Doc 2.3, ver. 1.1 | October 15, 2013

ALMA Cycle 2 Technical Handbook

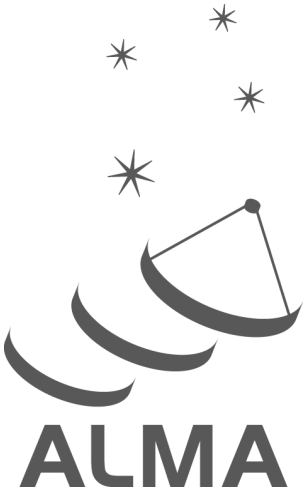

www.almascience.org

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

User Support:

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

Revision History:

Version	Date	Editors
0.1	2013-02-07	A. Lundgren
0.2	2013-04-04	A. Lundgren
0.3	2013-06-06	A. Lundgren
0.4	2013-07-18	A. Lundgren, M. Saito, A. Hales, S. Asayama, T. Kamazaki, J. Di Francesco, S. Leon, S. Komugi, H. Nagai, A. Biggs, B. Vila Vilaro, E. Fomalont, A. Richards
0.5	2013-08-12	A. Lundgren, E. Fomalont, B. Dent, J. Di Francesco, F. Stoehr, A. Hales, B. Vila Vilaro, S. Leon, A. Peck
0.6	2013-09-10	A. Lundgren, S. Komugi, E. Fomalont, S. Asayama, J. Di Francesco, S. Leon, F. Stoehr, M. Saito, A. Biggs, B. Dent, B. Vila Vilaro
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
1.1	2013-10-15	A. Lundgren, S. Leon, J. Hibbard, E. Müller, J. Di Francesco

Contributors

Eiji Akiyama, Paola Andreani, Yiping Ao, Denis Barkats, Andy Biggs, Crystal Brogan, James Chibueze, Stuartt Corder, Carlos De Breuck, Itziar De Gregorio, Bill Dent, James Di Francesco, Darrel Emerson, Ed Fomalont, Roberto Galvan-Madrid, Antonio Hales, Bunyo Hatsukade, Cinthya Herrera, John Hibbard, Todd Hunter, Akiko Kawamura, Takeshi Kamazaki, Eelco van Kampen, Rüdiger Kneissl, Shinya Komugi, Yasutaka Kurono, Mark Lacy, Stephane Leon, Harvey Liszt, Carol Lonsdale, Andreas Lundgren, Geoff Mathews, Jeff Mangum, Nuria Marcelino, Ivan Marti-Vidal, Rie Miura, Koh-Ichiro Morita, Erik Müller, Hiroshi Nagai, Lars-Åke Nyman, Sachiko Okumura, Rosita Paladino, Alison Peck, Dirk Petry, David Rabanus, Suzanna Randall, Mark Rawlings, Anita Richards, Masao Saito, Markus Schmalzl, Scott Schnee, Kimberly Scott, Kartik Sheth, Richard Simon, Hiroko Shinnaga, Tomas Stanke, Felix Stoehr, Kenichi Tatematsu, Ignacio Toledo, Baltasar Vila-Vilaro, Al Wootten, Martin Zwaan.

In publications, please refer to this document as:
A. Lundgren, 2013, ALMA Cycle 2 Technical Handbook Version 1.1, ALMA
ISBN 978-3-923524-66-2

Contents

List of Figures

This page was intentionally left almost blank

Chapter 1

Introduction

The Atacama Large Millimeter/Submillimeter Array (ALMA) is an aperture synthesis telescope that will consist of 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 2 is still within the ALMA Early Science period, and a reduced number of antennas, frequency bands, array configurations and observing modes will be available. Users should refer to the Capabilities section on the ALMA Science Portal at http://www.almascience.org/ for the latest information.

This Technical Handbook provides additional information on technical aspects and its limitation of the Cycle 2 setup for ALMA users, to a deeper level than what is described in the ALMA Early Science Primer.

The handbook is divided in three main sections: The hardware and correlator setup (chapters 2–5), the concepts and software (chapters 6–9), and finally the data quality and handling (chapters 10–13). Chapter 2 describes the ALMA array components: The 12-m Array and the ACA (Atacama Compact Array), which also is known as the Morita array. Chapter 3 describes the details of the 6 receivers offered in Cycle 2. Chapter 4 describes the correlators and the data processing taking place in these special purpose super computers. Chapter 5 describes how the spectral setup is done in the correlators. Chapter 6 gives a brief introduction to interferometry, while Chapter 7 applies these concepts to derive ALMA imaging considerations. Chapter 8 describes how projects are observed and Chapter 9 gives a brief overview how the sensitivities and integration times are calculated. Chapter 10 describes the data calibration criteria, and Chapter 11 describes the data quality assurance process. Chapter 12 describes the structure of the data, while Chapter 13 describes how the data is stored. The handbook ends with the Appendix, which contains supplemental material about concepts and hardware, such as the antenna, transporter, LO system, and calibrator source selection algorithm.

Chapter 2

Array Components

2.1 The ALMA Telescope

Upon completion, ALMA will be composed of 66 high-precision antennas. Fifty of these antennas will be 12 meter dishes in the 12-m Array, used for sensitive, high-resolution imaging. These will be complemented by the Atacama Compact array (ACA), composed of twelve closely spaced 7 meter antennas (7-m Array), and four 12 meter antennas for single-dish (or Total Power) observations (TP Array), to enhance wide-field imaging of extended structures. In Cycle 2, ALMA will cover most of the frequency range 84 to 720 GHz (from 0.42 to 3.6 mm), and in Full Operation, the frequencey coverage will be 31–950 GHz (from 0.32 to 10 mm).

The array is located on the Chajnantor plain of the Chilean Andes, a site that normally offers the exceptionally dry and clear sky conditions required to observe at millimeter and sub-millimeter wavelengths. 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 40 km from the town of San Pedro de Atacama at an altitude of 2900 meters. Science operations are conducted from the OSF and coordinated from the JAO Central office in Santiago. All three ALMA arrays are controlled via control software developed on the ALMA common software (ACS).

In the following sections we describe the main characteristics of each ALMA array. Unless otherwise noted, the description is appropriate for the fully completed ALMA. Not all capabilities may be available during the Early Science observing period. 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 2) will be published in the Capabilities section of the Proposers Guide. Complementary background information on ALMA and its capabilities for Early Science can be found in the document "A Primer for Early Science" which can be found on the link http://www.almascience.org/documents-and-tools/.

2.2 The 12-m Array

The 12m array consists of fifty 12 meter diameter antennas, which can be located on 170 foundations (pads) distributed over the Chajnantor (and Pampa la Bola) plateau. 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).

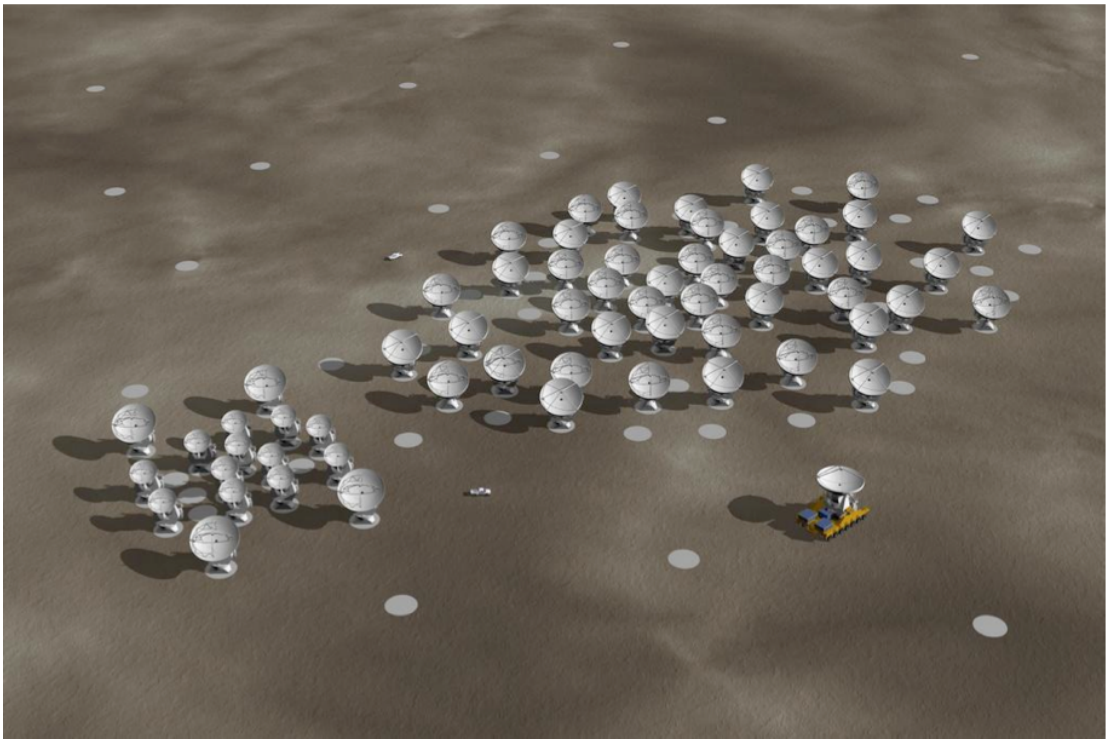

Figure 2.1: Artist's conception of the ALMA 12-m Array in its compact configuration, with the ACA toward the left, and an antenna transporter in the lower right. A few unoccupied pads can be seen, to which antennas of the 12-m Array can be moved by the transporter as the array is being reconfigured. At its most extended configuration, antennas in the 12-m Array will be about 16 km apart.

The antennas for the 12-m Array were designed and built in both Europe and North America (each providing 25 units), according to the stringent ALMA Antenna Performance specifications (see Appendix A.1). 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 above each 12-m Array element. The cryostat can contain up to ten cartridges (only 7 in ALMA construction), each covering one frequency band (see Chapter 3). The overall frequency range for which ALMA was designed is 31 GHz to 950 GHz. Only one band observes at any time, but up to three can be switched on simultaneously, and rapid switching between those is possible. Each cartridge (Chapter 3) receives two orthogonal polarizations and down-converts the signals to an intermediate frequency with eight GHz of bandwidth per polarization. Some cartridges are dual sideband (2SB; bands 3-8) and others are double sideband (DSB; bands 9 and 10).

The Local Oscillator signals (Section 3.1) are transmitted to the antennas on optical fibres with a round-trip measurement to correct for changes in the fiber length. There are four independent LO reference systems so that the 7-m Array, the Total Power array and two sub-sections of the 12-m Array (e.g., two 'sub-arrays') can make simultaneous independent observations. The 8-10 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 (A.7). Total data rates are therefore 96 Gbits/s per antenna. With formatting, the bit rate is 120 Gb/s.

On arrival at the central building, the data are recovered and processed in one of the two correlators: the 64-input Correlator and the ACA Correlator (see Chapter 4). All antennas can feed both correlators. The 64-input Correlator (Section 4.2) is normally used for the 12-m Array, but it can also take inputs from the 7-m Array or TP antennas. It is an XF correlator (cross-correlates first, then Fourier transforms), but the correlator proper is preceded by Tunable Filter Banks, which makes it a digital hybrid XF correlator or FXF. These can select sub-bands from the 2 GHz-wide basebands in a very flexible manner. From each of these, the correlator then generates 2016 cross-correlations and 64 auto-correlations (requiring 1.7×10^{16} operations per second). Either 2-, 3- or 4-bit resolution is used, and the sampling can be Nyquist or twice Nyquist. See Table 4.1 for an example table of modes. The correlated data are fed to a group of processors which do the transforms and carry out integration and data compression.

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

The 12-m Array configurations have been designed so that in Full Operations 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[1]. 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 "Largest Angular Structure" for that array - see Section 7.4), ALMA has been designed to include the ACA. In Full Operations, the ACA will be 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

[1]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.

auto-correlations. The 7-m Array samples baselines from 8.5 m to 33 m, bridging the baseline sampling gap between the 12-m Array and the TP Array.

The ACA is controlled via control software developed on the ALMA common software (ACS) and is operated in a similar fashion to the 12-m Array. To achieve this unified operation, the ACA system is as compatible with the 12-m Array as possible at the level of hardware, interface, data, and observing modes. In Full Operations, the standard observing modes for the TP Array will include spectral line and continuum observations with raster or Lissajous on-the-fly (OTF) scans, or position switching. The raw time-series signals from the ACA antennas are processed in the ACA Correlator (see Section 4.3) 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 meter diameter antennas designed and built by East Asia to the ALMA specifications (A.1). 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 ∼0.6 GB/s per baseband.

No baseline coverage from even the most compact configuration of the 12-m Array is obtained for spacings smaller than 15 m. The array configurations of the 7-m Array are designed to fill missing spacings in from about 9 m to ∼30 m (7.6). The specific 7-m Array configurations available for Cycle 2 will be published in the Proposers Guide "Capabilities".

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.1). The specification of the TP antennas is almost identical to the ones for the 12m main/interferometric array. The TP antennas are located on pads surrounding the 7 m array. The TP antennas will be fitted with a nutating subreflector (aka wobbler or chopper). Those components will allow corrections for rapid atmospheric fluctuations that would otherwise limit continuum observations. The nutator provides the capability to switch rapidly the optical axis of an antenna's primary beam between a target source and nearby reference positions. If the switching is sufficiently faster than significant variations in the atmospheric emission noise, the atmospheric brightness is the same for both the target and reference sources so the differential signal measures the flux for the source being observed.

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. Section 8.6 describes the observing mode and capabilities offered for the TP Array in Cycle 2.

Due to the small 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. Since the 7-m Array is quite compact, atmospheric phase fluctuations will be smaller than for the 12-m Array. The WVRs installed on the TP Array antennas that surround the 7-m Array will be used for these corrections.

Chapter 3

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 3.1 and the specifications are listed in Table 3.1. In Cycle 2, Band 3, 4, 6, 7, 8 and 9 will be available (see available frequency and wavelength ranges for these bands in Table 3.1). The receivers are described in more detail in the following sections as well as in the references listed in Table 3.2.

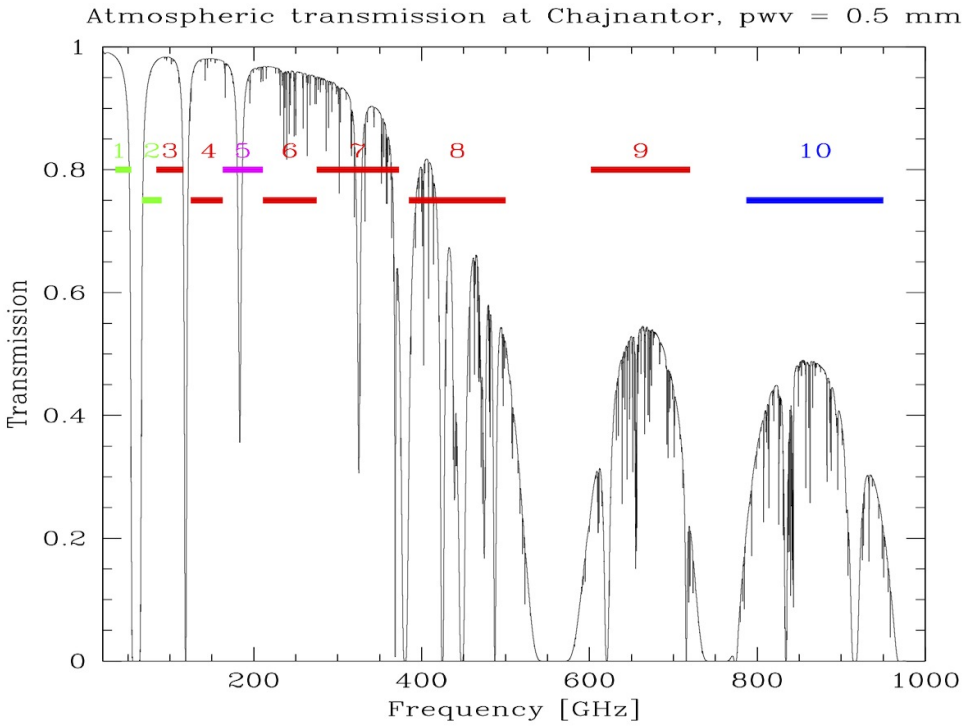

Figure 3.1: ALMA 10 receiver bands. Receiver bands for Cycle 2 are shown in red superimposed on a zenith atmospheric transparency plot at the AOS for 0.5 mm of PWV.

The ALMA receivers in each antenna are situated in a single front-end assembly (see Appendix, Section A.4). The front-end assembly consists of a large cryostat containing the receiver cold cartridge assemblies (including SIS mixers and LO injections) and the IF and LO room-temperature electronics of each band (the warm cartridge assembly, WCA). The cryostat is kept at a temperature of 4 K through a closed-cycle cooling system. The Amplitude Calibration Device (ACD) is mounted above the front end (see Appendix A.5). Each receiver

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

Table 3.1: Receiver Characteristics. *Notes to Table:* **1.** Frequency range is the maximum available, at the extreme upper and lower limits of the IF passband. For FDM mode, the coverage is a bit smaller. **2.** Sideband modes: SSB means single sideband receiver, 2SB means dual sideband receiver where the two sidebands are available simultaneously, DSB means double sideband receiver. See text for details. **3.** Usable IF range is extended to allow simultaneous observations of multiple lines. However, the autocorrelation noise performance is degraded by a factor of up to about 1.5 below 5.5 GHz (Section 3.2.3) **4.** Maximum instantaneous IF bandwidth: As both upper and lower sidebands both pass through the same IF bandwidth but are subsequently separated, the effective signal bandwidth given in this column for 2SB receivers is twice the actual IF filter bandwidth. In addition, this is per polarization, so the total effective bandwidth for each receiver is then another factor of 2 higher. Note that the effects of the anti-aliasing filters have been included (see Section 5.5). **5.** In future Cycles, the maximum bandwidth will double in cross-correlation mode, because both sidebands can be separated and correlated using 90-degree phase switching. **6.** List of the minimum specification of the SSB receiver temperature ($T_{\rm rx}$), unless otherwise noted, is shown. These values are the average over the IF band. The sections on individual receiver bands describe the real values measured, which in many cases are better than specifications. **7.** The specification for Band 3 receivers is $T_{\rm rx}$<41 K at LO=104 GHz, and $T_{\rm rx}$<45K for any other valid LO setting. Both values should be the average over all four IFs and 4 GHz bandwidth.

Topic	Author/Year	Meeting proceedings	ADS identifier
B3	Claude et al. 2008	SPIE 7020	2008SPIE.7020E..33C
B4	Asayama et al. 2008	19th Intl Symp Space Terahertz Tech	2008stt..conf..244A
B6	Kerr et al. 2004	15th Intl Symp Space Terahertz Tech	2004stt..conf...55K
B7	Mahieu et al. 2005	16th Intl Symp Space Terahertz Tech	2005stt..conf...99M
B8	Sekimoto et al. 2008	19th Intl Symp Space Terahertz Tech	2008stt..conf..253S
B9	Baryshev et al. 2007	18th Intl Symp Space Terahertz Tech	2007stt..conf..164B
Optics	Rudolf et al. 2007	IEEE Trans. on Antennas & Propagation	2007ITAP...55.2966R
WVR	Emrich et al. 2009	20th Intl Symp Space Terahertz Tech	2009stt..conf..174E

Table 3.2: Conference proceedings papers describing the receiver bands, optics and the water vapor radiometer.

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

To avoid overloading the cryostat cooler, only three bands can be switched on at a time. From a hardware point of view, it takes only about 1.5 seconds to switch between these bands, but in reality, switching between phase calibrator and science source currently takes, depending on situation, from 30 seconds to a few minutes (but future improvements is expected to decrease this overhead). For bands that are not switched on, the time to fully thermally-stabilize them from an off state is up to 60 minutes - this is mainly to ensure the optimum flat bandpass shape. All of the receivers are mounted off-axis to avoid extra rotating band-selection mirrors, which necessitates an offset of the antenna to change band. This means that only one receiver can be used at a given time.

3.1 Local Oscillators and IF ranges

The observed sky-frequencies need to be down-converted to frequency bands between 0-2 GHz to send the signals to the correlator. The frequency down-conversion involves a set of Local Oscillators (LOs). The LO and IF systems are described in detail in the Appendix (A.7).

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, 7, and 8 have an output range of 4–8 GHz, Band 6 a range of 6–10 GHz and Band 9 a range of 4–12 GHz (Table 3.1). The possible sky frequency ranges covered by each receiver with the first Local Oscillator (LO1) set to a frequency F_{LO1} are:

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

where IF_{lo} and IF_{hi} are the lower and upper IF ranges in the "IF Range" column of Table 3.1, and the IF bandwidth (per sideband) is IF_{hi}–IF_{lo}. This is illustrated in Figure 3.2. Note that the maximum IF bandwidth in Table 3.1 may be a few percent less than the IF range in Table 3.1 (see Section 5.5).

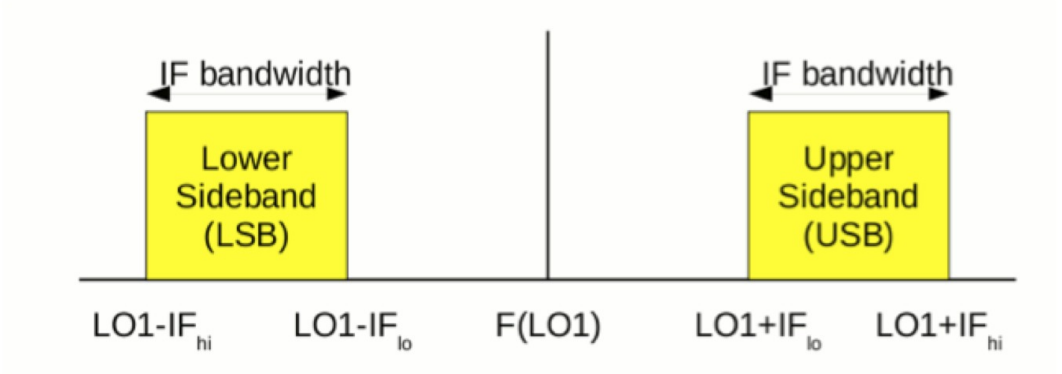

Figure 3.2: IF ranges for the two sidebands in a heterodyne receiver.

3.2 The Cycle 2 Receivers

The Band 3, 4, 6, 7, and 8 receivers are dual-sideband (2SB) receivers, where both the upper and lower sidebands are provided separately and simultaneously. There are 4 outputs from each of the receivers, comprising the upper

and lower sidebands in each of the two polarizations. Each output has a bandwidth of 4 GHz (reduced to an effective total bandwidth of 3.75 GHz due to the anti-aliasing filters, etc., see Section 5.4). The mixers give 10 dB or more unwanted sideband rejection, which is adequate for reducing the degradation of S/N from noise in the unwanted sideband, but not adequate for suppressing astronomical signals in the unwanted sideband. Further suppression is performed by offsetting LO1 and LO2 (and eventually the tunable filter LO, TFB LO) by small and opposite amounts, which depend on the antenna, such that the signals from two antennas in the image sideband do not correlate.

The Band 9 receivers are double-sideband (DSB) receivers, where the IF contains noise and signals from both sidebands. They only have two outputs, one per polarization. However, the IF effective bandwidth is 7.5 GHz per sideband (after passing through the IF processing units), so the total instantaneous bandwidth is the same as Bands 3, 4, 6, 7 and 8. In Cycle 2, only one sideband per spectral window will be correlated, and the other rejected using LO offsetting, as mentioned above. This does not remove the noise from the rejected sideband. The noise of the sideband that is kept will be twice that of the DSB noise level. In the future, suitable phase switching will be introduced in the correlator, and both sidebands can be correlated and processed independently, thus doubling the effective system bandwidth.

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

3.2.1 Band 3 receiver

Band 3 is the lowest frequency band available in Cycle 2, 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 3.3). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the SIS mixers.

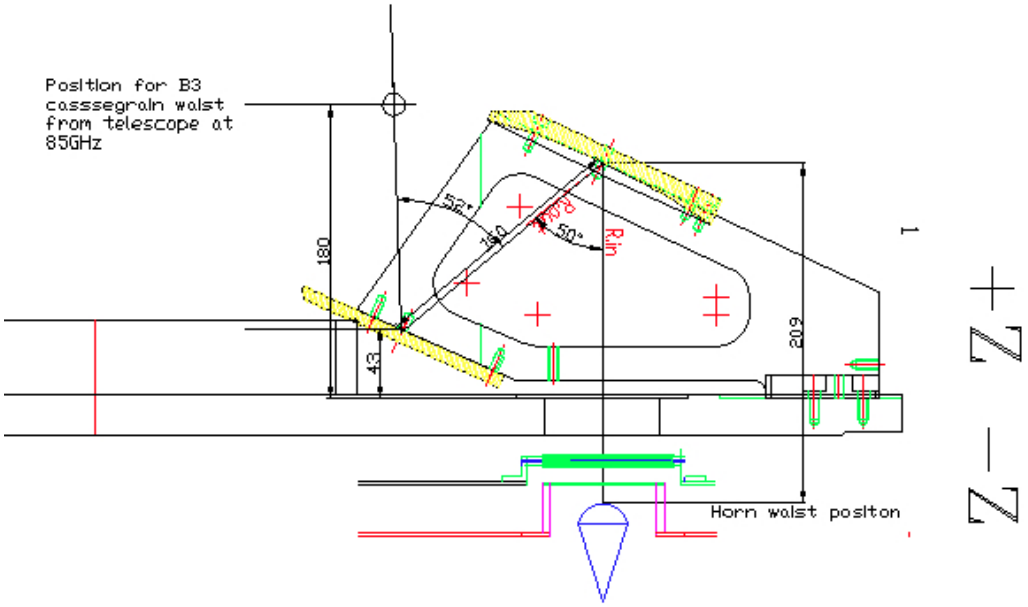

Figure 3.3: Input optics for Band 3, showing the warm pickoff mirrors. The location of the antenna beam from the secondary mirror is shown by the solid line, and the Cassegrain focus is shown by the small circle to the upper right.

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

in Figure 3.4. The Cold Cartridge Assembly (CCA) contains the cold optics, OMT, SIS mixers and the low-noise HEMT first IF amplifiers. At room temperature, the Warm Cartridge Assembly (WCA) includes further IF amplification and the Local Oscillator covering 92–108 GHz.

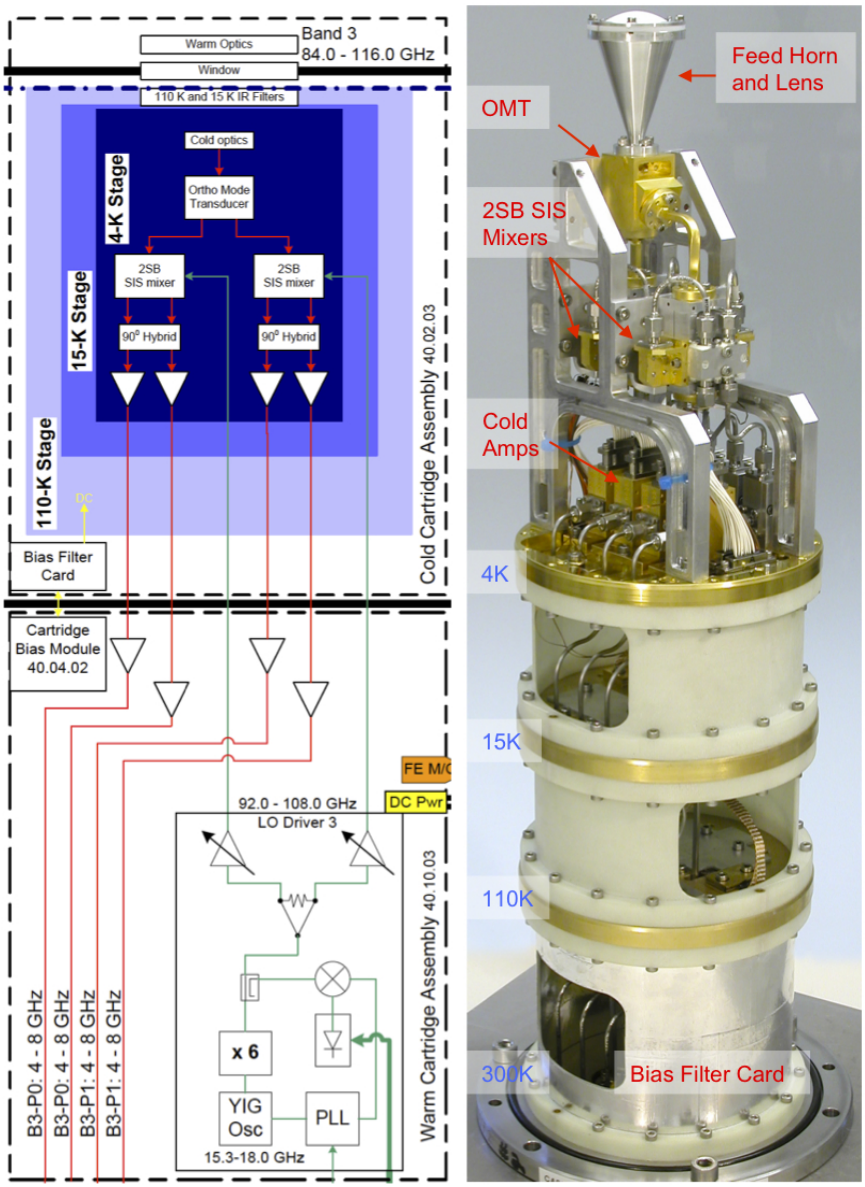

Figure 3.4: Block diagram of the Band 3 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 3 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB mixers. The Band 3 cartridges were constructed in Canada at NRC-HIA, Victoria.

The specification for the Band 3 receiver noise performance ($T_{\rm rx}$) is <41 K at LO=104 GHz, and <45 K for any other valid LO setting.The atmospheric transmission over most of Band 3 is very high, even with a large PWV (Figure 3.5) which means observations in Band 3 can, in principle, take place with 10 mm or more of PWV. The resulting system temperature ($T_{\rm sys}$) shows the expected rise at the higher end, due to an atmospheric oxygen line (Figure 3.6).

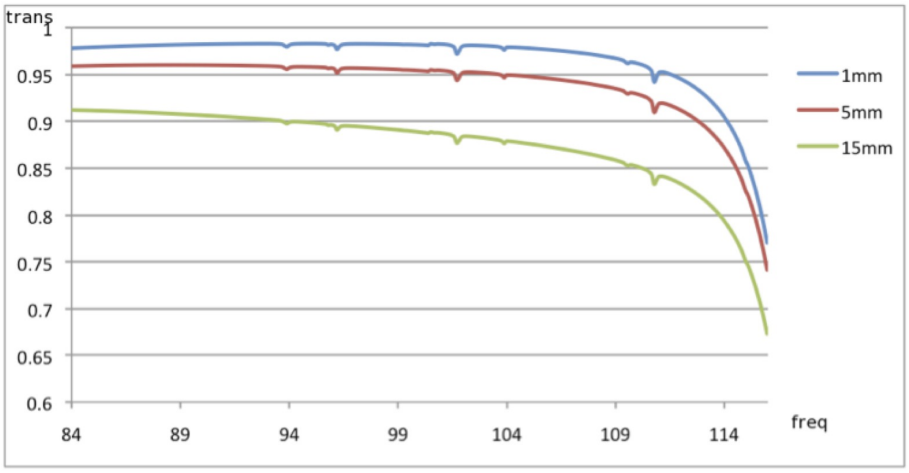

Figure 3.5: Band 3 zenith transmission for 1, 5 and 15 mm of PWV. Frequency is in GHz.

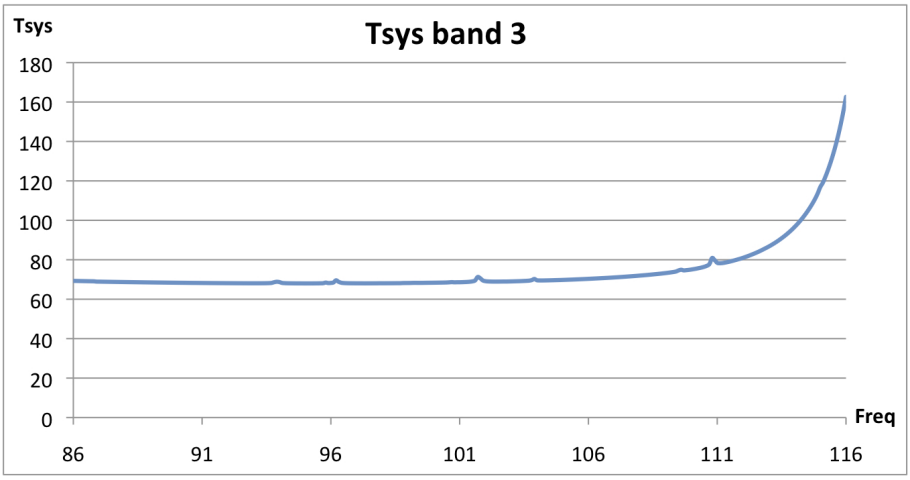

Figure 3.6: Typical system temperature ($T_{\rm sys}$) at zenith for Band 3 with 0.9 mm of PWV. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

3.2.2 Band 4 receiver

The Band 4 receiver covers the 125 to163 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 3.7). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the 2SB SIS mixers.

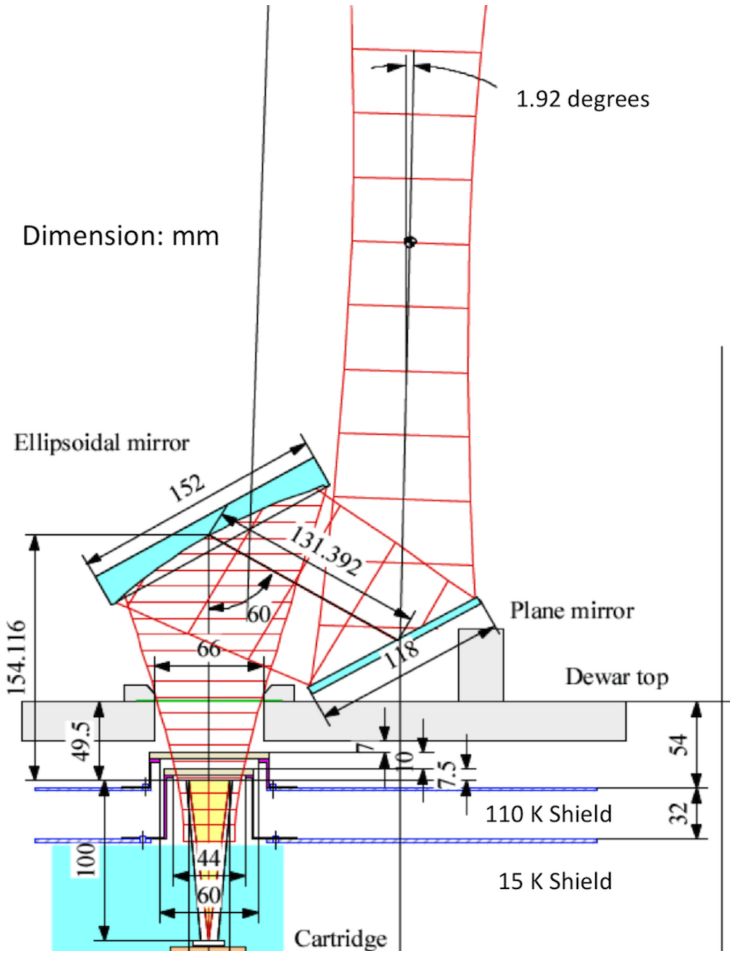

Figure 3.7: Optical layout of the Band 4. Red indicates Band 4 Gaussian beam of 5 w.

A block diagram of the Band 4 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 3.8. The Band 4 CCA contains a feed horn, an OMT as a polarization splitter, 2SB SIS mixer assemblies, cold IF amplifiers, isolators, and LO frequency doublers. The RF signal is down converted to 4–8 GHz using a 2SB mixer unit.

The atmospheric transmission in Band 4 is shown in Figure 3.9 for three typical PWV values. Most observations in Band 4 will be done with PWV <5 mm. The specification for Band 4 receiver noise performance ($T_{\rm rx}$) is <51 K over 80% of the band, and <82 K over the whole band (SSB $T_{\rm rx}$). However, the performance of the receiver is considerably better than 50 K over the band. The resulting system temperatures ($T_{\rm sys}$) for 1 mm PWV are shown in Figure 3.10.

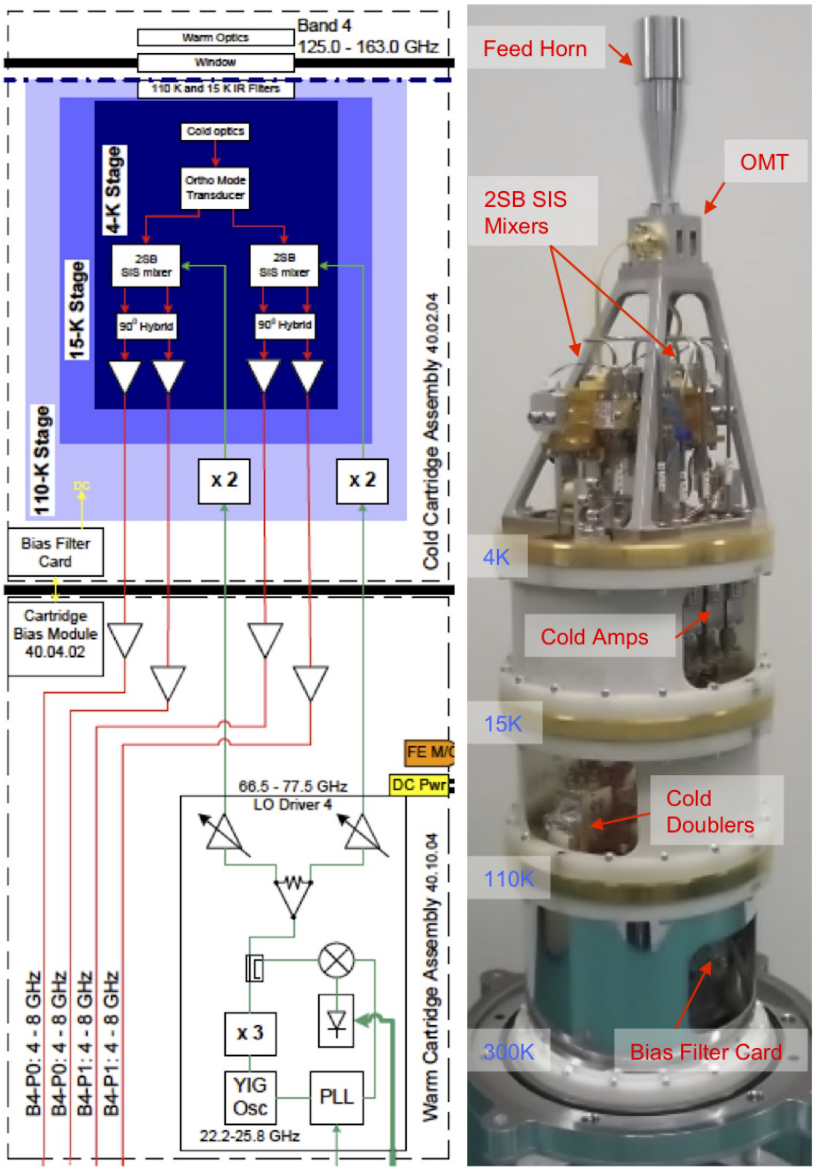

Figure 3.8: Block diagram of the Band 4 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 4 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 4 cartridges are constructed at the Advanced Technology Center (ATC).

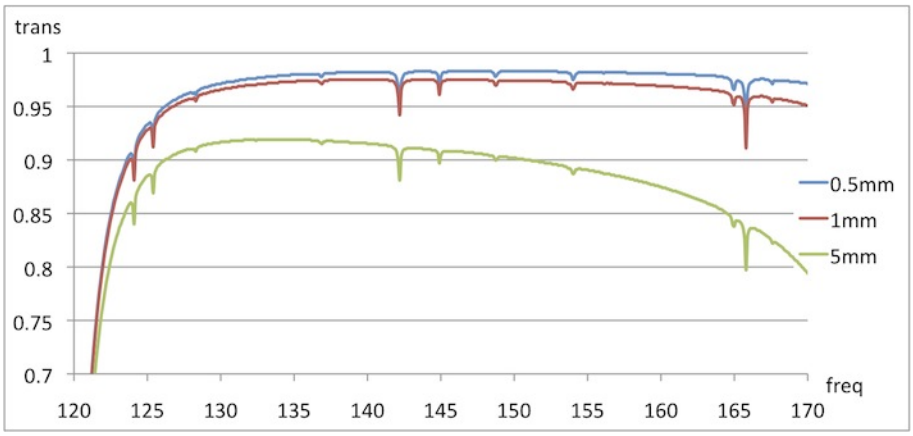

Figure 3.9: Band 4 zenith transmission for 0.5, 1 and 5 mm of PWV. Frequency is in GHz.

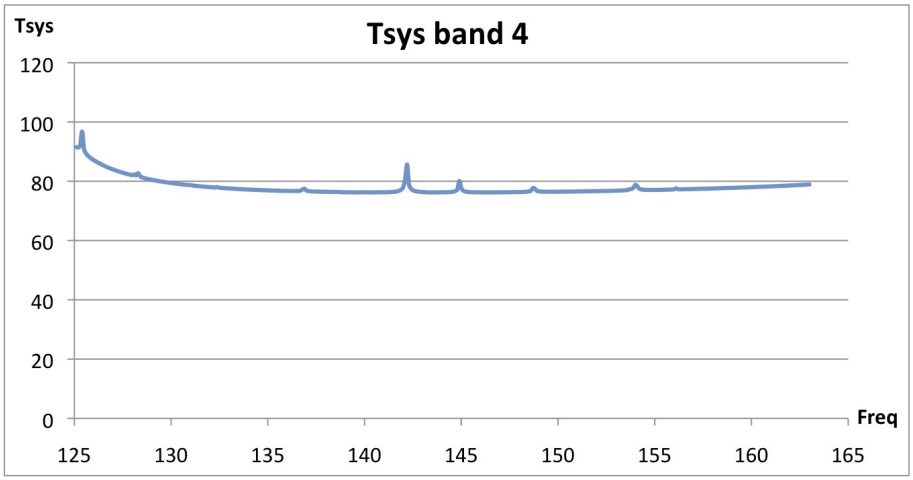

Figure 3.10: Typical system temperature ($T_{\rm sys}$) at zenith for Band 4 with 0.9 mm of PWV. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

3.2.3 Band 6 receiver

The Band 6 receiver covers a frequency range of 211.0–275.0 GHz (the 1.3 mm atmospheric window). This receiver has a window with a pair of off-axis ellipsoidal mirrors inside the cryostat (Figure 3.11). A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the SIS mixers. A block diagram of the Band 6 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 3.12.

