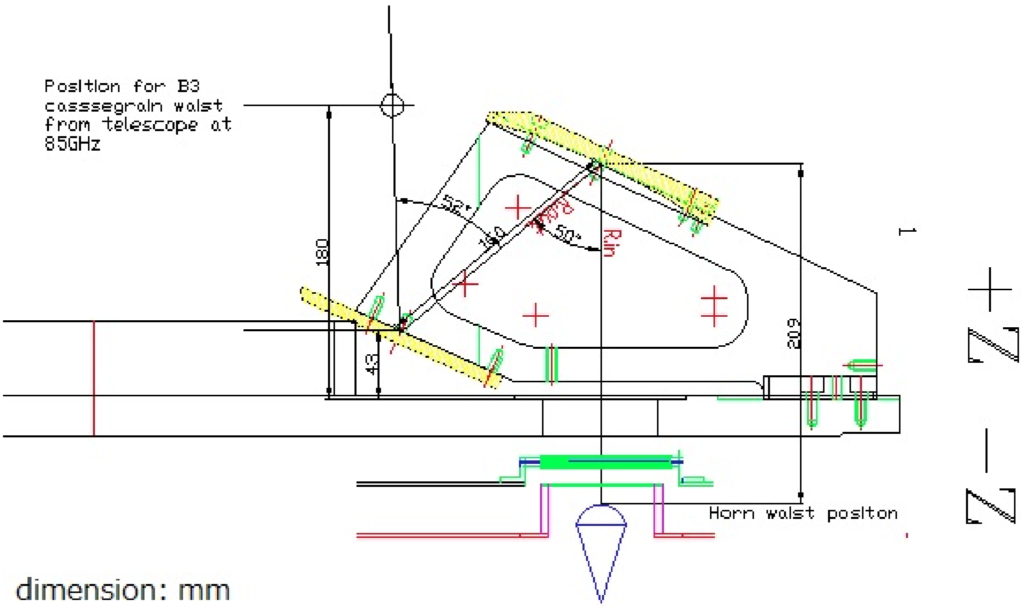

Figure 4.4: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the solid line, and the Cassegrain focus is shown by the small circle to the upper right.

A block diagram of the Band 3 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.5. The Cold Cartridge Assembly (CCA) contains the corrugated feedhorn, 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 acceptance criteria for the Band 3 receiver noise performance ($T_{\rm rx}$) is <41 K at LO1=105 GHz, and <45 K for any other valid LO setting (the OT assumes 45 K). The achieved performance is typically somewhat better than this, with $T_{\rm rx}$ ~37 K. The atmospheric transmission over most of Band 3 is very high, even with a large PWV (Figure 4.6) which means observations in Band 3 can (and do) take place with up to 10 mm or more of PWV. The system temperature ($T_{\rm sys}$) shows the expected rise at the higher end, due to an atmospheric oxygen line (Figure 4.7), so that the increase is mostly independent of the PWV.

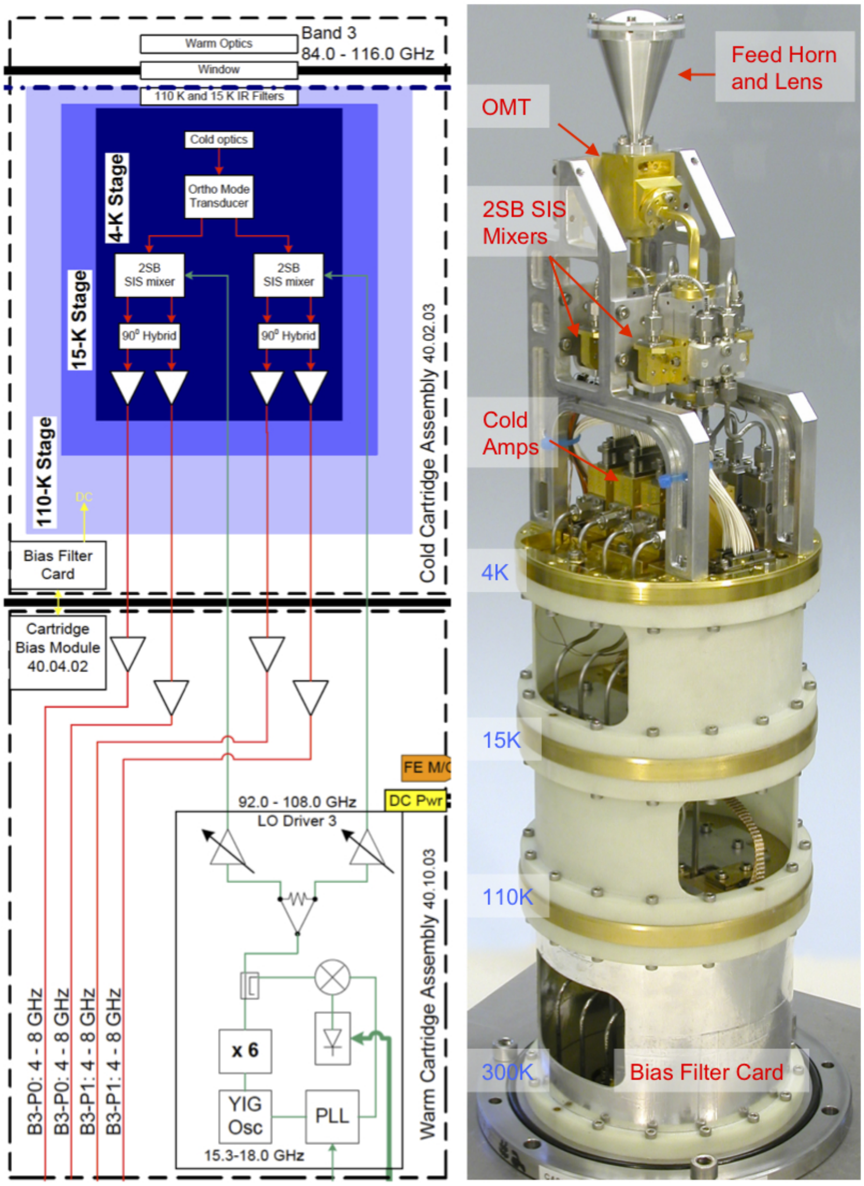

Figure 4.5: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges were constructed in Canada at NRC-HIA, Victoria.

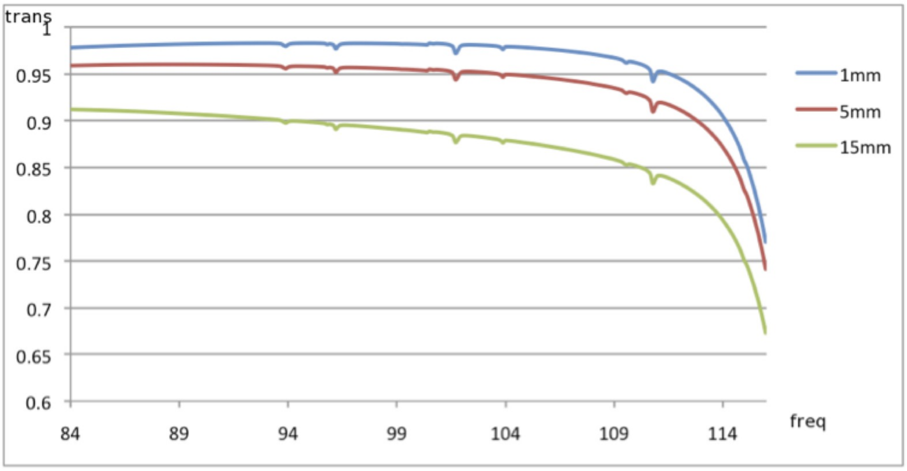

Figure 4.6: Band 3 zenith transmission for 1, 5 and 15 mm of PWV. Frequency is in GHz.

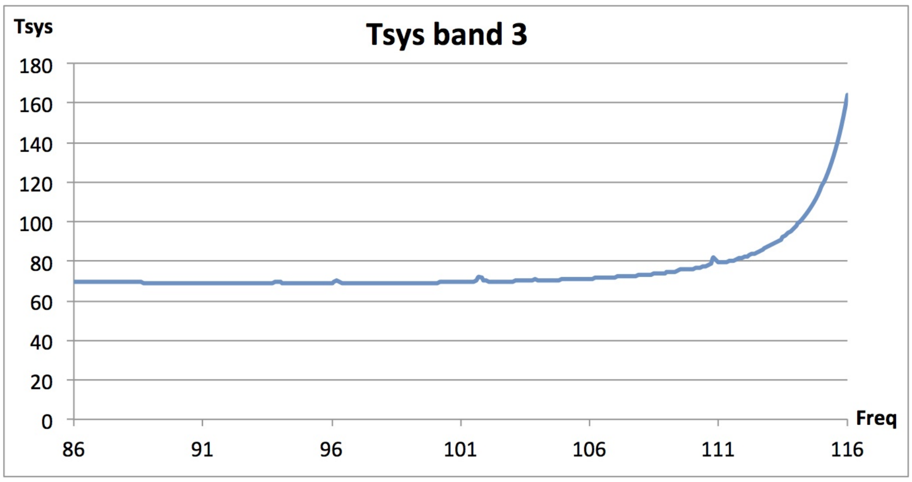

Figure 4.7: Typical system temperature (T_{sys}) at zenith for Band 3 with 1.262 mm of PWV. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.2 Band 4 Receiver

The Band 4 receiver covers the 125 to 163 GHz spectral window (in the 2 mm atmospheric window). The signal collected by the telescope is focused to the Band 4 cartridge using a set of warm mirrors (Figure 4.8). A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the 2SB SIS mixers.

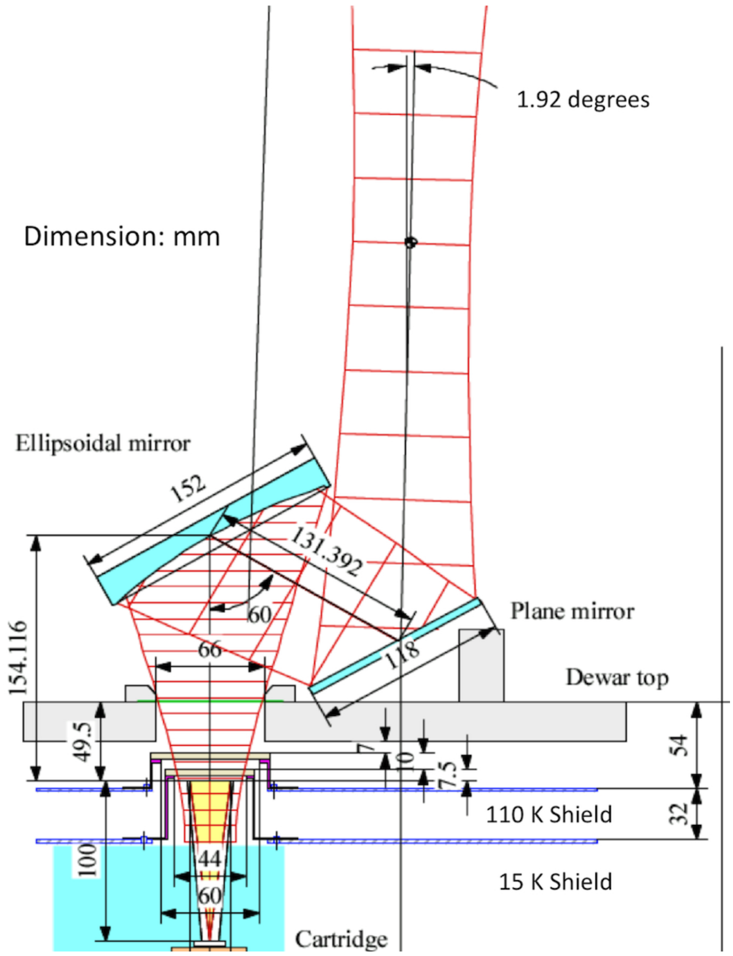

Figure 4.8: Optical layout of the Band 4. Red indicates Band 4 Gaussian beam of 5 times the beam width.

A block diagram of the Band 4 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.9. The Band 4 CCA contains a feed horn, an OMT as a polarization splitter, 2SB SIS mixer assemblies, cold IF amplifiers, isolators, and LO frequency doublers. The RF signal is down converted to 4–8 GHz using a 2SB mixer unit.

The atmospheric transmission in Band 4 is shown in Figure 4.10 for three typical PWV values. Most observations in Band 4 will be done with PWV $<$ 5 mm. The specification for Band 4 receiver noise performance ($T_{\rm rx}$) is $<$51 K over 80% of the band, and $<$82 K over the whole band (SSB). The OT assumes 51 K. The resulting system temperatures ($T_{\rm sys}$) for 1.262 mm PWV are shown in Figure 4.11. The on-array performance of the receivers is often more typically $\sim$40 K over most of the band, so the achieved $T_{\rm sys}$ may be slightly lower.

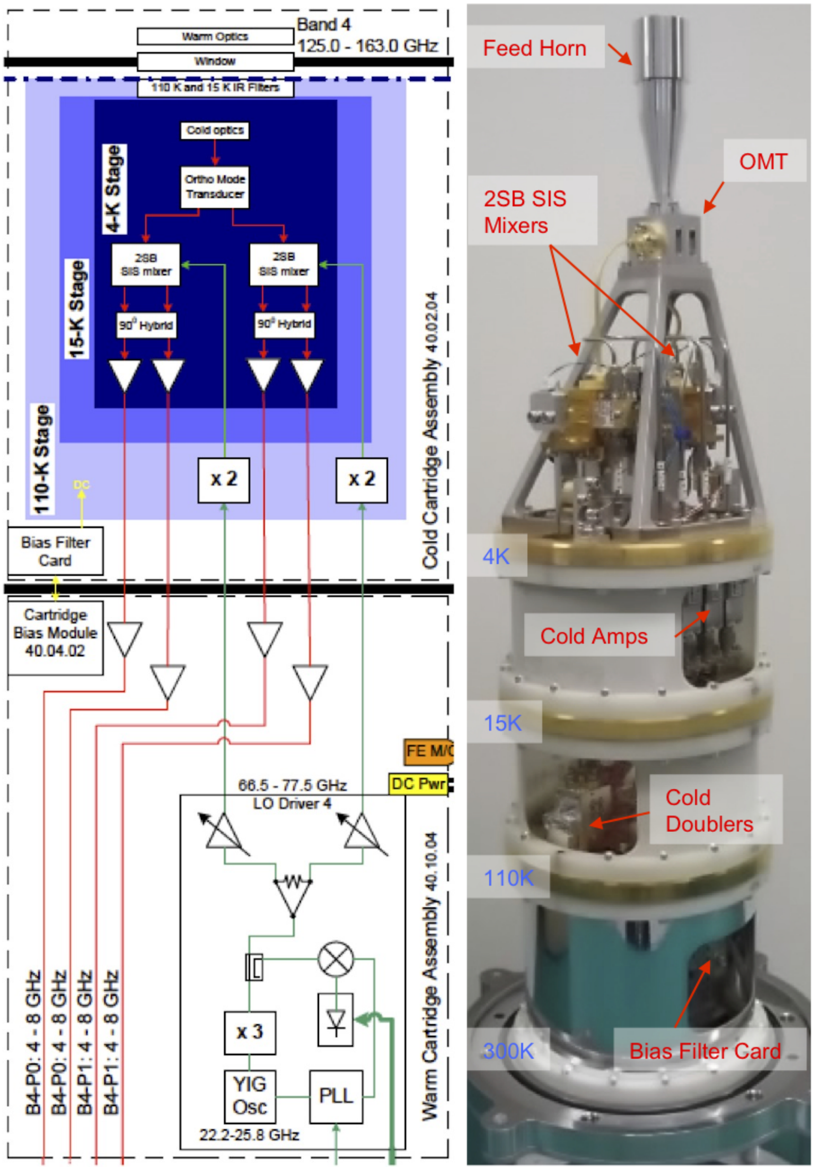

Figure 4.9: Block diagram of the Band 4 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 4 CCA. A single feedhorn feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The LO is generated in the WCA, with the final x2 multiplier on the 110 K stage.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

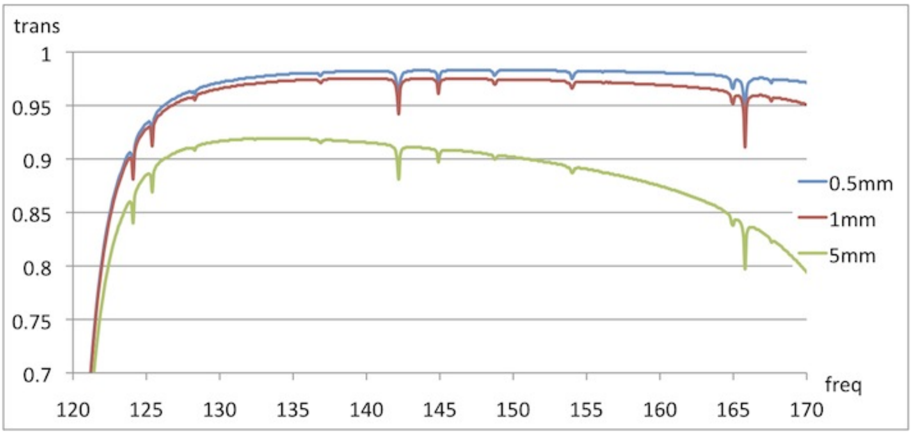

Figure 4.10: Band 4 zenith transmission for 0.5, 1 and 5 mm of PWV. Frequency is in GHz.

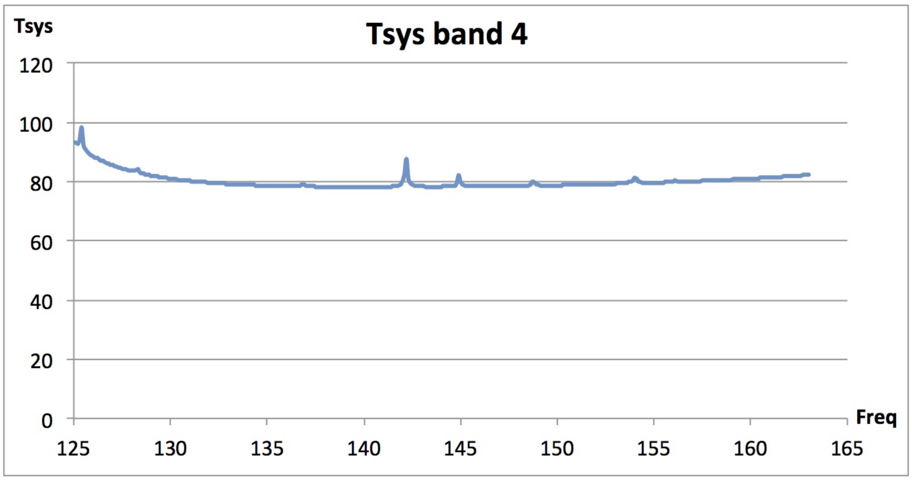

Figure 4.11: System temperature ($T_{\rm sys}$) at zenith for Band 4 with 1.262 mm of PWV. $T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.3 Band 5 Receiver

The Band 5 receiver covers the 163 to 211 GHz spectral window (in the 1.42-1.84 mm atmospheric window). The signal collected by the telescope is focused to the Band 5 cartridge using a set of cold mirrors (Figure 4.12). A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the 2SB SIS mixers.

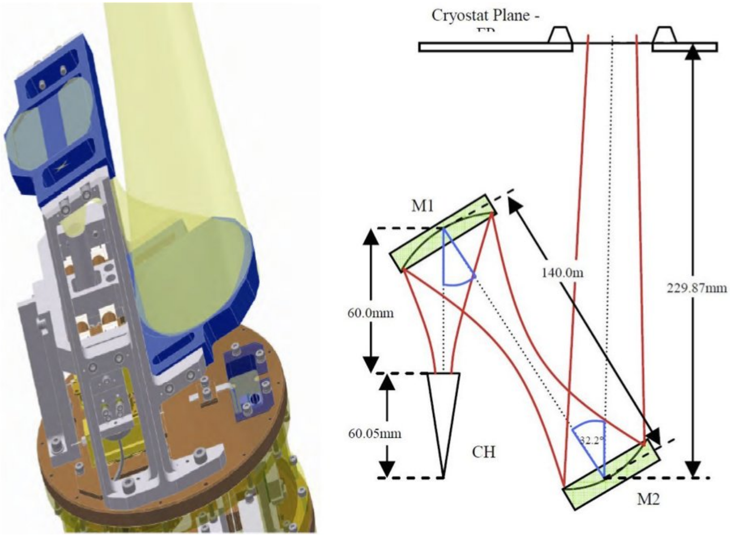

Figure 4.12: Optical layout of the Band 5. Placing the mixer assembly along the side of M 1 and streamlined layout of the IF system are the key to fulfil the requirements on $5w_o$ optical beam clearance.

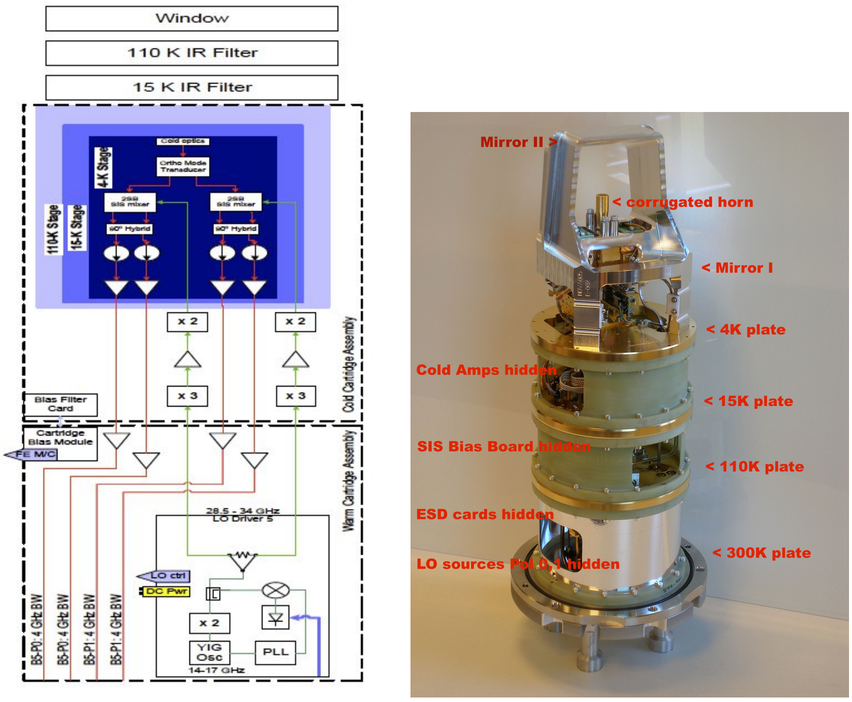

Figure 4.13: Block diagram of the Band 5 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 5 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The active multiplier (x3, a power amplifier, then x2) is mounted between the 300 K and 110 K stage.

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

The atmospheric transmission in Band 5 is shown in Figure 4.14 for three typical PWV values. The specification for Band 5 receiver noise performance ($T_{\rm rx}$) is <65 K over 80% of the band, and <105 K over the whole band (SSB $T_{\rm rx}$). The OT assumes 55 K. However, the performance of the receiver is considerably better than 50 K over the band. The resulting system temperatures ($T_{\rm sys}$) are shown in Figure 4.15.

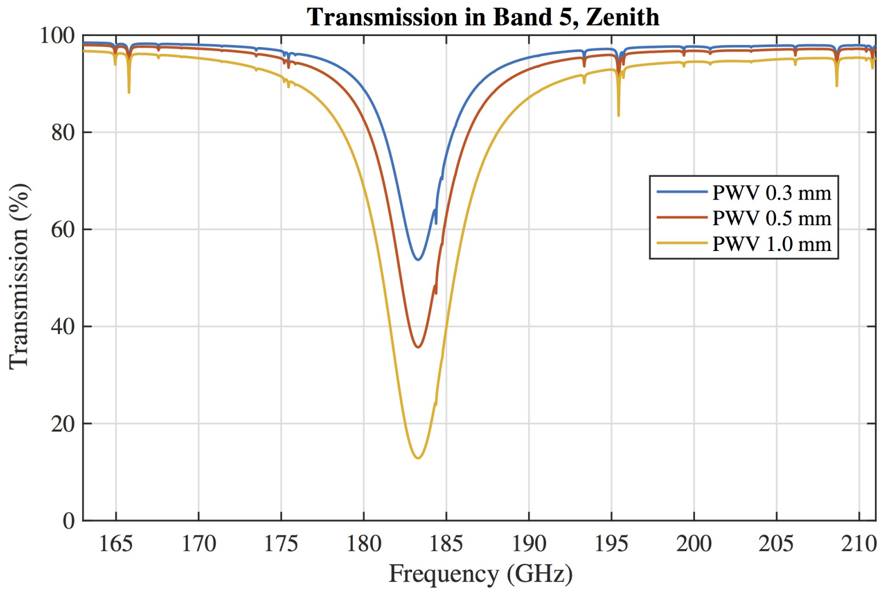

Figure 4.14: Band 5 zenith transmission for 0.5, 1 and 5 mm of PWV. Frequency is in GHz.

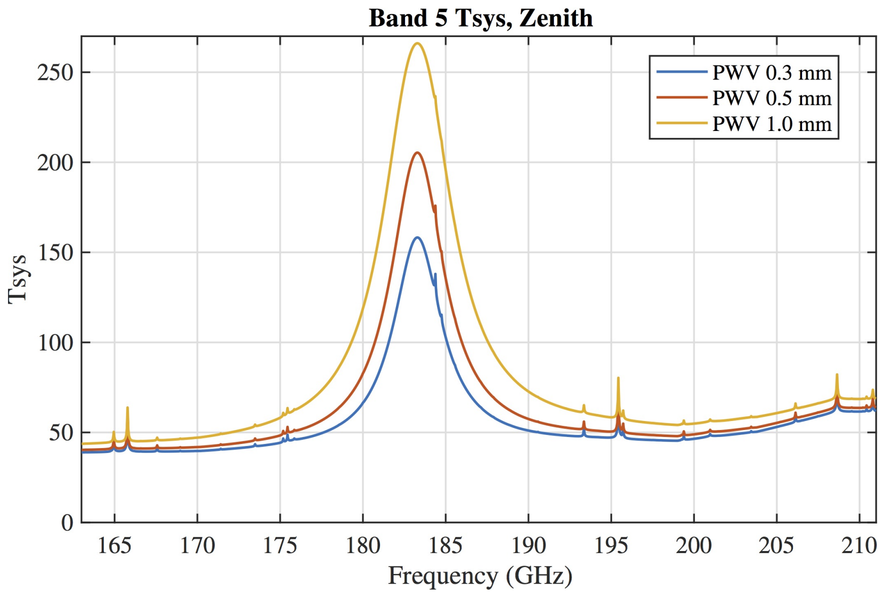

Figure 4.15: Typical system temperature ($T_{\rm sys}$) at zenith for Band 5 for three typical PWV values. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.)

4.2.4 Band 6 Receiver

The Band 6 receiver covers a frequency range of 211.0–275.0 GHz (the 1.3 mm atmospheric window). This receiver has a window with a pair of off-axis ellipsoidal mirrors inside the cryostat (Figure 4.16). A single feedhorn feeds an OMT which splits the two linear polarizations and feeds the SIS mixers. A block diagram of the Band 6 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.17.

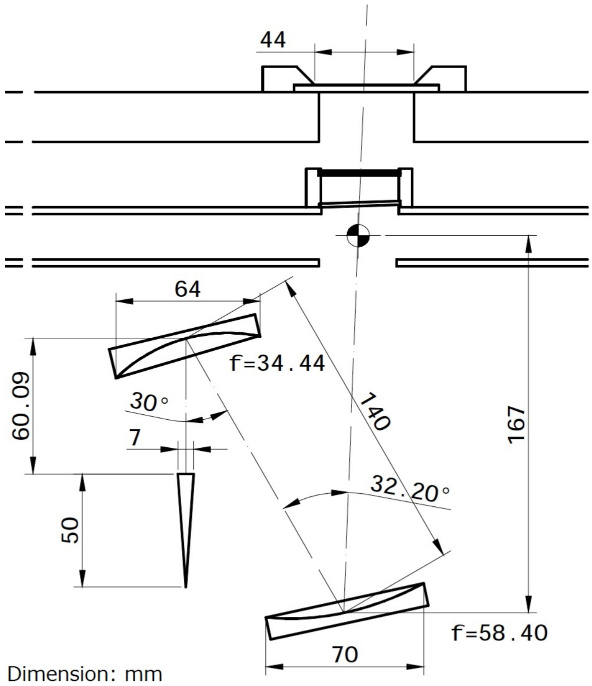

Figure 4.16: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocking filter. Note the slightly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

The Band 6 IF frequency has been chosen to allow for multiple simultaneous line observations[1]; it covers the range 5.0–10.0 GHz. There is ∼10-25% excess receiver noise below 5.5 GHz due to LO1 (although the increase of $T_{\rm sys}$ is less), however this multi-transition setup is still considerably more efficient than observing each line separately. However, it is recommended that for continuum observations, the IF range 6-10 GHz is used. Also, it should be noted that the full range 5–10 GHz cannot be completely sampled because of the limited 4 GHz width of the two basebands per polarization.

The atmospheric transmission in Band 6 is shown in Figure 4.18 for three typical PWV values. Most of the narrow absorption lines are from ozone.

The specification for Band 6 receiver noise performance ($T_{\rm rx}$) is <83 K over 80% of the band, and <136 K over the whole band (SSB $T_{\rm rx}$). The achieved on-array results are considerably better, typically ∼40–50 K over most of the band. The OT assumes 55 K. The resulting system temperatures ($T_{\rm sys}$) for 1.262 mm PWV are shown in Figure 4.19.

[1] Specifically, the ^{12}CO/^{13}CO/C^{18}O $J=2$ -1 combination at 230.538/220.398/219.560 GHz, which has a minimum separation of 10.14 GHz and requires the IF to reach to 5.0 GHz in order to cover all three lines.

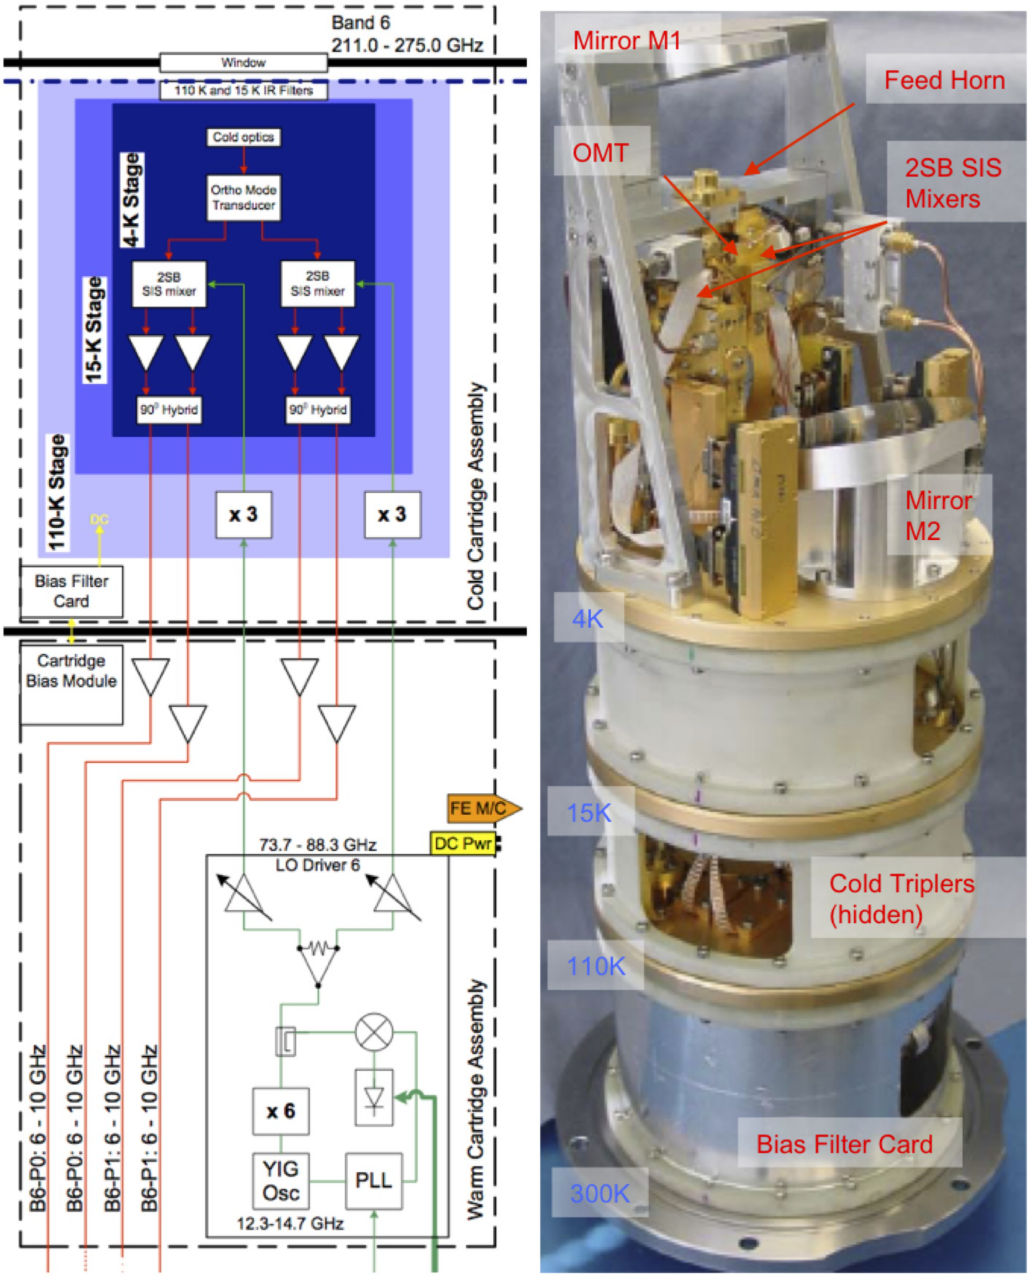

Figure 4.17: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra $\times 3$ multiplier inside the cryostat. The Band 6 cartridges were built at NRAO, Charlottesville. Note that the IF output range is actually 5-10 GHz; the range shown is the one recommended for continuum observations (see text).

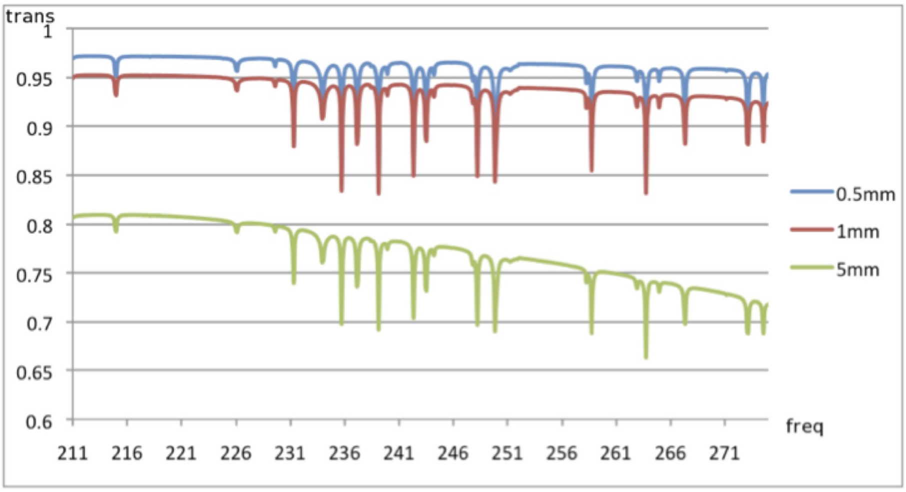

Figure 4.18: Band 6 zenith transmission for PWV=0.5, 1 and 5 mm. Frequency is in GHz. Most of the narrow absorption lines are from ozone.

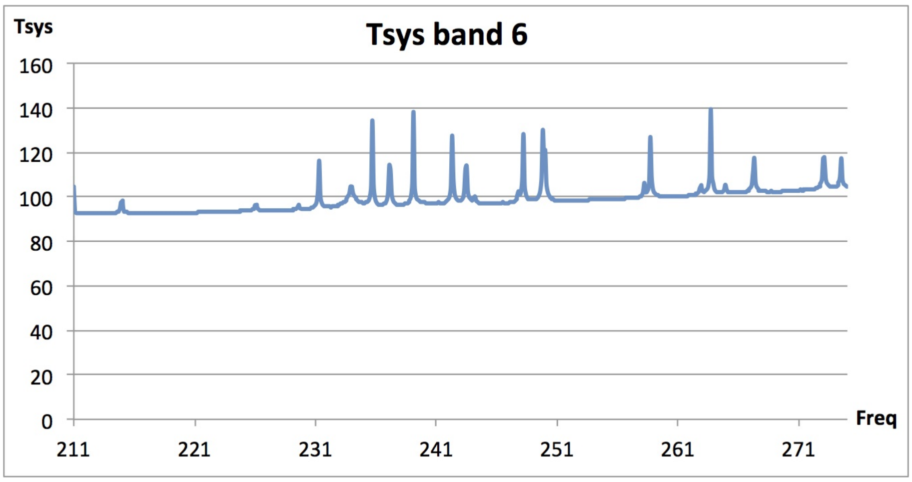

Figure 4.19: Typical $T_{\rm sys}$ at zenith for Band 6 with 1.262 mm PWV, based on measured values of the receiver temperatures. $T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.5 Band 7 Receiver

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

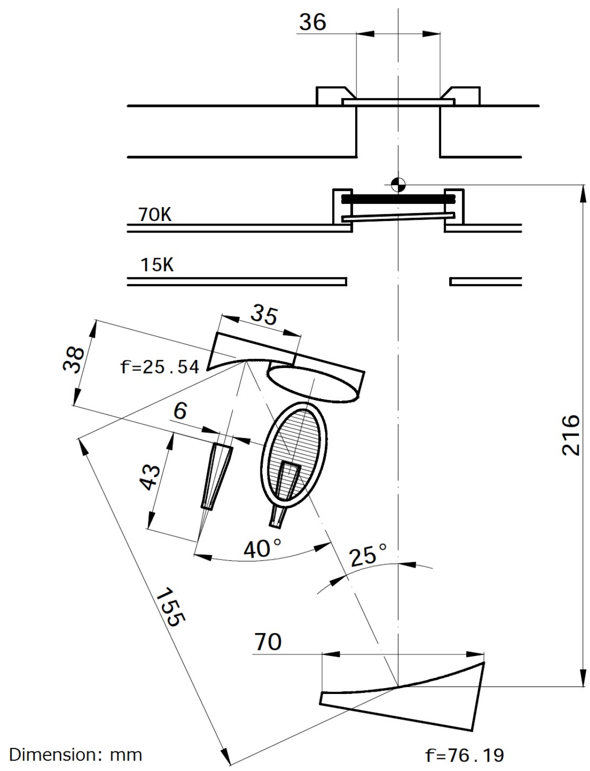

Figure 4.20: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid.

The atmospheric transmission in Band 7 is shown in Figure 4.22 for three typical PWV values. The specification of the Band 7 receiver noise temperature is $T_{\rm rx}$ <147 K over 80% of the range and <219 K over the whole tuning range, except at the upper end of the band (370-373 GHz), where the specifications are <300 K SSB. However, the performance of the receiver as measured in the lab and on the array is considerably better than this, with typically 65 K achieved, in mid-band. The OT assumes 75 K over the whole band. The resulting system temperatures ($T_{\rm sys}$) for 1 mm PWV are shown in Figure 4.23. Note that the atmospheric transmission (and hence $T_{\rm sys}$) at frequencies below 300 GHz is considerably better than that of the top half of Band 7; in that respect the performance is closer to that of Band 6. Also to maximise continuum sensitivity, the deep water absorption lines around 325 GHz should be avoided if possible.

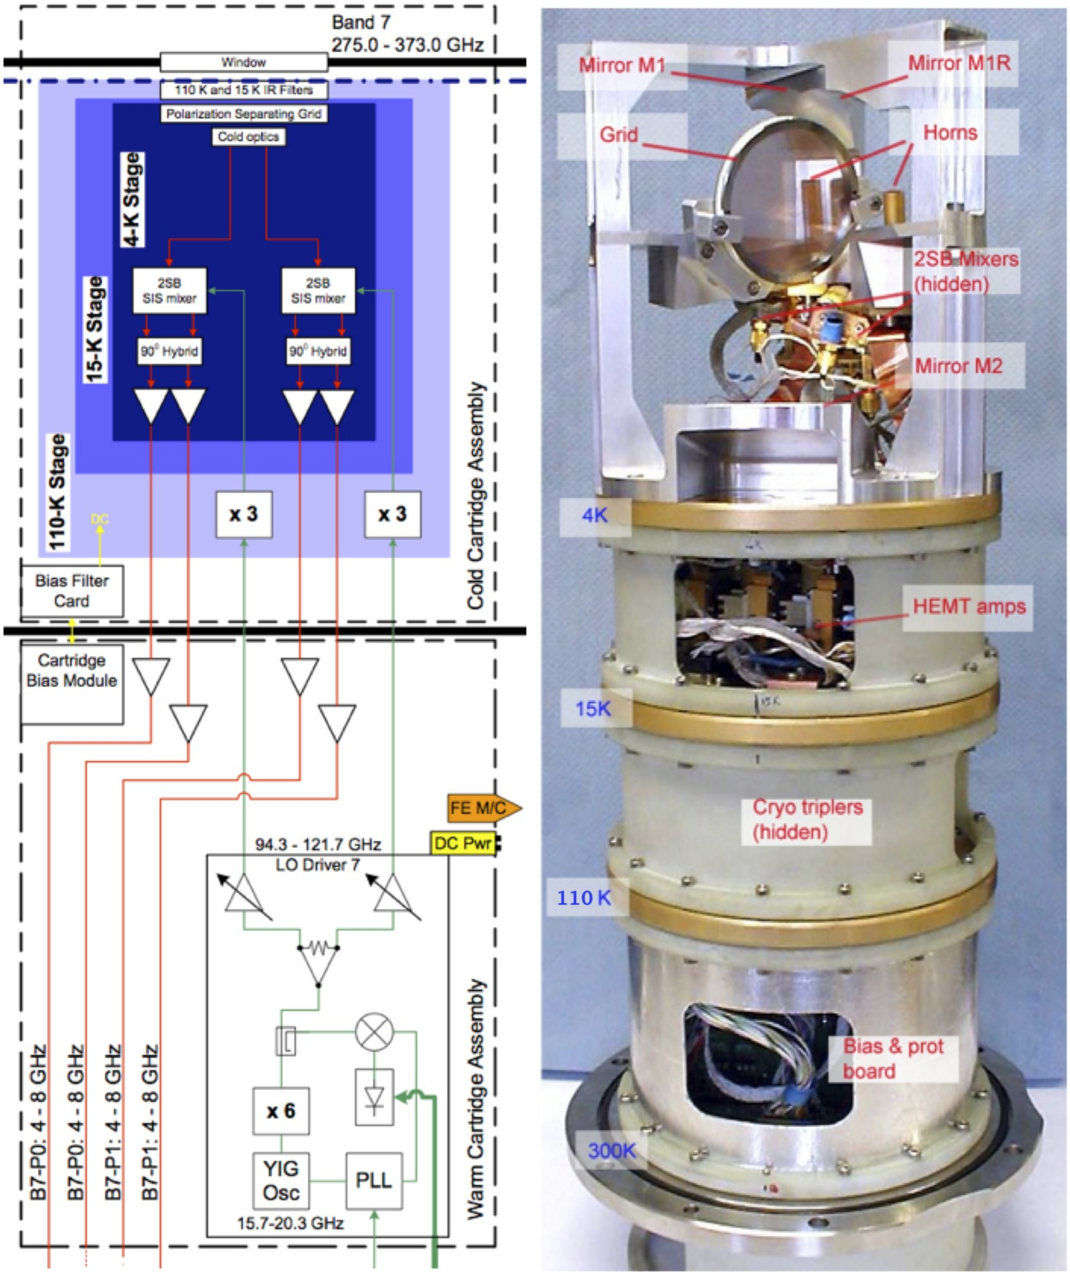

Figure 4.21: Band 7 front-end receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built at IRAM in France.

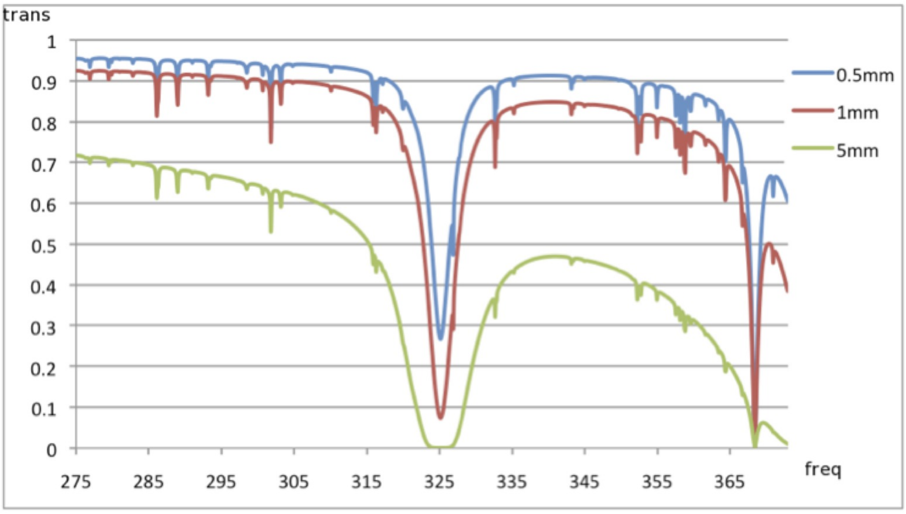

Figure 4.22: Band 7 atmospheric zenith transmission for PWV=0.5, 1.0 and 5.0 mm. Frequency is in GHz. The deep atmospheric absorption at 325 GHz is due to water, and the narrower feature at 369 GHz is due to oxygen.

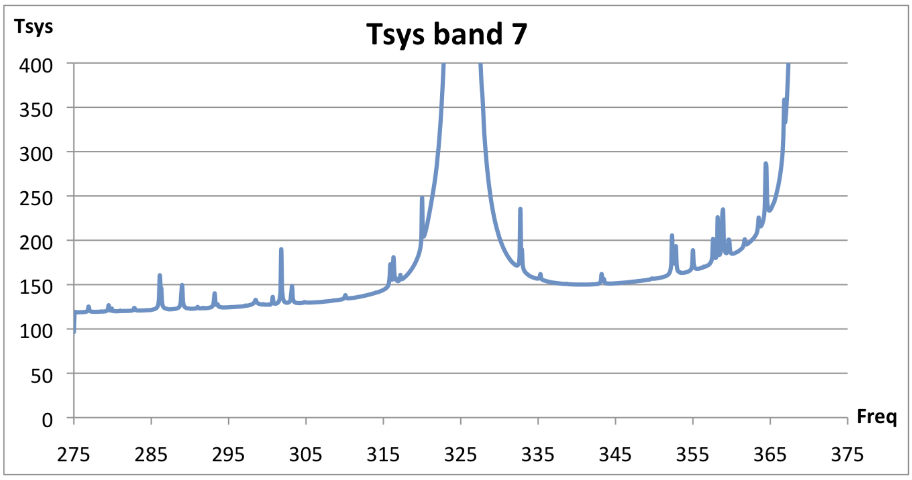

Figure 4.23: Typical T_{sys} at zenith for Band 7 with PWV=0.913 mm. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.6 Band 8 Receiver

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

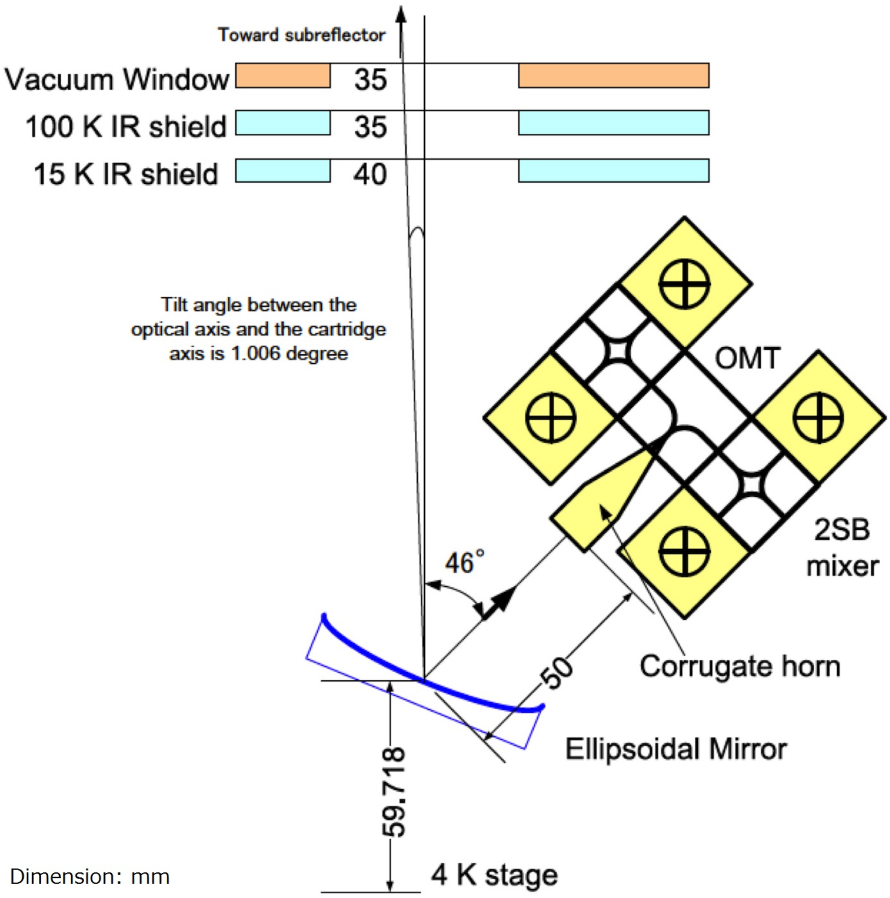

Figure 4.24: Optical layout of the Band 8 receiver.

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

The atmospheric transmission in Band 8 is shown in Figure 4.26 for three typical PWV values. The specification of the Band 8 receiver noise temperature is $T_{\rm rx}$ <196 K over 80% of the range and <292 K over the whole tuning range. However, the performance of the receiver as measured in the lab and on-array is considerably better than this, with typical values of 70-120 K. The OT assumes 150 K. The resulting system temperatures ($T_{\rm sys}$) for 0.472 mm PWV are shown in Figure 4.27. Note that transmission in the lower part of the band is almost as high as Band 7, whereas at the high-frequency end of the band (e.g. around the CI line at 492 GHz) the transmission is as low as the centre of Band 9.

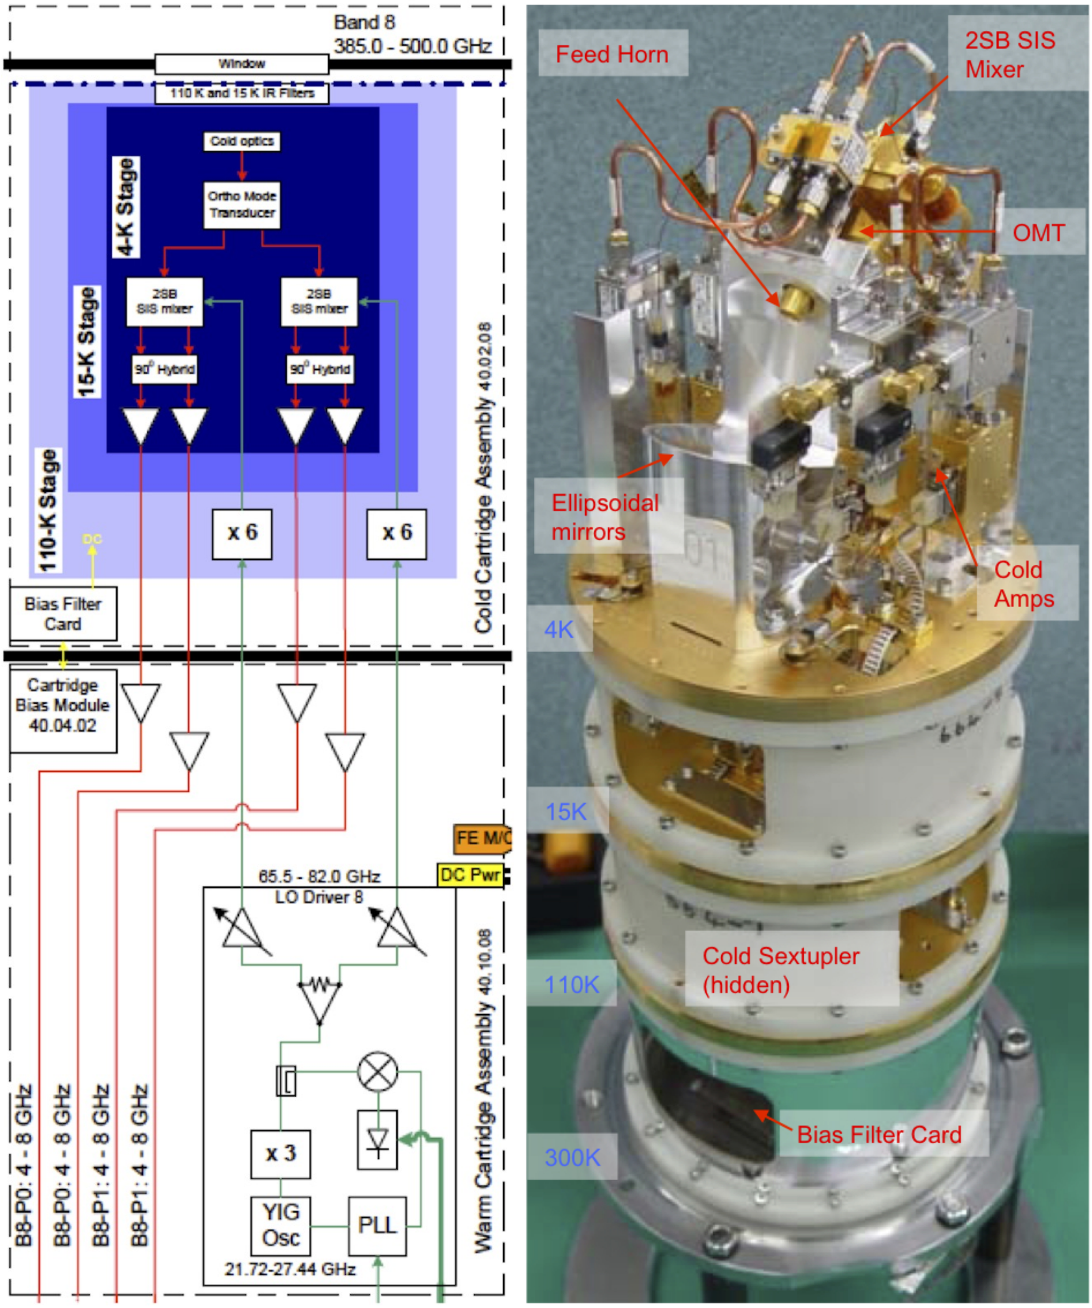

Figure 4.25: Block diagram of the Band 8 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 8 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 8 cartridges were constructed at the Advanced Technology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

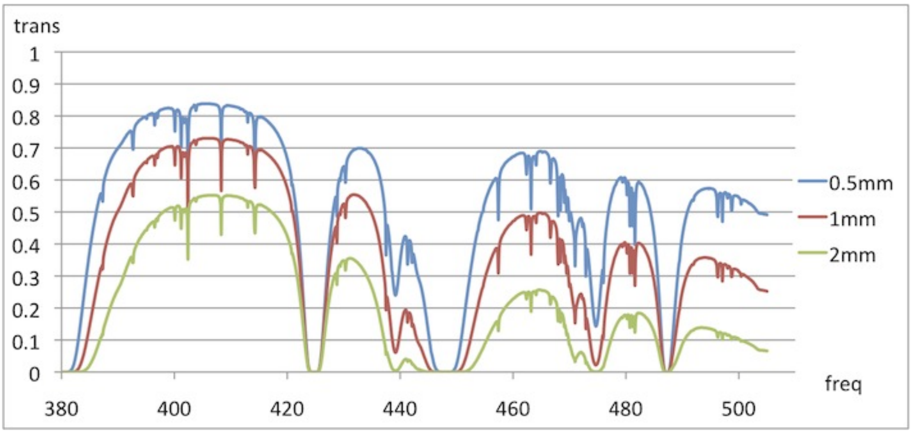

Figure 4.26: Band 8 atmospheric zenith transmission for PWV=0.5, 1.0 and 2.0 mm. Frequency is in GHz. The atmosphere in the Band 8 frequency range has some deep absorption by water and oxygen.

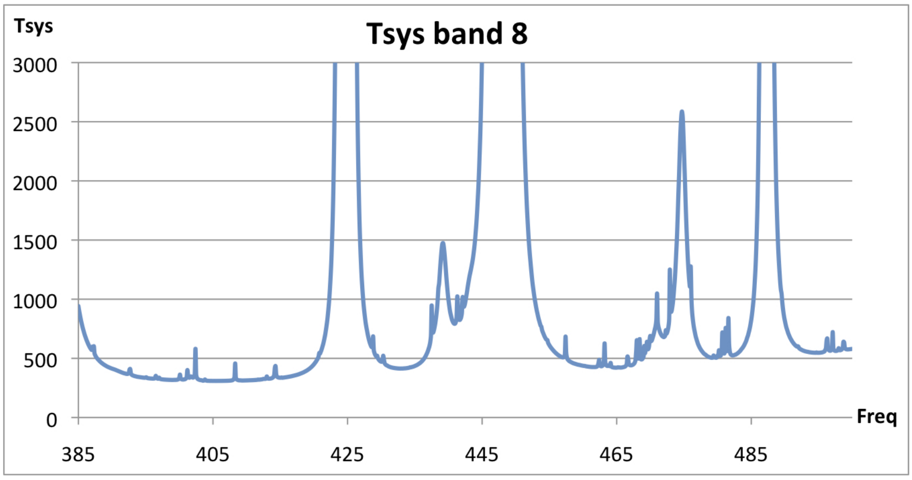

Figure 4.27: Typical T_{sys} at zenith for Band 8 with PWV=0.472 mm. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.7 Band 9 Receiver

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

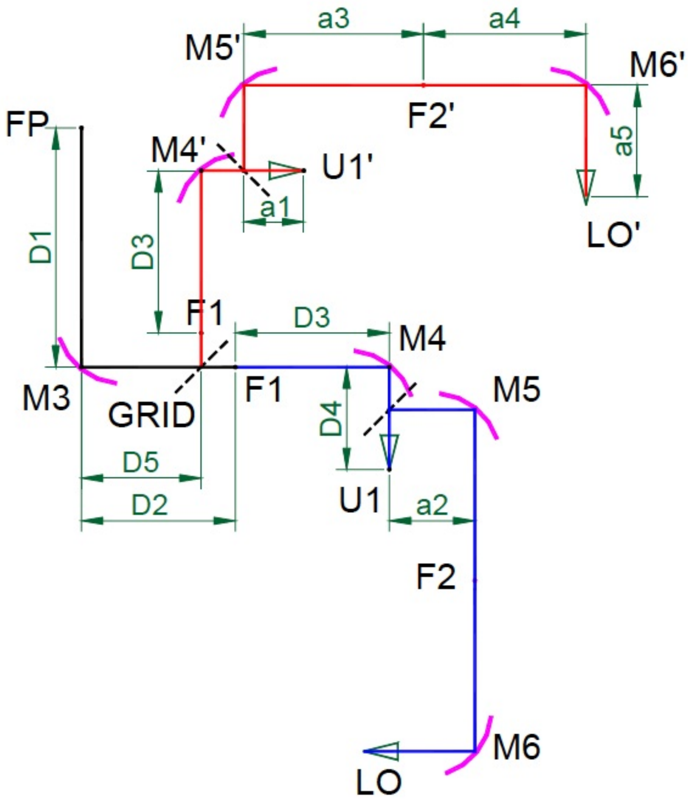

Figure 4.28: Basic Band 9 optics layout. The signal path is symmetrical for the vertical polarization (P0) and the horizontal polarization (P1). For P1 the path is as follows: a telescope focal point followed by mirror M3, grid, mirror M4, a beam splitter for LO insertion and finally the mixer horn U1.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

The mixers are double sideband (DSB), and therefore additional techniques must be employed during the observations to either separate the sidebands or reject the unwanted sideband. LO offsetting can 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 so that the signal is smeared over a larger bandwidth increasing noise incoherently. In Cycle 5, 90-degree phase switching can also be used to correlate both USB and LSB simultaneously (see Chapter 6). The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 6.4), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 4.29.

The Band 9 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.30 for three low values of PWV. The specifications for the receiver are $T_{\rm rx}$ <175 K over 80% of the band and <261 K over all the band. However, the performance is considerably better than this, with engineering and typical on-array values for $T_{\rm rx}$ of 65-120 K. The OT assumes 110 K. Figure 4.31 shows the expected $T_{\rm sys}$ for 0.472 mm of PWV, over most of the band given the OT default expected receiver noise, however, the achieved value is extremely dependent on the line-of-sight PWV. Phase stability also limits when observations can be made, therefore, most observations in Band 9 are done at night (most commonly during austral winter). As well as having a lower atmospheric transmission and a less stable atmosphere, Band 9 observing provides several

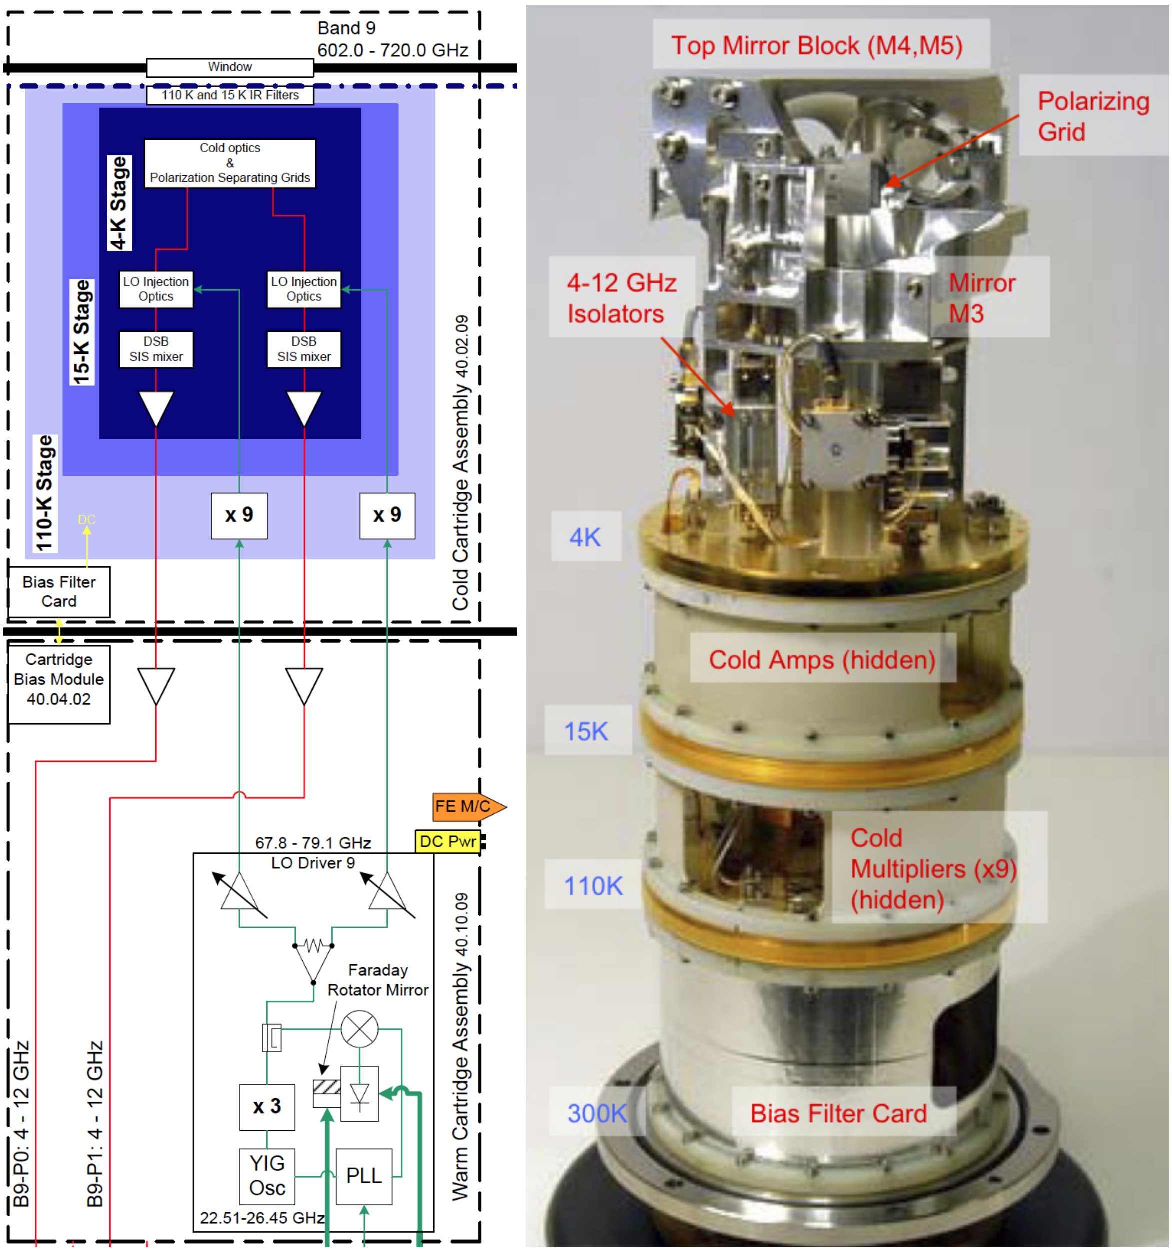

Figure 4.29: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The extra Faraday rotation mirror in the LO system is part of the ALMA fibre-optic Line Length Corrector system (see Appendix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

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

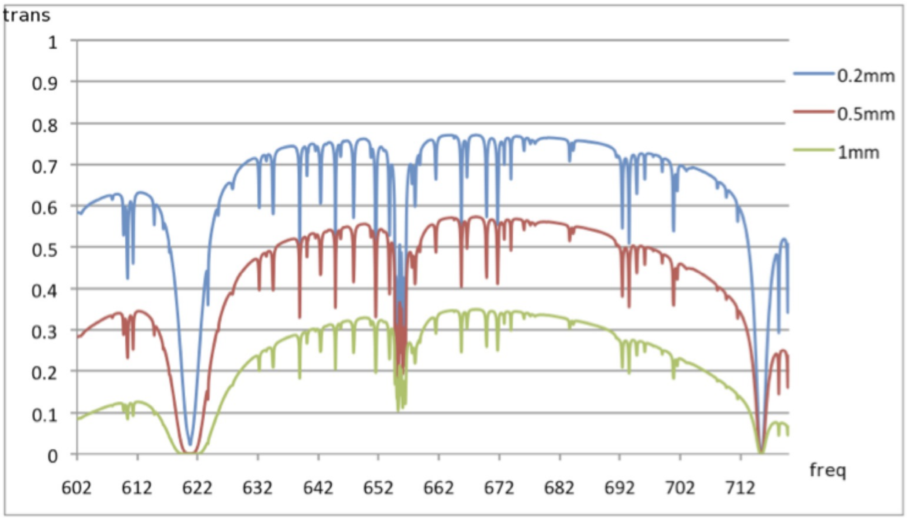

Figure 4.30: Band 9 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

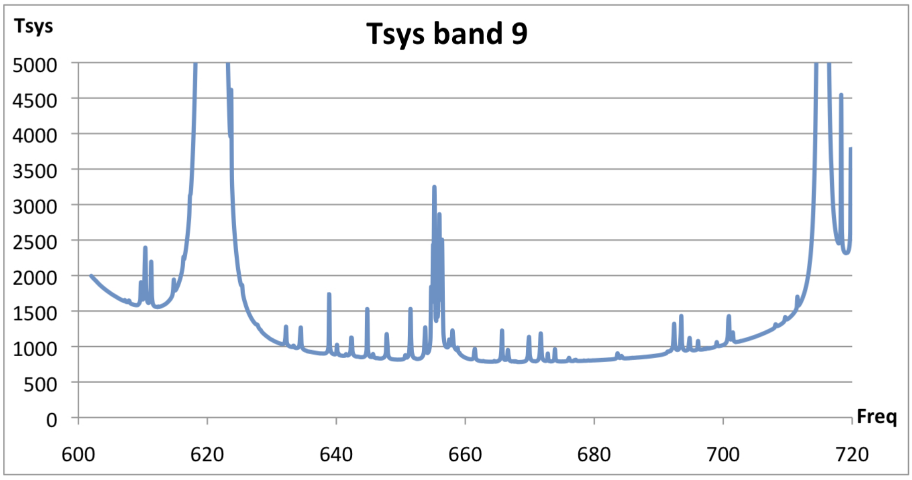

Figure 4.31: Typical T_{sys} at zenith for Band 9 with PWV = 0.472 mm. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

4.2.8 Band 10 Receiver

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

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

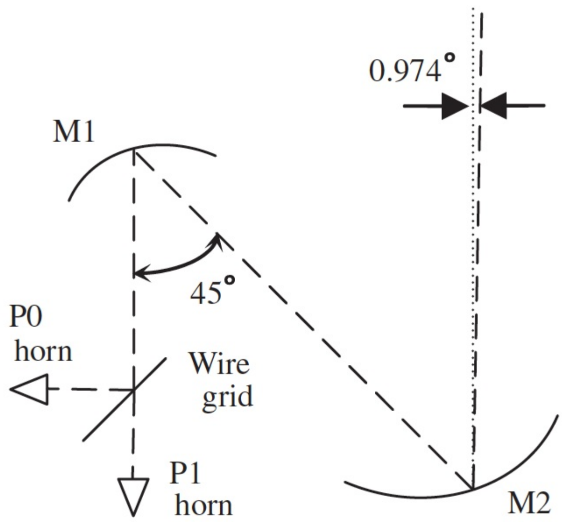

Figure 4.32: Schematic of ALMA Band 10 optics. ALMA Band 10 optics are composed of two elliptical mirrors, M1 and M2, a wire grid and two corrugated horns.The wire grid is used to separate the two linear polarizations, P0 and P1, and it is located after the two elliptical mirrors to minimize the number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

The Band 10 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 4.34 for three low values of PWV. The specifications for this receiver are $T_{\rm rx}$ <230 K over 80% of the band and <344 K over all the band. Typical on-array $T_{\rm rx}$ values are 175-275 K. The OT assumes 230 K. Figure 4.35 shows the expected $T_{\rm sys}$ for 0.472 mm of PWV, over most of the band given the OT-assumed receiver noise. Phase stability also limits when observations can be made. Also the antenna surface accuracy has a significant effect on the aperture efficiency at this high-frequency band. Therefore, most observations in Band 10 will be done at night (and mostly during austral winter), when the sky and the system is most stable. As well as having a lower atmospheric transmission and a less stable atmosphere, Band 10 observing provides the most challenges for observing: finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing for the relatively small primary beam (although pointing is done in Band 6), and the need for the highest level of stability in the rest of the system.

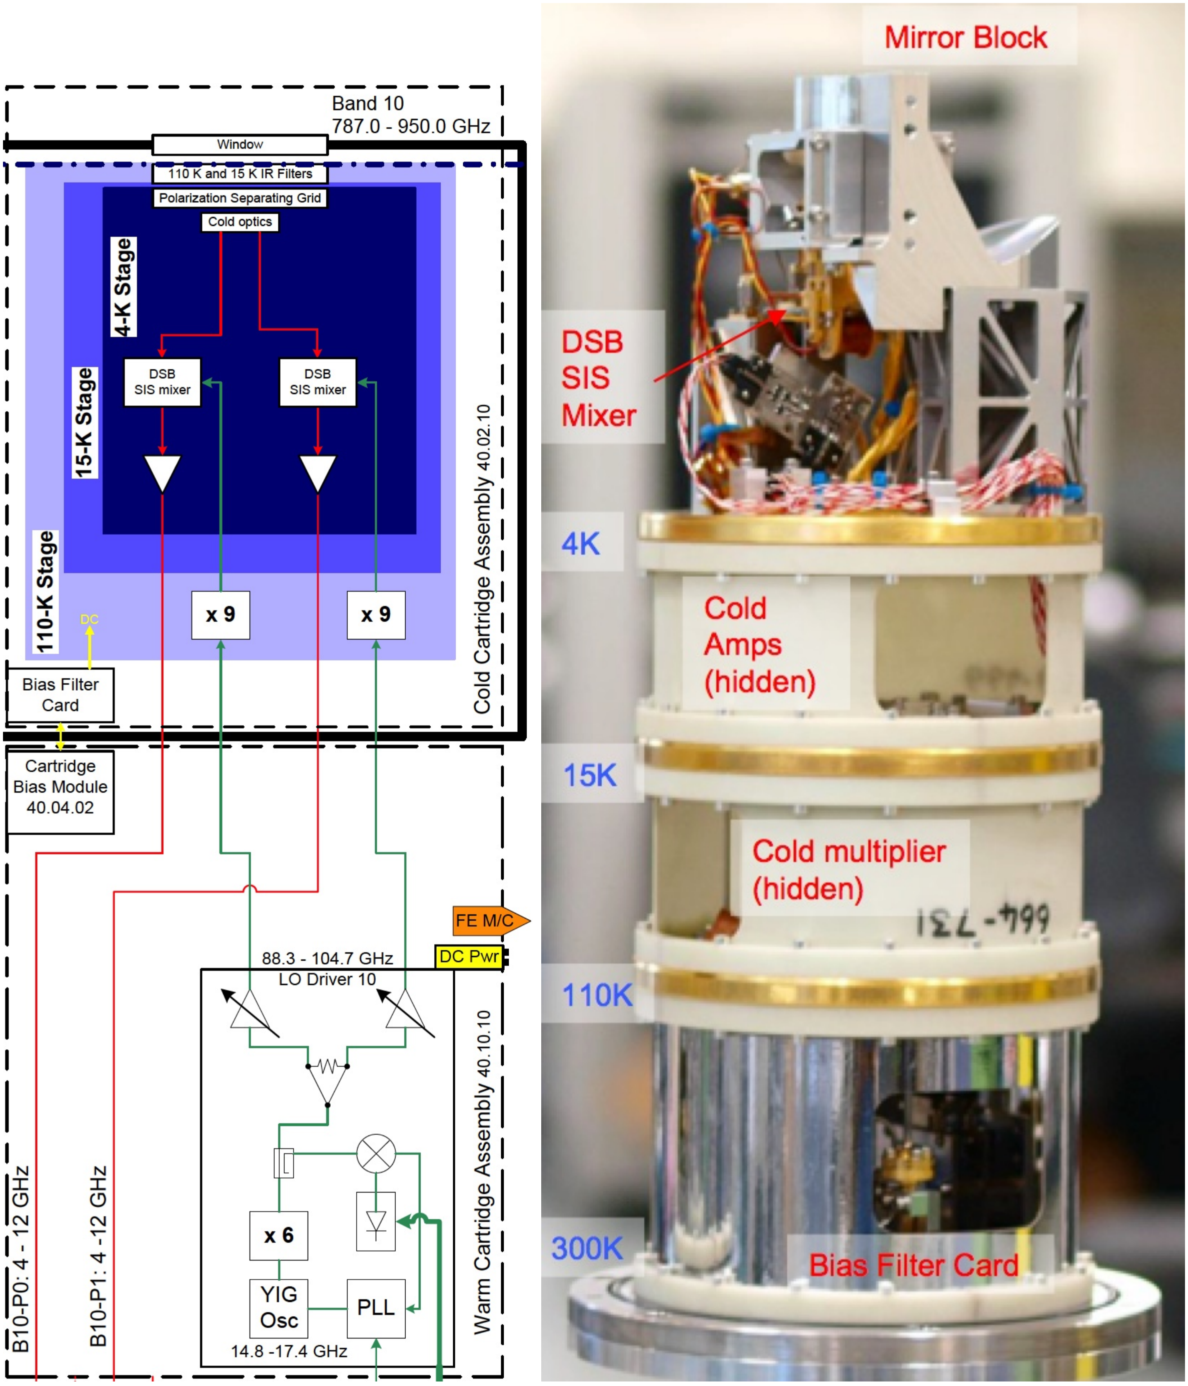

Figure 4.33: Block diagram of Band 10 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 10 cartridges were constructed at the Advanced Technology Center (ATC).

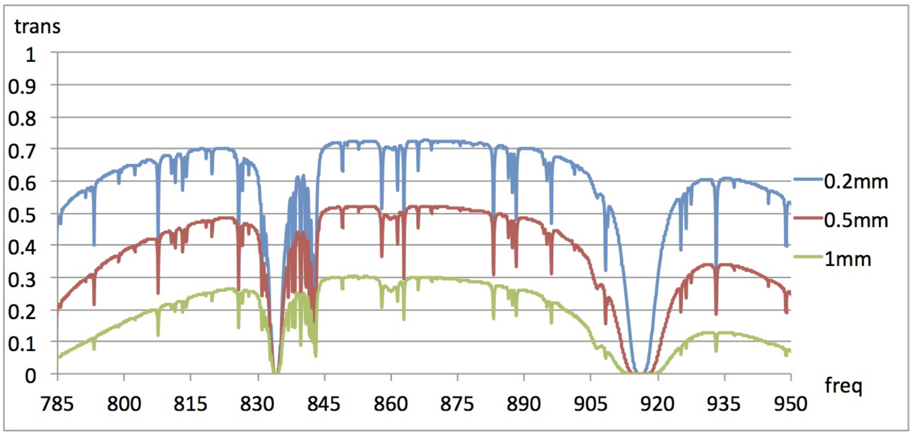

Figure 4.34: Band 10 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

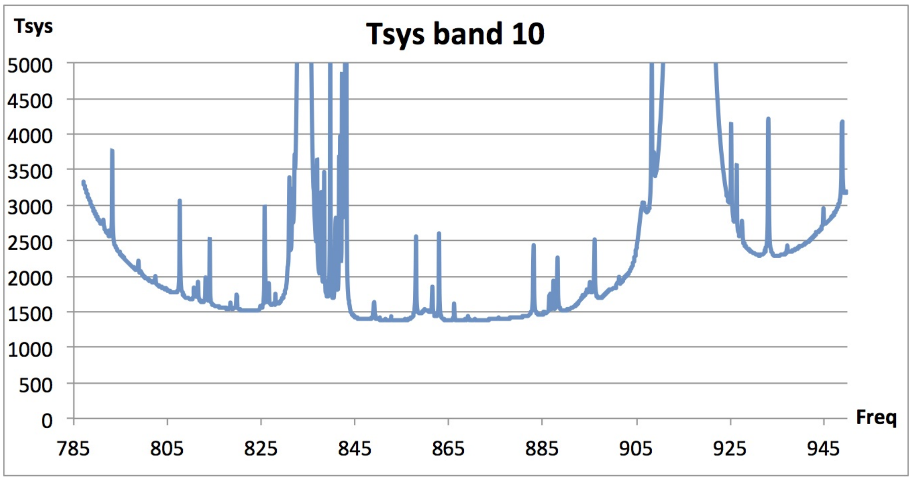

Figure 4.35: Typical T_{sys} at zenith for Band 10 with PWV = 0.472 mm. T_{sys} was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included. Temperature is given in Kelvin.

Chapter 5

The Correlators

A correlator can be viewed as the virtual focal plane of an interferometer array where all voltage-based signals from all individual antennas are processed to derive both the cross-correlation products from all independent antenna pairs and the autocorrelation function of each antenna. It delivers complex fringe visibilities which, once calibrated in amplitude and phase, allow the ALMA users to synthesize astronomical images. The correlator also provides pre-correlation delay and phase tracking functions to adjust the response to the wavefronts of received signals in order to maintain the coherence of the complex visibilities. Walsh-switching modulations for sideband separation (90° phase switching) and spurious-signal suppression (180° phase switching) are demodulated in the correlator.[1]

All signals received by ALMA antennas are processed in one of two correlators: the 64-input Correlator (also known as the Baseline Correlator or BLC) and the ACA Correlator. The 64-input Correlator is used primarily for the main 12-m Array, while the ACA Correlator is used for the Morita Array comprised of the ACA 7-m Array and the Total Power Array (i.e. single-dish), respectively[2]. Both correlators run simultaneously and independently. Thus, while the 12-m Array observes an object using the 64-input Correlator, the ACA Correlator can be used with the 7-m Array and/or the Total Power Array observing either the same or a different object.

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

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

This chapter addresses capabilities of the correlators to be offered for Cycle 5 observations. The specifications of the correlators determine a lot of observational performance such as bandwidth, spectral resolution, time resolution, and polarimetry. The phase tracking performance, including online WVR correction, provides coherence and improves phase stability in correlated data delivered to users. Imperfect correction for the non-linear response in the correlators invokes systematic errors in complex visibilities. The response of digital signal processing is also presented in this chapter.

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

[2] Crossbar switching allows for some flexibility in this arrangement.

5.1 The 64-input Correlator

The 64-input Correlator employs hybrid design, also known as the FXF[3] system (Escoffier et al. 2007, A&A 462, 801), that brings a 32× spectral resolution of the traditional lag (XF) correlators. It operates in two basic modes, Time Division Mode (TDM) — equivalent to an XF correlator with a wide bandwidth and a coarse spectral resolution for mainly continuum observations, and Frequency Division Mode (FDM) with fine spectral resolutions for spectral-line observations. A simplified overview diagram of the 64-input Correlator is shown in Figure 5.1. It consists of four quadrants, all of which are available for Cycle 5. Each quadrant can handle a 2 GHz dual-polarization BB for up to 64 antennas[4]. The full set of four quadrants is capable of accepting four BBs to cover a total 8 GHz bandwidth with a dual polarization.

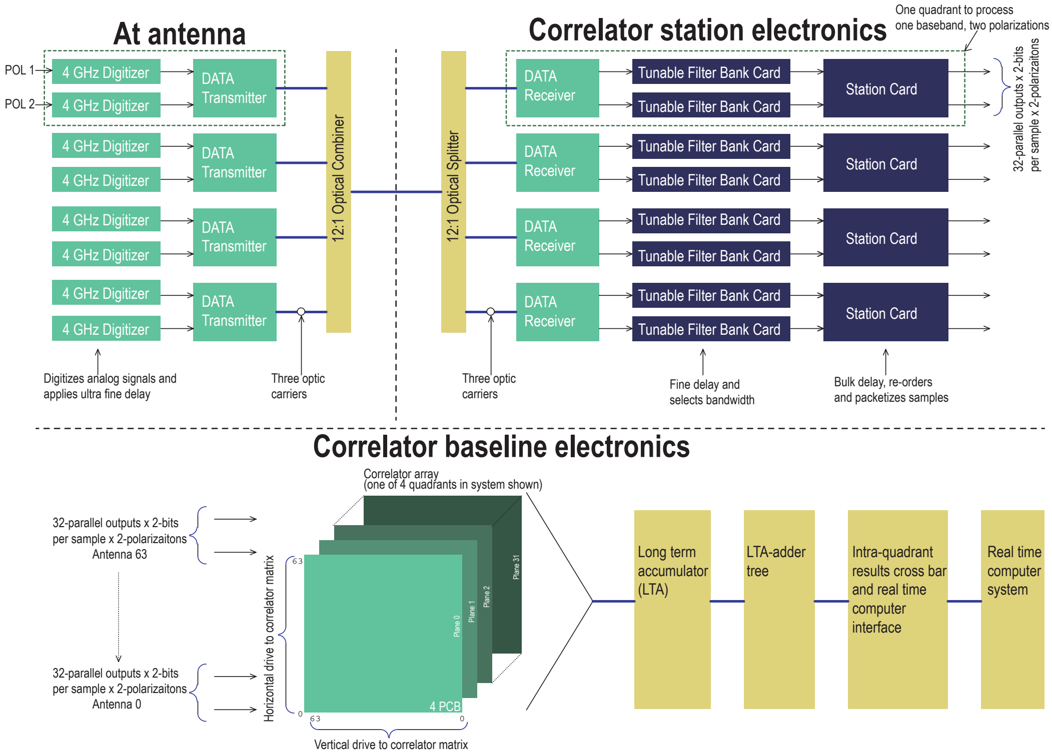

Figure 5.1: Overview diagram of the ALMA digitizers, data transmission system and 64-input Correlator (Escoffier et al. 2007, A&A 462, 801). The digitized data from the individual BBs are transferred to the Tunable Filter Banks (TFBs) and station cards (top right) for the parallelization and processing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

5.1.1 TDM mode

TDM is mostly used for continuum observations. The simplicity, compared with FDM, offers advantages of a lower data rate and better linearity. Therefore, it is used for standard setups such as pointing, focus, delay calibration, system temperature measurements, sideband ratio measurements, etc.

The full 2 GHz BB is directly sent to the correlator bypassing the Tunable Filter Banks (TFBs). The correlator cuts off the Least Significant Bit (LSB) to reduce quantization levels from 3- to 2-bits per sample

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

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

Figure 5.2: The ALMA 64-input Correlator. This view shows lights glowing on some of the racks of the correlator in the ALMA Array Operations Site Technical Building and shows one of four quadrants of the correlator. Credit: ALMA (ESO/NAOJ/NRAO), S. Argandoña

(see Section 5.3).

The TDM mode provides a Spectral Window (SPW)[5] that consists of up to $256/N_{\rm pol}$ channels per BB, where $N_{\rm pol}$ is the number of polarization products per BB[6]. As the full 2000 MHz BB is covered, this requires some truncation of the band-edge channels in offline data processing yielding 1875 MHz of usable bandwidth – see Section 6.4.

5.1.2 FDM mode

FDM is adequate for spectral-line observations that require a spectral resolution higher than that of TDM. Each 2-GHz-bandwidth digitized signal stream is processed in a TFB card, where digital filtering and frequency conversion with the LO4 are implemented in a Field Programmable Gate Array (FPGA). In a TFB there are 32 independent 62.5 MHz bandpass filters or sub-bands. To avoid edge artifacts, sub-band center frequencies are set such that their edges overlap and the effective bandwidth is thus reduced by a factor of 15/16, that is, 58.59375 MHz of effective bandwidth per sub-band. The correlator software stitches together contiguous sub-bands to output a seamless cross power spectrum to form an spw with as many as $\frac{15}{16} \times 8192/N_{\rm pol}$ channels. The number of channels in each spw can be reduced by averaging 2, 4, 8, or 16 channels into one to accommodate the maximum data rate. Table 5.1 lists the spectral setups for Stokes I ($N_{\rm pol} = 2$) observations. The number of valid sub-bands to form an spw is selectable from 32, 16, 8, 4, 2, or 1. That is, spw bandwidths of 1875, 937.5, 468.75, 234.375, 117.1875, and 58.59375 MHz, respectively. It is possible to have multiple (up to 4) spws in the same BB. While different spectral setups can be set for spws in different BBs, all of spws in the same BB must have the same channel spacing.

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

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

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

See also Section 5.5.2 for detail about spectral setting.

5.1.3 Correlation and Realtime Processing

The correlation cards perform the multiply-and-add operations to produce the correlation functions at a clock rate of 125 MHz (4 GHz samples demultiplexed by 32). Four quadrants handle four BBs. A quadrant of the correlator consists of 32 planes of 64×64 256-lag correlator circuits, and it yields auto-correlations and cross-correlations for 64 antennas[7]. In FDM there are always 8192 spectral points per BB per polarization, and higher spectral resolutions can be achieved by reducing the number of sub-bands (i.e. bandwidth) used to form an spw. It is possible to set different modes in different quadrants (i.e. BBs); for example, while one BB is set in TDM, other BBs can be set in FDM. For spectral setup details, see Chapter 6. The Long Term Accumulator (LTA; see Figure 5.1) takes short 1 ms or 16 ms integrations from the correlator circuits and provides longer term integration. Further time averaging through multiple dump intervals is performed in the Correlator Data Processor (CDP) computers. See also Section 5.5.3 about time resolution. Section 5.5 describes in detail the correlator data processing.

5.2 The ACA Correlator

The ACA Correlator is dedicated to observations with the Atacama Compact Array that consists of the 7-m Array with twelve 7 m antennas and the Total Power Array (TP Array) with four 12 m antennas. The ACA Correlator is based on the FX[8] design in which the incoming time-domain data stream is converted into frequency-domain spectrum via an FFT (Fast Fourier Transform) module before cross multiplication to form the power spectrum. The FFT part always accepts a full 2 GHz bandwidth and outputs 524288 channel spectra with a frequency resolution of 3.815 kHz. This design enables flexible spectral handling as shown in Figure 5.6 producing multiple spws with different channel spacings by averaging multiple channels (spectral binning). Since the spectral resolution function is different from that of the 64-input Correlator, a frequency profile synthesis (FPS) is performed in the Correlator Data Processor computer so that the outputs of both correlators are matched (Kamazaki et al. 2008, ALMA Memo 580).

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

After receiving the digitized BB signals in the Digital Transmission System Receiver (DTS-Rx), the DTS-Rx and FFT Processor (DFP) modules compensate for geometrical delays between antennas and perform the 2^{20}-point FFT that produces a 2^{19}-point complex spectrum[9] (hereafter, voltage spectrum) for every BB per antenna, with a channel separation of 3.815 kHz (= 2 GHz $\div$ 524288 channels) in a 16-bit complex integer form. The 16-bit complex voltage spectra are re-quantized into a 4-bit complex integer and sent to the Correlation and Integration Processor (CIP) modules.

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

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

The overall hardware design of the ACA Correlator is shown in Figure 5.5. The ACA Correlator is equipped with high-speed FPGA chips rather than Application Specific Integrated Circuit (ASIC) chips used in the

[7] 2016 cross-correlations (real and imaginary) and 64 auto-correlations (real) totals 4096 real correlations.

[8] F and X stand for Fourier transform and correlation, respectively.

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

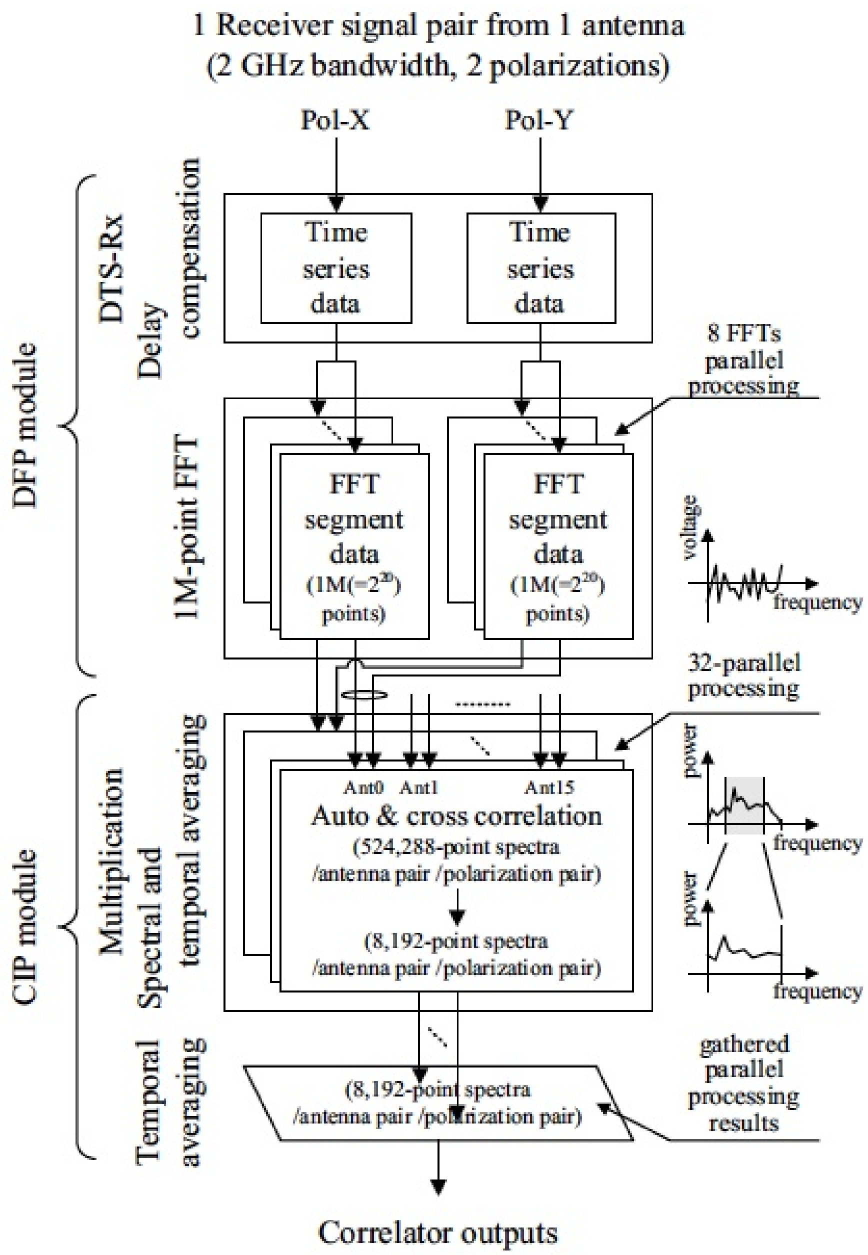

Figure 5.3: Block diagram of the ACA Correlator. One of four quadrants is shown. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. FFT is performed in an FPGA delivering signals in the spectral domain which are then multiplied. The auto and cross-correlation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

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

64-input Correlator. For technical detail, see Kamazaki et al. 2012, PASJ 64, 29.

5.2.1 The 7-m Array

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

5.2.2 The Total Power Array

The TP (Total Power) Array consists of four 12 m antennas in single-dish observation mode, for the purpose of obtaining zero-spacing visibilities to recover the spatially-integrated flux densities that are not obtained with cross-correlations. In the TP Array observations, the ACA Correlator generates only parallel polarization auto-correlations of XX and YY for Cycle 5. Cross-polarization products are not taken in Cycle 5. Single-dish dedicated observing modes such as position switching, raster scan, and OTF will be performed with the TP Array. Fast switching ($\sim$ 10 Hz) using the nutating subreflector will not be available during Cycle 5.

5.3 Digitizers

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

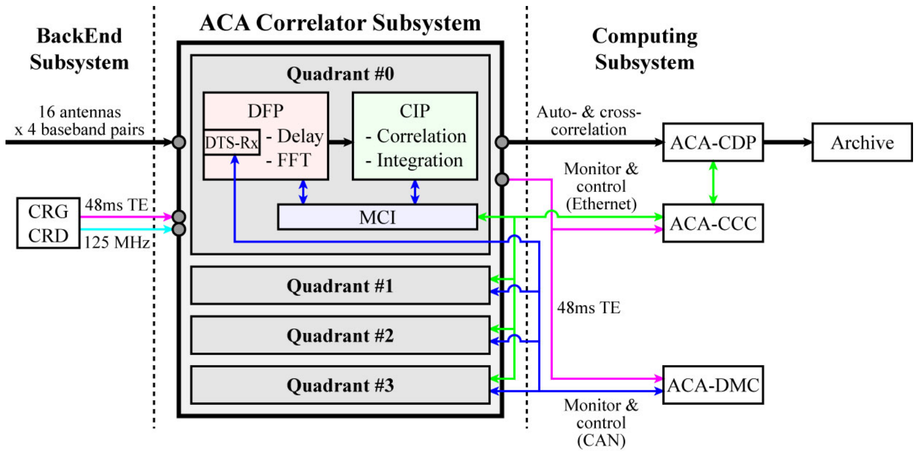

Figure 5.5: Overview diagram of the ACA Correlator. The signals from the BackEnd Subsystem are processed in the DFP and CIP modules and output to the Computing Subsystem (see text and Figure 5.3 for details of these). The modules are controlled by the Monitor and Control Interface (MCI), which communicates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

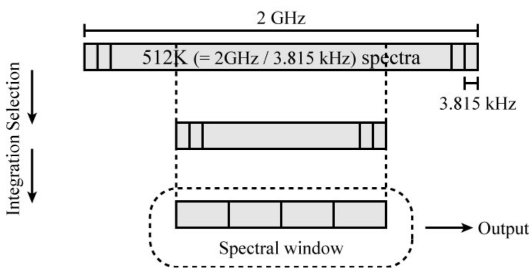

Figure 5.6: Data capture sequence of the ACA Correlator. The design enables output with multiple spws with different channel spacings by averaging multiple channels (spectral binning). The frequency range of an spw is selectable from the 524288-point (512K) spectra across 2 GHz bandwidth.

Each antenna BE is equipped with 8 digitizers to accept 8 signal streams of 4 BBs × 2 polarizations. To sample and quantize the 2-4 GHz BB signals, the ALMA digitizer samples at the Nyquist frequency of 4 GHz and uses 8 quantization voltage levels (3-bit per sample). However, the 64-input Correlator reduces the quantization levels to 4 (2-bit). Signals through the TFBs are re-quantized in 4-bit before the correlation process. The ACA Correlator handles 3-bit digital form in FFT and the spectrum is requantized in 4-bit before cross multiplication. These quantization processes affect the sensitivity and linearity in the power spectra. Compared to an ideal analog system, the 64-input correlator has a sensitivity loss of 12% (TDM) or 15.5% (FDM) due to the quantization including re-quantization. Similarly the ACA Correlator has a 4.8% sensitivity loss. These numbers suppose that the digitizer input power is in the expected range and that the digitizer adjustments are optimal. See also Section 5.6.1 about the quantization loss.

5.4 Online WVR Correction

The Water Vapor Radiometer (WVR; see Appendix A.6 for technical description) is a device mounted on each of the 12 m antennas to measure the amount of Precipitable Water Vapor (PWV) in order to correct for pathlength fluctuations in the troposphere. The WVR data are recorded together with the uncorrected visibilities, and in Cycle 5 (as in previous cycles), the WVR corrections are made offline in CASA. However, the CDP of the 64-input Correlator can also apply a realtime correction of the path length fluctuation for each dump duration (online correction). This online corrected data is currently used for QA0 checks whenever possible. Recording both WVR corrected and uncorrected results impacts data rate and storage, and it needs to be dealt with by means of longer integration time.

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

5.5 Capabilities of the Correlators

5.5.1 Polarization

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

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

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

The 64-input Correlator forms cross-polarization products of $\langle XY^*\rangle$ and $\langle YX^*\rangle$ only when polarimetry is required. This is obtained at the expense of less spectral channels. Although the ACA Correlator is capable of polarimetry, the cross-polarization products of the 7-m Array or the TP Array observations will not be offered for Cycle 5.

5.5.2 Spectral Setup

Both correlators are capable of accepting eight 2 GHz bandwidth signal streams consisting of four BBs and two polarizations, and to output multiple spws within the bandwidths.

Each spw yields a continuous spectrum composed of uniformly spaced spectral channels. Table 5.1 summarizes the maximum number of spectral channels and the channel spacing of the correlators in dual linear polarization mode (XX and YY). The maximum number of channels per spw would by doubled (and the channel spacing halved) for single polarization (XX) observations. Conversely, for full polarization observations (XX, XY, YX, and YY), the maximum number of channels is halved and the channel spacing correspondingly doubled as compared to Table 5.1. The bandwidth and the number of channels available from the ACA correlator are identical through the application of FPS (Frequency Profile Synthesis). Note that the channel spacing is not the same as the spectral resolution because of the effects of the applied weighting function, as listed in Table 5.2 and shown in Figure 5.7. See also Chapter 6 about multiple-spw setup.

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

Table 5.1: Correlator modes and spectral performances per BB for dual parallel polarization (XX and YY). *Usable bandwidth excluding band edges is $\sim$ 1875 MHz.

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

Table 5.2: Spectral resolutions (FWHM) and the maximum sidelobe levels for weighting functions. Default is the Hanning window.

The spectral profile which the correlators output is a convolution of the true spectrum with the spectral resolution function shown in Figure 5.7. The spectral resolution function is given by Fourier transform of the weighting function applied to the correlation function in the lag domain. The spectral resolution is often characterized by the FWHM (full width at half maximum) of the spectral resolution function. Note that the channel spacing in Table 5.2 is not the same as the spectral resolution, which depends on the weighting function in the lag domain.

Uniform weighting in the lag domain results in the sinc function with the FWHM of 1.21× the channel spacing and yields spectral sidelobes. This 'ringing' phenomenon affects the spectra when a narrow line, interference spike, or strong edge channels are present. Alternative weighting functions[10] are applicable to suppress the spectral sidelobes as listed in Table 5.2. Note that the sidelobes do not matter in most of astronomical observations where a line profile spreads over several spectral channels. The choice of a weighting function is a trade off between resolutions and sidelobes. The default weighting function is the Hanning function, which gives the spectral resolution of 2 × the channel spacing and the maximum sidelobe level of −2.6%. If a different weighting function is required in a specific science goal, the proposer should justify it in the technical case, and select the setup in the Phase 2 SB.

Spectral channel averaging is available to bin or average spectral channels in the CDP. Channels can be averaged together; factors of N = 2, 4, 8, or 16 are available. The main purpose is to reduce the data rate to the archive and the total data volume. It provides a broader spread of correlator functionality between the current TDM (which has only 128 channels in dual polarization) and full FDM (with 3840 channels in dual polarization mode). It might be quite acceptable for those who need something with more resolution than TDM, but where the FDM channels at the full resolution are unnecessary. Table 5.3 shows the resolutions (in kHz)

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

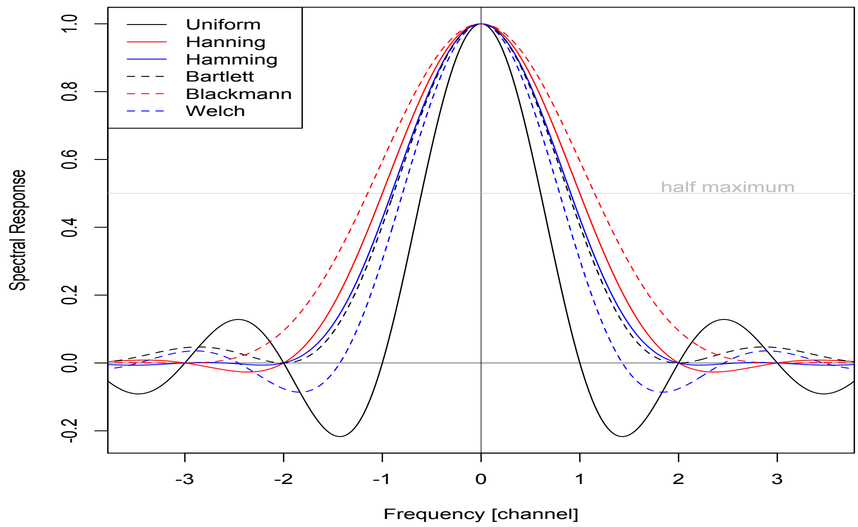

Figure 5.7: Spectral resolution functions corresponding to various weighting functions

for different values of N, using Hanning weighting, in the different bandwidth modes. The channel spacings are in parentheses. In Cycle 5, the default value is $N = 2$.

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

Table 5.3: Spectral resolution and channel spacing (in parentheses) in kHz for different correlator bandwidth modes (left column) and for different spectral channel averaging factors (columns, $N = 1$ to 16), using Hanning smoothing. The number of channels can be reduced from 3840 (for the unaveraged case, $N = 1$) down to 240 (for $N = 16$). As N increases, the spectral resolution functions of adjoining channels are combined and then the spectral resolution approaches the channel spacing. Values are given for the 2 polarization case. For Cycle 5 the default value for spectral averaging will be 2.

Note that the default Hanning window function gives a resolution 2 times the channel spacing, so using $N = 2$ (cutting the number of spectral channels from 3840 to 1920) results in negligible loss of final resolution. It is recommended that unless the maximum spectral resolution is required by the observations, to reduce the number of channels when feasible. This is selected in Phase 2 of the SB creation. However, note that this is a non-reversible operation. For Cycle 5 the default value for spectral averaging will be 2. The proposer must specifically request and justify the use of $N = 1$ for spectral averaging.

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

5.5.3 Time Resolution

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

1. Accumulator#1: implemented in the correlator itself. Data are accumulated without applying a proper averaging factor. The accumulation period is called dump duration.

2. Accumulator#2: implemented in the CDP, this stage first truncates $\sim 5\%$ band edges for each spw and then averages the whole spectral channels into one. Those channel averages are then further averaged in time during a number of dump duration periods. This average period is called channel average duration, which must be an integer multiple of the dump duration.

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

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

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

where $N_{\rm ant}$, $N_{\rm BB}$, $N_{\rm ch}$ are the number of antennas, BBs, and total number of spectral channels (for 1, 2, or 4 polarization products), respectively. $N_{\rm byte}$ (bytes/visibility) is the byte size of the real or imaginary part of a visibility, and $T_{\rm int}$ is the integration duration in the Accumulator#3. For example, an array configuration with $N_{\rm ant} = 42$, a spectral setup with $N_{\rm BB} = 4$, $N_{\rm ch} = 3840 \times 2$ polarization products (XX and YY), $T_{\rm int} = 6$ sec, and assuming $N_{\rm byte} = 2$ yields a data rate of 18.1 MB/sec.

In Cycle 5, ALMA regulates $R_{\rm int}$ at 40 MB/sec. The ALMA Observing Tool (OT) estimates the data rate for each science goal and warns if it exceeds the limit. In this case, it is required to consider the optimal trade between time and spectral resolutions, number of polarization products, number of BBs, and number of antennas in the array.

Some specific considerations about different time duration follows:

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

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

- Integration duration of spectral data with the ACA Correlator

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

- Channel-averaged duration

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

- Integration duration of online WVR Correction

  If online WVR corrections are requested by a given spectral configuration, then the integration duration $T_{\rm int}$ must be equal to or longer than that of the WVR system (1.152 sec). Given that for Cycle 5 neither the 7-m Array nor the TP Array uses online WVR, the above restriction applies to the 64-input correlator only.

A pulsar gating function is not implemented in the ALMA correlators. Timing sensitive analysis must be taken offline under the time resolution limited by the actual integration duration, $T_{\rm int}$, used during an observation.

5.5.4 Phase Switching

The 180° phase switching capability is implemented to suppress spurious signals and IF cross talks between the LO1 and the digitizers[11]. Starting with Cycle 5, a new capability has been introduced to separate, in real-time, upper and lower sidebands of DSB receivers (Bands 9 and 10). The expedient is based on 90° phase switching at the LO1, and accumulation in different memory areas (bins) in the correlator[12]. Walsh function series are employed to make antenna-based switching patterns orthogonal to each other. The Walsh sequences permutes every 16 ms, and the full cycle completes after 2048 ms to achieve orthogonality for all antennas in the array. The dump duration, during a side band-separation observation, must be an integer number of the complete Walsh cycle. See also Appendix B.4.3.

The 180° phase switching

This function was already applied before Cycle 4. The phase switching capability (see principles in Figure 5.9) is realized by phase modulation in the LO1 and demodulation in the DTS transmitter module after digitization.

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

Table 5.4: Cross correlation of a baseline under 180° phase switching. The received signal of antenna 2, S_2, is modulated by 180° by the LO1 and demodulated in the DTS transmitter module at each antenna. Spurious signals, N, in the IF or cross talks are not modulated via the LO1 and thus are canceled via demodulation and accumulation of data in states 0 and 1.

The 90° phase switching

This function aims to separate USB and LSB for the DSB receivers of Bands 9 and 10. When the LO phase of antenna 2 is offset by 90°, the phases of USB and LSB signals in IF will be $-\phi_{\rm LO} - \frac{\pi}{2}$ and $\phi_{\rm LO} + \frac{\pi}{2}$, respectively. The 90°-demodulation is taken by swapping real and imaginary parts of the cross correlations in the CDP. After demodulation, the sign of LSB correlation will be inverted with respect to USB. Thus, the average and the difference of cross correlations in two states gives the USB and the LSB visibilities, respectively.

[11] Note that this scheme doesn't work for autocorrelations.
[12] See Thompson, Moran, and Swenson, "Interferometry and Synthesis in Radio Astronomy (2nd edition)", §6.1 and §7.5.

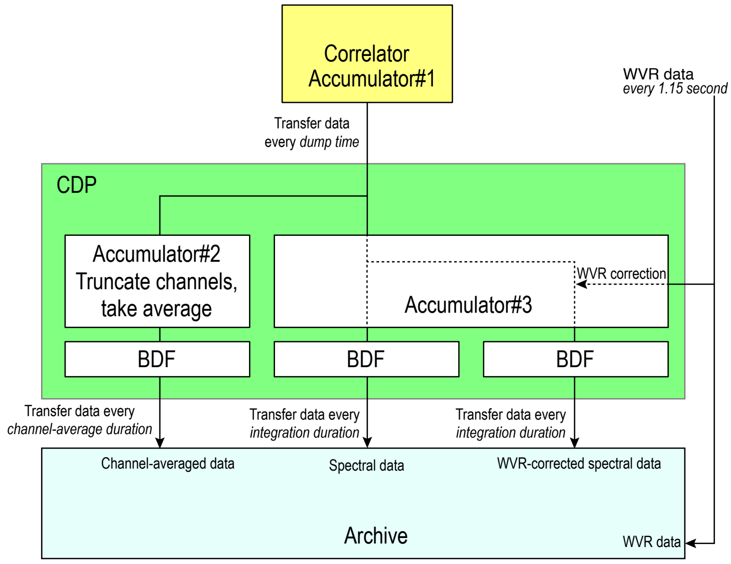

Figure 5.8: Basic data processing and accumulation/averaging steps between the correlator and archive. The CDP applies channel truncation, spectral averaging, and time averaging for the channel-averaged data. It also applies online WVR correction for the WVR-corrected spectral data. They are packed in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

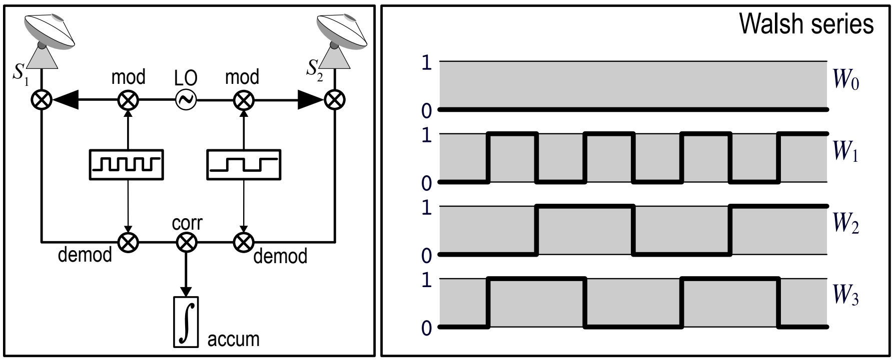

Figure 5.9: Schematic diagram of phase switching. The LO1 signal phase is modulated with the the Walsh series pattern and demodulated in the DTS before cross correlation. The Walsh function series are employed to generate switching patterns that are orthogonal to each other.

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

Table 5.5: Cross correlation of a baseline under 90° phase switching. The suffixes of $1, 2, U, L$ stand for antenna 1 and 2, and USB and LSB, respectively. The LO phase of antenna 2 is modulated by 90°. After 90°-demodulation and correlation, the cross correlations of USB and LSB have opposite signs. Thus, the average and the difference of cross correlations in two states gives the USB and the LSB visibilities, respectively.

5.6 Practical Performance

5.6.1 Sensitivity

A digitizer adds quantization noise to its input analog signal, with a consequent signal-to-noise reduction or sensitivity loss. The ALMA digitizer employs 3-bit (8-level) quantization[13], and additional re-quantization processes are applied in the correlators.

In Cycle 5, the 64-input Correlator TDM mode makes use of 2-bit quantization that yields a quantization efficiency of $\eta_{\rm Q} = 0.88$. In the FDM case, digital filtering is applied to 3-bit quantized signals followed by 2-bit re-quantization. The two stages of quantization would cause the net quantization efficiency of $\eta_{\rm Q} = 0.96 \times 0.88 = 0.85$ if they were independent (see mode chart tables of Escoffier et al., ALMA Memo 556), nevertheless, the real efficiency can be slightly better than the multiplication (Iguchi et al. 2005, PASJ 57, 259)[14].

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

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

5.6.2 Linearity

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

Figure 5.10 shows the Van Vleck relationships in 2-, 3-, and 4-bit quantizations. The correlation coefficients of quantized signals keep adequate linearity for small correlation coefficients[15]. This indicates that for most sources the cross-correlation response is effectively linear.

However, auto-correlations are affected by the non-linearity because the correlation coefficient at zero lag must be unity by definition. The non-linearity in auto-correlation power spectra influences system temperature measurements and the TP Array observations. In both correlators, the CDPs measure total power of digitized signals and determine the Van Vleck relationship to correct the (cross) power spectra.

5.7 Final data product - the ASDM

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

In the ASDM, the binary data are saved in a structure of spws. All of the data in a single spw must share the same frequency setup, including the number and width of spectral channels, and the integration time. The observed spws will be a combination of the science spectral windows set up by the proposers in

---

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

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

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

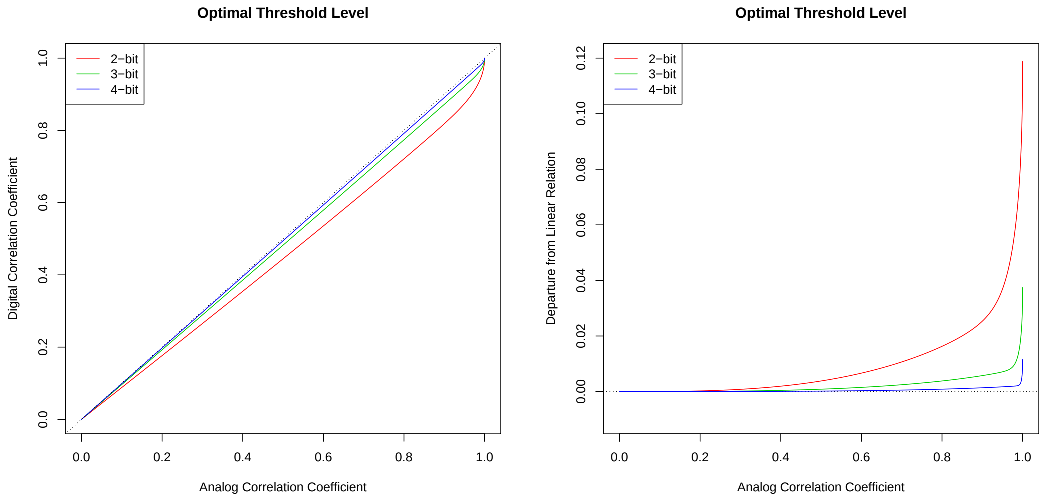

Figure 5.10: Van Vleck relationship of 2-, 3-, and 4-bit quantizations at the optimal signal power level and threshold voltages. (Left): Correlation coefficient of digitized signals as a function of analog correlation coefficients, ρ. (Right): Departures of the Van Vleck relationship from linear relation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

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 online telescope calibration system (TELCAL) and for realtime diagnostic purposes (QuickLook), and are typically not used downstream in the data reduction. Overall, this can lead to ASDMs with a large number of spectral windows. For example, a typical science observation in FDM mode can have more than 25 spectral windows. Every scan/spw combination has an “intent” associated with it that indicates its purpose (pointing, system temperature, science, etc). The intent can be used in CASA to decode how to utilize each spectral window in the data reduction process.

Chapter 6

Spectral Setups

This chapter describes the frequency setup of ALMA. In particular it shows how spectral setups are defined from the user viewpoint and how these then set up the ALMA hardware. It summarises the functions and operation of the local oscillator (LO) and Intermediate Frequency (IF) systems, as well as outlining some of the other points pertaining to the frequency setup. For those interested in more details of the LO and IF components and how the hardware works, please jump to Appendix B. The main changes for Cycle 5 are the addition of the Band 5 spectral coverage, and the inclusion of 90-degree phase switching.