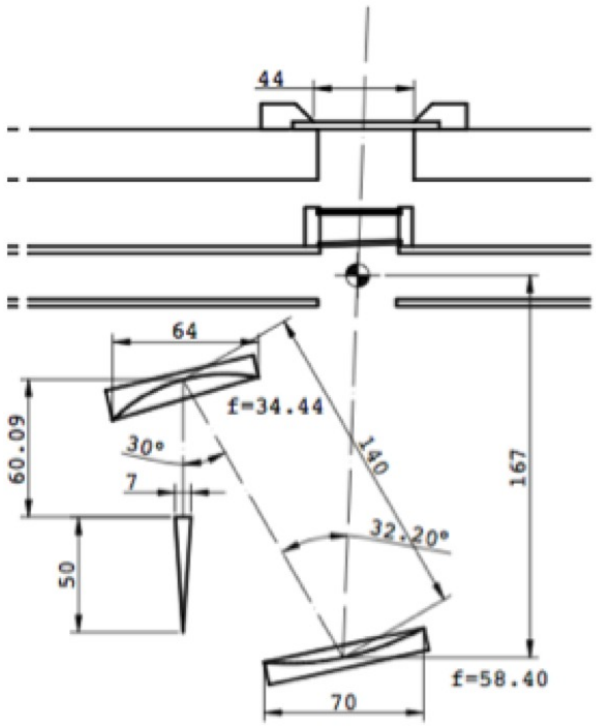

Figure 3.11: Band 6 cold off-axis ellipsoidal mirrors feeding the single feedhorn. The off-axis beam from the telescope secondary mirror (shown by the dashed line) feeds directly through the cryostat window, and the Cassegrain focus is just inside the inner infrared blocker. Note the slightly inclined inner window, designed to minimize standing waves.

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

The atmospheric transmission in Band 6 is shown in Figure 3.13 for three typical PWV values. Most of the narrow absorption lines are from ozone. Most observations in Band 6 will be done with PWV <5 mm.

The specification for Band 6 receiver noise performance ($T_{\rm rx}$) is <83 K over 80% of the band, and <138 K over the whole band (SSB $T_{\rm rx}$). The measured results are considerably better, typically 50 K over most of the band. The resulting system temperatures ($T_{\rm sys}$) for 1 mm PWV are shown in Figure 3.14.

[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 to cover all three lines

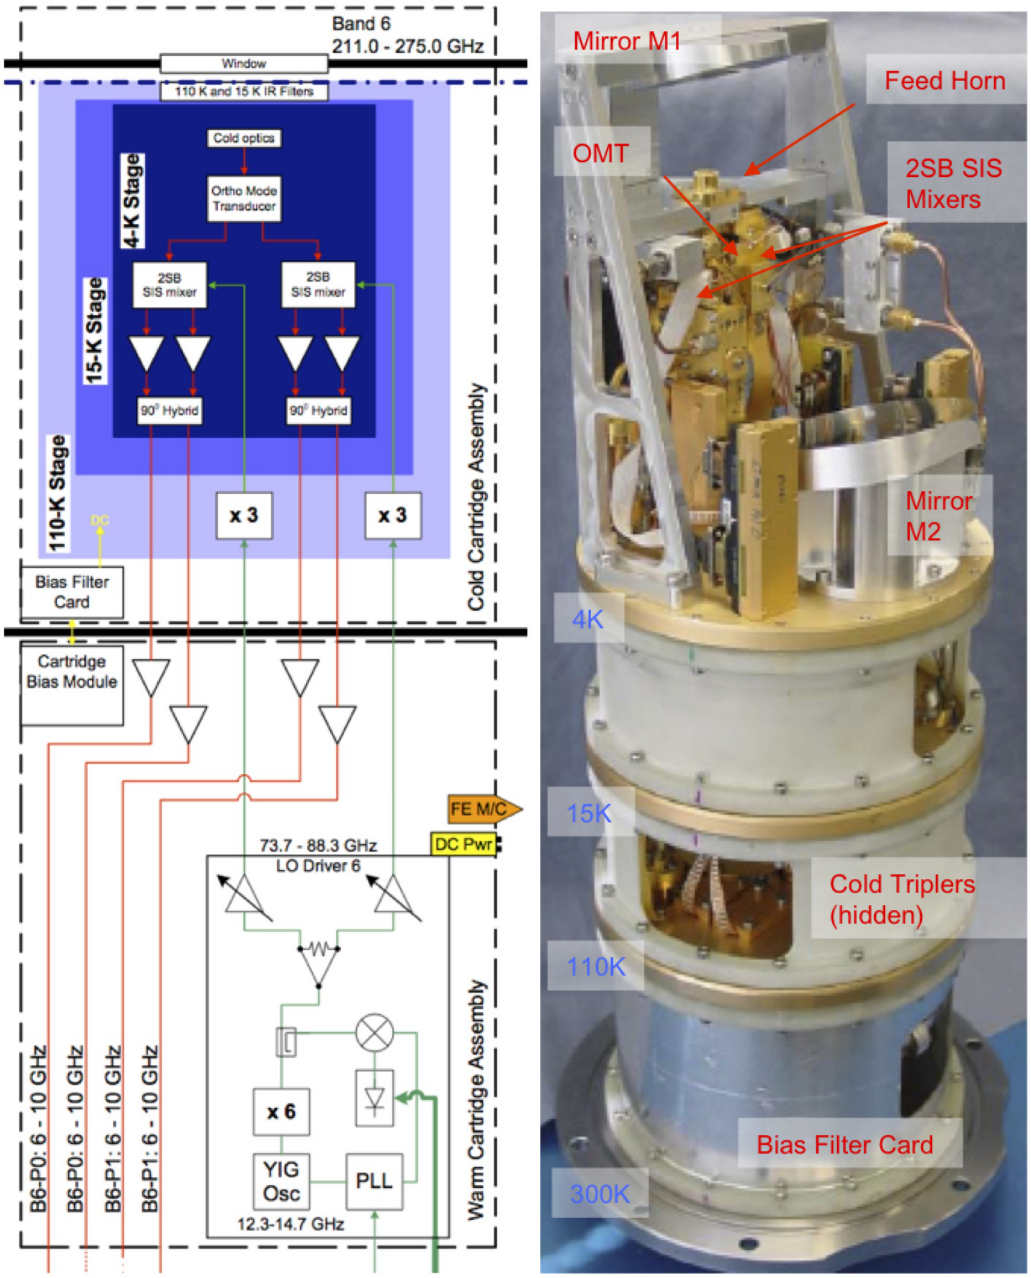

Figure 3.12: Band 6 receiver block diagram, and (right) image of cartridge. Note the OMT used to split the polarizations feeding the two 2SB mixers. The LO around 80 GHz requires an extra x3 multiplier inside the cryostat. The Band 6 cartridges were built at NRAO, Charlottesville. Note that the IF output range is actually 5-10 GHz. The range shown is the one recommended for continuum observations (see text).

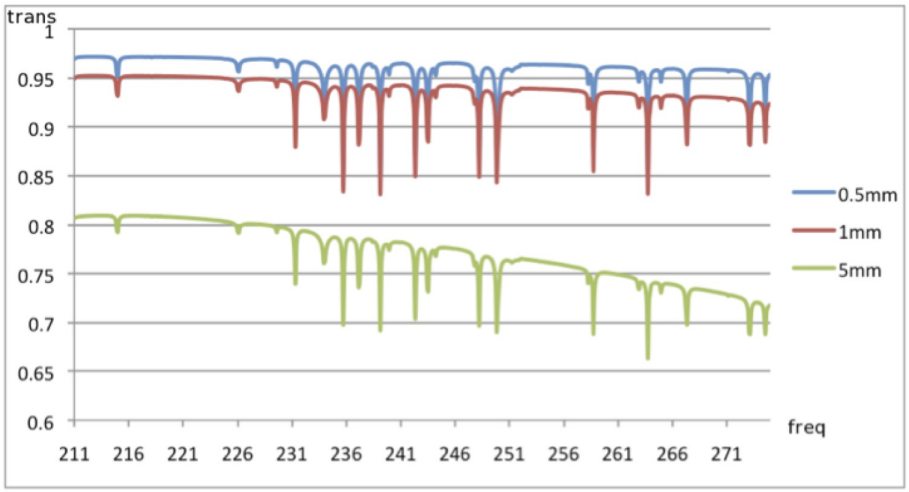

Figure 3.13: Band 6 zenith transmission for PWV=0.5, 1 and 5 mm. Frequency is in GHz.

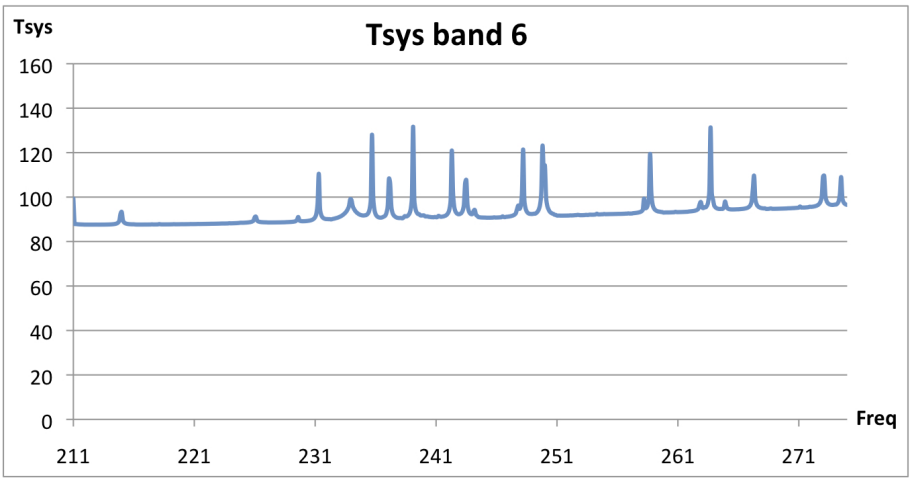

Figure 3.14: Typical $T_{\rm sys}$ at zenith for Band 6 with 0.9 mm PWV, based on measured values of the receiver temperatures. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

3.2.4 Band 7 receiver

The Band 7 receiver covers the frequency range 275–373 GHz (the 0.85 mm atmospheric window). It has a similar cold optics design as Band 6, but uses a wire-grid polarization splitter instead of an OMT (Figure 3.15). A block diagram of the Band 7 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 3.16.

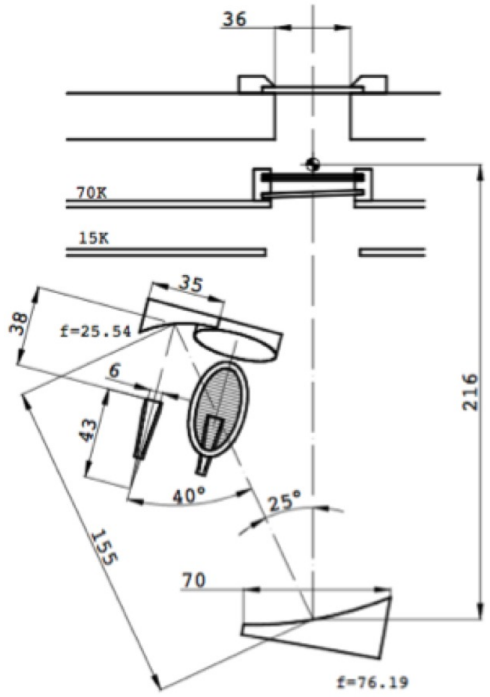

Figure 3.15: Band 7 cold optics arrangement, showing the off-axis ellipsoidal mirrors and the polarization splitter wire grid. The Band 9 cold optics arrangement is very similar this, although the actual numbers differ.

The atmospheric transmission in Band 7 is shown in Figure 3.17 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 <221 K over the whole tuning range, except at the upper end of the band (370-373 GHz), where the specifications are <300 K SSB. However, the performance of the receiver as measured in the lab is considerably better than this. The resulting system temperatures ($T_{\rm sys}$) for 1 mm PWV are shown in Figure 3.18. Note that the atmospheric transmission (and hence $T_{\rm sys}$) at frequencies below 300GHz is considerably better than that of the top half of Band 7; in that respect the performance is closer to that of Band 6.

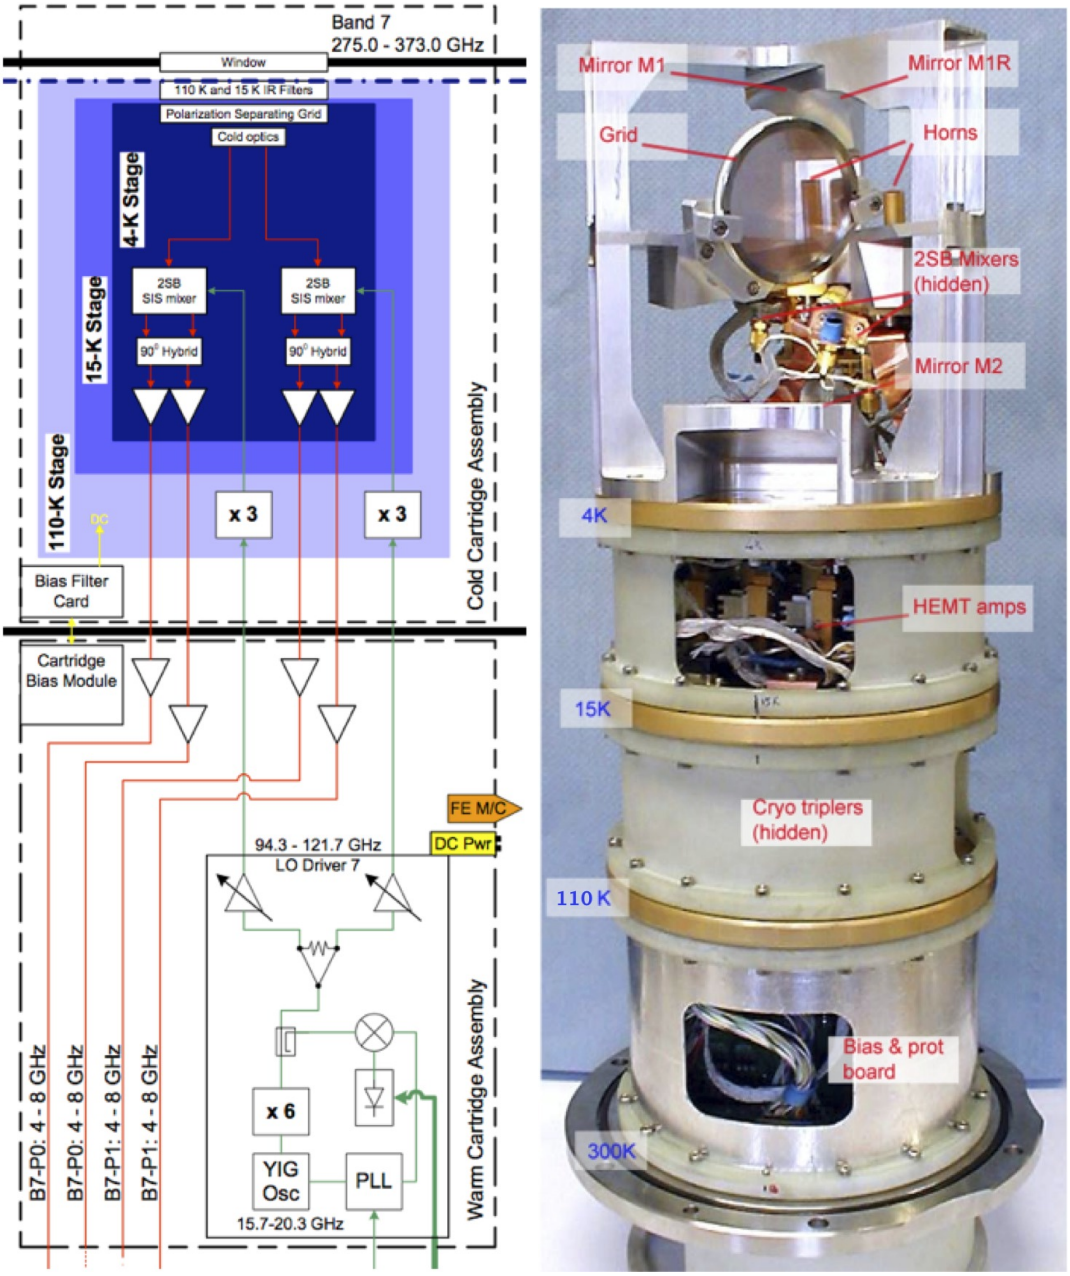

Figure 3.16: Band 7 front-end receiver block diagram, and (right) annotated image of the Band 7 cartridge. Note the polarization-splitting grid and LO injection in the cold optics above the mixers. The Band 7 cartridges were built at IRAM in France.

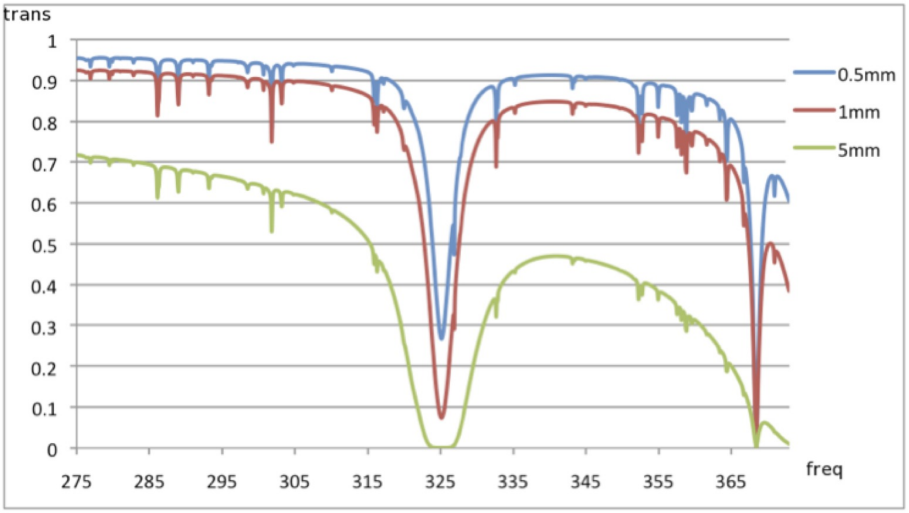

Figure 3.17: Band 7 atmospheric zenith transmission for PWV=0.5, 1.0 and 5.0 mm. Frequency is in GHz. The deep atmospheric absorption at 325 GHz is due to water, and the less prominent absorption feature at 369 GHz is due to oxygen.

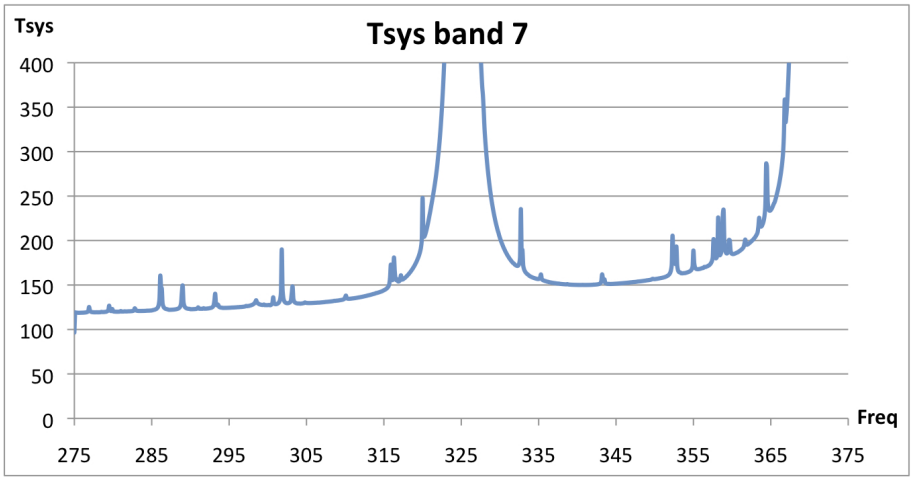

Figure 3.18: Typical $T_{\rm sys}$ at zenith for Band 7 with PWV=0.9 mm. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

3.2.5 Band 8 receiver

Band 8 covers the frequency range 385-500 GHz (650 μm atmospheric window). The cryogenic optics of this receiver adopts a single mirror to couple a feed horn in front of an SIS mixer block to the sub-reflector. A single feedhorn feeds an ortho-mode-transducer (OMT) which splits the two linear polarizations and feeds the 2SB SIS mixers (Figure 3.19).

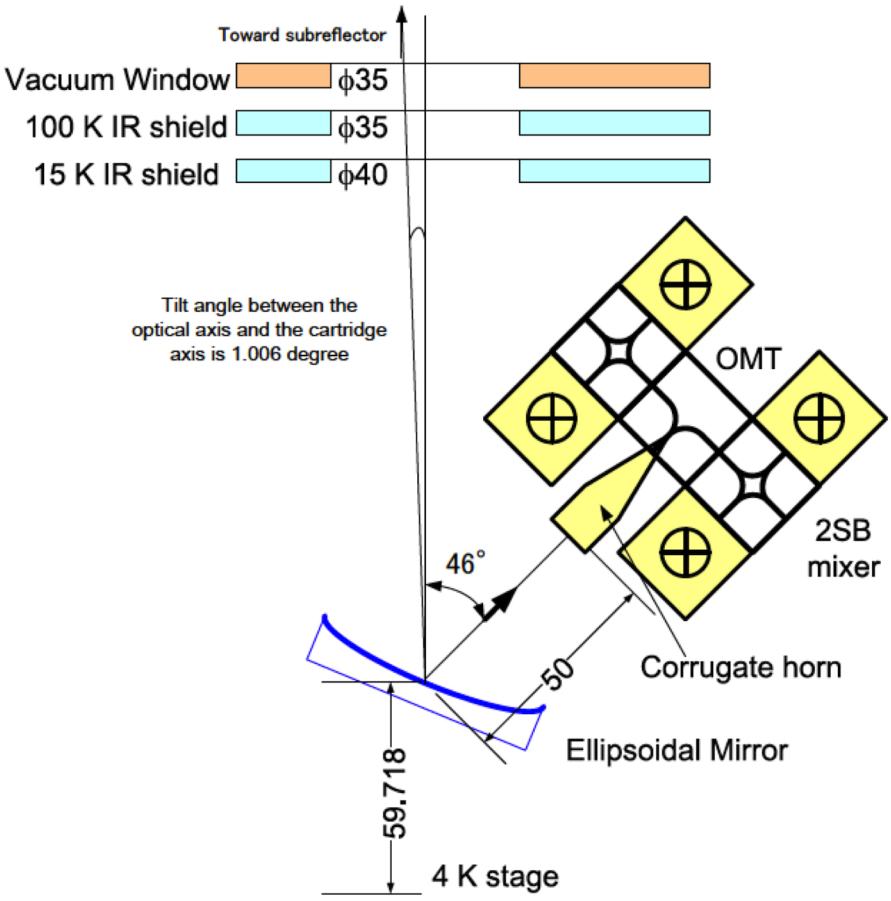

Figure 3.19: Optical layout of the Band 8.

A block diagram of the Band 8 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 3.20. The Band 8 CCA consists of a cold optics, a feed horn, an OMT, 2SB SIS mixers assemblies, cold IF amplifiers, isolators, and LO frequency sextuplers.

The atmospheric transmission in Band 8 is shown in Figure 3.21 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 is considerably better than this. The resulting system temperatures ($T_{\rm sys}$) for 1 mm PWV are shown in Figure 3.22.

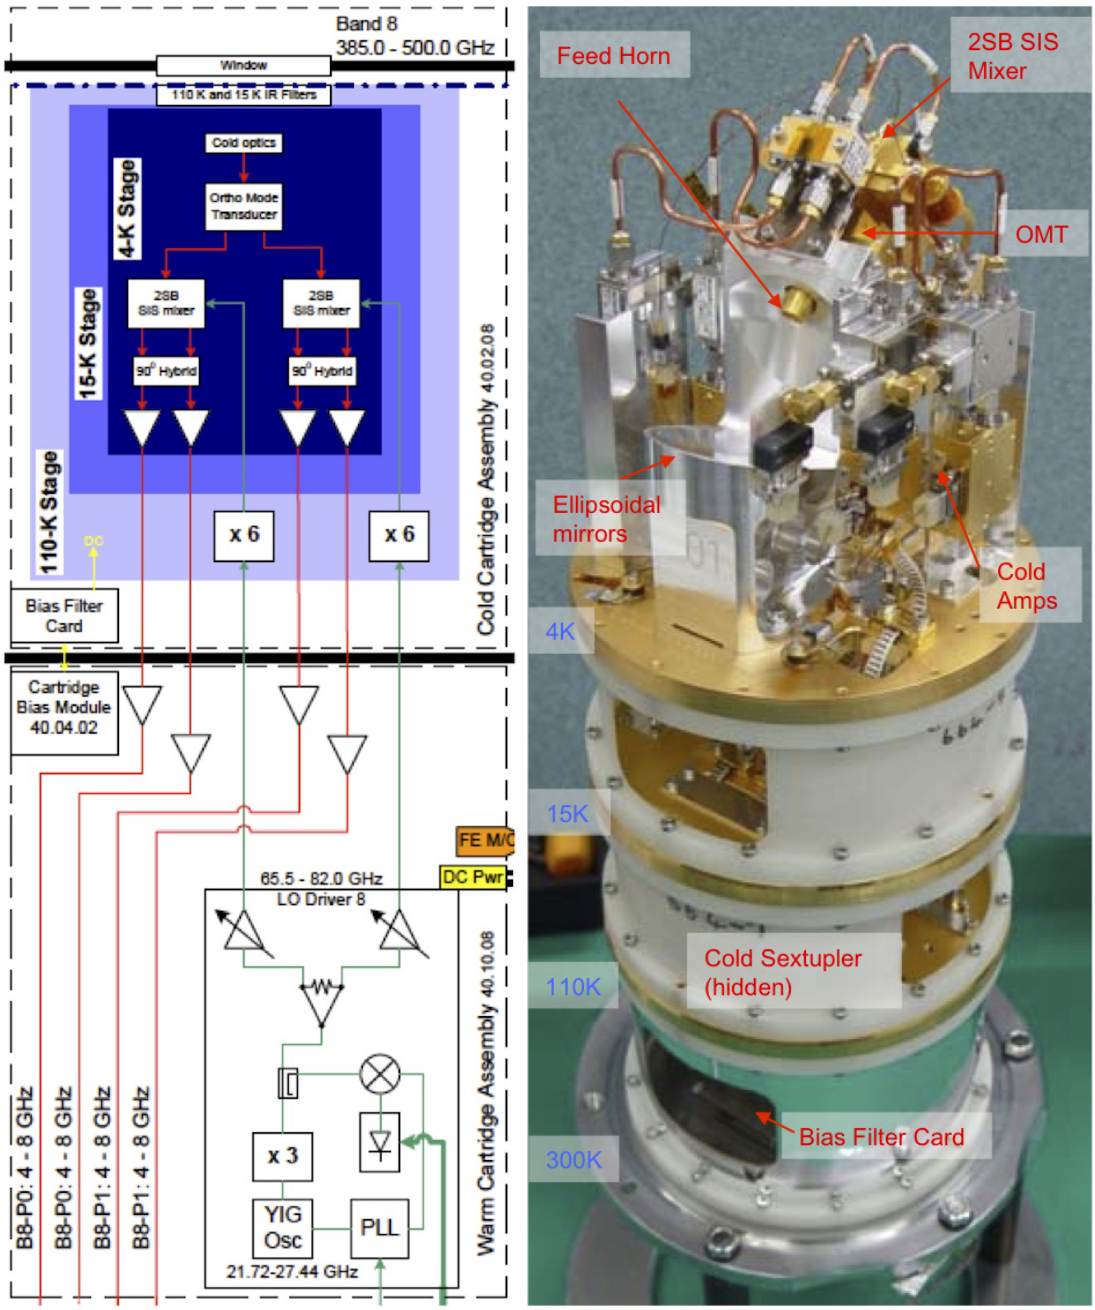

Figure 3.20: Block diagram of the Band 8 receiver (left) including CCA (upper) and WCA (lower). Right image shows a Band 8 CCA. Note the single feedhorn which feeds the OMT, splitting the two polarization signals for the 2SB SIS mixers. The Band 8 cartridges are constructed at the Advanced Technology Center (ATC).

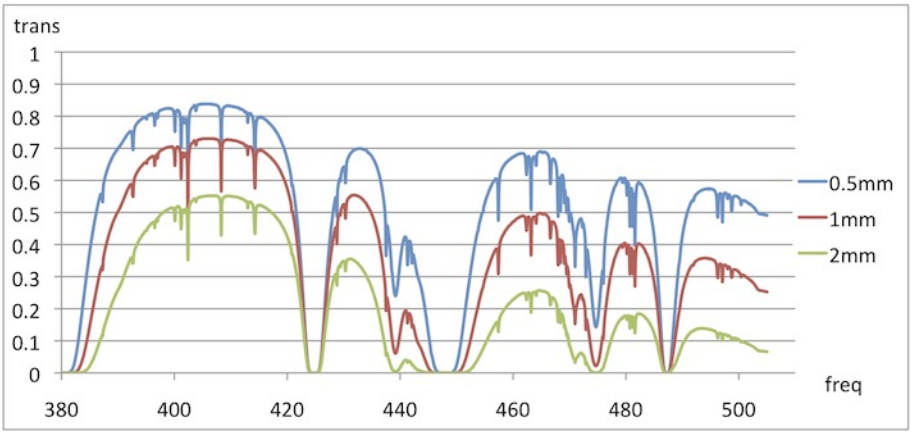

Figure 3.21: Band 8 atmospheric zenith transmission for PWV=0.5, 1.0 and 2.0 mm. Frequency is in GHz. The atmosphere in the Band 8 frequency range has some deep absorption by water and oxygen.

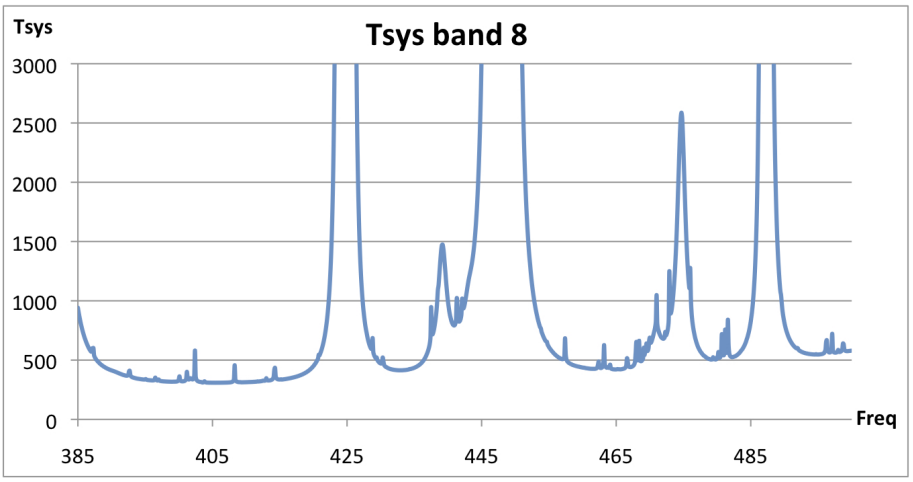

Figure 3.22: Typical $T_{\rm sys}$ at zenith for Band 8 with PWV=0.5 mm. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

3.2.6 Band 9 receiver

Band 9 covers the frequency range 602-720 GHz (450 μm atmospheric window). It uses a wire grid to separate the two orthogonal polarizations, as well as to provide the LO injection scheme (Figure 3.23).

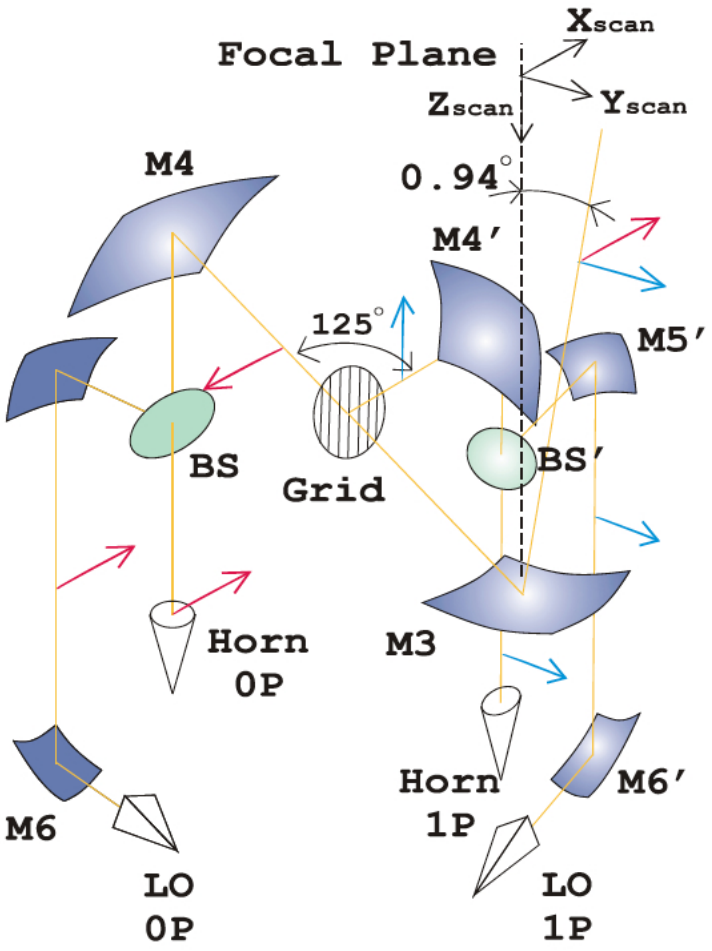

Figure 3.23: Optical layout of the Band 9.

The mixers are double sideband (DSB), and therefore additional techniques must be employed during the observations to either separate the sidebands or reject the unwanted sideband. In Cycle 2, LO offsetting will be used to reject one of the two sidebands, which can be chosen independently for each spectral window. Note that LO offsetting does not reject the noise from the unwanted sideband, it simply moves any correlated signal to a high fringe rate so that the signal is smeared over a larger bandwidth increasing noise incoherently. The IF bandwidth in this receiver is 8 GHz per polarization (7.5 GHz effective bandwidth after the IF Processor units, see Section 5.5), covering 4-12 GHz. A block diagram of the Band 9 receiver, including the cold cartridge and warm cartridge assembly, is shown in Figure 3.24.

The Band 9 atmospheric transmission is significantly dependent on the PWV, as illustrated in Figure 3.25 for 3 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, and Figure 3.26 shows the expected $T_{\rm sys}$ for 0.5 mm of PWV, over most of the band given the expected receiver noise. As well as having a lower atmospheric transmission and a less stable atmosphere, Band 9 observing provides several challenges for observing: finding sufficiently bright calibrators (most QSOs are relatively faint at this frequency), requiring accurate pointing for the relatively small primary beam, and the need for the highest level of stability in the

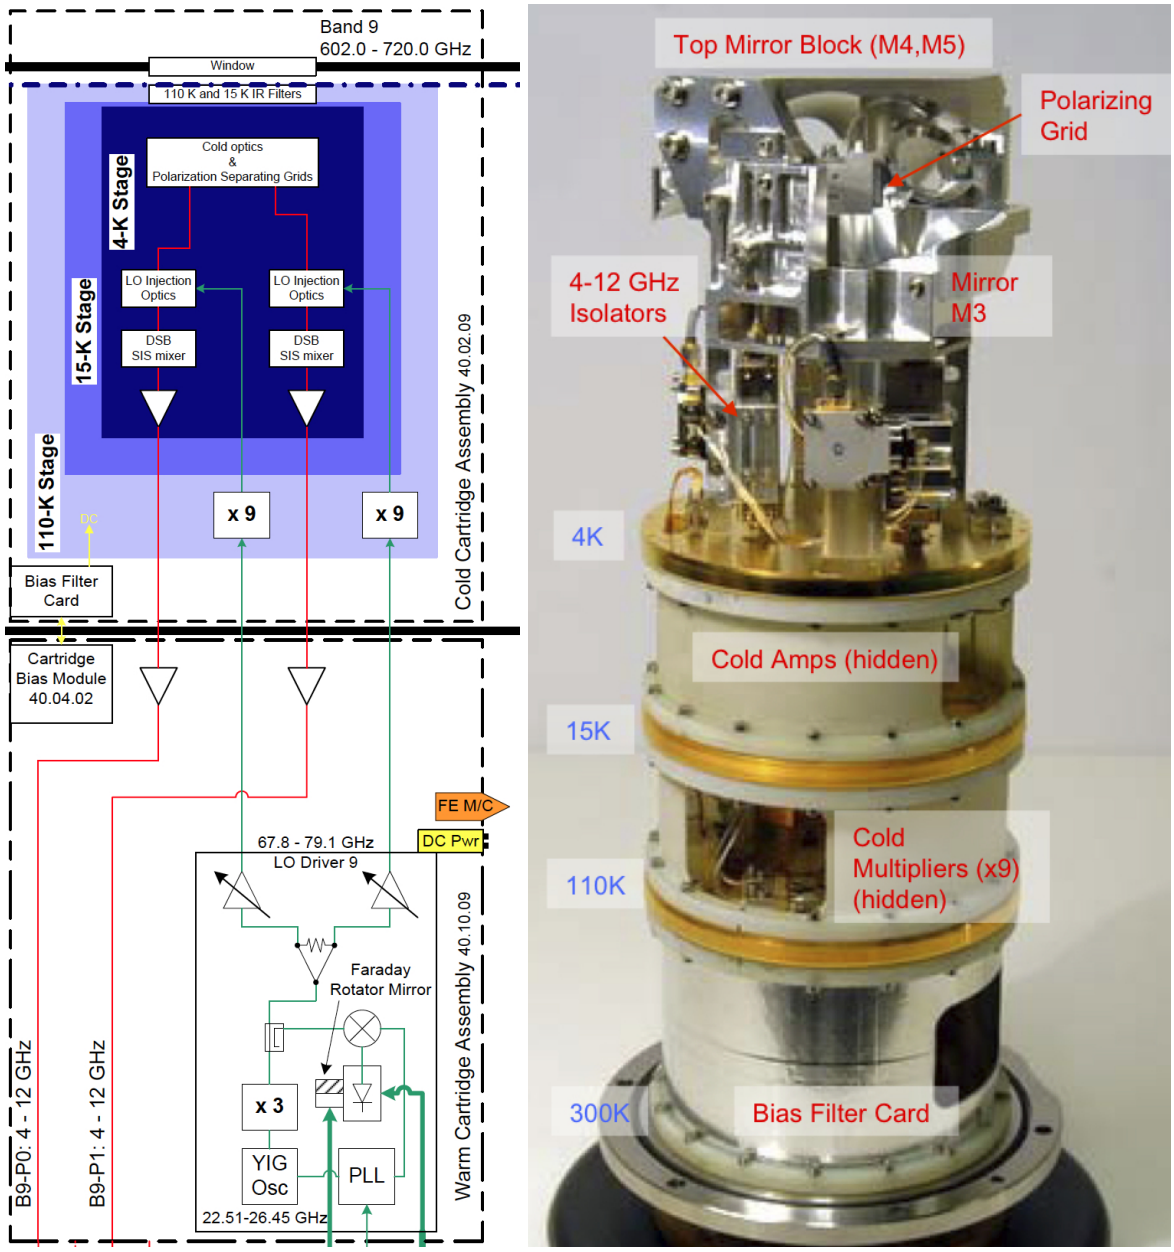

Figure 3.24: Block diagram of Band 9 cartridge (left) and a schematic image (right). Note that there are only two IF outputs, one from each polarization in this DSB receiver. The Band 9 receiver was built at SRON in the Netherlands.

rest of the system.

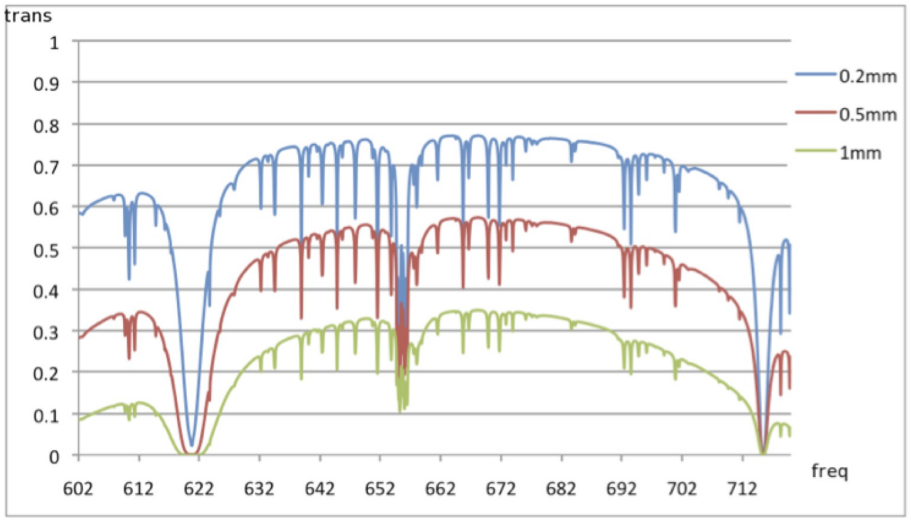

Figure 3.25: Band 9 zenith transmission for PWV = 0.2, 0.5 and 1 mm. Frequency is in GHz.

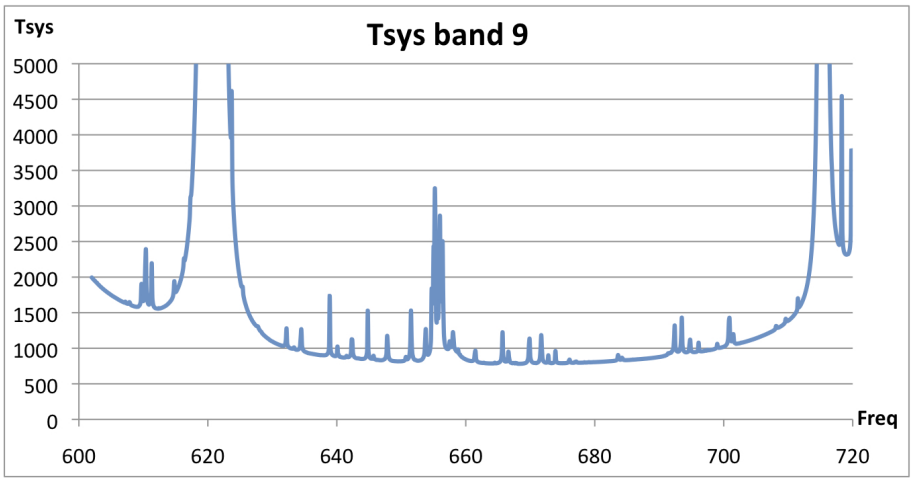

Figure 3.26: Typical $T_{\rm sys}$ at zenith for Band 9 with PWV=0.5 mm. ($T_{\rm sys}$ was computed using only the receiver temperature values adopted in the OT and the atmospheric contribution. No spill-over or background terms have been included.)