To the ALMA system, a spectral setup includes the hardware LO, IF and correlator settings, such that each *Spectral Window* (spw) covers the desired lines and/or continuum frequencies. To the end-user, however, the spectral setup is normally defined in the Observing Tool (OT) just in terms of the desired lines or observing frequencies, spectral window bandwidths and spectral resolutions: there is no need for the user to worry about the details of each hardware setting. For full details of the OT and how to use it, see the user manual and reference manuals, available from the ALMA website[1] (and also in the OT itself). In this chapter, we try to reconcile these two viewpoints.

6.1 Introduction

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

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

Using the OT during both the proposal preparation (Phase 1) and scheduling block (SB) creation (Phase 2), the user chooses the frequencies (or spectral lines) to be observed in each spw and their required bandwidths or spectral resolutions. During the spectral tuning, there are effectively four different LOs set up by the system (Figure 6.1):

- LO1 which sets the frontend tuning frequency. This has continuous coverage.

[1] http://almascience.org/documents-and-tools/

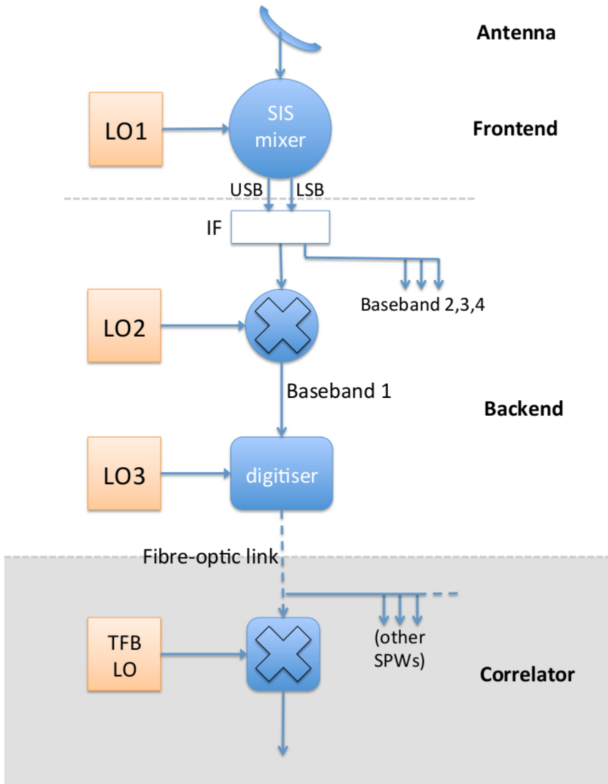

Figure 6.1: Summary block diagram of ALMA signal path and LO system for one spectral window, one baseband, one polarisation. The Backend IF system is located in each antenna, and the digitised signal for each baseband/polarisation is fed to the correlator in the central technical building through buried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

- LO2 which positions the basebands within the frontend receiver IF (each baseband uses a different LO2, hence there are four LO2s). This does not have continuous frequency coverage, as it is generated by a 125 MHz frequency comb plus a 20-42.5 MHz continuous frequency synthesizer.
- LO3 which is the clock frequency of the digitizers (fixed at 4 GHz),
- LO4 (also known as the tunable filterbank LO, or TFB LO), which is a digital LO synthesized in the correlators allowing positioning of the spectral windows within each baseband. Each spw has effectively a separate LO4.

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

Cycle 5 has similar restrictions on the spectral setups and tuning as Cycle 4 (see Table 6.2 in Section 6.8).

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

The main limits are that the edges of the 2 GHz-wide basebands and cannot lie outside the receiver tuning range listed in Table 4.1, and the edges of the (untruncated) spectral windows cannot lie outside the 2 GHz-wide basebands. The settings of each baseband in the correlator are independent, so the bandwidth, resolution and correlator mode as well as the value of LO4 can be different for each baseband. For example, it is possible to have a 58 MHz-wide spw centered on a particular line in baseband 1, and simultaneously use a broadband time division multiplexing mode (or TDM mode - see correlator chapter) for all the other basebands. It is possible to have multiple FDM spectral windows *within* each baseband, described in Section 6.3.2 below; however, currently the underlying resolution of all spws *within the same baseband* has to be equal. The spectral setups of the ACA Correlator and the 64-input Correlator are - to the end user - effectively the same, with the ACA Correlator having the same allowable spectral functions as the 64-input Correlator.

6.2 Frequency Definitions

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

Center Frequency (Rest) (f_{SPW}) This is set by the user in the OT Phase 1, and is the central frequency of the spw in the requested rest frame of the source. Note that the source rest frame is selectable, but is commonly set to the local kinematic standard of rest (LSRK), which is the conventional local standard of rest based on the average motion of the Sun with respect to the solar neighborhood (see Section 6.7).

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

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

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

$$f_{spw} = f_{BB} \pm (f_{offset} - 3.0GHz) \tag{6.1}$$

where the sign depends on the observing sideband.

6.3 Spectral Setups

The wide IF bandwidth and tuning ability allows for routine simultaneous imaging of multiple lines. Some examples (with the approximate line frequencies in GHz) are shown in Table 6.1 (for a source of redshift zero). Note that in many cases the lines will not necessarily appear in the center of the spws (for example, in the Band 6 combination given). When <4 spectral windows are required for the primary lines, the others can be set up to cover fainter lines or to observe the continuum, potentially in TDM mode. The continuum spw frequencies need to be setup manually, and one method of doing this is described in the next section. The selection of secondary lines to be observed can be done using the OT spectral interface. In the case of continuum spws, to maximize the sensitivity, the widest bandwidth mode should be chosen (i.e. in TDM, or 1.875 GHz in FDM - preferably with some channel-averaging, see Section 6.4). Also the continuum spws should cover as much of the IF band as possible. Not only will this maximize the continuum SNR on the science target, but these continuum data can be used to improve phase and amplitude calibration. Note that there is no gain in sensitivity by overlapping spws in frequency, since the input signal is the same in the overlap region; the OT uses the unique aggregate bandwidth in calculating the sensitivity, taking any spw overlap into account.

6.3.1 Observing Frequencies for Continuum

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

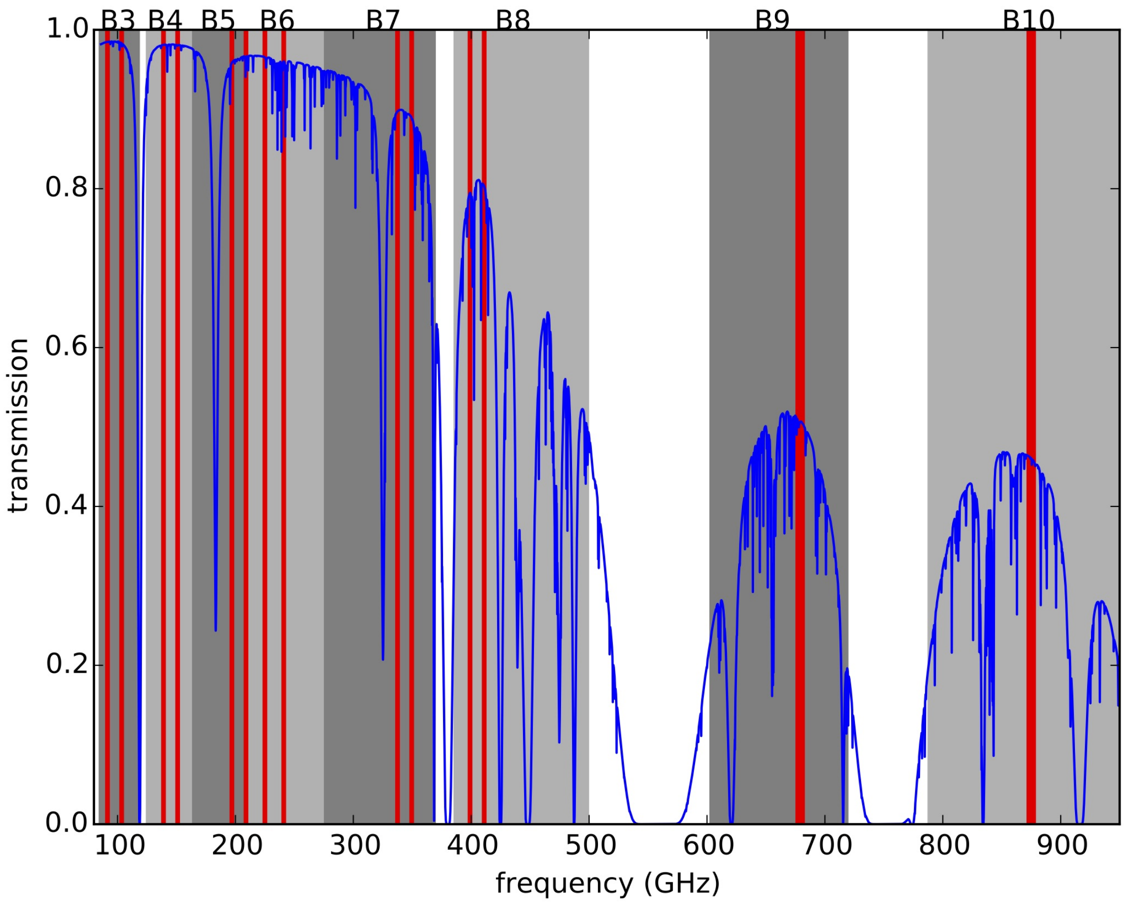

Figure 6.2: ALMA bands available in Cycle 5 (Bands 3 - 10), showing the frequency of the standard continuum settings as red shading. This gives the coverage of both USB and LSB, except for Bands 9 and 10, which have 8 GHz of bandwidth using the USB only. The available frequency coverage of each band is shown by the grey shading, and atmospheric transmission for 0.6 mm of PWV is given by the blue line.

If mixed line+continuum operation is desired, for example with a single line in BB1 observed at frequency f (GHz), the other BBs (2-4) can be set up for continuum observing using TDM mode[3]. For 2SB receivers, normally this would have BB2 in the same sideband (offset by 2 GHz from f) and BB3 and BB4 in the opposite sideband (offset by the front end center IF frequency, f_{IF}). To maximise sensitivity the continuum BBs should cover the maximum aggregate unique bandwidth, i.e. they should not overlap in frequency. If sb is the sign of the sideband of BB1 ($sb = +1$ for USB, $sb = -1$ for LSB), then the four BBs could be set up as follows:

- BB1: f *(BB covering the primary line)*

[3]Note - this case assumes one spw per BB

Band/setup	**Species/transition**	**Freq. (GHz)**	**Sideb.**	**bandwidth**	**BB**	**spw**	**Notes**
3 a	**Standard cont.**	97.5[1]	dual	TDM			LO1=97.5
3 b *	HCO^+ 1-0	89.188	LSB	58 MHz	1	1	HCO^+/HCN/H_2CO
-	HCN 1-0	88.632	LSB	58 MHz	2	1	
-	CH_3OH/H_2CO	101.293/101.333	USB	125 MHz	3	1	
-	continuum spw	101.3	USB	TDM	4	1	
3 c *	CO 1-0	115.271	USB	58 MHz	1	1	CO/CN/$C^{17}O$
-	CN	113.499	USB	58 MHz	1	2	
-	$C^{17}O$ 1-0	112.359	USB	117 MHz	2	1	
-	continuum spw	102.5	LSB	TDM	3	1	
-	continuum spw	100.5	LSB	TDM	4	1	
4 a	**Standard cont.**	145.0[1]	dual	TDM			LO1=145.0
4 b	CS 3-2	146.969	LSB	117 MHz	1	1	CS/DCO^+
-	DCO^+ 2-1	144.077	LSB	117 MHz	2	1	
-	SO2 4(2,2)-4(1,3)	146.605	LSB	117 MHz	3	1	
-	H_2CO 2(0,2)-1(0,1)	145.603	LSB	117 MHz	4	1	
5 a	**Standard cont.**	203.0	dual	TDM			LO1=203.0
5 b	H_2O 3(1,3)-2(2,0)	183.31	USB	250 MHz	1	1	
-	continuum spw	181.4	USB	TDM	2	1	
-	H_2S 1(1,0)-1(0,1)	168.763	LSB	250 MHz	3	1	
-	continuum spw	170.8	LSB	TDM	4	1	
6 a	**Standard cont.**	233.0[1]	dual	TDM			LO1=233.0
6 b *	^{12}CO 2-1	230.538	USB	117 MHz	1	1	J=2-1 CO isotopes
-	continuum spw	231.6	USB	TDM	2	1	
-	$C^{18}O$ 2-1	219.560	LSB	117 MHz	3	1	
-	^{13}CO 2-1	220.399	LSB	117 MHz	3	2	
-	continuum spw	218.5	LSB	TDM	4	1	
7 a	**Standard cont.**	343.5[1]	dual	TDM			LO1=343.5
7 b *	continuum spw	343.0	LSB	TDM	1	1	CO/HCO^+/HCN
-	^{12}CO 3-2	345.796	LSB	469 MHz	2	1	1/2-corr
-	$HC^{15}N$	344.200	LSB	234 MHz	2	2	1/4-corr
-	$H^{13}CN$	345.340	LSB	234 MHz	2	3	1/4-corr
-	HCN 4-3	354.505	USB	469 MHz	3	1	
-	HCO^+ 4-3	356.734	USB	117 MHz	4	1	1/4-corr
-	NHD_2 2(2,0)-2(1,2)	356.230	USB	234 MHz	4	2	1/2-corr
7 c	^{12}CO 3-2	345.796	USB	58 MHz	1	1	J=3-2 CO/^{13}CO
-	continuum spw	344.8	USB	TDM	2	1	
-	^{13}CO 3-2	330.588	LSB	58 MHz	3	1	
-	continuum spw	331.6	LSB	TDM	4	1	
8 a	**Standard cont.**	405.0[1]	dual	TDM			LO1=405.0
8 b	CI 3P1-3P0	492.160	USB	117 MHz	1	1	CI/^{13}CI
-	CS 10-9	489.751	USB	117 MHz	2	1	
-	continuum spw	477.5	LSB	TDM	3	1	
-	continuum spw	479.5	LSB	TDM	4	1	
9 a	**Standard cont.**	679.0	USB	TDM			LO1=671.0
9 b	^{12}CO 6-5	691.472	USB	469 MHz	1	1	CO/CS
-	CS 14-13	685.436	USB	117 MHz	2	1	
-	H_2S	687.303	USB	234 MHz	3	1	
-	$C^{17}O$ 6-5	674.009	LSB	469 MHz	4	1	
10 a	**Standard cont.**	875.0	USB	TDM			LO1=867.0
10 b *	CO 7-6	806.652	LSB	469 MHz	1	1	CO/HCO^+
-	HCO^+ 9-8	802.458	LSB	117 MHz	2	1	
-	continuum spw	805.5	USB	TDM	3	1	
-	continuum spw	803.65	USB	TDM	4	1	

Table 6.1: Examples of spectral setups possible in Cycle 5. This includes the standard continuum-only setups for each band, and some multiple line/continuum configurations. *Notes:* **1.** Frequency for standard continuum setups are the mean observing frequency of all spws, ≈LO1 in dual-sideband (2SB) receivers, so this frequency is not actually covered if both sidebands are used. See Figure 6.2. * Based on a template spectral setup released with Cycle 5 version of the OT.

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

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

6.3.2 Multiple Spectral Windows in the Same Baseband

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

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

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

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

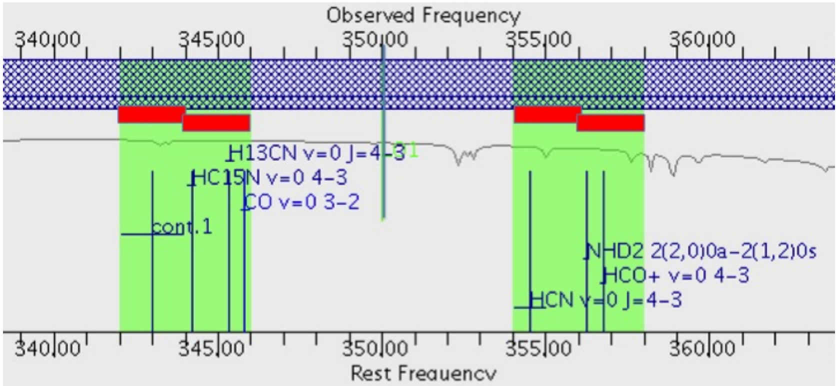

Figure 6.3: Illustration of a frequency setup, based on the OT spectral display. Green areas are the IF ranges, horizontal red bars are the 2 GHz-wide basebands (BB1-BB4), and smaller horizontal lines represent the spectral windows (up to 3 per baseband in this example). The frequency of LO1 is shown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

6.3.3 Spectral Setups for Lines near the Edge of the Bands

The restriction that the baseband coverage cannot fall outside the maximum or minimum tuning range of the receiver is an issue for certain lines at the edge of the tuning range. One obvious example is the 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 6.1). A setup using the wide bandwidth modes (TDM or FDM/1875 MHz) centered at the line frequency with zero redshift will not validate in the OT, because some of the basebands will fall outside the maximum Band 3 frequency (116.0 GHz). Narrower modes will not go over the edge and are allowed.

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

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

Bands 9 and 10 use double-sideband (DSB) receivers, so the image sideband can either be suppressed using LO offsetting (currently using LO1 and LO2), or both sidebands can simultaneously be recorded using 90-degree phase switching. These can only be used in interferometric mode - see Section B.4.2.

Observing a single sideband

The choice of which sideband a particular baseband is configured to observe (and which sideband is suppressed) depends only on the relative sign of the LO1 and LO2 offset. This can be arbitrarily different for different basebands. So it is possible to set up two basebands at approximately the same LO1/LO2 frequency, but observe one line in the upper sideband in one baseband, and another in the lower sideband using the next baseband, just by having different signs in the LO2 offsetting. However, in this case it might be better to use 90-degree switching (see next section) to record both using the same baseband. Note that different spws in the *same* baseband must have the same sideband, as image suppression uses LO2, which is common to the whole baseband.

Observing both idebands - 90-degree switching

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

6.4 Usable Bandwidth

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

6.5 Spectral Resolution and Channel-averaging

The spectral resolution of each spw is set by a combination of the inherent resolution/bandwidth setup in the correlator, the correlator weighting function (also known as the smoothing function) and the channel-averaging (also known as spectral-averaging) factor (which is carried out at a late stage of processing in the correlator). This is described in the correlator Section 5.5.2. The default FDM setup from the OT in Cycle 5 uses Hanning smoothing with a channel-averaging factor of 2 (see Section 5.5.2). This is changed from previous Cycles, which

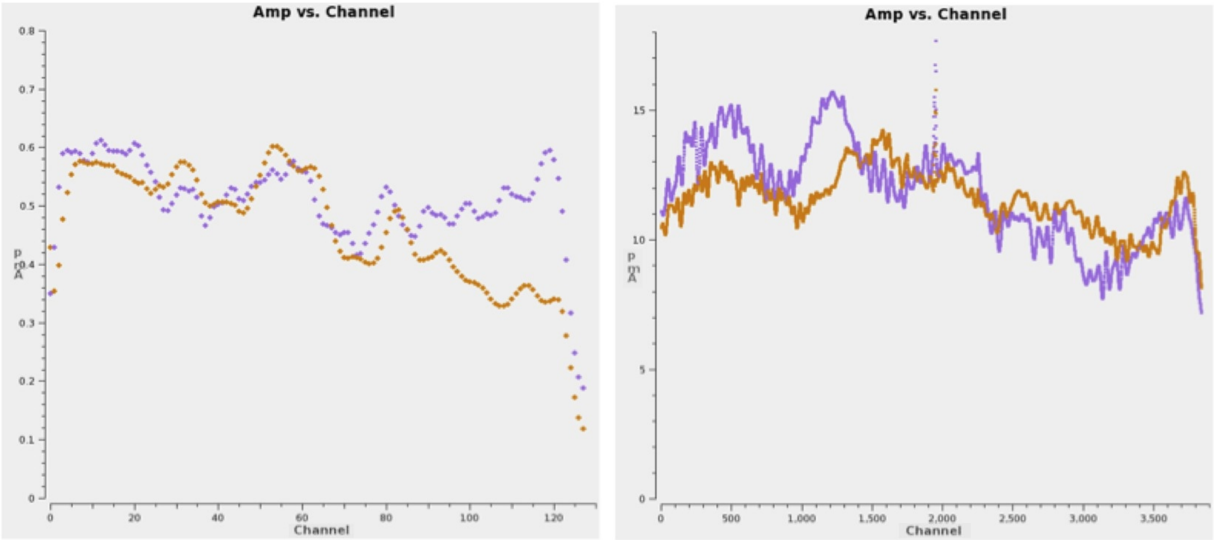

Figure 6.4: Comparison of TDM (left) and FDM (right) autocorrelation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in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

had a channel-average factor of 1, and has been introduced to reduce the size of the FDM datasets with negligible loss of spectral resolution. Hanning smoothing plus a channel-average of 2 results in a spectral resolution equal to the *final* channel spacing x 1.15 - only ∼15% worse than Hanning without channel-averaging but has the advantage of halving the data rate. If the line is well resolved, with ≥3 channels covering the linewidth, then this will not cause a problem It is recommended to keep this default, in order to reduce the data rate and size of the final product; this is more critical where fast dump times are needed, e.g. for long baseline observing. If the extra 15% resolution is required, a justification for using channel-average=1 should be made. Note that for TDM mode, Hanning smoothing without channel-averaging (1) is the default.

6.6 Spurious Signals

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

6.7 Doppler Setting and Velocity Reference Frames

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

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

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

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

If a target velocity and reference frame other than topocentric is selected in the OT, at the start of the observation, the velocity of the science target in the chosen reference frame *and* the velocity of the observatory relative to the chosen reference frame are used to set the center frequencies of the science spws[5]. For example, for a source of radial velocity 100 km/s in the LSRK frame, the difference between rest and sky frequency in the OT will be the equivalent of 100 km/s. At runtime, the extra velocity from the Earth orbit and rotation (up to ±30.3 km/s and ±0.46 km/s), plus the motion of the LSRK (20.0 km/s toward RA=18 hrs, $\delta = 30°$ for epoch 1900.0) will be included. This means that the LO solutions are slightly different at runtime compared with those from the OT, although this is normally transparent to the user[6]. It is also possible to just make the correction for the motion of the observatory relative to the chosen reference frame, ignoring the source velocity in that frame. The frequency setting for the science target will normally also be used in the same execution for both the bandpass and amplitude calibration.

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

The user can specify either the velocity or the redshift of the target. Note that velocities can be defined in three different ways, resulting in different conversions to frequency. These differences only become significant at high velocities. The velocity formulae are:

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

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

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

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

6.8 Limitations and Rules for Spectral Setups in Cycle 5

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

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

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

Rule	Description
1.	LO1 must lie within the LO tuning ranges given in Table 6.1.
2.	No part of the 2.0 GHz-wide basebands can extend over the edge of the IF passband. So, the baseband centers cannot be closer than 1.0 GHz to the IF passband edge. For example, with a 4.0-8.0 GHz IF range, the baseband center frequency must lie between 5.0-7.0 GHz. The system actually does allow a small extension of the edges of the basebands over the IF edges, to cope with the differential Doppler shifts in the different basebands, but this is small ($<$ 1 MHz at Band 6) and is transparent to the OT user (see Section 6.7).
3.	For 2SB receivers (Bands 3-8), the number of basebands in one sideband can only be 0, 1, 2 or 4 (i.e. not 3). For DSB receivers (Bands 9 and 10), there is no such restriction (the number can be 0, 1, 2, 3 or 4).
4.	For $>$ 2 basebands, BB1 and BB2 should be in the same sideband, as should BB3 and BB4.
5.	No part of the full nominal bandwidth of the spw can extend over the edge of the 2 GHz-wide baseband. For a mode with nominal bandwidth B (e.g. 62.5 MHz), that means the spw center IF frequency (a.k.a. Center Offset Frequency in the OT, Phase 2) must be $>$(2000+B/2) and $<$(4000-B/2) MHz. The current version of the OT forces this restriction. For 2 GHz FDM and TDM modes, this means that the spw must be at the center of the baseband. However there is a further restriction on this, as noted in the next rule.
6.	The spw usable bandwidth (ie 15/16 of the nominal spw bandwidth of a multiple of 62.5 MHz) should be in an allowed region of the baseband. This is in addition to (5). In practice this means that the required range of the spw should normally be inside the range $\sim$2050 - 3950 MHz (i.e. $>$50 MHz from the edges of the 2-4 GHz second IF) to ensure that the edge of the anti-aliasing filter does not significantly affect the IF power. In practice it is possible to extend some fraction of the spw to $<$50 MHz of the IF edges; this is necessary for the Band 6 setup 6 b in Table 6.1, where the ^{12}CO and ^{13}CO lines are only 70 MHz from the IF band edges.
7.	The line frequency does not need to be in the center of the spw in Phase 1 of the OT: if a line like ^{12}CO 1-0 is requested, it will generate an SB with the correct TFB LO offset, as long as rule 5 is obeyed (which may require using a narrower spw). This is mentioned in more detail in Section 6.2
8.	Only 2-bit, Nyquist sampling is allowed in the correlator.
9.	It is possible to have multiple targets with different redshifts within the same Science Goal in the OT. For SGs including sources with more than one redshift, all the observations must be achievable using five or fewer tunings within the same receiver band, considering the source redshifts and, in the case of spectral lines, the line widths and configuration of spectral windows.
10.	The number of spws per baseband can be 1, 2, 3 or 4. For 3 spws, the correlator resources per spw should be set to 1/4, so only 3/4 of the available resources are used in this case.
11.	All the spws *within the same baseband* are required to have the same correlator channel width (before channel-averaging). An individual spw within a baseband may occupy 1, 1/2 or 1/4 of the resources available in the baseband and the sum of the fractional resources within one baseband must be $\leq$1. The correlator resources are proportional to the number of correlator spectral channels.
12.	With DSB receivers (Band 9 and 10), 90-degree phase switching is allowed, giving both sidebands simultaneously. This is only available for TDM and FDM 1875 MHz bandwidth modes.

Table 6.2: Rules for spectral setups.

FWHM

Chapter 7

Imaging with ALMA

7.1 Introduction

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

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

An ALMA observation can be represented by N discrete points in the uv-plane $\mathcal{V}_k(u_k, v_k)$ where $k = 1, \ldots, N$ and N can reach a few million. In addition to the visibilities $\mathcal{V}_k$ obtained, a weighted visibility function $\mathcal{V}^W$ can be defined (see Briggs, Schwab & Sramek, 1999) that enables the resulting synthesized beam size and the sensitivity to be controlled:

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

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

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

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

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

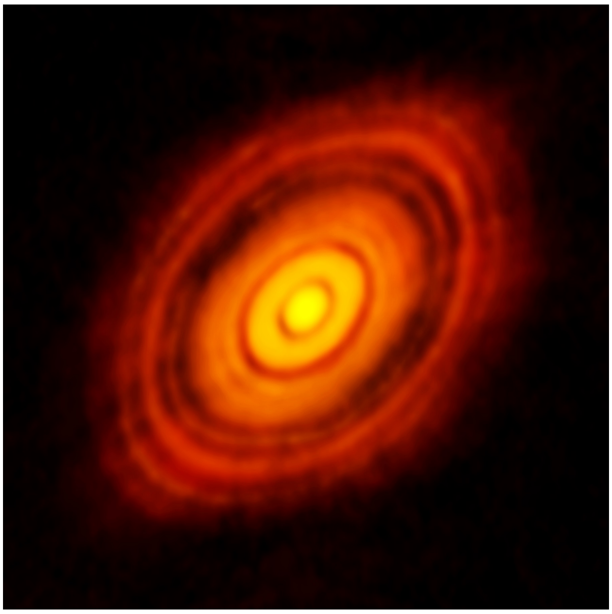

Figure 7.1: This is one of the sharpest images ever taken by ALMA. It shows the protoplanetary disc surrounding the young star HL Tauri. These new ALMA observations reveal substructures within the disc that have never been seen before and even show the possible positions of planets forming in the dark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

7.2 Cycle 5 Configurations

In Cycle 5, depending on the range of angular scales required, an ALMA project will obtain data from one or two 12-m Array configurations, and potentially also from the 7-m Array. If yet larger angular scale information is required, single-dish data from the Total Power (TP) Array will be obtained as well when possible (for instance, the TP Array does not currently support continuum science observing). Imaging consists of combining the interferometric data together using an appropriate weighting function; de-convolving the interferometer synthesized beam (the point spread function (PSF), or also called "dirty beam") which results from finite sampling of the *uv*-plane; and, if TP Array data exist, gridding them into an image cube and combining it with the de-convolved interferometer image cube using the feathering technique. We describe in this chapter some of the aspects of imaging that should be considered to obtain the best results from ALMA observations.

During Cycle 5, the 12-m Array - with (at least) 43 antennas - will be arranged in ten different configurations. The maximum baselines achievable will be 16.2 km for Bands 3-6, 8.5 km for Band 7 and 3.6 km for Bands 8-10. The 7-m Array, with (at least) ten antennas, will be available in only one configuration. Four extended configurations will be available in Cycle 5: C43-7, C43-8, C43-9, and C43-10. The most extended configuration - C43-10 - will be offered alone; it cannot be combined with other configurations. The other three extended configurations - C43-7, C43-8 and C43-9 - are offered in combination with other configurations, as in previous cycles. Tables 7.1 and 7.2 give the basic properties of the 12-m Array and 7-m Array configurations. In addition, Figures 7.2, 7.3 and 7.4 show antenna locations for the configurations of the 12- and 7-m Arrays. The three TP Array single-dish 12 m antennas will be in fixed positions. **It should be noted that these configurations (especially for the 12-m Array configurations) and the corresponding resolutions/max angular scales are representative of the actual configurations that will be in use.** A variety of factors including, but not limited to, the existence of hybrid configurations during antenna reconfiguration periods and the precise antenna pad availability schedule will likely lead to deviations from the values shown here. From the user's point of view configurations are not directly selected, instead the required angular resolution (θ_{res}) and the largest angular scale (θ_{LAS}) that need to be recovered are entered in the OT. Projects will then be observed using array configurations chosen to achieve the specified goals. More details about array combination are provided in Section 7.8.

	Band	3	4	5[*]	6	7	8	9	10
	Frequency (GHz)	100	150	185	230	345	460	650	870
Configuration									
7-m	θ_{res} (arcsec)	12.5	8.35	6.77	5.45	3.63	2.72	1.93	1.44
	θ_{MRS} (arcsec)	66.7	44.5	36.1	29.0	19.3	14.5	10.3	7.67
C43-1	θ_{res} (arcsec)	3.38	2.25	1.83	1.47	0.98	0.735	0.52	0.389
	θ_{MRS} (arcsec)	28.5	19.0	15.4	12.4	8.25	6.19	4.38	3.27
C43-2	θ_{res} (arcsec)	2.3	1.53	1.24	0.999	0.666	0.499	0.353	0.264
	θ_{MRS} (arcsec)	22.6	15.0	12.2	9.81	6.54	4.9	3.47	2.59
C43-3	θ_{res} (arcsec)	1.42	0.943	0.765	0.615	0.41	0.308	0.218	0.163
	θ_{MRS} (arcsec)	16.2	10.8	8.73	7.02	4.68	3.51	2.48	1.86
C43-4	θ_{res} (arcsec)	0.918	0.612	0.496	0.399	0.266	0.2	0.141	0.106
	θ_{MRS} (arcsec)	11.2	7.5	6.08	4.89	3.26	2.44	1.73	1.29
C43-5	θ_{res} (arcsec)	0.545	0.363	0.295	0.237	0.158	0.118	0.0838	0.0626
	θ_{MRS} (arcsec)	6.7	4.47	3.62	2.91	1.94	1.46	1.03	0.77
C43-6	θ_{res} (arcsec)	0.306	0.204	N/A	0.133	0.0887	0.0665	0.0471	0.0352
	θ_{MRS} (arcsec)	4.11	2.74	N/A	1.78	1.19	0.892	0.632	0.472
C43-7	θ_{res} (arcsec)	0.211	0.141	N/A	0.0917	0.0612	0.0459	0.0325	0.0243
	θ_{MRS} (arcsec)	2.58	1.72	N/A	1.12	0.749	0.562	0.398	0.297
C43-8	θ_{res} (arcsec)	0.096	0.064	N/A	0.0417	0.0278	-	-	-
	θ_{MRS} (arcsec)	1.42	0.947	N/A	0.618	0.412	-	-	-
C43-9	θ_{res} (arcsec)	0.057	0.038	N/A	0.0248	-	-	-	-
	θ_{MRS} (arcsec)	0.814	0.543	N/A	0.354	-	-	-	-
C43-10	θ_{res} (arcsec)	0.042	0.028	N/A	0.0183	-	-	-	-
	θ_{MRS} (arcsec)	0.496	0.331	N/A	0.216	-	-	-	-

Table 7.1: Resolution (θ_{res}) and maximum recoverable scale (θ_{MRS}) for the 7-m Array and 12-m Array configurations available during Cycle 5 as a function of a representative frequency in a band. The value of θ_{MRS} is computed using L05 from Table 7.2 and Equation 7.7; the value of θ_{res} is the mean size of the interferometric beam obtained through simulation with CASA, using Briggs uv-plane weighting with *robust=0.5*. (This value of *robust* offers a compromise between natural and uniform.) The computations were done for a source at zenith; for sources transiting at lower elevations, the North-South angular measures will increase proportional to 1/sin(ELEVATION). **Note:** Band 5 is planned to become available March 2018. Consequently, according to the current planning the long baselines configuration C43-6 to C43-10 will not be available for Band 5 in Cycle 5.

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

Table 7.2: Basic parameters of the 7-m Array configuration and the ten 12-m Array configurations offered during Cycle 5. The baselines are projected for a transiting source (HA = ± 1h) at a declination of $-23°$. Note that C43-8 will not be available for Bands 8-10, and C43-9 and C43-10 will not be available for Bands 7-10.

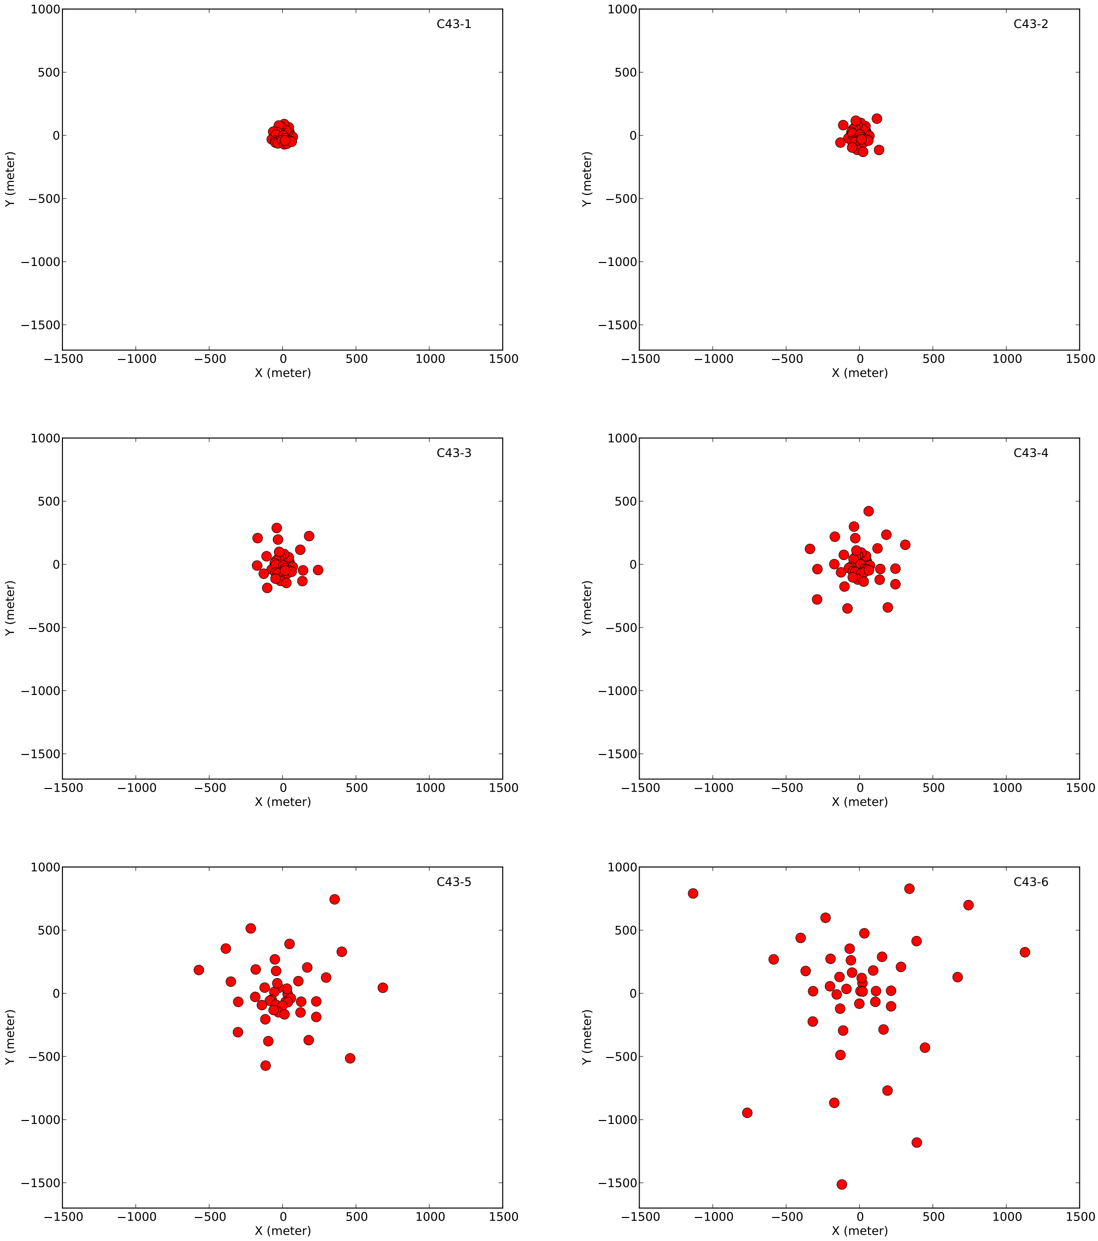

Figure 7.2: Representative 12-m Array compact configurations for Cycle 5.

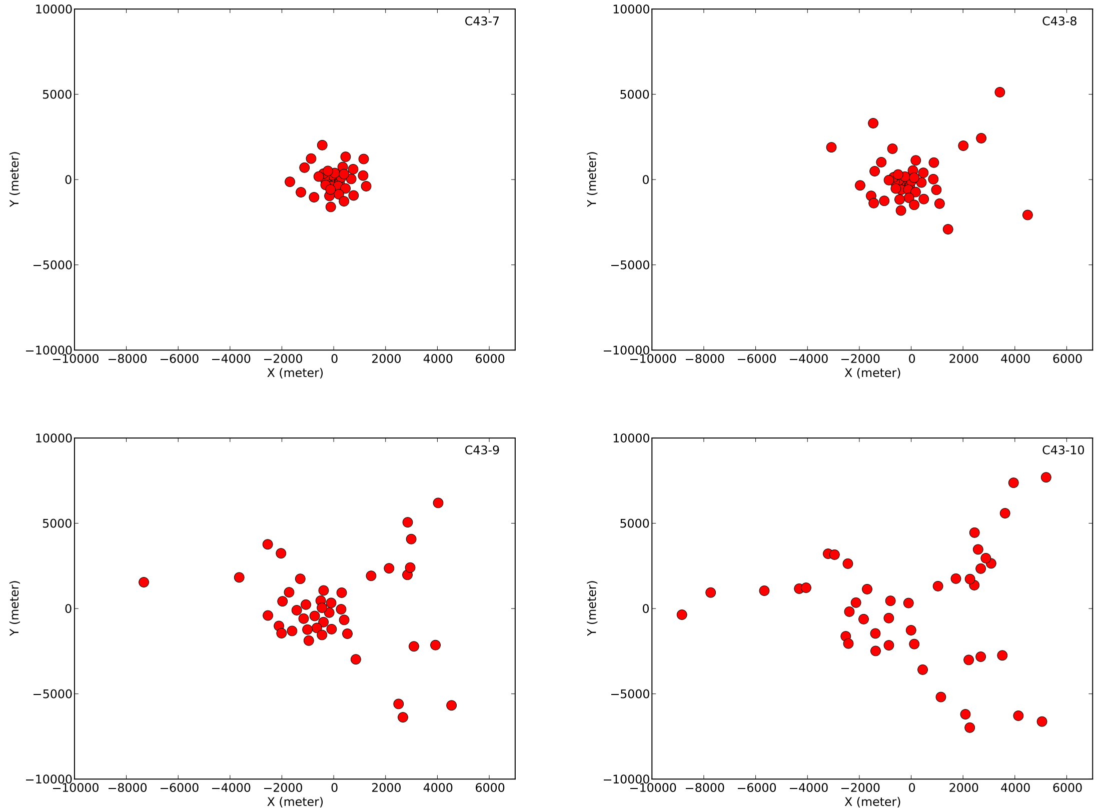

Figure 7.3: Representative 12-m Array extended configurations for Cycle 5.

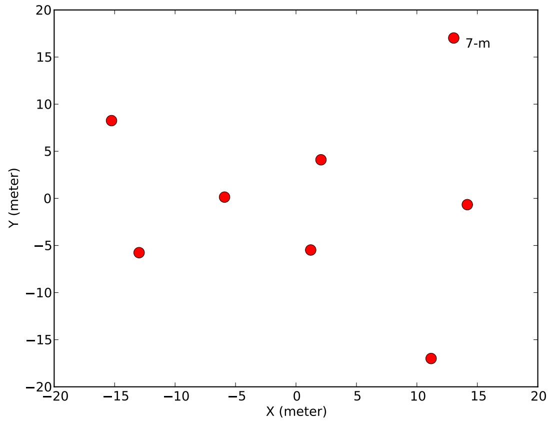

Figure 7.4: Representative 7-m Array configuration for Cycle 5.

7.3 Shadowing

ALMA is located at latitude=-23.02917°, longitude=-67.754649°. During observations, particularly in very compact configurations like those of the 7-m Array or C43-1, one antenna's view of the source can be partially blocked by another, making the data from that antenna unusable. This phenomenon is known as shadowing and can be quite severe for sources that are observed at very low elevation. Targets as far North as declination +47°, corresponding to a maximum source elevation at Chajnantor of ∼20°, can in principle be observed from the ALMA site, but shadowing by adjacent antennas becomes an increasing problem at low elevations. For example, Figure 7.5 shows the percentage of data that will be shadowed (the "shadowing fraction") when sources of various declinations are observed in the most compact ALMA configurations. As it can be seen, the shadowing fraction can be as large as 55% for sources observed with the most compact 12-m Array configuration (C43-1). The imaging capability, as well as the time on source, will necessarily be limited for such northern sources, especially at the higher frequencies. Shadowing depends on the antenna configuration. Given the short baselines in the ACA configuration, sources with declinations less than −70° or greater than +25° are subject to significant shadowing. For the 12-m Array, shadowing becomes significant (> 5 %) in the most compact configuration for sources with declination lower than −65° or higher than +20°.

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

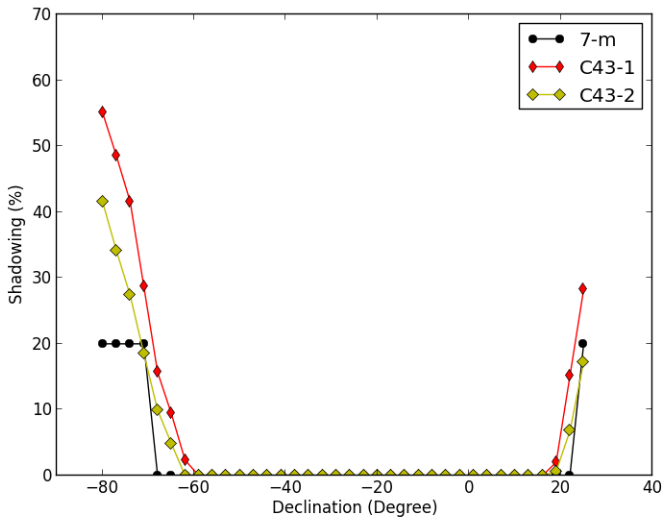

Figure 7.5: Shadowing fraction (%) for the most compact configurations as a function of declination (for a 1-hour observation at transit).

7.4 Resolution and Beam Shape

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

$$\theta_{res} = \frac{k\lambda}{L_{max}} \; [radians] \tag{7.3}$$

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

$$\theta_{res} \approx \frac{0.574\lambda}{L_{80}} \; [radians] \tag{7.4}$$

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

An important consideration when imaging is the uv-plane weighting scheme (k) used when Fourier transforming the visibilities. Natural weighting gives the highest-sensitivity image as each visibility is assigned a

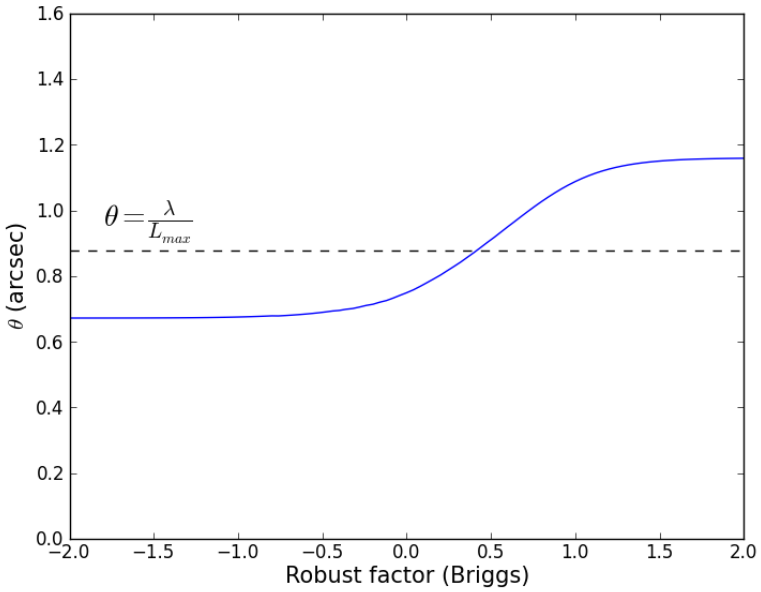

Figure 7.6: Angular resolution achieved using different values of the CASA *robust* parameter for a 1-hour observation at 100 GHz and a declination of -23 deg in the C43-4 configuration. Note that $robust = -2$ is close to uniform weighting and $robust = 2$ is close to natural weighting. The dotted line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

weight based on its intrinsic uncertainty. Uniform weighting, on the other hand, modifies the weights in such a way that each part of the *uv*-plane contributes equally to the final image. If this were not done, the shorter baselines would dominate as they have a higher density in the *uv*-plane. Therefore, uniform weighting degrades the sensitivity (the intrinsic weights are no longer used) but increases the angular resolution as the longer baselines have a higher weight.

To bridge the extremes of natural and uniform weighting, Briggs (1995) defined a continuous scheme that uses a "robustness" parameter R (this parameter is called *robust* in CASA). In CASA, uniform weighting is close to $robust = -2$ and natural weighting is close to $robust = 2$. Figure 7.6 shows the angular resolution achieved for an observation with the C43-4 configuration at 100 GHz using *robust* values between -2 and 2. As can be seen, the angular resolution varies from $0.7''$ ($robust = -2$) to $1.2''$ ($robust = 2$). The angular resolutions presented here for ALMA were all computed using CASA simulations with Briggs weighting and $robust = 0.5$.

The synthesized beam shape, which is the Fourier transform of the *uv*-plane sampling during the observation(s), is a function of the source declination. In addition to shadowing, sources that must be observed at low elevations also have shorter *projected* North-South baselines and thus the shape of the *uv*-plane sampling distribution becomes more elongated and the beam shape more elliptical. For example, Figure 7.7 shows the different beam shapes for sources observed at declinations of $-70°$ (similar to the SMC and LMC) and $-30°$. Also, Figure 7.8 shows the *uv*-coverage for these sources, revealing the elongation of the *uv*-plane sampling distribution, together with the large fraction of shadowing, for the more southern sources. To mitigate this effect, one can convolve the resulting images to a more circular (albeit larger) beam after deconvolution.

Figure 7.9 shows the minor and major axes of the synthesized beam widths (θ_{res}) for each array configuration as a function of source declination for a 2-hour observation at 100 GHz. Figure 7.10 shows the geometrical mean of the major and minor axes of θ_{res} at the same frequency. The beam width scales with λ, but bear in mind that not all configurations can be used with the higher-frequency bands, e.g., C43-8 is not available at Bands 8-10.

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

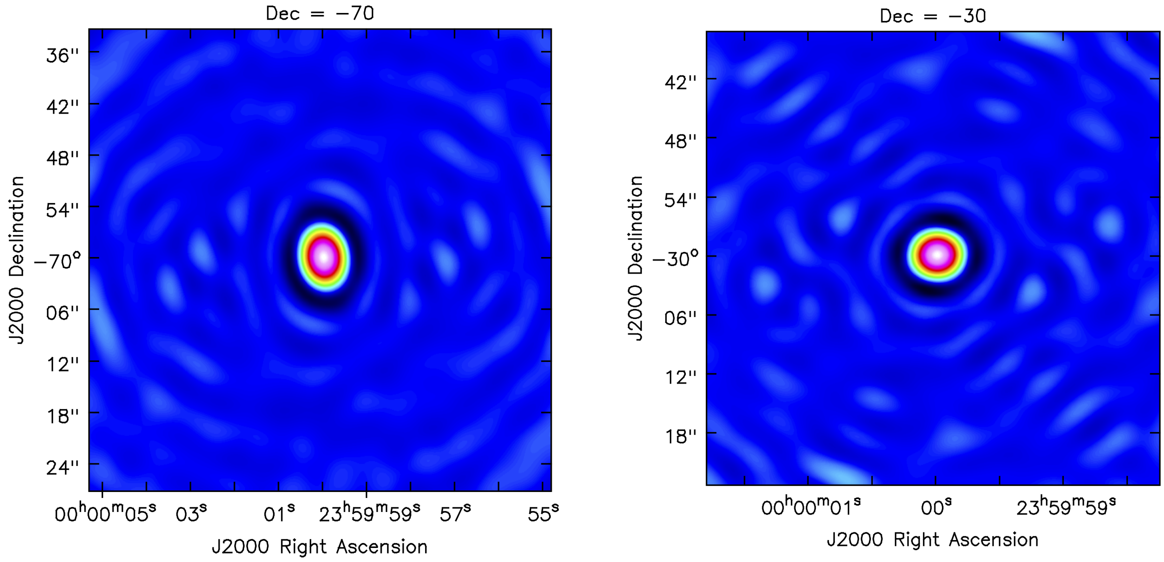

Figure 7.7: Beam shape for configuration C43-1 with a 2-hour observation of a transiting source at a declination of either −70° (*left*) or −30° (*right*).

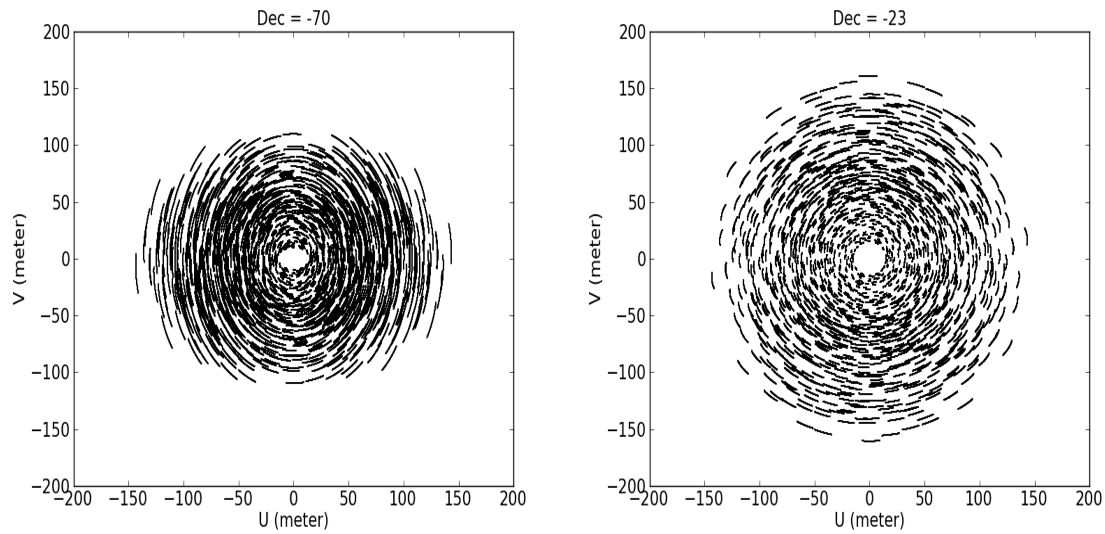

Figure 7.8: *uv*-plane coverage for configuration C43-1 with a 1-hour observation of a transiting source at a declination of either −70° (*left*) or −23° (*right*). For the source with a declination of −70°, the shadowing fraction represents 25.7% of the total.

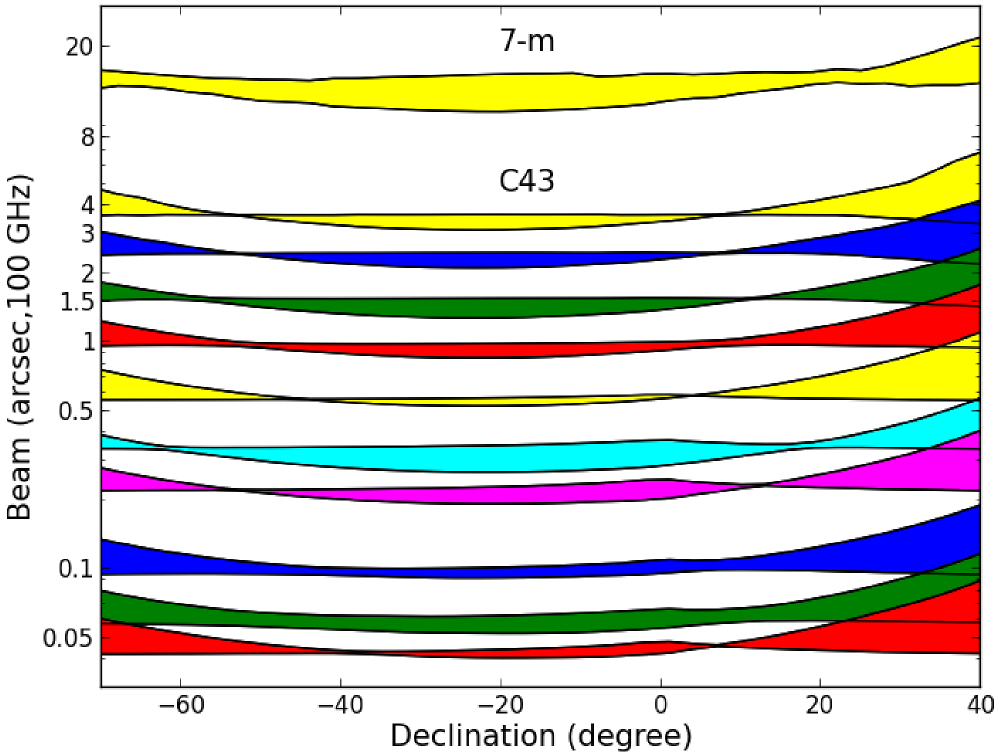

Figure 7.9: Cycle 5 angular resolutions as a function of source declination. Each color shows, for a particular configuration, the range of the major and minor axes of the synthesized beams expected from a 2-hour observation at 100 GHz and at transit. 12-m Array configurations are arranged with C43-1 at the top and C43-10 at the bottom.

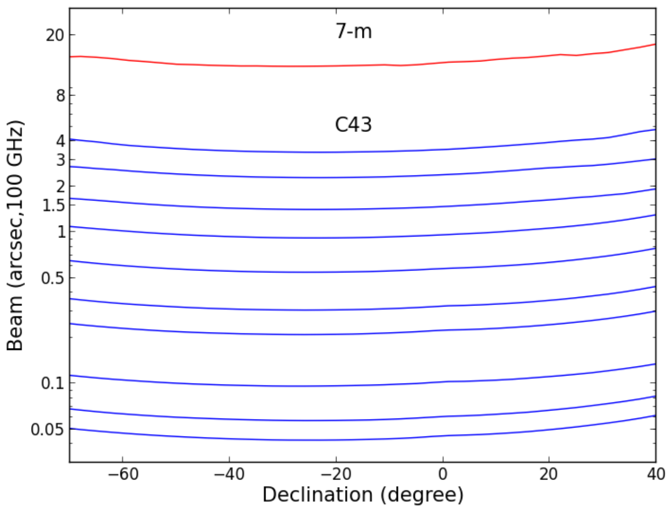

Figure 7.10: Geometrical mean of the major and minor axes of the synthesized beams as a function of source declination. These correspond to a 2-hour observation at 100 GHz and at transit. 12-m Array configurations are arranged with C43-1 at the top and C43-10 at the bottom. The 7-m Array is also shown.

determined based on the goals of the scientific case (required angular resolution, sensitivity, etc). In cases where more precise information is required, CASA simulations using the representative configurations, actual target declination, *uv*-coverage, and clean[1] parameters can be used to predict the resolution and sensitivity for a specific science use case. One should clean the images carefully, and also add noise (and realistic systematic errors, if possible) to simulations, to see what kind of residual sidelobes are expected. (Cleaning perfect data does not show the limitations that one might obtain with real data.)