Chapter 4

The Correlators

The observed sky frequencies need to be down-converted to a lower frequency before the signal is digitally sampled and sent to the correlators. The frequency down-conversion in the IF system involves a set of Local Oscillators (LOs) and mixers, described in detail in Section A.7. The outputs from the IF system, known as "basebands", cover 2-4 GHz. These analog baseband signals are sampled, digitized and transferred via fiberoptic cable to one of the two correlators, available for Cycle 2: the 64-input Correlator (also known as the baseline correlator) and the ACA Correlator.

For Cycle 2, both have identical characteristics as far as the 'end user' is concerned. They can handle up to 4 baseband pairs (defined as one of the basebands in each polarization) per antenna simultaneously. The 64-input Correlator is used primarily for the main 12-m Array, and the ACA Correlator used primarily with the ACA 7-m Array and Total Power Array (TPA) Array. However, crossbar switching allows for some flexibility in this arrangement. Both correlators can be run simultaneously or independently. This allows for example the 12-m Array to observe one object using the 64-input Correlator, while the ACA Correlator is used with the 7-m Array to observe either the same or a different object. In addition, it is possible to take TP autocorrelation data simultaneously using a separate sub-array of the ACA Correlator.

4.1 Digitizers

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

- The first stage of digitization is in the backend (BE) at the antenna. This is a 3-bit sampler (8 levels), and the digital signal is converted to a parallel datastream for transmission over fiberoptic cable to the correlators in the AOS technical building.

- The second stage of sampling is from the output of the station cards in the 64-input Correlator feeding the correlator boards themselves. This can use 2, 3 or 4 bits sampling. In the latter cases, the number of channels is reduced. In Cycle 2, only 2-bit sampling is available. In the ACA Correlator, 4-bit resampling is always used at the second stage, which is between the Fourier Transform part and the correlation multiplication part.

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

Because of the limited number of bits used in sampling, there is a slight increase in the noise on the signal from the digitization (the quantization noise). This noise depends on the number of bits and the level into the samplers; at the optimum setting, the loss with 2-bit sampling is 12% (the mode available in Cycle 2), which decreases to 1.1% for a 4-bit sampler (to be available for the 64-input Correlator in the future).

4.2 The 64-input Correlator

The 64-input Correlator is of hybrid design that enhances the performance of more traditional lag correlators (XF)[1], and is known as an FXF system. It is described in detail in Escoffier et al. (2007, A&A 462, 801). It operates in two basic modes, TDM (Time Division Mode) for low resolution wideband continuum observations, and FDM (Frequency Division Mode) for higher spectral resolutions. A simplified overview diagram of the 64-input Correlator is shown in Figure 4.1. It consists of 4 quadrants, all of which are available for Cycle 2. Each quadrant can handle a full baseband pair for up to 64 antennas for a total bandwidth of up to 8 GHz per antenna per polarization.

4.2.1 FDM mode

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

For Cycle 2, the TFBs in each baseband are stitched together in the correlator data processing to give up to 4 Spectral Window (spw) per baseband[2]. This provides nominal bandwidths from 62.5 to 2000.0 MHz, depending on the number of sub-bands used[3]. The full list of Cycle 2 correlator modes for dual-polarization operation is given in Table 4.1; this includes the nominal bandwidth internal to the system and the available bandwidth that the end-user sees in their spectrum. Normally the stitching is reliable and the individual sub-bands cannot be discerned in the final spectrum.

At the output of the TFBs, the signals are re-quantized to 2, 3 or 4 bits (for Cycle 2, only 2-bits; see Section 4.1) for correlation. In this case, FDM modes provide up to 7680 channels per baseband pair. For N polarization products, the number of channels is 7680/N, so for the most common operating mode (dual-polarization FDM) this gives 3840 channels (Table 4.1). FDM is used mostly (although not exclusively) for spectral line observing. Additional averaging of spectral channels is available in multiples of 2, to reduce the data rate in FDM mode. For example, in FDM dual-polarization mode, the number of channels per spw could be 3840 (the maximum), 1920, 960, 480, or 240. This is described in detail in § 4.4.3.

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

4.2.2 TDM mode

In TDM mode, the TFBs are bypassed and the full 2 GHz baseband is passed to the correlator. For Cycle 2, only 2 of original 3 bits are used for the correlation by dropping the least significant bit (LSB) (see Section 4.1). The signal is distributed in time over several correlator chips, each of which performs the cross-correlation in parallel

[1] in XF or FXF, X stands for correlation and F for Fourier Transform
[2] In future observing Cycles, the sub-bands will be independently tunable for full flexibility, allowing multiple spws per baseband
[3] Note that the available bandwidth is 15/16 of the nominal bandwidth

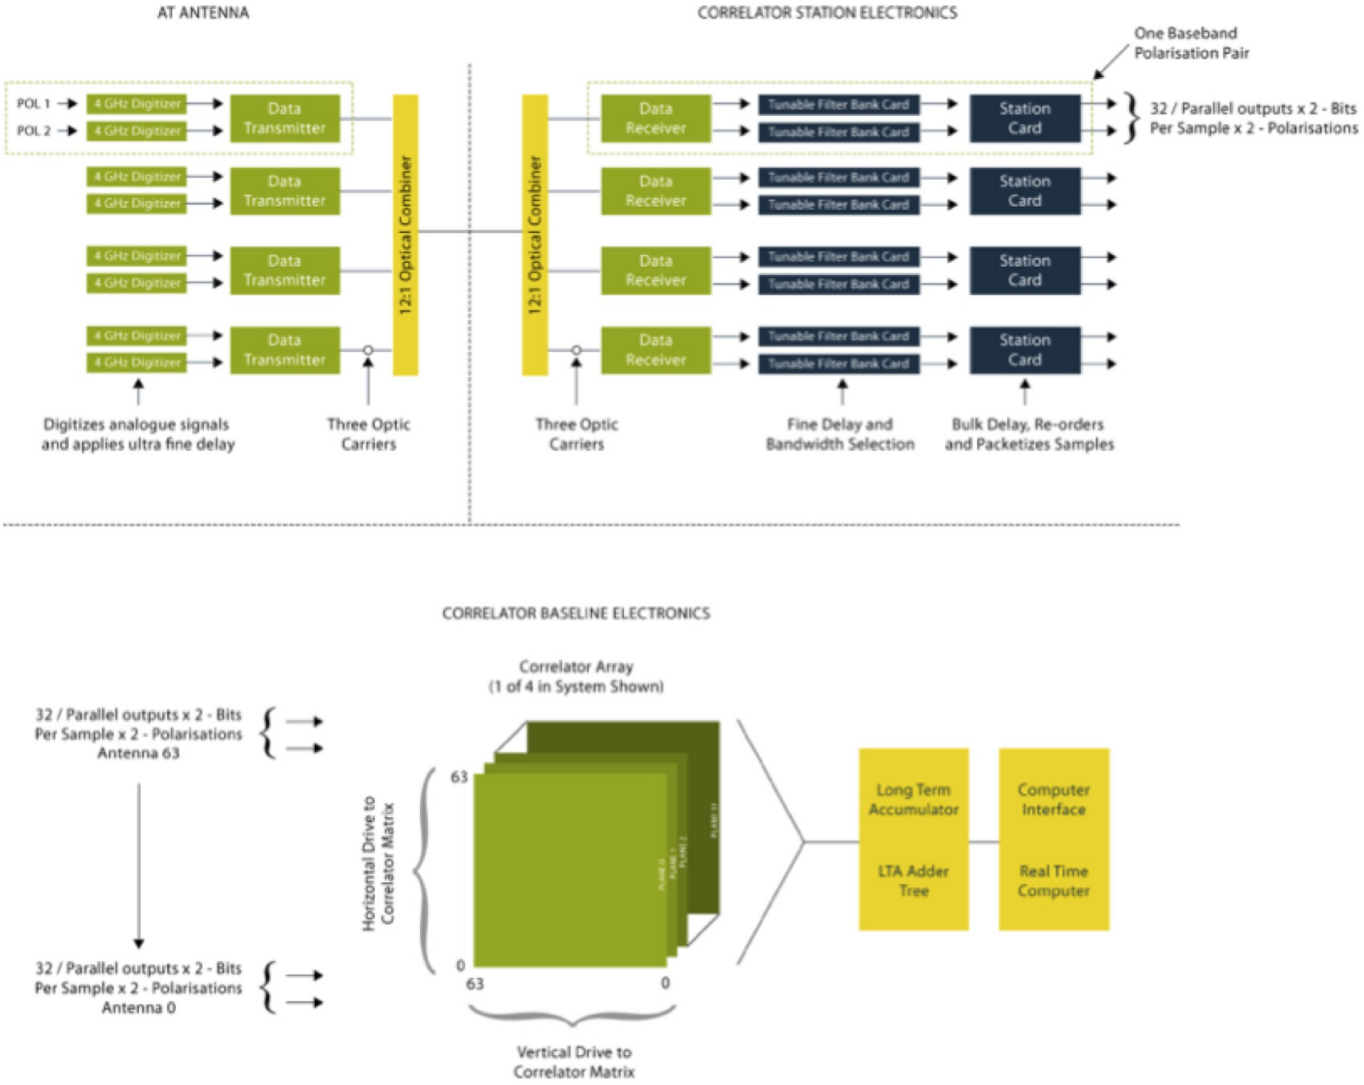

Figure 4.1: Overview diagram of the ALMA digitizers, data transmission system and 64-input Correlator. The digitized data from the individual basebands are transferred to the tunable filter banks (TFB) and station cards (top right) for the parallelization and processing required on each antenna datastream. The matrix of correlator boards (lower half) does the cross-correlation between antenna stations, and the resulting data is then integrated in the Long-term Accumulators (LTAs, lower right).

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

Table 4.1: Correlator modes (in dual polarization). Note that the channel spacing given here is *not* the same as the spectral resolution because of the applied weighting function; for calculating spectral resolution, see Section 4.4.3.

for some fraction of the time interval. The correlations are then re-combined to get the fully time-sampled data. The TDM mode provides up to 256 channels per baseband (for N polarization products, the number of channels are reduced to 256/N), and no edge channels are dropped in the correlator. As the full 2000 MHz baseband is covered, this requires some truncation of the edge channels in offline data processing – see Section 5.5. TDM is used mostly for continuum observing. It has the advantage of having a lower data rate than FDM, and is therefore used for pointing, focus, delay, system temperature, sideband ratio and other calibration observations where high-resolution is not required.

4.2.3 Correlation and realtime processing

In the cross-correlation modules, the multiply-and-add operations are performed at a clock rate of 125 MHz (4 GHz samples divided by 32). Each board (Figure 4.2) contains 64 4-k lag correlator chips, connected in a matrix to provide the required complex correlations. Cross-correlation data then goes to a time integrator (the Long Term Accumulator, LTA – Figure 4.1) that adds the data up to an integer multiple of 16 ms for cross-correlations and 1 ms for auto-correlations. The data can then be further averaged over a dump period in the correlator Data Processor computers (CDP).

In Cycle 2, each baseband pair is fed to one correlator quadrant, which can have a bandwidth mode given in 4.1. It will be possible to set the correlators to different modes in different basebands. So, for example, one baseband can be set to TDM for the continuum, and another baseband can be set to a high-resolution FDM mode to study an individual line. More details of spectral setups are given in Chapter 5.

Section 4.4 describes further the processing of the data in the correlators.

4.3 The ACA Correlator

The ACA Correlator is using the FX technique, whereby the incoming digitized data stream is first Fourier-Transformed ("F") in realtime before the correlation multiplications ("X") are performed. All processing is done with high-speed FPGA chips, rather than custom chips as used in the 64-input Correlator. An overview block diagram is shown in Figure 4.3, and a detailed description of the ACA Correlator is given in Kamazaki et al., 2012 (PASJ, 64). Although the internal operation is different, the correlator is designed to mimic the performance capabilities of the 64-input Correlator in bandwidth, resolution and other properties. The ACA Correlator does not use sub-bands. It does allow for further flexibility of modes, but this is not available in Cycle 2.

A detailed block diagram of the ACA Correlator itself is shown in Figure 4.4. Similar to the 64-input Correlator, each quadrant of the ACA Correlator processes a baseband pair; so the complete system can process up to sixteen antennas independently of the 64-input Correlator. After receiving the digitized baseband

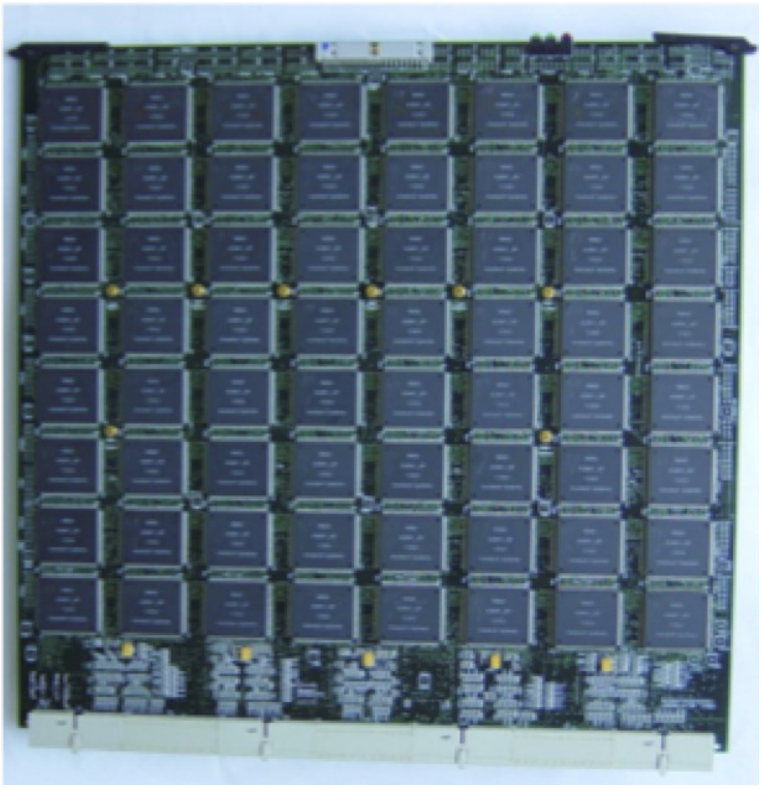

Figure 4.2: Correlator board for the 64-input Correlator, with the 8x8 array of custom correlator chips, each with 4k lags. There are 128 such boards per correlator quadrant, adding up to a total of 32768 chips.

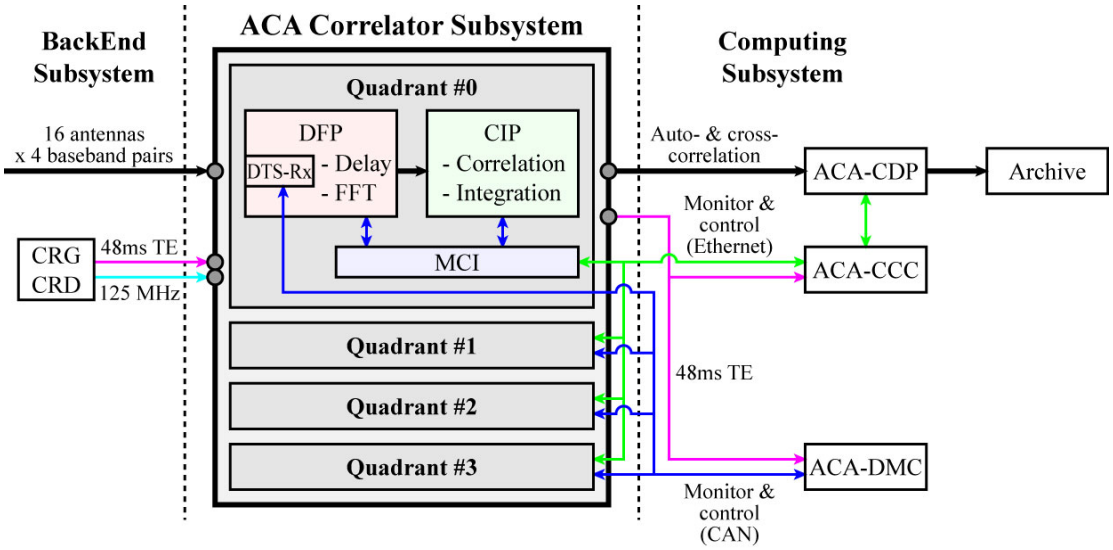

Figure 4.3: Overview diagram of the ACA Correlator. The signals from the BackEnd Subsystem are processed in the DFP (DTS-Rx and FFT Processor) and CIP (Correlation and Integration Processor) modules and output to the Computing Subsystem (see text and Figure 4.4 for details of these). The modules are controlled by the MCI (Monitor and Control Interface), which communicates with the correlator control computer (ACA-CCC). The DTS-Rx is connected with the monitor and control computer (ACA-DMC) for the monitor and control compatibility with the 64-input Correlator.

signals in the DTS-Rx (Digital Transmission System Receiver), the DFP (DTS-Rx and FFT Processor) modules compensate for geometrical delays between antennas and performs the 1M(= 2^{20})-point FFT (Fast Fourier Transform). This results in 512K(=2^{19}) 16-bit complex values (hereafter, voltage spectra) for every baseband per antenna, giving a frequency resolution of 3.815 kHz (= 2 GHz ÷ 512K-point). All the voltage spectra are re-quantized to 4-bit complex from 16-bit complex and sent to the CIP (Correlation and Integration Processor) modules for the calculation of power spectra.

In the CIP modules, the voltage spectra in the frequency ranges specified by observers are multiplied with each other between antennas. Combining signals from different antennas produces the cross-correlation (XC, cross-power) spectra, whilst signals from the same antenna produce auto-correlation (AC, auto-power) spectra. The same receiver polarizations are always combined, i.e. XX and, if present, YY; the cross-hand polarizations XY and YX can, optionally, also be formed. The XC and AC spectra are then integrated to give the user-specified channel separation and output to the post-processing computers (ACA-CDP) in the Computing subsystem through optical fibers.

Once the data have been accumulated, the CDP processing proceeds in the same way as the 64-input Correlator, and further spectral and temporal integration are performed, before data is sent to the archive. The maximum data rate from the ACA Correlator is 3.6 MB/s, independent of the number of antennas. However, the average data rate to the archive will be lower during typical observing, because of overheads and the use of TDM modes during some observing.

The ACA Correlator supports flexible spectral configurations, which are compatible with those of the 64-input Correlator. However, it does employ tunable filter banks (TFBs) and produces the FDM mode in another way. As previously mentioned, the ACA Correlator adopts a conventional FX-architecture and always performs spectroscopy of inputs signals using FFT at the highest frequency resolution inside the correlator. Hence, it can easily output multiple spws with different frequency ranges and resolutions by selecting the highest resolution spectra and summing them up as shown in Figure 4.5. In the spectral integration, two types of integration methods are implemented in the ACA Correlator. One is non-weighted spectral integration, which just adds up spectral data without weighting in the frequency domain. The other is frequency profile synthesis, which performs convolution with a weighting function in the frequency domain (Kamazaki et al. 2008, ALMA Memo No.580). The latter is necessary to make the frequency profile of the ACA Correlator the same as the 64-input Correlator. In Cycle 2, only the latter function is available, using the default Hanning smoothing.

4.4 Correlator data processing

Parameters of the data processing in the CDP (the Correlator Data Processor) and later are effectively identical for the 64-input and the ACA Correlator, outlined below.

4.4.1 Integration time intervals, channel-average and spectral data

The correlators produce two datasets: the spectral data, and the channel-averaged data. These are written to the archive from the CDP at different rates (see Figure 4.7). The channel-averaged data is a single complex number for each spw. In the case of FDM, the data is truncated in the CDP before averaging (only the central 15/16 are used). For TDM no truncation is applied. The primary purpose of the channel-averaged data is to provide a smaller dataset for TELCAL, the ALMA online processing software, to compute real-time telescope calibration corrections. It is stored in the ASDM (ALMA Science Data Model) dataset as a separate spw (see 4.4.5). Note that this should not be used as science continuum data in normal observing, but the continuum should be constructed offline in CASA (Common Astronomy Software Applications) using the appropriate portions of the spectral data. The spectral data contains the requested number of channels for each polarization product. The time intervals involved at this stage of the observing and data acquisition are also shown in Figure 4.7. In detail, they are:

Dump duration: The internal time period inside the correlator (16 ms for cross correlations and multiples of 1 ms for auto-correlations) over which data is accumulated before sending to the CDP. Data can be

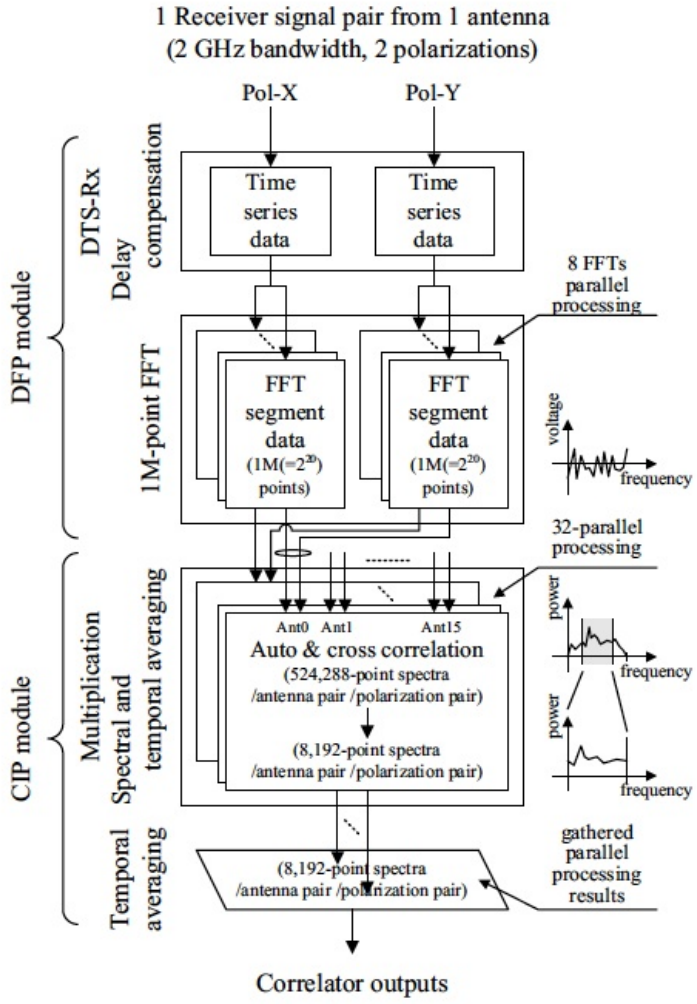

Figure 4.4: Block diagram of ACA Correlator. Time-series data from each antenna are divided in the time domain and processed in an 8-way parallel stream. A Fast Fourier Transform (FFT) is performed using FPGA, resulting in signals in the spectral domain which can then be correlated. The auto and cross-correlation data are then accumulated in time, and the parallel streams averaged together in the CIP (Correlation and Integration Processor) module. The correlation spectra are then fed to the ACA-CDP for further accumulation and processing.

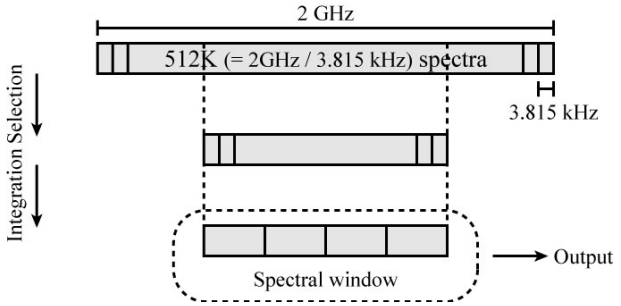

Figure 4.5: Data capture sequence for the ACA Correlator. A spectral window is available to select a frequency range from 512K spectra, whose total bandwidth is 2 GHz and frequency resolution is 3.815 kHz, in the Cycle 2 mode. Selected spectra are integrated and then sent to the output.

Figure 4.6: A quadrant of the ACA Correlator. The left and middle racks compose a quadrant, which processes a baseband pair from sixteen antennas. The right rack is one of four computing racks, which have correlator control, post processing computers, network instruments and so on.

corrected for atmospheric phase fluctuations using the WVR correction once every dump duration in the CDP, or offline in CASA once every Spectral Integration Time (see below). During observations with the 64-input Correlator in Cycle 2, the dump time should be a multiple of 48 ms in FDM and 32ms in TDM. A dump duration of 480 or 960ms is typically used, and it should not be more than the dumptime of the WVR system. It is always 16ms on the ACA Correlator.

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

Spectral integration time: The spectral integration time (also known as the integration duration) is the shortest integration time before the full-resolution spectral data is written to the ASDM. It is also the shortest time interval that can be selected (and potentially self-calibrated) in CASA. The spectral integration time must be a multiple of the channel average time. For Cycle 2 observations using FDM, a Spectral Integration time of several seconds is used to avoid data rate problems (see 4.4.4). For TDM observations, which have fewer channels, a faster rate can be used.

Subscan duration: A further time interval, not shown in Figure 4.7, is the subscan duration or time. In an observation, this will effectively be the shortest time interval where no parameter is changed in the system. For example, this could be a single integration on a source, one point of a mosaic, one scan of a raster, integration at one frequency, or on one of the loads, etc. The subscan duration must be a multiple of the spectral integration time. Typically a subscan might be 30 or 60 seconds; the ASDM might therefore contain 5-10 spectral integration times per subscan, resulting in 5-10 rows of data per baseline/polarisation/baseband product.

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

4.4.2 Online WVR correction

Semi-realtime correction of pathlength fluctuations are done in the CDP using the WVR data from each antenna. The advantage of doing WVR correction in realtime is that it can be performed relatively rapidly (up to the

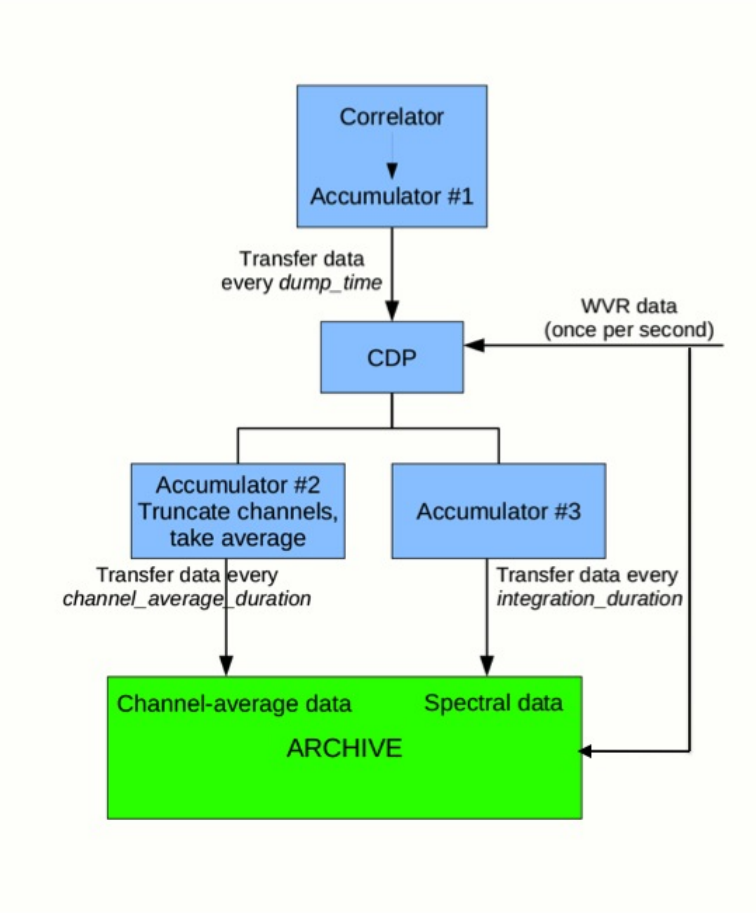

Figure 4.7: Basic data processing and accumulation steps in the correlator and archive. The timing intervals shown are described in the text.

largest value of the dump duration or the WVR chopper wheel rotation and readout period, approximately 1 second), thus enabling corrections to be made for faster variations in atmospheric transmission. At the same time, corrected data can be integrated in the CDP, allowing data to be transferred from the correlator to the archive at a slower rate. Note that once the data is combined, the online WVR correction is non-reversible.

For some time during Cycle 2, it is expected that both the online-corrected and the uncorrected data will be written to ASDM once every Spectral Integration Time (although note that this will double the overall data rate; see 4.4.4). So, as in Cycle 1, it will be possible to do the WVR correction offline on the slower uncorrected data (or even not to do any WVR corrections). But it is expected that the online WVR corrections will track the sky more accurately, and this data will eventually be used directly. The optimal method of WVR correction, potentially including a combination of on-line and off-line corrections, is still under investigation.

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

4.4.3 Spectral resolution, smoothing and channel averaging

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

It is possible to select various weighting functions in the correlation (lag) dimension (see Table 4.2) which will affect the resolution. Without a weighting function (listed as 'Uniform' in the table), the spectral response will be a sinc function (the Fourier transform of the unweighted top-hat function). This has a resolution of 1.2 times the channel spacing, but has undesirably-high spectral ripples or sidelobes, resulting in 'ringing' in the spectra when a narrow line, interference spike, or strong edge channels are present. Different levels of smoothing can be used to suppress these, but they also degrade the spectral resolution[4]. The default weighting function is the Hanning function, which gives a spectral resolution 2.0 × the channel spacing; in this case, the largest sidelobe response for an unresolved signal is 2.6% (see Table 4.2). Note that this level of sidelobe will not normally be noticeable in astronomical spectra, as most lines will normally be resolved over several channels. If a different smoothing function is required in a specific science proposal, this should be justified in the technical case, and this can be selected in the Phase 2 setup of the SB.

Mode	Resolution (FWHM) (channels)	Max spectral sidelobe
Uniform (no weighting)	1.2	-0.22
Hanning (default)	2.0	-0.026
Hamming	1.81	-0.0069
Bartlett	1.77	+0.047
Blackmann	2.230	+0.0012
Welch	1.59	-0.086

Table 4.2: Spectral Resolution and sidelobe levels for different smoothing functions.

Channel averaging is available to bin or average spectral channels on the CDP. Channels can be averaged together in factors of N=2, 4, 8, or 16. The main purpose is to reduce the data rate to the archive and the total data volume. It provides a broader spread of correlator functionality between the current TDM (which has only 128 channels in dual polarization) and full FDM (with 3840 channels in dual polarization mode). It might be quite acceptable for those that need something with more resolution than TDM, but where the FDM channels at the full resolution are unnecessary. Table 4.3 shows the resolutions (in kHz) for different values of N, using Hanning smoothing, in the different bandwidth modes. The channel spacings are in brackets. N=1 is the default un-binned case, where the resolution is 2× the channel spacing.

Note that the default Hanning spectral smoothing function (see above) gives a resolution 2x the maximum,

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

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

Table 4.3: Spectral resolution and channel spacing (in brackets) in kHz for different correlator bandwidth modes (left column) and for different channel averaging factors (columns, N=1 to 16), using Hanning smoothing. The number of channels can be reduced from 3840 (for the un-averaged case, N=1) down to 240 (for N=16). Values are given for the 2-polarization case.

so using N=2 (cutting the number of spectral channels from 3840 to 1920) results in negligible loss of final resolution. It is recommended that unless the maximum spectral resolution is required by the observations, then the number of channels be reduced when feasible. This is selected in Phase 2 of the SB creation. However, note that this is a non-reversible operation!

4.4.4 Correlator speed and data rates

The maximum data rate from the 64-input Correlator in Cycle 2 will be similar to Cycle 1, or about 17 Mbytes/sec (for 32 antennas). However, it is intended that this will be gradually increased with software upgrades to the design specification of 66 Mbytes/s (using 64 antennas). For the ACA Correlator the maximum data rate is 3.6Mbytes/s. These rates set limitations on the shortest value of the spectral integration time in the correlator, particularly in FDM mode, where the number of channels without spectral averaging is 32 times larger than TDM mode. The FDM raw data rate (in Mbytes/sec) is approximately

$$N_A(N_A-1)/2 \times N_{bb} \times 8192\ channels \times (4/10^6) Mbytes/I_T \tag{4.1}$$

where N_A is the number of antennas, N_{bb} is the number of basebands (normally 4) and I_T the spectral integration time. Note that this rate will double if both WVR online-corrected and uncorrected data is required from the correlator. As in cycle 1, the usual FDM spectral integration times set in the SBs are typically expected to be 6 seconds; this may be increased if the on-line WVR correction proves to give accurate and reliable path-length corrections. So under most circumstances the maximum data rate limitations should not have any impact on the science that can be done in Cycle 2, as it is expected that these high values will seldom be reached.

Note that the *average* data rate to the archive will be considerably lower than the maximum correlator data rate, because of observing overheads and the fact that some data taking within an SB will be in TDM rather than FDM mode (eg pointing and atmospheric calibration). The maximum average data rate to the archive is currently 6.6Mbyte/s (see Chapter 13).

4.4.5 Final data product - the ASDM

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

In the ASDM, the binary data are saved into data structures called "spectral windows" (spws). All of the data in a single spw must share the same frequency setup, including the number and width of spectral channels,

and the integration time. The observed spws will be a combination of the "science spectral windows" set up by the proposers in Phase 1 of the OT, and additional spws from observations needed for calibration (pointing, and sometimes system temperature) set up during Phase 2. Additionally, the WVR data are stored in a spectral window with 4 channels around the water line at 183 GHz. In the ASDM, except for the WVR spectral window, each requested spectral window maps into two output spws in the data: one with the requested dimensions of N channels per polarization product (for example 128 or 3840), and a second "channel averaged" version with one channel per polarization product (this averaging is done in the correlator). The "channel averaged" data are used by the on-line telescope calibration system (TELCAL) and for real time diagnostic purposes (QuickLook), and are typically not used downstream in the data reduction. Overall, this can lead to ASDMs with a large number of spectral windows. For example, a typical science observation in FDM mode can have upwards of 25 spectral windows. Luckily every scan/spw combination has an "intent" associated with it that indicates its purpose (pointing, system temperature, science, etc). This intent can be used in CASA to decode how to utilize each spectral window in the data reduction process.

Chapter 5

Spectral Setups

Creating a spectral setup effectively consists of setting the local oscillators and correlator in the system such that the spectral windows cover the desired lines and/or continuum emission. The spectral setup is defined using the Observing Tool. For full details of the OT, see the user manual and reference manuals, available from the ALMA website[1] (and also in the OT itself). In this section we describe how the OT information is used to set up the ALMA system.

An overview of the ALMA local oscillator (LO) and intermediate frequency (IF) system is given in Appendix A.7.1, and illustrated in Figure A.12. As a reminder, the main points are that the front-end receivers use LO1 to generate an IF typically from 4-8GHz. Up to four 2-GHz wide *Basebands* can be placed within this IF range using LO2, and *Spectral Windows* (spws) can be placed within each Baseband. Each Spectral Window (spw) forms a final contiguous spectrum, from 62MHz up to 2 GHz wide. Using the OT during both proposal preparation by the PI (Phase 1) and scheduling block creation by ALMA staff (phase 2), the user chooses the frequencies (or transitions) to be observed and the required spectral window bandwidths or spectral resolutions, and the OT will look for the best solution for the settings of the LOs, correlator etc. There are effectively four different LOs in the system which are set up during the spectral tuning:

- LO1 which sets the front-end tuning frequency
- LO2 which positions the basebands within the front-end receiver IF (each baseband uses a different LO2, hence there are four LO2s)
- LO3 which is the clock frequency of the digitizers (fixed at 4 GHz)
- LO4 (also known as the tunable filterbank LO, or TFB LO), which is a digital LO synthesised in the correlators allowing positioning of the spectral windows within each baseband. Each spw has effectively a separate LO4

The OT and the realtime system have a tuning algorithm that attempts to find the best tuning solution for all these LOs based on the required observing frequencies. Because the LOs themselves are generated by combinations of frequency synthesizers, several different tuning solutions may be possible, and the algorithm picks the best one. Alternatively, no solutions may be possible, and the OT will flag this. A more detailed description of the LO operation is given in Section A.7. Apart from one significant capability increase (see below), Cycle 2 has similar restrictions on the spectral setups and tuning as Cycle 1 (see Table 5.2 in section 5.8). The main limits are that the edges of the 2 GHz basebands cannot lie outside the receiver tuning range listed in Table 3.1, and the edges of the (untruncated) spectral windows cannot lie outside the 2 GHz-wide basebands [2]. The settings of the spws in each baseband are independent, so the resolution and correlator mode as well as the value of LO4 can be different in each baseband. So for example it is possible to have a 62.5 MHz-wide spw centered on a particular line in spw1, and simultaneously use a broadband time division multiplexing

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

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

mode (or TDM mode - see correlator chapter) for all the other spws. Cycle 2 takes this flexibility a step further: it allows for multiple spectral windows within each baseband, described in Section 5.2.2 below. Also, for Cycle 2 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 spectral functions as the 64-input Correlator.

5.1 Frequency definitions

There are several frequencies defined in the system, in the OT phase 1 and phase 2. Most important to the user are:

Center Frequency (Rest) This is set by the user in OT phase 1, and is the central frequency of the spw in the rest frame of the source.

Center Frequency (Sky) This is seen in the OT phase 1, and is the central frequency of the spw in the coordinate frame selected. So it includes the source velocity in the OT, but not the velocity shift of the coordinate frame with respect to the Earth (which is not known until runtime).

Baseband Actual Center Frequency This is the frequency center of the baseband. In OT phase 2, both the Sky and Rest Baseband center frequencies are seen. Unless the spws are centered on the basebands, these will not be the same as the spw Center Frequencies, because of the TFB LOs.

Center Offset Frequency In OT phase 2 terminology, this is known as the TFB LO offset frequency. The center (rest) frequency of the spw is given in terms of this and the center frequency of the baseband by: $(spwCenterFrequency(Rest) = (BasebandActualCenterFrequency) \pm (CenterOffsetFrequency - 3.0GHz)$.

5.2 Spectral setups for multiple lines

The wide IF bandwidth and tuning ability allows for simultaneous imaging of multiple lines. Some examples (with the approximate line frequencies in GHz) are shown in Table 5.1 (for a source of redshift zero). Note that in many cases the lines will not necessarily appear in the center of the spws (for example, in the Band 6 combination given). When <4 spectral windows are required for the primary lines, the others can be set up to cover fainter lines or to observe the continuum, potentially in TDM mode. The selection of secondary lines to be observed can be done using the OT spectral interface. In the case of continuum spws, to maximize the sensitivity, the widest bandwidth mode should be chosen in these spws (i.e. 1.875GHz in FDM, or 2 GHz in TDM - giving a usable bandwidth of ∼1.9 GHz – see Section 5.5). Also the continuum spws should cover as much of the IF band as possible, so they should ideally not overlap in frequency. Not only will this maximize the continuum SNR on the science target, but these continuum data may eventually be used to improve phase and amplitude calibration. The continuum spw frequencies need to be setup manually, and one method of doing this is described below.

5.2.1 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 line frequency) the OT has standard optimized frequencies for each band; these are noted in Table 5.1 and illustrated by the red in Figure 5.1.

If mixed line+continuum operation is desired, for example with a single line in spw 1 observed at frequency f (GHz), the other spws (2-4) can be set up for continuum. For 2SB receivers, normally this would have spw2 in the same sideband (offset by 2 GHz from f) and spw3 and 4 in the opposite sideband (offset by the front-end center IF frequency, f_{IF}). The continuum spws should cover the maximum aggregate unique bandwidth. If *sb*

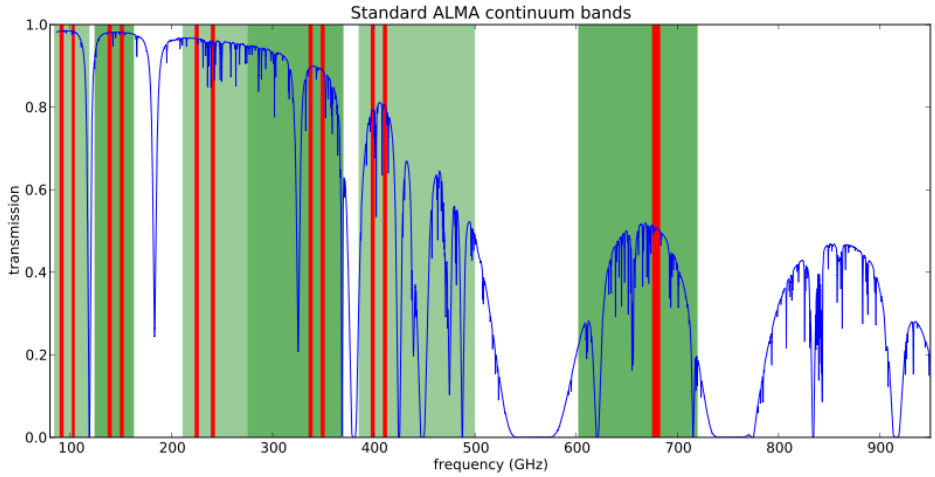

Figure 5.1: ALMA bands available in Cycle 2 (bands 3, 4, 6, 7, 8 & 9), showing the frequency of the standard continuum settings as red shading. This gives the coverage of both USB and LSB, except for Band 9, which has 8GHz of bandwidth in USB only. The available frequency coverage of each band is shown by the green shading, and atmospheric transmission for 0.6mm of PWV is given by the blue line.

is the sign of the sideband of spw1 ($sb = +1$ for USB, $sb = -1$ for LSB), then the 4 spws would be set up as follows:

- spw1: f (spw covering the primary line)
- spw2: $f - 2.0$ (in the same sideband, with a continuum spw 2 GHz below the primary line)
- spw3: $f - 2.sb.f_{IF}$ (in the opposite sideband)
- spw4: $f - 2.sb.f_{IF} - 2.0$ (in the opposite sideband)

f_{IF} is the front-end IF frequency (6.0 GHz for Band 3, 4, 7 & 8, and normally 8.0 GHz for Band 6 and 9). Note that this is an approximate rule - if possible, adjustment of the continuum spws (for example, choosing whether the opposite sideband is LSB or USB) should be done using the OT spectral display, to avoid deep atmospheric absorption features. This is particularly important at Band 9 and near the water lines around 325GHz in Band 7 (see Figure 5.1).