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

7.5 Response and Snapshot

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

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

If S were a continuous function, e.g., a Gaussian, the synthesized beam would also be a Gaussian i.e., a central peak with a smoothly decreasing response away from the center. Given the finite number of baselines, however, the sampling function is an ensemble of Dirac functions and the Fourier transform of this produces a central peak surrounded by a complex pattern of sidelobes. This "dirty beam" response is a consequence of the gaps in the *uv*-plane, the sidelobes becoming increasingly prominent as the gaps increase (Gibbs phenomenon). As the projected baselines (i.e., the baseline lengths and orientations as seen from the source) change as a function of time as the Earth rotates, the longer a source is observed for, the greater the *uv* coverage and the smaller the sidelobes. Short integrations are still valuable, particularly when using many of the fifty 12-m Array antennas, however, if angular structure and sensitivity are less important. Such integrations are sometimes called *snapshots* , and they produce obvious sidelobes which can be mitigated by applying a *uv-taper* , a down-weighting of the visibilities on longer baselines, during the Fourier transform. Table 7.3 gives the level of sidelobes for each configuration, for a 1-hour observation of a source at Dec=$-23°$.

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

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

Figures 7.11 and 7.12 show examples of a *uv*-plane sampling distribution and cleaned image for a snapshot of 1-minute duration and a longer integration of 1-hour duration expected with the configuration C43-1. As shown in Figure 7.11, the *uv*-coverage in 1 minute is quite uniformly sampled, but much less dense than that of the 1-hour integration. In addition, the cleaned images of the snapshot and the longer integration (see Figure 7.12)

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

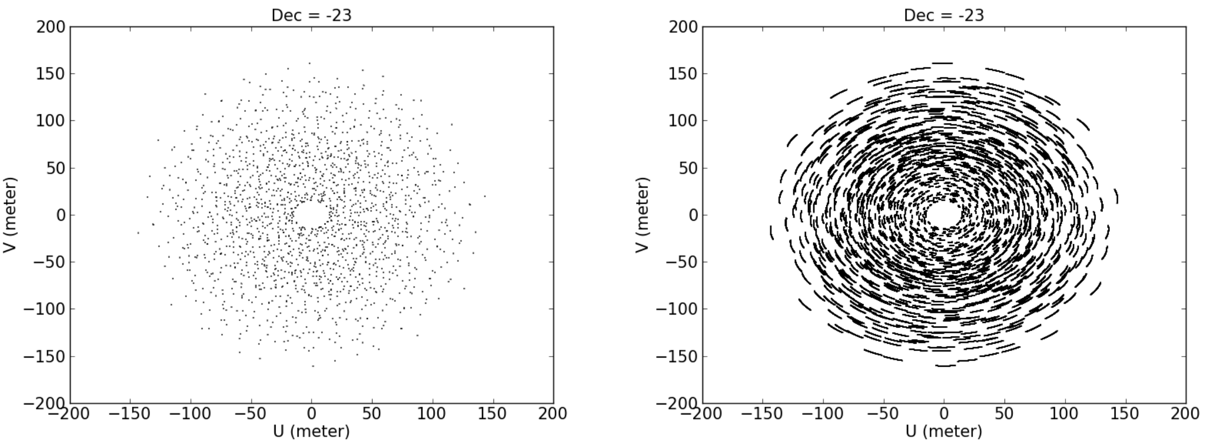

Figure 7.11: *uv*-plane sampling distributions of a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the C43-1 configuration to observe a source at a declination of −23°.

are similar in terms of angular resolution and the apparent differences are quite small. The main difference between the two (besides sensitivity) is in the dirty beam: the sidelobe level for the snapshot will be much higher than for the 1-hour integration. The snapshot image will then require more careful cleaning in order to avoid introducing spurious sources from the strong sidelobes. Condon et al. (1998) gave a very comprehensive description of how this issue impacted the 20-cm NRAO VLA Sky Survey (NVSS).

A useful practice to disentangle sidelobe effects from real point sources, especially with relatively strong point sources, is to perform CASA simulations using a component list of the strongest sources together with the actual array configuration. This test allows one to estimate the sidelobe fingerprint left by the strong point sources after deconvolution (cleaning), although the residuals will likely be higher in practice due to imperfect calibration. It is worthwhile to use the interactive mode during the deconvolution so that residuals can be monitored.

7.6 Spatial Scale Filtering

As described in Chapter 3, an interferometer measures the Fourier components of the sky brightness distribution in an area of the *uv*-plane defined by the array configuration used to observe the target. Since the *uv*-plane is a representation of the power of the sky brightness distribution per angular scale, an interferometer observation acts like a passband filter in the spatial domain. An astronomical source is then filtered if its Fourier transform has substantial power on angular scales outside the region of the *uv*-plane sampled by a given configuration. To illustrate this concept, the annularly averaged amplitudes of the Fourier transforms of three uniform disks with sizes of 5″, 10″ and 20″ are shown in Figure 7.13 for an observation at 100 GHz. The smallest uniform disk is closest to a point source and so it has large amplitudes up to a baseline of 180 m. Meanwhile, the most extended disk has large amplitudes only up to ∼40 m. Therefore, an array with baselines larger than 40 m will not be sensitive to emission on angular scales larger than ∼20″, and will filter out most of the emission from such an extended disk. An important consequence of such filtering is that an interferometer only detects a fraction of the total flux density for sources with emission on size scales larger than its shortest baseline. Indeed, if the source only has structures on size scales larger than the shortest observed baselines, one can "resolve-out" the source entirely. It is to ameliorate the effects of spatial filtering for extended sources that the 7-m and TP Arrays are available in addition to the 12-m Array configurations.

The maximum recoverable scale (θ_{MRS}, see Equation 3.27) is the largest angular structure to which a given array is sensitive. It is determined by the length L_{min} of the shortest baseline in the array. In principle this baseline measures spatial scales with a period of λ/L_{min}. In fact the sensitivity of this measurement to such large

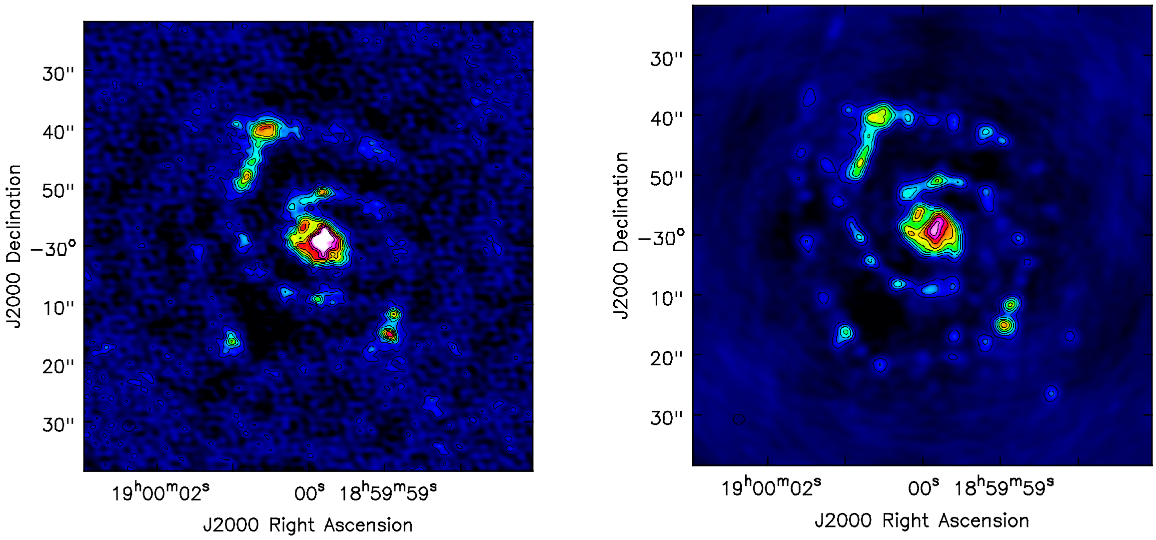

Figure 7.12: Images obtained from a model ALMA observation with 1-minute integration (*left*) and 1-hour integration (*right*), using the C43-1 configuration to observe a source at declination of −23°. Black contours at 100, 200, 300, 500, 700 and 900 mJy beam^{-1} are overlaid.

structures is not very good, so a smaller value of θ_{MRS} is typically adopted. Indeed, the exact filtering depends on precisely how the structure of the large scale emission maps to the (lack-of) short-baseline uv-coverage and is best determined by simulations. ALMA has adopted a criterion of measuring 10% of the total flux density of a uniform disk, which yields:

$$\theta_{MRS} \approx \frac{0.6\ \lambda}{L_{min}}\ [radians] \tag{7.6}$$

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

$$\theta_{MRS} \approx \frac{0.983\ \lambda}{L_5}\ [radians] \tag{7.7}$$

where λ is the observing wavelength in meters, and L_5 is the 5$^{\text{th}}$ percentile of the *uv*-distance in meters.

Table 7.1 lists the θ_{MRS} for the various Cycle 5 array configurations. Note that these expressions assume good *uv*-coverage. This situation is typically the norm for ALMA 12-m Array observations, but may not be for very short 7-m Array observations due to the relatively small number of antennas.

The spatial filtering of an interferometer observation is a serious issue that must be considered carefully for each science case. Ideally, the range of spatial frequencies over which the source has interesting structure will be known *a priori* so that the appropriate configuration can be used during observations. Selecting a specific range of angular scales has the effect of discarding all emission from those angular scales outside that range. If very large spatial scales (relative to those measured by the 12-m and 7-m Arrays) are important, the combination of the TP Array, 7-m Array, and 12-m Array should be used, as discussed in Section 7.8.

Figure 7.14 shows the annularly averaged visibility amplitudes *vs.* *uv*-distance of an example astronomical source, M51. This example is taken from an H-α image, but it is used to indicate the flux expected at 100 GHz. The amplitudes of the visibilities can be approximated as a power law of *uv*-distance with a negative index. This

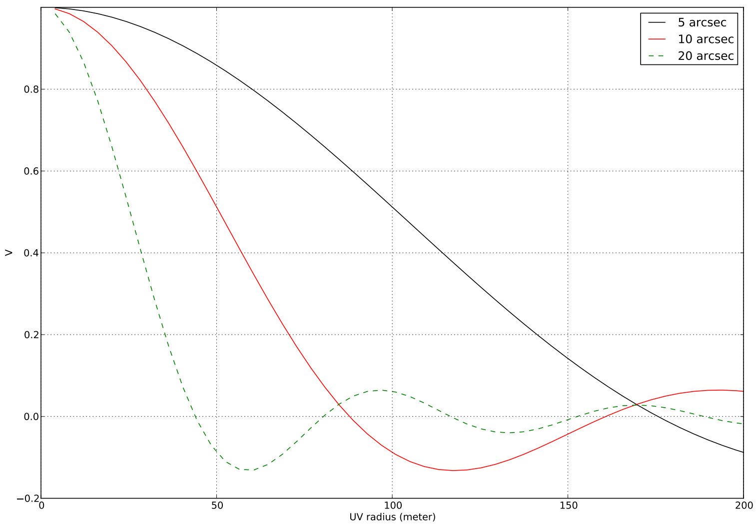

Figure 7.13: Annularly averaged amplitudes of the expected visibilities of three model uniform disks (annular averages) at 100 GHz as a function of uv-distance.

indicates that most of the power is located in larger scale structures and that power decreases rapidly at smaller scales. This result shows that the flux that would be received by the 7-m Array would be much higher than that received by the 12-m Array with extended configurations (e.g., C43-6). The only case where flux is independent of the sampling in the uv-plane is for point sources, which have the same amplitude for all visibilities (i.e., the Fourier transform of a Dirac function). In our M51 example, the visibility amplitude detected by the 7-m Array (at $k_{uv} \approx 5k\lambda$) is ∼23 Jy, whereas that detected by the configuration C43-6 ($k_{uv} \approx 200k\lambda$) is only ∼0.9 Jy. Ideally, ALMA users should use simulations to estimate the distribution of power at different length scales for targets they wish to observe.

7.7 Mosaicing

The field of view of a single interferometer pointing is determined by the antenna primary beam. A uniformly illuminated, circular aperture will have a beam width at half power ($FWHM$) of $1.02\frac{\lambda}{D}$. It will also have very high side lobes due to the abrupt truncation of the antenna illumination pattern; indeed, these side lobes would be particularly problematic for single-dish observations. Most radio and millimeter receivers illuminate their antennas with approximately Gaussian illumination patterns that smoothly go down to -10 dB to -15 dB response at the edges of the dish. The ALMA feedhorns were designed to illuminate the dish with a -12 dB edge taper to provide a nearly Gaussian primary beam with low sidelobes while preserving as much of the resolution and sensitivity as possible. This tradeoff is a fundamental aspect of radio telescope design. The resulting $FWHM$ of the ALMA 12 m antennas is measured to be $FWHM = 1.13\frac{\lambda}{D} = 58''$ at 100 GHz, and it varies inversely with frequency. For example, Figure 7.15 shows a Gaussian profile that approximates the shape of the primary beam at 112 GHz with $FWHM = 52''$. As described in Chapter 3, the primary beam attenuation can be corrected as the last step of imaging. In addition to correcting the sky distribution of the signal to its correct value, this correction, however, also increases the noise (i.e., the signal-to-noise is unchanged). For example at the radius of the FWHM, the noise will be increased by a factor of two compared to that at the phase center. Moreover, beyond a certain map size, the antenna is not sensitive and it is necessary to observe several adjacent pointings, i.e., a mosaic, to recover the sky emission.

Observing a mosaic with ALMA is needed if a map size larger than approximately the $FWHM$ of the

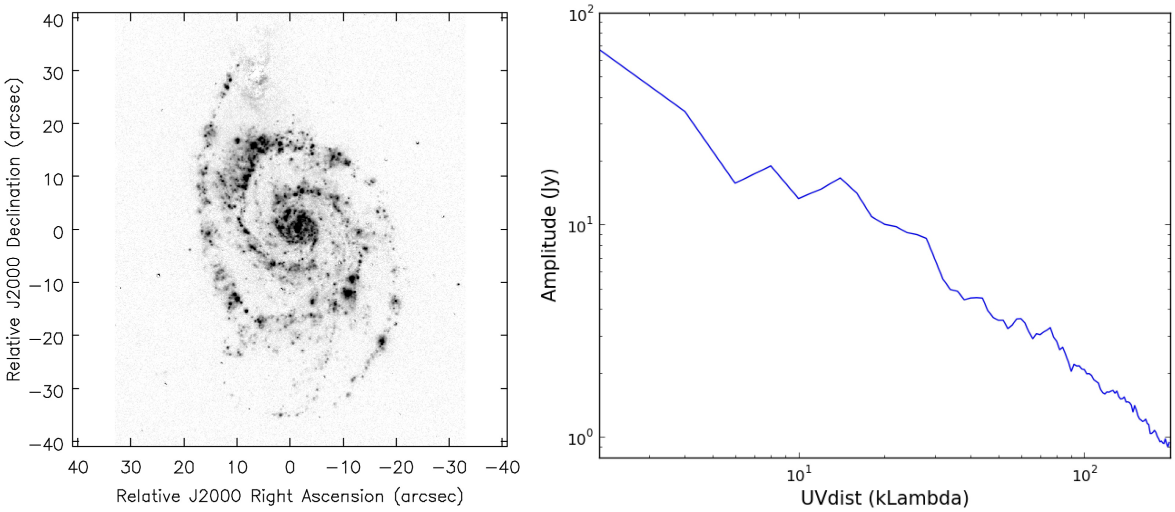

Figure 7.14: Image of Hα emission from M51 used as a model of emission at 100 GHz (*left*) and the expected visibility amplitudes with *uv*-distance (*right*).

primary beam is required. The default mosaic pointing pattern used by ALMA is a "fully sampled" hexagonal grid with equilateral triangles whose vertices are separated by $\theta_{hex} = \frac{\lambda}{D\sqrt{3}} = 0.511 \times FWHM$. With this spacing, the mosaic will sample the emission at the Nyquist spatial frequency. Note that a hexagonal mosaic has spacing θ_{hex} along a row (e.g., in right ascension) and $\frac{\sqrt{3}}{2}\theta_{hex}$ between rows (e.g., in declination). Figure 7.16 gives an example of such a mosaic. To estimate the number of pointings (N_p) necessary to cover an area of $L_X \times L_Y$ using this hexagonal pattern, the following expressions can be used:

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

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

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

7.8 Multi-array and Multi-configuration Observations

As shown in Table 7.1, different 12-m Array configurations provide different angular resolutions, θ_{res}, which are a function of the longest baselines of the configuration. Similarly, the maximum recoverable scales , θ_{MRS}, of each configuration depend on the shortest baselines present. To achieve the requirements entered by the PI in the OT, θ_{res} and θ_{LAS}, multiple configurations may be needed. Note that θ_{LAS} is a property of the science

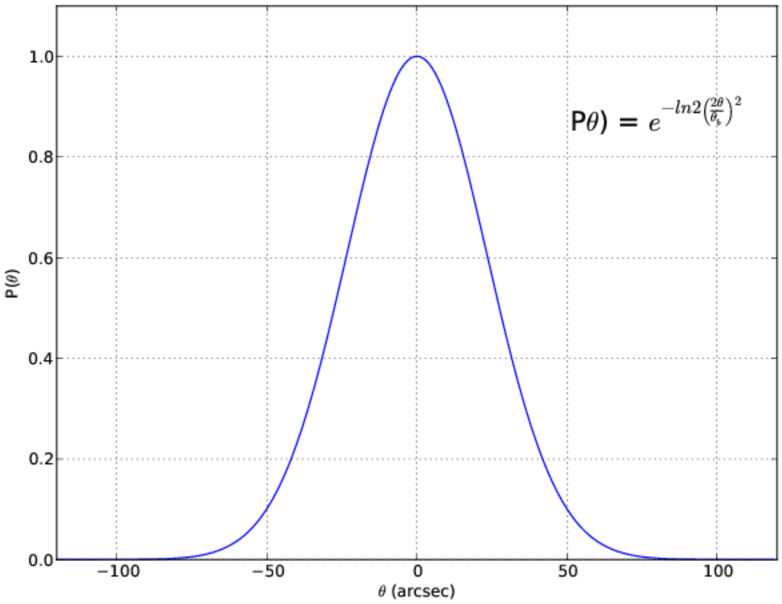

Figure 7.15: Approximation of the primary beam profile of an ALMA 12 m antenna at 112 GHz with FWHM = 52″.

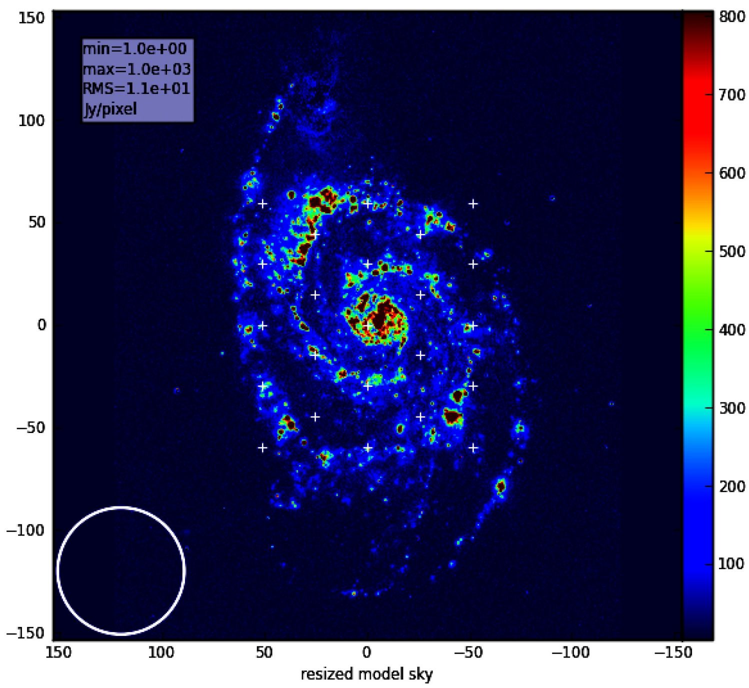

Figure 7.16: An example of mosaicing with a field of 2 arcmin extent at 100 GHz using a hexagonal pattern with Nyquist sampling (white crosses).

target, while θ_{MRS} is a property of an array configuration. In particular, the most compact configuration for a project must ensure that θ_{MRS} is larger than or equal to the θ_{LAS} of the science target.

Based on the user-entered values of θ_{res} and θ_{LAS}, the OT will automatically attempt to choose configurations that will produce a final image with the requested scientific goals. Note, however, that there are significant restrictions in this process. For example, at most two 12-m Array configurations are allowed and the most extended configurations (C43-10) cannot be combined with any more compact configurations. If such restrictions result in the requested θ_{LAS} not being achievable, the OT will return a validation error and θ_{LAS} must be modified. In contrast, a new possibility for Cycle 5 is "stand-alone ACA", i.e. observations that use the 7-m Array, with the TP Array added if required, but without 12-m Array observations. It is not possible to request observations that only require the TP array.

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

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

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

Table 7.4 also shows the relative integration times that will be used by each array. Broadly speaking, because of their smaller collecting areas, the TP and 7-m Arrays require more on-source integration time than the 12-m Array. The needed time ratios have been calculated by requiring similar sensitivity in the region of the *uv* plane where configurations or arrays provide overlapping baselines; details of this calculation are in Mason & Brogan (2013). The actual time ratios adopted for ALMA observations are guided by the calculations outlined in this memo, but are subject to further operational constraints. Cycle 5 implements again *configuration-dependent* time ratios. That is, the time multiplier for, e.g., the 7-m Array depends on the 12 m configuration with which it is being combined. A trend can be seen whereby the more compact the 12-m Array configuration, the more 7-m Array time is required. Increasingly compact 12-m Array configurations contain more sensitivity on baselines similar to those provided by the 7-m Array, which must in turn observe for longer to contribute significantly to the final image. The TP Array's time multiplier is fixed at 1.7 times the 7-m Array integration time.

As the compact configurations or arrays do require extra observing time, it is important to consider and justify the requested value of θ_{LAS}. An additional consideration when requiring the 7-m Array (with only 10 antennas) is that it provides limited *uv* coverage for snapshot observations. For Cycle 5, at least one hour of unique 7-m Array hour-angle coverage is recommended to provide sufficient *uv*-coverage for good image combination with data from the 12-m Array, and snapshot observations are therefore strongly discouraged. A comprehensive plot about the array combination is shown in Fig. 15 of the ALMA Primer for Cycle 5.

The process whereby the OT chooses 12-m Array configurations is done assuming the exact antenna positions that make up the ten representative configurations that have been selected for Cycle 5. In reality, the configuration available to the scheduling software at the time of Scheduling Block (SB) execution rarely conforms to this ideal, due to the fact that the array is often being reconfigured, i.e., antennas are being moved from one pad to another, a process that takes some time. Therefore, it is usually the case that the available

θ_{res} (arcsec)	θ_{LAS} (arcsec)	Array combination	Time ratios	Total Time
0.042	< 0.496	C43-10	1	$1.0 \times \Delta_{extended}$
0.042	> 0.496	-	-	-
0.057	< 0.814	C43-9	1	$1.0 \times \Delta_{extended}$
0.057	0.814-4.11	C43-9 + C43-6	1 : 0.5	$1.5 \times \Delta_{extended}$
0.057	> 4.11	-	-	-
0.096	< 1.42	C43-8	1	$1.0 \times \Delta_{extended}$
0.096	1.42-6.7	C43-8 + C43-5	1 : 0.5	$1.5 \times \Delta_{extended}$
0.096	> 6.7	-	-	-
0.211	< 2.58	C43-7	1	$1.0 \times \Delta_{extended}$
0.211	2.58-11.2	C43-7 + C43-4	1 : 0.4	$1.4 \times \Delta_{extended}$
0.211	> 11.2	-	-	-
0.306	< 4.11	C43-6	1	$1.0 \times \Delta_{extended}$
0.306	4.11-16.2	C43-6 + C43-3	1 : 0.4	$1.4 \times \Delta_{extended}$
0.306	16.2-66.7	C43-6 + C43-3 + 7-m	1 : 0.4 : 2.3	$3.7 \times \Delta_{extended}$
0.306	> 66.7	C43-6 + C43-3 + 7-m + TP	1 : 0.4 : 2.3 : 3.9	$5.3 \times \Delta_{extended}$
0.545	< 6.7	C43-5	1	$1.0 \times \Delta_{extended}$
0.545	6.7-22.6	C43-5 + C43-2	1 : 0.5	$1.5 \times \Delta_{extended}$
0.545	22.6-66.7	C43-5 + C43-2 + 7-m	1 : 0.5 : 4.6	$6.1 \times \Delta_{extended}$
0.545	> 66.7	C43-5 + C43-2 + 7-m + TP	1 : 0.5 : 4.6 : 7.8	$9.3 \times \Delta_{extended}$
0.918	< 11.2	C43-4	1	$1.0 \times \Delta_{extended}$
0.918	11.2-28.5	C43-4 + C43-1	1 : 0.4	$1.4 \times \Delta_{extended}$
0.918	28.5-66.7	C43-4 + C43-1 + 7-m	1 : 0.4 : 5	$6.4 \times \Delta_{extended}$
0.918	> 66.7	C43-4 + C43-1 + 7-m + TP	1 : 0.4 : 5 : 8.5	$9.9 \times \Delta_{extended}$
1.42	< 16.2	C43-3	1	$1.0 \times \Delta_{extended}$
1.42	16.2-66.7	C43-3 + 7-m	1 : 5	$6 \times \Delta_{extended}$
1.42	> 66.7	C43-3 + 7-m + TP	1 : 5 : 8.5	$9.5 \times \Delta_{extended}$
2.3	< 22.6	C43-2	1	$1.0 \times \Delta_{extended}$
2.3	22.6-66.7	C43-2 + 7-m	1 : 5	$6.0 \times \Delta_{extended}$
2.3	> 66.7	C43-2 + 7-m + TP	1 : 5 : 8.5	$9.5 \times \Delta_{extended}$
3.38	< 28.5	C43-1	1	$1.0 \times \Delta_{extended}$
3.38	28.5-66.7	C43-1 + 7-m	1 : 5	$6.0 \times \Delta_{extended}$
3.38	> 66.7	C43-1 + 7-m + TP	1 : 5 : 8.5	$9.5 \times \Delta_{extended}$
12.5	< 66.7	7-m	1	$1.0 \times \Delta_{extended}$
12.5	> 66.7	7-m + TP	1 : 1.7	$1.7 \times \Delta_{extended}$

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

configuration's properties, θ_{MRS} and θ_{res}, do not exactly conform to those assumed by the OT. Given that the available configuration provides a similar range of spatial scales to that which was assumed, SB execution may still be possible provided that the other parameters (weather, project rank, LST range and executive balance) are also appropriate. This flexibility in scheduling is vital to the efficient execution of SBs and is built into the second stage of ALMA's Quality Assurance process (QA2) such that all observed projects should meet the user's imaging requirements, within some tolerances. Note that some additional flexibility is provided by the assumption of Briggs' weighting during the imaging stage, a scheme which allows a tradeoff between angular resolution and sensitivity.

7.9 Multi-array and Multi-configuration Imaging

If a project requires observations with different arrays or configurations, the data from each will be processed individually and delivered if they pass quality assurance. When all observations are complete, they will also— on a best-efforts basis— be combined to make and deliver a final image, again if it passes quality assurance.

The combination procedure comprises either one or two steps, depending on whether total power data are collected. First, the interferometric data will be imaged together in a single, "joint" deconvolution. This can be done in CASA by passing all interferometric measurement sets— 7-m and one or more 12-m Array configurations— directly to the CLEAN task[2]. All 7 m and 12 m data delivered during Cycle 5 will be directly combinable. *Manually* calibrated data deliveries from earlier ALMA cycles that were calibrated using CASA versions earlier than 4.3 require an additional step to put the 7-m and 12-m Array data weights on an equal basis. *All* ALMA data calibrated by the ALMA pipeline, however, have correct weights, irrespective of Cycle number and CASA version. The steps to correct data weight issues for manually calibrated data from earlier Cycles are described in the "Data Weights and Combination" CASA guide[3]. Using the STATWT method is another easy and reliable approach assuming sufficient line-free spectrum can be identified.

If total power data are part of a project, then there is a second step to the combination procedure, which is to combine the deconvolved interferometric (7-m+12-m) image cube with the total power image cube. In Cycle 5 this combination will be done by "feathering" them together. Feathering is a commonly used technique in radio imaging for combining two images together by forming a weighted sum of their Fourier transforms. The steps are as follows:

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

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

As an example, the uv-coverages and uv-plane sampling distributions for a 1-hour observation with the C43-1 configuration and 5-hour observation with the 7-m Array are shown in Figure 7.17 and Figure 7.18, respectively. The Cycle 5 7-m Array configuration with ten 7 m antennas provides *uv* measurements in the range 8-32 m and takes over the 12-m Array configurations in the range $R_{(u,v)} < 15$ m. Using the Hα emission in M51 as a model, and given the spread of its emission over different length scales (see Figure 7.14), Figure 7.19 shows the images resulting from observing with a C43-compact configuration (C43-3), a compact and extended configuration (C43-3 + C43-6), and adding the 7-m Array. Natural weighting was used and deconvolution was performed using the CLEAN algorithm. The recovery of larger scales is quite noticeable with the inclusion of the 7-m Array data. Using only the C43-6 configuration at 100 GHz provides an angular resolution of 0.27″. Combined with the C43-3 configuration, the angular resolution is lowered to 0.31″. Finally, the angular resolution with the full combination (C43-3 + C43-6 + 7-m) is 0.32″, very similar to that of only the two 12-m Array configuration

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

[3]https://casaguides.nrao.edu/index.php/DataWeightsAndCombination

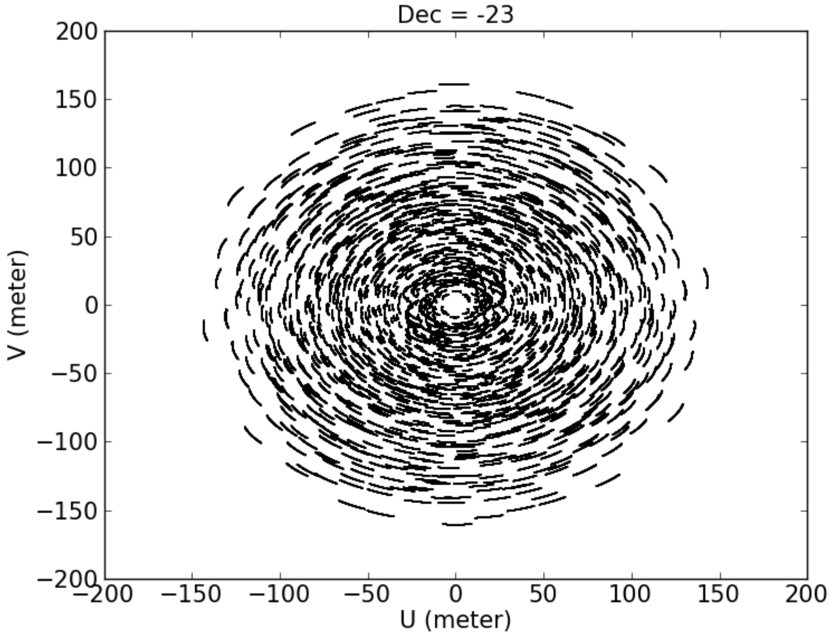

Figure 7.17: Expected uv-coverage for C43-1 (1 hour) and 7-m observations (5 hours)

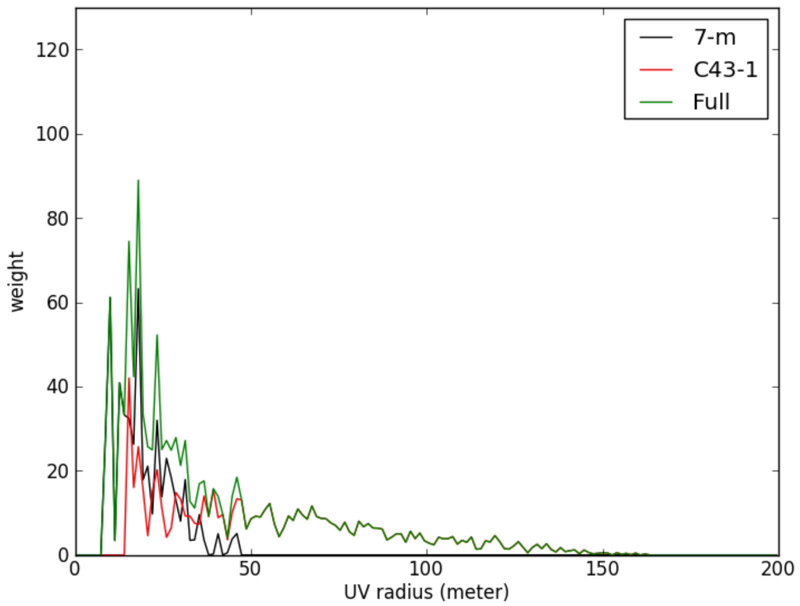

Figure 7.18: Radial density of uv-coverage with uv-distance using C43-1 (1 hour) and 7 m observations (5 hours)

data but 17% larger than the angular resolution obtained with the extended 12-m Array data alone. That difference is due to the respective *uv*-coverage and weight of the array configurations in the combined dataset.

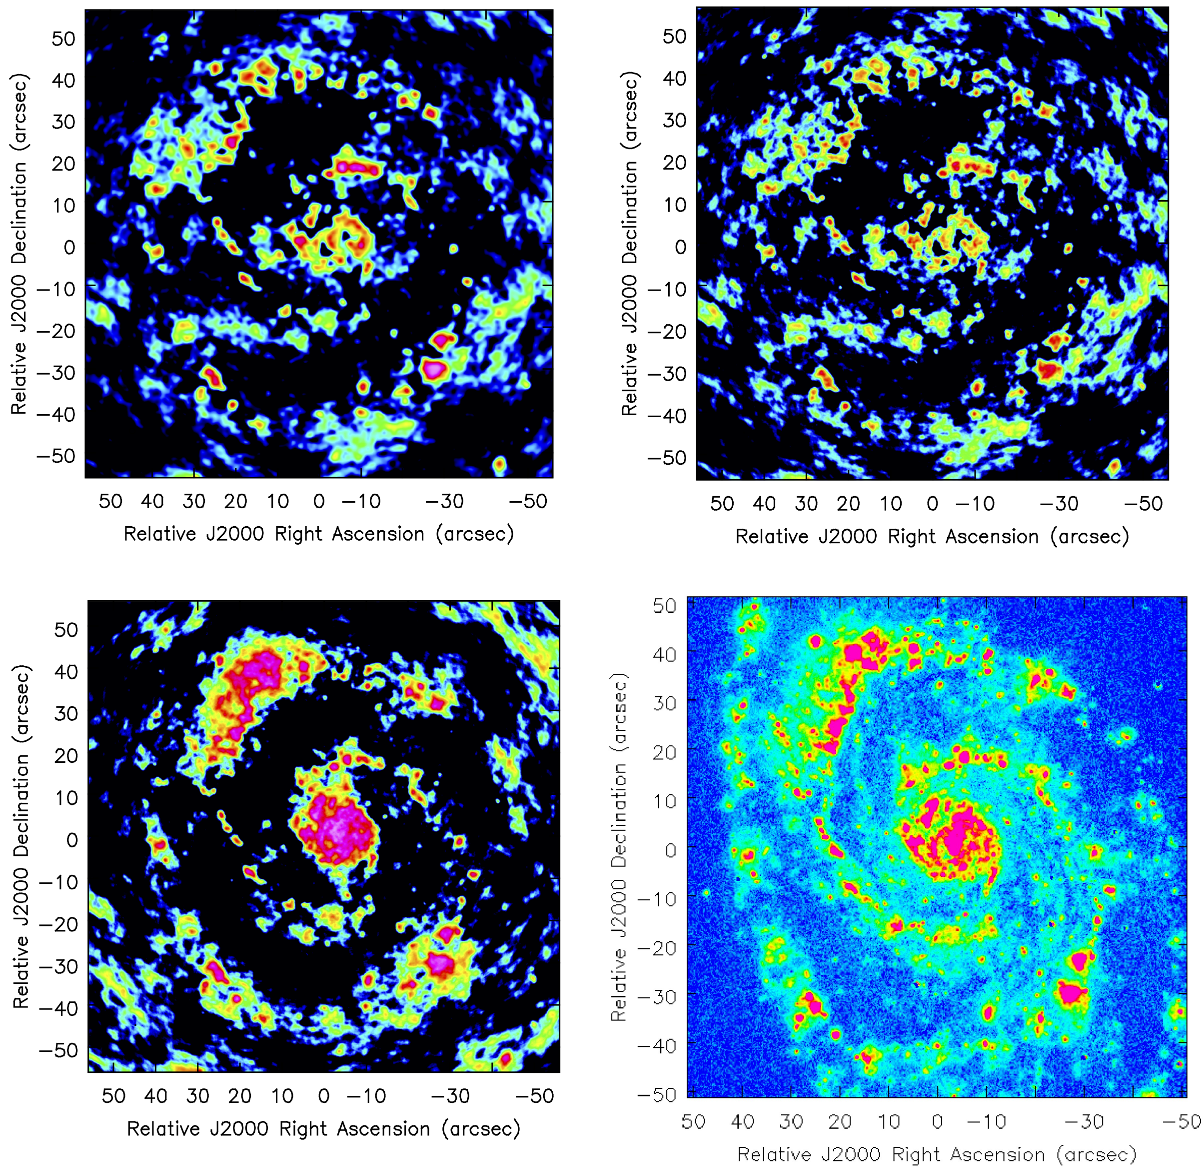

Figure 7.19: Images obtained using C43-3 (*top left*; 1 hour), C43-3 + C43-6 (*top right*; 1 + 3.3 hours) and C43-3 + C43-6 + 7-m (*bottom left*; 1 + 3.3 + 1.3 hours) Array combinations, and the image model itself (*bottom right*).

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

References

[1] Briggs, D. S. 1995, *High Fidelity Deconvolution of Moderately Resolved Sources* Ph.D. thesis, New Mexico Institute of Mining and Technology

[2] Briggs, D. S., Schwab, F.R., Sramek, R. A. 1999, *Synthesis Imaging in Radio Astronomy II*, ASP Conference Series, 180, 127

[3] Condon, J. J., R, W. D., Greisen, E. W., Yin, Q. F., Perley, R. A., Taylor, G. B., Broderick, J. J. 1998, AJ, 115, 1693

[4] Mason, B. & Brogan, C., 2013, Relative Integration Times for the ALMA Cycle 1 12-m, 7-m, and Total Power Arrays, ALMA memo 598

[5] Moriarty-Schieven, G. ed., 2013, Observing with ALMA: A Primer for Early Science, ALMA Doc. 2.1, ver. 1

[6] Stanimirovic, S. 2002, in Single-Dish Radio Astronomy: Techniques and Applications, edited by S. Stanimirovic, D. Altschuler, P. Goldsmith, & C. Salter, vol. 278 of Astronomical Society of the Pacific Conference Series, 375. astro-ph/0205329

Chapter 8

Observing Modes

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

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

8.1 Observing Project Structure

The organization of each science project is subdivided into a well-defined structure with clear hierarchical levels. At the bottom of this structure are the *Scheduling Blocks* (hereafter SB). An SB is the minimum set of instructions describing an ALMA observation that can be fully calibrated. The SBs to be executed during the observations are produced automatically by the ALMA Observing Tool (hereafter OT) during the Phase 2[1] stage of the project life cycle. The SB may be edited, using the OT, in Phase 2 after the proposal has been accepted for observation and also during the proposal execution if necessary[2]. Projects are broken down into a set of these fundamental units for flexibility, given the properties of the ALMA site and the continuously-varying status of the Observatory as a whole (including the weather), and to encapsulate the scientific objectives of the proposal into single entities.

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

Each SB typically executes set-ups, calibrations, and target observations within 1.5 hours. The end of an SB may be specified in terms of a maximum amount of time or when certain well-defined science goals have been reached as specified in the *Science Parameters* section of the SB. An SB is atomic in the sense that it cannot be stopped and re-started. Therefore, an SB runs to completion, fails, or is terminated by the *Astronomer on Duty* (or AoD). Given the limited length of the SBs, it is often necessary to observe it several times to achieve

[1]For a description of Phase 1 and Phase 2 see the Cycle 5 Proposer's Guide
[2]SBs can also be created in an ad-hoc manner for commissioning purposes

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

the required sensitivity. The current maximum time (or end execution) for a particular SB, has been chosen based on statistical measurements about the stability of the system. A system failure will prevent the online software from flushing the data into the archive losing the observed project data. Thus, executing SBs for more than two hours becomes unsafe from the operational perspective. To optimize this scheme, a *session* was implemented, which is the continuous execution of a single SB until a certain goal has been reached (see detailed description later on). Sessions have the added benefit that some of the calibrations can be shared among the different elections, which increases the overall observing efficiency. For this cycle, sessions are primarily used for polarization observations.

If the SB is the smallest entity used for observing, the Observing Unit Set (hereafter OUS) is the smallest unit for data processing[3]. The SB/Member OUS/Group OUS are the smallest structures that hold science observations that need to be observed/processed/combined together. These data hierarchies therefore maximize observing flexibility and ensure that data gets processed as early and as fast as possible. In the OT, observations are divided into different Science Goals (see the ALMA Proposer's Guide). In order to follow this structure, each Group OUS is attached to a Science Goal OUS.

All SBs have to be processed together to produce calibrated science products (e.g. images or data cubes). In most cases, there will be exactly one SB in each Member OUS, but the latter may hold multiple executions of this SB. For polarization observations, all the SBs that belong to the same session should be grouped together for data processing. Unfortunately, this is not explicitly described in the SB.

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

Pipeline processing happens at the Member OUS level, as soon as all observations of a Member OUS are completed (see Chapter 13 for more details). It is expected that in future cycles, processing will also happen at the Group OUS level, in case the Group OUS contains several Member OUSs which have already been processed. The only other event that triggers data reduction is the end of an observing season, when all the Member OUSs and potentially Group OUSs are reduced, irrespective of their degree of completion. As polarization observations do not currently fit this scheme, the data reduction will be done manually by the ALMA staff.

8.2 Program Execution

Once a given SB has been selected for execution (by the Scheduler subsystem or by the Astronomer on Duty), it is read into the online Control software, particularly by the Science Software Requirements subsystem (hereafter SSR). The SSR subsystem commands the lower level Control software to create an Execution Block (EB) structure that is attached to the SB. As many SBs will be executed several times, a number of EBs may exist for a given SB. Each EB contains a record of the parameters and conditions under which the SB was executed along with references to the acquired data. The internal hierarchical structure of the EB is also shown in Figure 8.1. The SSR subsystem constructs a sequential series of scans for each of the required calibrations and command Control to execute them. Each scan[4] execution is in fact carried out by breaking it down into a series of subscans, each of which is itself broken into a series of integrations (the correlator software only understands subscan sequences). Although commands are issued at the scan/subscan level, the correlator output corresponds to a particular integration. In general, each calibration observation consists of a scan containing several subscans (e.g. the 5 points of pointing calibration constitutes a scan with five subscans). Similarly, the integration time

[3]There are two types of OUS, the Member OUS and the Group OUS. Member OUS is the set of all SBs needed to achieve (part of) a Science Goal, and usually it is associated to a specific array configuration. Group OUS contain all Member OUS needed to achieve the science goals of a complete science proposal, and it can include Member OUS corresponding to different ALMA configurations and/or arrays.

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

on a single science source between phase calibrations can consist of one scan comprised of a number of subscans. To optimize the execution, scans may be organized in scan sequences which are passed to Control for execution. Scans and scan sequences can be of arbitrary length, depending on the characteristics of a given observation. But subscans are recommended to be 30 seconds or less. Integrations tend to be on the order of 1 to 10 seconds where the final value has to be an integer multiple of the correlator dump time. These values can be specified during Phase 2 SBs construction, particularly in the SB target parameters section. Calibration results from the Telescope Calibration subsystem (TelCal) and QuickLook (QL) pipeline are usually attached to a scan where one example is an antenna pointing result. The SSR and the Control subsystem are responsible for the creation of all the metadata needed downstream for data processing.

8.3 The Observing Process

For a given project, a set of targets are specified in the SB (for more details, see the Cycle 5 OT User Manual). These targets can correspond to either calibration or science executions. Targets are organized in observing groups where the first group (group 1) is always the initial calibration group, with subsequent groups detailing the science observations[5] All Science targets and relevant calibrators within a group are observed before the next group is started. An SB can have multiple groups.

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

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

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

The user has several options to select optimal calibrators. He/she can let the OT set up default queries to the Calibrator Database in the ALMA Archive which will be used to select appropriate calibrators at run time (*system-defined calibration* in the OT). This is the recommended mode. He/she can also enter specific calibration sources or set up the queries using alternative values for the parameters, but this carries some risk (for example, calibrators will not be observed during the execution of a group if they are not visible at the time of the observation), and thus must be fully justified in the Technical Justification of the proposal (*user-defined calibration* in the OT).

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

8.3.1 The Source Selection Algorithm

The ALMA system calibration for a given scheduling block is done using astronomical sources selected at runtime. To correctly select the appropriate calibration sources, selection criteria were implemented in the SSR using the following steps:

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

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

- Bandpass calibrator. The bandpass calibrator integration time defined in the SB (typically 5-15 minutes, depending on e.g. band and spectral setup) and a number of dynamically determined parameters (e.g. the number of antennas in the array and the estimated system temperature) are used to determine a flux density threshold for the catalog search that would correspond to a minimum SNR per antenna of 50. If no suitable sources are found, the SNR limit is lowered. Other factors that are taken into account are elevation and shadowing. The search radius is typically 45 degrees from the query center, but if no suitable calibrator is found the search radius is increased.
- Phase calibrator. A signal to noise of 15 given is used in the flux estimation. The required integration time is derived either using the widest spectral window bandwidth (usually 2 GHz), or the integrated bandwidth when only narrower spectral set-ups are used. A 15 degrees from the query center search radius is used as default (smaller cycle times and radii are used for the highest frequencies and longest baselines, as described in Chapter 10. However, the value is taken from the SB and the closest source to the target, from the selected list, is returned.
- Flux calibrator. A solar system object is searched for first, then the grid cal sources are searched. The first observable grid source is picked up. A detailed list of the Solar System objects and grid sources used for flux calibration is available in Chapter 10.

8.4 Single Field Interferometry

Single field interferometry is the most basic form of observation that ALMA supports. It consists of standard calibration scans associated with the constituent targets, and science observations of a single field (primary beam). A typical observation will start with a bandpass calibration. The bandpass observation is executed to measure the spectral response of the system, and thus should be done on a bright source with simple spectral properties, such as a bright quasar with no emission or absorption lines and a reasonably flat spectrum.

The flux scale calibration will be taken next, which is intended to obtain the observed flux of a well known source, such as a solar system object. The observed flux will be used to compare with the established model fluxes of these objects to obtain the scaling factor to be applied to all other sources in the SB. Ideally these flux scale calibrator sources should be small with respect to the synthesized beam so as not to resolve the source structure and create uncertainties in the flux calibration. In practice many solar system objects may be moderately resolved. In such cases only a subset of the antennas may be used to estimate the flux based on accurate models of the objects available in the analysis software. From this result, a scaling factor is derived which can then be applied to the data at the data reduction stage. If no solar system object is observable, the SSR will pick the first available grid source.

The interferometer is a device for measuring the spatial coherence function (Clark 1999), the phase information of the sources must be preserved and discerned from phase variations at the individual antenna elements. This is achieved with the phase calibration during the course of an observation. Since the phase is expected to change much more rapidly in time than the amplitude, these sources will be observed more frequently than the other calibrators. Because the phase varies on small scales on the sky, the calibrator must be quite close to the science target. The phase calibration observations are taken before and after observations of the science target, and the phase correction will be interpolated in time when applying to the science target. As the atmosphere fluctuates rapidly (especially at higher frequencies), radiometric observations of atmospheric water lines are done to correct for these additional phase variations. These corrections can be applied online and offline (see Section 10.4.2 for further details).

8.5 Pointed Mosaic Observations

Pointed mosaic observations are also offered for this Cycle. This mode enables a single Science Goal to cover a field of view larger than the primary beam by making observations of multiple single fields that overlap by an amount specified in the proposal. Up to 150 pointings are possible in a single SB. This limitation is set by the maximum execution time for a single SB and the necessity to finish all fields at least once within an execution. In pointed mosaic observations, each of the fields will be assigned different field IDs in the data, but the same source ID. Thus all fields will share their bandpass, amplitude and phase calibrations which are done as single fields. For larger fields that cannot be covered by 150 pointings, multiple SBs may need to be defined. The mosaic is arranged as a single scan composed by a number of subscans corresponding to the individual rows in the mosaic (see Section 5.5). Both the Science target specific last mosaic pointing position and index within the Science target list are handled by the SSR to ensure the next scan begins on the proper Science target and proper offset position. The mosaic observations can be set up by specifying *a Rectangular Field* in the OT. If the user wishes to execute several small mosaics within an SB, the *custom-mosaic* under *Multiple Pointings* should be selected.

8.6 Single Dish (Total Power) Observations

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

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

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

The OTF map is observed as a series of raster rows, scanning in the coordinate system specified in the SB. In Cycle 5, scans are taken either in longitudinal or latitudinal directions as specified in the OT[6]. In later cycles, scans may be taken in the two perpendicular directions in turn, to minimize scanning effects. The reference positions, assumed to be positions which are free of emission, are specified either in absolute coordinates or as offsets from the map center. By default, the OTF map will cover an area one half of the full beam larger than the interferometric observation on all sides. This will ensure that undersampling at the map edges does not affect the data combination process with the interferometric data. A raster row in these observations is defined as a subscan, with a maximum length of a scan (consisting of some number of subscans) of 600 seconds. The reference position is observed as specified by its cycle time during the science target scans. Pointing is calibrated on a bright calibrator near the Science target with a frequency indicated in the SB, and atmospheric calibrations are taken every 10 minutes at the reference position to measure the system temperature.

The calibrated TP map of the science target will be in units of Kelvin, on the antenna temperature (T_a^*) scale. Since all TP observations are expected to be combined with 12-m and 7-m Array data which come in calibrated units of Jy/beam, the TP data must also be converted into these units. This conversion from Kelvin to Jy/beam requires knowledge of the main beam efficiency η_{mb} and the beam size θ. These are measured separately for each project by obtaining a continuum map of a bright quasar or a planet with known flux[7]. These "amplitude calibration" maps are reduced in the same way as the science observations, then the emission will be compared with their model or observed flux (taken from the most recent calibrator survey measurements in the case of quasars) to calculate the Kelvin to Jansky/beam conversion factor, which will then be applied to the science observations.

8.7 Polarization

The ALMA antennas have receivers with linearly polarized feeds followed by a waveguide with a polarization splitter[8]

In this way, the incoming radiation is separated into two orthogonal components (X and Y) which are down-converted and digitized independently. For each baseline, the digital signals are cross-correlated at the correlator where the outputs are the four cross-correlations XX, YY, XY, and YX (or V_{xx}, V_{yy}, V_{xy}, and V_{yx}). These four cross-correlations, as a function of the Stokes parameters, are ideally given by

$$V_{xx} = I + Q \tag{8.1}$$
$$V_{xy} = U + iV \tag{8.2}$$
$$V_{yx} = U - iV \tag{8.3}$$
$$V_{yy} = I - Q \tag{8.4}$$

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

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

[7]It is expected that during Cycle 5, these measurements will only be needed for the high-frequency ALMA bands, while for bands up to Band 7, the values will be automatically derived from fits to the aperture efficiency of each antenna as a function of time, elevation, outside temperature, surface accuracy, etc., which will be regularly monitored by the observatory.

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

$$I = \frac{V_{xx} + V_{yy}}{2} \tag{8.5}$$
$$Q = \frac{V_{xx} - V_{yy}}{2} \tag{8.6}$$
$$U = \frac{V_{xy} + V_{yx}}{2} \tag{8.7}$$
$$V = \frac{V_{xy} - V_{yx}}{2i} \tag{8.8}$$

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

1. The splitting of the incoming radiation into orthogonal components is not perfect and small projections of one component into the other are produced. This is called the *instrumental polarization* or $D - terms$. The instrumental polarization is an antenna based quantity which also depends on the frequency and the band (cartridge design and external optics). Additionally, this quantity is measured in the frame of the antenna which is an elevation and azimuth frame (Alt/Az) and thus, it rotates with respect to the frame of the sky. This rotation introduces an angular dependance into the visibilities as a function of the parallactic angle (ψ). By design, the instrumental polarization is small (a few percent), but not negligible.
2. The signal path followed by each of the individual X and Y polarizations is slightly different, which introduces a small delay in the signal that needs to be accounted for. It has also been observed that the usage of the different local oscillators also introduces an offset in the XY-Phase.
3. The usual off-line calibration procedure arbitrarily sets the phase of a reference antenna to zero; this results in the inverse of its X-Y offset being imprinted on the cross-hand phases of the other antennas.
4. Other effects, such as position within the primary beam (off-axis polarization), are still under commissioning. Thus, only on-axis polarization is offered for Cycle 5 (see below for a detailed explanation.)

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

$$V_{XX} = (I + Q_\psi) + (U_\psi + iV)d^*_{X_j} + d_{X_i}(U_\psi - iV) + d_{X_i}(I - Q_\psi)d^*_{Xj} \tag{8.9}$$
$$V_{XY} = (I + Q_\psi)d^*_{Y_j} + (U_\psi + iV) + d_{X_i}(U_\psi - iV)d^*_{Y_j} + d_{X_i}(I - Q_\psi) \tag{8.10}$$
$$V_{YX} = d_{Y_i}(I + Q_\psi) + d_{Y_i}(U_\psi + iV)d^*_{X_j} + (U_\psi - iV) + (I - Q_\psi)d^*_{X_j} \tag{8.11}$$
$$V_{YY} = d_{Y_i}(I + Q_\psi)d^*_{Y_j} + d_{Y_i}(U_\psi + iV) + (U_\psi - iV)d^*_{Y_j} + (I - Q_\psi) \tag{8.12}$$

where d_{X_j} are the $D-terms$ as a function of polarization and antenna, the asterisk denotes complex conjugates, and U_ψ and Q_ψ are the Stokes parameters as a function of the parallactic angle. Figure 8.2 shows an example of $D - terms$ at Band 3. The $D - term$ level is typically a few percent at Bands 3, 4, 5, 6, and 7 on axis **TBC** with some variations over frequency[9]. Without any $D - term$ calibration, an unpolarized source may appear to be polarized at the 1% level. The most straightforward way to calibrate the $D - terms$ is to observe an unpolarized source. The cross-hands output will be purely $D - terms$ since Stokes Q, U, and V will have no signal. However, bright unpolarized sources are rarely found.

ALMA therefore normally uses an unresolved quasar as a polarization calibration source. These are monitored to ensure they are bright enough, and the exact polarization is determined during data reduction provided

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

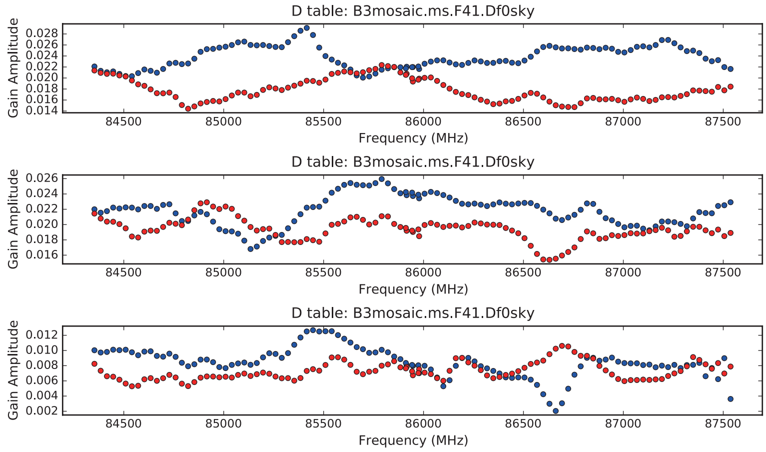

Figure 8.2: The $D - term$ plots of DA42 (top), DV03 (middle), and PM01 (bottom). The vertical axis is the fraction of the input signal voltage in one polarization that leaks into the output of the other polarization in voltage units and the horizontal axis is the frequency in MHz. The blue and red symbols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

that the source is observed over a wide enough range ($> 60°$) of parallactic angles to separate the effects of the source polarization, the $D - terms$ and the cross-hands phase spectra and delay. Fig 8.3 shows that the rate of change of parallactic angle is fastest near transit and slowest for low elevation sources. The $D - terms$ may have a slight elevation dependence, and therefore it is favorable for the polarization calibrator to be close to the science target. The ALMA Observatory will choose appropriate calibrators, but the users must take into account these limitations when planning the observation.

With the current calibration scheme, linear polarization imaging of a compact source on-axis, at the level of 0.1% in fractional polarization is feasible. The accuracy of absolute polarization position angle will be nominally 6 degrees (based on initial polarization mode verification results), with a contribution of 2 degrees coming from the specification on the orientation of the receiver feeds, and an additional scaling of the error with the number of antennas in the final image. Of course, these levels of accuracy will only be reached if there is sufficient signal-to-noise in the polarised emission. Although a Stokes V (circular polarization) image can be produced with reduction and analysis software (CASA), this capability is still under commissioning. The ALMA Observatory cannot guarantee the accuracy of Stokes V images at the current cycle. Thus, the users may investigate the Stokes V emission at their own risk.

As the $D - term$ component also arises from the off-axis geometry of feed horn, antenna illumination, and the alignment of optics, it will also vary across the primary beam pattern. Generally, the $D-term$ level becomes larger when increasing the offset from the beam center. This is so-called *off-axis* instrumental polarization. In Cycle 5, the off-axis instrumental polarization calibration will not be employed. In Cycle 5, polarization imaging with accuracy better than 0.3% can be achieved within the inner 1/3 of the primary beam FWHM and this determines the largest acceptable angular size of sources which can be observed in full polarization. Users must justify that their science goals can be met by sources exceeding these lower limits (of 0.1% polarization for a point source in the center of the field or 0.3% within up to the inner 1/3 of the FWHM).

For this cycle, polarization observations have been extended to FDM correlation mode at arbitrary frequencies and frequency resolutions within three of offered Bands (3, 6, and 7). For the two bands for which we offer polarization observations for the first time in this Cycle, i.e. Bands 4 and 5, the observations will only be possible at the following frequencies: Band 4: 127, 129, 139 and 141 GHz, and Band 5: 196, 198, 208, 210 GHz, respectively. Figure 8.4 shows the *D-term* solutions of a representative antenna (DV22) obtained during commissioning. Figure 8.5 shows a zoomed-in view in frequency of the *D-term* solutions shown in Figure 8.4. These figures show that the spectral shape of the *D-term* solutions is smooth and no spur like structure is seen

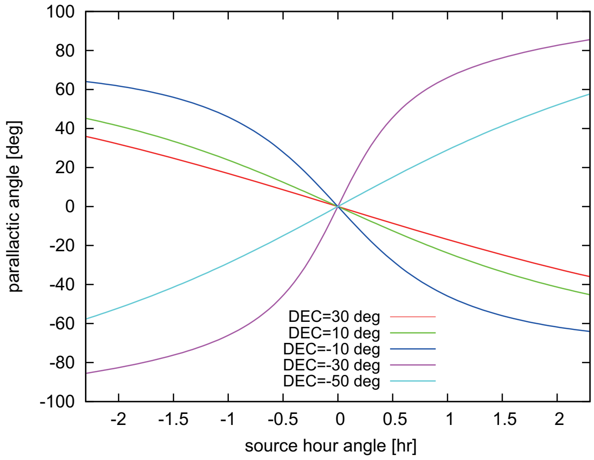

Figure 8.3: Parallactic angle (ψ) plot as a function of hour angle for the declination of 30, 10, -10, -30, and -50 degrees.

even in highest frequency resolution mode (30.5 kHz at the right bottom panel of Figure 8.5). Figures 8.4 and 8.5 clearly demonstrate that the instrumental polarization is mostly spectral resolution independent for these datasets. No mosaic, ACA, circular polarization and TP Array observations are offered. Further, a minimum execution time of 3 hours will be imposed to ensure sufficient parallactic angle coverage for calibration.

8.7.1 Sessions

Though not directly related to polarization iteself, sessions are the observing scheme designed to observe polarization projects with ALMA. In order to allow the execution of long programs and to avoid the execution of long scheduling blocks (more than 1.5 hours), the concept of a session was implemented into the SSR. Given the current stability of the system, it was considered an unnecessary risk to have more than 1.5 hours scheduling blocks, because in the event of a crash all data are lost. Thus, a session scheme was designed as an alternative. A session is defined as the continuous execution of the same SB until the scientific criteria are met. A session will manage the cycle time of each of the calibrators from the starting point of the session, i.e. the first execution of the SB. In this way, calibrations are only executed when needed avoiding unnecessary observations giving an additional level of optimization for ALMA. By default, calibrations such as flux calibration will be done once per project and bandpass every hour saving observing time. Also, the phase calibration science target loop will be interrupted (preserving phase calibration) when an additional calibration is-needed (e.g. polarization)[10] The current mode for sessions in ALMA is the observation of polarization projects. Because for an Alt/Az antenna the frame of the sky rotates, the calibration of the instrumental polarization requires sampling a strong compact source as a function of parallactic angle approximately every 30 minutes. The session will remember the last time the polarization calibrator was observed and interleave the calibration when needed. In general, one can achieve the parallactic angle coverage by running the polarization SB 2 or 3 times, which will give between 3 to 4 hours of observation, which is a perfect example of the session scheme. Other cases in which the session scheme might be useful are large mosaics, surveys (multi-target), and large single dish raster maps, but those cases are not offered in session mode for Cycle 5. The implementation of sessions is done by keeping information in memory about the previous execution of the SB. This is done through an interface provided by Control software to access a persistent component which is used to store and retrieve this information.

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

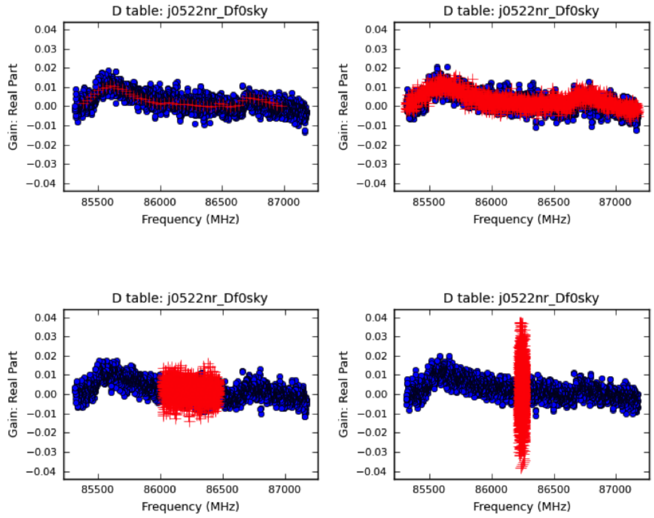

Figure 8.4: A frequency vs real part of the *D-terms* plot for an antenna (DV22) obtained towards J0522-364. Blue circles and red crosses are showing the D-term obtained using a setup of FDM with 2 GHz bandwidth and using a setup with higher spectral resolutions (FDM 500 and 62.5 MHz bandwidth respectively).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

8.8 Multiple Region Modes

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

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

This mode is useful when a number of line features which require high spectral resolution are spread across the IF bandwidth. Since the filters have a unit width of 62.5 MHz, if the user chooses a mode with 250 MHz total bandwidth, it is possible to place four separate windows, each with 62.5 MHz width, anywhere within the 2 GHz baseband. In Cycle 5, a maximum of four spectral windows per baseband will be offered.

8.9 Observation of Ephemeris Objects

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

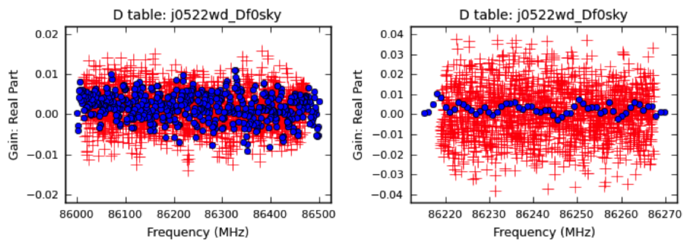

Figure 8.5: A close-up view in frequency space of the real part of D-term shown in Figure 8.4. Blue circles and red crosses are showing the *D-terms* obtained from setup #1 (FDM, 2 GHz bandwidth) and setup #2 (FDM, 500 and 62.5 MHz bandwidth respectively).

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

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

8.10 Solar Observations

Solar observing was part of the original science case for ALMA. It has recently been thoroughly tested and verified and released for PI science observations. In Cycle 5, solar observing with ALMA will be supported, albeit with limitations. In particular, the following conditions apply to solar observing in Cycle 5:

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

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

- Interferometric observations must be done using the TDM mode (see section 5.1.1); frequencies are fixed to 2 GHz-wide spectral windows centered on the frequencies shown in the table above. The spectral-line FDM observing mode is not offered for solar observations (Section 5.1.2).
- Simultaneous observations with Bands 3 and 6 are not offered: each execution block can only include one band.
- Observations will be carried out while the array is in a compact configuration corresponding to C43-1, 2, or 3.

- Because the WVR receivers are saturated when dishes point at the Sun, on-line WVR phase correction will not be applied and off-line WVR correction for on-source (solar) data is not possible.
- Interferometric observations will use a combined array comprising both 12 m and 7 m dishes and will be processed with the baseline correlator (Section 5.1).
- To minimize shadowing of 7 m antennas, observations will be carried out between 13:00 UT and 20:00 UT.
- Mosaicing of larger FOVs can be carried out with up to 150 different pointing offsets relative to the phase center specified in the ephemeris, typically in a rectangular pattern.
- One target (one ephemeris file) is allowed for each science goal.
- A single observation (an execution block) cannot exceed 2 hours, which will include the time overheads for bandpass and flux calibration. These calibration overheads amount to about 25 mins.
- The integration time of interferometric observations is fixed to 1 seconds.
- Observations may be performed using dual linear polarization (XX, YY) or single polarization (XX) correlations; full polarization measurements are not currently offered for solar observations.
- Total-power full-Sun fast-scanning single-dish observations are offered as context and to recover the largest angular scales for interferometric observations; proposals requesting only total-power single-dish observations will not be accepted in Cycle 5.
- Single-dish total-power mapping will be "monomode" (double-circle scan pattern) with a 2400 arcsec-diameter circular FOV centered on the solar disk (see below).
- Single-dish full-Sun will not be executed when the Sun is at elevations above 70^o because the required fast-scan azimuth slew speeds are too high.
- The time cadence of full-sun images obtained from total power observations is fixed to about 7 minutes for Band 3 and 10 minutes for Band 6.

8.10.1 Solar Observing Modes

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

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

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

Two so-called "mixer de-tuned" or "mixer de-biased" settings have been adopted for solar observations in Band 3 and Band 6 for Cycle 5. These are referred to as solar observing modes MD1 and MD2.

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

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

8.10.2 Array Observations

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

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

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

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

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

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

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

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

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

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

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

8.10.3 Single Dish Mapping

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

While various types of scan patterns have been developed and tested for ALMA dishes, Cycle 5 supports the use of a "double circle" pattern, which maps a circular region on the sky. The double-circle pattern is particularly well suited for full-disk mapping because its coverage matches the shape of the solar disk and it repeatedly revisits the region of the center of the disk, allowing atmospheric opacity variations to be corrected. Standard observing procedures include focus and pointing checks on suitable sources prior to the fast-scan mapping.

8.11 VLBI Observing Mode

The Very Long Baseline Interferometry (VLBI) Observing Mode (VOM) is a variant of the standard interferometry mode with some additional capabilities to allow ALMA to participate in global VLBI networks operating at millimeter and submillimeter wavelengths. Since Cycle 5, ALMA VLBI mode observing will be offered

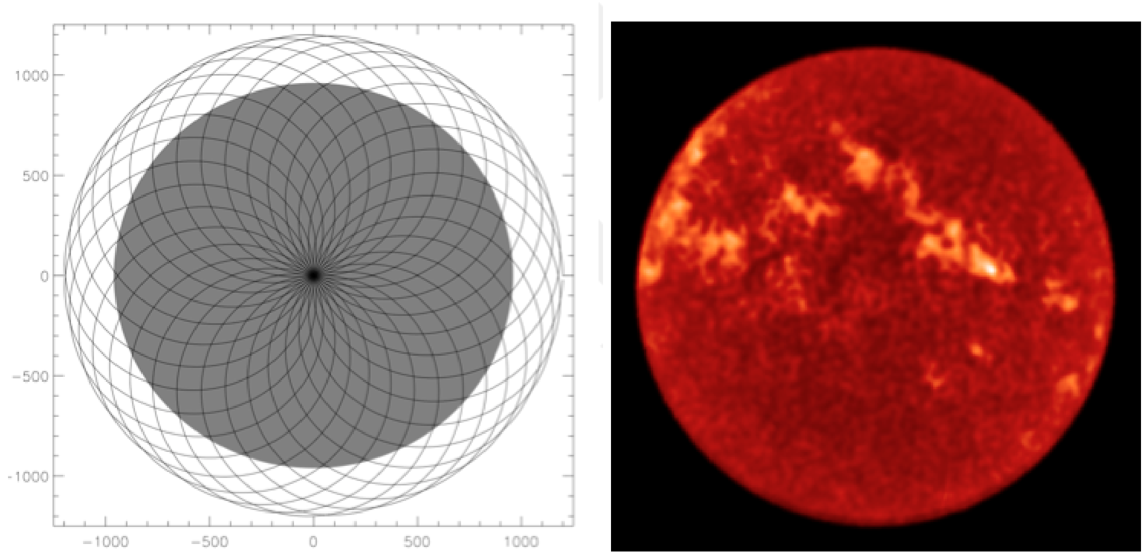

Figure 8.6: Left: Schematic illustration of a double-circle scan pattern used to map the full disk of the Sun. Actual patterns sample the source more densely to ensure that it is no less dilute than Nyquist sampling. Right: Example of a fast-scan map of the Sun obtained in Band 6.

in Band 3 in conjunction with the the Global Millimeter VLBI Array (GMVA)[11] and in Band 6 in conjunction with the Event Horizon Telescope (EHT) network[12]. For details about proposing for these opportunities, see the ALMA Cycle 5 Call for Proposals See `https://science.nrao.edu/observing/call-for-proposals/1mm-vlbi-cycle5/` .

For all VOM observations in Cycle 5, ALMA will be operated as a phased array of 12 m antennas (most likely containing between 35 and 43 phased antennas). For example, with 37 12 m antennas, this is an equivalent diameter of 73 m. From Section 9.2.1, Table 9.3 the aperture efficiencies in Band 3 and Band 6 are 0.71 and 0.68, and $T_{\rm sys}$ is 70 K and 100 K, respectively. Since $Gain \equiv A_{\rm eff}/(2k_B)$ and $SEFD \equiv T_{\rm sys}/Gain$,

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

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

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

8.11.1 General VLBI Considerations

The supporting VLBI networks (i.e., the GMVA for Band 3 and the EHT for Band 6), together with ALMA, place certain restrictions on the observing process; these are largely driven by the complexity of orchestrating the simultaneous, reliable execution of a common VLBI schedule, as well as seeing to the calibration requirements of the individual telescopes that comprise the global array. In practice, for the ALMA VLBI observer there will

[11] See `http://www3.mpifr-bonn.mpg.de/div/vlbi/globalmm/`
[12] See `http://www.eventhorizontelescope.org/`

be very little difference between planning an observation in the two observing bands, aside from a few minor details (*i.e.* differing bandwidths and frequencies).

Successful proposals will be passed to a network Scheduler who builds the common VLBI schedules (encoded in a so-called VLBI EXperiment, or VEX file built with `SCHED`[13]) for a particular observing campaign. In these schedules, some number of contiguous hours will be devoted to each set of science targets and necessary VLBI calibrators.

The schedule will be worked out between the Scheduler and the PI well in advance of the observation. Most sites participating in the VLBI network also have a "friend" of VLBI who assists the Scheduler in working out details specific to that site. Schedules from previous campaigns are available from the aforementioned GMVA website and it is expected that schedules for Cycle 5 with ALMA with either the GMVA or the EHTC will be similar.

In general, the required time on the VLBI science target(s) and VLBI calibrators will be broken into a number of VLBI scans of several minutes duration, separated by gaps on order of several minutes to allow time for local calibrations (pointing, system temperature measurements, *etc.*) and antenna slew time sufficient to meet the needs of each observatory. VLBI observations that include ALMA will require that the gaps between VLBI scans are of sufficient length to allow for the performance of the calibrations normally required for standard interferometric observing, including flux, gain, bandpass, and polarimetric calibrators (see below). Since all VLBI science targets for Cycle 5 are required to be sufficiently bright to allow phase self-calibration ($\geq$500 mJy), observations of a gain calibrator will be required only every 20-30 minutes to allow calibration of the amplitudes.

It is important to recognize that most VLBI stations currently record *circular* polarization (either left, right or dual), in contrast to ALMA which records dual linear polarizations. Consequently, *every VLBI mode observation with ALMA must treated as a polarization observation* in order that a transformation from linear to circular polarization can be made correctly during post-correlation processing. A special tool, `PolConvert` has been written to be used with the standard VLBI correlation software, `DiFX`[14], which converts ALMA data to a circular basis, greatly simplifying the analysis. In particular, this requires a minimum session duration of 3 hours (session length, not time on science target) in order to achieve adequate parallactic angle coverage on the polarization calibrator sources for the computation of polarimetric "leakage" or D terms.

Aside from differences in polarization, other VLBI observatories all have receiver capabilities that differ substantially from ALMA, especially with regard to bandwidth. In particular, none of the GMVA and EHT sites are currently able to match ALMA's full recording bandwidth. An additional complication is that the sample rate used at ALMA is non-traditional for VLBI, so some care must be taken in the experiment set-up to ensure that the VLBI data can be correlated with the other stations. The VOM spectral set-ups chosen for each band in Cycle 5 were selected to address these constraints and are therefore fixed. For Band 3 (at 3mm) there are 4 1.875 GHz bands centered at 86.124, 88.124, 98.124 and 100.124 GHz, respectively. The GMVA sites will be tuned to a 256 MHz band centered on 86.124, so only the first band is available for VLBI. In Band 6 (at 1mm) the bands are centered at 213.1, 215.1, 227.1 and 229.1 GHz, respectively. For Cycle 5, the EHTC is planning to offer 64 Gbps recording at all sites and will therefore be able to record all four bands and match for the first time ALMA's full frequency coverage and bandwidth.

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

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

[13] `http://www.aoc.nrao.edu/software/sched`
[14] `http://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/difx/start`

available to the observer (and uploaded to the ALMA Archive), as discussed in Section 8.11.2.

8.11.2 ALMA Considerations for VLBI

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

In order to participate in a VLBI observation, there are two special requirements at ALMA. The first is that the ALMA control system must tune receivers appropriately and point to the targets specified in the VEX schedule at the appropriate times. The second is that the signals from the ALMA antennas must be coherently summed, decimated, and recorded on media suitable for delivery to the common VLBI correlator. In order to form this coherent sum, the (arbitrary) phases of the signals at the antennas need to be adjusted during the observation so that they can be added "in phase". The necessary signal processing for this takes place within special cards and circuitry within the ALMA correlator as well as with special software within the ALMA control system which were developed by the ALMA Phasing Project (APP)[15] Details of this are presented in the following sections, with particular emphasis on those aspects which are essential either for the proposer or data analyst to be aware of. It begins with a discussion of the VOM Scheduling block in the next section, followed by information about the scan sequences used by the VOM. That is followed by a section with some important details about how the ALMA Phasing System (APS) works (Section 8.11.2). Finally Section 8.11.2 makes a few comments about the analsys of VLBI data (a full discussion would be considerably outside the scope of this document).

VLBI Observing Mode Schedule Block and Execution

As with other modes, VOM observations are executed with a scheduling block which provides observation-specific details to an SSR observing script. In the case of the VOM, the script is called `StandardVLBI.py`, although, as noted elsewhere, VLBI is a non-standard observing mode in Cycle 5.

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

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

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

[15] An NSF MRI/ALMA development fund project.

```
VLBI Scan 2016y089d07h50m00s on Target-X for 300s
perform some ALMA calibration
VLBI Scan 2016y089d08h00m00s on Target-X for 300s
. . .
perform final ALMA calibrations
```

VLBI Scan Sequence

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

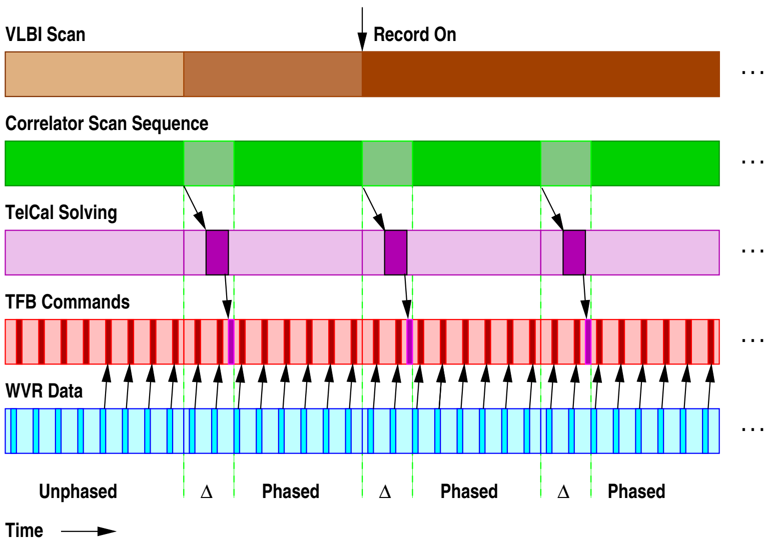

Figure 8.7: VLBI Scans in the Phasing System. Each VLBI scan is partitioned into "subscans" for correlation and "slow" timescale processing (seconds)in TelCal. WVR adjustments are made in the CDP on a "fast" timescale (every second). See text for full discussion and explanation.

While the correlator scans are shown as time-contiguous in Figure 8.7 (green bar), there are in fact small gaps to allow the data to be dumped out for subsequent correlator processing. Even so, the solutions from TelCal (purple) do not arrive at the TFBs (red) until after the correlator scan has started—the time of application is noted and the early part of the scan is excluded from the subsequent phase calculation. Usually one phase application is needed to get an acceptable solution, so the recording should start no sooner than the start of the third correlator subscan. These parameters are all programmable, and judicious choices for them are made when the schedule block is created. For Cycle 5, the subscans are likely to be 16 s long with 4 s (channel average) integrations; there are then about 30 s of observations prior to the start of the recording. Note that the ALMA correlator does its processing by subscans—each subscan must be dumped and processed by the correlator and ultimately archived—with a gap of several seconds between subscans. On the other hand, the signals from the antennas flow continuously to the summation logic and recorders, so there are no comparable gaps in the VLBI data.

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

VLBI Phasing System

Moving beyond the timing of the phasing system, there are several other aspects of the system that an ALMA VLBI observer should be familiar with. All ALMA observations operate with an array of antennas (*i.e.* those available to a scheduled observation). Of these, a portion are controlled by the phasing system to form the coherent sum signal which is used for the VLBI recordings. The previously mentioned reference antenna is merely a designated member of this "phased array". At least two antennas from the active array (designated "comparison antennas") are held outside the phased array. These comparison antennas are thus available for diagnosing the performance and efficiency of the phased array system. Finally, after construction the summed signal is decimated to two bits per sample (as is the case with all of the ALMA signals in the VOM), so it is possible to have the ALMA correlator correlate the sum signal during online processing. This is done by co-opting some antenna that is not used by the VOM and placing the sum signals into the correlator as if it had come from this "sum antenna". The sum signal therefore appears as antenna "APP001" in any ASDM file containing ALMA data acquired with the VOM. There are a few things to note about the data and metadata for this antenna:

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

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

A final important note concerns delays between antennas. The ALMA control system estimates the total delays between all antennas and adjusts the signals through a variety of techniques (in hardware and software) so that the visibilities found in the data set have essentially zero delay on each baseline, (i.e. there is no slope in a plot of visibility phase as a function of frequency on any baseline.) Unfortunately, a (significant) component of the delay is removed with a frequency dependent phase rotation in the online correlation processing, not in the hardware, so there is some significant delay present in the signals at the TFBs where the phase corrections are made and correspondingly in the logic which creates the "sum antenna" signal. Thus in order for the phasing system to work properly, it is necessary to turn off this component of software delay correction and to instead correct it in the phasing calculations. To do this, the observing band is subdivided into a relatively large number of channel averages (at least 8) and the normal phase-solving procedure produces different, independent phases that can be applied to the TFBs corresponding to each channel average. This does remove most of the delay, but it does leave a small residual of phase-slope within each channel average. The only restriction is that the source needs to be bright enough in each of the channels so that a usable phase solution can be found. The flux limit if 0.5 Jy in Cycle 5 ensures that this is the case.

The activities of the phasing loop are recorded in an ASDM table as described above.

VLBI Observing Mode Analysis

As pointed out in Section 8.11, any VOM observation actually involves two concurrent interferometric observations: one on the ALMA scale that results in correlated data products between all participating ALMA antennas and one on a global scale, where ALMA serves as a station in a global VLBI array. A description of the reduction and analysis of ALMA VLBI data is largely beyond the scope of this document[16]. This process is, however, informed by the analysis of the ALMA-scale observation with regard to observatory performance (SEFD, system temperature, *etc.*), which for the most part may be obtained in the usual way. However, since the VLBI signal that is recorded is that of the "sum" antenna whose properties in turn depend on the phasing performance, a few details should be pointed out. The most significant point is that the effective area of the "ALMA dish" is roughly proportional to the square root of the number of phased antennas. (In Cycle 5 only 12 m dishes are to be used, and they have similar effective areas. If the 7 m dishes were to be included the relation is more complicated.) In practice, a number of effects conspire to lower performance from this ideal relation[17] to about 90% of what would be expected in a perfect system. A second point here is that the sum antenna signal is decimated to two bits from the individual antenna signals which were also decimated to two bits. Thus there is an additional reduction in sensitivity by 0.88 that must be taken into account in the VLBI analysis. In any case, the net effective area of the ALMA "sum" dish may be calculated from the ASDM on a per-scan basis.

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

In each table, there are temporal bounds on validity (*i.e.* the range of time within the scan of stable phase) and entries to list the disposition of the antennas amongst the various categories (*i.e.* phased antennas, reference antenna, and comparison antennas) and whether this represents a change from the previous scan. The table primarily reports a number of phase values, N_v, which is dependent on the "packing mode". In Cycle 5, there is one value per channel average per polarization per antenna. Additionally, the `CalAppPhase` table reports lower-level details of the phasing system, including online per-antenna assessments from TelCal of the antennas drawn from the quality of the phasing solution. This table may be used in conjunction with the ALMA data for a particular scan to calculate the phasing efficiency and thus to make an appropriate correction for the absolute VLBI correlation amplitudes.

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

8.12 High Frequency Observing

High-Frequency Observations (specifically Bands 8, 9, and 10) require some special observing techniques, mostly because of the combined ill effects of lower atmospheric transmission and fainter and more widely spaced calibrators. These techniques, which are just addenda to standard observing modes, will be progressively

[16]A VLBI analysis handbook is in preparation.
[17]ALMA-05.11.63.03-0001-A-REP.pdf

implemented in the ALMA observations for Cycles 4 and beyond, after some extensive testing is carried out. It is expected that all of them will improve not only the quality of the final data products, but also the overall observing efficiency at high frequencies. For reference, a brief description of the nomenclature that would be used by Contact Scientists and the ALMA project in general whenever they are discussed, reported or communicated to PIs is added:

- **High Frequency Cone Searches and Calibrator Surveys:** Observations with the high frequency ALMA bands require that the availability of calibrators near the target source are checked near the time of the planned observations. Searches for phase calibrators and check sources (see Chapter 10. details) within 10 degrees of the science targets (i.e., the so-called "Cone Searches") in each SB will be carried out by ALMA staff within 90 days of the high frequency observations. The results of those measurements will be immediately ingested into the ALMA Calibrator Catalogue, so that they are available for queries at the time of the actual observations.

- **Bandwidth Switching (BWSW):** For high frequency observations (Bands 8, 9, 10) with spectral set-ups including only narrow band spectral windows, the PI will be asked to add an additional wide band spectral window. This approach increases the accuracy of the gain calibration, allowing a solution across the combined spectral windows and/or transferring of the wide-window solutions to the narrow spectral windows.Whenever adding a spectral window is not possible (i.e., projects that have occupied all the spectral windows with narrow-band set-ups), it is possible to use bandwidth switching, in which the phase calibrator is observed with wider spectral windows for increased signal to noise. The wide-to-narrow spectral window complex gain offsets are measured by observations of the bandpass calibrator, allowing the phase cal solutions to be transferred. For more details, please read Chapter 10.

- **Sideband separation by 90 degree Walsh phase switching (Bands 9 and 10):** The Band 9 and 10 receivers are of DSB design, i.e. do not separate or supress either sideband, and signals from both sidebands are superposed on a single IF output per polarization. Prior to Cycle 5 the only option available for interferometric observing has been to use orthogonal LO frequency switching sequences to suppress the contributions of one sideband to the visibilities. From Cycle 5 a means of sideband separation is also offered, resulting in twice the output bandwidth and number of channels. The separation uses orthogonal sequences of 90 degree phase shifts applied at the first LO in each antenna, combined with synchronised accumulation over the sequence duration of the three phase combination states of the cross-correlations of each baseline by the correlators, and subsequent combination to produce the two sideband spectra in the correlator data processing software. The sequences are of length 128, with switching on 16 ms ticks, and therefore total length of 2048 ms. This mode therefore adds a constraint that the dump and integration durations must be a multiple of 2048 ms. Due to an unrelated requirement that subscans must start and end on 48 ms Timing Event (TE) boundaries, and 2048 ms not being a multiple of 48 ms, the allowed values of subscan durations are somewhat limited, as will be seen when adjusting times in the OT. In Cycle 5 only a single pair of spectral windows (LSB+USB) in each baseband is offered, and either the TDM (2 GHz, 256/Npol channels) or widest bandwidth FDM (1.875 GHz, 7680/Npol channels) modes must be used (additionally for FDM, the spw should not be moved from the center of the baseband). The modes in the four basebands remain independent, e.g. one baseband can use FDM to observe a spectral line, while the other basebands use TDM to detect continuum.

Chapter 9

ALMA Sensitivity Calculator

The main tool for calculating the sensitivity of ALMA is the ALMA Sensitivity Calculator (ASC). This is an application contained within the ALMA Observing Tool (OT) that allows a user (via a GUI) to experiment with various sensitivity options using the same methods that the OT employs internally to calculate time estimates *based on parameters entered into a project's Science Goals.*

Although the user may experiment with various sensitivity options (PWV octile, number of antennas, etc.) in the ASC, the final time estimate for a project calculated by the OT is not influenced by user-supplied entries in the ASC. For instance, the OT will always assume a fixed number of antennas for the particular Cycle, and will always use the standard precipitable water vapor (PWV) octile that is appropriate to the frequency of observation.

The ASC is also available as a web page that is linked on the ALMA Science Portal[1]. This stand-alone ASC differs from the GUI version incorporated in the OT in that it is possible to enter some parameters that are outside the ranges that are applicable to the current observing cycle.

9.1 Calculating the System Temperature

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

9.1.1 Sky Temperature

The OT's estimate of both the atmospheric zenith opacity, τ_0, and the sky temperature, $T_{\rm sky}$, are calculated using the Atmospheric Transmission at Microwaves (ATM) code[2]. This provides values of the opacity and the atmospheric "output radiance", in steps of 100 MHz, for the seven different octiles of PWV. The sky temperature is converted from the radiance using the Planck function and includes the contribution due to the CMB.

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

[1] `http://almascience.org/`

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

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

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

the emission from the atmosphere can be approximated as

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

Inserting the ATM values of $T_{\rm sky}(z = 0)$ and τ_0 into Equation 9.1 allows the mean physical temperature of the atmosphere, $T_{\rm atm}$, to be measured i.e.

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

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

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

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

$$T_{\rm n} = T \times \left(\frac{h\nu/kT}{e^{h\nu/kT} - 1}\right) \tag{9.4}$$

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

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

9.1.2 CMB Temperature

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

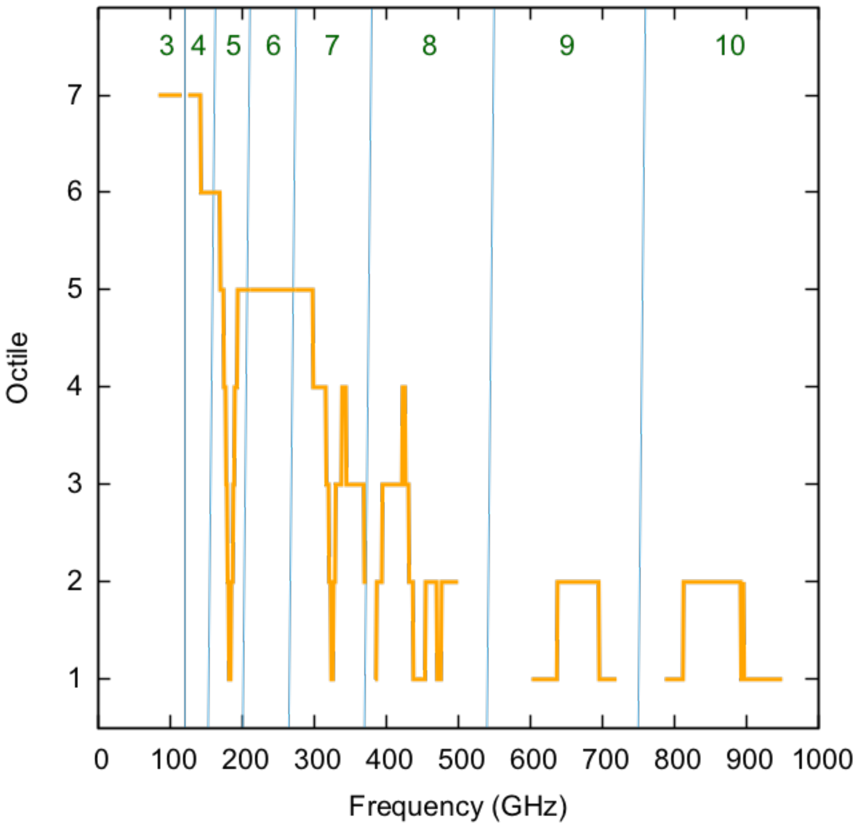

Figure 9.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown at the top of the plot. The water line at 183 GHz (Band 5) is particularly prominent. In general, higher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

9.1.3 Receiver Temperatures

For ALMA Bands 4 and 10, the calculator currently only uses the specifications for the receiver temperatures (required over 80% of the receiver bandwidth) and not the actual measured values. For ALMA Bands 3, 5, 6, 7, 8 and 9, however, typical values measured in the laboratory are used as these are usually significantly better than the specifications. The measured values are somewhat conservative and so are in between what might be expected at the middle and edges of the bands. The values used in the ASC are given in Table 9.2.

Note that single sideband noise temperatures are reported for Bands 1-8 and double sideband temperatures for Bands 9 and 10.

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

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

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

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

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

9.1.4 Ambient Temperature

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

9.1.5 DSB Receivers

Due to the way that the astronomical signals are down-converted to an intermediate frequency, every heterodyne radio receiver simultaneously detects radiation from two sidebands. This means that, if nothing were done to prevent it, a spectral window processed by the correlator would contain two sets of astronomical signals mixed in with one another, one from the "signal" sideband and an undesirable one from the "image" sideband. In the case of 2SB receivers , the contribution from the image sideband (emission from the source and noise) is suppressed to a very high degree. However, for DSB receivers (Bands 9 and 10) it is only possible to remove the source contribution (either by LO offsetting or 90-degree Walsh switching) and so the noise cannot be neglected.

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

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

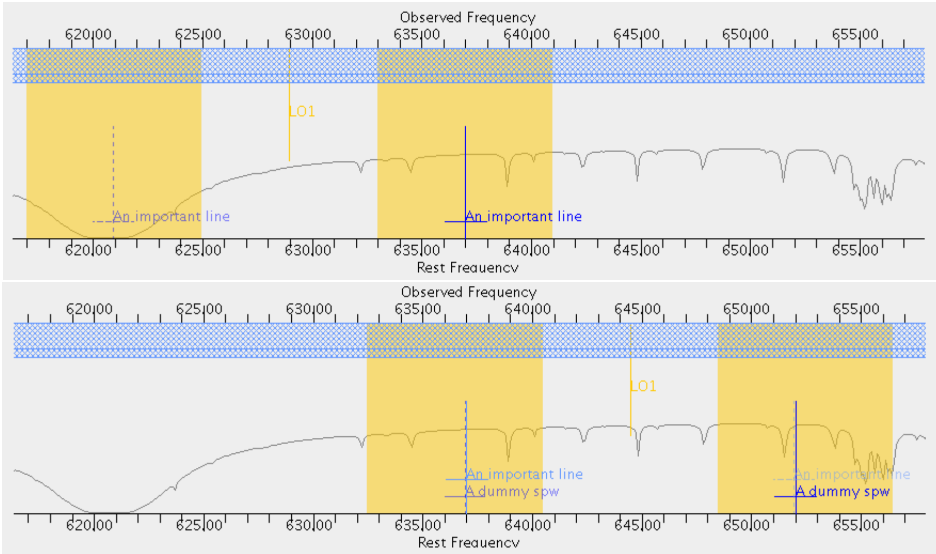

Figure 9.2: A Band 9 spectral setup as displayed in the spectral visual editor in the OT. The top figure shows an example of how the image equivalent of a single spectral window can fall into an area of poor atmospheric transmission, leading to much higher $T_{\rm sys}$ than necessary. Placing a dummy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

9.1.6 System Temperature (OT version)

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

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

where

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

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

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

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

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

9.1.7 System Temperature (Stand-alone ASC and GUI (OT) version)

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

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

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

9.2 The Sensitivity Calculation

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

Under certain circumstances it will not be possible to achieve the theoretical sensitivity, particularly if the source being observed is very bright and/or the u, v coverage relatively sparse. In these situations the image can be dynamic-range limited due to residual source flux scattered across the map as a result of calibration errors and imperfect source deconvolution.

9.2.1 12-m and 7-m Arrays

When dealing with the 12-m and 7-m Arrays, the point-source sensitivity, $\sigma_{\rm S}$, is given by the standard equation:

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

The various parameters are:

- w_r – robust weighting factor. Pipeline imaging and subsequent QA2 assessment is performed assuming that the visibilities are weighted using robust weighting, specifically a Briggs robustness factor of 0.5. Simulations have shown that this factor is equal to 1.1.
- $A_{\rm eff}$ – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency where $\eta_{\rm ap} = R_0 \exp(-16\,\pi^2\sigma^2/\lambda^2)$ and σ is the rms surface accuracy of the antenna. The latter is set to the design goal of 25 μm and 20 μm for the 12 and 7 m antennas respectively[3]. R_o is the product of a number of different efficiencies and is equal to 0.72. See Table 9.2 for values of antenna efficiencies and effective areas in various ALMA bands.
- f_s – shadowing fraction. For the more compact 12 m configurations and the ACA 7-m Array, antennas can block the field-of-view of other antennas in the array and thus reduce the total collecting area. The shadowing fraction is a function of source declination as shown in Fig. 7.5.
- $\eta_{\rm q}$ – 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.

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

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

Table 9.3: Aperture efficiencies at typical continuum frequencies for both the 12 and 7 m antennas. The effective area, A_{eff}, is equal to η_{ap} multiplied by the physical area of the dish i.e. 113.1 m^2 and 38.5 m^2 for the 12 and 7 m antennas respectively.

- η_{c} – correlator efficiency. This depends on the correlator (64-input or ACA) and correlator mode, although the efficiency of all 64-input Correlator modes is equal to 0.88. The ACA efficiencies *do* depend on the mode, but this is only taken into account in the OT, not by the stand-alone ASC (see below) which therefore also assumes a value of 0.88.
- N – number of antennas. This defaults to 43 for the 12-m Array and 10 for the 7-m Array.
- n_{p} – number of polarizations. $n_{\text{p}} = 1$ for single polarization and $n_{\text{p}} = 2$ for dual and full polarization observations.
- $\Delta\nu$ – resolution element width. As already mentioned, this should be equal to 7.5 GHz for continuum observations. This is due to the maximum usable bandwidth of a spectral window being limited to 1.875 GHz by the anti-aliasing filter through which the baseband signal passes. $n_{\text{p}}\,\Delta\nu$ is often referred to as the effective bandwidth.
- t_{int} – total on-source integration time.

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

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

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

$$\Omega = \frac{\pi\theta^2}{4\ln 2}. \tag{9.10}$$

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

9.2.2 Total Power Array

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

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

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

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

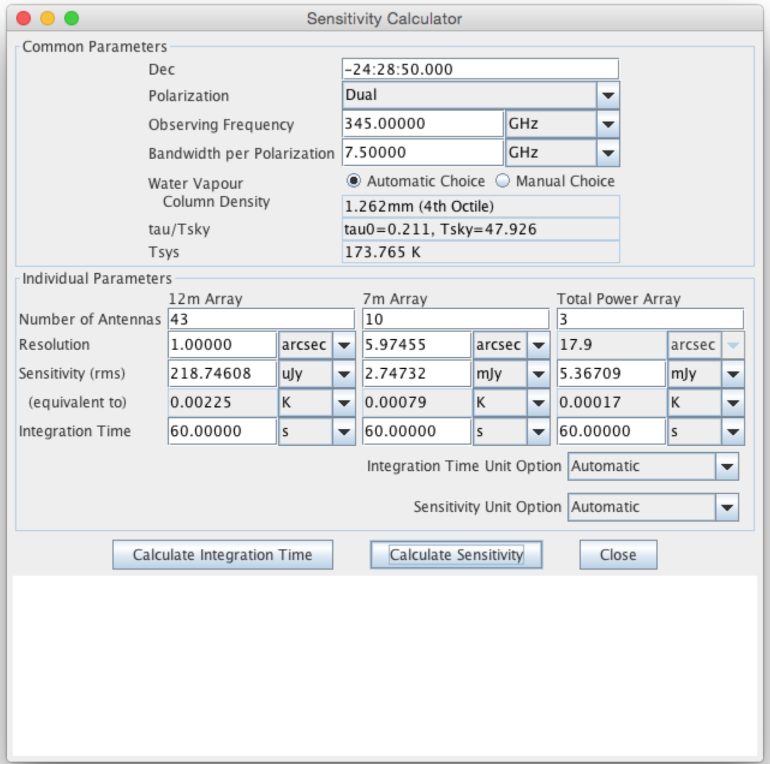

Figure 9.3: Screenshot of the GUI version of the ALMA Sensitivity Calculator as implemented in the ALMA Observing Tool. The white area at the bottom is for displaying error messages i.e. parameters out of bounds. The example here shows the achievable sensitivity for all three arrays for an on-source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

9.3 User Interface

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

- Source declination – this is used to calculate the maximum elevation of the observation and thus the minimum airmass i.e. the ASC assumes that the source is transiting. It is also used to calculate the amount of shadowing that is likely to affect the ACA 7 m and the smaller 12 m configurations.
- Observing frequency – this sets the receiver temperature, antenna efficiency and the PWV octile.
- Bandwidth per polarization – this otherwise straightforward parameter should be set to 7.5 GHz for continuum observations (see Section 9.2.1). For spectral line observations, it is usually set to the frequency/velocity resolution that one requires in one's spectrum.
- Water column density (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 43 from the 12-m Array, 10 from the 7-m Array and 3 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, $T_{\rm sky}$ (including the correction for the source elevation) and $T_{\rm sys}$ that correspond to the entered frequency and PWV.

9.4 Total Time Estimates

Note that the time calculated by the ALMA Sensitivity Calculator does not account for telescope overheads (calibration, software and hardware latencies, etc.) and therefore the time is always assumed to be the true on-source time. Chapter 5 of the OT User Manual should be consulted for details on how the total time required to observe and calibrate an ALMA project is calculated.

Chapter 10

Calibration and Calibration Strategies

10.1 The True and Observed Visibility Function

10.1.1 The Data Format and the Correlator Model

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

The *true complex visibility function*, $\mathcal{V}^t$, is related to the celestial radio emission angular and spectral distribution, and the relationship, an approximate Fourier transform, is described in the Chapter 7. The goal of the present chapter is to describe methods to retrieve the true complex visibility function from the observed complex visibility function.

The fastest time dependence of the observed visibility function is produced by the diurnal motion of the radio source which changes the phase difference of the signals in each antenna pair. This response is quasi-sinusoidal and is called the *natural fringe*, often with a rate of several kilohertz. This fast response can be approximated by assuming a position in the sky, σ_0, called the *phase center*, the positions of the appropriate antenna-pair and the observing frequency, and of course the earth rotation rate. This algorithm is called the *correlator delay model* and is described in more detail in Chapter 5. Other, smaller contributions to the visibility phases; such as the atmosphere delay over each antenna and phase or delay changes in the antenna electronic system, can be also included, if well-known. After this correlator model is subtracted from the natural fringes, the (residual) visibility function is a much slower function of time and can then be averaged over many seconds, with no loss of information about the angular distribution of the observed celestial object. The resultant image position of the celestial object will then be obtained with respect to the assumed phase center position. The most convenient location of the phase center is antenna pointing and delay center positions used for the observation.

10.1.2 The Visibility Function Dependence

The observed visibility function, $\mathcal{V}^o$, associated with each baseline pair (i,j), with each scan (see Section 8.2) and set of band frequencies is

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

where:

i, j is baseline pair between antenna i and j,

t_k is each time-stamp in the scan,

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

where s is a spectral window (spw),

and c is the channel index within each spw (from 1 to a maximum of 4096),

p_q represents up to a maximum of the four linear polarizations products (XX,XY,YX,YY).

The complex visibility function is often described by its amplitude, $\mathcal{A}^o$, and its phase (ϕ^0) components since virtually all of the calibrations change its amplitude and phase somewhat independently, rather than real and imaginary parts independently.

The *autocorrelation* data, for the case where $i = j$ is also contained in the data set. The calibration and use of this data are given in Section 8.6 and Chapter 13.

10.2 Main Interferometric Calibration Strategies

10.2.1 Phase Referencing

Most ALMA interferometric observations use phase referencing as the major calibration method. This method interleaves scans between the science target and a nearby calibrator source that has an accurately known position, spectrum and flux density, and is sufficiently small in angular diameter. Since the true visibility of the quasar can be modelled, its observed visibility indicates the entire changes induced by the ALMA system, including that of the atmosphere. If the calibrator and target are only several degrees apart and a typical switching time between them is several minutes, then the measured induced changes from the calibrator data will be approximately that for the target, hence its true visibility can be estimated.

10.2.2 Antenna-based Calibrations

Nearly all of the changes from the true visibility function to the observed visibility function are associated with changes in the signal flow in each antenna, from the radio wave path in the troposphere above that antenna, through all of the antenna electronic processing, then into the correlator that combines each antenna-pair of signals to obtain the visibility function for each baseline. Hence, the calibration equation that describes these changes to each baseline visibility function can be written solely in terms of the changes in the two antennas associated with the baseline pair[1]:

$$\begin{aligned} V^o_{(i,j)}(t_k, \nu_c) = G_i(t_k, \nu_c) * \tilde{G}_j(t_k, \nu_c)\ V^t_{(i,j)}(t_k, \nu_c) \\ +noise + g_{(i,j)}(t_k, \nu_c)\ V^t_{(i,j)}(t_k, \nu_c) \end{aligned} \quad (10.2)$$

where $G_i(t_k, \nu_c)$ is the antenna-based complex gain calibration as a function of time and frequency. This is the normal antenna-based calibration equation. The second line in equation 10.2 contains the effect of noise and any additional baseline-dependent calibrations, and these will be ignored for the most part. The polarization dependence (the calibration relationship between the X- and Y-polarization) is discussed in Section 10.5.

The complex antenna gain can be separated into the antenna-based amplitude gain, G and the antenna-based phase, ψ component,

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

where $a_i(t_k, \nu_c)$ is the antenna-based amplitude correction and $\psi[t_k, \nu)$ is the antenna-based phase corrections, both as a function of time t_k and frequency ν_c for each polarization. The symbol j is the imaginary unit. The antenna-based gains and phases are the major calibration quantities that are described in this chapter.

[1]The radio emission has two polarizations which are can be detected in ALMA into two linear parts, X and Y. The equations are for either polarization. For polarization interferometry, additional terms between the X and Y polarizations are discussed later.

10.2.3 Major Calibrations Affecting Interferometric Data

The modifications of the information in the radio waves from a celestial source occur at many locations along the path from the source to the data archive. A list of components associated with signal changes, based on G. Moellenbrock's 2014 Socorro Summer School Lecture[2] is:

I = Ionosphere: Opposite rotation of linear polarization position angle caused by propagation through the ionosphere—not significant at frequencies above 70 GHz.

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

P = Antenna Voltage pattern: The relative response across the reception area of the antenna beam which can be lightly modified because of pointing and focus errors, and by the quality of the antenna surface.

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

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

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

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

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

D = Instrumental delay offset of the signals between antennas.

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

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

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

Do not be intimidated by the above list. The most important ALMA calibrations deal with the temporal gain G and the bandpass response B. The calibration process is significantly simplified because the temporal dependence (GT) and frequency dependence (BF) for nearly all instrumental changes are only *lightly coupled*, so the variations in time, t_k and frequency ν_k can be determined independently to first order:

$$G_i(t_k, \nu_c) \approx GT(t_k) * BF(\nu_c) \tag{10.4}$$

10.2.4 The Organization of the Calibrations

There are three major stages of calibrations in ALMA.

A-priori Calibrations. These are calibrations that are associated with relatively fixed properties of the array; for example the antenna positions, pointing, focus and delay models, and calibration parameters of various devices throughout the system. These are measured periodically by the ALMA staff using methodologies briefly outlined below.

[2]https://science.nrao.edu/science/meetings/2014/14th-synthesis-imaging-workshop/lectures-files/nbsp;nbsp;nbsp;MoellenbrockCalibration2014_FINAL.pdf

Online parameters and calibrations. Various calibrations and corrections are performed during observations. Some happen completely transparently, such as regulating the fiber path length from antenna to correlator by the Line Length Correctors (LLCs), and small pointing corrections by metrology systems in each antenna to compensate structural deformation by uneven temperatures, wind loading etc.

Some explicit calibration scans are included in the observing sequence, which are clearly visible in the ASDM and can be re-processed offline if needed. Most observations include pointing calibration scans, after which the measured pointing offsets (only collimation terms) are applied as an increment to the antenna pointing models. Some observations similarly contain focus calibration scans, although in most cases focus is calibrated in separate observations throughout the day to allow more flexible scheduling of these calibrations. Atmosphere calibration scans are included, from which $T_{\rm sys}$ spectra are computed online and provided in the ASDM for application in data reduction offline. Auxiliary calibration results are also computed online from various scans primarily for quality assurance purposes, for example phase variability statistics from scans on bright calibrators. When online WVR correction is used, parameters to facilitate the simple linear approximation of path length with WVR channel temperatures are also computed from many scans and updated in the correlator software.

Offline calibrations. The offline calibrations are those determined after the completion of an observation. These calibrations are associated with special-purpose *calibrator* scans that are embedded within the sequence of scans. The three major types are: (1) the removal of the frequency dependence across the measured bandwidth, B, using a strong calibrator with no spectral features, known as the *bandpass calibrator*; (2) the removal of remaining amplitude and phase temporal changes during the experiment, G, using the phase referencing technique; (3) the conversion of the measured visibility amplitude to Jansky flux density units, best accomplished by a short observation of a source of known flux density.

A detailed description of each one of these major stages are given in Sections 10.3 to 10.5. Special calibration techniques are outlined in Section 10.6. Following each calibration item, its corresponding category according to the classification in Section 10.2.3 is given in parenthesis.

10.3 ALMA-Supplied a-priori Calibrations

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

10.3.1 Antenna Beam/Surface Characteristics (P)

The relative sensitivity of each antenna as a function of azimuth and elevation from the pointing direction is called *primary beam pattern*, which is described in Section 3.2. This sensitivity correction is needed for large angular source emission, and for multi-pointings observations (mosaicing) with overlapping fields. The measurements of the pattern are made by mapping the antenna beam using a bright radio source, either using single dish scanning around a strong source or with interferometric observations at many points over an area that extends well outside of the main primary beam area. The measurements are processed in special-purpose software that estimates the accurate antenna surface topology, and adjustments of the antenna panel supports can be made if the surface is not within specifications. The measured primary beam patterns for each antenna-type are available in the imaging software packages and applied when producing the sky-corrected images (see Chapters 3 and 7 for more details).

10.3.2 Antenna Pointing Models (P)

All-sky pointing of the array elements at a nominal accuracy of $2''$ RMS is achieved by a combination of metrology systems inside the antennas themselves to account for variable deformations, and static models maintained by the observatory. The static models are dominated by six terms: azimuth encoder/pad offset ("IA" term, of order 1000″), elevation encoder/collimation offset ("IE" term, order of few hundred arcsec), azimuth collimation offset ("CA", order few tens of arcsec), tilt in North-South and East-West directions "AN" and "AW" terms, order of ten arcsec), and elevation sag ("HECE") term, order of ten arcsec). The full pointing model has 17 to 19 terms, most of which are at the arc-second level. The pointing differences between receiver bands require only a collimation offset (CA, IE terms) with respect to a reference band (currently Band 3).

The complete pointing models are checked and adjusted on weekly timescales, primarily by making "all-sky pointing runs". These observations carry out a series of around 50 pointing calibration scans in Band 3 on bright quasars of accurately known position, well distributed over the sky. The total observing time is about one hour and each pointing measurement scan consists of a measurement of the quasar visibility amplitude at five subscans at different antenna pointing offsets (depending on frequency and antenna diameter) from the assumed pointing direction. From these data, the pointing offset for each quasar scan for each antenna can be determined to an accuracy of around 0.3″. The fit of these offsets to the full pointing model has a typical rms error over the sky of $< 2''$. Results from pointing calibration scans in science and other observations are also used to supplement dedicated measurements for maintaining the pointing models.

The pointing models are stored in the Telescope Monitor and Configuration Data Base (TMCDB) which includes a history of all changes. Currently only night time pointing measurements are considered in these *static* models. During daytime observations, the pointing changes with time due to differential antenna heating, and the modeling of the pointing is not yet well-characterized.

After each antenna relocation, the pointing changes are primarily confined to the azimuth offset and tilt terms (IA, AN, AW). The tilt is measured after the relocation by spinning the antenna in azimuth and fitting sinusoids to the two axes outputs of inclinometers located above the azimuth bearing. This gravitational tilt is converted to an astronomical tilt by adding a local gravity vector (which varies over a 5-km distance across the array due to the effect of the nearby mountains). With the tilt updated in the pointing models, the azimuth offset is determined by fitting total power data recorded while the antenna performs azimuth strokes across a planet or bright quasar. After this the all-sky pointing is typically accurate to within a few arcseconds, and it is later refined as described previously.

In all science executions, additional short pointing measurements are made as needed using a bright quasar near each of the science targets and calibrators. The pointing observations are currently made in the target observing band up to Band 8, but in Band 7 for Bands 9 and 10 due to insufficient sensitivity to quasars. Up to Cycle 4, pointing scans were made at a fixed tuning for each band, but from Cycle 5 it is hoped to use the same first LO frequency as the science tuning in order to reduce re-tuning overheads where atmospheric transmission makes this optimization feasible. The criteria for whether to do pointing calibration scans are more strict at higher frequencies. These additional pointing offsets are at the arcsecond level, and are applied in a temporary *auxiliary pointing model*, which is added to the master static TMCDB pointing model for each antenna. With these residual offsets included, the pointing accuracy for all scans is intended to be within 0.6″ RMS. The total pointing model used for any scan (sum of primary and auxiliary models) is stored in a `PointingModel` table in the ASDM datasets.

10.3.3 Focus Models (P)

As for pointing, static models for antenna focus (subreflector position) are stored in the TMCDB, and offsets applied while observing as needed. Unlike for pointing, we do not use active metrology systems in the antennas themselves to correct the focus due to unpredictable path length changes that may result.

The focus models only require six main terms, plus offsets for each receiver band. The models contain an offset for each axis which describes the optimum subreflector position for Band 7[3] at elevation 50° and temperature 0° C ("XR", "YR", "ZR", typically a few mm in size). In the axial (Z) and elevation-oriented direction (Y) there is a coefficient of sag with elevation ("ZS", "YS", up to a few mm in size and depending on antenna type). For the Z axis there is a temperature coefficient ("ZT"), which well tracks the night time (thermal

[3]Band 7 is used as focus reference for three reasons: it typically gives the smallest measurement uncertainty in an individual focus scan, Band 7 is located with the high frequency receivers near the optical axis, and Band 9 and 10 observations tend to depend on focus scans in Band 7 to guarantee good S/N. The bottom of Band 7, around 290 GHz is used to optimize measurement accuracy for a wide range of observing conditions.

equilibrium) variation of the focus. The band offsets are mostly small, around 0.1 mm, with the exception of the axial focus of Band 4 which differs by around 1 mm from the other bands as a compromise to avoid having to use a lens on the feed horn as has to be done for Band 3 and below to save space. These models are maintained by long- term compilation of focus measurements in all receiver bands for the antennas, with care taken with regard to hardware changes which may change the focus characteristics.

Focus measurements are traditionally made in a way similar to those described for pointing measurements, except the subreflector is offset (in one axis per scan) relative to the nominal model position. The subreflector position in the axis being offset which yields maximum gain is solved for each antenna from the cross-correlation data. A holographic technique using interferometric beam maps which simultaneously derives the three axis offsets in addition to pointing and illumination offsets, which does not require the subreflector to move, is also available, but is not commonly used yet due to concerns over dependence on atmospheric path length stability which the traditional technique is largely immune to, and requirements on calibrator fluxes for higher frequencies. These measurements are made by the observatory as needed to maintain the models, and routinely while observing to track non-equilibrium focus variation during daytime and sunrise/sunset and to correct for any errors in the TMCDB model. Most focus measurement observations only perform a scan for the axial (Z) axis, as the antenna gains are much more affected by the Z-axis variations than variations in the lateral axes (usually the measurement uncertainty from a single measurement scan in the lateral axes is at the same level as the error by just using the TMCDB model).

The focus model used during an observation is stored in a `FocusModel` table in the ASDM datasets.

Further improvements to focus calibration are desirable. As for pointing, performance in daytime and around sunrise/sunset is problematic and relies on frequent measurement and correction. Correlations between focus changes and temperature sensors at various locations in the antenna structures are being studied in the hope of improving the daytime modelling, which would improve daytime high frequency performance and improve daytime observing efficiency. In addition to the translational axes, the subreflectors can also be tilted, which will give a small reduction in spillover for the low frequency bands that are far from the optical axis (Bands 1–6). This capability may be implemented in future Cycles.

10.3.4 Antenna Delay Offsets (D)

As for pointing and focus, models for the antenna delays are stored in the TMCDB. These include an overall pad delay of up to tens of microsec for the most remote stations, plus delays for each signal path combination, each value typically up to a few hundred picosec. When an antenna is integrated after relocation to a different pad, the delay is obtained a-priori to within a few tens of nanosec using a hardware timing mechanism of the digital transmission system (DTS), and after that the remaining delay error is found by interferometric observation of a very bright continuum source in Band 3 with a medium spectral resolution FDM mode, which can reliably measure delay errors up to a few hundred nanoseconds. The signal path delays are obtained by low spectral resolution observation of a continuum source in a variety of tunings to cover all the signal path combinations.

Checks of the delays are usually performed at the start of an observing session, probing all of the backend signal paths using four scans in Band 3 (an "IF Delay" observation), and for the frontend bands that are to be used by making a scan in that band.

The typical spread of delay errors in science observations is around 50 ps, corresponding to 36° of phase slope over a 2 GHz baseband. This is what you should expect to see in the phase component of bandpass calibration solutions, as described below. The primary reason for maintaining such low errors is to allow spectral-averaged data to be used for processing calibration scans online (pointing, focus, phase statistics etc.). Repeatability of delay measurements with ALMA is actually at the 1 ps level, even over multiple days. However, in reality the "delay" measurements include a component which is just the phase vs. frequency response of the receiver systems, which differs from tuning to tuning, and results in a spread of some tens of picoseconds even if delays are set to 1ps accuracy at a nominal tuning. The receiver phase response spectra also mean that the delay measurement errors derived from residual phases after the fit are typically around 20 ps, and far greater than the repeatability.