5.2.2 Multiple spectral windows in the same baseband

A new capability in Cycle 2 is the ability to have multiple spws in the same baseband. This is useful for projects where many lines need to be observed simultaneously at high spectral resolution. In cases where 4 lines or fewer are observed, it will normally be better to place each line in a different baseband (as was done in Cycle 1). However, for >4 lines, a separate high resolution spw for each line could be useful. There are some restrictions: each spectral window within the same baseband must have the *same* resolution [3], and the correlator resources used for one baseband (effectively the number of spectral channels) will be distributed among the spectral windows. Thus 2 spws per baseband will each have 1/2 the number of spectral channels in each spw in that baseband, etc. Multiple spws are set up in the phase 1 OT simply by adding more than one spectral line in the same baseband. This will create spws with a correlator fraction less than 1 (so for 4 spws, each has a correlator fraction of 1/4). The number of spws may be increased up to 4 in cycle 2, however, the total number of spectral channels in each baseband is limited. So doubling the number of spws in one baseband will only give half the

[3] although, as before, spectral windows in *different* basebands can have different resolutions

Band	Species/transition	Freq. (GHz)	Sideb.	bandwidth	bb.spw[2]	Notes
3a	**Standard cont.**	97.5[1]	dual	2 GHz/TDM	1-4	LO1=97.5
3b	HCO^+ 1-0	89.188	LSB	62.5 MHz	1.1	HCO^+/HCN/H_2CO
-	HCN 1-0	88.632	LSB	62.5 MHz	2.1	
-	CH_3OH/H_2CO	101.293/	USB	125 MHz	3.1	
-		/101.333				
-	continuum spw	101.3	USB	2 GHz/TDM	4.1	
4a	**Standard cont.**	145.0[1]	dual	2 GHz/TDM	1-4	LO1=145.0
4b	CS 3-2	146.969	LSB	125 MHz	1.1	CS/DCO^+
-	DCO^+ 2-1	144.077	LSB	125 MHz	2.1	
-	SO2 4(2,2)-4(1,3)	146.605	LSB	125 MHz	3.1	
-	H_2CO 2(0,2)-1(0,1)	145.603	LSB	125 MHz	4.1	
6a	**Standard cont.**	233.0[1]	dual	2 GHz/TDM	1-4	LO1=233.0
6b	^{12}CO 2-1	230.538	USB	125 MHz	2.1	multiple J=2-1 CO isotopes
-	continuum spw	231.6	USB	2 GHz/TDM	1(1)	
-	$C^{18}O$ 2-1	219.560	LSB	250 MHz	3.1	
-	^{13}CO 2-1	220.399	LSB	125 MHz	4.1	
7a	**Standard cont.**	343.5[1]	dual	2 GHz/TDM	1-4	LO1=343.5
7b	^{12}CO 3-2	345.796	LSB	62.5 MHz	1.1	CO/HCO^+/HCN
-	$HC^{15}N$	344.200	LSB	62.5 MHz	1.2	
-	$HC^{15}N$	345.340	LSB	62.5 MHz	1.3	
-	continuum spw	343.200	LSB	2 GHz/TDM	2(1)	
-	HCO^+ 4-3	356.734	USB	62.5 MHz	3.1	
-	$H^{15}CN$ 4-3	355.440	USB	62.5 MHz	3.2	
-	HCN 4-3	354.505	USB	125 MHz	4.1	
7c	^{12}CO 3-2	345.796	USB	62.5 MHz	1.1	J=3-2 CO + isotope
-	continuum spw	344.8	USB	2 GHz/TDM	2.1	
-	^{13}CO 3-2	330.588	LSB	62.5 MHz	3.1	
-	continuum spw	331.6	LSB	2 GHz/TDM	4.1	
8a	**Standard cont.**	405.0[1]	dual	2 GHz/TDM	1-4	LO1=405.0
8b	CI 3P1-3P0	492.160	USB	125	1.1	CI/13CI
-	CS 10-9	489.751	USB	125 MHz	2.1	
-	continuum spw	477.5	LSB	2 GHz/TDM	3.1	
-	continuum spw	479.5	LSB	2 GHz/TDM	4.1	
9a	**Standard cont.**	679.0	USB	2 GHz/TDM	1-4	LO1=671.0
9b	^{12}CO 6-5	691.472	USB	500 MHz	1.1	CO/CS
-	CS 14-13	685.436	USB	500 MHz	2.1	
-	H_2S	687.303	USB	500 MHz	3.1	
-	$C^{17}O$ 6-5	674.009	LSB	500 MHz	4.1	

Table 5.1: Some examples of spectral setups possible in Cycle 2. This includes the standard continuum-only setups as well as some multiple line/continuum configurations. *Notes to Table:* **1.** Frequency given for these standard continuum setups are the mean observing frequency of all the spws, which is approximately the frequency of LO1 in dual-sideband (2SB) receivers. See Figure 5.1. **2.** Baseband (bb) and spectral window (spw) numbers are noted as bb.spw

number of channels per spw. This means either lower resolution or a reduced bandwidth per spw. So it may not be worth having many spws in one baseband, but instead just use a broader single spw covering all the lines of interest. For cycle 2, up to 4 spws are allowed per baseband.

Figure 5.2, which is adapted from the phase II Spectral Editor of the Observing Tool (OT), illustrates a more complex spectral setup with multiple spws/baseband which is possible in Cycle 2. The lines observed are given in Table 5.1 example 7b. The blue hashed area represents the receiver tuning range (in this case, a section of Band 7), and the curved line represents the nominal atmospheric transmission for the chosen PWV. The LO1 is 350.0 GHz, and the upper and lower sidebands of the Band 7 receiver are shown as yellow shaded areas. The four basebands, illustrated in this case by the green horizontal bars, can be moved around, but only within the two sidebands. The spectral windows are also shown. In this example, baseband 2 is set for a TDM spw covering the whole 2 GHz baseband for a continuum measurement. In some of the other basebands, multiple spectral windows are used. Baseband 1 has 3 spws of 58MHz bandwidth and 960 channels per spw. At the same time, we are observing two lines in bb3 with 58MHz/1920 channels per spw and one in bb4 with 125MHz bandwidth and 3840 channels. The basebands (and spws) may partially overlap, as seen in this case with bb3 and bb4. Many combinations of spws and basebands can be set up in this way.

5.3 Spectral setups for lines near the edge of the bands

The restriction that the selected bandwidth cannot fall outside the maximum or minimum tuning range of the receiver is an issue for certain lines at the edge of the tuning range. One obvious example is the 12CO(J=1-0) line at a redshift of zero (rest frequency of 115.271 GHz)., This is close to the maximum tuning range of Band 3 (116.000 GHz, see Table 3.1). 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.

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.

5.4 Spectral setup in Band 9 and DSB considerations

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

5.5 Usable bandwidth

The IF system contains an anti-aliasing filter which limits the bandwidth of the basebands. Nominally this filter has -1dB points at 2.10 and 3.90 GHz, giving a maximum bandwidth of 1.8 GHz. However, the IF response is such that the usable bandwidth is slightly wider – i.e. closer to 1.9 GHz. In FDM mode, the correlator outputs a bandwidth of 1.875 GHz, thus in FDM, the filters do not truncate the spectrum, and the full bandwidth in wideband mode can be used. In TDM the correlator outputs a bandwidth of 2.000 GHz, but typically the

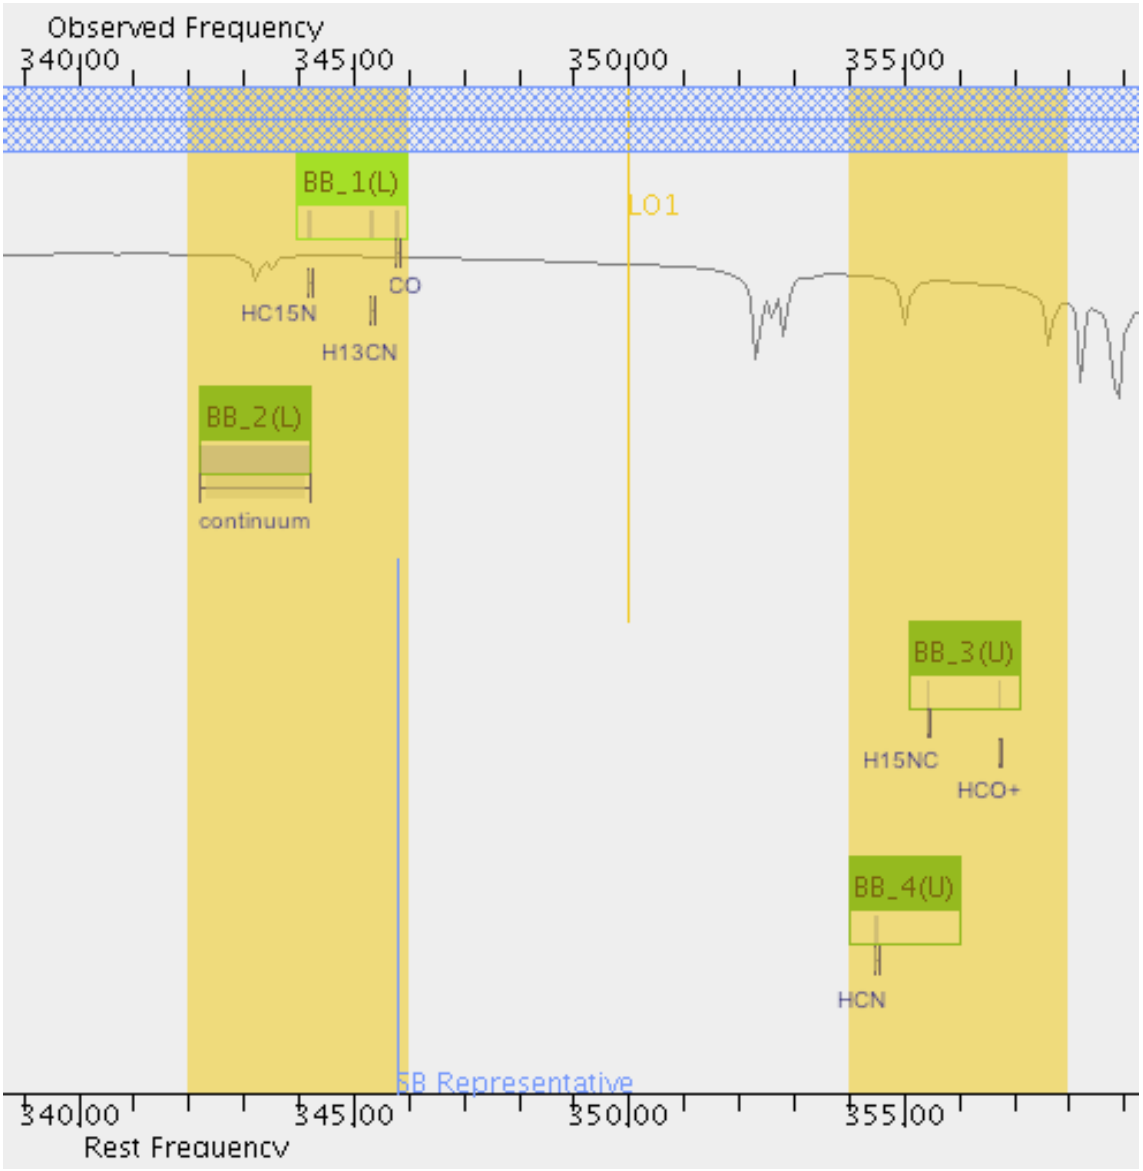

Figure 5.2: Illustration of a frequency setup, based on the OT spectral display. Yellow areas are the IF ranges, green bars are the 2 GHz-wide basebands (bb1-4), and smaller horizontal bars represent the spectral windows (1, 2 and 4 per baseband in this example). The frequency of LO1 is shown by the central vertical line. The blue hashed area shows the possible tuning range of the front end, and the curved line gives an indication of the atmospheric transmission. The bandwidths of the spws are illustrated by the widths of the horizontal bars. The lines observed are given in example 7b in Table 5.1.

edges of the spectra are affected by low power due to this filter and some ringing effects (see upper panel in Figure 5.3). It is recommended that 4 (in double-polarization) or 8 (single-polarization) channels are removed or flagged manually offline. This results in approximately the same usable bandwidth in both TDM and FDM modes and is illustrated in Figure 5.3.

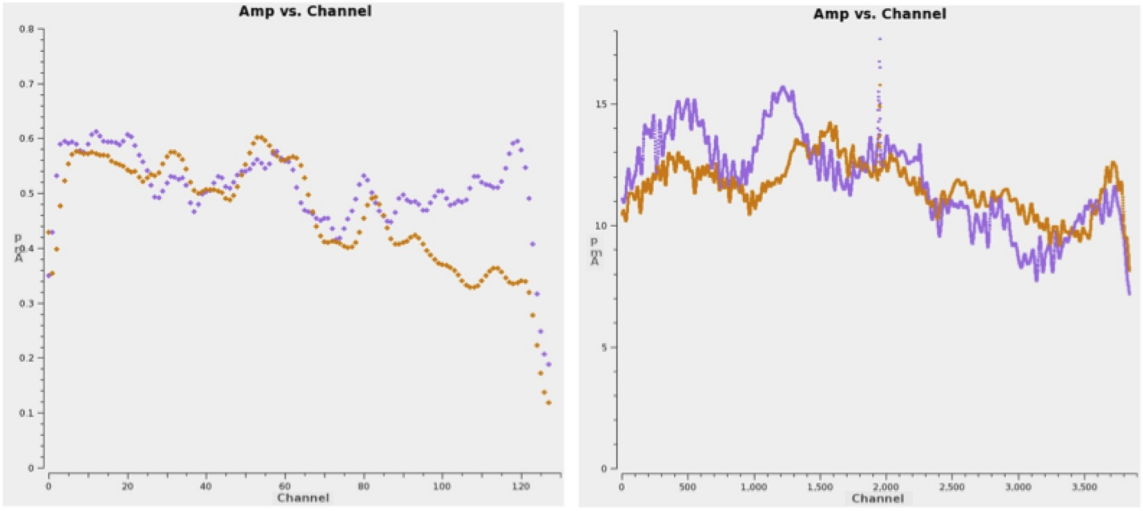

Figure 5.3: Comparison of TDM (left) and FDM (right) autocorrelation bandpass showing the dropoff in total power at the edges in the two modes. Colors represent the two polarizations from the example antenna. In TDM, 128 channels covering 2.0 GHz bandwidth are displayed, which illustrates the drop in power in the upper and lower 4 channels due to the anti-aliasing filter. The FDM spectrum has 3840 channels covering only the central 1.875 GHz, and the drop in power at the edges of this bandwidth is negligible (and comparable with the variations in the bandpass). (The narrow spike in the center of the FDM data is a test signal).

5.6 Spurious signals

Most spurious signals in the cross-correlation data are suppressed using Walsh-switching (see Section A.7.3). This effectively suppresses signals generated after the front ends, but will not reduce those coming in at the observing frequency. Results from tests with 180-degree Walsh switching show very few remaining spurious signals, although harmonics of the LO in the WVRs (ie multiples of 91.66 GHz) are seen. These are very narrow and can be flagged out during data reduction, but observing lines at these frequencies should be avoided if possible.

5.7 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 used in ALMA are:

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

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

LSRK Velocity with respect to the Kinematic local standard of rest, at $20 kms^{-1}$ in the direction 18^h, $+30^o$ [B1900.0].

If a target velocity and reference frame other than topocentric is selected in the OT, at the start of the observation, the velocity of the science target and the velocity of the observatory relative to the chosen reference frame are normally used to set the center frequencies of the science spws[4]. It is also possible to just do the correction for the motion of the observatory relative to the chosen reference frame, ignoring the source velocity. This frequency setting for the science target will normally also be used in the same observation for the bandpass and amplitude calibration. For sources with an external ephemeris file, the rate of change of distance between the target and the observatory taken from the ephemeris is used to compute the source velocity, and this is used to update the LOs regularly during the observation.

5.8 Limitations and rules for spectral setups in Cycle 2

Although the correlators will eventually allow a broad flexibility of spws in a single observation, these abilities are gradually being introduced and tested by CSV before release. Cycle 2 allows a broader subset of correlator capabilities compared with Cycle 1 and 0: the main new feature for Cycle 2 is multiple spws per baseband (see Section 5.2.2). The current rules for spectral setups are given in Table 5.2).

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

Rule	Details	Note
1.	LO1 must lie within the LO tuning ranges given in Table 3.1.	general
2.	No part of the 2.0 GHz-wide basebands can extend over the edge of the IF passband. This means that the baseband centers cannot be closer than 1.0 GHz to the IF passband edge. For example, for a 4.0-8.0 GHz IF range, the baseband center frequency must lie between 5.0-7.0 GHz. The system actually does allow a small extension of the edges of the basebands over the IF edges, to cope with the differential Doppler shifts in the different basebands, but this is small ($<$ 1 MHz at Band 6) and is transparent to the OT user.	general
3.	For 2SB receivers (Bands 3,4,6,7,8), the number of basebands in one sideband can only be 0, 1, 2 or 4 (i.e. not 3). For DSB receivers (Band 9), there is no such restriction (the number can be 0, 1, 2, 3 or 4).	general
4.	No part of the full nominal bandwidth of the spw can extend over the edge of the 2 GHz-wide baseband. For a mode with nominal bandwidth B (e.g. 62.5 MHz), that means the spw center IF frequency (a.k.a. Center Offset Frequency in the OT, phase II) 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.	general
5.	The spw usable bandwidth (ie 15/16 of the nominal spw bandwidth of multiple of 62.5 MHz) should be in an allowed region of the baseband. This is in addition to (4). In practice this means that the required region of the spw should normally be inside the range $\sim$2050 - 3950 MHz (i.e. $>$50MHz from the edges of the 2-4GHz second IF) to ensure that the edge of the anti-alias filter does not significantly affect the IF power. In practice it may be possible to extend some part of the spw to $<$50MHz of the IF edges; this is necessary for the Band 6 setup 6b in 5.1, where the ^{12}CO and ^{13}CO lines are only 70MHz from the IF band edges.	general
8.	The line frequency must be in the center of the spw, otherwise the scheduling block will not validate. This is mentioned in more detail in Section 5.3, for the case where a line is near the edge of the receiver band. It can be circumvented by entering an artificial line frequency.	
9.	Only 2-bit, Nyquist sampling is allowed.	Cycle 2
10.	It is possible to have multiple targets with different redshifts within the same Science Goal in the OT. For SGs including sources with more than one redshift, all the observations must be achievable using five or fewer tunings within the same receiver band, considering the source redshifts and, in the case of spectral lines, the line widths and configuration of spectral windows.	Cycle 2
11.	Multiple spectral windows are allowed in a single baseband. However all the spws in the same baseband are required to having the same bandwidth and resolution. An individual spw within a baseband may occupy 1, 1/2 or 1/4 of the resources (i.e. spectral channels) available in the baseband and the sum of all the fractional resources used within one baseband must be $\leq$1.	New for Cycle 2

Table 5.2: Rules for spectral setups. The note describes whether these are general restrictions, or restrictions in Cycle 2.

Chapter 6

Principles and concepts of interferometry

6.1 Introduction

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

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

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

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

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

By varying the projected distance between two antennas, and having several pairs of antennas in an array, a range of discrete angular scales can be sampled. An ensemble of the data, i.e., sinusoids of various amplitude and phase, can be then "summed" (via the Fourier transform) to produce an image of the sky brightness distribution. How well this image reflects the actual sky brightness distribution depends on how completely angular scales have been sampled. Interferometry, however, works extraordinarily well for observing intrinsically compact targets.

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

6.2 Single-dish Response

As in all astronomy, we define *brightness*, or equivalently *specific intensity*, I_ν, as the electromagnetic (EM) power δP within a range of frequencies (a band width) $\delta\nu$ received from a solid angle $\delta\Omega$ and intercepted by

surface area δA, i.e.,

$$\delta P = I_\nu \delta\Omega \delta A \delta\nu, \tag{6.1}$$

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.,

$$S_\nu = \int I_\nu \, d\Omega, \tag{6.2}$$

where S_ν has typical units of 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.,

$$P_{rec} = \frac{1}{2} I_\nu A_e \delta\Omega. \tag{6.3}$$

The coefficient of $1/2$ in Equation 6.3 comes from the fact that the a receiver is generally sensitive to only one mode of polarization. (The ALMA receivers have been designed to detect both modes, however.) As with all telescopes, antennas bring incident EM power to a focus after reflecting it off a primary surface. The antenna response, i.e., its sensitivity, is a summation of all EM power brought to the focus.

Antenna response is actually dependent on angle from the on-axis (i.e., pointing) direction of the antenna, due to self-diffraction. To demonstrate the angular dependence, Figure 6.1 shows in the top panel the case for EM power of wavelength λ arriving along the axis of an (unobstructed) antenna of diameter D. Since the source of the EM power is very distant, the EM power arrives at the primary surface essentially as plane-parallel wavefronts. Note that the antenna surface is parabolic in shape, so the path that each part of the front travels to the focus is constant. With zero path difference, the EM power arriving on-axis is coherently summed at the focus. This arrangement is only true, however, along the axis of the antenna. In the lower panel of Figure 6.1, 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 $1.22\lambda/D$ (radians), the path difference across the antenna diameter, the *aperture* will equal one wavelength of the incident emission. The combination of such emission at the focus leads to destructive interference at that angle.

As an illustration, Figure 6.2 shows an example antenna power response with angle in one dimension, for an ideal 12 m diameter ALMA antenna uniformly illuminated by emission of wavelength ≈ 0.85 mm (350 GHz). (By ideal, we mean the collecting surface is an unobstructed perfect parabola.) 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 beamsize* and it has a Half Power Beam Width (HPBW) given by:

$$\text{HPBW Primary Beam} \;=\; 1.02 \times \lambda/D, \tag{6.4}$$

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

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

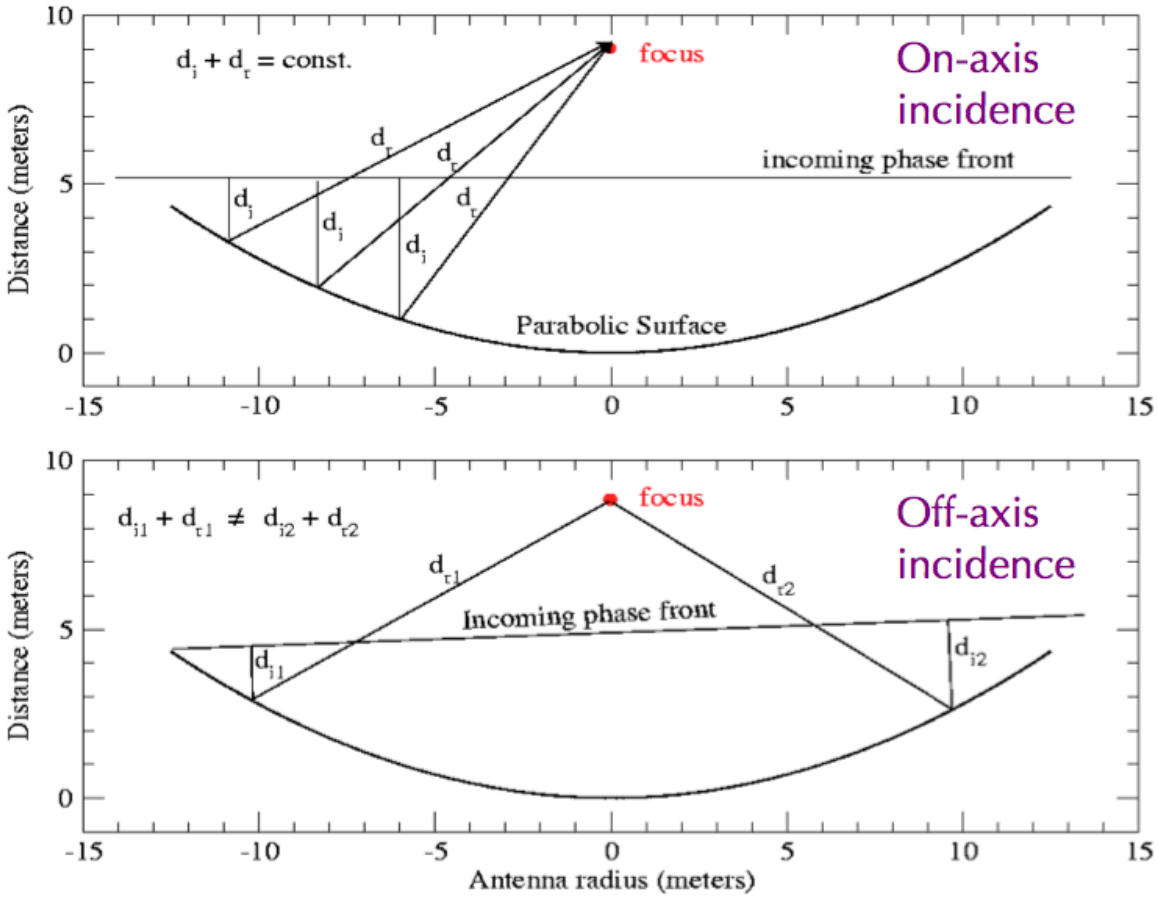

Figure 6.1: Schematic of an incoming plane-parallel wavefront reflecting off a antenna of diameter D and being brought to a focus. The top panel shows the case for a wavefront arriving on-axis. The bottom panel shows the case for a wavefront arriving off-axis. Note that the paths of incident EM power 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.

$$\text{BWFN Primary Beam } = 2.44 \times \lambda/D. \tag{6.5}$$

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, and called the *beam*. Figure 6.2 illustrates the antenna power response in one dimension; on the actual sky, the antenna power response is two-dimensional, and is obtained by rotating the function shown in Figure 6.2 about its central axis. Actual antenna power response can be altered by various effects, including the degree by which the secondary is illuminated, diffraction by the arms supporting the secondary, and surface imperfections. For ALMA, each 12-m antenna has a secondary reflector and support arms that block an effective area of 0.75 m diameter on the primary surface. The actual ALMA feedhorns were designed (via an illumination pattern with a -12dB taper) to provide an beam response that is nearly Gaussian but with low sidelobes, preserving as much resolution and sensitivity as possible. The actual ALMA 12-m antennas have measured primary beam HPWB 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 millimetre/radio wavelengths is essentially diffraction-limited.

Defining θ and φ as orthogonal directional variables (e.g., sky coordinates), we can define $I_\nu(\theta, \varphi)$ and $P_N(\theta, \varphi)$ as the directional functions of the sky brightness and the normalized antenna power, respectively. The 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 beam:

$$P_{rec} = \frac{1}{2} A_e \int_{4\pi} I_\nu(\theta, \varphi) P_N(\theta, \varphi) \delta\Omega. \tag{6.6}$$

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

$$\Omega_A = \int_{4\pi} P_N(\theta, \varphi). \tag{6.7}$$

6.3 Visibilities and Aperture Synthesis

For millimeter/radio astronomy, angular resolutions can be 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 of $\sim 14''$ at 850 μm, and the Arecibo 300-m diameter antenna has an angular resolution of $\sim 3'$ at 21 cm.

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 must contend, however, with accompanying losses of sensitivity, especially to information on various angular scales. In this section, we build on the concepts introduced previously to discuss aperture synthesis in more detail.

Earlier, we described how a plane-parallel wavefront arriving on-axis to an antenna is brought to a focus by a parabolic surface. Since there are no path differences in that case, EM power from across the antenna is brought together in phase at the focus. Now imagine that the parabolic surface is divided into N smaller

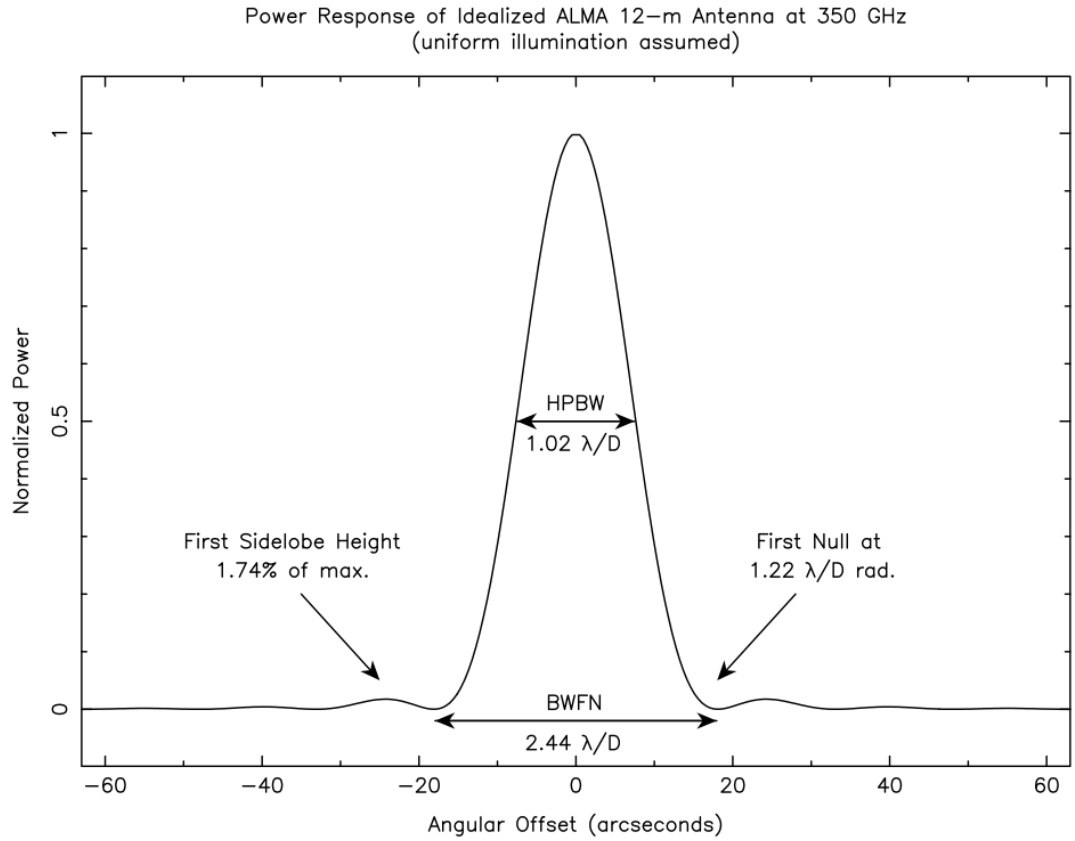

Figure 6.2: 1-D antenna power response for an ideal 12 m ALMA antenna at 300 GHz. The HPBW of the primary beam is $\sim$1.02 λ/D, and the angle of the first null, i.e., its resolution, is $\sim$1.22 λ/D. Note that the HPBW measured from actual 12 m ALMA antennas is $\sim$1.13 λ/D.

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.,

$$V(t) = \sum_i \Delta V_i(t) \tag{6.8}$$

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

$$\langle P \rangle \propto \langle (\sum \Delta V_i)^2 \rangle = \sum \sum \langle (\Delta V_i \Delta V_k) \rangle. \tag{6.9}$$

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

$$\langle P \rangle \propto \sum \langle \Delta V_i^2 \rangle + \sum \sum_{i \neq k} \langle \Delta V_i \Delta V_k \rangle. \tag{6.10}$$

The first and second sets of terms in Equation 6.10 are called *autocorrelation* or *cross-correlation* terms, respectively, since the voltages multiplied in each term are from either the same or different elements, respectively.

From Equation 6.10, we see that any measurement with a large filled-aperture telescope can be understood as being a sum in which each term depends on contributions from only two of the N elements. As long as the contributions from each element arrive at the focus in phase, *there is no need for the elements to be physically contiguous.* Generalizing, each cross-correlation term $\langle \Delta V_i \Delta V_k \rangle$ in Equation 6.10 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 correlation (i.e., multiplying) receiver. Moreover, if the source properties are unchanging, there is no need to measure all pairs at the same time. A given large filled-aperture telescope with N elements has $N(N-1)/2$ pairs of elements, and these could be observed sequentially to "synthesize" a measurement by a large filled-aperture telescope. Alternatively, numerous pairs of antennas can be distributed to positions at mutual 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-element interferometer.

Figure 6.3 shows a schematic picture of a two-element interferometer separated by distance b, known as a *baseline.* We can measure this distance in units of the observing wavelength, λ. In terms of familiar units of length, $b = L/\lambda$, where L is the distance between antennas and λ is the wavelength in the same unit, e.g., meters. Both antennas observe a common position s_o located at an angle θ from the meridian. The projected separation of the two antennas towards s_o (i.e., from the perspective of the source) is $u = b \cos\theta$. In this example, an on-axis wavefront incident to both telescopes reaches antenna 2 first and the wavefront reaches antenna 1 a little later, having traversed an extra path length of $b \cdot s_o = b \sin\theta$. In other words, emission received by antenna 1 experiences a *geometrical delay* relative to that received by antenna 2, where the time equals $\tau_g = b \cdot s_o / c$. To compensate for the geometrical delay, an artificial delay can be inserted into the signal path of antenna 2 (e.g., with cables) so that the signals from both antennas arrive coherently (i.e., with the same phase) at the correlator.

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

V_2, can be written in terms of the product of the voltage response of antenna 1, V_1, and a phase delay factor sinusoidally varying as a function of angle, i.e.,

$$V_2 = V_1 e^{2\pi i(ul)}. \tag{6.11}$$

Expanding to two directions, we can introduce β, a direction on the sky orthogonal to α. Also, we define $m = \sin\beta$ as the small angle analog to l in this new direction, and $v = b\cos\varphi$ where φ is the angle of the position s_o on the sky from the reference position orthogonal to θ. Finally, we define y as the extra path length introduced in this new direction, in units of the wavelength of emission, i.e., $y = v\sin\beta = vm$. With these changes, the two-directional voltage response of antenna 2 is:

$$V_2 = V_1 e^{2\pi i(ul+vm)}. \tag{6.12}$$

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

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

$$\langle V_1 V_2\rangle = \langle \iint V_1(l,m)dldm \iint V_2(l,m)dldm\rangle \tag{6.13}$$

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 6.13 can be simplified to:

$$\langle V_1 V_2\rangle = \langle \iint V_1(l,m)V_2(l,m)dldm\rangle \tag{6.14}$$

$$\langle V_1 V_2\rangle = \iint \langle V_1(l,m)V_2(l,m)\rangle dldm \tag{6.15}$$

$$\langle V_1 V_2\rangle = \iint \langle V_1(l,m)^2\rangle e^{2\pi i(ul+vm)}dldm \tag{6.16}$$

$$\langle V_1 V_2\rangle \propto \iint I(l,m)e^{2\pi i(ul+vm)}dldm \tag{6.17}$$

The correlator therefore measures a quantity known as the *complex visibility*, $\mathcal{V}$, which is formally the Fourier transform of its intensity distribution on the sky:

$$\mathcal{V}(u,v) = \iint I(l,m)e^{2\pi i(ul+vm)}dldm = Ae^{i\phi}. \tag{6.18}$$

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

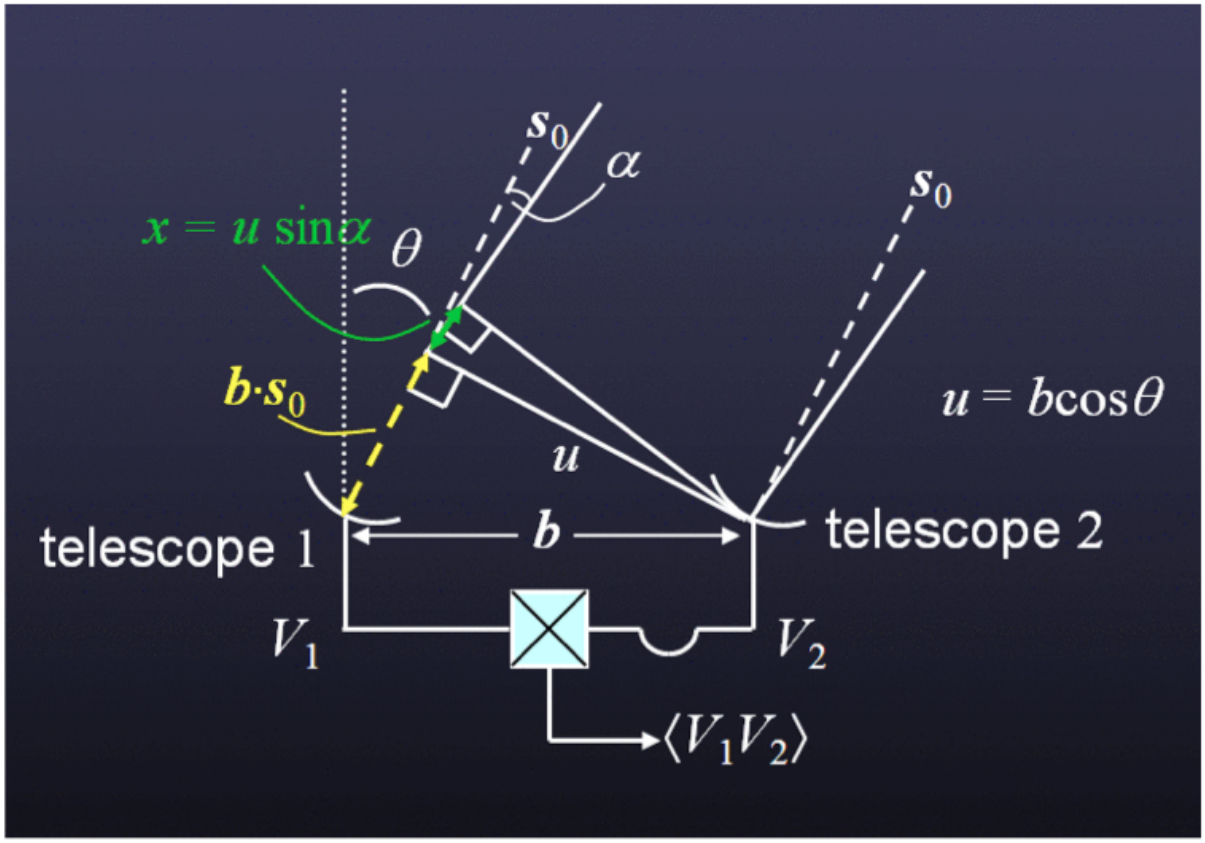

Figure 6.3: An ideal 1-D two-element interferometer consisting of two antennas, 1 and 2, separated by physical distance (i.e., a baseline) b. The antennas are both pointed towards a sky location given by s_o, which is at an angle θ from the meridian. The projected distance between the two antennas in that direction is thus $u = b\cos\theta$. The two antennas are connected to a correlator where the voltages detected from each are combined.

6.4 The Visibility or uv Plane

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

$$\mathcal{V}(u,v) = \iint I(l,m)e^{2\pi i(ul+vm)}dldm \tag{6.19}$$

$$I(l,m) = \iint \mathcal{V}(u,v)e^{-2\pi i(ul+vm)}dudv \tag{6.20}$$

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

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

$$\text{Fringe Spacing} = 1/u \text{ (radians)} = 1/(b\cos\theta) = \lambda/(L\cos\theta). \tag{6.21}$$

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 those spaced further apart measure emission on smaller scales. (Since fringe spacing also depends on the wavelength of emission, as b is measured in numbers of wavelengths, observing shorter or longer wavelengths also can sample smaller or larger scales, respectively.) These ideas can be easily generalized to two dimensions, with the fringe spacing and on-sky orientation depending on the relative magnitudes of u and v.

A given pair of antennas will only instantaneously sample a single scale of the sky brightness distribution. Given the E-W and N-S separations of the pair, a visibility in the uv-plane is measured. Since visibilities are actually Hermetian complex numbers, a single sampling gives two visibilities, one at (u,v) and its complex conjugate at $(-u,-v)$. To recover the true sky brightness distribution, however, knowledge of the distribution of visibilities across the uv-plane is needed. Improving coverage of visibilities over the uv-plane can be done in several ways. First, multiple antennas can be incorporated into an array, with each at 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 moves across the sky due to the Earth's rotation. Though the physical distances between the antennas does not change, its 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 visibilities across the uv-plane. Finally, antennas in the array may be arrangeable in several configurations so that pairs of antennas have different distances and can sample different parts of the uv-plane. The combination of these schemes can reasonably sample the uv-plane, yielding an image that can resemble the true sky brightness distribution.

6.5 Fields-of-View and Mosaics

Note that we have described an interferometer with idealized elements. Each element of an actual interferometer is, however, an antenna of finite diameter. As described in the preceding section, such antennas have their own

power response on the sky P_N (i.e., the Airy function = $\mathcal{A}(l,m)$). Indeed, the individual antenna response fundamentally limits the extent of an interferometric image made with a single pointing; in practice, the HPBW of the primary beam serves as the "field-of-view" of the single-pointing interferometric image. Moreover, $\mathcal{A}(l,m)$ is actually included formally in the correlator output:

$$\mathcal{V}(u,v) = \iint \mathcal{A}(l,m) I(l,m) e^{-2\pi i(ul+vm)} dl dm. \tag{6.22}$$

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 objects located far from the phase center.)

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

6.6 Spatial Filtering

Though the principles described in Sections 6.3-6.5 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 shortes baselines (i.e., lower than those sampled by the smallest baselines) results in an intrinsic lack of sensitivity to large-scale emission. It is important for ALMA users to understand these limitations.