10.3.5 Antenna Positions (M)

The relative positions among the ALMA antennas must be known to a small fraction of a wavelength in order for the delay model M (see Section 10.3.4) to be accurately calculated during an observation[4]. This corresponds to an accuracy of about 0.1 mm.

The antenna positions are determined in a two-step procedure. The first step is made within a day or two of the antenna relocation to a new pad. Since the a-priori value of an antenna position error after the move can be as large as 10 mm, its position must be updated to about 1 mm accuracy so it can be used for the additional calibrations to incorporate the antenna into the array. One of the initial calibrations, that of determining the antenna pointing parameters over the sky, can also be used to determine the antenna positions. The positions are stored in the TMCDB and are the most critical part of the delay model that determines the observed visibility phases for all observations. The antenna positions are the (x,y,z) distance of the intersection of the two axes of the antenna from the Earth geocentric position.

When all of the antenna relocations to an array configuration have been completed, the second step is to determine accurate antenna positions of about 0.1 mm accuracy. This is done with a one-hour *all-sky-delay* observation at 88 to 100 GHz of about 50 quasars distributed over the sky, and when the atmospheric phase stability is excellent.

For observations that are made before the most accurate antenna positions have been determined, there are offline procedures using the Pipeline or Manual reductions that compare the antenna values in the TMCDB used for the observation with values from the best all-sky delay observation, which are stored in an accessible computer file, to correct for the known antenna position offset during the observations.

The repeatability of antenna position measurements in long baseline configurations is under active investigation. Position variations of over a millimeter are often seen on a variety of timescales for the most distant antennas, which is thought to primarily be due to variation in zenith path delay through the troposphere above each antenna. When analyzing observations made in the longest baseline configurations, it is quite possible you will see effects of position errors at around 1 mm level for the most distant antennas (i.e. phase slope in phase calibrator scans after bandpass calibration).

10.4 On-Line Calibrations

10.4.1 Water Vapor Radiometer Measurements (M)

Fluctuations in the line-of-sight precipitible water vapor (PWV) of an antenna cause significant delay corrections, up to 0.3 mm/sec, corresponding to 30° of phase/sec at 90 GHz, and scales in phase linearly with frequency. The changes are driven by wind, and cover a large range of spatial scales from 5 m to many km. All 12 m ALMA antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures the emission of the atmospheric water line along the same path of the other ALMA bands, in four spectral frequencies near 183 GHz with a sampling of about 1 Hz. These radiometers are described in Nikolic, B., et al., 2013, A&A, 552, 104. The CM antennas do not have WVR receivers. Tests to using interpolated WVR data from the PM antennas that surround the CM array suggested any corrections on these short baselines are insignificant.

The conversion of the water vapor emission into appropriate delay changes (M) depends on the atmospheric temperature and pressure and is measured using two systems that are discussed in the following section.

10.4.2 System and Receiver Temperature(T)

At millimeter and submillimeter wavelengths, the additional random noise associated with a measurement has two major causes: the noise associated with the receiving system, $T_{\rm RX}$, and the noise associated with the

[4]To reiterate, the observed visibility phase is the difference between the observed quasi-sinuoisoid correlated signal between any two antennas with that of the correlator delay model. The difference should be small ($< 60°$) and not too variable ($< 10°$ per minute) for accurate calibration and imaging of celestial sources

atmosphere $T_{\rm sky}$, which acts like a black body emitter, and attenuates the astronomical signals. This sky noise is a strong function of frequency, elevation, the column of wet and dry constituents of the atmosphere, and the temperature of the atmosphere. The noise from the receivers can also vary for a given tuning with factors such as cryostat temperature and pointing direction with respect to the local magnetic field.

To measure these effective temperatures, the ALMA and ACA front ends are equipped with an Amplitude Calibration Device (ACD) that moves a hot and an ambient blackbody load in front of each receiver feed. The description of the system temperature electronics is given in Chapter 9, equation 1.7. The noise measurement from these two loads give the total system temperature, $T_{\rm sys}$, and the receiver temperature, $T_{\rm RX}$ for each spectral channel, and significantly effect the observation sensitivity, given in Chapter 9, equation 1.7. The accuracy of these measurements depends on the linearity of the autocorrelation data with input power, which currently limits these observations to the lower spectral resolution TDM mode (256 total channels, divided amongst 1, 2 or 4 polarization) of the baseline correlator[5].

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

The on-line measurements of T_{sys} and T_{RX} are stored in the `CalAtmosphere` table in the ASDM datasets, and are displayed in real-time at the telescope to show measurement problems immediately. The system temperature measurements can also be re-derived off-line from the original hot and cold load data if there were on-line conversion problems or if a modified processing algorithm is used. The T_{sys} spectra, suitably interpolated between time and sources, is then applied off-line to the correlation data. Assuming that the correlated data is in units of percentage correlation, multiplication by the T_{sys} will change the correlated data units to Kelvin.

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

10.4.3 Additional Antenna Pointing Checks (P)

As discussed in Section 10.3, corrections are determined for the antenna pointing after each antenna move and approximately once per week. The nominal all-sky pointing accuracy is about 2.0 arcsec. In order to obtain ALMA's 0.6″ RMS pointing accuracy specification, additional pointing observations may be made before each scan. To some extent these additional pointing calibrations compensate for shortcomings in the antenna metrology systems, and errors in the static models. These pointing measurements are automatically included in the observation schedules, with no PI input. At present a pointing calibration is triggered when moving to a target more than 25 degrees from the previous pointing calibration.

The results of the pointing measurements in each execution are stored in the `CalPointing` table in the ASDM datasets, and are displayed as they are produced during the observations to allow problems to be immediately seen on-line. These results can be used for off-line data flagging of antennas: if a correction is deemed too large; the measurement uncertainty is too high; or the pointing result is missing.

[5]In TDM mode the autocorrelation zero lags can be used to fully compensate for quantization effects on the autocorrelations. However, in FDM mode an extra filter bank is inserted before the correlations so an external measure of the total digitized power is needed. The current plan is to use the baseband square-law detector data to provide this extra input for Cycle 5.

10.5 Off-Line Calibrations

Off-line calibrations, those made after completion of an observation, are described in other chapters (QA0, QA2, Pipeline, Manual Reductions). The data format conversion of the ASDM to the measurement set, and its archive are also given in other chapters. In this chapter an overview of the calibrations and editing are described.

The compilation of the radio sources that are used for many of the off-line calibrations are found in the ALMA catalog that contains over 15,000 objects, south of $\delta = 45°$, nearly all of them small-diameter quasars. Information on calibrators near a specified target can be selected, plotted and downloaded to aid in the selection of the best calibrator(s) for any target[6].

The ALMA catalog is compiled from many major catalogs, mostly at frequencies between 1.4 and 23 GHz, over the last 30 years. The catalogs include VLA calibrator catalog, at20g catalog, vlbi (Petrov) catalog, CRATES 8.4 GHz compilation and the SMA catalog. Nearly all of the brightest 2000 sources in the ALMA catalog have been observed with VLBI techniques and thus have positions that are accurate to $< 0.002''$, all measured on the ICRF2 [7] coordinate frame. It is recommended that all objects observed with ALMA use their ICRF positions.

10.5.1 On-Line Flagging (F)

The on-line calibration system uses the results of many monitoring devices to indicate periods when data should be flagged; for example, when an antenna pointing or focus settings are offset more than a specified amount from the calculated settings, or when relevant local oscillators are not locked. These flags are stored in the archive data base in two ways, xml and binary flags. When converting from the ASDM format to the MS format, the two sets of flags (applyflags = T and bdfflags = T) can be inserted into the data or stored in a separate flag table for later use.

10.5.2 Calibration-based Flagging (F)

During most of the off-line calibrations, described below, additional time periods and antennas that require flagging are often found. Examples are: noisy edge channels in the spectral windows, low relative signal in one antenna compared with the others, and phase jumps in time. In most cases, these bad data periods can only be seen in the calibrator observations; hence, most flags associated with a calibration scan should be *appropriately* extended to other source scans, spws and baselines.

10.5.3 Bandpass Calibration (B)

The correction of the frequency response from flatness across each spw and polarization is the purpose of the bandpass calibration. First, there is a *residual* amplitude and phase variation with a characteristic length of tens of MHz of about 10% and 5 deg. These remain stable over many hours, and are nearly independent of the sky position of the radio source[8]. However, within about 5% of the band edge, larger changes occur and the appropriate channels should be flagged out. Second, an additional phase slope in the bandpass is produced by a offset delay from that assumed[9].

The residual bandpass for each spw and pol is determined from a 5 to 30 min scan of a bright calibrator source with a *known radio spectral index* and no spectral features. From the CASA task *bandpass* the relative amplitude and phase for each antenna/spw/pol as a function of frequency is determined. The nominal bandpass accuracy is $< 0.2\%$ and $< 0.5°$ over the spectral resolution requested. If an observation contains several tunings or a large Doppler change, the bandpass for each setting must be obtained.

[6] https://almascience.eso.org/sc/
[7] /www.iers.org/IERS/EN/DataProducts/ICRF/ICRF2/icrf2.html
[8] http://almascience.org/documents-and-tools/alma-technical-notes/ALMATechnicalNotes15_FINAL.pdf
[9] The major delay offset is corrected on-line once a day so that the residual phase delay slope is generally small

Lines at known frequencies from components in the atmosphere will produce apparent bandpass amplitude changes that are similar for all spw's and pols. These lines will also produce system temperature increases. These data can be used but caution is necessary, since the line amplitude does depend on the source elevation, so it may differ between the bandpass, calibrator and the target sources. Generally, the bandpass phase is not affected by atmospheric lines.

In Cycle 5, the bandpass calibrator is chosen from a set of 35 bright quasars, called *grid* sources and the list is given in Table 10.1

Grid Sources Used for Bandpass

J0006−0623	J0237+2848	J0238+1636	J0319+4130	J0334−4008
J0423−0120	J0510+1800	J0519−4546	J0522−3627L	J0538−4405
J0635−7516L	J0750+1231	J0854+2006	J1037−2934	J1058+0133
J1107−4449	J1146+3958	J1229+0203L	J1256−0547	J1337−1257
J1427−4206	J1517−2422	J1550+0527	J1617−5848	J1642+3948
J1733−1304	J1751+0939	J1924−2914	J2025+3343	J2056−4714
J2148+0657	J2232+1143	J2253+1608	J2258−2758	J2357−5311
L = large scale emission				

Table 10.1: The 35 Grid Sources Used by ALMA

10.5.4 Gain Calibration (G)

Nearly all ALMA observations use the phase referencing technique to measure the residual antenna-based temporal amplitude and phase variations. The weather conditions (transparency of the atmosphere and magnitude of the phase fluctuations) are monitored continuously so that the chosen PI observation can be successfully phase-referenced. The effectiveness of the on-line WVR-corrections is critical, often removing up to 75% of the short term phase changes. See Chapter 11 for more descriptions of the on-line observational checks.

The choice of calibrator is critical to the success of the reduction and imaging. Generally, the best calibrator is that closest to the target that can determine the antenna-based solutions with an SNR$>$ 15. The calibrator parameters needed for the optimum phase reference choice are given in Table 10.2. For many targets the calibrator choice can be automatically selected from the ALMA catalog and is described in Section 8.1. However, for long baseline and high frequency targets, the availability of a suitable calibrator is checked by the JAO-staff months ahead of the proposed observation. These additional observational checks are described in Section 10.6.

The CASA task *gaincal* is used to determine the antenna-based amplitude and phase gains from each phase calibrator scan. Several recommendations are: (1) Pick a reference antenna near the center of the array that has not been recently moved[10]. (2) If the calibrator is a relatively strong quasar, antenna-based solutions can be generated accurately for each spw and polarization. (3) For a weak phase calibrator for which more than 5% of the data are flagged using each spw solution separately, use the bandpass observation to align the phases of the eight data streams (4 spw, 2 pol's) which can then be added (combine='spw'; gaintype = 'T' to give an SNR increase of 2.8) before determining a solution. This will enable the use of calibrators that are significantly closer to the target.

After using gaincal in a manual or pipeline script, the plots from CASA tasks plotcal or plotms should be inspected. The ALMA pipeline can generate appropriate flags from anomalous calibrator values but the PI should inspect the flags since not all problems are found by the pipeline. Some guidelines in the interpretation of the calibration plots are:

(1) The antenna-based amplitudes should not vary by more than about 10% over the observation period, and not differ by more than about 20% among the antennas. For antennas with low gain values or obvious drop-outs, check on possible reasons: for example, is the T_{sys} for the antenna large? Are there comments in the observing logs about pointing errors, local-oscillator drop-outs?

[10]CASA helper task of es.getReference does this

Flux Limit-Avg Separation for Phase Calibrators

BAND	Frequency (GHz)	T_{sys} (K)	PWV (mm)	Flux (mJy)	AvgSep (deg)
3	86.2	83	5.186	4.6	1.0
3	115.3	173	5.186	9.7	1.4
4	147.0	86	1.796	4.9	1.2
5	192.0	100	1.500	5.3	1.3
6	230.5	102	1.262	6.0	1.4
7	345.8	151	0.658	9.5	1.6
8	461.0	431	0.472	30.0	1.9
8	492.2	638	0.472	46.0	2.2
9	624.2	4302	0.472	364.4	8.0
9	658.0	1721	0.472	152.9	4.6
9	691.5	1231	0.472	114.9	4.3
10	806.7	2405	0.472	271.3	7.1

Table 10.2: The minimum phase calibrator flux density in mJy needed for an antenna-based amplitude accuracy of 7% and a phase accuracy of 4 deg rms (15 SNR) with 120 seconds integration time, forty-three 12 m antennas, an elevation of 50 degrees, and no coherence loss. The total bandwidth, ν_{BW} in both polarizations of 7.5 GHz is assumed.

(2) The antenna-based phase difference between adjacent calibration scans should be relatively small. For example, if the antenna-based phase change between scans exceeds about 60°, interpolation to the intervening target scan becomes uncertain. This often occurs at longer baselines and high frequencies, and flagging of the target data during these intervals may be useful to improve the image quality. This is the criterion used by the *Go-noGo* system at the telescope to decide if observations are feasible for a project.

10.5.5 Flux Density Calibration (T)

The gain calibration determines the relative amplitude between the phase calibrator to an accuracy of $<$ 6% per antenna solution per scan. This accuracy depends on the strength of the calibrator and other array parameters. However, the amplitude accuracy of the image will be significantly better since it is based on many antenna/scan solutions, so that a 1% accuracy of the target, relative to the assumed flux density of the phase cal, is typical.

In order to determine the *absolute* flux density scale of an observation, the observing schedule should contain either a solar system object and/or a bandpass calibrator with a recently measured flux density from the 'Grid' source monitoring. The present list of solar system objects used as primary flux density references are: Uranus (standard, but large in angular size); Neptune, Mars (Large), Ganymede, and Callisto. Ceres, Titan, Juno, Vesta and Pallas are also used, but may be too weak at Bands 3 and 4. The flux density model accuracy of the solar system objects is 3% at Band 3, rising to 15% at Band 9.

Once the T_{sys} corrections are applied to all scans, the known flux density of the grid source and/or solar system object can then be transferred to the phase calibrator to an accuracy of a few percent since the ALMA antenna gains do not vary by more than this level. If more accurate relative and/or absolute flux density precision is needed for a source observed over time or over frequency, contact your local arc for additional calibrations that may be necessary.

Because of the excellent stability of the antenna gains and the number and quality of the antennas, the a-priori calibration from Kelvin to Jansky is possible. Tests are not underway to determine the conversion accuracy and stability.

10.5.6 Polarization Calibration (p,X)

The XY and YX (crossed-hand) true visibility functions, in addition to the XX and YY, are needed to determine the linear and circular polarization images of a source. However, these two cross-hand terms are contaminated by three effects that must be calibrated: 1) The delay and phase difference between the X and Y total intensity calibrations to remove the global delay and delay difference after the total intensity calibration; 2) the non-orthogonality for each antenna-spw of the X and Y feeds, and 3) the slight rotation of the antenna-feed orientation from that assumed. The last two effects are parameterized by *D-terms*, two complex values as a function of frequency for each antenna-spw.

The present ALMA linear polarization calibration method requires the addition of observations of a bright, polarized source during a science experiment. The accurate linear polarization of the quasar need not be known, only that it is more than about 2.0% polarized in order to have its true polarization signal dominate over that of the polarization error terms. In order to separate the quasar true polarization XY and YX from that caused by the three contaminants listed above, the polarization calibrator must be observed over a period of three hours (depends on declination) for its parallactic angle to rotate at least 60°. This means that the polarization calibration of a science target requires at least a three hour observation, even for target observation lengths that may require only a few minutes. Other polarization methods that require a shorter period of time are currently under investigation.

The D-terms vary considerably over the primary beam so that at present only sources with a limited angular extent can be accurately imaged. In Cycle 5 this size limit will remain limited to the 70% power point of the antenna primary beam. Both 12-m and 7-m Arrays can be used for polarization observations, but are limited to Bands 3, 4, 5, 6 and 7. Investigations are in progress to open up the area of polarization imaging, with mosaicing in all Stokes parameters an ultimate goal.

Circular polarization information can be determined from the four Stokes parameters (I,Q,U,V) that are measured. However, the calibration accuracy of V remains uncertain, especially at a location more than a few percent of the primary beam size from the beam center, and thus circular polarization is not currently supported. Please contact the ALMA helpdesk of your local arc for more information about any circular polarization results from your data before their publication.

10.6 Additional Calibration Concerns

10.6.1 Observations at Frequencies > 400 GHz

At high frequencies the number density of detectable quasar calibrators drops significantly because of the decrease in ALMA sensitivity and the negative spectral index of quasars. Thus, it is difficult to find a calibrator that is closer than 3° to a target needed for accurate phase referencing. In Cycle 4 and continuing at least into Cycle 5, JAO and ALEGRO (EU ARC Node in Leiden, Netherlands) plan and carry out short observations with part of the ALMA array in order to determine which of the 15,000 sources in the ALMA catalog are sufficiently strong to be used as a nearby calibrator for each PI target. These observations are called *Cone Searches*. With an observation of a few minutes, the suitability of a candidate calibrator can be ascertained (see Section 8.12)

10.6.2 Observations with Baselines > 3 km

The calibration of long baseline observations ($>$ 3 km) is similar to those at shorter spacings, with the main differences being that the phase calibrator is visited more often, and should be as close to the science target as possible. Starting with Cycle 4, *cone searches*, that were described above for the high frequency targets, have been used to find the closest detectable calibrator for a long baseline observation.

Do not propose for the longer baselines unless there is sufficient information about the small-scale structure in the target. Otherwise, you may have to smooth the data to detect the source at a lower resolution—which essentially removes the longer baseline data from the image.

Continuum Observations: Does the source have sufficiently small-scale structure in order to benefit from long baseline observations? For example, at 350 GHz with a 5 km array with 40 12 m antennas, the resolution is $0.06''$ and the rms image noise in a 600-sec integration over a bandwidth of 7.6 GHz is 0.1 mJy. Hence, unless there is sufficient substructure in the target at this angular size, the long baseline will dilute the extended source emission.

Narrow line emission at high resolution: The imaging of line emission with the longest baselines is challenging. For example, at 89 GHz with 0.25 km/s channels, 40 antennas, 3.5 hours on-source, and an angular resolution of 70 mas, the expected spectral rms noise per channel is 2.0 mJy/beam but only 63K in surface brightness sensitivity. Hence, a brightness of emission of at least 300K is needed for a solid detection.

10.6.3 Phase Calibration with Narrow Bandwidths

When observing a target within narrow spw bandwidths (typically $<$ 250 MHz), a suitable phase calibrator, especially at the highest frequencies, may not be detectable with a one to two minute calibrator scan. Two methods of calibrating these narrow bandwidths are available:

One wide spw: Low-SNR calibration. If the aggregate bandwidth in the combined spw's in the experiment has a bandwidth $>$ 1 GHz (usually because one of the spws is wide band), then there will be sufficient sensitivity to determine the antenna-based gains using the low SNR calibration. This method, which uses the bandpass scan to determine the spw phase differences, is described in Section 5, Gain Calibration.

Band-Width-Switching. If all of the spw's are narrow bandwidth, it is likely that a suitably nearby phase calibrator cannot be detected within scan, even after phasing up all of the spw's. In this case, the strategy for phase calibration is to observe the phase calibrator with wide-band spectral windows and the target in the desired narrow bandwidths by switching the spectral setup between the relevant scans.

The antenna-based amplitude ratio and phase difference between the wide-band spw's and the narrow band spw's is measured by observing a strong calibrator, also called a Differential Gain Calibrator (DGC), with both spw setups. In order to remove short-term phase changes, the measurement of the DGC is made in three consecutive scans with: wide band, narrow band, wide band. The temporal variation between the two wide band scans is used to remove the temporal variations in the intervening narrow band observation.

The advantage of the second strategy is that all of the target spw's can be used for narrow-band spectroscopy, whereas the first strategy uses one wide band spw for calibration purposes. But, the first is simpler to schedule, somewhat more efficient, and easier to calibrate.

10.6.4 Band-to-Band Calibration

For phase referencing at frequencies above 400 GHz, finding a calibrator closer than 5° to the target is rare (even after using cone-searches), whereas a closer calibrator at lower frequencies is often available. The goal of band-to-band phase transfer is to use observations of a phase calibrator at a lower ALMA Band, scale these antenna-based phase solutions (G) to the higher frequency band, and then apply this calibration to the target source, observed only at the higher frequency.

This method is being tested and may be available for Cycle 6. The main experimental problem is that the phase difference between the low and high frequencies must be determined at one specific time. Since ALMA cannot observe at two frequencies simultaneously, the phase difference measurement is contaminated by the atmospheric phase fluctuations between the high and low frequency observation. The minimum switching time is now about 70 sec, but could be decreased to 20 sec in the next year, permitting a much more accurate measurement of the instrumental phase difference between frequencies.

10.6.5 Check Source

Some observations (those at high frequency or long baselines) contain short observations of a weak calibrator source, called a *check source*, in order to provide an estimate of the quality of the observational calibration based on the check source image purity and position). For Cycle 5, a check source is included in all observations with angular resolution $< 0.25''$ and for all observations at a frequency greater than 400 GHz. All potential check sources are in the ALMA catalog; a typical choice is weaker than the phase cal but at a similar distance from the target. It is typically observed every third phase-target cycle with less integration time than the phase calibrator scan. After normal calibration of the data, an image of the check source provides a good indication of the image quality (its dynamic range) and positional accuracy (offset from the phase center) since the check source position is usually known to 2 mas.

10.6.6 Astrometric Observations

ALMA Nominal Positional Accuracy

The absolute positional accuracy of an ALMA image, made with normal phase referencing depends, on the target signal to noise and the array resolution, but also on the quality of the phase calibration that depends on the atmosphere delay fluctuations and the proximity of the phase calibrator to the target.

First, the signal-to-noise of the source image and the resolution of the observations provide a minimum obtainable astrometric precision. The relationship is given by

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

where

$\Delta p =$ Theoretical accuracy limit by SNR only,

$BSL =$ Maximum baseline length in km,

$FREQ =$ The observing frequency in GHz,

$SNR =$ The signal to noise at the source peak.

For example, at Band 7 (350 GHz), for the 1-km BSL configuration (C40-5), and for $SNR = 40$, the theoretical astrometric accuracy is 4.2 mas.

However, for a reasonably bright and compact source, the astrometric accuracy is often limited by the size and nature of the phase variations with time and over the array, rather than that from the SNR limit and array resolution. An estimate to the best positional accuracy of an ALMA observation with normal calibrations are: 1) from the SNR estimate given above; 2) 5% of the resolution; 3) or a limit of 3 mas, whichever is the largest. This limit is frequency independent.

Optimum Choice of Frequency and Configuration

The choice of configuration and observing frequency is important to the success of an astrometric observation, and depends on the spectral and angular properties of the radio source. Since virtually all astrometric phase errors are delay-like and scale with frequency, simply changing an observation from a low frequency to a higher frequency, other things being equal, will *not* produce better astrometric accuracy even with the higher resolution. Of course, for thermal objects, the increase in flux density with increasing frequency does increase the SNR with a probable gain in precision. On the other hand for a non-thermal source a lower frequency may give better positional accuracy.

The larger the configuration, the more accurate will be the astrometric results. But, the positional accuracy tends to increase only with the square-root of the maximum baseline length of the array for two reasons: (1) Most of the phase/delay errors also increase somewhat with baseline length, so the gain in astrometric accuracy varies less than linear with the configuration size: (2) If the object is not a point source, then its peak brightness will decrease with the increased resolution, limiting the astrometric accuracy from SNR considerations and with the possible change of the maximum brightness position as a function of resolution. Hence, choose the frequency and configuration that best suits the celestial object—where it is strong but not too resolved.

More Sophisticated Astrometric Calibrations

When the SNR of the object is sufficiently large (> 20) and point-like, increased astrometric precision to < 1 mas can be obtained with ALMA with somewhat more complicated calibration strategies. Some suggestions are given below, but please contact your local arc for more information about setting up your astrometric observations.

Check source: Make sure a check source is included in the observation plan. Even though observed typically only a few times per observation, for a shorter time than the phase calibrator, the check source image quality and its positional offset are an indication of the image quality and position accuracy of the target. This is minimal but useful information for assessing the accuracy of the target position.

Second Phase Calibrator: If the check source is sufficiently strong and close to the target, then it can be used as a second phase calibrator by obtaining antenna-based phases in the same manner as that for the main phase calibrator. How to use two or more phase calibrators to provide better astrometric accuracy should be discussed with your local arc astrometric expert. This multi-phase calibrator observation requires User-Specified Calibration and is therefore a non-standard mode. One example is the use of two calibrators on approximately opposite sides of the target. The antenna-based phases derived from each of the calibrators can be averaged to determine an estimate of the phase calibration near that of the position of the target. Experiments of this type of calibration can produce astrometric precision to 0.5 mas with baselines 7 km or longer.

Phase Stability: Request better than average phase stability at the frequency of the experiment, although keep in mind that the observatory cannot guarantee this. While large phase fluctuations are a detriment to all ALMA observations, the effect on positional accuracy is more serious than that for image quality. For example, image distortions can be reduced using self-calibration techniques, whereas it does not improve the astrometric precision.

10.6.7 Self-Calibration

If the target has significant emission at the longer spacings, then self-calibration should be considered. This method simply uses the CASA gaincal task on the target data using the phase-referenced image (input image model from clean) in order to determine the residual phase errors between the data and model.

To determine if self-calibration is possible, the correlated flux density of the target at the longer spacings should be at least that needed to detect a phase calibrator, which is listed in Table 10.2 For example at Band 7, with 7.5 GHz bandwidth, the correlated flux density of the target must be > 125 mJy at the longer baselines in order for self-calibration to be successful. Such precise knowledge about the target may not be known, but the PI should have some indication of small-scale structure to indicate whether self-calibration is a possibility. When in doubt try the self-calibration. It will be obvious if the SNR of the solutions are poor, especially for antennas near the end of the array.

10.6.8 Total-Power Calibrations

For sources that are significantly more extended than the 7 m primary beam size, the 12-m TP Array (the PM antennas each used for scanning over a source) are used to make on-the-fly raster observations that cover the appropriate target area. The observing modes are discussed in Chapter 8 and the imaging in Chapter 7.

T_{sys} Calibration: The data are calibrated into a brightness temperature in units of Kelvin using the equation $Ta* = T_{sys}(ON-OFF)/OFF$, where T_{sys} is the system temperature and ON and OFF are, respectively, the measurement on-source (during the raster scanning) and off-source (observations towards a defined reference position which should contain only background emission).

Baseline Subtraction: After applying the T_{sys} calibration, the off-source regions within the on-the-fly area should be near zero emission for all spectral channels. However, atmospheric changes during the observations can produce a small change across the spectrum at these off-source positions. These off-source offsets can be used to correct the data from all raster positions. More details are given in Section 8.6.

Flux Calibration: In order to convert the single-dish data from units of Kelvin to units of Jy/beam, a conversion factor and the effective primary beam size, after the gridding used to determine the total power image, is applied. These corrections are stable over long periods of time and are measured by the ALMA staff as needed, and applied in the total power reduction scripts.

10.6.9 Acknowledgments

Many people contributed to the checking, revising and improving of Chapter 10. I thank: Yanett Contreras, Akihiko Hirota, Neil Phillips, Carmen Toribio, and to Catherine Vlahakis for her detailed and germane comments.

Chapter 11

Quality Assurance

The goal of ALMA Quality Assurance (QA) is to ensure that a reliable final data product is delivered to the PI, that is, the product has reached the desired control parameters outlined in the science goals (or is as close to them as possible), it is calibrated to the desired accuracy, and calibration and imaging artifacts are mitigated as much as possible. The QA analysis will be based on a calibration plan that specifies which observations must be acquired and at which intervals to monitor system performance and environmental factors as they evolve with time. Furthermore, it will also be used in assessing the merging of data within each science goal taken with different configurations, the inclusion of 7-m Array and TP Array data, and the final image quality. Errors introduced by user supplied parameters, such as incorrect source coordinates, inadequate frequency setting (e.g. an incorrect redshift) or inadequate sensitivity limits (leading to an inadequate integration time or inadequate *uv* plane coverage) are outside the scope of the ALMA QA, unless the error occurred due to faulty information or tools provided by the Observatory.

To be more efficient in detecting problems, ALMA QA has been divided into several stages that mimic the main steps of the data flow. The broad classification of this multi-layered QA approach is:

QA0: Monitoring of calibrations and overall performance during observations

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

QA2: Full Calibration and generation of science products

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

The QA0, QA1 and QA2 stages will be handled by the Program Management Group (PMG) and the Data Management Group (DMG) (with contributions from ARC personnel) using the ALMA Quality Assurance (AQUA) Tool (see Section 11.6). Responsibility for data quality assurance in Chile rests with the Data Manager and his Deputy, within the Department of Science Operations, drawing upon the resources of the Program Management Group and the Data Management Group. The final output of the ALMA QA0-QA2 process is a "QA Report" per ObsUnitSet (Member or Group[1]) that summarizes all the relevant QA information for each of the different QA stages up to, and including, the final imaging[2]. This report is included in the data package delivered to the PI. The QA3 stage will be handled separately, by the ARCs, via JIRA[3] tickets created by the ARC personnel (see below). A more detailed description of the different stages of QA is given below.

[1]Group ObsUnitSet processing, that is, the combination of datasets taken in different configurations/arrays, will only be done sporadically during Cycle 5. It is expected that group processing will be implemented as a regular data reduction strategy in later ALMA cycles.

[2]For datasets that have been processed manually or those that required additional imaging, no imaging report will be provided during Cycle 5

[3]JIRA is a proprietary issue tracking product, commonly used for bug tracking, issue tracking, and project management

11.1 Cycle 5 Quality Assurance Goals

The ultimate goal for ALMA QA is that the delivered products are considered "Science Ready", suitable for publishing with little need for the user to reprocess. However, ALMA cannot guarantee that all the scheduled Cycle 5 projects are completed (especially those with grades B and C), or that specific criteria like sensitivity or angular resolution are precisely met, give the restrictions of the array configurations, system performance at the time of the observations, etc. The observatory will therefore attempt to meet the sensitivity and resolution stated by PIs in their Science Goals, within certain tolerances as described later in this chapter. In principle, each project component ("Schedule Block" or SB) is scheduled for the number of executions that are expected to reach these goals (based on nominal Array performance values). Additional executions may be needed during the observing cycle if specific executions do not pass QA0 or QA2 (see below).

Observers may need to invest their own time and expertise to ensure that the data products are of the appropriate quality and to re-reduce the raw data if the quality is not satisfactory. This may include the need to visit the relevant ARC or ARC node to get help and to assist with quality assurance and potential data reduction.

11.2 QA0

QA0 is a near-real-time verification of data quality. It deals with rapidly varying performance parameters (on timescales of an SB execution length or shorter) and thus has to be performed at the time of data taking. Assessment is performed by AoDs (Astronomers on Duty) at the OSF, with support from the Science Archive Content Managers, using the AQUA software tool, based on semi-real time output of the calibrations (obtained by the real-time TelCal ALMA software) as displayed by QuickLook and the "Calibration Summary" files that are produced at the end of each execution of an SB. This information is complemented with reports derived using Monitor and Control display tools to monitor specific parameters not directly tracked by the calibrations (e.g., total power level variations, weather parameters, etc). AQUA can currently generate the complete Calibration Summary file contents using the information available in the ASDM files for each Execution Block.

QA0 metrics/parameters have been selected to check the health of the whole signal path from the atmosphere down to the back-ends. These parameters can be grouped into the following categories:

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

Antenna Issues: Antenna Gain, Relative/Offset Pointing, Focus, Antenna Tracking, Geometric Shadowing, Nutators (unavailable for Cycle 5).

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

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

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

The tolerances for these parameters that have been adopted by ALMA for this Cycle are listed in Section 11.7. Apart from these, the metadata is checked for inconsistencies and the amplitude calibrator for recent flux measurements. Each SB execution is classified into three categories, i.e. QA0_PASS, QA0_SEMIPASS and QA0_FAIL, based on the criteria described in that Section. QA0_PASS are datasets that comply with all the QA0 criteria and that will be used for the final imaging. QA0_SEMIPASS are datasets that do not fulfill all the QA0 criteria, but contain data that is deemed of scientific value[4]. QA0_SEMIPASS data are not included in the final data products, but PIs can access those data and reduce them if they wish. QA0_FAIL datasets are those that are not included in the other two categories. Since they represent unusable/uncalibratable data, they are not visible to the PIs in the Project Tracker.

[4]For instance, the presence of usable calibrator data.

For operational purposes, individual SB executions during Cycle 5 will be counted as a normalized fraction of the theoretical sensitivity that a given execution should have reached based on the number of antennas that should be in an array in Cycle 5 (i.e., forty-three 12 m antennas on the main array, ten 7 m antennas on the Morita Array, and three TP 12 m antennas), and the average system temperatures expected at the frequency of the observations. Executions with *higher* RMS noise levels than the reference will have fractional executions between 0 and 1, while those with lower RMS noise levels, which can be achieved, for instance, with a larger number of antennas in the array, or with lower system temperatures, will have a fractional execution value above 1. The number of executions may be changed (i.e., increased/decreased) based on the total execution fraction of a given member OUS at a given point in time.

11.3 QA1

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

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 the same band

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

11.4 QA2

QA2 deals with QA at the level of data reduction and imaging using the Science Pipeline or performed interactively by the ALMA Data Reducers Team. It is only at the stage of data reduction that the science goals set by the PI can be compared with the actual value in the data products (i.e., RMS, angular resolution, SNR, dynamic range, etc). The official software for ALMA data reduction is CASA.

During Cycle 5, it is expected that, for the ALMA standard observing modes, the automated Pipeline (see Chapter 13) will be used for the calibration of the data and imaging. The Pipeline provides a Weblog page with detailed information on the quality of the calibration, which will provide the basis for Pipeline QA2. This information will also be made available to the AQUA tool during Cycle 5, and the weblogs will also be included in the data packages sent to the PIs. For manually-reduced datasets, the QA2 will be carried out with special purpose scripts whose outputs will be included in the data packages delivered to PIs. The QA2 metrics which determine the success of an observation are given in Section 11.7.

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

Calibration Issues:

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

Final Data Characterization:

- Longest baseline, visibility coverage and time on target (after flagging).
- Synthesised beam (spatial resolution) for specified weighting scheme; specific imaging requirements (if stated by the PI). See Section 11.7 for tolerances.
- Spectral resolution and channel ranges used to make sample images.
- RMS noise in target images: Values are compared with those predicted from data after flagging and with those requested. See Section 11.7 for tolerances.
- Residual imaging artifacts: Sidelobe Levels, effects of "missing spacings" and possible dynamic range limitations.
- Combinations of array configurations and/or total power, e.g. amplitude scale consistency.
- Mosaicing and/or contamination by bright sources outside FOV or aliasing in clean.
- Polarization purity: minimum fractional polarization detectable, and polarization angle accuracy (if relevant).

The data reduction scripts sent to the PI will include any modifications to the standard scripts which were found to be essential during QA2.

There are three possible QA2 states a reduced dataset can be placed into. QA2_PASS implies that the scientific goals, as defined by measurable parameters such as noise RMS, LAS and angular resolution, have been achieved within the specifications listed in section 11.7. QA2_SEMIPASS refers to those datasets that fall short of meeting the PI requested science goals, but are otherwise of good quality (or as good as possible with the observations that could be achieved). QA2_FAIL is a temporary state used during an observing cycle when an observation fails to meet the PIs goals by a significant margin (see Section 11.7 of this handbook) and needs to be scheduled for additional observations. If these data fail to obtain enough additional observations to pass QA2 by the end of an observing season and they are not grade A proposals, their QA2 state is changed to QA2_SEMIPASS and they are delivered to the PI.

11.5 QA3

QA3 is post-reduction evaluation of the data products delivered to the PIs. It is advisable that PIs check the data products themselves and report any problems that they find to their Contact Scientist via the ALMA Helpdesk (https://help.almascience.org/). The QA3 process will be triggered by the PI, the contact scientist of other ARC personnel. They will open an ALMA Helpdesk ticket reporting a problem with the data products which may reflect an underlying problem with the data, observing procedure or calibration. The ARC receiving the Helpdesk ticket will retrieve the data from the archive and evaluate the nature of the problem. The evaluation by the ARC should include an assessment on whether the problem is present only in a particular dataset or whether others taken under similar set-ups and conditions also show it. If the problem is deemed to reflect a problem with the performance of the array, the calibration or data reduction processes, or the QA process, the ARC will communicate their findings to the observatory, which will work on solving the problem in collaboration with the ARCs. The result will be communicated back to the reporting investigator. An extension of the proprietary period of delivered datasets with QA3 issues will be granted based on the policies in the Cycle 5 User Policies document.

11.6 The Quality Assurance Report

During Cycle 5 the Quality Assurance Reports of QA0, QA1 and basic QA2 will be accessible to the PIs using SnooPI. These reports will be generated using the AQUA tool. For Pipeline calibrated datasets, additional reports (Weblogs) will be sent to the PI as part of the data delivery. For non-standard observing modes that require manual data reduction, script-generated QA2 reports will be delivered instead.

The basic unit of a Report is the ObsUnitSet, which represents a scientific goal stated by the PI during Phase 1 (project creation and review). An ObsUnitSet will typically contain several executions of SBs. For each execution QA0 and QA1 reports are generated by the AoDs using the information available at the time of the observations, which includes TelCal outputs and other monitoring data (weather, total power levels, Corr GUI outputs, etc). A given execution is only cleared for reduction if it has passed both QA0 and QA1. There will be only one QA2 report for the whole ObsUnitSet[5] generated by the DMG/ARCs at the end of the data reduction process; this also has to be approved before the data products are delivered to the PI. It is expected that sometime during Cycle 5, a complete interface between the Pipeline and AQUA is made available. From that point onwards, complete QA Reports will be generated using the AQUA software and delivered as such (including the Weblog information). The final report per ObsUnitSet delivered to the PI would in that case be a concatenation of all the relevant QA0 and QA1 reports per execution with the QA2 report. Comments on each stage of the QA process (with supporting images, if required) would be added to the Report.

The standard policies for QA0 and QA1 failures are that the observations of those Execution Blocks that failed have to be repeated. Failures to pass QA2 may trigger additional observations if the achieved RMS or imaging parameters do not fullfil the QA2 pass criteria. Additional observations may not be possible in circumstances, such as projects with very tight weather, very specific configurations, or time constraints. If the available data are insufficient to reach the required sensitivity, but are otherwise of good quality, they will be released to the PI at the end of the cycle.

11.7 QA Criteria

For an execution to be considered QA0_PASS it must have a phase RMS (post-calibration) of $\leq$0.5 rad, enough calibrator data to be able to calculate necessary calibration terms (bandpass, amplitude, complex gain, and, if relevant, polarization), and enough science data to be useful (see below). Furthermore, the execution has to be "significant" in terms of the fractional execution (i.e., more than 0.2) OR in terms of sensitivity (i.e., more than 50% of expected Cycle 5 sensitivity per execution of a given SB). If it does not meet these criteria, but has some useful data (calibrator or source), it is classified as QA0 SemiPass. Such data are available through the ALMA archive, but they will not be used in the generation of the imaging data products. If neither of these hold, the execution is deemed to have no useful data and declared QA0_FAIL.

The QA0 pass/semipass/fail criteria that have been adopted by ALMA during Cycle 5 5 are based on the following:

- **Antennas**: QA0_SEMIPASS for fewer than 20 antennas available in the 12-m Array or fewer than 5 antennas in the 7-m Array or fewer than two Total Power Array antennas. Situations that will be considered to render an antenna unavailable include issues with the antenna itself (including large pointing scatter of $>1/10$ HPBW, and bad band focusing with offsets from model of $>1/5\lambda$ in Z and $>\lambda$ in X/Y), as well as issues with the receiver bands not being available at the antenna for a given observation.
- **Bandpass**: QA0_SEMIPASS if the bandpass calibrator signal is too weak, with amplitude wiggles across the spectral window of >3 dB on the autocorrelations (not due to atmospheric features) or strong CW signals, and the phase/gain calibrator cannot be used instead.
- **System Temperatures**: QA0_SEMIPASS if more than 20 antennas in the 12-m Array, or more than 3 antennas in the 7-m Array, or more than one Total Power antenna have $T_{\rm sys}$ values $>50\%$ of the others,

[5]As stated above, processing of group ObsUnitSets will not be regularly done during Cycle 5. PIs should assume that all statements here apply to member ObsUnitSets.

have differences between polarizations > 50%, or have $T_{\rm sys} > 2000$ K.

- **Gain**: QA0_SEMIPASS Phase RMS > 0.5 rad, on scans of the gain calibrator, for a significant fraction of the baselines after WVR phase correction (using all possible bandwidth in a given spw)
- **Execution**: QA0_SEMIPASS All datasets whose execution failed at some point and contain less than 20% of the expected execution time of a given SB execution.
- **Calibrations**: QA0_FAIL Datasets missing critical calibrations that render them uncalibratable. Execution Blocks only containing usable calibrator data shall be labelled QA0_SEMIPASS.
- **Storage**: QA0_FAIL Data that could not be read from the Archive
- **Time-Critical Observations**: QA0_SEMIPASS if observations have not been carried out at the requested dates and times.

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

For QA2 at the GOUS (Group OUS) level (or at the Member OUS, MOUS, level if no data combination is required), the main criteria are the achievement of the requested noise RMS in the images (it must be within 10%, 15% and 20% of the goal for Bands 3/4/5/6, 7/8, and 9/10, in flux units per beam, respectively), the synthesized beam shape[6] (the angular lengths of the major and minor axes of the synthesised beam must be within ±20% of the angular resolution requested by the PI) and the calibration quality (phase RMS of 5 degrees after calibration, and absolute flux scales within the ALMA specifications[7]). For the individual MOUS that must be combined at the GOUS level, the RMS values quoted above should be scaled by the relative sensitivities of the configurations, using the most compact 12-m Array configuration as reference.

[6] During Cycle 5, SBs shall be executed following the recommendations given by the OT, which will select the configuration that produces the closest values to the requested angular resolution.

[7] Absolute Calibration accuracy is regularly checked as an observatory task.

Chapter 12

Data Flow and Logical Data Structure

This chapter describes the data flow process from the observations until data is ingested into the ALMA Archive. It includes a brief description of all the main software subsystems involved in the data acquisition and archiving, as well as a summary of the metadata structure adopted by ALMA.

12.1 Data and Control Flow

This section describes the overall control of the ALMA system and the flow of data. A summary of the main actors and operations involved in the observations is shown in Figure 12.1.

Each of the gray boxes in Figure 12.1 represent an ALMA subsystem involved in the observations. The rest of the boxes, colored according to the actor involved, include labels for the actions performed either by external agents (actors) or by those subsystems.

A typical observing session would be started by the Telescope Operator interacting with the Executive subsystem via one of the dedicated control computers in the Control Room (the so-called "Standard Test Environment" or STE). The Executive subsystem is in charge of starting up the ALMA Common Software (ACS) and its Common Object Request Broker Architecture (CORBA)-based services and then initializing all of the various software subsystems involved in the observing and data storage process. Once all the components are ready, the Executive also handles asynchronous events from several of the subsystems and responds to them accordingly. Among the events, the Executive also publishes a list of error conditions to the attention of the operator and the requests for status of the Control, Telescope Calibration, Correlator and Scheduling subsystems.

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

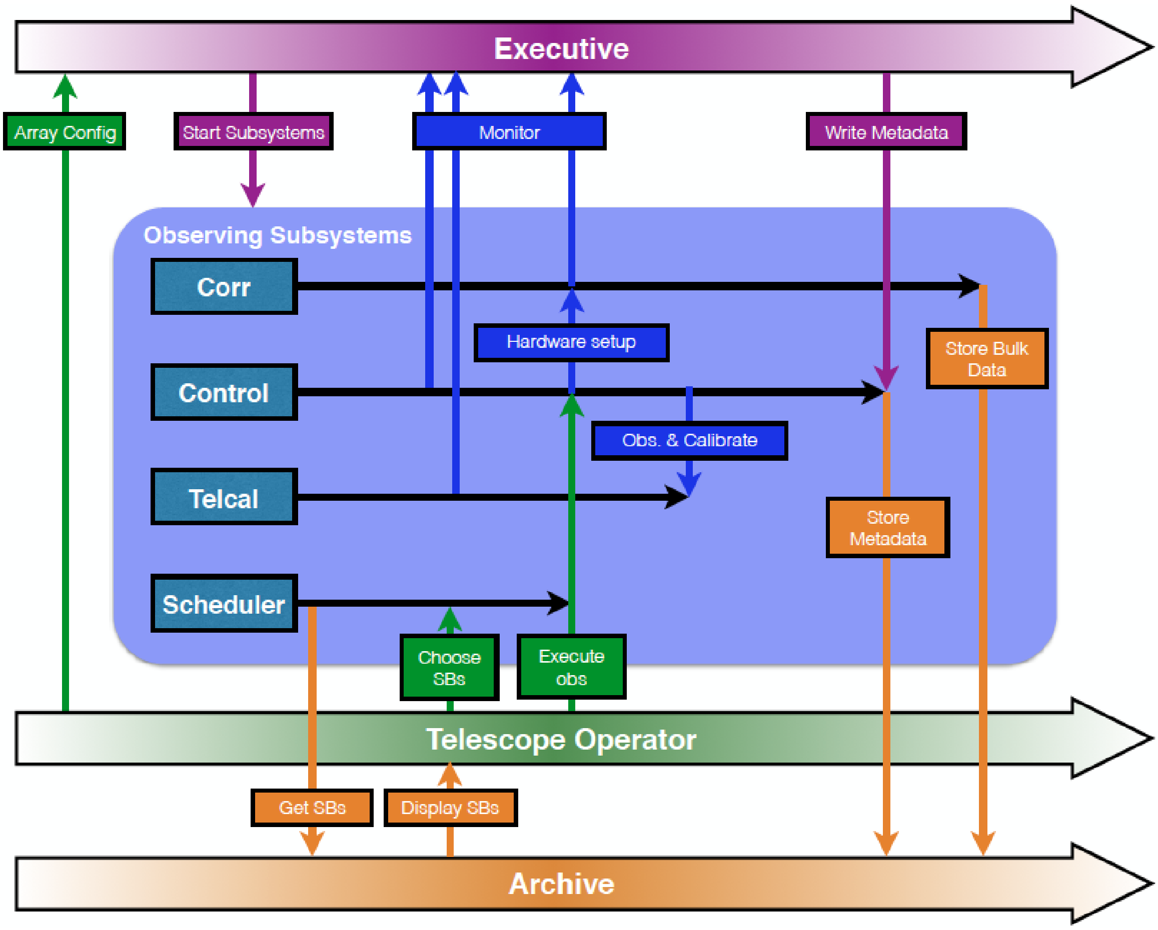

Figure 12.1: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

Once the execution of the SB has been selected from the Scheduler, the observation is set up. The SB contains which observing script to use. The observing scripts, which reside in the Science Observing Scripts module, uses the information from the SB as input to determine the necessary execution sequence. After the SB has been selected, the SSR scripts (1) use the information in the SB, including the observing script, to set up the observation and calibrations. The SSR script then sends this information to the Control subsystem as a sequence of scans and subscans. The scan and subscan specifications specify for each subscan the tuning, phase center and antenna pointing, calibration device position, intent metadata (to convey the purpose of the scan/subscan), and other parameters that may need controlling from the observing script level. Control executes the scan/subscan sequences by commanding all relevant hardware, the relevant correlator subsystem (Baseline (BL) or Atacama Compact Array (ACA)) and the Total Power Processor (TPP), resulting in raw data and metadata being sent to the archive subsystem and made available to the online Telescope Calibration (TelCal) subsystem. TelCal publishes results from calibration scans it reduces, which are sent to the archive as calibration tables, and received by the QuickLook GUI and other software to display results to the Operators and AoD to evaluate the observation progress. To carry out its function, Control has many interfaces to the instrumental hardware. It is in fact one of the truly real-time subsystems within ALMA because it is in charge of synchronization of the actions of all antennas (scanning, source acquisition, etc) and the correlator to within 48 ms (Timing Event (TE)). Control is also in charge of storage of data from all monitor points set in the hardware of the ALMA array.

Each run of an SB produces an Execution Block (EB) that is stored into the archive through two parallel paths[1] and is handled by a part of the Control software called the "Data Capture" module (see, for instance, the 2006 ADASS contribution by Hafok, Caillat and Mcmullin). The best way to describe the Data Capture module is as an interface between the real-time domain of the data taking and the storage side. As an interface, it captures and stores into the Archive all the relevant meta-data information pertaining to a complete description of the data and their supporting calibration. In addition, it also monitors datasets and condenses all that information into a set of XML tables. The contents of all these tables are defined in the ALMA Science Data Model (ASDM). Together with these tables, Data Capture also creates the relevant links of these metadata to the actual bulk data that the correlator stores directly into the archive. Furthermore, it also provides calibration data to the TelCal subsystem for calibration reduction and the Quicklook GUI for displaying in semi-real time. Finally, when the SB is finished, Data Capture is in charge of checking that all products representing the raw data have been produced and stored in the archive, and announces the completion of the SB to the Scheduler subsystem. It is clear from the list of roles above that Data Capture is a very complex module: it has to interface with the Correlator, coordinate the backend data with the metadata, and receive, store, and link together supporting observation information (source information, spectral set-up, etc) and monitoring data for data reduction (weather, pointing, etc). All this information originates in different hardware and software elements, each of which can be sampled at different rates and with limited view of the behaviour and state of the observing system.

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

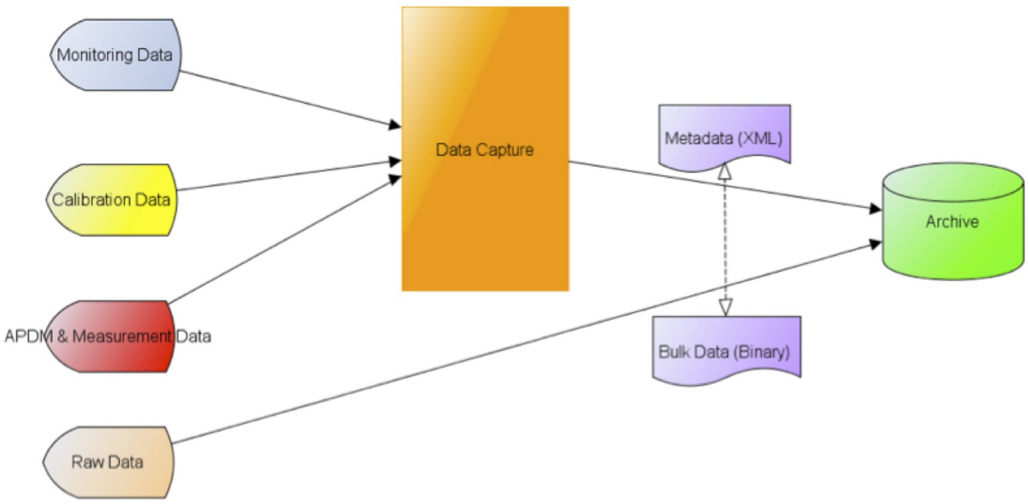

Figure 12.2: Data flow components.

[1] There are two types of data that are stored to the archive: meta-data (which is XML/text data stored by Data Capture) and bulk data (which is binary data stored by the correlators). The meta-data stored by Data Capture contains the relevant references to the bulk data, so that when a user retrieves an ASDM from the archive, the meta-data and corresponding bulk data are retrieved at the same time. For more info on this, please read Chapter 14 of this Handbook.

12.2 The ALMA Science Data Model

The ALMA Science Data Model[2] (ASDM; Viallefond, F. 2006, Astronomical Data Analysis Software and Systems XV, 351, 627) defines the collection of information recorded during an observation that is needed for scientific analysis. As described above it contains both bulk binary data and metadata (XML) organized in tables. The tables contain links to other XML tables and addresses pointing to the actual bulk data in the Archive. The ASDM contains 16 core tables that are common to all observing modes, and up to 23 additional tables that are only created for specific observations. On top of these, TelCal also creates associated tables whenever it processes any of the calibrations that can be done online. All tables are organized with a similar structure, with the columns listing the contents and the rows including the actual values. The core tables have been defined to outline some of the following: hardware characteristics, array configuration, antenna tracking, targets, auxiliary monitoring data, overall project and post-processing. A list of the core tables is shown in Figure 12.3. The term *referenced* in the tables means that a given table is linked to other tables via some references, that is, some of the information is shared between tables.

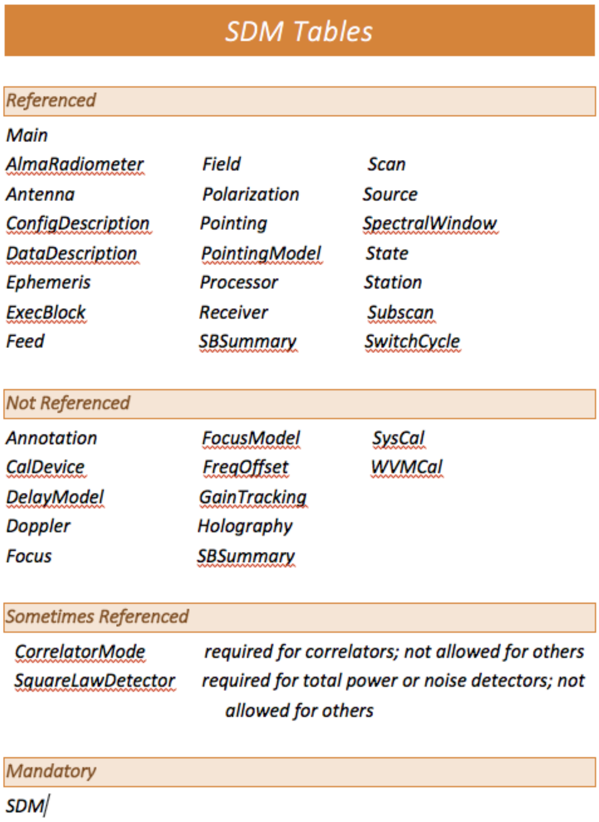

Figure 12.3: ASDM Tables. Outlined set of tables are the core ones (i.e., present in all ASDMs).

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

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

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

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

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

Chapter 13

Pipeline Processing

13.1 Overview

The ALMA Science Pipeline is used for the automated calibration and imaging of ALMA interferometric and single-dish data. It consists of modular calibration and imaging tasks within the Common Astronomy Software Applications (CASA) data reduction package that are selected and put together in a specific order based on standard prescriptions or "recipes". The standard ALMA recipes cover the processing of the standard ALMA observing modes, as defined in Section 5 of the Proposer's Guide for each ALMA proposal cycle. These recipes are submitted for execution on the computing systems at the Joint ALMA Observatory (JAO) and ALMA Regional Centers (ARCs) to produce standard products (calibration and flagging tables and FITS images), informative logs and scripts, and a "weblog" consisting of a collection of webpages with diagnostic messages, tables, figures, and "Quality Assurance" (QA) scores. These products are reviewed as part of the ALMA Quality Assurance process (see Chapter 11 of this handbook), and, if satisfactory, are stored into the ALMA Science Archive.

The Pipeline is data-driven and includes specialized algorithms or "heuristics" that automatically set task parameters or trigger different procedures based on the observing mode or quality of the data. Examples of heuristics are whether to flag certain data, or whether to smooth data before producing a calibration table.

This chapter describes the basic workflow for ALMA Pipeline Processing, and its use in ALMA Science Operations. Detailed information on the individual Pipeline tasks, heuristics, and QA scores are given in the Pipeline User's Guide and Pipeline Reference Manual, which are updated at the start of each public Pipeline software release (including the start of each ALMA observing cycle), and are available from the ALMA Science Portal.

The pipeline recipes are deterministic and should always give the same result for the same data. The documentation referenced above describes how to use the archived pipeline products to reproduce the standard ALMA calibration and/or imaging products. However, the CASA pipeline tasks are designed to be highly flexible, so that they can have the default inputs over-ridden with user-specified values, or be added, subtracted, or rearranged to produce alternative processing recipes. This enables a manual "mix and match" mode for data reduction and imaging that combines standard CASA pipeline tasks with other CASA commands or python code to produce scripts that are better tuned to the idiosyncrasies of a specific dataset. Some common "manual mode" modifications are presented in the Pipeline User's Guide.

13.2 Pipeline In ALMA Operations

13.2.1 Standard and Non-standard Observing Modes

One of the criteria for whether an ALMA Observing Mode is considered standard or non-standard is whether the ALMA Pipeline has been commissioned for that mode. The standard and non-standard modes for a given cycle are listed in the Call for Proposals for that cycle. Data associated with standard observing modes will be sent to the ALMA Pipeline for processing. Data from non-standard modes may be considered for pipeline processing (e.g. high frequency data with sufficiently bright calibrators), but as a rule will be sent instead for manual data reduction.

Since a project may have science goals that include both standard and non-standard observing modes, their resulting products may be a mix of both pipeline processed and manually processed data.

13.2.2 Pipeline Triggering

As described in Chapter 12 of this document, the structure of an ALMA project is broken down into a hierarchy of Observing Unit Sets (OUSs). There are OUSs at the Science Goal, Group, and Member levels. An OUS will include one or more datasets, each resulting from a successful execution of a Scheduling Block. These individual datasets, called Execution Blocks, are in ALMA Science Data Model (ASDM) format, and hence are often referred to as ASDMs. Each ASDM is assigned a Quality Assurance, level 0 (QA0) score (see Chapter 11 for more details). The QA0 scores are "QA0 Fail" (no useful data), "QA0 SemiPass" (some useful data but missing critical calibration or science data), or "QA0 Pass". Only ASDMs with "QA0 Pass" score are used for subsequent data processing.

The OUS level determines when the Pipeline is triggered and which products are produced. Member OUSs become eligible for data reduction after they transition into the "Fully Observed" state. Group OUS processing is triggered when all Member OUSs are completely processed and pass level 2 of Quality Assurance (QA2 - see Chapter 11). At the current time, the Pipeline is only commissioned for Member OUS processing; Group OUS combination is left to individual researchers.

The pipeline triggering conditions for projects with different Science Goal OUS structures are depicted in Figures 13.1 – 13.3. Figure 13.1 shows the conditions for a Science Goal consisting of a single Member OUSs, resulting in one pipeline processing request and one set of archived science products. Figure 13.2 shows the conditions for a Science Goal consisting of two Member OUS each under their own Group OUS, typical of multi-epoch observations. In this case, each epoch is processed independently by the pipeline, and separate sets of science products are archived. The final figure, Figure 13.3, depicts a Science Goal that includes two Member OUSs under a single Group OUS. This structure is typical of multi-array observations, where observations from different array components (distinct 12-m Array configurations, the 7-m Array, and/or the TP Array) each are grouped into separate Member OUSs, but gathered under a single Group OUS. In this case, each Member OUS is independently processed through the pipeline with its products separately archived and delivered. When all constituent Member OUSs have been successfully processed, the Group OUS becomes complete and ready for joint processing at the group level. When this level of processing is commissioned within the Pipeline, the calibration for each Member OUS will be restored and the data processed to produce an additional combined set of imaging products. This is expected to be available in ALMA Cycle 6 or shortly thereafter.

13.2.3 Pipeline Runs

A pipeline run consists of executing a pre-established sequence of commands known as a recipe. There are separate recipes for the following pipeline procedures:

- Calibration of standard-mode interferometric data.
- Imaging of interferometric data.
- Calibration+imaging of standard mode interferometric data.
- Calibration+imaging of standard-mode single-dish data.
- Calibration+imaging of "simple" ALMA Calibrator Survey data.

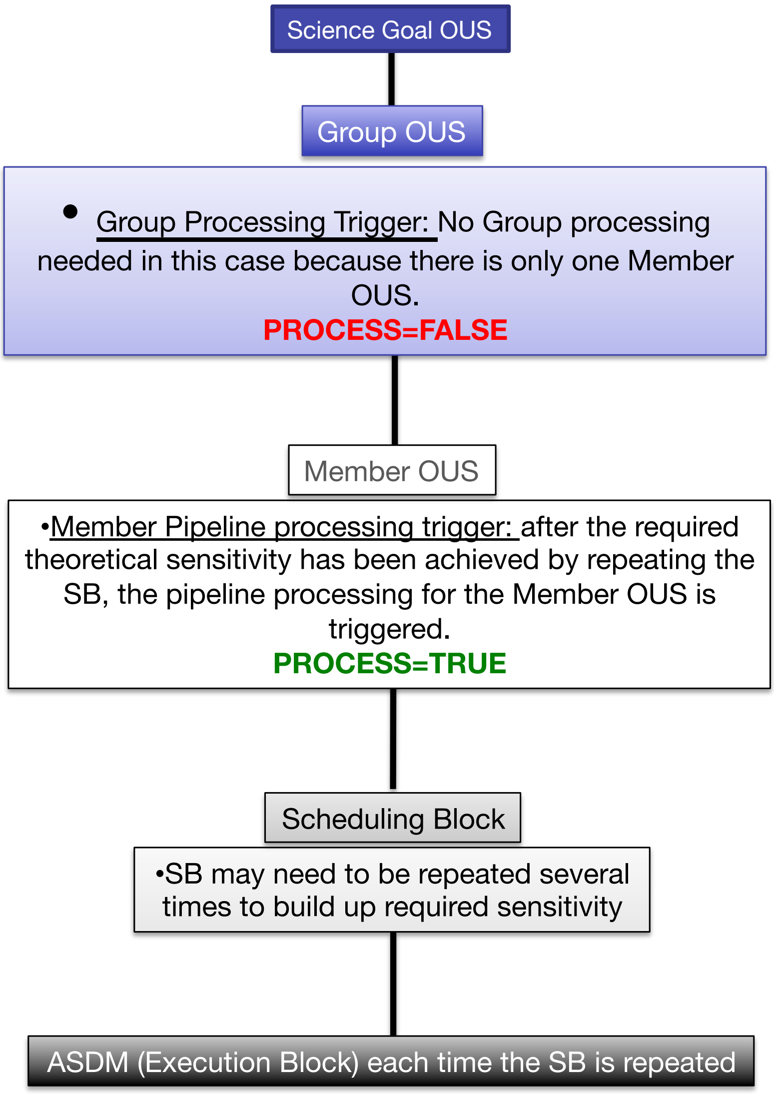

Figure 13.1: Pipeline processing of a Science Goal for single epoch, single array data. The same SB may need repeating N times to build up the requested sensitivity. This Science Goal would result in 1 Pipeline processing, with 1 set of science products delivered to PI.

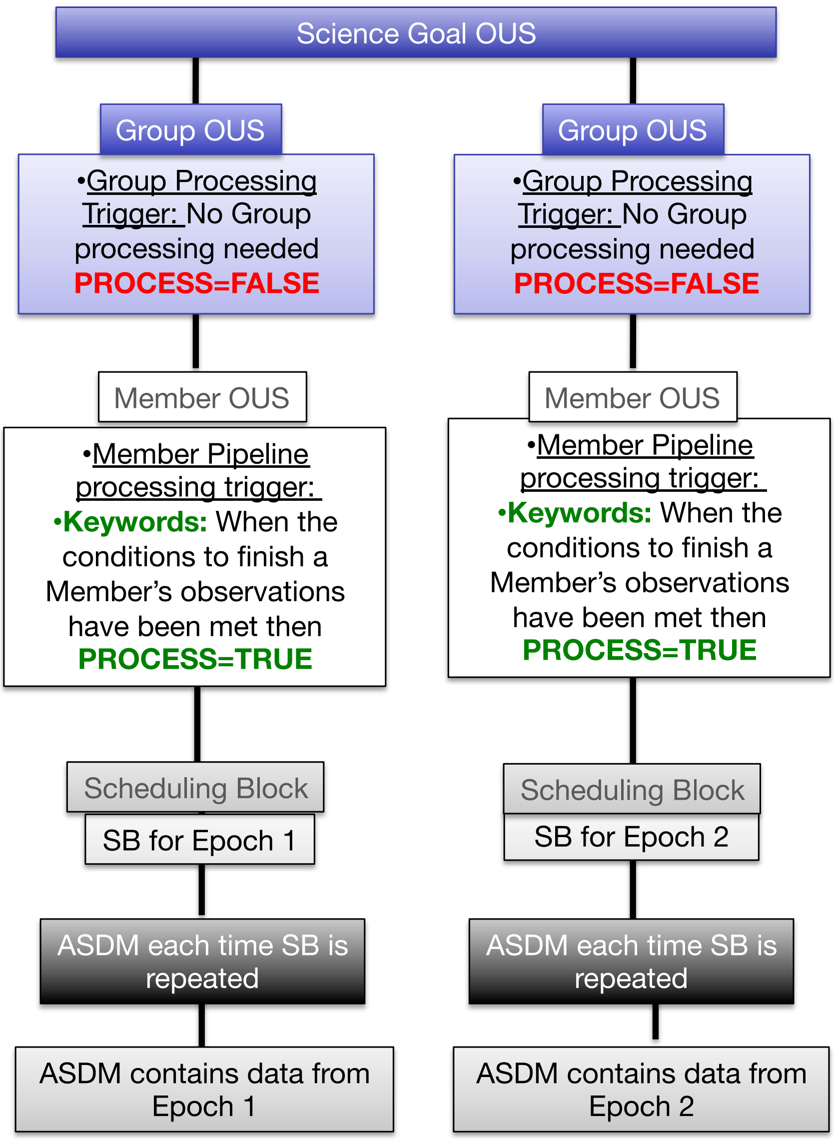

Figure 13.2: Pipeline processing of a Science Goal containing multi-epoch observations. Each epoch is processed independently by the Pipeline, with science products created for each epoch. Data from different epochs are not processed together to form a combined product (the Member OUS from each epoch are in different Group OUS). This Science Goal would result in two Pipeline processing runs, with two sets of science products archived and delivered.

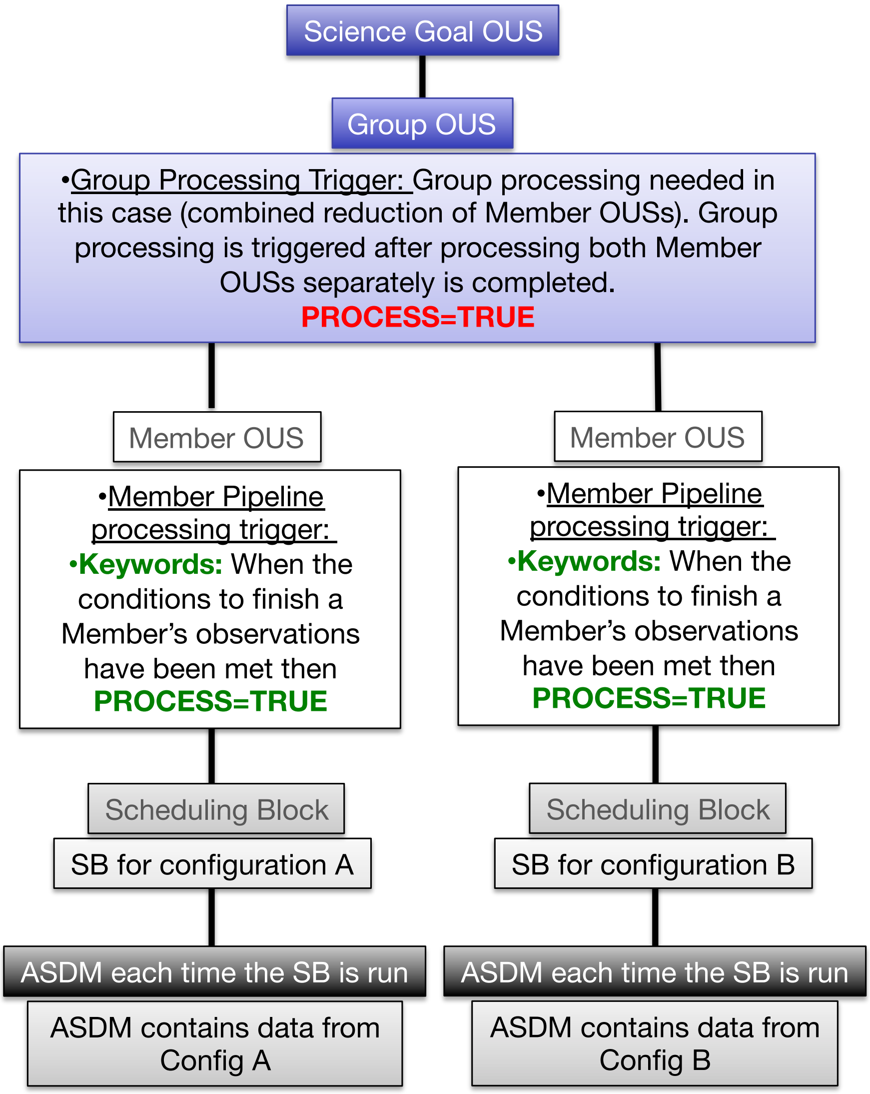

Figure 13.3: Pipeline processing of a Science Goal containing multiple ALMA 12-m Array configurations or containing data from different arrays (e.g. 12-m Array, 7-m Array or TP Array). The data from the different configurations/arrays are grouped into separate Member OUSs, each of which is processed independently, resulting in separate science products for each MOUS being archived and delivered. Once the processing of the MOUSs is complete, they are eligible for Group OUS processing, although this is not yet commissioned in the pipeline.

In the future, these recipes will be expanded as non-standard observing modes (e.g. polarization; solar observing) are commissioned in the pipeline and become standard.

At the end of each pipeline run, the results are reviewed and the success or failure of the pipeline results assessed. If the pipeline run was successful, the data move on to the next stage of processing using the appropriate recipe (e.g., from interferometric calibration onto interferometric imaging). If the pipeline run was not successful, it may be repeated after some manually introduced mitigations are applied (e.g., editing a template flagging file that the pipeline will automatically apply when the pipeline run is re-executed), or sent to manual processing (calibration and/or imaging) for treatment that is not available within the pipeline (e.g. self-calibration; using the phase calibrator to calculate a failed bandpass solution). As a result, the final delivered products may too be a mix of pipeline and manual efforts (e.g. pipeline calibrated but manually imaged or *vice versa*). Since single-dish data only has a single calibration+imaging recipe, single-dish products will never be a mix of pipeline and manual reductions.

13.3 Details on Pipeline Processing

As mentioned above, detailed information on the individual Pipeline tasks, heuristics, and QA scores are given in the Pipeline User's Guide and Pipeline Reference Manual, which are updated at the start of each public Pipeline software release (including the start of each ALMA observing cycle), and are available from the ALMA Science Portal. This subsection briefly describes a few aspects of pipeline operations that proposers may wish to be aware of.

13.3.1 Pipeline Runtime Database Queries

The ALMA Pipeline queries various ALMA databases for measurements that may have been updated since the data were initially obtained (such as calibrator fluxes from the ALMA Source Catalog; antenna positions from the Telescope Monitor and Control Database), or that must be provided outside the pipeline (Kelvin-to-Jansky conversion factors for single-dish observations).

At the current time, not all of these databases are available or fully enabled in the pipeline, and the updates may instead be done with observatory scripts that are not part of the official pipeline tasks, but are integrated with the pipeline recipes. The product of these scripts are text files that contain the updated information which are applied to the data during pipeline execution. These text files are included with the archived pipeline products.

13.3.2 Observatory "Fixes"

In previous cycles, issues with the data, metadata, or observing scripts were discovered after the start of the observing cycle that required special workarounds to modify the data and/or to include special calibrations not included in the current pipeline tasks. The modular design of the pipeline tasks allows these fixes to be incorporated into the pipeline processing by modifying the pipeline execution script to insert additional processing commands between the standard pipeline tasks, most commonly after the initial import of the raw ASDMs (similar to how the unavailable database queries described in the previous subsection are enabled). These "fixes" will be included in the processing files archived and delivered with the pipeline products so that they are properly applied when the data is reprocessed by user's.

13.3.3 Flagging Heuristics

The pipeline has a number of flagging heuristics, which are described in the pipeline documentation. However, these heuristics are only run on calibrators. The heuristics will extend the scope of the flags to flag all data or calibration tables associated with bad antennas or baselines, which will propagate to science targets. However, this will not catch bad data that is only associated with science targets. As such, investigators should carefully review the calibrated data for their science targets for any such data.

13.3.4 Treatment of Multiple Execution Observations

During data calibration, each ASDM is calibrated independently. In the calibration processing, the Pipeline runs sequentially through each pipeline task, applying that task to each ASDM. For example, the Pipeline first imports the data from each ASDM; it then performs deterministic flagging on each ASDM; it then moves onto the next task etc. This proceeds until all the flagging and calibration steps have been performed for each ASDM. As a consequence, the calibration and flagging tables are stored separately for each ASDM.

In the weblog, ASDMs that were obtained within the same observing "session" (continuous science observing period without a software restart) are grouped together. In the future, some calibrations may be "sessionized", meaning that calibrations obtained in one ASDM can be applied to other ASDMs in the same session, but this is not yet developed within the pipeline.

Imaging (of calibrators and science targets) is not done on a per-ASDM basis, but by including all ASDMs within the OUS.

13.3.5 Low SNR Heuristics in Interferometric Pipeline Processing

The interferometric pipeline has the following heuristics that are triggered if the signal-to-nose ratio (SNR) on a calibrator is too low.

- Using the Gain Solutions Obtained from the Aggregate of all Spectral Windows.

  If the SNR on the phase calibrator is too low, then the aggregate gain solution obtained by using all spectral windows is used to determine the time variable gain solutions. The bandpass calibrator observations are used to determine the phase offsets between the gain solutions for individual spectral windows and the aggregate solution.

- Additional Bandpass Pre-averaging.

  If the SNR on the bandpass calibrator is too low when calculating the prior gain calibration on a per-integration basis, then the time solution interval for this calculation is increased. If the SNR is too low when calculating the bandpass solution at high spectral resolution, then the frequency solution interval for this calculation is increased.

13.3.6 Total Power Pipeline Processing

There is a maximum of one single dish (Total Power) Science Scheduling Block per Member OUS. The Science SB will have one or multiple science targets, each with a corresponding rectangular observing field. It is advisable to make the mapped area large enough that there are emission-free regions around the edges of each field.

The single-dish pipeline recipe performs $T_{\rm sys}$ and blank-sky calibration, Kelvin-to-Jansky conversion (per antenna, per spectral window), spectral baseline subtraction, subsequent flagging, and imaging to process data for science targets. The Kelvin-to-Jansky conversion is not calculated by the single dish pipeline but instead extracted from an observatory database. This database is compiled from observations of standard amplitude calibrators (planets and bright quasars) using the 12 m antennas in the TP Array to derive the trends of the Gain (Jy/K) conversion factor for each of the antennas as a function of frequency, ambient temperature, elevation, and day-night dependence. As a consequence, the amplitude calibration observations are not included in the data products. Instead, a list of the applied Kelvin-to-Jansky factors is added to the data products.

During the imaging stages, the single-dish pipeline subtracts a low-order spectral baseline from each spectrum, using an automated line-identification algorithm to determine the channel ranges to be masked. Readers should refer to the Pipeline User's Guide for a detailed description of the line identification algorithm. This algorithm has been optimized to detect emission lines with a moderate width (wider than 100 channels) away from the edges of the spectral window. As a consequence, users should set up their spectral windows so that they are centered on the expected line frequencies and have sufficient bandwidth to provide an adequate number of line-free channels on either end. If the baselines are not flat due to inadequate line-identification and subsequent spectral baseline subtraction, the project will be sent to manual reduction.

13.3.7 Imaging – interferometric

The interferometric imaging pipeline is still relatively new. As such, the default pipeline products are expected to be reliably calibrated, but in many cases the images could be improved, in some cases significantly. For example, the continuum identification algorithm uses statistical measurements of the dirty cubes to identify frequency ranges that should be free from line emission. Its internal heuristics have been tuned to give reliable results over a wide parameter space; however they will not be optimized for every spectral window or science target, and investigators should anticipate that a more careful per-target, per-spw evaluation could improve their results. Additionally, the Cycle 5 pipeline will for the first time use an auto-masking algorithm in the imaging deconvolution (clean) step that is similarly tuned to work under a wide range of parameter space; more carefully defined masks and interactively defined cleaning thresholds may produce better results. Finally, the pipeline does not include science target self-calibration. Therefore, the pipeline imaging products of bright sources will be dynamic range limited.

The imaging pipeline is computationally demanding, and long imaging runs could impact the delivery of data to other investigators. As a result, there are automatic heuristics in the interferometric pipeline tasks that limit the generation of full-channel, full-source, full-field-of-view images if the estimated time to produce the full products is more than a few days. The pipeline imaging scripts can be modified after-the-fact to produce the missing imaging products. Investigators may wish to consider visiting their supporting ARCs if their home computing resources are inadequate. These requests should be submitted to the ALMA helpdesk.

13.4 Pipeline and the ALMA Quality Assurance Process

For each pipeline processing step, the Pipeline calculates Quality Assurance (QA) measures that are reported within the "Pipeline QA" section of the task weblog page, and a single QA score is reported for the entire processing step. These QA scores and any associated messages are primarily an indication of how successful that step was. They do not, by themselves, indicate whether the final products are deliverable.

That assessment (level 2 of the ALMA Quality Assurance process, QA2 – see Chapter 11) is made separately on the products from the pipeline using the ALMA Quality Assurance tool (AQUA), taking into account the requirements specified by the Primary Investigator when the proposal was submitted. Only data that pass (or semipass) QA2 will be ingested into the archive and delivered. The pipeline QA scores for each task are passed to the AQUA subsystem to aid in this assessment, as is the pipeline weblog, but the pipeline scores themselves do not determine the QA2 status of the OUS.

The main point of the pipeline QA scores and associated messages or warnings is to indicate whether any steps should be looked into more carefully, to possibly improve the processing and final products. It is important to note that data may pass QA2, but still benefit from a more careful manual review (e.g. data flagging, continuum identification, clean masking). The pipeline QA scores and weblog are tools to help researchers make this assessment. See the Pipeline User's Guide for more details.

13.5 Output of the Pipeline Processing

The main output from the pipeline are final images and cubes in FITS format, a weblog with informative plots and information on each pipeline task, and calibration and flagging tables and processing scripts that will either convert raw data into calibrated measurement sets, or run the pipeline tasks on the raw data to reproduce the pipeline products. There are also text files that can be edited to modify the pipeline run (e.g., by changing the flux scale or the continuum frequency ranges used for continuum subtraction - see the ALMA Pipeline User's Guide for details).

Regardless of how your data have been processed (manually or through the pipeline), there will be two types of archive products: the processed products (FITS images, scripts etc.), and the original raw visibility data. You can inspect/manipulate the processed FITS images in any package you like (although CASA is recommended). However, if you want to re-make images or combine data e.g. as in Figure 13.3, you will also need to download the raw visibility data and re-run the calibration scripts in CASA. Once you have regenerated the calibrated visibility data you can re-image (e.g. at a different spectral resolution) using the provided imaging script as a template. If you want to self-calibrate there are CASA Guides to help, or you can consult your local ARC through the helpdesk.

Chapter 14

Data Archiving

14.1 Introduction

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

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

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

An ALMA Science Archive Manual is available from the ALMA Science Portal (https://almascience.eso.org/documents-and-tools/cycle5/science-archive-manual).

14.2 Data Flow and Archive

Data from the correlator, together with monitor and weather data, are sent via dedicated optical fiber links (1-10 Gbit/s) to the OSF, where they are archived. A peak data rate of 66.6 MB/s can be sustained for short periods of time, i.e. days. This peak rate is a technical limitation of the data capture and data flow systems and will be imposed by the Observing Tool at the proposal validation stage.

The Pipeline processing system and Archive storage system have been designed to cope with an average data rate of 10% of the peak rate, i.e. 6.6 MB/s leading to a yearly amount of 200 TB of data. The volume of data products from the pipeline are expected to approximately equal the raw data volume, resulting in about 400 TB/yr being ingested into the ALMA archive.

The ALMA Archive at the OSF is designed to provide up to a year of temporary storage for the instrumental data (in the form of files in the "ALMA Science Data Model", or ASDM, format) and the monitoring data. The instrumental data are then transferred to the main archive at the SCO, where the pipeline is run and from where the data and pipeline products are distributed to the three ARCs. The process of copying the data to any of the archives involves a replication of the metadata (support data) and of the bulk data (ASDM and FITS files). All data transfer is done over the network. Metadata replication from SCO to the ARCs happens within a few seconds, transfer of bulk data can take longer (up to several hours), depending on the amount of data to transfer to the individual ARC.

PIs of large programs are expected to return processed data products to the project. These data will be transferred to JAO from where they will be ingested into the ASA.

Figure 14.1: The ALMA Archive is in the center of the ALMA data flow.

As soon as data is taken and has passed (or semipassed) the QA0 quality control step (see Chapter 11.2), its metadata are harvested from the AFA into the ASA and made available for search. This allows archival researchers to see which data will become public in the future. Metadata harvested include all the information needed to describe the observations, including date and time, source coordinates, frequency settings for each spectral window, and spectral and spatial resolution.

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

14.3 PI Data and Data Delegation

The unit of data delivery to the PI is the OUS. As soon as QA2 on an OUS has passed and the data products have been replicated to the user's home ARC, PIs will receive an email notification containing a link to their data, and supporting information. This sending of the notification to the PI triggers the start of the proprietary period of 12 months for standard proposals and 6 months for Director's Discretionary Time (DDT) proposals. Extensions of the proprietary period can be granted under certain circumstances. When a Group OUS consisting of several Member OUSs is processed (for example, a combined TP, 7-m and 12-m Array observation), the Group OUS products are released separately, with their own 12-month proprietary period.

The data deliveries consist of one or more bundles of raw (ASDM-format) data, and one "products" bundle tar file which includes the FITS files, logs, scripts, QA information and calibration tables.

Often PIs want to make their proprietary data available to collaborators, e.g. CoIs. To this end a data delegation service is available so that PIs do not have to give away their Science Portal password to anyone. Instead, PIs can give access rights to the data of a project to any registered ALMA user. To do so, PIs need to log into the Science Portal, go to their user profile page in the top right corner of the Science Portal page and then add delegates in the "Project delegation" tab.

14.4 Archive Query

At http://almascience.org/alma-data/archive users can query the holdings of the ALMA Science Archive (Figure 14.2). They then can download the data corresponding to their queries, if those data are public or if the users are authenticated and have the proper access rights to those data.

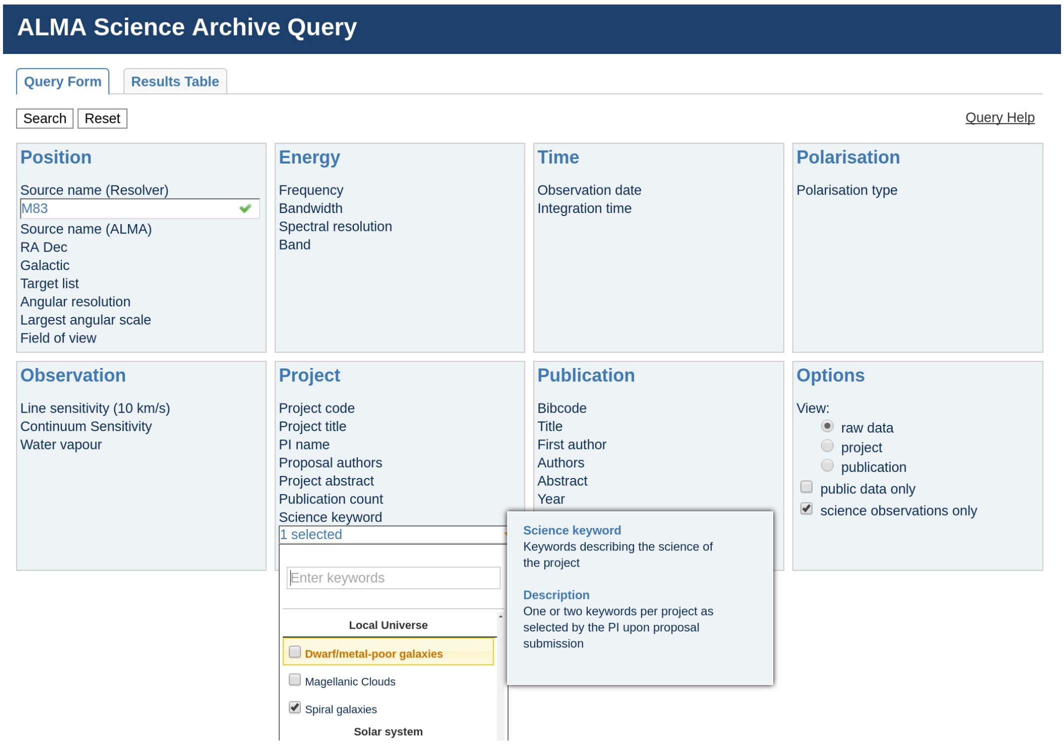

Figure 14.2: The ALMA Science Archive Query Form

Queries can be made by physical quantities along the Position, Energy, Time and Polarization axes. Help on how to query is provided in the tooltips of the query fields as well as through the “Query Help” link. Query constraints using standard operators for strings (*, ?) and numbers (>, <, ..) can be placed into the form fields. There is also a name resolver available for non-solar system objects (Sesame), which queries the Simbad, NED and Vizier databases. No operators can be used in the name resolver field.

By default, the Archive Query Interface will present users with the metadata of the observations, although they can choose to also see the metadata per project or per publication. On the results page, users can sort and subfilter the results and add or remove columns from the result table to narrow their search. The sky coverage of the observations is visible in a graphical window.

The results page also offers to download the results in VOTable, TSV (Tab-separated values) or CSV (Comma-separated values) format for further processing in tools like topcat[1]. The URL used for the exporting

[1] http://www.star.bris.ac.uk/~mbt/topcat/

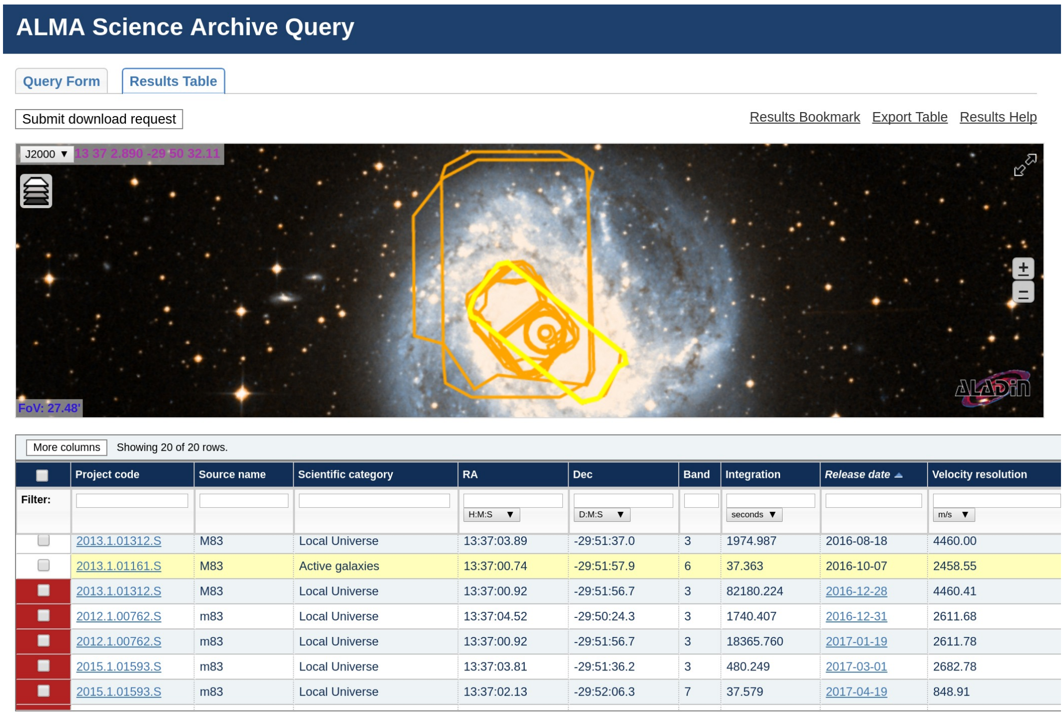

Figure 14.3: The ALMA Science Archive Results Page

of the query results can also be modified to access the ALMA Archive queries programmatically. Astroquery[2] can be used to encapsulate and simplify programmatic access to the ASA. Neither Topcat nor Astroquery are part of the offical ALMA software, and support of users of these tools will be limited to best efforts.

14.5 Request Handler

Once data of interest are defined, they can be selected via checkboxes on the results page and submitted to the ALMA Request Handler for download. The data are displayed in the full OUS hierarchy. This allows users to check immediately if there are data from other Member OUS or higher-level Group OUS products available. By default, only data products (i.e. FITS-format, and ancillary processing documentation) will be selected for download. Users who want to download the raw data as well are asked to select "include raw" before hitting the Download button or to select the desired data manually. The sources row can be expanded to show all the source names that are part of the package. Furthermore, a readme file is provided summarising the structure of the data and providing the full links to the OUS.

Four possibilities exist for the download itself:

- The first option is to use the download script. This script runs under Linux and MacOS and downloads files in parallel streams and is also adapted for downloads to a processing environment where no web browser is available.
- The second option is to use the download manager applet with the FireFox browser. This method is very convenient and allows for parallel downloads, too. It requires a Java Browser Plugin.
- It is also possible to use the download manager through Java Webstart technology.

[2] `https://astroquery.readthedocs.org/en/latest/alma/alma.html`

- Finally, a page with the links to all selected files can be displayed which then can be conveniently downloaded, e.g. using a browser plugin like "DownThemAll".

The Request Handler also allows selecting individual files for download. Only the data deliveries that users have permission to download can be selected. If they are not authorised to access any data delivery in the request, no "Download selected" button appears.

If the user is authenticated before requesting data for download, the request will be stored. This allows users to go back to previous requests. Note that these requests are stored only at the ARC the user is currently accessing. If data should be downloaded from a different ARC, then a new request has to be issued. For very large data requests PIs or archival researchers have the possibility to ask via the ALMA Helpdesk for data delivery on hard media, i.e. USB hard-disks. The ARCs may have different policies regarding the details of the shipping of the hard-disks.

Appendix A

Antennas

A.1 Design and Properties

ALMA is composed of a total of 66 antennas, 54 with a diameter of 12 m and 12 with a diameter of 7 m. The four 12 m antennas used for total power observations and the twelve 7 m antennas together form the Atacama Compact Array (ACA). The ALMA antennas are manufactured by three different contractors. These are VertexRSI (North America) which provided 25 12 m antennas, Alcatel Alenia Space European Industrial Engineering MT Aerospace (AEM, Europe), which also provided 25 12 m antennas and Mitsubishi Electric Corporation (MELCO; East Asia), which provided the four 12 m total power antennas and the twelve 7 m antennas (Figure A.1).

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

- Range of Ambient Temperatures: -20 °C $\leq T_{amb} \leq +20$ C
- Gradient of temperature: Delta(T_{amb}) $\leq$ 0.6/1.8 °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 $\leq$2.0 arcsec all-sky. Offset pointing $\leq$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 $\pm$5 mm. Maximum focal displacement (Z) of $\pm$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 5 have both 12 meter and 7 meter diameters, with the receivers mounted at the secondary (Cassegrain) focus. The 12 m dishes have a focal length of 4.8 meters, but the distance from the secondary focus to the plane of the subreflector of the 12 m antennas is 6000 mm, giving an effective focal ratio f/8, with an effective secondary focal length of 96 m and a plate scale of 2.15 arcsec per mm. The subreflector has a diameter of 750 mm. The 7 m dishes have a focal length, to the primary focus, of 2.572 meters. Given an effective focal ratio f/8, an effective secondary focal length is 56 m. The subreflector has a diameter of 457 mm.

The main reflectors of the ALMA 12 m and 7 m antennas are composed of individual panels. The size and number of panels varies between the different types of antennas:

VertexRSI: 264 panels spanning 8 rings with 12 (rings 1 and 2), 24 (rings 3 and 4), and 48 (rings 5 through 8) individual panels which are roughly a half-meter-square in area.

AEM: 120 panels spanning 5 rings with 8 (ring 1), 16 (ring 2), and 32 (rings 3 through 5) individual panels which are roughly one-meter-square in area.

Melco 12 m: 205 panels spanning 7 rings with 5 (ring 1), 20 (rings 2 and 3), and 40 (rings 4 through 7) individual panels which are roughly one-meter-square in area.

Melco 7 m: 88 panels spanning 5 rings with 4 (ring 1), 12 (ring 2), and 24 (rings 3 through 5) panels which are each roughly one-meter-square in area.

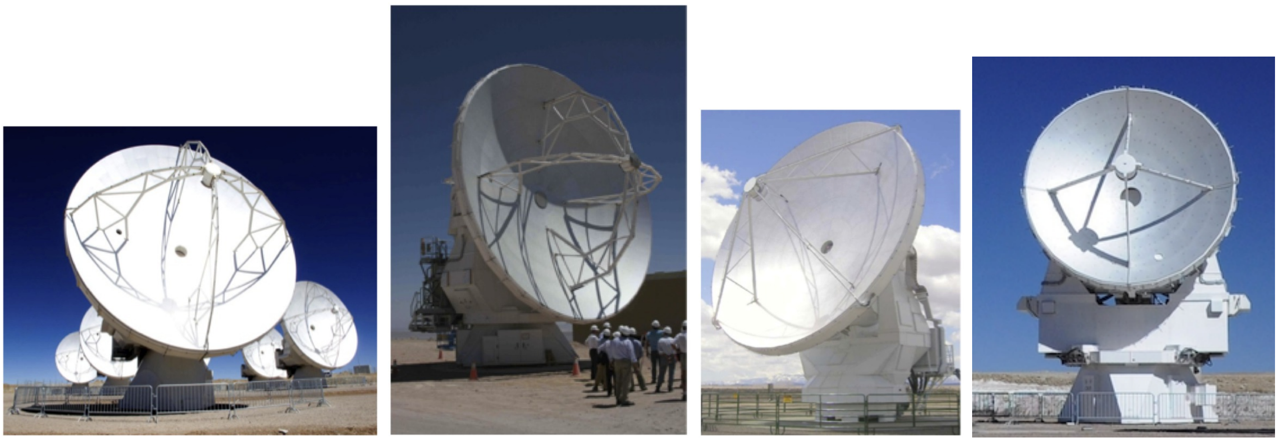

Figure A.1: The four different ALMA Antenna designs: Vertex 12 m, MELCO 12 m, AEM 12 m, and MELCO 7 m (from left to right).

Each panel has up to 5 adjustment screws, which can be used to optimize the surface accuracy of the individual antennas (based on holographic measurements). The surface of the panels are etched to scatter optical and near infrared solar radiation.

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

	BUS	Number Rings/ Panels	Panel Material	Quad type[1]	Cabin	Drive System[2]	Metrology System[3]
Vertex	CFRP Al Invar	8/264	Al	+	Steel	Gear	4 linear displacement sensors + 1 two-axis tiltmeter (above the azimuth bearing)
Melco 12 m	CFRP	7/205	Al	+	Steel	Direct	Reference Frame metrology
Melco 7 m	Steel	5/88	Al	+	Steel	Direct	Thermal (main dish), Reference Frame metrology
AEM	CFRP Invar	5/120	Nickel Rhodium	x	CFRP	Direct	86 thermal sensors + 2 tiltmeters in yoke arms

Table A.1: Design Properties of the Different ALMA Antennas.
Notes: **1** Shape of the quadrupod supporting the subreflector as seen looking along the optical axis of the antennas when they are pointed to the viewer. **2** A gear drive consists of a main motor driving a series of connected reduction gears (i.e., gearbox) that do the actual precision work. A direct drive system does not require such gears and takes the power directly. The direct drives used in ALMA antennas are magnetically supported. **3** Jointly used to correct in semi-real time the pointing of the antennas, under a wide range of environmental conditions, to meet the ALMA specifications.

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

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

The different antennas use a combination of steel, aluminium, Carbon Fiber Reinforced Polymer (CFRP) and Invar to achieve the best compromise between stiffness, robustness, smoothness, and low thermal expansion (see Table A.1 for a summary of properties). Common to all antennas is that they have a steel pedestal.

All antennas have built in metrology systems which allow thermal and wind deformations to be computed and corrected. For these purposes, the antennas are fitted with thermal sensors, linear sensors and inclinometers (tiltmeters).

The Vertex antennas have a drive system that is gear-driven whereas the AEM and MELCO antennas have magnetically supported direct drives.

The antennas are controlled using the ALMA Control Software (ACS). ACS sends instructions to the Antenna Bus Master (ABM) computer, which are then sent to the Antenna Control Unit (ACU) through a CAN bus.

A.2 Antenna Foundations

The antennas are placed on specially-designed concrete pads to guarantee stable orientation and location (Figure A.2). All antennas are attached to the pads at three points at the vertices of a triangle. The three points (inserts) are located on a circle centered at the antenna pad with a spacing of 120 degrees.

This interface guarantees a position repeatability error of the antenna, considered as a rigid body, not exceeding the values below:

- X/Y plane < 2 mm (peak to peak)
- Rotation around Z < 30 arcsec (peak to peak)
- Parallelism with respect to Z +/- 10 arcsec with respect to Zenith

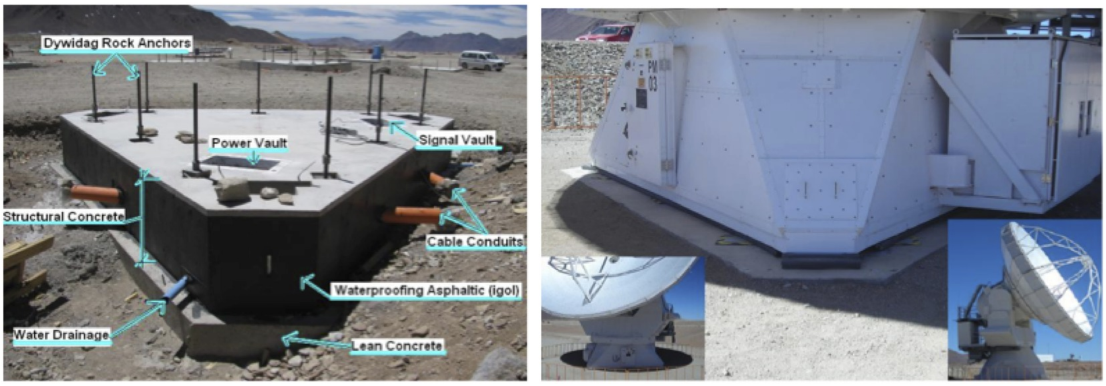

Figure A.2: Structure of an antenna pad (actual pad at the OSF) (left) and detail of antenna anchored to a pad (right).

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

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

Stiffness includes the inserts, the concrete pad and the soil. This does not include the kinematic mount lower part nor the foot of the antenna. The position of the pads are measured to a precision of 65 microns, and then monitored for stability for over two weeks. The pads are equipped with two vaults that contain the power, communication, Local Oscillator (LO) and data transmission cables that are connected once the antenna is placed on the pad.

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

A.3 Antenna Transportation

Antennas are moved from one pad to another using a specially-designed transporter (Figure A.3, righthand panel). ALMA has two of these vehicles. They are 20 meters long, 10 meters wide and 6 meters high, and each has 28 tires. The transporter positioning system performs a fine positioning of the antenna before setting it down on the foundation in the 3 in-plane degrees of freedom (x, y, rot -z) and 2 in-tilt (rot-x, rot-y). Adjustment in each of the 5 adjustment axes can be done independently. The adjustment range of the antenna positioning system compensates for the inaccuracy of the vehicle position with respect to the antenna foundation (which must be smaller than 10 cm) to achieve the required antenna positioning accuracy. The antennas can be positioned to within a few millimeters, ensuring accurate placement on the antenna foundation pads. More information on the transporters can be found on the ALMA EPO pages[1].

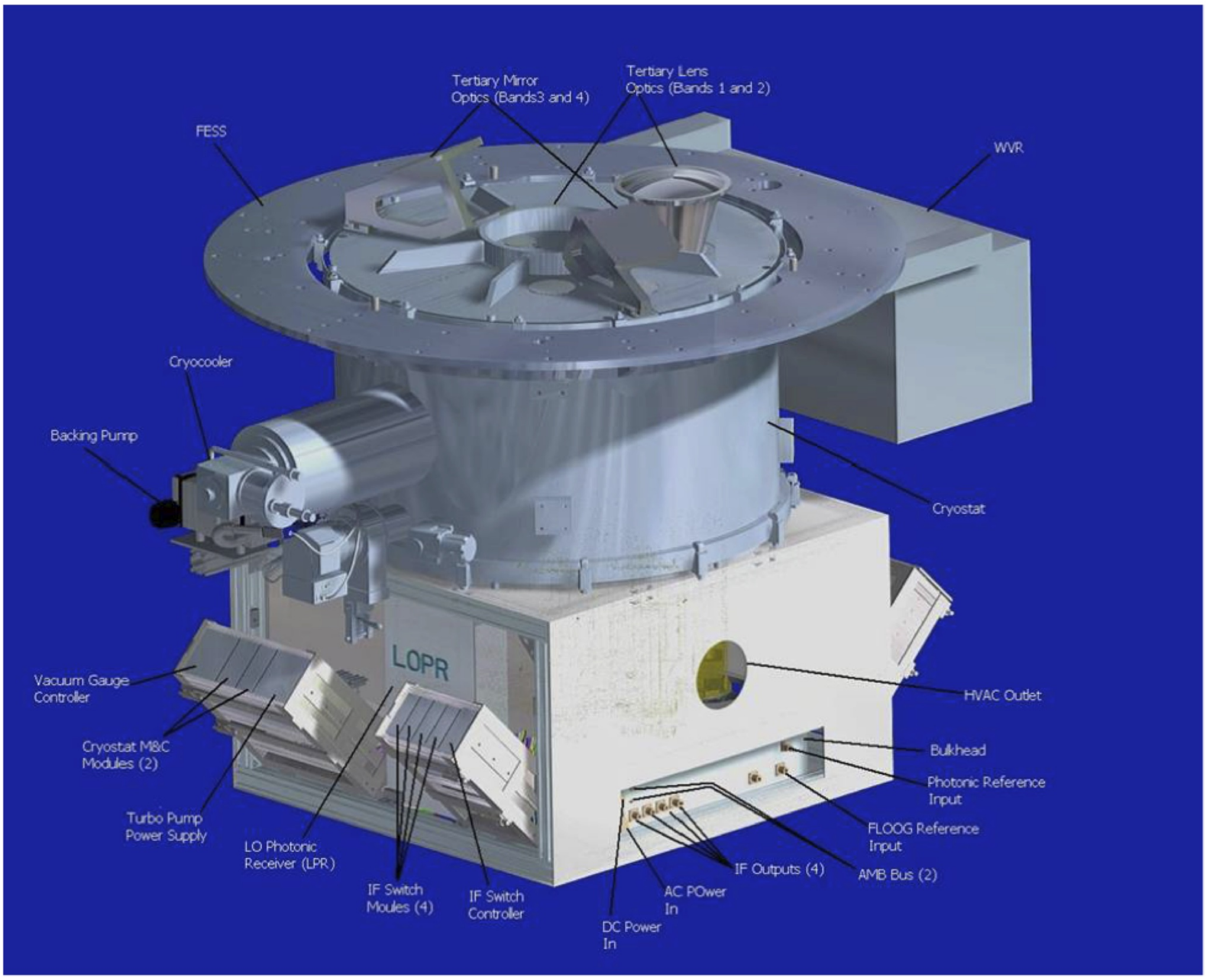

Figure A.4: Side view of ALMA front end showing cryostat assembly, with room temperature unit below.

A.4 Cryostat

The ALMA front end consists of a large closed-cycle 4 K cryostat containing individual cold cartridge assemblies (CCA) with mixers and LO injection for each band, along with room temperature electronics for the IF and LO for each band (the warm cartridge assembly, WCA) and fore-optics and entrance windows for each band.

[1] http://www.almaobservatory.org/en/technology/transporters

The water vapor radiometer (WVR) is mounted to one side of the cryostat using a pickoff mirror to direct the antenna beam into the WVR. The Amplitude Calibration Device (ACD) is mounted above the front end, and is described in Section A.5. Figures A.4 and A.5 show overviews of the front-end unit, with the cylindrical cryostat on top and the room temperature electronics beneath.

All of the receiver cartridges are in the same cryostat, with the mixers thermally-coupled to the same 3-stage Sumitomo cryocooler (Figure A.6). The three stages have nominal temperatures of 4 K, 15 K and 110 K. To avoid overloading the cooler, only three bands can be switched on at a time. It takes about 1 minute to switch between any of the bands that are switched on at a given time. For bands that are off, the time to fully thermally-stabilize them from an off state is 15 minutes – this is mainly to ensure a flat bandpass shape. All of the receivers are mounted off-axis to avoid extra rotating band-selection mirrors, which necessitates a pointing offset of the antenna to change band. The band pointing offsets are known and well-measured; the reference band for pointing is Band 6, and all offsets are with respect to this band. The four higher-frequency bands (Bands 7-10) are mounted close to the central boresight to minimize aberrations.

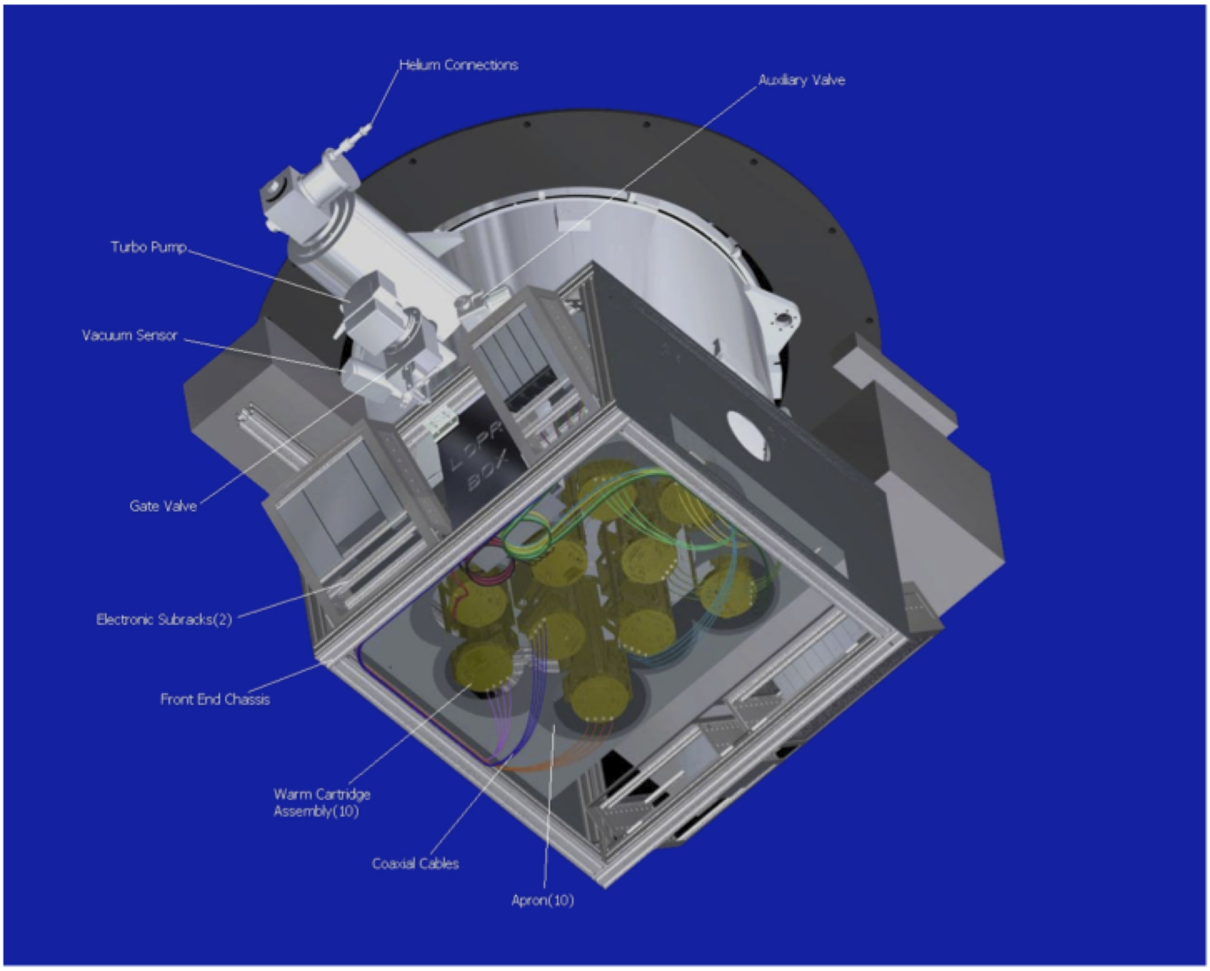

Figure A.5: Bottom view of ALMA front end, showing WCAs.

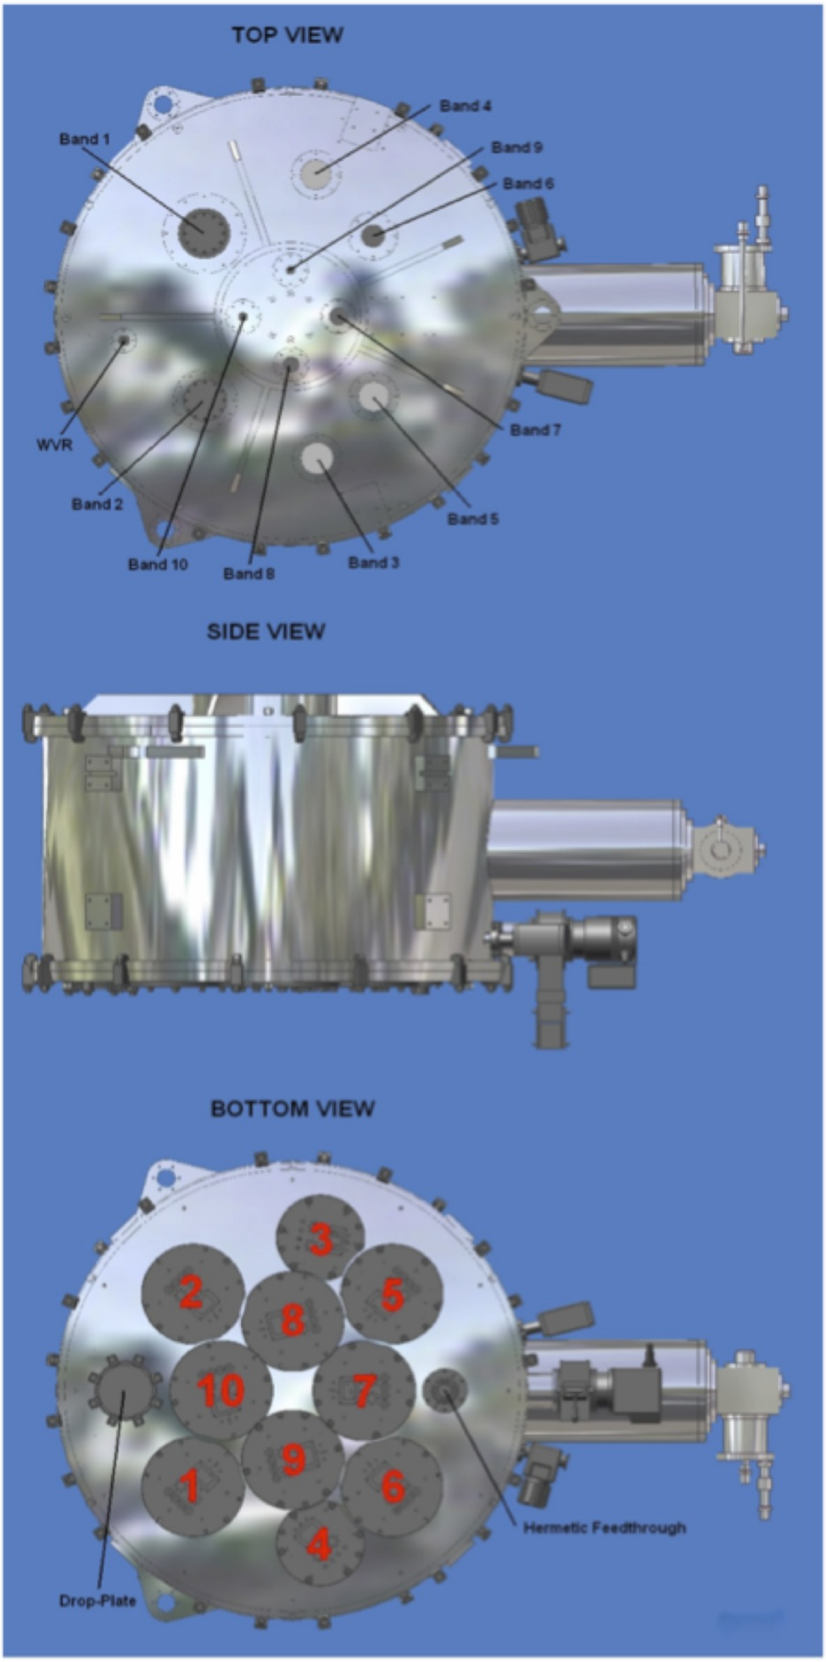

Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view).

A.5 Amplitude Calibration Device

The ALMA specification for relative amplitude calibration repeatability[2] has been set to be better than 1% for frequencies below 300 GHz and better than 3% for all other frequencies covered by the ALMA Front End. To achieve this goal, ALMA has adopted a two-load amplitude calibration approach.

The Amplitude Calibration Device (ACD) is located above the cryostat. It consists of a robotic arm attached to the top plate of the front end (Figure A.7). The arm holds two calibration loads, one at ambient (i.e., receiver cabin) temperature and the other one maintained at 80 °C (353 K). In addition, this arm also holds a solar filter to attenuate solar radiation during observations of the Sun. The arm is designed to allow the two loads to be placed in the path of any of the receiver beams (Figure A.8). Typically it takes 2 seconds to move the arm from the park position to the position where one of the loads is in the beam, and also 2 seconds to change between loads.