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

$$B(u,v) = \sum_{k=1}^{2M} \delta(u-u_k, v-v_k). \tag{6.23}$$

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

$$I^D(l,m) = FT^{-1}\{B(u,v)\mathcal{V}(u,v)\}. \tag{6.24}$$

Using the convolution theorem, we can rewrite Equation 6.24 as:

$$I^D(l,m) = b(l,m) * I(l,m)\mathcal{A}(l,m). \tag{6.25}$$

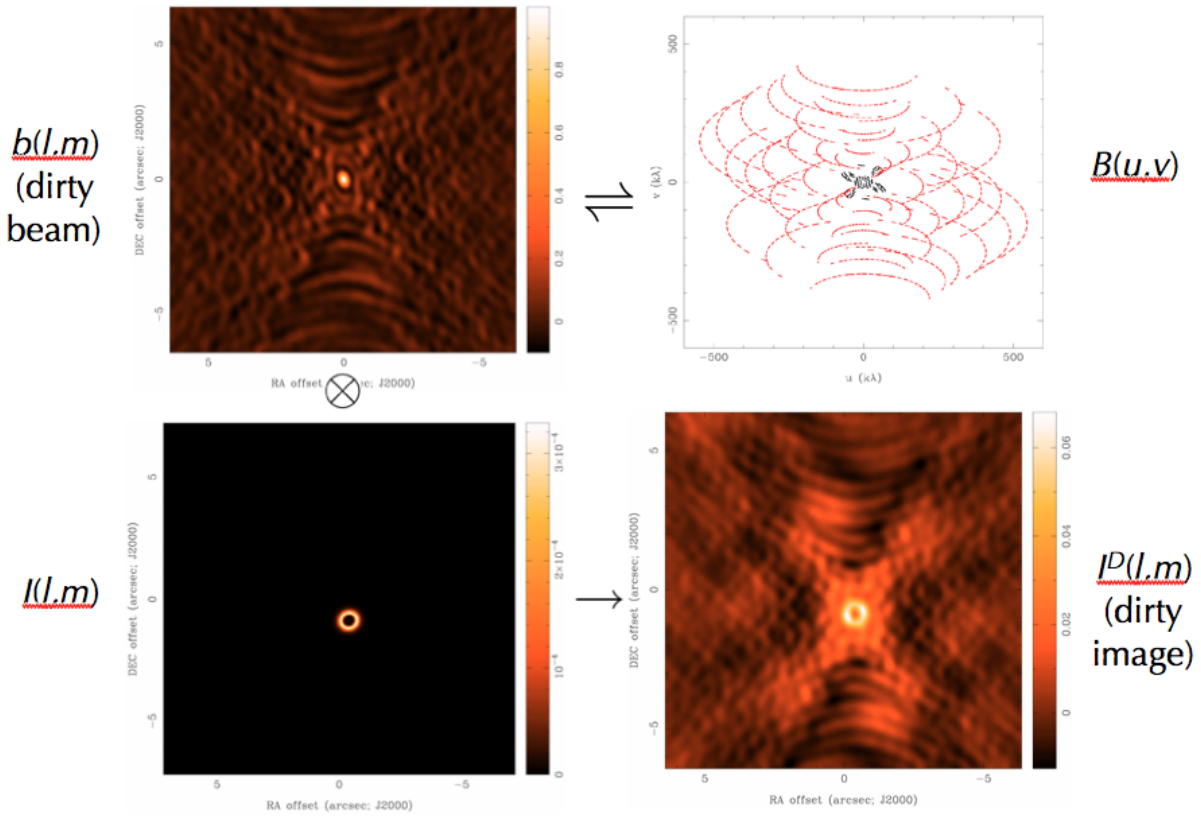

Figure 6.4: Imaging concepts. *Panel a (upper left)* Example of a dirty beam, $b(l, m)$. *Panel b (upper right)* The related ensemble of discrete points sampled in the uv-plane, $B(u, v)$. *Panel c (lower left)* Example of a true sky distribution, $I(l, m)$. *Panel d (lower right)*, The dirty image $I^D(l, m)$ resulting from observing $I(l, m)$ over the baselines of $B(u, v)$, or equivalently the convolution of $I(l, m)$ by $b(l, m)$. (We have ignored the antenna power response, $\mathcal{A}(l, m)$, in this illustration since it is much wider than the true sky brightness distribution.)

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

The measure of how similar an image is to the true sky distribution is sometimes referred to as *image fidelity*. (Formally, image fidelity is the ratio of the maximum brightness of an image to its 1 σ rms noise level.) Image fidelity depends on the specifics of coverage of the uv-plane sampled by the interferometer. Since the numbers of samples are necessarily finite and discrete, there are invariably gaps in any practical sampling of the uv-plane. These gaps mean that no information is obtained about the true sky brightness distribution on 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 (e.g., CLEAN and its variants) to minimize the effect of incomplete spatial frequency sampling (see next Chapter).

Though we have been speaking generally about true sky brightness distributions so far, a special note should be made for the case of point sources. Obviously, a point source is a distribution of emission that is not extended relative to the resolution of the observation. In this case, the morphology of the dirty image will equal that of the dirty beam. Moreover, the complex visibilities of the point source have the same amplitudes on all angular scales; no spatial filtering occurs. Of course, sources may appear point-like at low resolutions but may appear extended at higher-resolution observations.

Figure 6.4 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. 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*. (Note that two sets of tracks are identified in panel *b*; these are due to observations of the same antennas in two different configurations.) Panels *c* and *d* (lower pair) show respectively an example of a true sky distribution (actually, a model of a ring of emission) and the dirty image. The dirty image is the convolution of the true sky distribution by the dirty beam, and we can easily see how incomplete sampling of the uv-plane leads to the appearance of significant artifacts in the resulting dirty image.

The resolution of the dirty image is defined effectively by the compactness of the central feature of the dirty beam, e.g., half its FWBN. Since the structure of the dirty beam is generally more complicated than that of a single-dish antenna (e.g., the idealized beam shown in Figure 6.2), it is not so easy to measure FWBN. Instead, the resolution is typically approximated (i.e., obtained 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, broad distributions in one domain have narrow ones in the other, and vice versa. By analogy, a wide sampling of discrete points in the uv-plane by an interferometer, $B(u,v)$, yields a compact central beam feature in $b(l,m)$. Note that resolution is thus fundamentally limited by the extent of the longest baselines in a given configuration; the minimum scale discernible in the image (the resolution) is limited by these maximum baselines. A handy formula for the approximate resolution provided by an interferometer is:

$$\text{Interferometer Resolution} = \theta_{res} = k\,\lambda/L_{max}, \tag{6.26}$$

where k is a factor that depends on the imaging weighting function (typically $\sim 1''$; see Figure 7.8) and L_{max} is the longest baseline in the array.

Another important limitation of interferometric array observations is insensitivity to large angular structures. 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. In practice, the interferometer image has a *maximum recoverable scale* given by:

$$\text{Maximum Recoverable Scale} = \theta_{MRS} \approx 0.6\lambda/L_{min}, \tag{6.27}$$

where L_{min} is the minimum baseline in the array configuration. The smallest baseline possible in an array occurs when two antennas are adjacent to each other. Of course, the antennas cannot be moved physically closer together than their diameters. Note that in projection antennas can appear closer than their diameters. In those cases, found typically when observing low elevation sources in compact configurations, the antenna in front blocks the reception by the antenna in the rear, causing the latter to have reduced sensitivity. This situation is called *shadowing* and affected data are typically removed from the ensemble of observed visibilities.

Figure 6.5 illustrates the idea of spatial filtering using simulations of actual ALMA configurations. Here, panel *a* shows an optical image of the galaxy M51 we use as an example of a true sky brightness distribution, after changing the frequency of the image to 100 GHz and placing the image center at δ = -40°, a declination easily observable with ALMA. Panels *b*, *c*, and *d* show the recovered images obtained by simulating ALMA Cycle 1 (32-antenna) observations of the galaxy (using *simobserve*) only in the C32-6, C32-4, and C32-1 configurations respectively (and CLEANing). Angular resolutions of $0.55''$, $1.1''$, and $3.6''$ are obtained, respectively. (The corresponding maximum recoverable scales are $6.6''$, $14.4''$, and $36.4''$, respectively.) As can be seen, larger scale emission from the galaxy has been filtered out in the highest resolution C32-6 observations (panel *b*), leaving only the compact structures of its arms. On the other hand, these compact structures are not very discernible in the lowest resolution C32-1 observations (panel *d*). A reasonable compromise is found in the middle-resolution C32-4 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 (as available in later ALMA Cycles), would increase the fidelity of recovered images.

6.7 Multi-configuration Observations

As seen above, a given configuration with baselines ranging from L_{min} to L_{max} is sensitive to angular scales from θ_{MRS} to θ_{res}. As discussed above, sensitivity to a broader range of angular scales is possible by combining data obtained in different 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 the highest angular resolution. For sensitivity to extended structures, the more compact 12-m Array configurations can be used, or the more tightly clustered 7-m Array. Finally, the individual Total Power Array (TPA) antennas of the ACA can be used to map the largest angular structures and address the zero-spacing problem noted above. 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.

Data combination appears to work best when the signal-to-noise ratios (SNR) of the datasets are similar. Otherwise, information on scales covered by lower SNR data is relatively less reliable, making interpretation of the images difficult. Assuming the SNR of the individual datasets is high, combination is best done in the visibility domain rather than the image domain to minimize the effect of artifacts produced by aliasing (i.e., incomplete spatial frequency 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

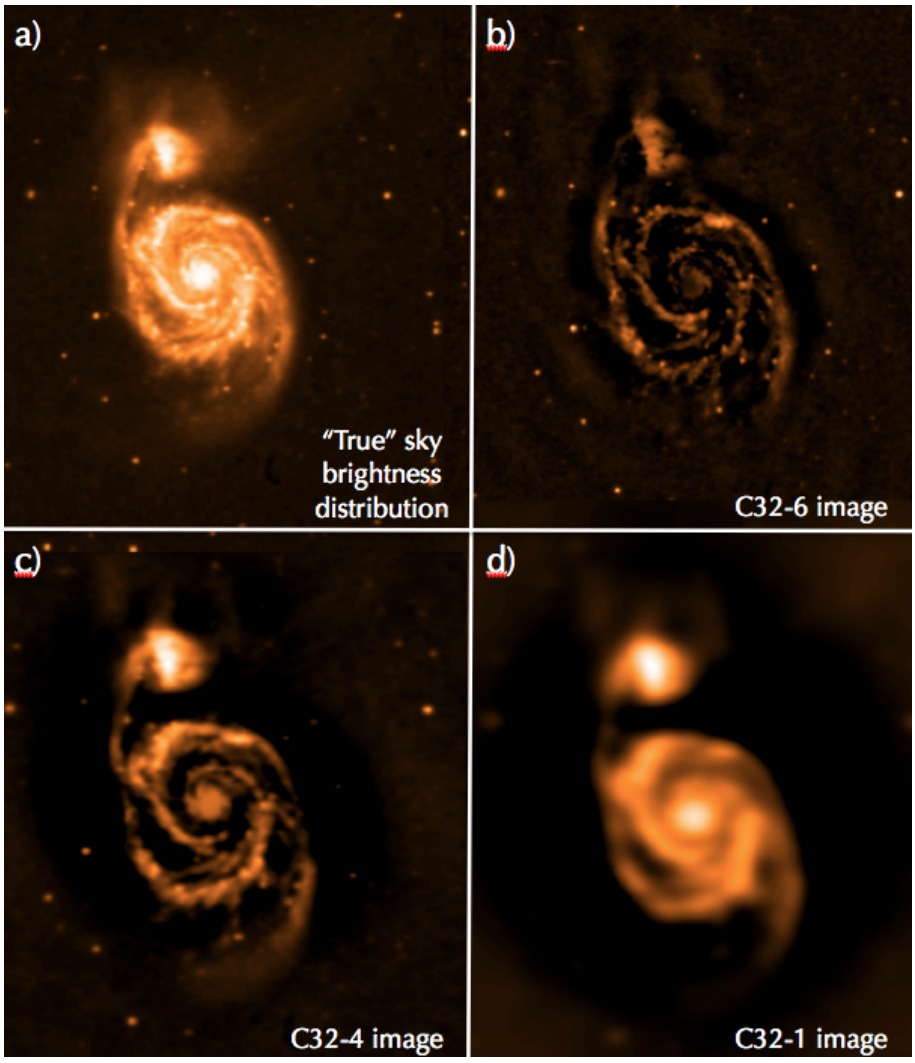

Figure 6.5: Examples of spatial filtering using *simobserve* and actual ALMA configurations. *Panel a (upper left)* An optical image of the galaxy M51 used as a template for a true sky brightness distribution for the simulations. The frequency of the emission has been changed to 100 GHz, and its declination has been changed to -40° to allow ALMA observations to be simulated. For the simulations, the galaxy was "observed" over a mosaic of 39 pointings, for 10 hours in total. The resulting dirty images were CLEANed. *Panel b (upper right)* The high-resolution image of the galaxy obtained when observed in the ALMA Cycle 1 C32-6 configuration. The resulting synthesized beam is $\sim 0.55''$ and the maximum recoverable scale is $6.6''$. *Panel c (lower left)* Medium-resolution image of the galaxy when observed in the ALMA Cycle 1 C32-4 configuration. The resulting synthesized beam is $\sim 1.1''$ and the maximum recoverable scale is $14.4''$. *Panel d (lower right)* Low-resolution image of the galaxy when observed in the ALMA Cycle 1 C32-1 configuration. The resulting synthesized beam is $\sim 3.8''$ and the maximum recoverable scale is $36.4''$.

Fourier transformed to produce a new (dirty) image. (Note that Figure 6.4b shows two distinct sets of samples of the *uv*-plane, one obtained with a compact configuration (in black) and one obtained with a more extended one (in red).)

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 a significant range of spatial frequencies in common. For example, a reasonable overlap in spatial frequency coverage allows (amplitude) calibration differences in between datasets to be minimized through re-scaling of the single-dish visibilities vs. the interferometric ones. As a rule of thumb, a reasonable overlap of spatial frequencies will occur if the single-dish data are obtained by an antenna twice the diameter of those used to obtained the interferometer data. Multiple interferometer pointings (i.e., mosaics), can also partially recover missing low spatial frequency information.

6.8 Units and Conversions

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

$$I_\nu(\theta,\varphi) = \frac{2k\nu^2}{c^2} T_B(\theta,\varphi). \tag{6.28}$$

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 6.2), where

$$S_\nu = \frac{2k\nu^2}{c^2} \int T_B d\Omega. \tag{6.29}$$

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

$$\left(\frac{T}{1\text{ K}}\right) = \left(\frac{S_\nu}{1\text{ Jy}}\right) \left[13.6 \left(\frac{300\text{ GHz}}{\nu}\right)^2 \left(\frac{1''}{\theta_{max}}\right) \left(\frac{1''}{\theta_{min}}\right)\right]. \tag{6.30}$$

In millimeter/radio astronomy, the power received from astronomical sources is typically so weak that a convenient unit to use for flux density is the *Jansky* (Jy), where 1 Jy = 10^{-26} W m^{-2} Hz^{-1} = 10^{-23} erg s^{-1} cm^{-2} Hz^{-1}. The ALMA Observing Tool (OT) 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 6.30, we see that an rms value in flux density (ΔS) can translate to an rms value in brightness temperature (ΔT), assuming a given synthesized beam size. Larger beam sizes correspond to lower ΔT, i.e., the surface brightness sensitivity increases. In turn, extended low surface brightness objects may be harder to detect at higher angular resolutions as the corresponding sensitivities may be too low. Typically, a compromise must be obtained between angular resolution and brightness sensitivity when planning interferometric observations.

Chapter 7

Imaging with ALMA

7.1 Introduction

As described in Chapter 6, and derived from the van Cittert-Zernike theorem (e.g., Rohls & Wilson 2004), there is a fundamental relationship among 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 will be constituted of discrete points in the uv-plane $\mathcal{V}_k(u_k, v_k)$ with k = 0.. N, and N reaching a few million. In addition to the visibilities $\mathcal{V}_k$ obtained, a weighted visibility function $\mathcal{V}^W$ can be defined (see Briggs, Swab & Sramek, 1999) to control better the synthesized beam and the sensitivity:

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

In Equation 7.2, R_k, T_k, and D_k are weights. R_k indicates the reliability of the k^{th} visibility measurement based on certain factors, e.g., the integration time, the system temperature, or the band width. (Note that these factors are external and not controlled by the imaging process.) T_k is used to taper the outer edge of the uv-coverage and suppress small-scale sidelobes and increase the beam width. Finally, D_k is used to offset the high concentration of measurements near the center of the uv-plane and to reduce the sidelobes caused by gaps in the uv-coverage.

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

- $D_k = 1$ is called natural weighting, and is based primarily on the density of visibilities in the uv-plane. Natural weighting yields maximum sensitivity and good angular resolution.
- $D_k = \frac{1}{N_s(k)}$ 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 on the density of visibilities in uv-plane, and gives the same weight to all visibilities. Uniform weighting allows maximum angular resolution, at the expense of sensitivity.

In Cycle 2, an ALMA project will consist of data from one or two 12-m Array configurations, and possibly additional data from the 7-m Array and the Total Power Array (TPA), both forming the Atacama Compact Array (ACA). Imaging then consists of combining all acquired data and using the most appropriate weighting

function to achieve the scientific goals. We describe in this chapter some of aspects of imaging to consider and obtain the best results from ALMA observations.

During Cycle 2, the 12-m Array, with 34 antennas, will be arranged in 7 different configurations. The maximum baselines achievable will be 1.5 km for Bands 3-6 and 1.0 km for Bands 8 and 9. The 7-m Array, with nine antennas, will have two slightly different configurations, a standard one and an alternate for low-elevation targets. Tables 7.1 and 7.2 give the basic properties of the 12-m Array and 7-m Array configurations. In addition, Figures 7.1 and 7.2 show representative antenna locations for these configurations of the 12-m Array and the 7-m Array. The two single-dish 12-m antennas of the TPA will be in a fixed position. Only representative configurations are shown in Figures 7.1 and 7.2 because configurations cannot be selected directly in the OT. Projects will be observed using array configurations chosen to achieve their goals of angular resolution (θ_{res}) and maximum recoverable scale (θ_{MRS}).

	Band	3	4	6	7	8	9
	Frequency (GHz)	100	150	230	345	460.	650
Configuration							
7-m	θ_{res} (arcsec)	19.26	12.84	8.38	5.58	4.19	2.96
	θ_{MRS} (arcsec)	41.7	27.8	18.1	12.1	9.1	6.4
7-m-NS	θ_{res} (arcsec)	18.80	12.53	8.17	5.45	4.09	2.89
	θ_{MRS} (arcsec)	44.2	29.4	19.2	12.8	9.6	6.8
C34-1	θ_{res} (arcsec)	3.73	2.49	1.62	1.08	0.81	0.57
	θ_{MRS} (arcsec)	26.1	17.4	11.4	7.6	5.7	4.0
C34-2	θ_{res} (arcsec)	2.04	1.36	0.89	0.59	0.44	0.31
	θ_{MRS} (arcsec)	26.3	17.5	11.4	7.6	5.7	4.0
C34-3	θ_{res} (arcsec)	1.40	0.93	0.61	0.40	0.30	0.21
	θ_{MRS} (arcsec)	18.0	12.0	7.8	5.2	3.9	2.8
C34-4	θ_{res} (arcsec)	1.11	0.74	0.48	0.32	0.24	0.17
	θ_{MRS} (arcsec)	18.0	12.0	7.8	5.2	3.9	2.8
C34-5	θ_{res} (arcsec)	0.75	0.50	0.33	0.22	0.16	0.12
	θ_{MRS} (arcsec)	14.4	9.6	6.3	4.2	3.1	2.2
C34-6	θ_{res} (arcsec)	0.57	0.38	0.25	0.16	0.12	0.09
	θ_{MRS} (arcsec)	9.1	6.1	4.0	2.6	2.0	1.4
C34-7	θ_{res} (arcsec)	0.41	0.27	0.18	0.12	-	-
	θ_{MRS} (arcsec)	9.1	6.1	4.0	2.6	-	-

Table 7.1: θ_{res} and θ_{MRS} for the 7-m Array and 12-m Array configurations available during Cycle 2 according to (a representative frequency in a) Band. The values of θ_{res} and θ_{MRS} are computed using minimum and maximum baselines from Table 7.2.

Configuration	7-m	7-m-NS	C34-1	C34-2	C34-3	C34-4	C34-5	C34-6	C34-7
Minimum baseline (meter)	8.9	8.4	14.2	14.1	20.6	20.6	25.8	40.6	40.6
Maximum baseline (meter)	32.1	32.9	165.6	303.6	442.7	558.2	820.2	1091.0	1507.9
RMS (meter)	19.5	20.9	81.4	127.3	186.6	241.8	341.3	502.8	645.3

Table 7.2: Basic parameters of the two 7-m Array configurations and the seven 12-m Array configurations during Cycle 2. The baselines are projected for a transiting source at Dec = -23. Note that C34-7 will not be available for Bands 8 or 9.

7.2 Shadowing and beam shape

During observations, particularly in very compact configurations like those of the 7-m Array or C34-1, one antenna can partially obscure another, making the data from that baseline unusable. This phenomenon, also known as shadowing, can be quite severe when the sources are very low in elevation. For example, Figure

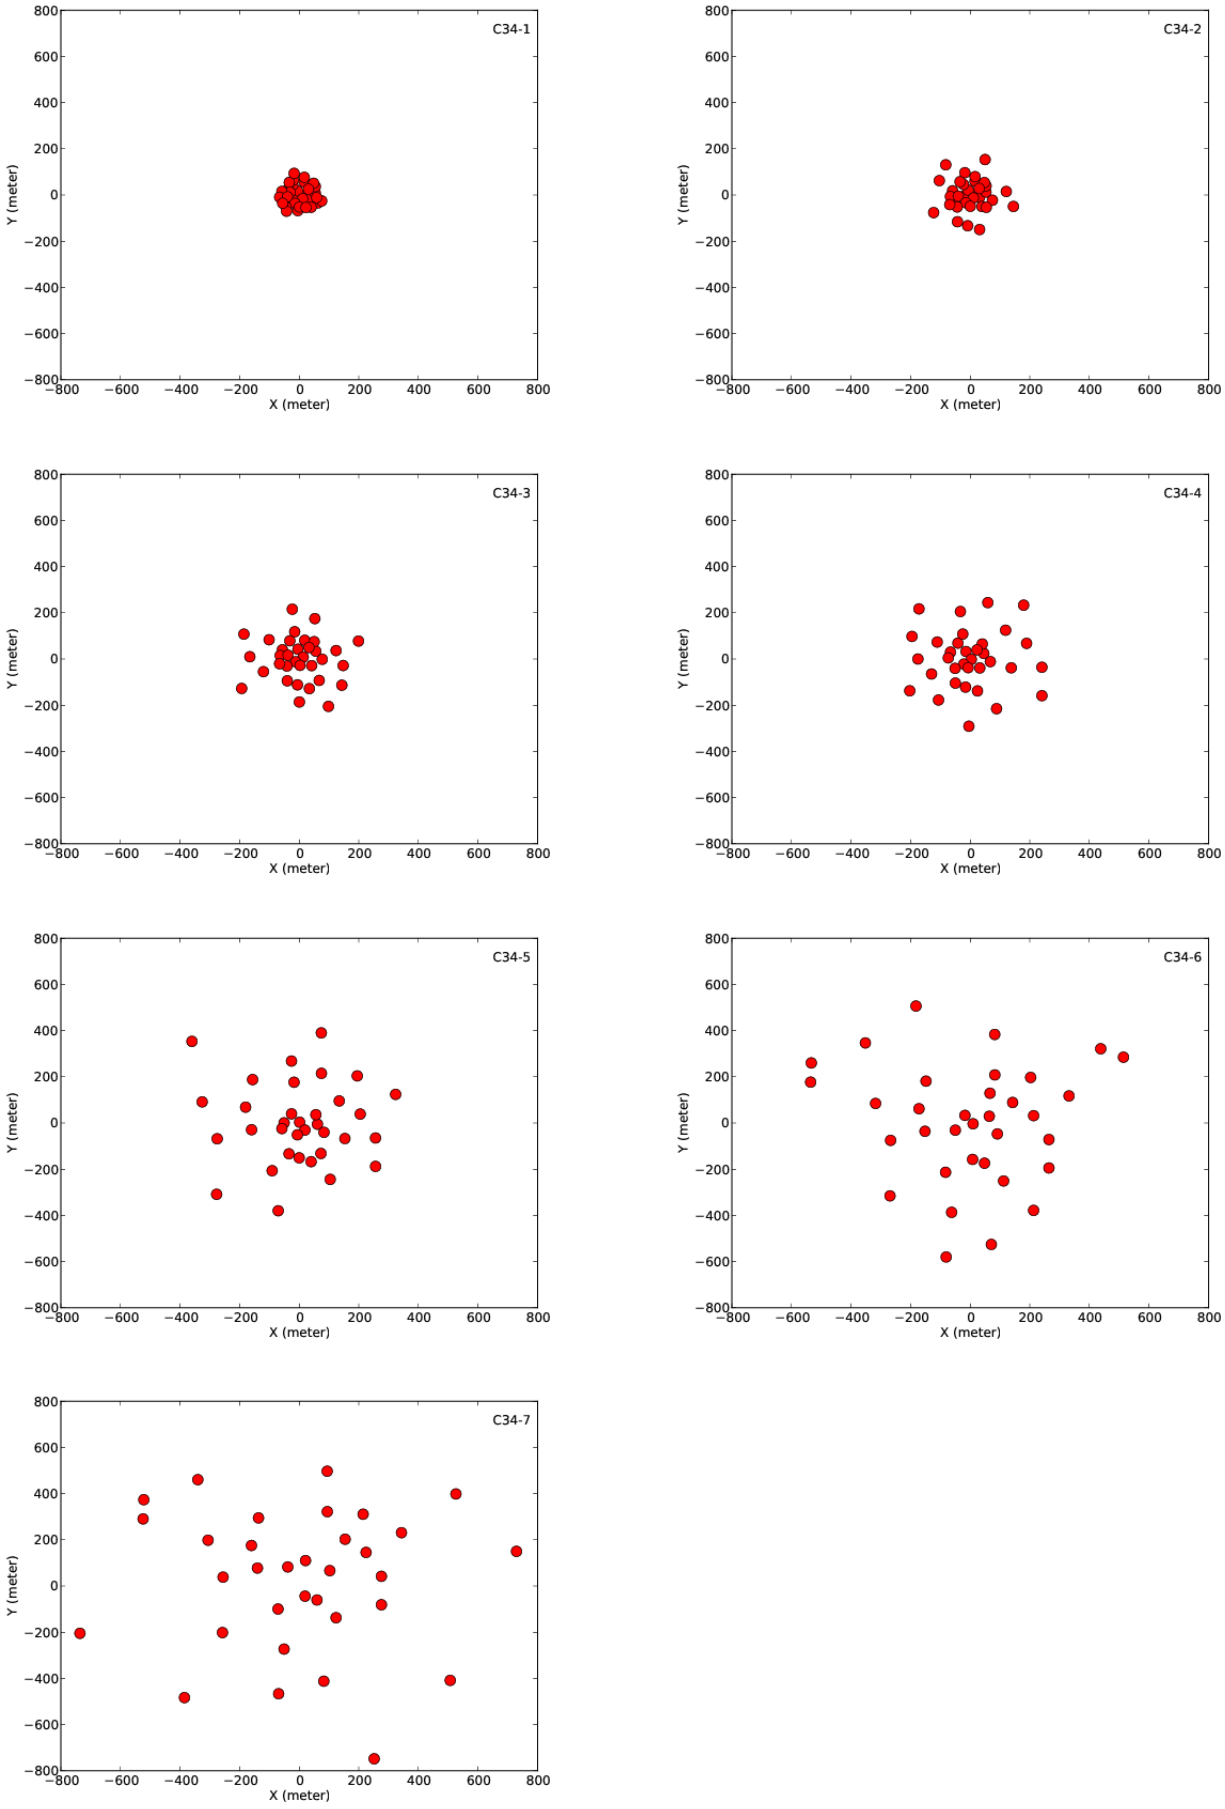

Figure 7.1: Representative configurations for Cycle 2: 12-m Array

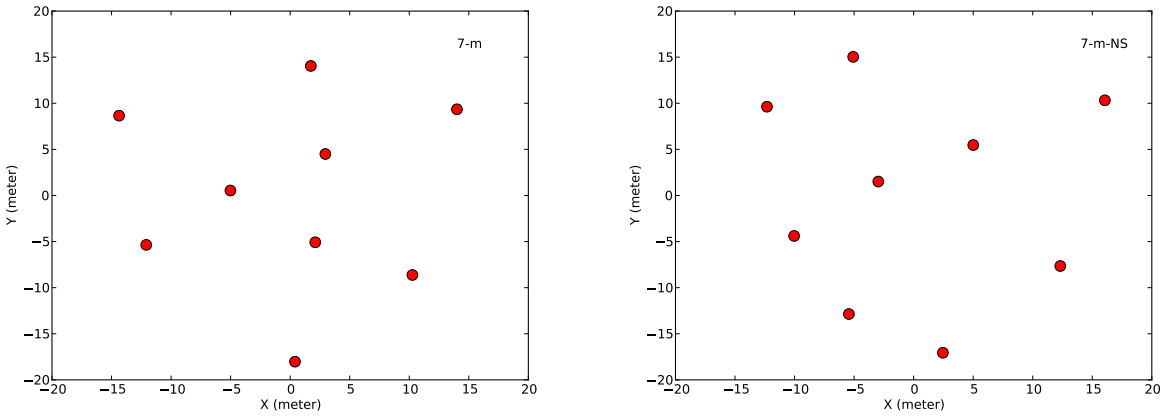

Figure 7.2: Representative configurations for Cycle 2: 7-m array (*left*) and 7-m-NS array (for Northern/Southern low elevation sources; *right*).

7.3 shows the percentage of data that will be shadowed (the "shadowing fraction") when sources of various declination are observed in the most compact ALMA configurations. As can be seen, the shadowing fraction can be as large as 40% for sources observed with the 7-m Array or the most compact 12-m Array configuration (C34-1). Note that the shadowing fraction of 7-m Array data of sources at extreme declinations can be reduced by up to a factor of 2 by using the 7-m-NS configuration. Using 7-m-NS array has some drawbacks, however. For example, Table 7.3 shows the sidelobe levels resulting from all configurations, and reveals that the sidelobe level from the 7-m-NS configuration is markedly worse than that for 7-m Array (39.4% vs. 24.0% for 1 hour integration). Nevertheless, a 3 hour integration around the transit would reduce the sidelobe level to 30.7% (18.8% for 7-m Array).

Configuration	7-m	7-m-NS	C34-1	C34-2	C34-3	C34-4	C34-5	C34-6	C34-7
Natural	23.2%	43.8%	3.9%	3.7%	8.8%	5.4%	12.8%	11.2%	12.4%
Briggs ($R = 0.5$)	24.0%	39.4%	8.9%	6.1%	7.2%	7.6%	12.0%	8.5%	11.5%
Uniform	30.2%	39.9%	16.8%	14.3%	17.1%	13.9%	16.0%	8.5%	10.9%

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

The synthesized beam shape, which is the Fourier transform of the *uv*-plane sampling during the observation(s), is affected by source elevation because of the decrease of projected baselines in the north-south direction. The shape of the *uv*-plane sampling distribution becomes more elongated and consequently the beam shape is more elliptical. For example, Figure 7.4 shows the different beam shapes for sources observed at declinations of -70 degrees (similar to the SMC and LMC) and -30 degrees. Also, Figure 7.5 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 source. The ALMA user should keep in mind that during the imaging process it is possible to use in CASA an ad hoc isotropic restoring beam.

The angular resolution θ_{res} (see Equation 6.26) of an array can be estimated with the following equation but may differ slightly from the final spatial resolution after the deconvolution of the uv plane ("cleaning"):

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

$$\approx \frac{61800}{L_{max}\nu} \; [arcseconds] \tag{7.4}$$

where where k is a factor $\sim$1 that depends on the imaging weighting function, L_{max} is the longest baseline in

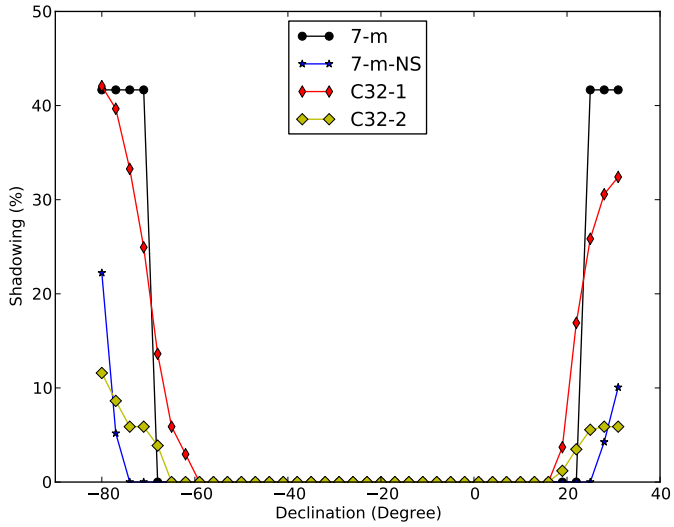

Figure 7.3: Shadowing fraction (%) for the most compact configurations with respect to declination.

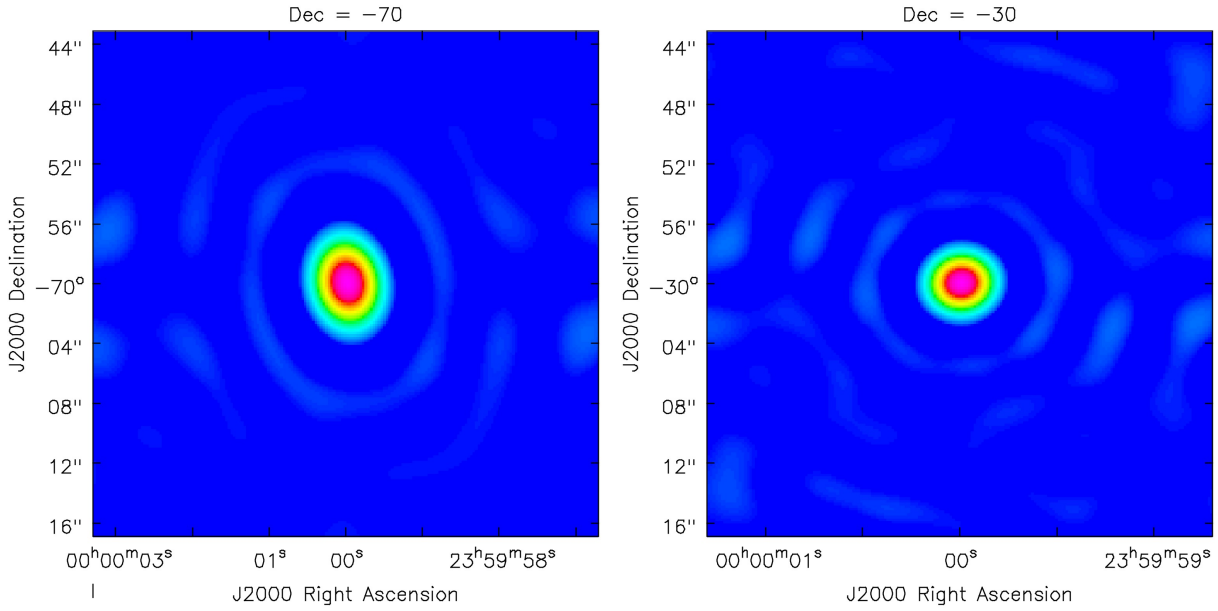

Figure 7.4: Beam shape for the configuration C34-1 with a 2 hour observation of a transiting source at a declination of either -70 degrees (*left*) or -30 degrees (*right*).

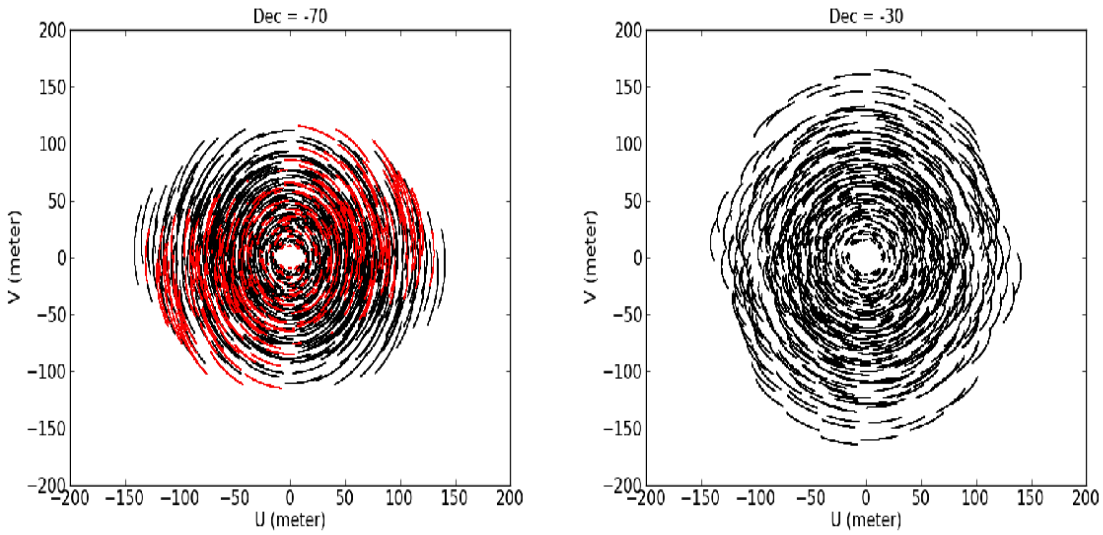

Figure 7.5: *uv*-plane coverage for the configuration C34-1 with a 2 hour observation of a transiting source at a declination of either -70 degrees (*left*) or -30 degrees (*right*). For the source at decl. = -70 degrees, the red color indicates shadowed measurements (21% of the total).

meters, λ is the observing wavelength in meters, and ν the observing frequency in GHz.

Figure 7.6 shows the minor and major axes of the synthesized beam widths (θ_{res}) in the different array configurations and source declinations for a 1 hour observation at 100 GHz. Furthermore, Figure 7.7 shows the geometrical mean of the major and minor axes of θ_{res}. The beam width scales with λ (though Band 8 or 9 projects will not be able to use the C34-7 configuration in Cycle 2 due to its largest 1.5 km extent).

An important consideration when imaging is the weighting scheme used on the visibilities when Fourier transforming them. The angular resolution indicated here for ALMA were computed using CASA simulations with natural weighting. Uniform weighting, however, yields the highest angular resolution by setting the same weight to all the visibilities, regardless of their density, but the final sensitivity is degraded. Natural weighting yields lower angular resolution than uniform weighting since it takes into account the densities of visibilities in the *uv*-plane, and these are typically higher at smaller distances from the origin in the *uv*-plane due to greater numbers of smaller baselines in a given observation. Bridging the extremes of natural and uniform weighting, Briggs (1995) defined a continuous scheme that describes weighting using a robust parameter R. Uniform weighting is close to R = -2 and natural weighting is close to R = 2. Figure 7.8 shows the angular resolution achieved for an observation with the C34-4 configuration at 100 GHz using R between -2 and 2, the range allowed in CASA. As can be seen, the angular resolution varies from 0.87″ (R = -2) to 1.36″ (R = 2). Note that in the OT, angular resolution is computed using a uniform weighting but sensitivity is computed using a natural weighting. Thus, the final angular resolution achieved for a project with nominal sensitivity may be slightly degraded by the use of natural weighting. The most appropriate value of R to use when imaging a dataset must be determined empirically, following the goals of the scientific case (angular resolution, sensitivity, etc.)

7.3 Response and snapshot

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

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

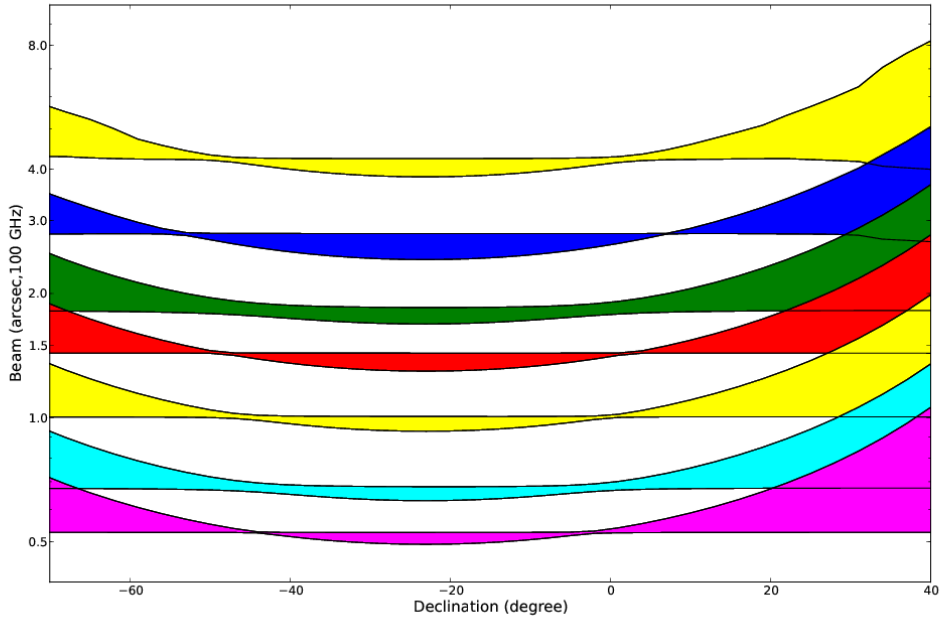

Figure 7.6: Beam shapes for Cycle 2: Colors fill the ranges of the major and minor axes of synthesized beams expected from a 1 hour observation at 100 GHz in each configuration.

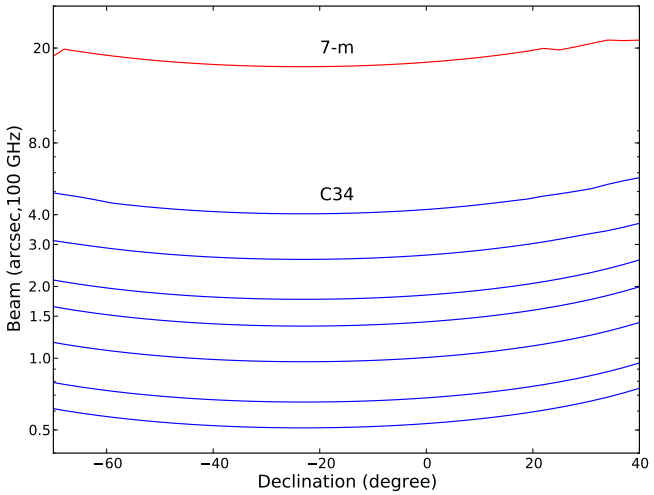

Figure 7.7: Geometrical means of the major and minor axes of the synthesized beams expected for 1 hour observations at 100 GHz for Cycle 2 configurations.