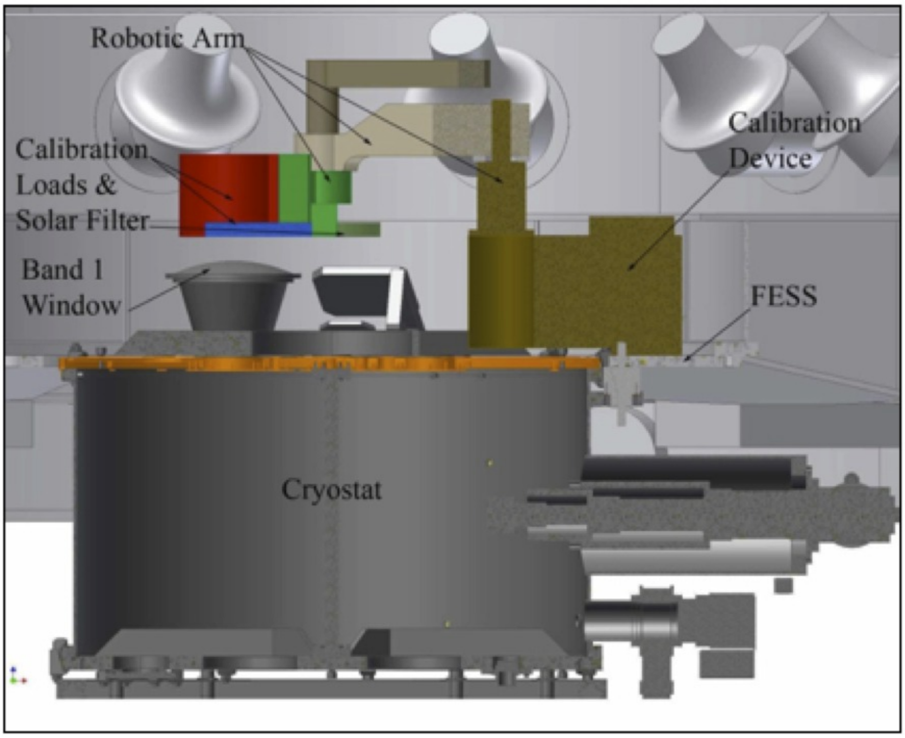

Figure A.7: Lateral view of the ACD on top of the ALMA front ends.

To accurately calibrate radio astronomical data to a temperature scale, the actual brightness of the two loads has to be precisely known. Critical to this calibration precision is the coupling of the load to the beam of a given band. This coupling must be very good at any telescope elevation and free of reflections of the load emission. This is because any reflection from the loads back into the cryostat would be terminated at a different temperature and would cause standing waves. Both loads have thus been designed so that the actual effective brightness temperature and that computed from the measured physical temperature (with sensors embedded in the loads) using known emissivities differ by, at most, ±0.3 K and ±1.0 K for the "ambient" and "hot" loads, respectively. This requirement also sets a limit to the fluctuations and departure from the set temperature that are allowed for the "hot" load. Furthermore, the return loss specifications for these loads are -60 dB and -56 dB, respectively.

[2]"Calibration Repeatability" means being able to make repeated measurements of the same flux densities (or brightness temperatures) for the same source under different conditions (weather, telescope elevations, front-end status, etc.).

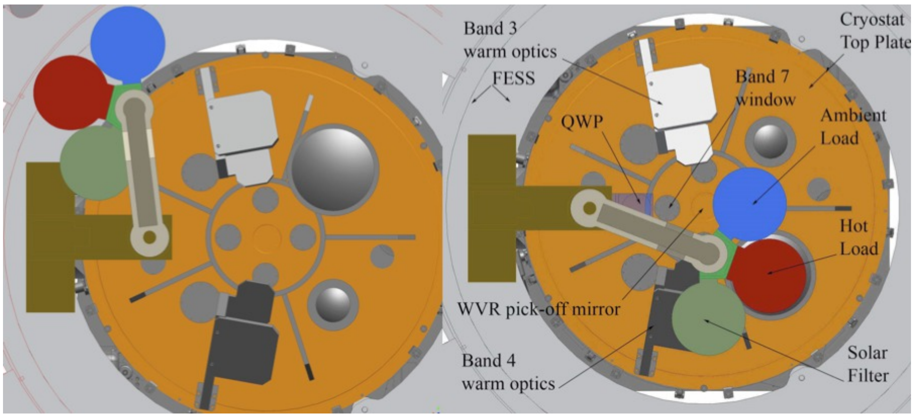

Figure A.8: Top view of an ALMA front end showing the robotic arm of the ACD retracted during normal observations or on top of one of the front-end inserts for calibration. The current design has been improved by placing all the loads in a wheel.

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:

$$V_L = \alpha V_{L_1} + (1-\alpha)V_{L_2} \tag{A.1}$$

where α is the weighting factor, and V_{L_1}, V_{L2} 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:

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

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

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:

$$T_{Sys} = T_{Cal}\frac{V_{Sky}}{V_L - V_{Sky}} \tag{A.4}$$

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 5.

A.6 Water Vapor Radiometers

In the mm and submm regions, variations in the water vapor distribution in the troposphere that move across an interferometer cause phase fluctuations that degrade the measurements. ALMA uses the so-called "Water Vapor Radiometry" technique to correct for these phase fluctuations. Water Vapor Radiometry involves estimating the excess propagation path amount due to water vapor along a given line-of-sight by measuring the brightness temperature of the sky at frequencies near the atmospheric water vapor resonances. These temperatures can then be transformed into a path length and the difference between any pair of antennas in the array gives the final phase fluctuations to be corrected for a given baseline. ALMA has implemented this technique by placing a Water Vapor Radiometer (WVR) on each 12 m antenna (the 7 m antennas do not have WVRs). For the WVRs to be effective, the measurements have to be taken with a cadence that is fast enough to map the actual variations in the atmosphere. The relevant shortest timescale is the antenna diameter divided by the wind speed as the path delay is averaged over the whole antenna beam and cannot therefore be corrected at any finer time resolution than that. The effective diameter is about 10 m for the ALMA antennas and the relevant windspeed is usually 10 m/s or a bit less so the fastest necessary sampling speed is 1 Hz. On timescales shorter than this 1 Hz timescale, the water vapor path fluctuations are expected to lead to small apparent pointing fluctuations which are analogous to the seeing effects in single-aperture optical telescopes. ALMA selected the 183 GHz line because it is quite bright and allows a more compact design than would the 22 GHz water line. It was decided to measure the temperature of the 183 GHz line in four regions offset from the center using filters of different bandwidths. The positions of the filters are indicated as blue boxes superimposed on the profile of the water vapor line in Figure A.9. The sensitivity specification for the WVRs is 0.08–0.1 K per channel RMS.

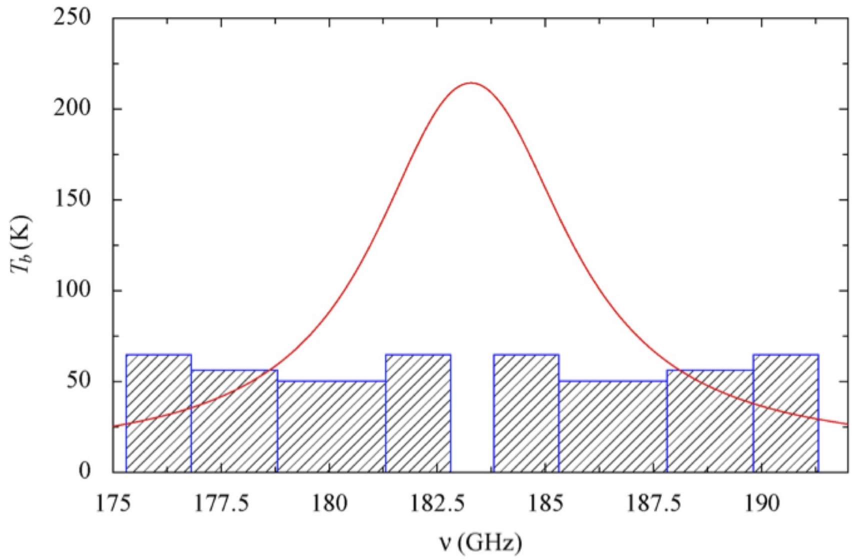

Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

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

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

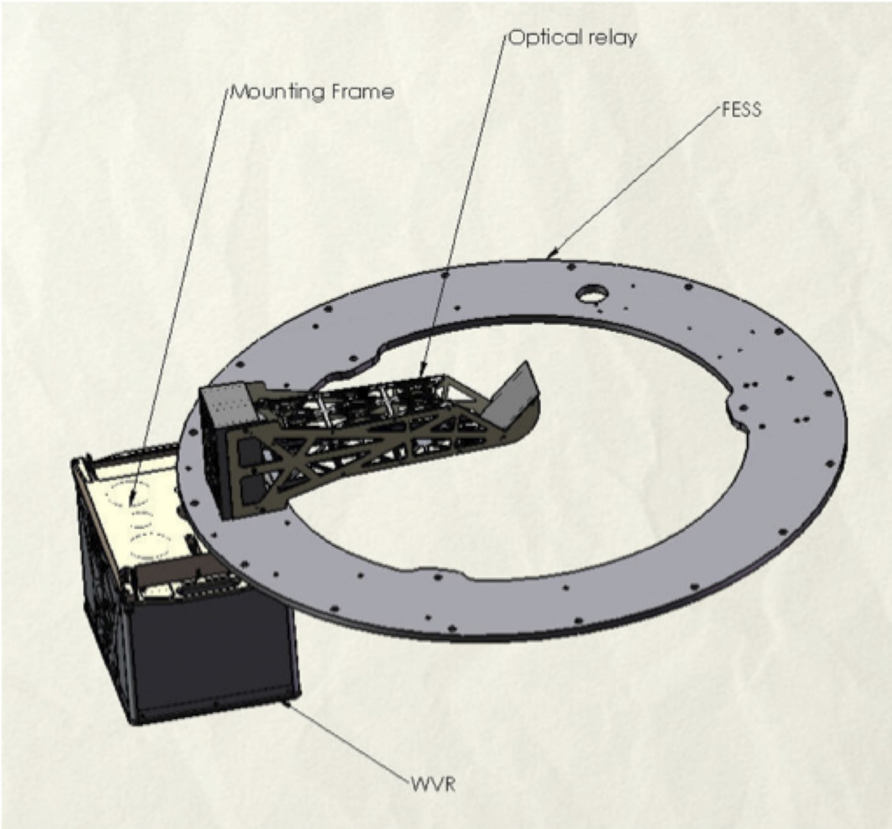

Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

atmosphere (i.e., the precipitable water vapor). There are also variations due to the changes in bulk ambient temperature at different heights above the observatory. It is expected that these could become significant during day time and some techniques are being currently studied to try to measure them (including thermal sounders of the atmosphere that use the profiles of the emission of the oxygen molecules). The brightness temperature variations of the sky that the WVRs have to detect are sometimes quite small, so the quality of the receiving system becomes very important. In fact, the current specification for the ALMA WVRs is that they need to allow corrections of the path fluctuations (in μm):

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

where w is the precipitable water vapor (PWV) content in mm along the line of sight, and L_{raw} the total fluctuations observed at any given time. Therefore, this formula includes the expected error of about 2% in measuring the total fluctuations, and states the total resulting path errors after correction (L_{corr}). For a 1 mm PWV, the residual term in the formula would be 20 μm. The stability specification for the WVRs is very stringent (0.1 K peak-to-peak over 10 minutes and 10 degree tilts). To achieve this, a Dicke-switching-radiometer approach was adopted. The input into the mixer is switched periodically (5.35 Hz) between two calibrated loads (the "cold" and "hot" loads at 293 K and 351 K, respectively), and the sky using a rotating vane embedded in the light path as shown in Figure A.11.

Calibration of the measurements is done following the usual method for a 2-load system. The ratios of the output powers when observing the "hot" and "cold" loads can be used to determine the receiver temperatures. Furthermore, these output powers from the loads are also used to extrapolate to a virtual load that has a brightness temperature similar to that of the atmosphere. The specification for the absolute accuracy of the calibration is 2 K (maximum error). The mixer system is an un-cooled DSB Schottky diode pumped by an

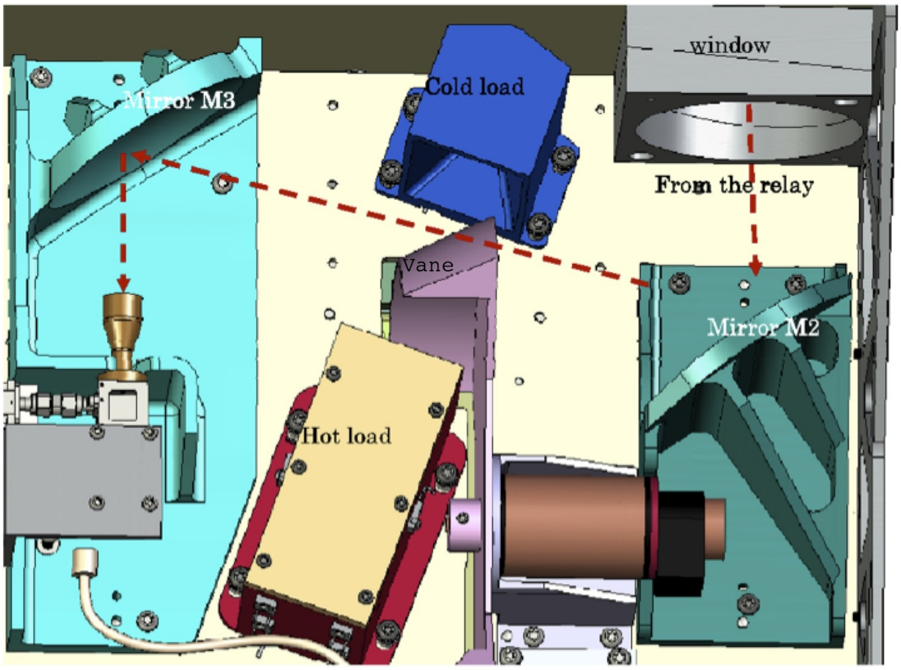

Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

LO at 15 GHz that undergoes 2 stages of multiplication. The receiver noise temperature is about 1000 K. After amplification, the IF signal is split into four complete chains (one per filter) and a bandpass filter is applied to select the four desired sampling regions in the profile of the water vapor emission line. In each IF chain, the signal is detected with diodes and after a Voltage-to-Frequency conversion, sent to the Control section for accumulation and control. There is a possibility of LO leakage out of the WVRs that could affect the ALMA receivers in the same antenna and others nearby. To avoid coherence, all the WVRs are tuned to a frequency slightly different (offsets by consecutive integer multiples of 10 kHz up to the total number of WVRs available). The final products sent to the ALMA Control system are time-stamped, calibrated measurements of the brightness temperatures in the 4 filter regions. The path length error due to the PWV can be calculated from these brightness temperature measurements and used to correct the data. Corrections at the scales of the sampling rates of the WVRs are possible at the correlator and refinements for longer timescales can also be done offline in CASA using the wvrgcal tool. Currently both streams of data are being recorded (WVR phase corrected and WVR non-corrected) in order to rigorously assess effectiveness.

Appendix B

The LO and IF System

In this Appendix, the signal path, LO chain, and the relation between these systems and the observer's spectral setups are described. To the system, a spectral setup effectively consists of the settings of the local oscillators and correlator in the system such that each spectral window (spw) covers the desired lines and/or continuum frequencies. To the end-user, the spectral setup is normally defined in the Observing Tool (OT) just in terms of the observing frequencies and spectral resolutions, and there is no need to worry about the details of each LO setting. For full details of the OT and how to use it, see the OT User and Reference Manuals, available from the ALMA website[1] (and also in the OT itself).

The following sections show how the LO system works. For those only interested in the spectral setups and not the details of the components in ALMA, please see Chapter 6.

B.1 Functions of the LO and IF System

In the signal path from the front end to correlator, ALMA uses three frequency conversions. The associated LO and IF systems perform multiple functions:

1. Down-conversion of the sky frequencies to basebands in the range 2–4 GHz, which then alias down to 0–2 GHz for digitization.

2. Amplification and adjustment of the correct power levels into the digitizers.

3. Adjustment of the spw center frequencies (in the Correlator FDM modes) within the basebands. This is actually done in the correlator using the TFB LO, but can effectively be treated as a 4th stage of the LO system.

4. Application of frequency corrections for fringe rotation, and compensation for the slight differences in the Doppler shifts at each antenna due to the differential line-of-sight velocities with respect to the target.

5. Provision of geometric delay corrections.

6. Separation of the signal and image sideband. In the case of DSB receivers, selection of the wanted sideband(s). This is done through frequency offsets and phase modulation at each antenna using Walsh patterns.

7. Suppression of spurious signals and reduction of the effects of DC drifts in the samplers. This is done using phase modulation of the LOs using Walsh patterns.

[1] http://almascience.org/documents-and-tools/

Frequency down-conversion therefore effectively occurs in four stages: two hardware Local Oscillators (LO1 and LO2), a 4 GHz sampler/LO and a digital LO synthesised in the tunable filterbanks (TFBs) in the Correlator[2]. Section 6 shows how to setup the system to observe spectral lines (particularly multiple spectral lines) and continuum. Some other aspects of frequency setups are then discussed, including the usable bandwidth, spurious signals, and rules and limitations pertaining to this observing Cycle. An overview of the LO and IF operation in ALMA is given in Sections B.2 and in Section B.3, the hardware and how the LO frequencies are synthesised and distributed around ALMA are described.

B.2 Summary of Operation

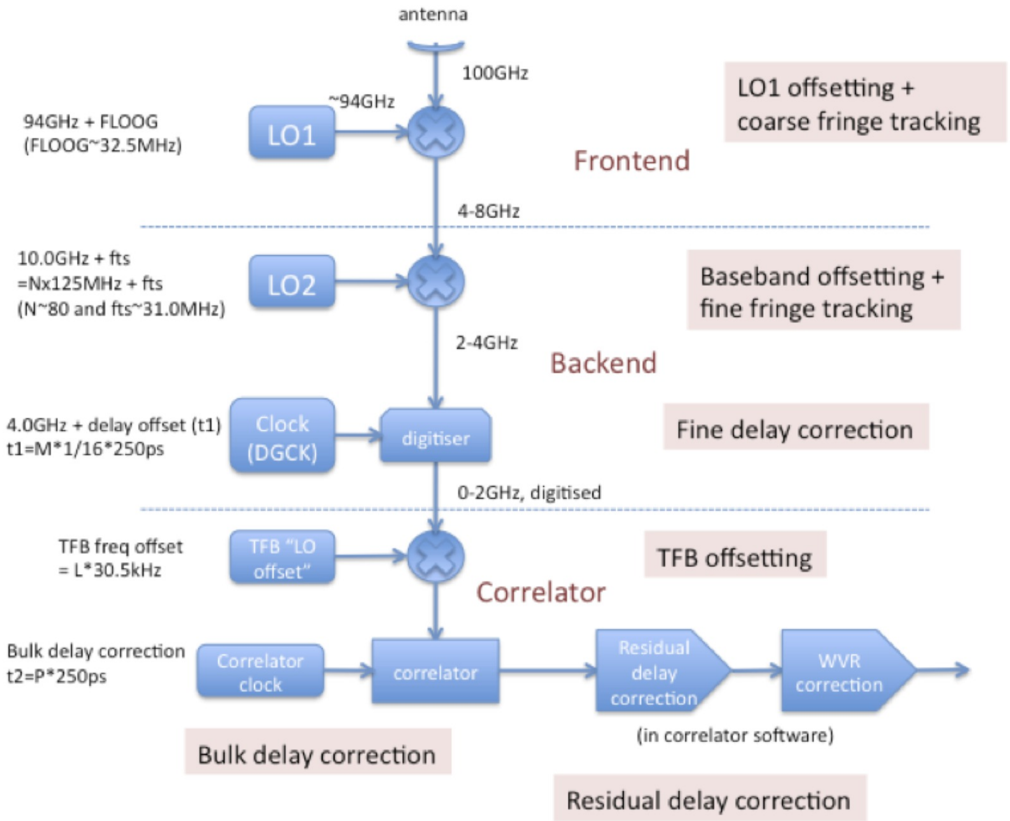

Figure B.1: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the front end, back end, and correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz seen in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized tuning steps, such as LO2 (which can be changed in a multiple ("N") of 125 MHz plus a finely-adjustable offset of "fts"), the TFB LO (which uses a multiple "L" of 30.5 kHz) and the Bulk Delay Correction (which has steps of 250 ps, with a factor of "P"). See text for descriptions of each stage.

Figure B.1 shows a simplified block diagram of the ALMA LO/IF system, showing example setups for an observing frequency centered on 100 GHz. Referring to this diagram, the system operates in the following way:

1. The front-end mixer uses LO1 to downconvert the observing frequency into an IF range covering up to 4-12 GHz. This wide range is needed to cover the IFs of all the ALMA bands, since the mixers for Bands

[2]Note that the ACA correlator is designed to appear like the 64-input Correlator to the end user, although it does not use TFBs in the same way as the BLC

3, 4, 7 and 8 have an output IF of 4-8 GHz, Band 6 a range of 5-10 GHz and Bands 9 and 10 a range of 4-12 GHz. Over most of the front-end tuning range, LO1 and the front-end mixer can be used in upper or lower sideband; although at the edges of the tuning band, only one sideband is possible. LO1 consists of a common reference component for all antennas, plus a smaller offset component generated in the FLOOG (First LO Offset Generator) which is different for each antenna (see LO1 Section B.3.4). The FLOOG is used to perform coarse fringe tracking (i.e. rough correction for the small offsets in the observing frequency at each antenna), to offset the LO1 frequencies slightly to suppress internally-generated interference, and for sideband separation or selecting the sideband. It is also used to offset the LO1 phase in conjunction with a Walsh switching pattern on the antennas. A 180 degree phase offset is used to remove spurious signals and DC offset while a 90 degree phase offset is used for sideband suppression and, in the future, for sideband separation on the DSB receivers (Bands 9 and 10).

2. In the back end (BE), the IF processor (IFP) splits the IF into basebands, each with a frequency range of 2-4 GHz, via a set of filters and tunable second LOs (LO2) (see IFS/IFP in Section B.3.5). LO2 is used to offset the individual baseband frequencies within the IF range. The LO2 and second mixer only operates in LSB, with a possible LO2 tuning range of 8-14 GHz. The LO2 signal itself is generated by a course synthesiser which can be set only in steps of 125.0 MHz, plus a second fine-tuned synthesiser (fts) which provides an offset in the range 20.0-42.5 MHz (marked as "fts" in Figure B.1). The limited fts range and the 125 MHz quantization means that LO2 setting is not fully contiguous; consequently there can be up to ∼30 MHz difference between the desired and the set value. With a single baseband, this can be compensated by a suitable offset of LO1, but with multiple basebands this is not always possible. So without additional correction, setups with multiple basebands could have the requested lines offset from the spw center by up to 30 MHz. However, the remaining differences in the different spws are compensated by applying an opposite offset to the TFBLOs (LO4 - see below.)[3]. An algorithm used by the OT and the realtime system generates the best LO tuning "solution" for LO1, LO2 and LO4 which minimizes the offset of the requested observing frequencies from the centers of the spws. Other uses of LO2 are that the finely-tunable fts is used for fine fringe tracking and LO2 can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband.

3. The 2-4 GHz analog IF signal from the second mixer in the IFPs is digitised (or sampled) with a 4.0 GHz clock (DGCK). A fine delay (or time) offset is applied to this clock in units of 1/16 of the clock period (250 ps) (the "fine delay correction").

4. In the FDM correlator mode, up to 32 digital filters (known as TFBs, or "Tunable Filterbanks") are applied to each digitised baseband signal, each of which can be individually adjusted across the baseband frequency (the TFB offsetting). This is effectively applying a digital LO (the TFBLO, or LO4), which is adjustable in steps of 30.517578125 kHz[4] and allows the spectral windows to be moved around within the basebands. At Phase 2, the TFB is centered on the baseband if the TFB "offset" is set to the default of 3000.0 MHz; it can be moved up to +/-900 MHz from that frequency, the range depending on the spw bandwidth. The TFB outputs are resampled and sent to the correlator. The TFBLO can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband. Finally the correlator software is used to perform the finest level of residual delay correction.

B.3 Frequency Generation and Distribution in ALMA

Figure B.2 shows a summary of the main units involved in the LO generation and distribution. The LOs are generated by the Central LO (CLO) (Section B.3.1) in the AOS Technical Building (lower half of diagram). A fiber-optic system is used to distribute these signals out to the antennas (Section B.3.1) incorporating a realtime path length correction system (Section B.3.3). In the antennas, the important outputs are LO1 (FE 1st LO) (Section B.3.4), LO2 in the IF Processer (Section B.3.5) and the digitizer clock (DGCK). All of these

[3]This is done automatically when the OT generates a spectral setup in an SB from a proposal. However, it is repeated at runtime. See https://safe.nrao.edu/wiki/pub/ALMA/AlmaLamaMemos/lamaMemo808.pdf for more details on the tuning algorithm

[4]The Phase 2 OT has an "adjust" button which quantizes the value entered by the user by this unit

are required in each antenna, shown in the upper part of Figure B.2[5]. In the following subsections, some of these components are described.

B.3.1 Reference and LO Signal Generation

The Central Local Oscillator (CLO) generates and distributes the reference, timing and LO signals to all ALMA components to ensure that antenna movement, electronics, and data acquisition are synchronized. These signals are distributed to the antennas through optical fiber using the light of three infrared lasers. The ALMA frequency and phase standard is a Hydrogen maser (installed in 2014), known as the Master Frequency Standard (MFS, lower right), which produces a signal at a frequency of 5 MHz. This is fed into the Central Reference Generator (CRG) module, which produces several signals at multiples of the 5 MHz signal. The 125 MHz signal becomes one of the standards used by many components in the ALMA system. At the AOS Technical building it is used by the Slave Lasers in the Laser Synthesizer modules (Section B.3.2). At the antennas, it is fed into the FLOOG, the Digital Clock (DGCK, see Section B.3.6), the IFPs, and the LO2 Synthesizers. The 2 GHz signal goes into the DTS cards at the antennas. All the reference signals are modulated onto a 1532 nm IR laser in the Central Reference Distributor (CRD) module. The CRD has an internal 48 ms (known as a TE, or Timing Event) clock that is also modulated into the same signal, and is used to synchronize many of the hardware events in the observatory. The modulated 1532 nm signal is sent to an optical distributor (with 80 outputs), the Low Frequency Reference Distributor (LFRD), that feeds it into the Sub Array Switch (SAS) modules, where it is merged with the signals from the Master and Slave lasers (see Section B.3.2). The reference and LO signals are fed to the antennas through the fiber-optic distribution system (see Section B.3.3).

B.3.2 LO1 Signal Generation and Distribution

LO1 is distributed through fiber-optics and regenerated photonically in each antenna front end by mixing the two infrared laser carriers (known as the Master (ML) and Slave Lasers (SL)) to produce a fixed frequency for all the antennas. Figure B.3 shows a block diagram of this part of the LO system. There are six Slave Lasers (Laser Synthesizers, LSs), each containing fpour tunable units, that can produce six different LO1 frequencies which allow simultaneous observations at different frequencies with different subsets of the array (multiple arrays or sub-arrays), and may be used for observing modes that will benefit from rapid switching between frequencies (e.g., band-to-band transfer or spectral scans).

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 six Laser Synthesisers (LS) fed into the six Central Variable Reference modules, which the LSs use to control 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 Warm Cartridge Assemblies (WCAs) in the front end (B.3.4). It is used by the software to set up the front-end observing frequency. With six LSs it is possible to generate six separate LO1 frequencies.
- The Photonic Reference Distribution (PRD) feeds the optical signals to the Sub Array Switch (SAS) module (one per antenna), which chooses the sub array to connect to.

Figure B.4[6] shows the three laser signals after combination in the Sub Array Switches (SAS). The Master and Slave laser signals have wavelengths of about 1556 nm and the laser carrier signal for the reference signals from the CRD has a wavelength of 1532 nm. The signals are distributed via a single-mode fiber optic line to

[5]Note: The figure needs an update. The frequency of 25 MHz displayed in blocks LORR and LORRD, LFRD, CRD, and CRG needs to be replaced by 125 MHz. In addition, IFP received a 125 MHz signal from LORR.
[6]This figure needs an update. The Slave Laser (variable) wavelengh is ∼1557 nm.

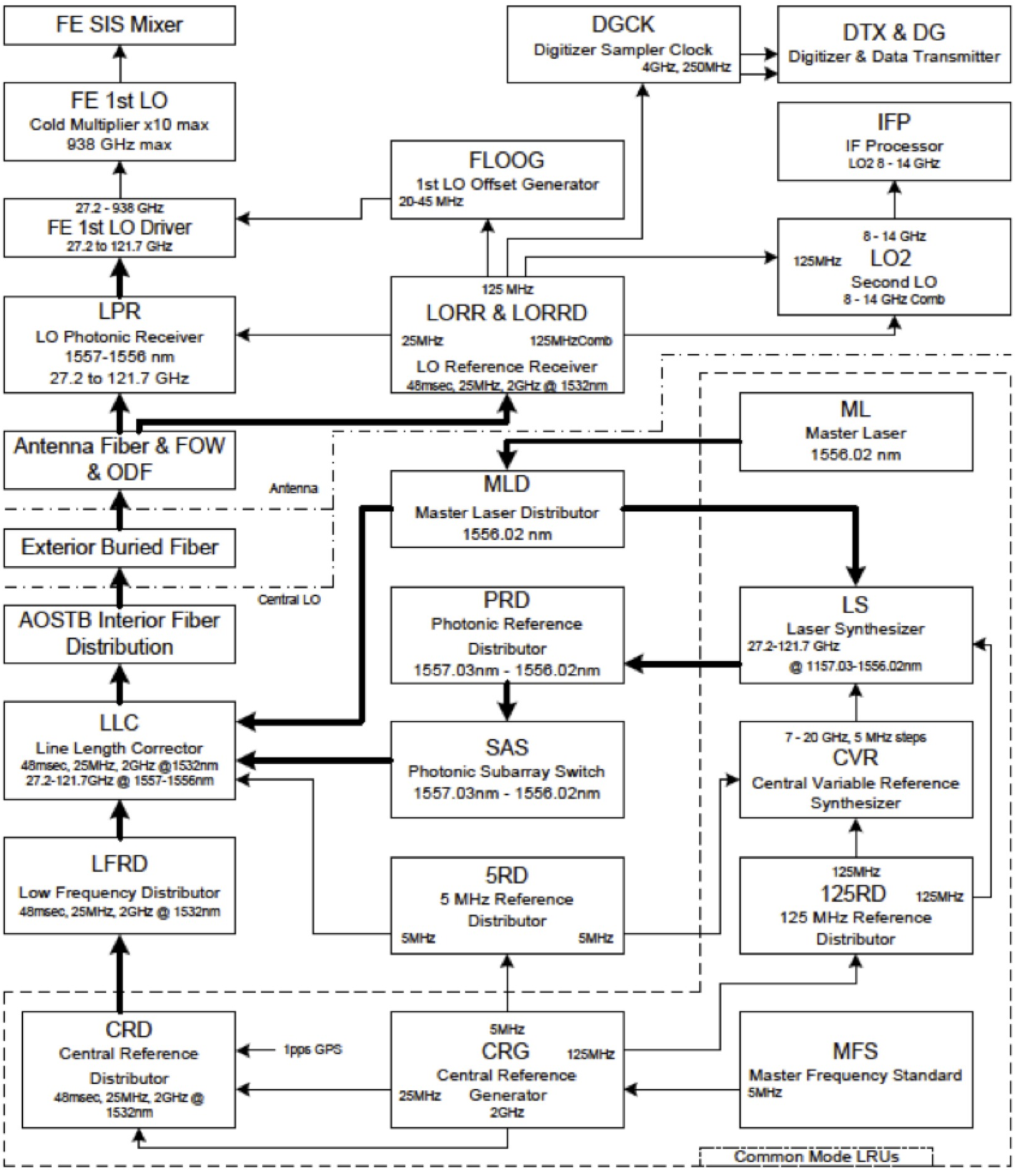

Figure B.2: Summary block diagram of the LO distribution system. The components in the lower section of the diagram (below the the dashed-dotted line) form the Central LO (CLO), located in the AOS Technical Building. The components above are located in each Antenna (only one antenna is depicted in this diagram). These are linked by the exterior buried fibers linking the Technical Building with the antenna pads (shown middle-left of the diagram). Thin lines with arrows represent cable distribution, thick lines represent fiber-optic distribution.

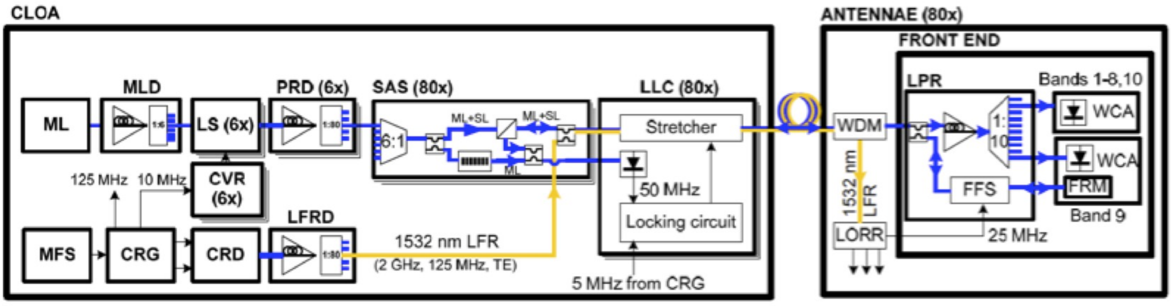

Figure B.3: LO block diagram, showing the Central LO (CLO) and the LO section in the WCA in each front end. Yellow lines represent common synchronization and reference signals fed to all antennas, blue lines are LO signals which are individual to each antenna. For a description of acronyms, see text.

each of the antennas. The fibers are distributed in buried trenches, and fed into the Cassegrain cabin on each antenna through Az and El fiber wraps. All are fed through Line Length Correctors (LLCs), which are used to correct for changes in the optical fibers. The LLCs are described in B.3.3 below.

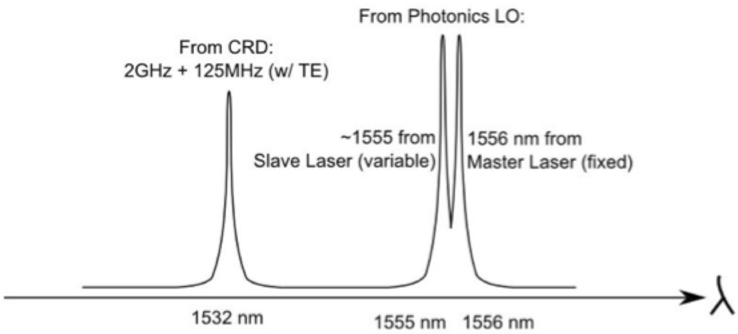

Figure B.4: The ALMA fiber signals. The 1532nm carrier contains the frequency reference signals, and the 1556/1555 nm carriers are used to remotely generate LO1. Within each antenna, the optical signals are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing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

B.3.3 Line Length Corrections

The LO Reference signals are generated at the AOS Technical Building and need to be distributed via optical fibers to all the antennas. In order to guarantee that the phase of the LO signals is stable during the observations for fibers of up to 15 km in length, compensation for changes has to be done in real time. The method adopted by ALMA is based on a round-trip optical interferometer. Phase fluctuations for an optical fiber transmission system are mainly caused by thermal expansion of the fiber and mechanical stresses, which produce birefringent effects and changes in the absolute polarization of the signals. These changes, in turn, cause differential group propagation delays (PDM) that show up as LO phase jitter. The method implemented by ALMA to correct for this is known as the Line Length Corrector (LLC). Part of the LLC can be seen in Figure B.3, and a more detailed block diagram of the system is shown in Figure B.5.

The two-wavelength laser synthesizer signal (master and slave lasers) is adjusted in polarization and mixed at the SubArray Switch (SAS) and then passed through a 3-port polarizing beam splitter assembly (PBS). The polarization is aligned so that all the light passes through the beamsplitter. It then passes through a piezo-driven fiber stretcher assembly and the fiber to the antenna. At the antenna end there is a 3-dB coupler, so that

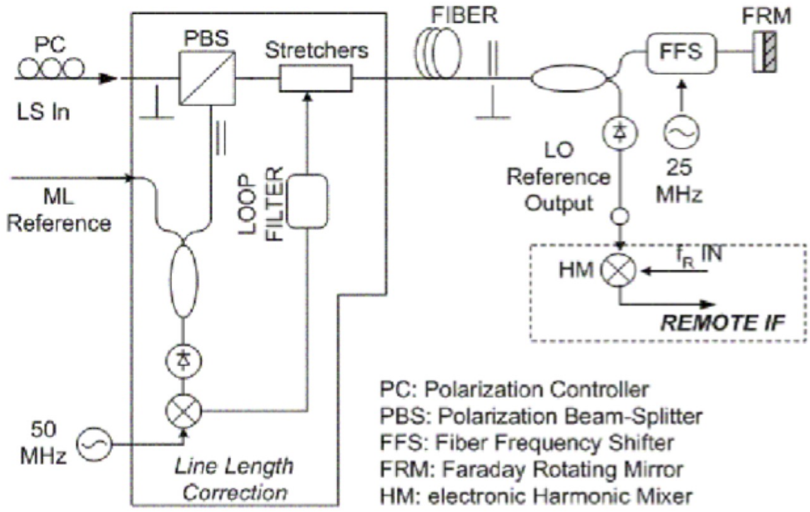

Figure B.5: Block diagram of the Line Length Corrector system for ALMA. The FRM (Faraday Rotation Mirror), shown upper right, is located in the Band 9 cartridge of every front-end. It rotates the polarisation of the incoming light, and the resulting reflected signal is fed back though the buried optical fiber, to be compared with the outgoing signal. This allows the optical path length to be adjusted in a closed loop using the fiber stretchers.

half of the light goes to the turnaround assembly and half to the photomixer in each WCA. The turnaround assembly consists of a fiber frequency shifter (located at the LO Photonic Receiver module) and a Faraday Rotator mirror located within the WCA of specifically the Band 9 cartridge in each front end. The frequency of the signal traveling back to the AOS technical building receives thus twice a frequency shift of 25 MHz, thus it comes back offset by 50 MHz from the original. The Faraday rotator reflects the signal but turns its polarization angle by 90 degrees to the incident polarization. This means that the outgoing and returning light is orthogonal everywhere along the fiber between the PBS and the Faraday Mirror. Back at the PBS, the returning signal is sent to a third port where it is mixed with a sample of the Master Laser reference signal in a low-frequency photodetector. This results in an output at the 50 MHz offset frequency. This output is compared in a phase detector with a 50 MHz reference signal and the phase of the whole loop is kept constant by a servo driving the fiber stretchers.

The current stretchers can cover ranges up to 5 mm in two modes. A “slow” mode (about 10Hz) copes with the large deformations (about 3 mm, allowing for some headroom at the ends of the ranges) and a “fast” response mode (about 1 kHz) copes with the small range variations (about 0.1 mm). The LLCs are reset to mid-range at the start of every SB execution.

B.3.4 The First Local Oscillator (LO1)

The reference signal required to tune LO1 in the receivers is obtained as the difference of the wavelengths of two phase-locked infrared lasers, the Master and Slave lasers. The Master Laser (ML) has a fixed wavelength of 1556 nm and the tunable Slave Laser (SL) is offset from this; both are generated in the CLO (see Section B.3.2). The offset frequency can be anywhere in the range 27 to 38 and 65 to 122 GHz. The beat note from the two lasers constitutes the Photonic LO Reference; the LO1 reference signal is generated from this by photomixers located in the Warm Cartridge Assembly (WCA) of each receiver. This reference signal is used to drive a YIG (Yttrium Iron Garnet) oscillator operating at frequencies around 10–30 GHz (the exact range depending on the band), via a Phase Locked Loop (PLL) circuit. This produces LO1 for the SIS mixers via two sets of multipliers (see example Figure B.6 for Band 7). The same photonic reference signal is distributed to all antennas in the same sub-array. However, to correct for different delay rates required in different antennas, the First LO Offset Generator (FLOOG) in each antenna generates a small but variable (and different) offset frequency in the range 20–45 MHz which is also fed into each PLL. The FLOOGs for all the antennas are continuously tracked during an observation.

B.3.5 LO2 and the IF Processor Units and IF Switch

The output of each front-end cartridge is connected to an IF Switch unit (IFS) situated in the front end, which selects between bands, provides some amplification, and has variable attenuators to set the output levels. The four (or two) outputs from the IF switch unit are fed into two IF Processor units (IFP), one per orthogonal polarization. Figure B.7 shows a basic block diagram of one IF Processor (only one polarisation channel is shown). The Bands 3, 4, 6, 7 and 8 receivers are dual-sideband (2SB), where both the upper and lower sideband signals are provided separately and simultaneously. So there are four outputs from each receiver cartridge in these bands, two sidebands times two polarisations. Each output has an IF bandwidth of up to 4 GHz. For Bands 9 and 10, the receivers are double-sideband (DSB), where the mixer produces a downconverted output from signals in *both* USB and LSB. These bands have only two outputs, one per polarization, but the signal IF bandwidth of these DSB receivers is 8 GHz per output.

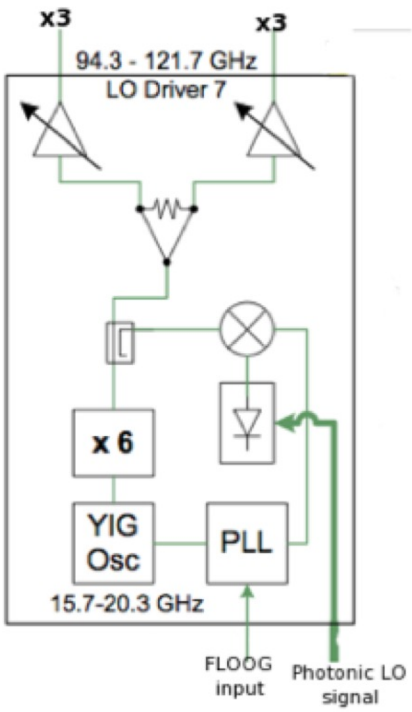

Figure B.6: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier, not shown here, is used to generate the LO1 frequency at 282.9 – 365.1 GHz (in this case, an x3 multiplier). The photonic LO signal (green) feeds a photomixer which creates a beat signal between the ML and SL frequencies. This is mixed with a fraction of the LO from the YIG (x6), and the difference frequency is used in the PLL. The FLOOG also generates a small offset frequency for the PLL, which is different for each antenna. See text for details.

The IF processors divide the incoming IF bands from both sidebands into four 2 GHz basebands and downconvert them to the 2-4 GHz range using the second LO (LO2). Since each baseband is fed by a separate LO2, it is possible to locate them at different frequencies within the IF bandwidth of the receiver (see Chapter 6 and Table 6.2 for limitations). The LO2s are common to both mixer polarizations which means that both polarizations will have the same spectral setups.

The LO2s are digitally-tuned YIG oscillators with a range of 8-14 GHz. LO2 is generated from a harmonic of 125 MHz, plus a fine-tuned synthesiser (fts) of range 20-42.5 MHz, added or subtracted depending on the lock sideband selected by the software. Note that this does not give continuous LO2 coverage, and has to be compensated elsewhere in the LO system.

The IFP unit has 0.5 dB stepped attenuators and Total Power detectors for tuning/optimization of the IF power levels into the digital samplers; these levels are set up at the start of each scan. It is important to note that the switch network layout in the IFP means it is NOT possible to select IF configurations with one baseband in one sideband and three in the other (except for DSB receivers, where this is done using sideband selection). The IFP has anti-alias filters, one set of which is switchable depending on whether the IF range in use is in the upper or lower part of the IF band. As well as downconversion, the LO2s can also be used for sideband separation when combined with the first LO (Section B.1).

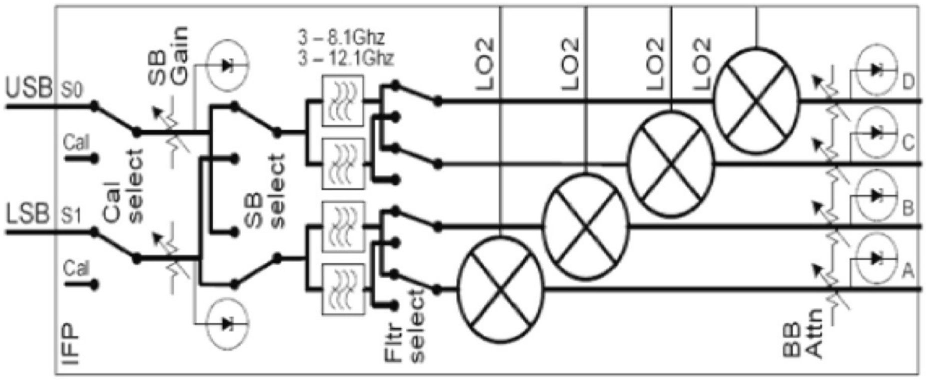

Figure B.7: Block diagram of one polarisation channel of the IF Processor. This has two IF inputs (at left), and feeds 4 baseband (BB) outputs to the digitisers (to the right of the diagram).

The IF processor also has anti-aliasing filters, which define the 2 GHz baseband width and remove out-of-band signals (Section B.3.6). This results in the higher noise levels on the upper and lower 50-100 MHz of channels in the TDM correlator mode (see Section 6.4). These filters cause a decrease in the effective IF range to approximately 1.875 GHz.

B.3.6 Digitization and Signal Transmission

The outputs of the IF Processor units are fed into the Data Transmission System Transmitter (DTX), that include digitizers and formatters to convert the signals to optical wavelengths for transmission via optical fibers. There are four DTX units per antenna, each one handling data for a given baseband pair (i.e., the same 2 GHz baseband from each of the two orthogonal polarizations). Each baseband is digitized by a separate digitizer at 4 GHz (i.e., Nyquist sampling for a 2 GHz bandwidth), quantizing each sample into 3 bits (8 levels) per polarization, so that a total of 6 bits must be transferred per baseband pair. The digitized signal is then transferred to the formatter part that packages the data in frames of equal size. The output of each DTX module is fed to three optical fibers, each transporting 2 bits, and the signal leaves the antenna after passing through a Fiber Optic Multiplexer (FOM). All DTX modules are fed with reference/timing signals from an associated Digital Clock (DGCK), which is also used to do the fine delay tracking.

The outputs of the DTX are sent, via the optical fibers, to the AOS Technical Building where the process is inverted (conversion from optical to digital signal) at the DRXs (Data Transmission System Receiver modules), before the signals are sent to the correlator. Delay corrections due to changes in the length of the optical fibers are done using metadata information to realign the frames sent from the transmitting side at the antenna

(DTX) and the receiving side at the Technical Building (DRX). Figure B.8 shows a block diagram of a single DTX module.

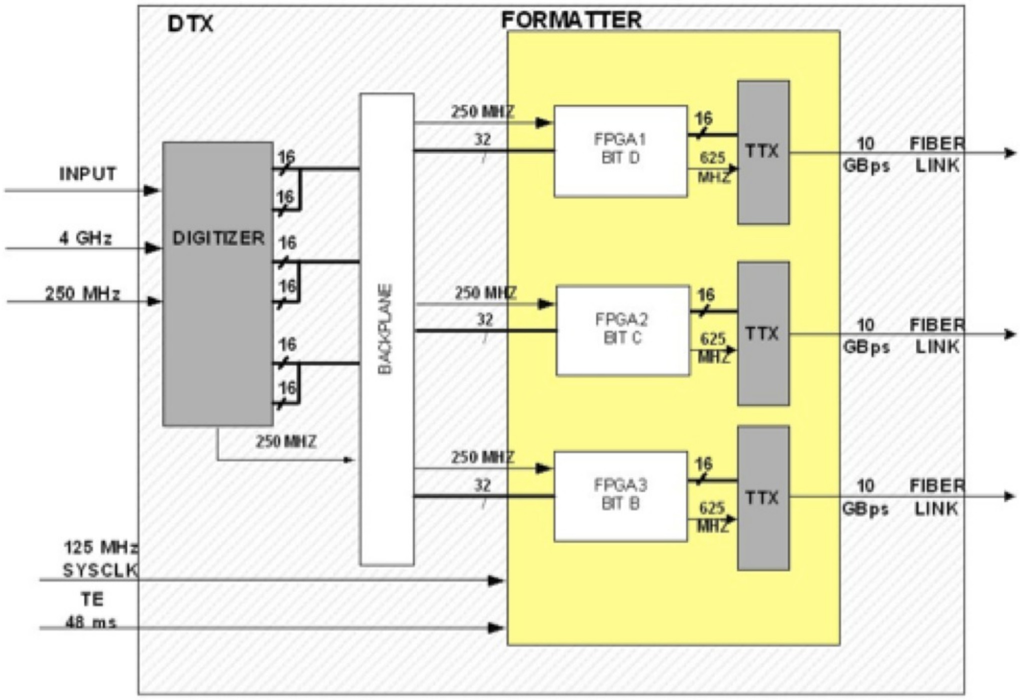

Figure B.8: The Data Transmission System Transmitter (DTX) module, as used in each antenna.

B.4 Other functions of the LO/IF

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

B.4.1 Delay Corrections

ALMA handles delay corrections via the "Delay Server" software package. It computes the corrections for all the different components involved with a cadence of one minute and distributes them buffered. The three main components along the data flow chain where the corrections are applied are: the First LO Offset Generator (FLOOG), the Digital Clock (DGCK) and the correlator (see Figure B.1). Fringe tracking is done at the FLOOG by slightly offsetting the frequency of the LO1 signal. Currently, the delay handled by the FLOOG is in steps of 250 ps. The FLOOG is also used for phase and frequency switching for suppression and separation of sidebands, and for rejection of internally-generated interference, described in the next subsections.

Fine delay corrections are handled by the DGCK that feeds the corrections into the four DTS modules in each antenna. The delay correction resolution of these is 1/16 of the FLOOGs (i.e., 1/16th of 250 ps). The bulk delay correction is handled by the Correlator in integer multiples of the 250 ps units. On top of these corrections, the correlator also handles the "residual" delay corrections at much higher temporal resolution (<250 ps/16) by applying a linear phase gradient across the passband after correlation. Also, the correlator applies relative delay corrections between all the basebands and polarizations of a given ALMA band receiver. Currently, the first baseband of the X polarization is used as reference.

B.4.2 Sideband Suppression - LO Offsetting

Some of the ALMA receivers (e.g. Bands 3, 4, 6, 7, and 8) are inherently single sideband (SSB), either through having a mixer or quasioptic design which rejects the unwanted sideband[7]. Their intrinsic sideband rejection is typically only about 10-15 dB, which, although adequate for rejection of the unwanted sky noise, is not enough to remove strong lines from the other (image) sideband. Others receivers (Bands 9, 10) are double sideband (DSB), and the relative response of the two sidebands may not be equal, significantly affecting calibration. Accordingly, additional schemes are necessary for more effective removal of the unwanted sideband (known as sideband suppression), and for correlation of both sidebands independently (or sideband separation - see next section). Sideband suppression in ALMA is done using the FLOOG, and either LO2 (2LO offsetting) or a combination of LO2 and LO4 (3LO offsetting). A small frequency offset F_o is added to LO1 and subtracted from the other LOs, so that while the signal sideband remains at the same frequency, the image sideband is shifted $2F_o$ away from its nominal value. A different value of F_o is applied at every antenna (the offsets are defined using a Walsh pattern), so that all signal sidebands are at the same frequency, but all image sidebands are at slightly different frequencies and no longer correlate.

Note that each of the basebands has an independent LO2 and LO4. So by setting the sign of the offset in (LO2+LO4) differently, each baseband can be set up to observe in a different sideband.

For single-dish observing, such interferometric sideband rejection methods cannot be used, and a frequency scanning method is under development which will allow image rejection.

B.4.3 Interference Rejection - 180 degree Phase Switching

The FLOOG is additionally used to reject spurious signals prior to digitization by applying 180 deg phase switching according to orthogonal Walsh function patterns, with a pattern cycle time of 16 ms. The Walsh pattern is different on each antenna, and is demodulated by a sign change within the DTS; as a result, the wanted signals correlate, and the unwanted signals are canceled out. This rejects spurious signals generated in the system between the receiver and the sampler, and also suppresses sampler DC offsets. 180-Walsh is a default setup for all observations. For further details please see Section 5.5.4.

B.4.4 Sideband Separation - 90 degree Walsh Switching

It is possible to apply a 90 deg phase switch in the FLOOG and in the correlator processing, allowing correlation of the upper and lower sidebands separately. For Bands 9 and 10 (DSB) it will effectively double the bandwidth from 8 GHz to 16 GHz per polarization from the DSB receivers. For further details please see Section 5.5.4.

[7]Although they have mixers to allow both sidebands to be observed separately and simultaneously

Appendix C

Acronym Dictionary

2SB	Dual Sideband
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
APP	ALMA Phasing Project
APS	ALMA Phasing System
AQUA	ALMA Quality Assurance software
ARC	ALMA Regional Center
ASA	ALMA Science Archive
ASC	ALMA Sensitivity Calculator
ASDM	ALMA Science Data Model
ASIC	Application Specific Integrated Circuit
AZ	Azimuth
BB	Baseband
BDF	Binary Data Format
BE	Backend
BL	Baseline
BLC	BaseLine Correlator
BWFN	Beam Width between First Nulls
BWSW	Bandwidth Switching
CASA	Common Astronomy Software Applications package
CCA	Cold Cartridge Assemblies
CCC	Correlator Control Computer
CDP	Correlator Data Processor
CFRP	Carbon Fiber Reinforced Plastic
CLO	Central Local Oscillator
CLT	Chilean Local Time
CORBA	Common Object Request Broker Architecture
CRD	CentralReference Distributor
CRG	Central Reference Generator
CSV	Commissioning and Science Verification
CW	Continuous Wave
DC	Direct Current
DEC	Declination
DFP	DTS-Rx and FFT Processor
DGC	Differential Gain Calibrator

DGCK	Digital Clock
DMG	Data Management Group within DSO
DRX	Data Receiver module
DSB	Double Sideband
DSO	Division of Science Operations
DTS	Data Transmission System
DTS-Rx	Data Transmission System Receiver
DTX	Data Transmitter module
EB	Execution Block
EHT	Event Horizon Telescope
EL	Elevation
EM	Electromagnetic
EPO	Education and Public Outreach
ES	Early Science
ESO	European Southern Observatory
FDM	Frequency Division Mode
FE	Frontend
FITS	Flexible Image Transport System
FLOOG	First LO Offset Generator
FOM	Fiber Optic Multiplexer
FOV	Field of View
FPGA	Field-Programmable Gate Array
FT	Fourier Transform
FWHM	Full Width Half Maximum
FWHP	Full Width to Half Power
FWBN	Full Width Between the Nulls
FXF	Filtering, Correlation, and Fourier transform type correlator
GMVA	Global Millimiter VLBI Array
GPS	Global Positioning System
HA	Hour Angle
HEMT	High Electron Mobility Transistor
HPBW	Half Power Beam Width
IF	Intermediate Frequency
IFP	Intermediate Frequency Processor
IRAM	Institut de Radioastronomie Millimetrique
JCMT	James Clerke Maxwell Telescope
JPL	Jet Propulsion Laboratory
LAS	Largest Angular Resolution
LFRD	Low Frequency Reference Distributor
LLC	Line Length Corrector
LO	Local Oscillator
LO1	First LO
LO2	Second LO
LO3	Digitizer Clock Third LO
LO4	Tunable Filterbank LO
LORR	LO Reference Receiver
LS	Laser Synthesizer
LSB	Lower Sideband
LSRK	Local Kinematics Standard of Rest
MD	Mixer De-tuned
LTA	Long Term Accumulator
MCI	Monitor and Control Interfce
MFS	Master Frequency Standard

ML	Master Laser
MLD	Master Laser Distributor
MRS	Maximum Recoverable Scale
MS	Measurement Set
NGAS	New Generation Archive System
NRAO	National Radio Astronomy Observatory
NVSS	NRAO VLA Sky Survey
OMC	Operator Monitoring and Control
OMT	Ortho-mode Transducer
OSF	Operations Support Facility
OST	Observation Support Tool
OT	Observing Tool
OTF	On the Fly
OUS	Observing Unit Set
PBS	Polarization Beam Splitter
PDM	Propagation Delay Measure
PI	Principal Investigator
PLL	Phase Lock Loop
PMG	Program Management Group within DSO
PRD	Photonic Reference Distribution
PSF	Point Spread Function
PWV	Precipitable Water Vapor
QA	Quality Assurance
QA0	Quality Assurance Level 0
QA1	Quality Assurance Level 1
QA2	Quality Assurance Level 2
QA3	Quality Assurance Level 3
QL	QuickLook pipeline
RA	Right Ascension
RF	Radio Frequency
RMS	Root Mean Square
SAS	Sub Array Switch
SB	Scheduling Block
SCO	Santiago Central Office
SD	Single Dish
SED	Spectral Energy Distribution
SIS	Superconductor-Insulator-Superconductor Mixer
SL	Slave Laser
SNR	Signal-to-Noise Ratio
SPW	SPectral Window
SRON	Stichting Ruimte Onderzoek Nederland (Netherlands Institute for Space Research)
SSB	Single Sideband
SSR	Science Software Requirements
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable Filterbanks
TFB LO	Local Oscillator at the Tunable Filterbanks
TMCDB	Telescope Monitor and Configuration DataBase
TP	Total Power
Tsys	System Temperature

$T_{\rm rx}$	Receiver Temperature
USB	Upper Sideband
VLA	Very Large Array
VOM	VLBI Observing Mode
VO	Virtual Observatory
WCA	Warm Cartridge Assembly
WVR	Water Vapor Radiometer
XF	Correlation-Fourier Transform Type Correlator
YIG	Yttrium-Iron Garnet Oscillator

Index

The Atacama Large Millimeter/submillimeter Array (ALMA), an international astronomy facility, is a partnership of the European Organization for Astronomical Research in the Southern Hemisphere (ESO), the U.S. National Science Foundation (NSF) and the National Institutes of Natural Sciences (NINS) of Japan in cooperation with the Republic of Chile. ALMA is funded by ESO on behalf of its Member States, by NSF in cooperation with the National Research Council of Canada (NRC) and the National Science Council of Taiwan (NSC) and by NINS in cooperation with the Academia Sinica (AS) in Taiwan and the Korea Astronomy and Space Science Institute (KASI).

ALMA construction and operations are led by ESO on behalf of its Member States; by the National Radio Astronomy Observatory (NRAO), managed by Associated Universities, Inc. (AUI), on behalf of North America; and by the National Astronomical Observatory of Japan (NAOJ) on behalf of East Asia. The Joint ALMA Observatory (JAO) provides the unified leadership and management of the construction, commissioning and operation of ALMA.