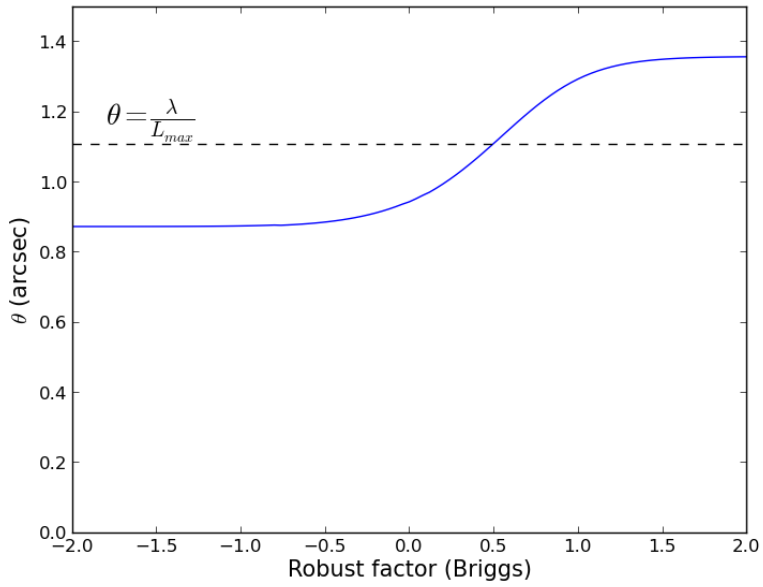

Figure 7.8: Angular resolution achieved using different values of the robust parameter R for a 1 hour observation at 100 GHz in the C34-4 configuration. Note that a R = -2 is close to uniform weighting and R = 2 is close to natural weighting.

If S were a smooth function, e.g., a Gaussian, the synthesized beam would have just smooth wings. In practice, given the finite number of baselines, the sampling function is an ensemble of Dirac functions, and the synthesized beam has aliasing in the form of sidelobes. Observing a target as the Earth rotates can allow the *uv*-plane to be sampled more extensively, as the projected baselines change with source position. Short integrations are still valuable, 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 longer baselines, during the Fourier transform. Figures 7.9 and 7.10 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 C34-4. As shown on Figure 7.9, 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 7.10) are similar in term 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 reaches about 30% for the snapshot whereas it is only 6% for the 1 hour integration. The snapshot image will then require a comparatively more careful cleaning process 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 check directly the sidelobe fingerprint left by the strong point sources after deconvolution. It is worthwhile to use the interactive mode during the deconvolution so that residuals can be monitored.

7.4 Spatial scale Filtering

As described in Chapter 6, an interferometer detects Fourier components of the sky brightness distribution in an area of the *uv*-plane defined by the array configurations used to observe a target. Since the *uv*-plane is a

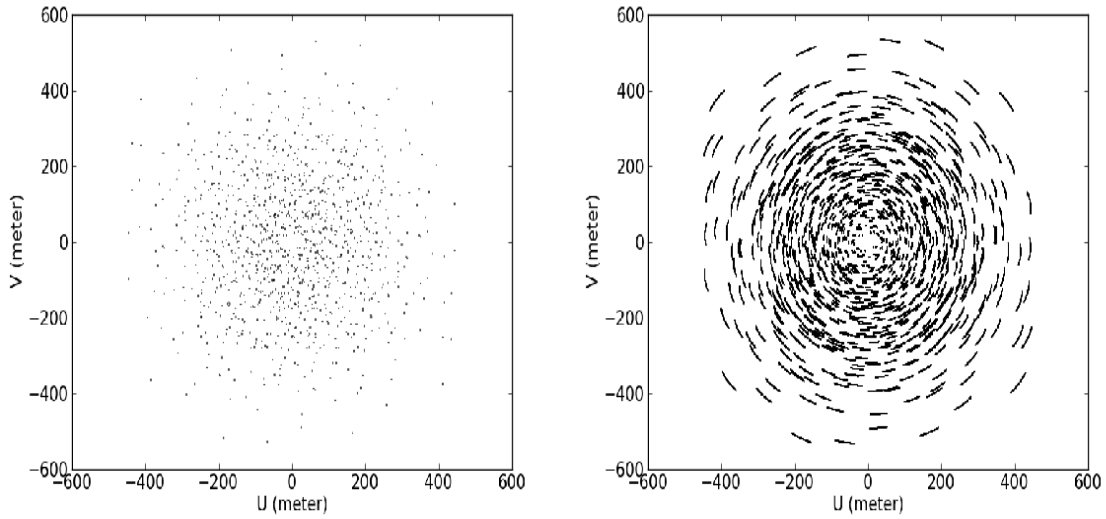

Figure 7.9: *uv*-plane sampling distributions of a model ALMA observation with 1 minute integration (*left*) and 1 hour integration (*right*), using the C34-4 configuration to observe a source at declination = -30 degrees.

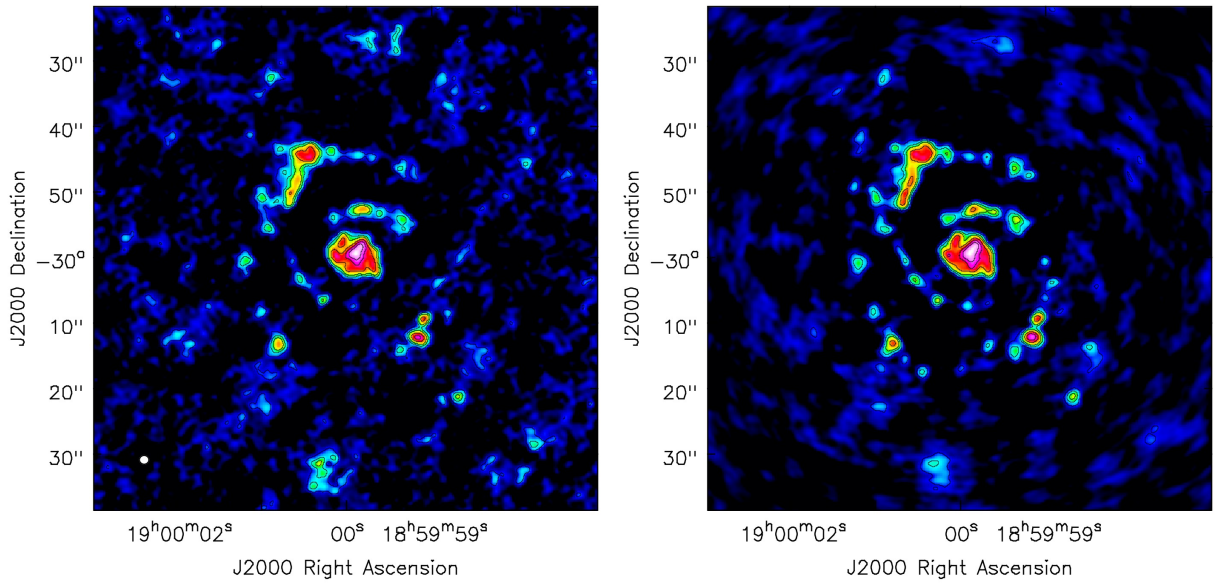

Figure 7.10: Images obtained from a model ALMA observation with 1 minute integration (*left*) and 1 hour integration (*right*), using the C34-4 configuration to observe a source at declination = -30 degrees. The black lines are for the levels 100, 200, 300, 500, 700 and 900 mJy beam^{-1}

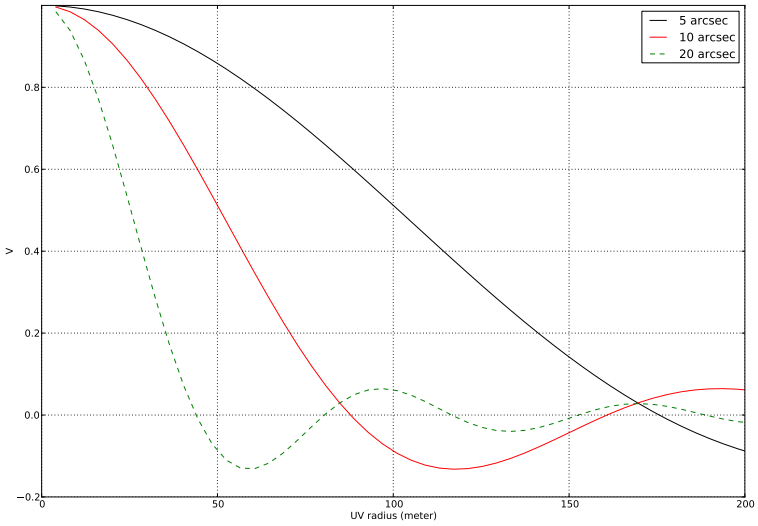

Figure 7.11: Expected visibilities of three model uniform disks at 100 GHz with *uv*-distance.

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 domains of the *uv*-plane sampled by a given configuration. To illustrate this concept, the Fourier transforms of three uniform disks with sizes of $5''$, $10''$ and $20''$ are shown in Figure 7.11 for an observation at 100 GHz. (For Figure 7.11, the amplitudes shown are those annularly averaged in the *uv*-plane, and are plotted against *uv-distance*, the length in the *uv* plane from the origin.) The smallest uniform disk is closest to a point source and so it has large amplitudes up to a *uv*-distance of 180 m. Meanwhile, the most extended disk has large amplitudes only up to ∼40 m. Therefore, an array with baselines larger than 40 m will not be sensitive to emission on angular scales larger than $\sim 20''$, and will filter out most of the emission from such an extended disk. An important consequence of such filtering is that an interferometer only detects a fraction of the flux. Only a single-dish observation, sensitive to the visibility at the origin of the *uv*-plane, can detect the total flux.

The maximum recoverable scale (θ_{MRS}, see Equation 6.27) measures the largest angular structure to which a given array is sensitive. The θ_{MRS} is proportional to the inverse of the Nyquist frequency (i.e., twice the minimum baseline of a given configuration). Since this frequency represents a very low sensitivity to the power from a source observed with the minimum baseline, it is usual to adopt a less stringent limit so that some power is detected at such small *uv*-distances. ALMA has adopted the 10% power sensitivity level, which corresponds to 1.2 times the inverse of the Nyquist frequency, yielding:

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

$$\approx \frac{37100}{L_{min}\nu}\ [arcseconds] \tag{7.7}$$

where L_{min} is the shortest baseline in meters, λ is the observing wavelength in meters, and ν the observing frequency in GHz. Table 7.1 lists the θ_{MRS} for the different Cycle 2 array configurations. In addition, Figure 7.12 shows the fractions of flux from a uniform disk lost according to angular scale for two array configurations, together with the analytical expectation for θ_{MRS}. The most compact configuration, C34-1, can measure up to 50% of the total flux for angular scales smaller than $20''$ at 100 GHz.

The spatial filtering of an interferometer observation is a serious issue that must be addressed carefully

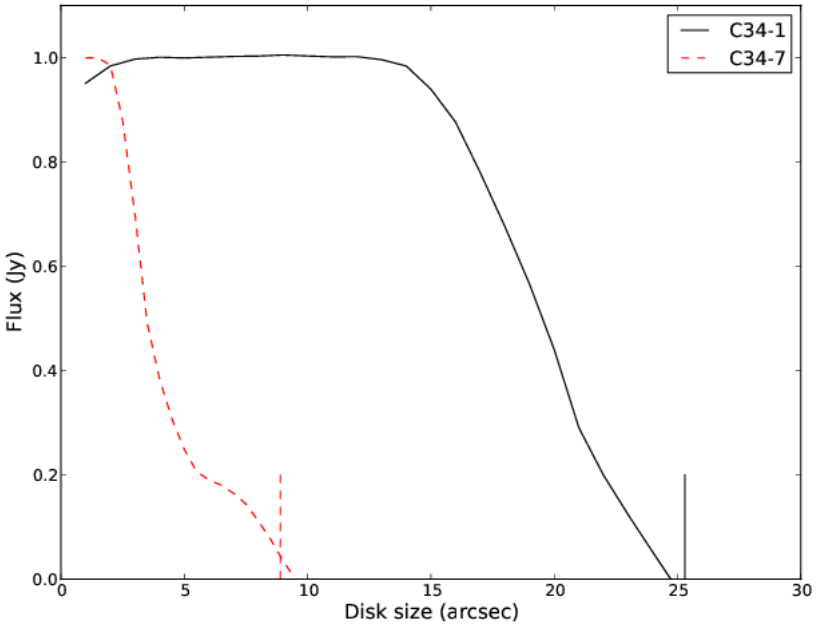

Figure 7.12: Fractions of flux expected from a uniform disk observed in two different array configurations, C34-1 and C34-7, at 100 GHz. The vertical ticks indicate the respective analytical expectation of θ_{MRS}.

for each scientific case. Ideally, the range of the *uv*-plane where a source emits should be known *a priori* so the appropriate configuration can be used during observations. Selecting a specific range of angular scales has the effect of discarding all emission from those angular scales not in that range. Eventually, if all scales are important, the combination of the TPA, 7-m Array, and 12-m Array should be used, as discussed in Sect. 7.6. Figure 7.13 shows the visibility amplitudes with *uv*-distance of a typical astronomical source, M51, observed in Hα but which we use as indicative of 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, indicating that most of the power is located in larger scale structures, and that that power decreases rapidly with at smaller scales. This result shows that the flux that would be received by the 7-m Array would be much brighter than that received by the 12-m Array with extended configurations (e.g., C34-7). 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 example of M51, the visibility amplitude detected by the 7-m Array (at $k_{uv} \approx 5k\lambda$) is $\sim$23 Jy, whereas that detected by the configuration C34-7 ($k_{uv} \approx 200k\lambda$) is only $\sim$0.9 Jy. Ideally, ALMA users should estimate through simulations the distribution of power at different different length scales for targets they wish to observe.

7.5 Mosaicing

The field-of-view of a single interferometer pointing is given by the HPBW of the primary beam, i.e., $1.02\frac{\lambda}{D}$. Figure 7.14 shows an approximation of the primary beam at 100 GHz for an ideal 12-m antenna with HPBW = $52''$. (The actual HPBW of the ALMA 12-m antennas appears to be $1.13\frac{\lambda}{D} = 58''$.) As described in Chapter 6, the primary beam attenuation for a single-pointing observation can be corrected as the last step of imaging, but this correction produces non-uniform noise. What is called the "ideal antenna" in the current text would have terrible sidelobes and be unusable for single dish operation. The ALMA feedhorns were purposefully designed to provide an illumination pattern with a -12dB taper, in order to provide a nearly Gaussian beam response with low sidelobes while preserving as much of the resolution and sensitivity as possible. This is the fundamental

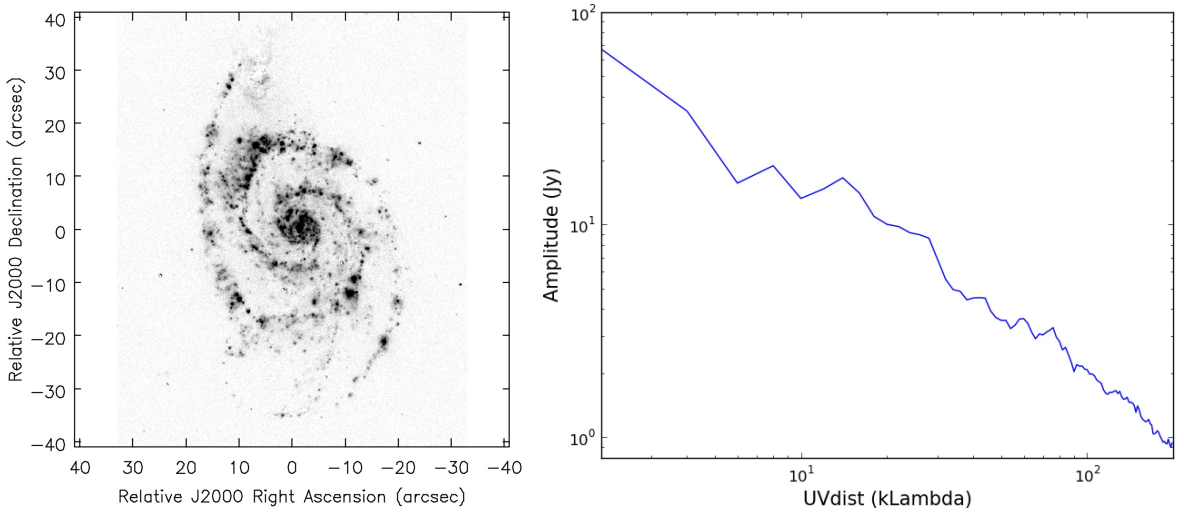

Figure 7.13: Image of Hα emission from M51 used as a model of emission at 100 GHz (*left*) and the expected visibility amplitudes with *uv*-distance (*right*).

tradeoff of radio telescope design.

As an example, Figure 7.15 shows the image and the fluxes recovered from a set of 1 Jy model sources spread over the primary beam, revealing the limitation produced by a single pointing. 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 twice the HPBW of the primary beam is requested. The pointing pattern used by ALMA to cover a such a map is an hexagonal grid with equilateral triangles whose vertices are separated by $\theta_{hex} = \frac{\theta_B}{1.2\sqrt{3}}$ to sample the emission at the Nyquist frequency. An ALMA user should keep in mind that a hexagonal mosaic has spacing θ_{hex} along a row (e.g., in R.A.) and $\frac{\sqrt{3}}{2}\theta_{hex}$ between rows (e.g., in decl.) Typically, long and short rows of pointings alternate, but to fill completely a rectangular area one extra pointing in the long rows is needed.

To estimate the number of pointings (N_p) necessary to cover an area of $L_X \times L_Y$, the following expressions can be used:

$$\begin{aligned} N_X &= (int)\left(\frac{1.2\times\sqrt{3}\times L_X}{\theta_B} + 1.\right) \\ N_Y &= (int)\left(\frac{1.2\times 2\times L_Y}{\theta_B} + 1.\right) \\ N_p &\approx (2N_X - 1.) \times \frac{N_Y}{2.} \end{aligned}$$

The OT can be used to set up easily a mosaic of adjacent pointings to cover a wide area at a chosen spacing. 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 should differ slightly from the OT. The ALMA user should keep in mind that mosaicing with the 7-m Array is more efficient because its 7 m antenna diameters have relatively larger HPBWs.

7.6 Combining 12-m and 7-m data

As shown in Table 7.1, the different 12-m Array configurations can achieve different angular resolutions, θ_{res}, depending mainly on the maximum baseline, and different maximum recoverable scales, θ_{MRS}, depending mainly on the minimum baseline. Different situations can arise for a scientific project according to its requirements of θ_{res} and θ_{MRS} which are set by the PI through the OT during proposal preparation. As shown by Figure

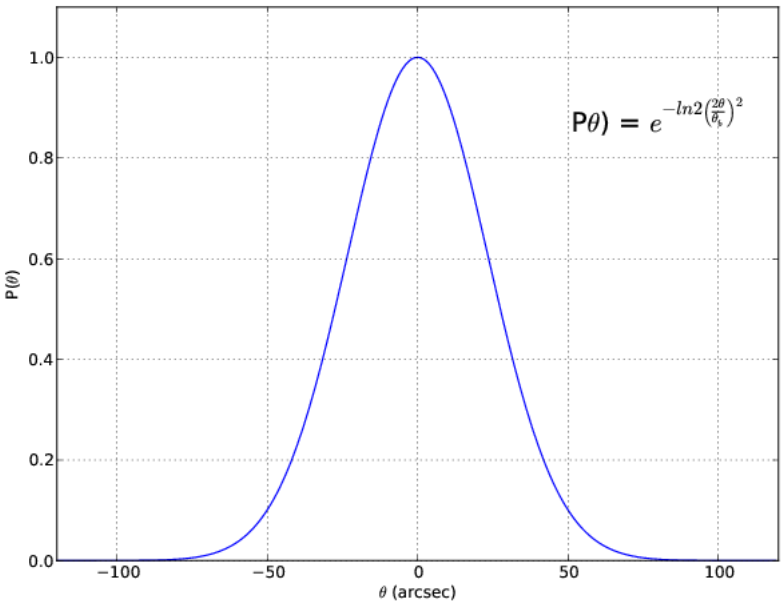

Figure 7.14: Approximation of the primary beam of a 12-m antenna at 100 GHz with HPBW = 52″.

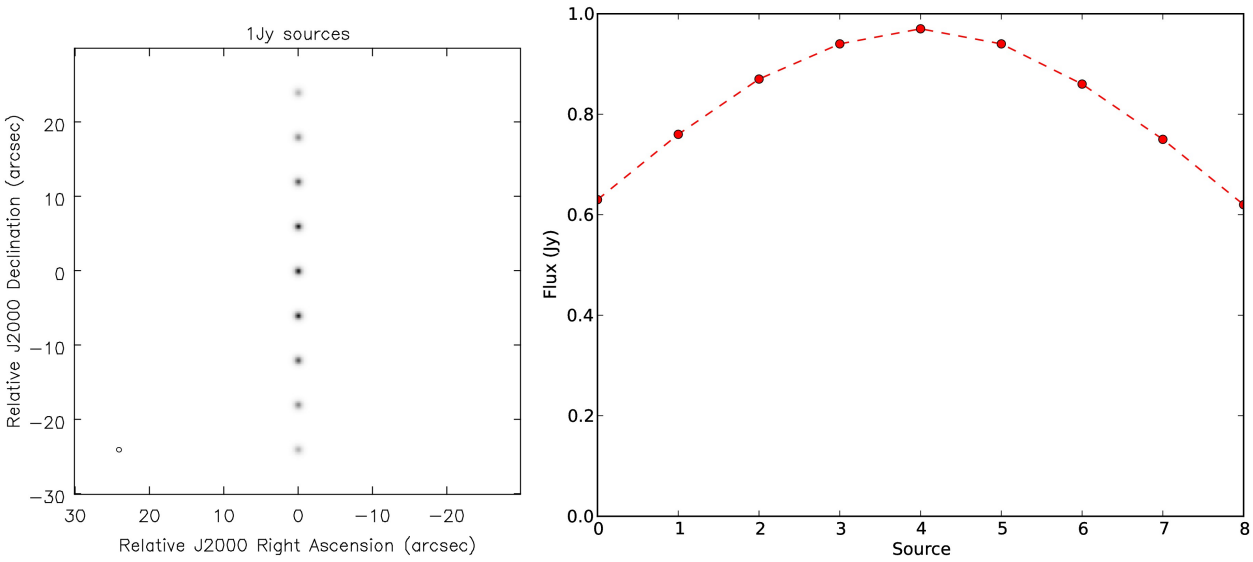

Figure 7.15: (*left*) Simulation of imaging nine model Gaussian sources (of 1″ size and 1 Jy flux) at 100 GHz using the C34-6 configuration C34-6 to observe over 2 hours. The primary beam attenuation has not been not corrected. (*right*) The measured flux of each source starting from the north.

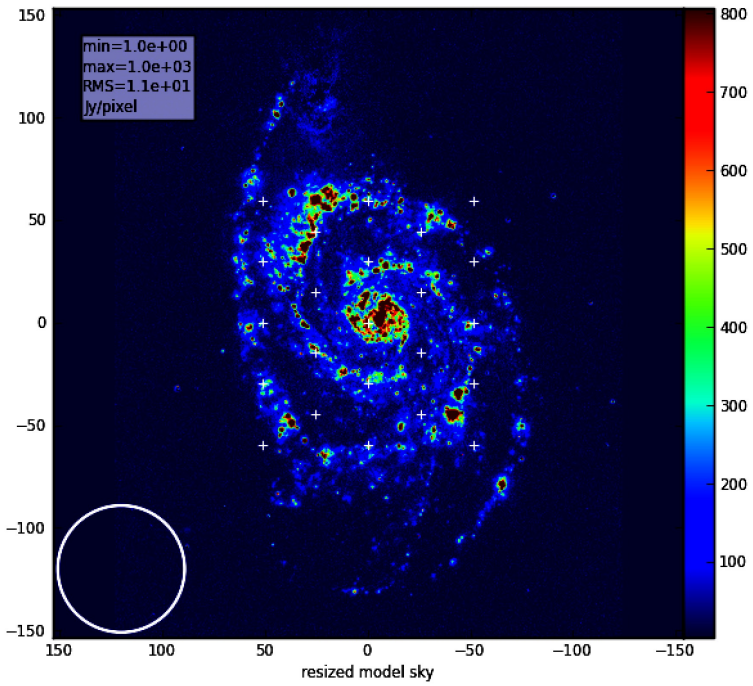

Figure 7.16: An example of mosaicing with a field of 2 arcmin at 100 GHz using an hexagonal pattern with a Nyquist sampling (white crosses).

7.13, however, all length scales have power contributing to the sky brightness distribution. The PI should be very careful to set the angular properties $\{\theta_{res}, \theta_{MRS}\}$ of their proposed ALMA observations since they have a strong impact on the way they will be planned and executed. Below is a summary of different situations with their possible consequences:

1. $\{\theta_{res},\theta_{MRS}\}$ are listed in Table 7.1 for a single 12-m Array configuration. In this simplest case, only one 12-m Array configuration is necessary.
2. $\{\theta_{res},\theta_{MRS}\}$ are listed in Table 7.1 but for different 12-m Array configurations. To achieve that specification, two 12-m Array configurations will be used and the data will be combined.
3. $\{\theta_{res},\theta_{MRS}\}$ are not listed in Table 7.1 because θ_{MRS} is larger than all those listed for the 12-m Array configurations. This situation means that short baselines are necessary and the use of 7-m Array is required. Again, different situations can arise:
   (a) The angular resolution θ_{res} is small enough to be achievable with a compact 12-m Array configuration, overlapping with the 7-m Array in the shorter baselines. The project will be performed with the following combination: TP + 7-m + C34-compact.
   (b) The angular resolution θ_{res} is too large to be achievable with a compact 12-m Array configuration but needs a more extended one. The project will be performed with the following combination: TP + 7-m + C34-compact + C34-extended.

The ALMA user proposing a project of case 1 and 2 has access only to the choice of $\{\theta_{res}, \theta_{MRS}\}$. The ALMA staff will schedule the appropriate array configuration(s) to achieve the angular properties requested by the project. As discussed above, the angular resolution depends on the weighting scheme used during the Fourier transform. For similar reasons, the combination of data from different configurations can alter the angular resolution and provide a slightly lower one than the most extended configuration taken alone. In that case, the data will pass QA2 (see Chapter 11) with a 10% tolerance on the requested pair $\{\theta_{res}, \theta_{MRS}\}$.

In the case of the set of 12-m Array configurations given earlier, which are only representative of the actual configurations used in Cycle 2, the following separation based on the *uv*-coverage can be drawn: C34-1, C34-2, and C34-3 are considered compact, C34-4, C34-5, C34-6, and C34-7 are considered extended. The extended configurations cannot be combined directly with the 7-m Array configurations because of their *uv*-coverage which would lead to "holes" in the obtained visibility distribution.

In case 3, the ALMA user has the choice to use the OT to select the ACA or not. Again, the choice of one or two 12-m Array configurations will be taken by the ALMA staff on the basis of the desired angular resolution. It is very important to note in the proposal the impact of θ_{MRS} which can multiply by a factor of 5.0 the total time needed for a given project with respect to the 12-m Array. The TPA and the 7-m Array antennas require more time simply because of their smaller collecting area. A strong justification should be provided for their use.

In Table 7.4, we summarize the different array combinations available during Cycle 2. The ALMA user should keep in mind that in Cycle 2 the ACA is not available as a stand-alone array and its use should be considered only in concert with the 12-m Array. A comprehensive plot about the array combination is shown in the Figure 15 of the ALMA Primer for Cycle 2 (Moriarty-Schiven et al., 2013).

With only 9 7m antennas the 7-m array have a limited (u, v) coverage for a snapshot observation. For Cycle 2, to get a reasonable image fidelity it is suggested that if the 7-m array time is less than 1 hour, the PI may need to increase it to reach that observing time. It will allow to get a sufficient *uv* coverage and a good S/N to obtain a sufficient image quality. It is particularly important to check the level of S/N reached by the 7-m array observations to be able to combine satisfactorily all the data together.

θ_{res} (arcsec)	θ_{MRS} (arcsec)	Array combinations for a specified $\{\theta_{res}$,$\theta_{MRS}\}$	Time ratios	Total Time
0.41-0.57	< 9.1	C34-7	1	$1.0 \times \Delta_{extended}$
0.41-0.57	9.1-18.0	C34-7 + C34-3	1 : 0.5	$1.5 \times \Delta_{extended}$
0.41-0.57	18.0-41.7	C34-7 + C34-3 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.41-0.57	> 41.7	C34-7 + C34-3 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
0.57-0.75	< 9.1	C34-6	1	$1.0 \times \Delta_{extended}$
0.57-0.75	9.1-18.0	C34-6 + C34-3	1 : 0.5	$1.5 \times \Delta_{extended}$
0.57-0.75	18.0-41.7	C34-6 + C34-3 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.57-0.75	> 41.7	C34-6 + C34-3 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
0.75-1.11	< 14.4	C34-5	1	$1.0 \times \Delta_{extended}$
0.75-1.11	14.4-26.3	C34-5 + C34-2	1 : 0.5	$1.5 \times \Delta_{extended}$
0.75-1.11	26.3-41.7	C34-5 + C34-2 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
0.75-1.11	> 41.7	C34-5 + C34-2 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
1.11-1.40	< 18.0	C34-4	1	$1.0 \times \Delta_{extended}$
1.11-1.40	18.0-26.1	C34-4 + C34-1	1 : 0.5	$1.5 \times \Delta_{extended}$
1.11-1.40	26.1-41.7	C34-4 + C34-1 + 7-m	1 : 0.5 : 2	$3.5 \times \Delta_{extended}$
1.11-1.40	> 41.7	C34-4 + C34-1 + 7-m + TP	1 : 0.5 : 2 : 4	$5.5 \times \Delta_{extended}$
1.40-2.04	< 18.0	C34-3	1	$1.0 \times \Delta_{extended}$
1.40-2.04	18.0-41.7	C34-3 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
1.40-2.04	> 41.7	C34-3 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$
2.04-3.73	< 26.3	C34-2	1	$1.0 \times \Delta_{extended}$
2.04-3.73	26.3-41.7	C34-2 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
2.04-3.73	> 41.7	C34-2 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$
3.73-7.46	< 26.1	C34-1	1	$1.0 \times \Delta_{extended}$
3.73-7.46	26.1-41.7	C34-1 + 7-m	1 : 2	$3.0 \times \Delta_{extended}$
3.73-7.46	> 41.7	C34-1 + 7-m + TP	1 : 2 : 4	$5.0 \times \Delta_{extended}$

Table 7.4: Array combination with the corresponding $\{\theta_{res}, \theta_{MRS}\}$ conditions for an observation at 100 GHz. As in the OT, the angular resolution is computed from the most extended configuration. The actual one obtained with combined configurations can be 50% lower due to different weighting (see text).

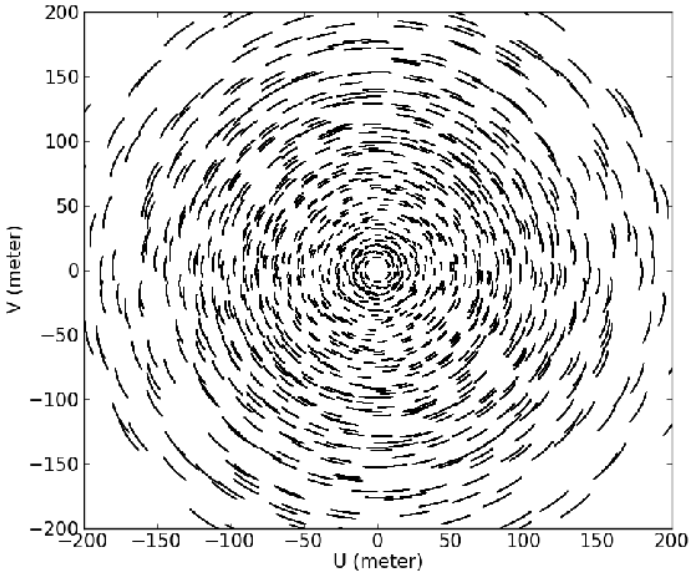

Figure 7.17: Expected *uv*-coverage for C34-2 (1 hour) and 7-m observations (3 hours).

An observation of 1 hour duration with the C34-2 configuration and 3 hours duration with the 7-m Array is shown in Figure 7.17 (*uv*-coverage) and Figure 7.18 (*uv*-plane sampling distribution). A weighting of $k = \frac{7^2}{12^2}$ has been applied to the 7-m Array data because of their smaller antenna diameters. The Cycle 2 7-m Array configuration with nine 7-m antennas provides *uv* measurements in the range 8-20 m and takes over the 12-m Array configurations in the range $R_{(u,v)} < 15$ m. A time ratio of 2-3 between 7-m and 12-m Array will be used in Cycle 2. The exact time ratio between the different array will be decided according to the nature of the target because the brightness distribution with respect to scale length can vary from one type of source to another. Nevertheless, a typical time ratio for a full (or partial) combination [TP : 7-m : C34-compact : C34-extended] will be [4 : 2 : 0.5 : 1].

Using the Hα emission in M51 as a model, and given the spread of its emission over different length scales (see 7.13), Figure 7.19 shows the images resulting from observing with a C34-compact configuration (C34-3), a compact and extended configuration (C34-3 + C34-7), and adding the 7-m. Natural weighting was used and deconvolution was performed using the CLEAN algorithm. The recovery of larger scales is quite noticeable with the use of the 7-m Array. Using only the C34-7 configuration at 100 GHz provides an angular resolution of $0.45''$. Combined with the C34-3 configuration, the angular resolution is lowered to $0.65''$. Finally, the angular resolution with the full combination (C34-3 + C34-7 + 7-m) is 0.67", very similar to that of only the two 12-m Array configurations but 45% larger than that obtained with the extended 12-m Array alone. That difference is due to the respective *uv*-coverage and weight of the array configurations in the combined dataset. Changing the weighting scheme (e.g., to uniform) can restore the highest angular resolution at the expense of sensitivity.

[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

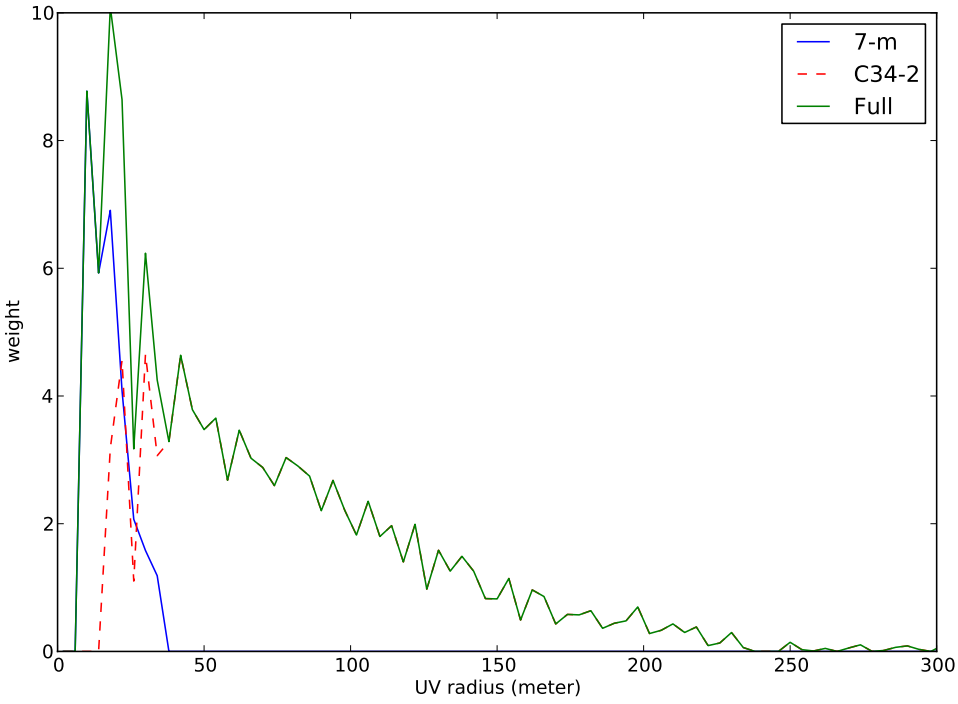

Figure 7.18: Radial density of *uv*-coverage with *uv*-distance using C34-2 (1 hour) and 7-m observations (3 hours).

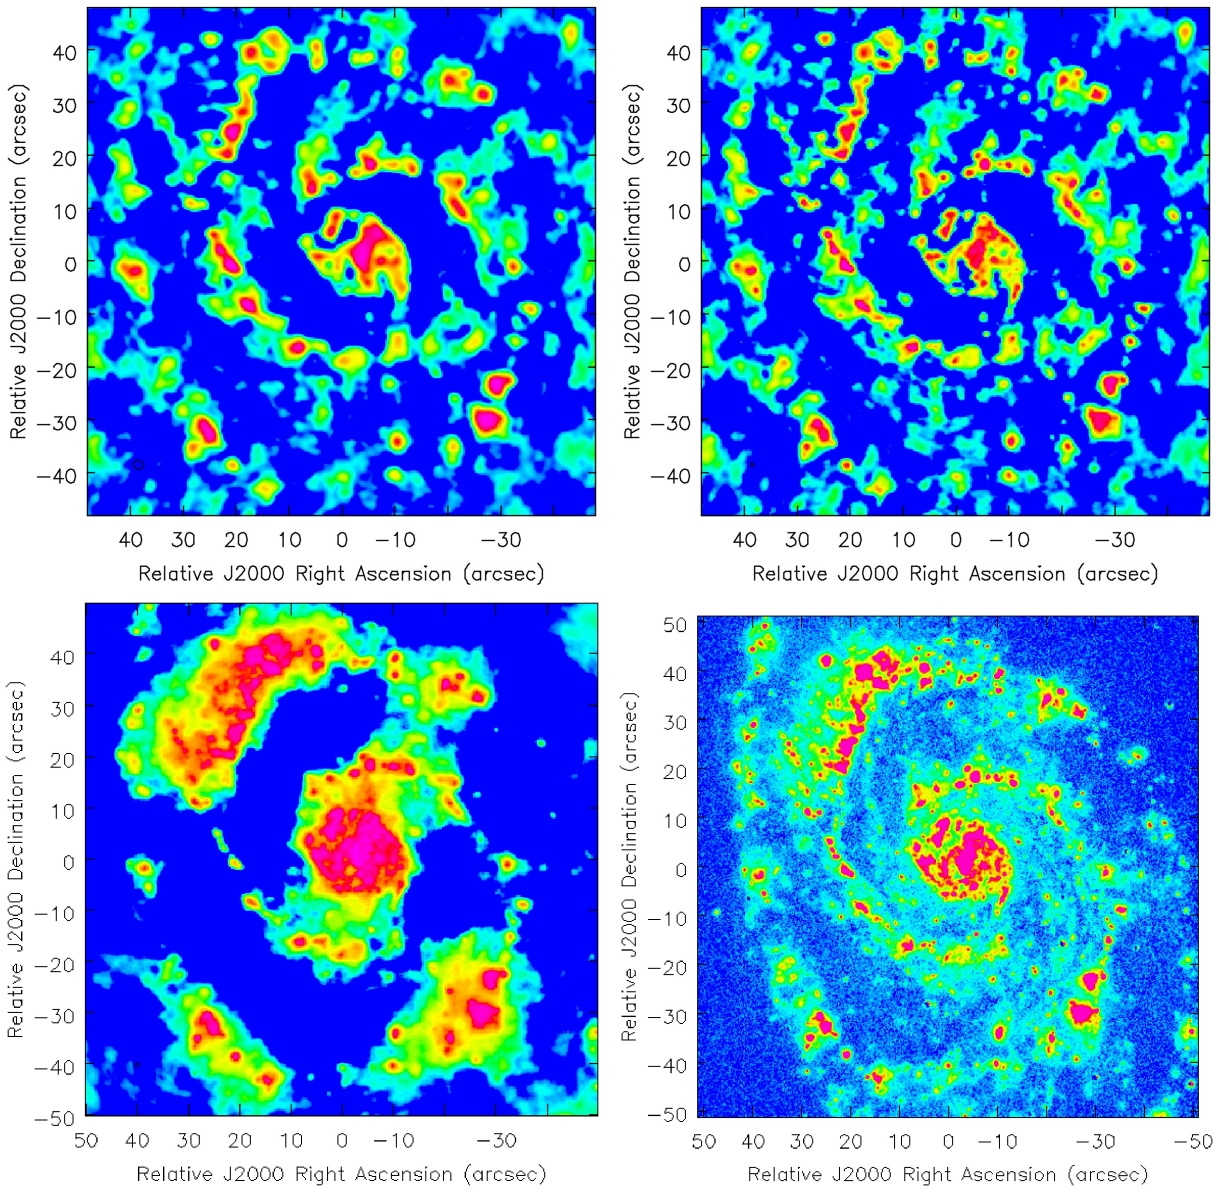

Figure 7.19: Images obtained using C34-3 (*top left*; 1 hour), C34-3 + C34-7 (*top right*; 1 + 2.5 hours) and C34-3 + C34-7 + 7-m (*bottom left*; 2 + 1 + 2.5 hours) array combinations, and the image model itself (*bottom right*).

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

[5] Rohlfs, K, Wilson, T. L. 2004. *Tools of Radio Astronomy*, Fourth Edition, ISBN 3-540-40387-6 Springer

Chapter 8

Observing Modes

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 science observations of each project are subdivided into a data structure that always has the same hierarchical levels.

At the bottom of this structure are the "Scheduling Blocks" (hereafter SB). They are the smallest sequence of observing instructions that can be scheduled. SBs are produced by the OT (including in Phase 1[1]) and may be edited using the OT in Phase 2 after proposal acceptance and before/during proposal execution. Projects are broken down into a set of these fundamental units to take maximum advantage of the properties of the ALMA site and the continuously-varying status of the observatory as a whole, including the weather.

An SB contains a very large amount of information about how and what should be observed, most obviously the positions and velocities of the science targets, details of the correlator setup and for how long the target should be observed. However, the SB itself does not control the observations. Instead, the SB calls an observing script which is appropriate to the observing mode required e.g., single-dish, mosaicing, and this then reads those parameters from the SB which it needs to perform the observation. The SB, for example, 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 measurements of system temperature).

Each SB has an execution time of typically 1 hour. Within this time limit, the set-ups, calibrations and target observations to ensure that the acquired data can be properly calibrated and used in the production of the final data product are carried out. The end of an SB may be specified in terms of a maximum amount of time or when certain well-defined science goals have been reached. An SB is atomic in the sense that it cannot be re-started in the middle. Therefore, an SB runs to completion, fails, or is terminated by the telescope operators. Given the limited length of the SBs, it often will be necessary to observe the SB several times to obtain a final image. If two science observations have to be observed together, e.g. because they share the same calibrator observations or because they need the same UV coverage or because they need to be simultaneous in time, then they end up in the same SB, otherwise, they will be grouped into separate SBs.

Whereas the SB is the smallest entity used for observing, the Observing Unit Set (hereafter OUS) is the smallest unit for data processing.

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

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

All SBs that have to be processed together to produce calibrated science products for the science sources contained in them are grouped together into one Member OUS. In most cases, there will be exactly one SB in each Member OUS, but the latter may hold multiple executions of this 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, the they will be placed into different Group OUSes. A typical case of Member OUSes that belong to the same Group OUS are observations of the 12-m Array, the 7-m Array and the TP array. According to the method above, these observations would be in three separate Member OUSes, but in the same Group OUS as one final image will be produced.

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

The SB/Member OUS/Group OUS are the smallest data structures that hold science observations that need to be observed/processed/combined together. This data hierarchy therefore maximizes observing flexibility and ensures that data gets processed as early and as fast as possible.

In the OT, observations are divided into different Science Goals (see Cycle 2 Proposer's Guide). In order to follow this structure, each Group OUS is attached to a Science Goal OUS. Each Science Goal contains only one Science Goal OUS.

8.2 Program Execution

Once a given SB has been selected for execution (by the Scheduler subsystem or by the Array Operator/AoD), it is read into the system and converted into a sequence of observations by the Control software. The Control subsystem creates an Execution Block (EB) structure that is attached to a given 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 Control subsystem constructs and executes a sequential series of scans. Each scan execution is in fact carried out by breaking it down into a series of subscans, each of which is itself broken into a series of integrations. Although commands are issued at the scan/subscan level, the correlator output corresponds to a particular integration. In general, each calibration observation consists of a scan containing several subscans. Similarly, the duration of the time on a single science source between phase calibrations can consist of one scan comprised of several subscans. Scans can be of arbitrary length, depending on the frequency of source changes or phase calibrations as dictated by the observing band and configuration, but subscans are recommended to be 30 seconds or less. Integrations tend to be on the order of one to 10 seconds within these subscans, but may be smaller for Total Power array observations. Calibration results from the Telescope Calibration subsystem (TelCal) and QuickLook (QL) pipeline are usually attached to a subscan. Some of these results, however, may be accumulated over many subscans, and the resulting object will be attached to the scan to which the subscans belong. The Control subsystem is responsible for the creation of all metadata needed downstream for data processing, for the EB, the scans, subscans and integration objects.

8.3 The Observing Process

For a given project, all observing Targets that are to be executed must be included in a Group within the SB. Each SB can include multiple Groups. The first Group (Group 1) of an SB is always the initial calibration group, with subsequent groups detailing the science observations. All Science targets and relevant calibrators within a Group are observed before the next Group is started. It is assumed that all the Science targets in a given group are sufficiently close as to share the same Phase and Pointing calibrators. 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 final 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-pointing interferometry, grouped source executions and pointed mosaics, use the same ALMA observing script, called the "standard interferometry" script. It is designed to provide well-calibrated observations of discrete sources. A set of necessary calibration measurements (e.g. amplitude, bandpass, sideband ratio) specified in group 1 are performed at the beginning of the observation sequence if appropriate sources are available, otherwise this information is collected at the end.

Pointing is verified before the amplitude and bandpass calibrators are observed (unless they are close together on the sky), and again before the main observations of the science target and phase calibrator cycle. Within a group, the Science targets are each observed in turn until observation of the primary phase calibrator (the calibrator with shortest cycle time in the group) is required. A typical cycle time for the phase calibrator may be 7-10 minutes. This process is repeated until the observing requirements are met, or the SB reaches its time limit. Any additional ("secondary") phase calibrators are observed as specified in the scheduling block. Atmospheric calibration to determine system temperature is performed during each iteration of primary phase calibration at high frequencies (Bands 7 and 9), and once every two primary phase calibrations at low frequencies (Bands 3 and 6). For a description of what each calibration entails, see Section 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 specify the particular calibrators or set up the queries based on different criteria, 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 cabliration" in the OT).

8.4 Single Field Interferometry

Single field interferometry is the most basic form of observation that can be taken. It consists of standard calibration scans associated with the constituent targets, and science observations of a single field (primary beam).

A typical observation will start with a bandpass calibration. The bandpass observation is executed to measure the spectral response of the system, and thus should be done on a bright source with simple spectral properties, such as a bright quasar with no emission or absorption lines and a reasonably flat spectrum.

The amplitude calibration will be taken next, which is intended to obtain the observed flux of a well known source, such as a solar system object. The observed flux will be used to compare with the established model fluxes of these objects to obtain the scaling factor to be applied to all other sources in the SB. Ideally these amplitude calibration sources should be small with respect to the synthesized beam so as not to resolve the source 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 all antennas.

The phase calibrator is used to correct for changes in the atmosphere which result in variations in phase during the course of the observations. 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.

Multiple fields can be observed separately (each field being a separate single pointing) in an SB by specifying "Multiple Pointings" in the OT.

8.5 Pointed Mosaic Observations

Pointed mosaic observations will again be offered in Cycle 2. 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. For larger fields that cannot be covered by 150 pointings, multiple SBs may need to be defined. For a scan executed on a Science target which includes a pointed mosaic, each position in the mosaic is observed for one subscan duration (see Section 4.4) until the time for the scan expires. Both the Science target specific last mosaic pointing position and index within the Science target list are stored to ensure the next scan begins on the proper Science target and proper offset position.

The mosaic observations can be set up by specifying "1 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 Zero Spacing Observations

The purpose of adding data from single-dish observations using autocorrelation is to recover large scale emission from the science target that may be spatially filtered with even the shortest baselines of the 7-m Array. For this reason, these observations are referred to as "zero spacing." For convenience, these observations are also sometimes referred to as "total power" (TP), although in practice they are taken in autocorrelation mode rather than using a total power detector, which means that it is possible to match the spectral resolution to that of the interferometric observations to facilitate combining the data. Four 12 m antennas connected to the ACA Correlator are available for this purpose. For cycle 2, continuum observations is expected to be offered in fast scanning mode.

The TP observations differ from the 12-m Array and ACA 7-m Array observations in that they use an observing script called the "standard single-dish" script. This script shares many features of the standard interferometry script because the calibrations for the zero spacing observations (pointing, sideband ratio and focus measurements) are taken interferometrically. The SB of a TP observation consists of a Group of calibrations followed by On-The-Fly (OTF) observations of a rectangular area on the science target, with periodic reference position and calibration observations. For the OTF and reference position integrations, only the autocorrelation data will be written to the ASDM to minimize data rate and size, whereas calibrations use cross-correlation information for analysis.

Most of the calibration scans are executed in Group 2 due to limitations in the observing script in cycle 1. Group 1 will consist only of a pointing calibration scan. Group 2 contains another round of pointing calibration, then a sideband ratio measurement. The sideband ratio measurement will be accompanied by a delay calibration (and atmospheric calibration), which is necessary to realize the execution of the 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 1, scans are taken either in longitudinal or latitudinal directions as specified in the OT. In later cycles, scans may be taken in the two perpendicular directions in turn, to minimize scanning effects. The reference positions, assumed to be positions which are free of emission, are specified either in absolute coordinates or as offsets from the map center. By default, the OTF map will cover an area one half of the full beam larger than the interferometric observation on all sides. This will ensure that undersampling at the map edges does not affect the data combination process with the interferometric data. A raster row in these observations is defined as a subscan, and the length of a scan (consisting of some number of subscans) is currently hard-coded to 60 seconds. The reference position is observed as specified by its cycle time during the science target scans. Pointing is calibrated on a bright calibrator near the Science target with a frequency indicated in the SB, and atmospheric calibrations are taken every 15 minutes at the reference position to measure the system temperature.

The calibrated TP map of the science target will be in units of Kelvin, on the antenna temperature (Ta*) 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. 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 obtain η_{mb} and θ. These values are used to calculate the Kelvin to Jansky/beam conversion factor, which will then be applied to the science observations.

8.7 Polarization

ALMA antennas can receive two linear polarizations (X and Y) simultaneously and output four cross correlations (XX, YY, XY, and YX). These four cross correlations are given by

$$\begin{aligned} X_iX_j &= I + Q\cos 2\psi + U\sin 2\psi \\ Y_iY_j &= I - Q\cos 2\psi - U\sin 2\psi \\ X_iY_j &= -Q\sin 2\psi - U\cos 2\psi + iV \\ Y_iX_j &= -Q\sin 2\psi - U\cos 2\psi - iV, \end{aligned}$$

where i and j are antenna numbers, I, Q, U, and V are Stokes parameters, and ψ is the parallactic angle. Using all these correlations, one can obtain Stokes I, Q, U, and V visibilities, which allows one to investigate the source polarization property[2]. In addition to the standard calibration for XX and YY correlations (hereafter "parallel-hand" correlations) such as system temperature, wvr, gain, and bandpass calibrations, the unknown delay and phase offsets between X and Y correlations (hereafter "cross-hand" correlations) and instrumental polarization (or polarization leakage) calibrations must be applied. The cross-hand phase and the instrumental polarization might be stable over days, and one could use the solutions derived from other experiments. However, for the time being, these calibrations are encouraged to be done in each execution.

To permit cross-hand delay and phase calibrations, a bright linearly-polarized (more than a few percent) calibration source must be observed. This calibration source is usually substituted by the instrumental polarization calibration source described below.

Instrumental polarization mainly arises from the imperfect polarization separation at the frontend. The instrumental polarization appears in the cross-hand correlations and corrupts the source polarizations. The cross-hand correlations can be written as

$$\begin{aligned} X_iY_j &= -Q\sin 2\psi - U\cos 2\psi + iV + D^*_{Y_j}\left(I + Q\cos 2\psi + U\sin 2\psi\right) + D_{X_i}\left(I - Q\cos 2\psi - U\sin 2\psi\right) \\ Y_iX_j &= -Q\sin 2\psi - U\cos 2\psi - iV + D_{Y_i}\left(I + Q\cos 2\psi + U\sin 2\psi\right) + D^*_{X_j}\left(I - Q\cos 2\psi - U\sin 2\psi\right), \end{aligned}$$

where D_X is the polarization leakage from Y polarization to X polarization and D_Y is the polarization leakage from X polarization to Y polarization. The asterisk denotes complex conjugates. These leakage terms are usually called D-terms. Figure 8.2 shows an example of D-terms at Band 3. The D-term level is typically few percent at bands 3, 6, and 7 on axis with some variations over frequency. Without any D-term calibration, an unpolarized source may appear to be polarized at the 1% level. The most straightforward way to calibrate the D-terms is to observe an unpolarized source. The cross-hand output will be purely D-terms since Stokes Q, U, and V are equal to zero. However, bright unpolarized sources are rarely found. Alternatively, it is possible to separate the source polarization and the D-terms by observing a polarized source with little structure over a wide range of parallactic angles. Many quasars suffice this condition. The D-term calibrator should be observed for at least 4 (*to be decided*) scans covering a wide range (> 45 degree, *to be decided*) of parallactic angles. The rotation of the parallactic angle mostly occurs right before and after the transit, but low declination sources do not rotate much (see Figure 8.3). The D-terms may have slight elevation dependence, and therefore it is favorable for the

[2] $\sqrt{(Q^2+U^2)}$: linear polarization flux density, $0.5\tan^{-1}(U/Q)$: linear polarization position angle, V: circular polarization flux density

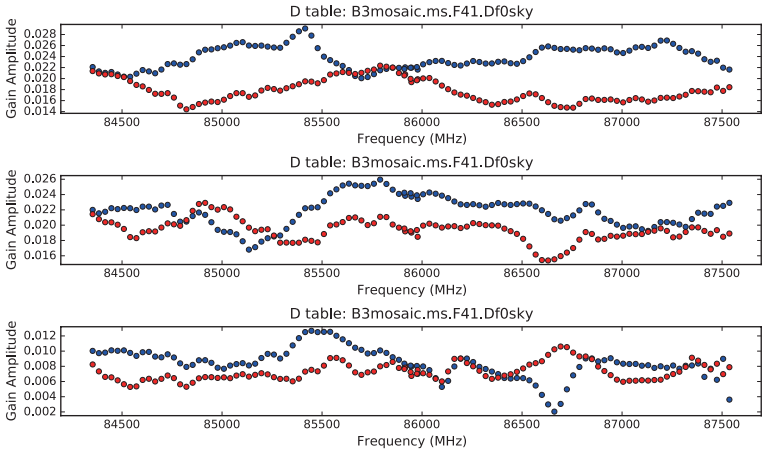

Figure 8.2: The D-term plots of DA42 (top), DV03 (middle), and PM01 (bottom). The vertical axis is the fraction of the input signal voltage in one polarization that leaks into the output of the other polarization in voltage unit and the horizontal axis is the frequency in MHz. The blue and red symbols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

D-term calibrator to be close to the science target (less than 40 degree). The ALMA observatory will choose appropriate calibrators (*to be decided*), but the users must take into account these limitations when planning the observation. The users can check the calibrator candidates from the ALMA polarization calibrator database.

With D-term calibration, linear polarization imaging of a compact source at the level of 0.1% is feasible. The accuracy of absolute polarization position angle will be nominally 6 degrees. The users will be encouraged to see the polarized flux and polarization position angle of the calibrator in the data reduction process and check the consistency with those taken by the polarization calibrator survey nearby in time. Although Stokes V (circular polarization) image can be produced with CASA, this capability is still under commissioning. The ALMA observatory cannot guarantee the accuracy of Stokes V image in Cycle 2. The users may investigate the Stokes V emission at their own risk.

The D-term component arise from the off-axis geometry of feed horn, antenna illumination, or the alignment of optics, and will vary across the primary beam pattern. Generally, the D-term level becomes larger with increasing the offset from the beam center. This is so-called “off-axis" instrumental polarization. In cycle 2, the off-axis instrumental polarization calibration will not be employed. Therefore, the acceptable angular size for the observing source is limited where the off-axis instrumental polarization level does not significantly become large. For this reason, the largest angular size of the source must be less than FWHM/3 of the primary beam. Polarization imaging with the level of better than 0.3% will be achievable in this area. The users must justify that the source polarization level of science target is well higher than these accuracies.

In Cycle 2, the polarization observation is limited to continuum source with the TDM mode. No mosaic, ACA, and TP array observations are offered. Further, a minimum execution time of 3 hours will be imposed to ensure that sufficient coverage in parallactic angle for calculation of the D-term.

8.8 Multiple Region Modes

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

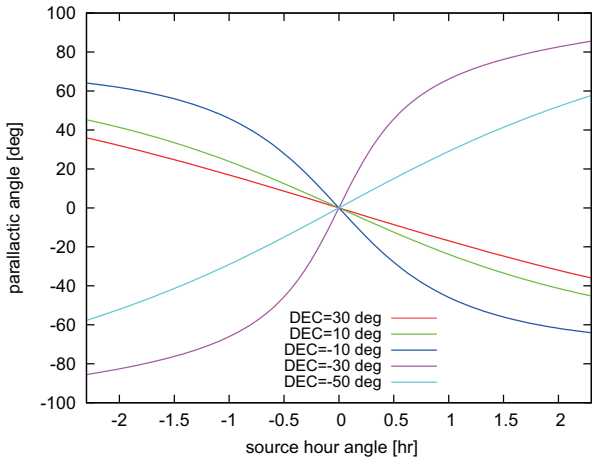

Figure 8.3: Parallactic angle plot as a function of hour angle for the declination of 30, 10, -10, -30, and -50 degree.

- 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 the same for all spectral windows.

This mode is useful when a number of line features which require high spectral resolution are spread across the IF bandwidth. Since the filters have a unit width of 62.5 MHz, if the user chooses a mode with 250 MHz total bandwidth, it is possible to place 4 separate windows, each with 62.5 MHz width, anywhere within the 2 GHz baseband. In Cycle 2, a maximum of 4 spectral windows per baseband will be offered.

8.9 Observation of Ephemeris Objects

Observation of solar system objects (with the exception of the Sun) is supported as in Cycle 1. Several well known solar system objects including planets, satellites and asteroids can be selected from a pulldown menu in the Observing Tool. For other sources including non-sidereal objects, an external ephemeris file can be supplied as an input. The ephemeris file must be in JPL Horizons format. A typical ephemeris file may consist of the date (time), Right Ascension, Declination, range and range rate, for example:

```
*******************************************************************************
 Date__(UT)__HR:MN     R.A._(ICRF/J2000.0)_DEC            delta      deldot
*******************************************************************************
$$SOE
 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. Solar observations will also be available at a later date.

Chapter 9

ALMA Sensitivity Calculator

The main tool for calculating the sensitivity of ALMA is the ALMA Sensitivity Calculator (ASC). This is an application contained within the ALMA Observing Tool (OT) that allows a user to experiment with various sensitivity options, and which is also used by the OT to calculate its time estimates based on the parameters entered into a project's Science Goals. The same tool is also available as a Java applet in the ALMA Science Portal[1].

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

9.1 Calculating system temperature

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

$$T_{\rm sys} = \frac{1+g}{\eta_{\rm eff} e^{-\tau_0 \sec z}} \Big(T_{\rm rx} + \eta_{\rm eff} T_{\rm sky} + (1-\eta_{\rm eff}) T_{\rm amb} \Big) \tag{9.1}$$

The various terms are

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

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

Table 9.1: Receiver temperatures (and their specifications) assumed in the ASC as a function of ALMA band. For most of the bands we are currently assuming the ALMA specification for the receiver temperature that should be achieved across 80% of the band, $T_{\rm rx,spec}$. 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).

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

$T_{\rm sky}$ and $T_{\rm amb}$ are 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.

$$T_{\rm n} = T \times \left(\frac{h\nu/kT}{e^{h\nu/kT} - 1} \right) \tag{9.2}$$

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

The terms $\eta_{\rm eff}$ and $e^{-\tau_0 \sec z}$ both attenuate the source signal and we thus divide through by them to obtain a measure of the system noise that is relative to the unattenuated source.

The temperature of the Cosmic Microwave Background is not explicitly included in Equation 9.1 as it is included in $T_{\rm sky}$ (see Section 9.1.2).

9.1.1 Receiver temperatures

For many of the ALMA bands, the calculator currently only uses the specifications for the receiver temperatures (required over 80% of the receiver bandwidth) and not the actual measured values. However, for the Early Science bands (3, 6, 7 and 9), typical values measured in the laboratory are used as these are usually significantly better than the specifications. The measured values are somewhat conservative and so are in between what we might expect at the middle and edges of the bands. The values used in the ASC are given in Table 9.1. 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).

Table 9.2: 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.

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

9.1.2 Sky temperature

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

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

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

where $T_{\rm atm}$ is the mean physical temperature of the atmosphere. This can be calculated using Equation 9.3 and from there the sky temperature as a function of zenith angle:

$$T_{\rm sky,z} = T_{\rm sky,z=0}\frac{(1 - e^{-\tau_0 \sec z})}{(1 - e^{-\tau_0})} \tag{9.4}$$

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 9.2.

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

9.1.3 Ambient temperature

This is essentially spillover from the sidelobes of the antenna beam corresponding to emission from the ground and the telescope itself. This is held constant at 270 K (median value as measured from many years of monitoring data at the ALMA site). The value used by the ASC is corrected according to Equation 9.2.

[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.

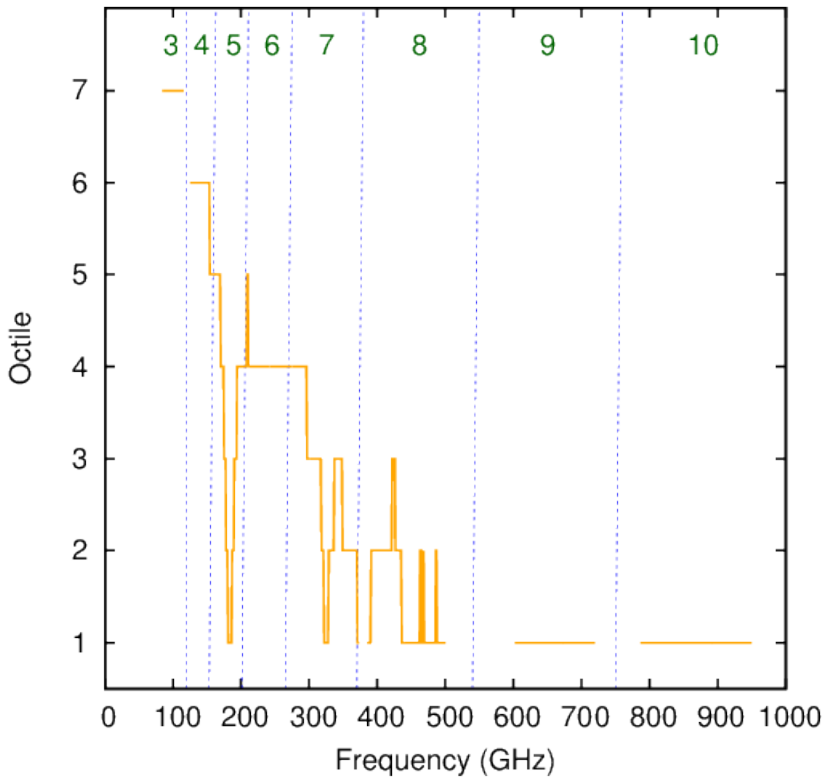

Figure 9.1: Plot of PWV octile assumed by the ASC as a function of frequency, for a source declination of zero degrees. The vertical lines separate different bands, the numbers of which are shown at the top of the plot. The water line at 183 GHz (band 5) is particularly prominent. In general, higher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

9.2 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 residual pointing and focus error.

9.2.1 12-m and 7-m Arrays

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

$$\sigma_{\rm S} = \frac{2\,k\,T_{\rm sys}}{\eta_{\rm q}\eta_{\rm c}A_{\rm eff}\sqrt{N(N-1)\,n_{\rm p}\,\Delta\nu\,t_{\rm int}}}. \tag{9.5}$$

The various symbols are

- $A_{\rm eff}$ – effective area. This is equal to the geometrical area of the antenna multiplied by the aperture efficiency ($\eta_{\rm ap}$). The latter is given by the Ruze formula i.e. $\eta_{\rm ap} = R_0 \exp\left(-16\,\pi^2\sigma^2/\lambda^2\right)$ where σ is the rms surface accuracy of the antenna – the specification of 25 μm and 20 μm for the 12-m and 7 m antennas respectively is currently used[3]. R_o is equal to 0.72. See Table 9.3 for values of antenna efficiencies and effective areas in various ALMA bands

[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.

Table 9.3: Aperture efficiencies at typical continuum frequencies for both the 12 and 7 m antennas. The effective area, $A_{\rm eff}$, is equal to the efficiency 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.

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

- $\eta_{\rm q}$ – 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
- $\eta_{\rm c}$ – correlator efficiency. This depends on the correlator (64-input or ACA Correlator) and correlator mode, although the efficiency of all 64-input Correlator modes available at Cycle 2 is equal to 0.88. The ACA efficiency *will* depend on the mode, but this is only taken into account in the OT, not by the Java applet (see below) which therefore also assumes a value of 0.88
- N – number of antennas. This defaults to 34 for the 12-m and nine for the 7-m Array
- $n_{\rm p}$ – number of polarizations. $n_{\rm p} = 1$ for single-polarization and $n_{\rm p} = 2$ 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. $n_{\rm p}\,\Delta\nu$ is often referred to as the effective bandwidth
- $t_{\rm int}$ – total on-source integration time.

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

$$\sigma_{\rm T} = \frac{\sigma_{\rm S}\lambda^2}{2k\ \Omega} \tag{9.6}$$

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

$$\Omega = \frac{\pi\theta^2}{4\ln 2}. \tag{9.7}$$

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

$$\sigma_{\rm TP} = \frac{2\,k\,T_{\rm sys}}{\eta_{\rm q}\eta_{\rm c}A_{\rm eff}\sqrt{N\,n_{\rm p}\,\Delta\nu\,t_{\rm int}}}. \tag{9.8}$$

This is the same as Equation 9.5, apart from there only being a factor of $\sqrt{N}$ in the denominator, where N is the total number of total power antennas.

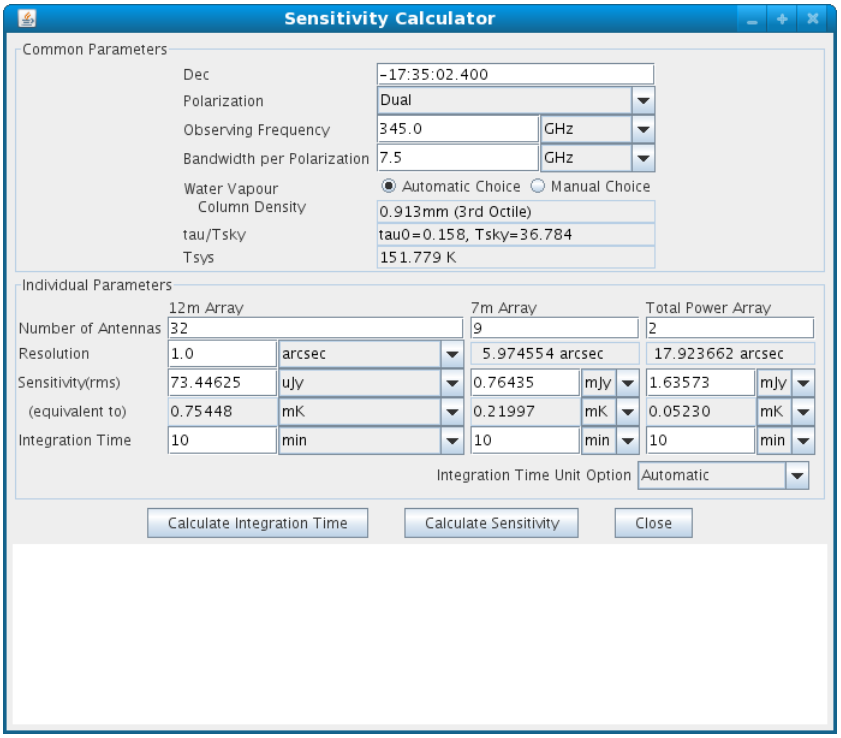

Figure 9.2: Screenshot of the GUI version of the ALMA Sensitivity Calculator as implemented in the ALMA Observing Tool. The white area at the bottom is for displaying error messages i.e. parameters out of bounds. The example here shows the achievable sensitivity for all three arrays for an on-source time of 10 minutes.

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 available at Cycle 2.

9.3 User Interface

The main way that a user interacts with the Calculator is through a GUI in the OT or via a Java applet on a web page - both are essentially identical. By entering various parameters, the time required to achieve a particular sensitivity (in either Jy or K) can be calculated, or vice versa. The inputs that affect the sensitivity or time are given below; a screenshot of the OT's GUI version is shown in Fig. 9.2.

- Source declination – this is used to calculate the maximum elevation of the observation and thus the minimum airmass i.e. the ASC assumes that the source is transiting.
- Observing frequency – this sets the receiver temperature, antenna efficiency and the atmospheric opacity.
- 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 34 from the 12-m Array, nine from the 7-m Array and two from the TP Array.
- Angular resolution – this affects the time estimates when sensitivities are specified in temperature units. The calculator will not perform any calculations when Kelvin 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

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

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

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

10.1 Long-Term Effects

10.1.1 All-Sky Pointing

These observations are used to correct the overall mechanical and orientation imperfections of each antenna & pad assembly. These imperfections include tilts of the azimuth (hereafter AZ) axis, zero-points of the encoders, encoder run-outs and tilts in AZ and elevation (hereafter EL), non-perpendicularity of the AZ and EL axes, sagging of the sub-reflector structure, differential sagging of the telescope structure in AZ between two successive antenna supports, plate offsets among ALMA Bands, etc. The preferred method of measurement is a series of interferometric observations of 50-100 pointing sources (mostly unresolved continuum sources) covering a wide range in AZ and EL. A 14-18 parameter pointing model (with well-known angular dependencies) is fitted to the data and used to derive the required corrections to the nominal pointing positions to correct for all of these structural effects. It is customary to define one of the ALMA bands as pointing reference (currently Band 3 for all antennas) and define the other Band pointing models as offsets from the reference[1]. The ALMA antennas have an all-sky pointing accuracy of $2''$ RMS under primary operation conditions.

[1] A full pointing model for each antenna is created at Band 3. The collimation offsets (CA and IE) are then measured for all other bands w.r.t. Band 3. However, whenever a pointing observation is carried out during operations, it is actually measured at Band 6 because that Band is closer to the focal plane center than Band 3 and doesn't have external tertiary optics.

10.1.2 Focus Models

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

10.1.3 Baseline Vectors

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

10.1.4 Cable Delay

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

10.1.5 Antenna Characteristics

Surface Measurements/Adjustments

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

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

Beam Patterns

Accurate knowledge of the primary beam patterns, for all ALMA bands, polarizations and ALMA antenna types[2], is required for high-quality imaging and for assessing telescope performance. Interferometric or SD (Single Dish) observations of bright unresolved sources, obtained by scanning the antenna to be measured in square patterns across the source, are usually used. These allow for the derivation of high SNR maps of the antenna total power beams. Currently the model beam patterns are used in the offline reduction instead of the measured results, but CASA has the capability to import them if necessary. For observation planning, the current best estimate of the ALMA primary beam profile is a Gaussian with a FWHM = $1.13\lambda/D$, where D=12 m (or 7 m for the ACA).

10.2 Short-Term Effects

10.2.1 Offset Pointing

To be able to center sources within the primary beam field of view at the highest frequencies offered by ALMA, a pointing accuracy for all the 12 m antennas of about $0.6''$ is needed. This is clearly not possible with the All-Sky Pointing model, and must be measured separately. This is achieved with interferometric observations of nearby pointing calibrators (within 4 degrees of intended target). Since some of the calibrators needed within an SB can be widely separated from the science target, i.e. bandpass and amplitude, it is customary to carry out offset pointing observations for those calibrators, too. Once the science target observations are started, usually one pointing calibration is enough for the typical duration of an SB ($\approx$ 1hr). These offset pointing measurements will only update the two collimation offsets (AZ, EL) of the antenna pointing model. Due to the reduced number of suitable pointing sources at the highest frequencies (particularly ALMA Bands 7 and 9), pointing calibrators will usually be observed in a lower frequency band (Band 3/6) and these observations along with relative focal plane offsets between receivers will be used to update the offset pointing model.

10.2.2 Bandpass

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

[2]There are three different designs of the ALMA 12 m antennas, which results in different beam sidelobe properties

simple formula for the time required for a bandpass observation can be derived assuming that the required spectral SNR is that of the target source:

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

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

10.2.3 WVR Corrections

Fluctuations in the line-of-sight Precipitable Water Vapor (PWV) of two antennas in a given baseline can cause significant decorrelation, especially at the high frequency bands of ALMA. Fluctuations in PWV levels are driven by wind, and can have rapid and strong variations over a large range of spatial scales. All ALMA 12 m antennas are equipped with a Water Vapor Radiometer (Dicke-type) that measures the emissivity of the atmospheric water line within four spectral bands near 183 GHz at a rate of 1 Hz. From these measurements, atmospheric models are used to derive the amount of PWV in the line-of-sight and from the difference between the values for the two antennas in a baseline, the phase fluctuations at the observing frequency. The PWV measurements are also used to derive sky opacities at the frequencies of the science observations when combined with the ATM atmospheric model. It is expected that the phase fluctuation corrections will be applied automatically by the correlator during Cycle 2. Since the ACA is very compact, the effect of the atmospheric PWV fluctuations is quite smaller than that for the 12-m Array. For Cycle 2, several WVR phase correction schemes are being explored (including not correcting it, because of the short spacings involved in the ACA), and a final decision on which to use will happen during Cycle 1. A more detailed description of the WVR hardware, specifications and operation can be found in Section A.6.

10.2.4 Gain (Amplitude & Phase)

The amplitude and phase transfer properties of the atmosphere and receiving system vary more slowly than the PWV fluctuations but vary on timescales of about 5-10 minutes and are still rapid compared to a typical SB. They are corrected using nearby calibrators (within 10 degrees of target). Usually, observations of the science targets are done in cycles straddled by observations of the gain calibrators. Since the WVR on the ALMA 12 m antennas help to correct the phase fluctuations at short timescales (i.e., 1 second to 2 minutes), the gain calibrations can be done more sparsely (every one to a few minutes). Faster cycle times between the gain calibrators and the target sources will only be needed for special observations at long baselines and under very adverse weather conditions. Ideally the calibrators should be unresolved for simplicity of reduction, but CASA can also cope with resolved calibrators if necessary (using models). Also, band-to-band phase transfer techniques are possible for example using gain observations at Bands 3 or 6 to calibrated science target data at Bands 7 or 9. This type of calibration requires observation, once per SB, of at least one calibrator at both the high and low bands to calibrate the band to band phase difference. Note that band-to-band phase transfer techniques are only needed when there are no sufficiently bright/nearby calibrators available for the high frequency bands. For Cycle 2, the fact that full polarization observations of compact sources will be offered will allow for the solving of intrinsic polarization of the gain calibrators, in case this is needed to improve accuracy. The observatory, however, will strive to minimize these errors by using calibrators with small polarization as much as possible. An estimate of the on-source observing time required to reach a given phase RMS is given by:

$$\tau(sec) = 9000 \left(\frac{T_{sys}}{100\,\mathrm{K}}\right)^2 \frac{31}{N_{ant}-1} \left(\frac{12\,\mathrm{m}}{D_{ant}}\right)^4 \frac{2\,\mathrm{GHz}}{\Delta\nu} \left(\frac{0.1\,\mathrm{rad}}{\sigma_\phi}\right)^2 \left(\frac{1\,\mathrm{mJy}}{S_\nu}\right)^2 \exp\left(\frac{0.1\epsilon}{25\,\mu\mathrm{m}} \frac{\nu}{100\,\mathrm{GHz}}\right)^2 \tag{10.2}$$

where τ is the required time in sec, T_{sys} the system temperature, N_{ant} the number of antennas, D_{ant} the diameter of one of the antennas in the array, $\Delta\nu$ the total bandwidth used in determining the phase (it is

usually the whole bandwidth of a given spectral window), S_ν the flux of the calibrator at the frequency of the observations, ϵ the surface RMS of an antenna (25 μm for the 12-m Array and 20 μm for the 7-m Array), ν the frequency of the observations and σ_ϕ the desired phase RMS. See also Table A.3, which shows some representative values.

10.2.5 System and receiver temperature

At millimeter and submillimeter wavelengths, the atmosphere both attenuates astronomical signals and acts as a black body emitter that adds additional noise to any measurements. This effect is a strong function of frequency, elevation, the column of wet and dry constituents of the atmosphere, and the temperature of the atmosphere. To measure this effect, the ALMA front ends are equipped with an Amplitude Calibration Device (ACD) which consists of a robotic arm with two "loads" of known emissivity, one at the antenna cabin temperature and the other heated to 353 K (i.e., 80 C). Consecutive observations of these two loads and an observation of the sky are used by ALMA to measure the $T_{\rm sys}$ (see description of the measurements in Section A.5.1). Because of its dependence on the possibly changing ambient temperature and elevation, $T_{\rm sys}$ needs to be tracked every 5-15 minutes depending on frequency. The specifications for $T_{\rm sys}$ calibrations are to reach 1% repeatability[3] for ALMA bands up to 7 and 3% repeatability at higher frequencies.

10.2.6 Sideband Ratio

ALMA receivers for Cycle 1 include types 2SB (i.e., the receiver outputs two separate sidebands) and DSB (i.e., the sidebands are superposed at output). Both types can have significant leakage of signals from one of the sidebands to the other, affecting the flux calibrations of the results (for 2SB receivers the specification is 10dB). Some techniques (i.e., LO offsetting) will be used during the interferometric observations to significantly reduce the leakage, but the residual has to be measured. The Sideband Ratio will also be measured interferometrically for the Total Power Array antennas. Since the ALMA receivers have very stable tuning properties, the Sideband Ratio (SBR) will usually only have to be measured once at the start of an SB using sources with well-known spectral properties. The results are automatically applied to all the $T_{\rm sys}$ calibrations during the observations.

10.2.7 Flux

Calibration of the data using the $T_{\rm sys}$ method above only yields an approximate absolute calibration of the actual flux density of a given source. In order to derive a more accurate absolute flux calibration scale, a measurement of a source with known flux density and structure is required for each spectral setup, each time an SB is executed. For observations, Solar System bodies [4] in conjunction with models for their size and brightness temperature are most often used for absolute flux calibration, though frequently monitored quasars (up to Band 7) and the photosphere of stars (currently being evaluated) are also options. A measurement of a flux calibrator is needed once per tuning (or per SB if high accuracy is desired). The absolute calibration accuracy for Cycle 2 is 5% for Band 3 and 4, 10% for Bands 6 and 7, and 20% for Bands 8 and 9.

10.2.8 Polarization

For accurate polarization imaging, the crossed-hand correlations, XY, YX, must be accurately calibrated. Hence, in addition to the parallel-hand calibrations described in the previous sections of this chapter, observations of a strongly linearly-polarized calibrator ($> 3\%$) must be added to the schedule. In Cycle-2, to determine the three polarization-type calibrations (see section 8.6), the SB should be at least three hours in duration with about

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

[4]All major planets and some satellites and asteroids are used, including Io, Callisto, Europa, Ganymede, Titan, Ceres, Pallas, Juno and Vesta. They are selected depending on availability, frequency and angular resolution of the observations, intrinsic flux, etc.

five calibrator observations uniformly spaced over the SB to obtain large parallactic angle coverage. A list of about 15 calibrators will be available to the OT when making the SB and many will be bandpass calibrators. In later cycles, the procedure may change and some of the polarization calibrations could be included in the 'Long-term Effects'. The software in casapy needed for the polarization calibrations are available and tested. For Cycle-2 only continuum linearly-polarized observations at specific frequencies in bands 3, 6 and 7 will be supported. No mosaicing, ACA and TP-array antennas can be used, hence the angular-extent limitation of the emission is HPBW/3.

10.3 Calibrator Selection

Calibrators for short-term effects will be dynamically selected at the time of the observations, in particular for bandpass, gain, amplitude/flux and pointing. ALMA has a master catalogue of possible calibrators, including all their measurements, and structure and polarization properties. Selection of the most convenient calibration at the required frequency, ALMA configuration and observing mode, follows the rules stated in Appendix A.8. A subset of the calibrators are periodically monitored at several standard frequencies to be used as flux standards. It is expected that the contents of the Calibrator Catalogue will keep increasing in volume of measurements and target numbers during the expected operation of ALMA.

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, that the product has reached the desired control parameters outlined in the science goals, that is calibrated to the desired accuracy and is free of calibration or imaging artifacts. For Cycle 2 there are several restrictions that may make it a bit more difficult to achieve this goal. It has therefore been decided that QA will also be done on a "best effort" basis, covering all the major issues affecting data quality. The QA analysis will be based on a calibration plan that specifies which observations must be acquired and at which intervals to monitor system performance and environmental factors as they evolve with time. Furthermore, it will tackle issues related to the merging of data for each science goal taken with different configurations, the inclusion of Compact Array and single-dish data, and the ultimate image quality. Errors introduced by user-supplied parameters, such as incorrect source coordinates, inadequate frequency setting (e.g. an incorrect redshift) 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: Standard data reduction process

QA3: Any additional (project-specific) data reduction after QA2

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

11.1 QA0

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

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

Assessment is performed by AoDs (Astronomers on Duty) at the OSF, using the AQUA software tool, based on semi-real time output of the calibrations (obtained by the real-time TelCal ALMA software) as displayed by QuickLook and the "Calibration Summary" files that are produced at the end of each SB observation or sequence of SB repeats. This information is complemented with reports derived using Monitor and Control display tools to monitor specific parameters not directly tracked by the calibrations (e.g., total power level variations, weather parameters, etc). QA0 metrics/parameters have been selected to check the health of the whole signal path from the atmosphere down to the back-ends. These parameters can be grouped into the following categories:

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

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

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

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

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

The tolerances for these parameters that have been adopted by ALMA for this Cycle are listed in Section 11.6.

11.2 QA1

QA1 tackles array performance parameters which vary slowly (on timescales longer than a week). They are measured by AoDs executing standard calibration SBs created as specified by the Calibration Plan. The QA1-related parameters will, in general, be measured at predefined periods during the month as "Observatory Tasks", 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 (DMG). The product is a set of parameters with errors that are ingested into the TMCDB (up-to-date view of the parameters used in any observation), so that they can then be used during observations.

The tasks that fall into this category are:

Array Calibrations: Baseline measurements, Delays

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

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

11.3 QA2

QA2 deals with QA at the level of data reduction using the Science Pipeline or performed semi-interactively by the ALMA Data Reducers Team. It is only at the stage of data reduction and analysis of data products that some of the science goals set by the PI can be compared with the results (i.e., SNR, dynamic range, etc). During Cycle 2, it is expected that, for the basic ALMA standard observing modes, the automated Pipeline will be used for data reduction and QA2 parameter measurements. These QA2 parameters will be placed onto a Weblog page generated by the Pipeline, which will provide the basis for Pipeline QA2. The QA2 metrics which determine the success of an observation are given in Section 11.6.

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

Calibration Issues:

- $T_{\rm sys}$ and its accuracy: median values and outliers.
- PWV during observations and improvements from applying Water Vapour Radiometry corrections.
- Flux scale calibration: is the absolute flux scale less certain than usual?
- Bandpass calibration: antennas with average gain corrections exceeding 10% or outliers exceeding $3\sigma_{\rm rms}$; rms amplitude and phase across the phase reference source bandpass after applying all corrections.
- Phase transfer and astrometry: Cycle time and sky separation between phase reference and target; typical and extreme (unflagged) phase differences between phase reference scans.
- Phase reference source phase and amplitude rms as a function of time (or, ideally, map noise).
- Calibration consistency if multiple arrays and/or single-dish data are combined (any potential issues)
- Polarization calibration: improvement in rms after applying leakage (D-term) corrections; scatter in calibrated polarization position angle for suitable calibrator; any residual Stokes V in calibrator with no known circular polarization (or other suitable metrics...)

Reduction Process:

- Data flagging: Amount of flagging, reason if substantial.
- Flux densities and dispersions determined for bandpass and phase-reference sources.
- Imaging and cleaning: any issues; self-calibration recommendations; continuum subtraction if relevant

Final Data products:

- Longest baseline, visibility coverage and time on target (after flagging).
- Synthesised beam (spatial resolution) for specified weighting scheme; any recommendations for imaging.
- Spectral resolution and channel ranges used to make sample images.
- rms noise in target images: standard (all-data) imaging (this may be atypically high if there are lines); in images of specified spectral regions. These values are compared with those predicted from data after flagging and with those requested.
- Residual artifacts: Sidelobe Levels, effects of "missing spacings" and possible dynamic range limitations.
- Issues arising from combinations of array configurations and/or total power, e.g. amplitude scale consistency.
- Issues arising from mosaicing and/or contamination by bright sources outside FOV or aliasing in clean.
- Polarization purity, confidence limits (minimum fractional polarization, polarization angle accuracy) if relevant.

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

11.4 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 Helpdesk.

The QA3 process will be triggered by PIs (or ARC personnel) once there exists a Helpdesk ticket reporting any problems with the data products which may reflect an underlying problem with the data, observing procedure or calibration. The ARC receiving the Helpdesk report will retrieve the data from the archive and evaluate the nature of the problem. The assessment by the ARC should include 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 DSO, which will work on solving the problem in collaboration with the ARCs. The result will be communicated back to the reporting investigator.

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

If data which have already been delivered are likely to be affected significantly, the Contact Scientist will be informed to contact the PI (via Helpdesk) with the suggested solution. If the proprietary period for a project ends whilst a problem is outstanding, only unaffected data will be released until the problem is fixed.

If it becomes apparent that the problems are due to deviations from observing parameters approved by the PI during Phase 2, the project may be considered for re-observing. Issues include errors in target position, mosaic spacing, time sampling, etc., or other changes not made for essential operational reasons (e.g. substituting a calibrator found to be too faint).

If the PI reports the issues within 2 months of the data delivery, the proprietary period will be extended by the time interval from the original delivery until the corrected data are delivered. If the issue is reported later than that, the proprietary period will be extended only by the time interval between reporting the problem until the corrected data are delivered.

11.5 The Quality Assurance Report

The Quality Assurance Reports that will be delivered to the PIs will be generated using the ALMA Quality Assurance (AQUA) package developed for this purpose. The basic unit of a Report is the ObsUnitSet, which represents a scientific goal stated by the PI during Phase 1 (project creation and review). An ObsUnitSet will typically contain several executions of SBs. For each execution QA0 and QA1 reports are generated by the AoDs using the information available at the time of the observations, which includes TelCal outputs and other monitoring data (weather, total power levels, Corr GUI outputs, etc). A given execution is only cleared for reduction if it has passed both QA0 and QA1. There will be only one QA2 report for the whole ObsUnitSet generated by the DMG at the end of the data reduction process; this also has to be approved before the data products are delivered to the PI. The final report per ObsUnitSet delivered to the PI will 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) will be added to the Report.

The standard policies for QA0 and QA1 failures are that the observations of those ExecBlocks that failed have to be repeated. Failures to pass QA2 may trigger additional observations if the achieved RMS or imaging parameters are quite far from what has been requested by the PIs. Additional observations may not be possible in exceptional circumstances, such as projects with very tight weather or time constraints. If the available data are insufficient to reach the required sensitivity, but are otherwise of an acceptable standard, they may be released to the PI.

11.6 Pass/Fail Criteria

The QA0 pass/fail criteria that have been adopted by ALMA during Cycle 2 are based on the following:

- **Antennas**: Fewer than 26 antennas available in the 12-m Array or fewer than 7 antennas in the 7-m Array or fewer than 2 Total Power Array antennas. States that will be considered to render an antenna unavailable include issues with the antenna itself (including large pointing scatter of $> 1/10$ HPBW, and bad band focusing with offsets from model of $> 1/5\lambda$ in Z and $>\lambda$ in X/Y), as well as issues with the receiver bands not being available at the antenna for a given observation.

- **Bandpass**: Bandpass calibrator signal too weak, with amplitude wiggles across the spectral window of > 3 dB on the autocorrelations (not due to atmospheric features) or strong CW signals

- **System Temperatures**: More than 6 antennas in the 12-m Array, or more than 2 antennas in the 7-m Array, or more than 1 Total Power antenna with $T_{\rm sys}$ values $> 50\%$ of the others, with differences between polarizations $> 50\%$ for a given antenna or with $T_{\rm sys} > 2000$ K

- **Gain**: 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**: All datasets whose execution failed at some point and contain less than 20% of the expected execution time of a given SB execution

- **Calibrations**: Datasets missing critical calibrations that render them uncalibratable

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

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

For QA2, the main criteria are the achievement of the requested noise RMS in the images (it must be within 10% of the goal), the synthesized beam shape and the calibration quality (phase RMS and absolute accuracy).

Chapter 12

Logical Data Structure and Data Flow

This chapter describes the the data flow process from the observations until data is ingested into the ALMA Archive. It includes a brief description of all the main software subsystems involved in the data acquisition and archiving, as well as a summary of the metadata structure adopted by ALMA.

12.1 Data and Control Flow

This section describes the overall control of the ALMA system and the flow of data. A summary of the main actors and operations involved in the observations is shown in Figure 12.1.

Each of the gray boxes in Figure 12.1 represent an ALMA subsystem involved in the observations. The rest of the boxes, colored according to the actor involved, include labels for the actions performed either by external agents (actors) or by those subsystems.

A typical observing session would be started by the Telescope Operator interacting with the Executive subsystem via one of the dedicated control computers in the Control Room (the so-called "Standard Test Environment" or STE). The Executive subsystem is in charge of starting up the ALMA Common Software (ACS) and its CORBA-based services and then initializing all of the various software subsystems involved in the observing and data storage process. This is done in several cycles to solve interdependencies between the different software components. Once all the components are ready, the Executive also handles asynchronous events from several of the subsystems and responds to them accordingly. Among the events that it listens to, it is worth mentioning the publication of error conditions to the attention of the operator, and the requests for display of the Control, Telescope Calibration and QuickLook subsystems.

The actual observations start by manually creating an array, which means selecting all the antennas that will be involved in the observations. An SB is then selected from the list provided by the Scheduler subsystem, and the execution is started. All the SBs for a given observing 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 have two options for using the Scheduler during Early Science. They can either let the Scheduler suggest possible SBs to execute, or they can carry out targeted searches of the local database. For Early Science, the Scheduler can produce a ranked list of optimum SBs to execute next based on weather conditions and forecast, hardware 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 it can run an automatic sequence of SBs without disagreeing with what an AoD would select.

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

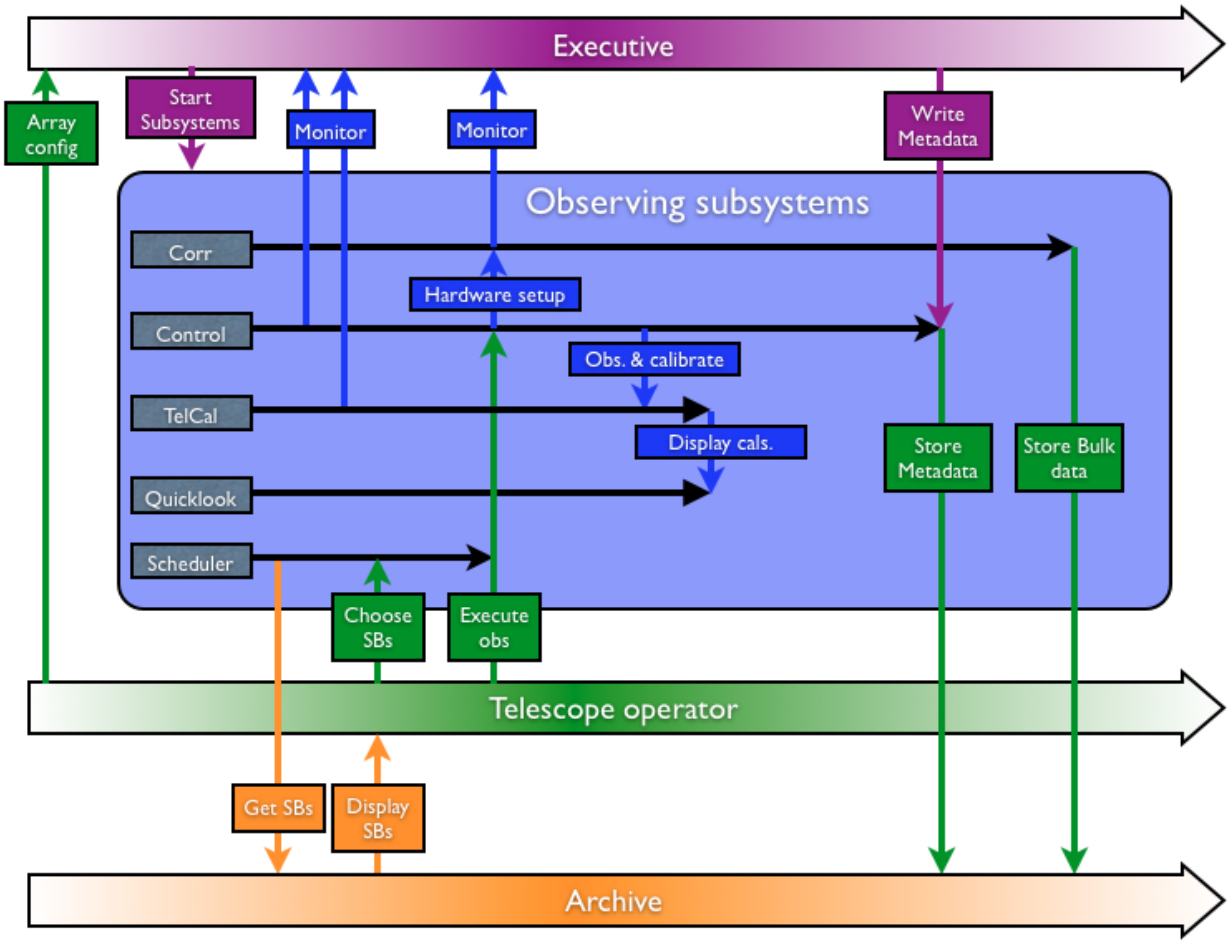

Figure 12.1: Main actors and roles during observations with ALMA. The horizontal direction represents time evolution.

Control executes the SB by commanding all relevant hardware and the correlator, resulting in raw data and metadata being made available to the Telescope Calibration (TelCal) and Quick-Look subsystems that report the results, progress and quality of the calibrations carried out during the observations. To carry out its function, Control has many interfaces to the instrumental hardware. It is in fact one of the truly real-time subsystems within ALMA because it is in charge of synchronization of the actions of all antennas (scanning, source acquisition, etc) and correlator to within 48ms (TE). Control is also in charge of storage of data from all monitor points set in the hardware of the ALMA array.

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

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

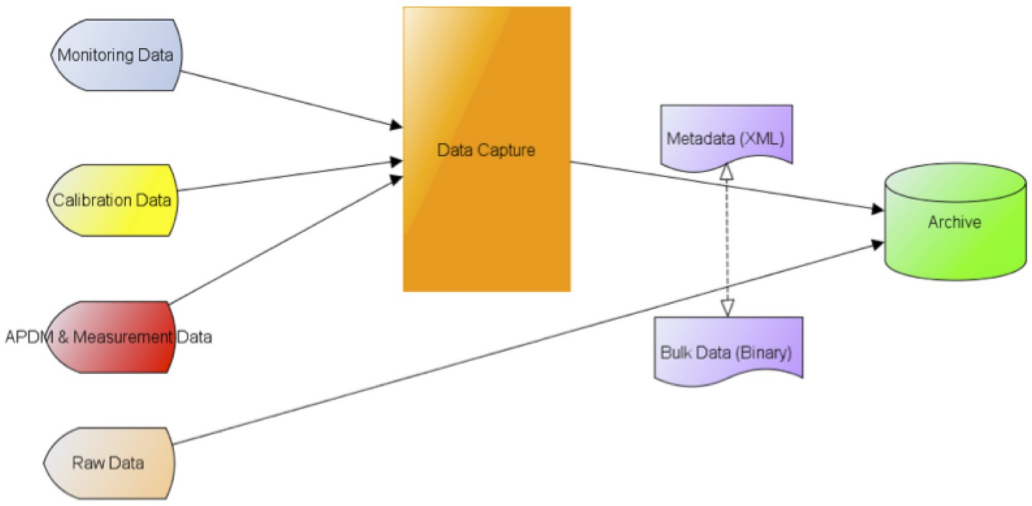

Figure 12.2: Data flow components.

12.2 The ALMA Science Data Model (ASDM)

The 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, the online calibration system (TelCal) also creates associated tables whenever it processes any on the calibrations that can be done on-line. All tables are organized with a similar structure, with the columns listing the contents and the rows including the actual values. The core tables have been defined to outline some of the following: hardware characteristics, array configuration, antenna tracking, the targets, auxiliary monitoring data, overall project and post-processing. A list of the core tables is shown in Figure 12.3. The term *referenced* in the tables means that a given table is linked to other tables via some references, that is, that some of the information is shared between tables.

SDM Tables

Referenced:

Main		
Antenna	Field	SpectralWindow
ConfigDescription	Pointing	State
DataDescription	PointingModel	Station
ExecBlock	Receiver	Subscan
Feed	Scan	SwitchCycle

Not referenced:

AlmaRadiometer	Focus	SBSummary
Annotation	FocusModel	Source
CalDevice	FreqOffset	SourceParameter
DelayModel	GainTracking	SpaceCraftOrbit
Doppler	Holography	SysCal
Ephemeris	Polarization *(required in MS)*	WVMCal

Sometimes referenced:

Beam	required for single dish or mosaicked data
CorrelatorMode	required for correlators; not allowed for others
SquareLawDetector	required for total power or noise detectors; not allowed for others

Mandatory:

SDM

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

The associated tables produced by TelCal all have a name starting by "Cal" and then a self-explanatory string on the type of calibration they are associated with. The list of these associated tables is being upgraded as new observing modes and calibrations become available (see Section 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.

Most users interact with the ASDMs via CASA. 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.

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

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

Chapter 13

Data Archiving

13.1 Introduction

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

13.2 Data Flow

Data from the correlator, together with additional monitor and weather data, are sent via dedicated optical fiber links (1-10 Gbit/s) to the OSF, where they are archived. The system has been designed to cope with 6.6 MB/s average data rates and 66.6 MB/s peak rates for short periods of time, i.e. days. The peak rate will be a hard limit imposed by the Observing Tool at the proposal validation stage.

The archive at the OSF is designed to provide up to a year of temporary storage for the instrumental data (in the form of files in the "ALMA Science Data Model", or ASDM, format) and the monitoring data. The instrumental data are then transferred to the main archive at the SCO, where the pipeline is run and from where the data and pipeline products are distributed to the three ALMA Regional Centers (ARCs) in North America, Europe and Japan (Figure 13.5). 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), see Figure 13.6. All data transfer is done over the network. Metadata replication from SCO to the ARCs happens within a few seconds, transfer of bulk data can take longer (up to several hours), depending on the amount of data to transfer to the individual ARC.

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

The ALMA archive consists of two parts - the ALMA front-end archive (hereafter AFA), which is used for storing all data and metadata of the observations and observing projects, and the ALMA Science Archive (hereafter ASA), which stores a subset of the science metadata in a relational database to enable fast science queries and VO access. The architecture is based on the Next Generation Archive System (NGAS) with Oracle technology for replicating the metadata (see Figure 13.1).

As soon as data is taken and has passed the QA0 quality control step (see Chapter 11.1), its metadata are harvested from the AFA into the ASA and made available for search. This allows archival researchers to see which data will become public in the future. Metadata harvested includes all the information needed to describe the observations, including date and time, target pointings, 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, too, get written into the AFA and the corresponding science metadata are written into the ASA.

13.3 PI Data delivery and data delegation

The unit of data delivery to the PI is the ObsUnitSet (OUS). PIs will receive a notification that data of one of their OUS is available to them for download as soon as QA2 on that OUS has passed and the data products have been replicated to the users' home ARC. This sending of the notification to the PI triggers the start of the proprietary period of usually 12 months for these data (extensions of the proprietary period can be granted under certain circumstances). Only the PI will receive the notification. 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 will consist of the raw data, the FITS files (4D cubes and/or 2D images), logs, scripts, QA information and calibration tables. In contrast to Cycle 0 and 1, Cycle 2 allows access to individual product files. ASDMs are downloadable separately.

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.

13.4 Archive Query

At http://almascience.org/alma-data/archive users can query the holdings of the ALMA Science Archive (Figure 13.2). 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.

Efforts have been made to allow users to query the Science Archive by physical quantities – along the Position, Energy, Time and Polarization axes – as much as possible. 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 science products (Figure 13.3), although they can choose to also see the metadata of the raw data. Users also can choose to be presented the results per observation, per project or per ObsUnitSet. On the results page (Figure 13.4), users can sort and subfilter the results and add or remove columns from the result table to narrow their search.

13.5 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.

Two possibilities exist for the download itself: the first option is to use the download manager (hit "Select all" and "Download selected"). This method is very convenient and allows for parallel downloads. It requires a Java Browser Plugin. The second option is to use the download script that is provided on the same page. This option might be better adapted for downloads to a processing environment where no web browser is available.

The Request Handler also allows to select individual files for download. Only the data deliveries that users have permission to download can be selected. If they are not authorized to access any data delivery in the request, no "Download selected" button appears and the download script will be empty.

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 for data delivery on hard media, i.e. USB hard-disks. The ARCs can have different policies regarding the details of the shipping of the hard-disks.

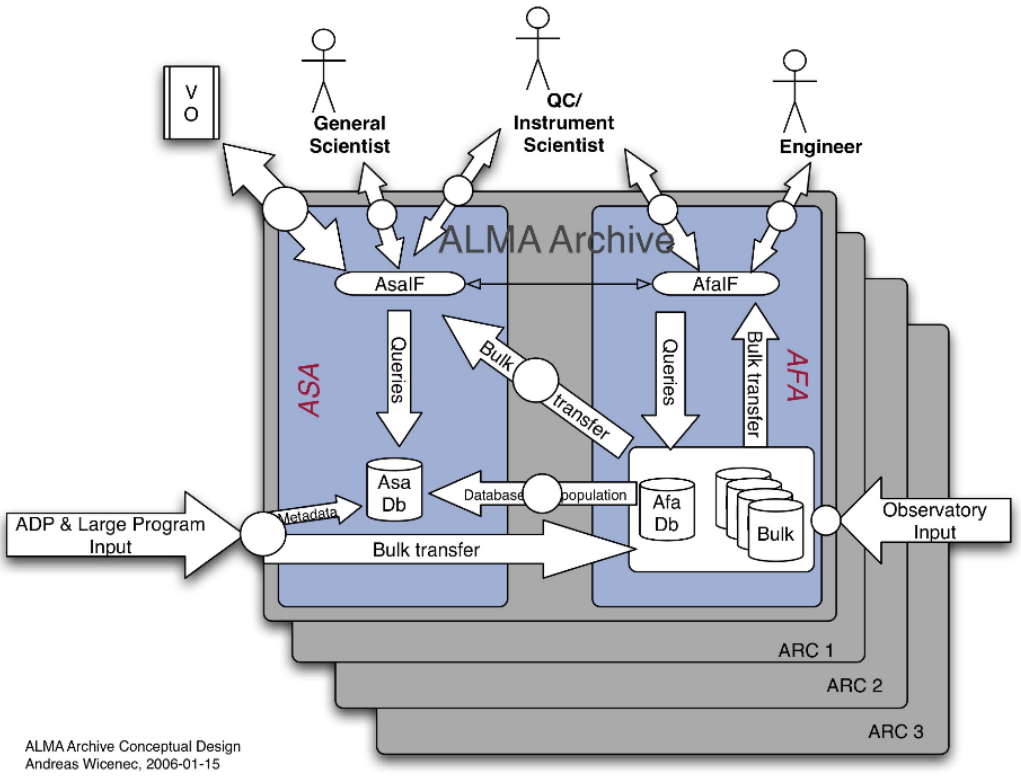

Figure 13.1: Archive design (front end & back end) at the OSF to store metadata and raw and monitor data.

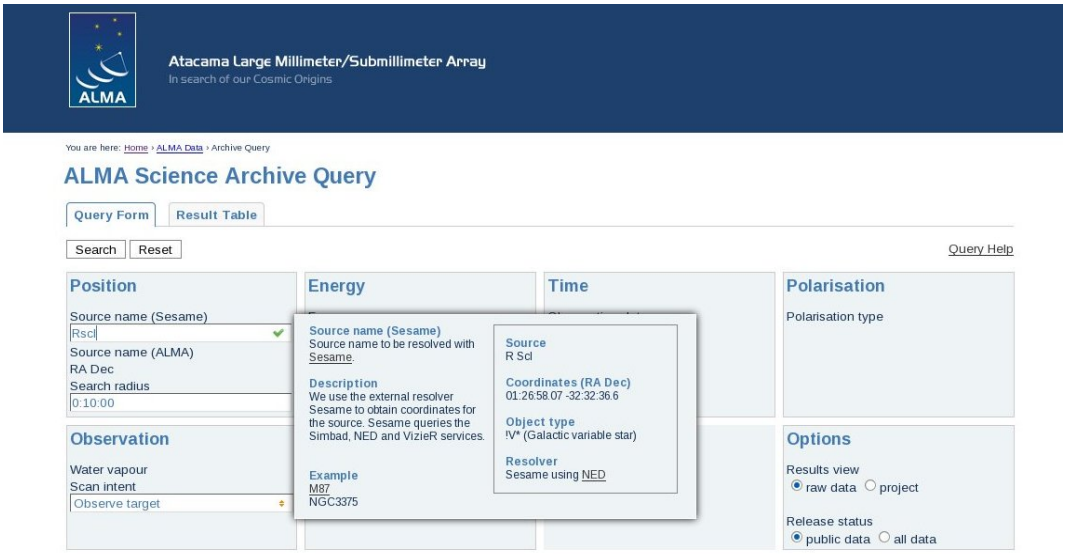

Figure 13.2: The ALMA Science Archive Query Form

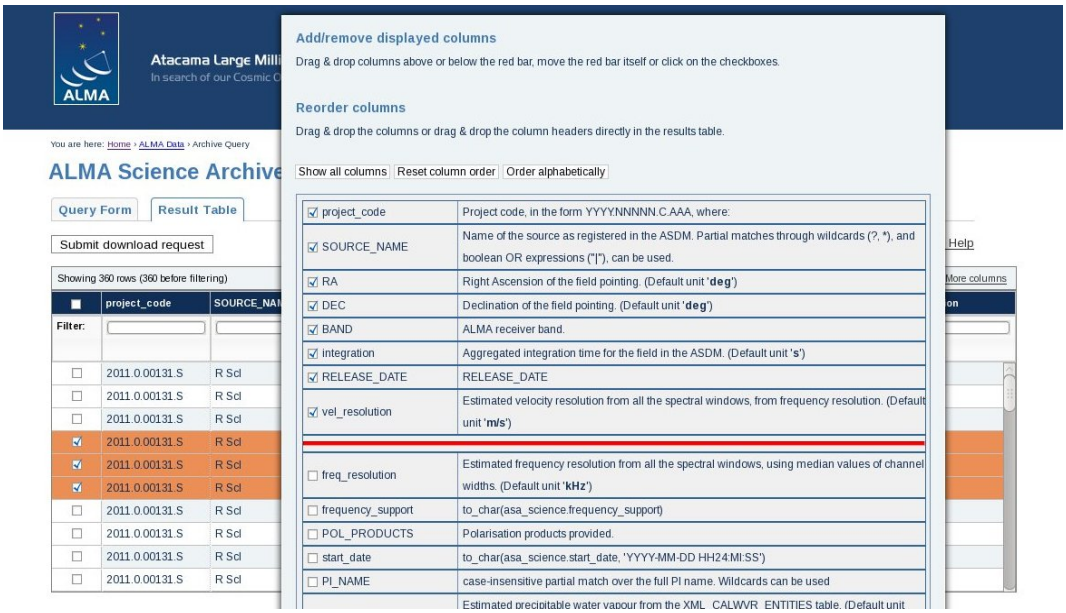

Figure 13.3: The ALMA Science Archive Result Page

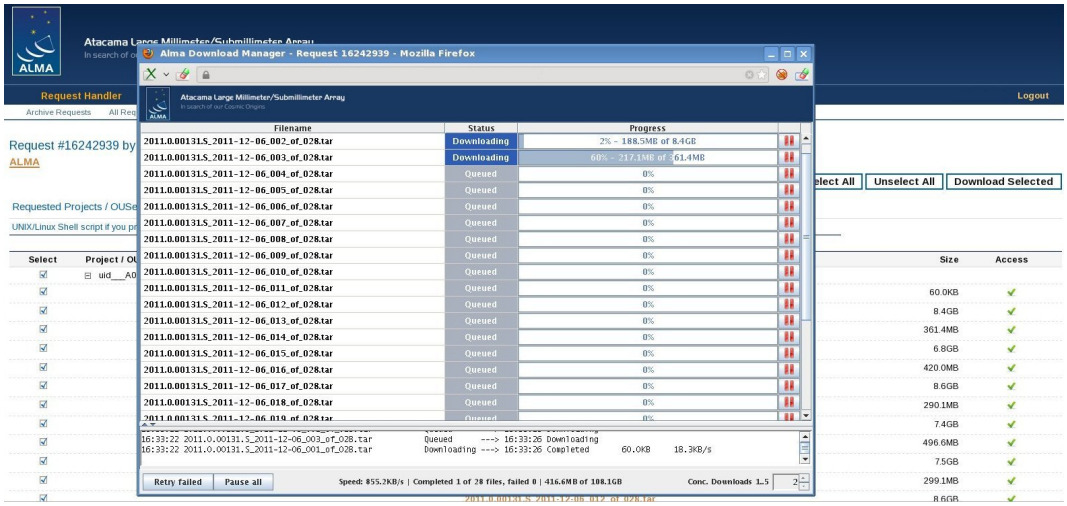

Figure 13.4: The ALMA Science Archive Request Handler

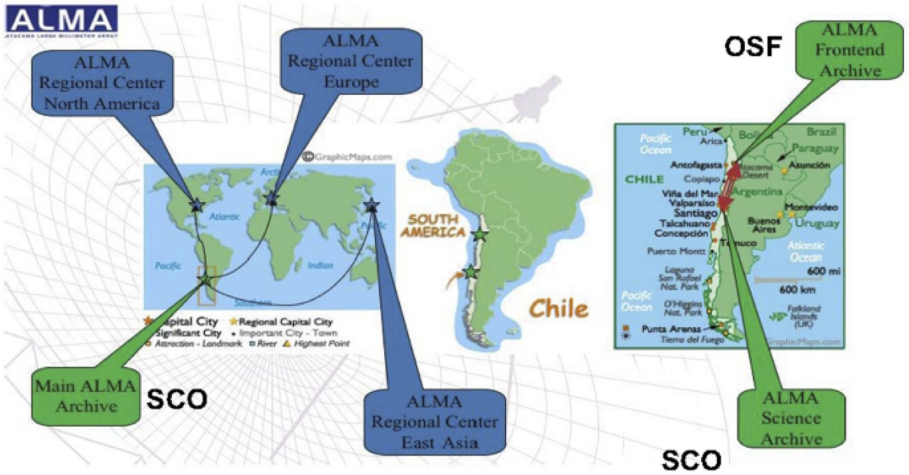

Figure 13.5: Location of the ALMA archives.

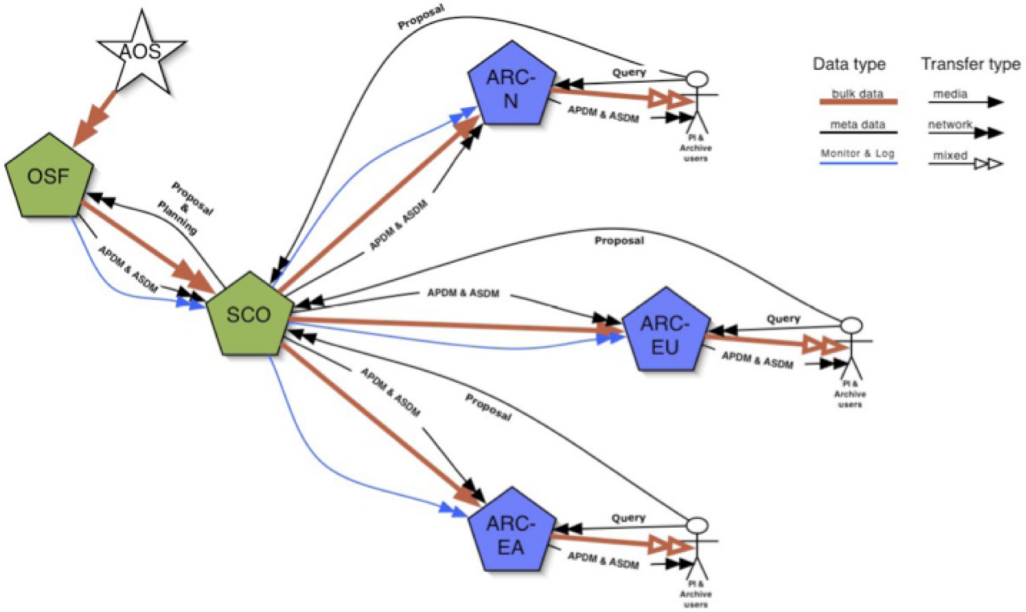

za

Figure 13.6: Data flow from the AOS down to the ARCs.

Appendix A

Appendix

A.1 Antennas

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

All antennas have been designed to meet very stringent ALMA performance criteria, and to successfully operate under the extreme environmental conditions at the Array Operation Site (AOS), i.e. strong winds, large temperature ranges and gradients, solar irradiation and snow. The primary operating conditions 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 2 have both 12 meter and 7 meter diameters, with the receivers mounted at the secondary (Cassegrain) focus. The 12 m dishes have a focal length of 4.8 meters, but the distance from the secondary focus to the plane of the subreflector of the 12 m antennas is 6000 mm, giving an effective focal ratio f/8, with an effective secondary focal length of 96 m and a plate scale of 2.15/arcsec per mm. The subreflector has a diameter of 750 mm. The 7 m dishes have a focal length, to the primary focus, of 2.572 meters. Given an effective focal ratio f/8, an effective secondary focal length is 56 m. The subreflector has a diameter of 457 mm.

The main reflectors of the ALMA 12 m and 7 m antennas are composed of individual panels. The size and number of panels varies between the different types of antennas:

VertexRSI: 264 panels spanning 8 rings with 12 (rings 1 and 2), 24 (rings 3 and 4), and 48 (rings 5 through 8) individual panels which are roughly a half-meter-square in area.

AEM: 120 panels spanning 5 rings with 8 (ring 1), 16 (ring 2), and 32 (rings 3 through 5) individual panels which are roughly one-meter-square in area.

Melco 12 m: 205 panels spanning 7 rings with 5 (ring 1), 20 (rings 2 and 3), and 40 (rings 4 through 7) individual panels which are roughly one-meter-square in area.

Melco 7 m: 88 panels spanning 5 rings with 4 (ring 1), 12 (ring 2), and 24 (rings 3 through 5) panels which are each roughly one-meter-square in area.

Each panel has up to 5 adjustment screws, which can be used to optimize the surface accuracy of the individual antennas (based on holographic measurements). The surface of the panels are etched to scatter optical and near infrared solar radiation.

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

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

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

The different antennas use a combination of steel, aluminum, Carbon Fiber Reinforced Plastic (CFRP) and Invar to achieve the best compromise between stiffness, robustness, smoothness, and low thermal expansion (see Table A.1 for a summary of properties). Common to all antennas is that they have a steel pedestal.

All antennas have builtin metrology systems, which allow thermal and wind deformations to be computed and corrected. For these purposes, the antennas are fitted with thermal sensors, linear sensors and inclinometers (tiltmeters).

The Vertex antennas have a drive system that is gear-driven whereas the AEM and MELCO antennas have magnetically supported direct drives.

The antennas are controlled using the ALMA Control Software (ACS). ACS sends instructions to the Antenna Bus Master (ABM) computer, which are then sent to the Antenna Control Unit (ACU) through a CAN bus.

	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 of such gears and takes the power directly. The direct drives used in ALMA antennas are magnetically supported. **3** Jointly used to correct in semi-real time the pointing of the antennas, under a wide range of environmental conditions, to meet the ALMA specifications.

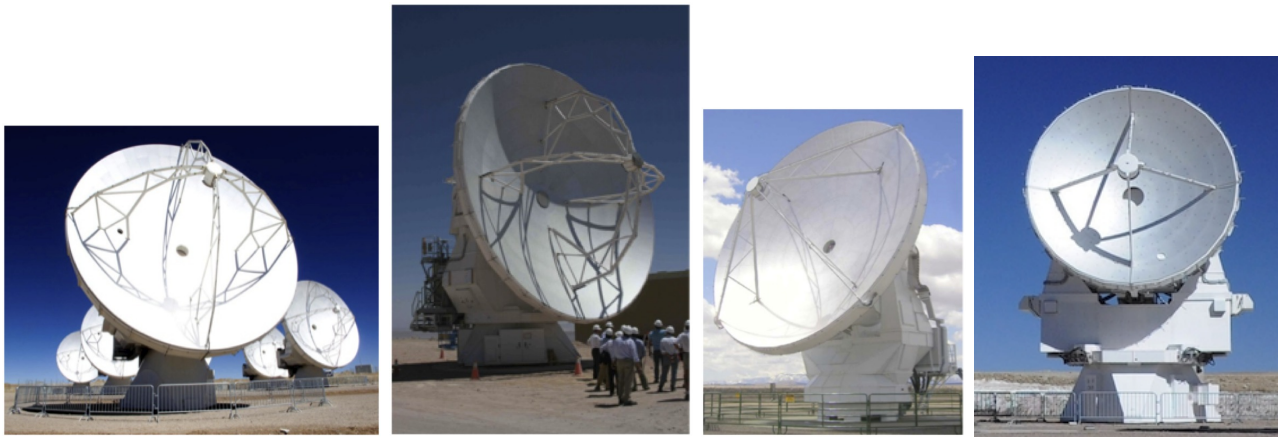

Figure A.1: The four different ALMA Antenna designs: Vertex 12 m, MELCO 12 m, AEM 12 m, and MELCO 7 m (from left to right).

A.2 Antenna Foundations

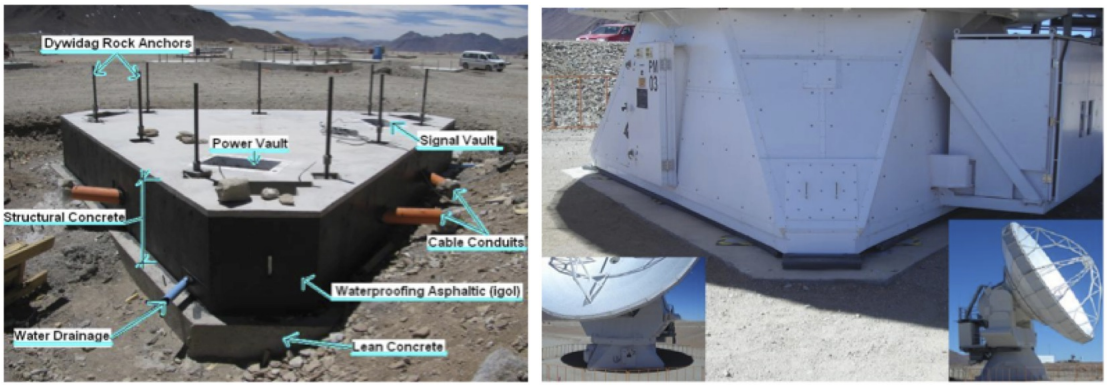

Figure A.2: Structure of an antenna pad (actual pad at the OSF) (left) and detail of antenna anchored to a pad (right).

The antennas are placed on specially-designed concrete pads to guarantee stable orientation and location (Figure A.2). All antennas are attached to the pads at three points at the vertices of a triangle. The three points (inserts) are located on a circle centered at the antenna pad with a spacing of 120 degrees.

This interface guarantees a position repeatability error of the antenna, considered as a rigid body, not exceeding the values below:

- X/Y plane $<$ 2 mm (peak to peak)
- Rotation around Z $<$ 30 arcsec (peak to peak)
- Parallelism with respect to Z +/- 10 arcsec with respect to Zenith

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

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

This stiffness includes the inserts, the concrete pad and the soil. This does neither include the kinematic mount lower part nor it includes the foot of the antenna. The position of the pads are measured to a precision of 65 microns, and then monitored for stability for over two weeks. The pads are equipped with two vaults that contain the power, communication, Local Oscillator (LO) and data transmission cables that are connected once the antenna is placed on the pad.

A.3 Antenna Transportation

Antennas are moved from one pad to another using a specially-designed transporter (Figure A.3, righthand panel). ALMA has two of these vehicles. They are 20 meters long, 10 meters wide and 6 meters high, and each has 28 tires. The transporter positioning system performs a fine positioning of the antenna before setting it down on the foundation in the 3 in-plane degrees of freedom (x, y, rot -z) and in tilt (rot-x, rot-y). Adjustment in each of the 5 adjustment axes can be done independently. The adjustment range of the antenna positioning system compensates for the inaccuracy of the vehicle position with respect to the antenna foundation (which must be smaller 10 cm) 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].

[1] http://www.almaobservatory.org/en/technology/transporters

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

A.4 Cryostat

The ALMA front end consists of a large closed-cycle 4 K cryostat containing individual cold cartridge assemblies (CCA) with mixers and LO injection for each band, along with room temperature electronics for the IF and LO for each band (the warm cartridge assembly, WCA) and fore-optics and entrance windows for each band. The water vapor radiometer (WVR) is mounted to one side of the cryostat using a pickoff mirror to direct the antenna beam into the WVR. The Amplitude Calibration Device (ACD) is mounted above the front end, and is described in Section A.5. Figure A.4 and A.5 show overviews of the front-end unit, with the cylindrical cryostat on top and the room temperature electronics beneath.

All of the receiver cartridges are in the same cryostat, with the mixers thermally-coupled to the same 3-stage Sumitomo cryocooler (Figure A.6). The three stages have nominal temperatures of 4 K, 15 K and 110 K. To avoid overloading the cooler, only three bands can be switched on at a time. It takes about 1 minute to switch between any of the bands that are switched on at a given time. For bands that are off, the time to fully thermally-stabilize them from an off state is 15 minutes – this is mainly to ensure a flat bandpass shape. All of the receivers are mounted off-axis to avoid extra rotating band-selection mirrors, which necessitates a pointing offset of the antenna to change band. The band pointing offsets are known and well-measured; the reference band for pointing is Band 6, and all offsets are with respect to this band. The four higher-frequency bands (Bands 7-10) are mounted close to the central boresight to minimize aberrations.

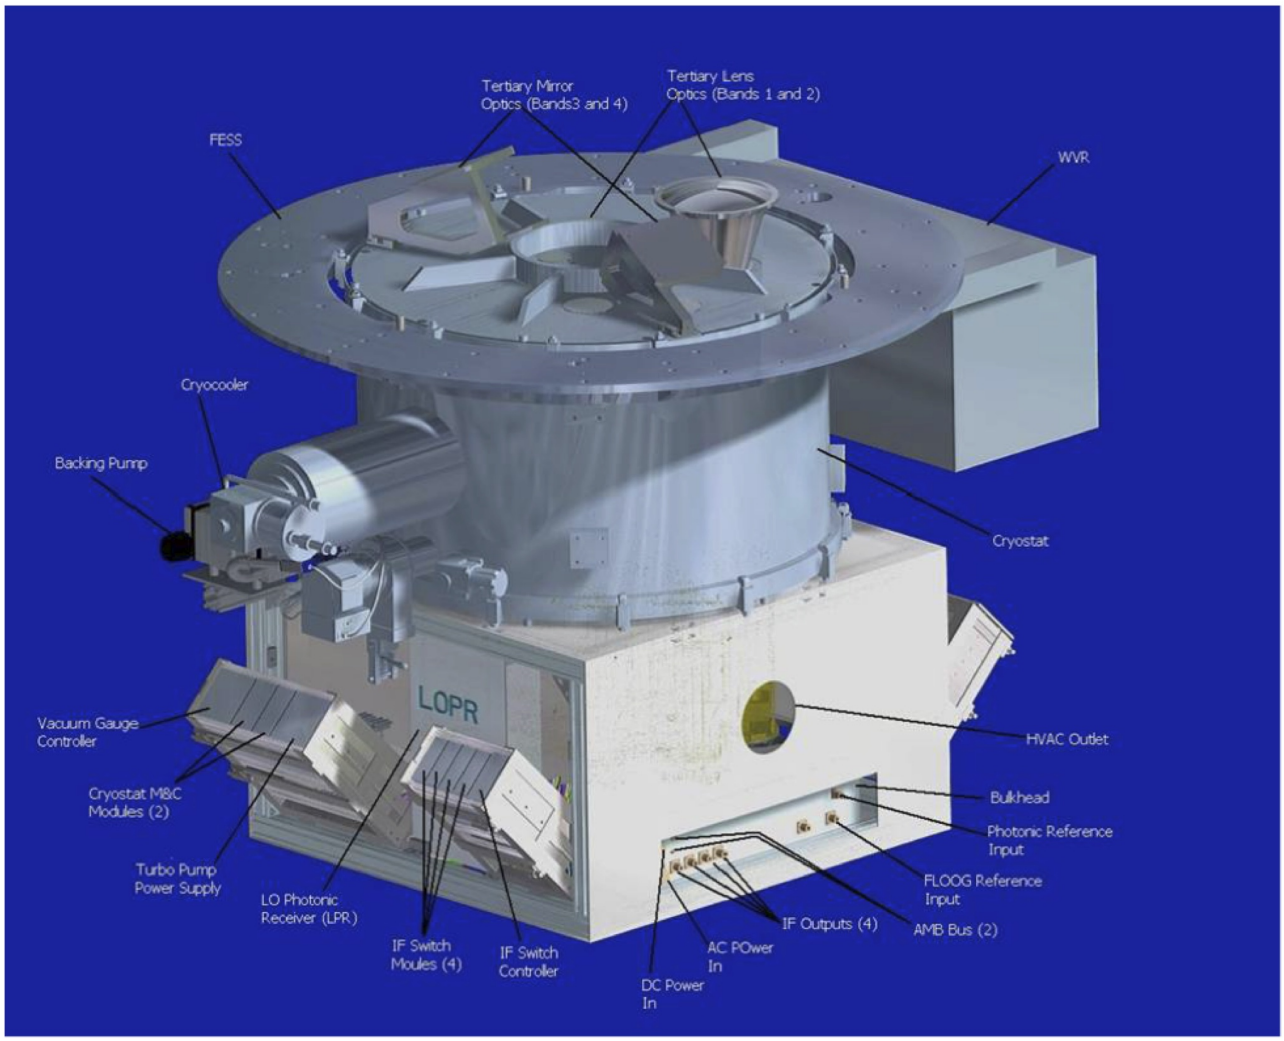

Figure A.4: Side view of ALMA front end showing cryostat assembly, with room temperature unit below.

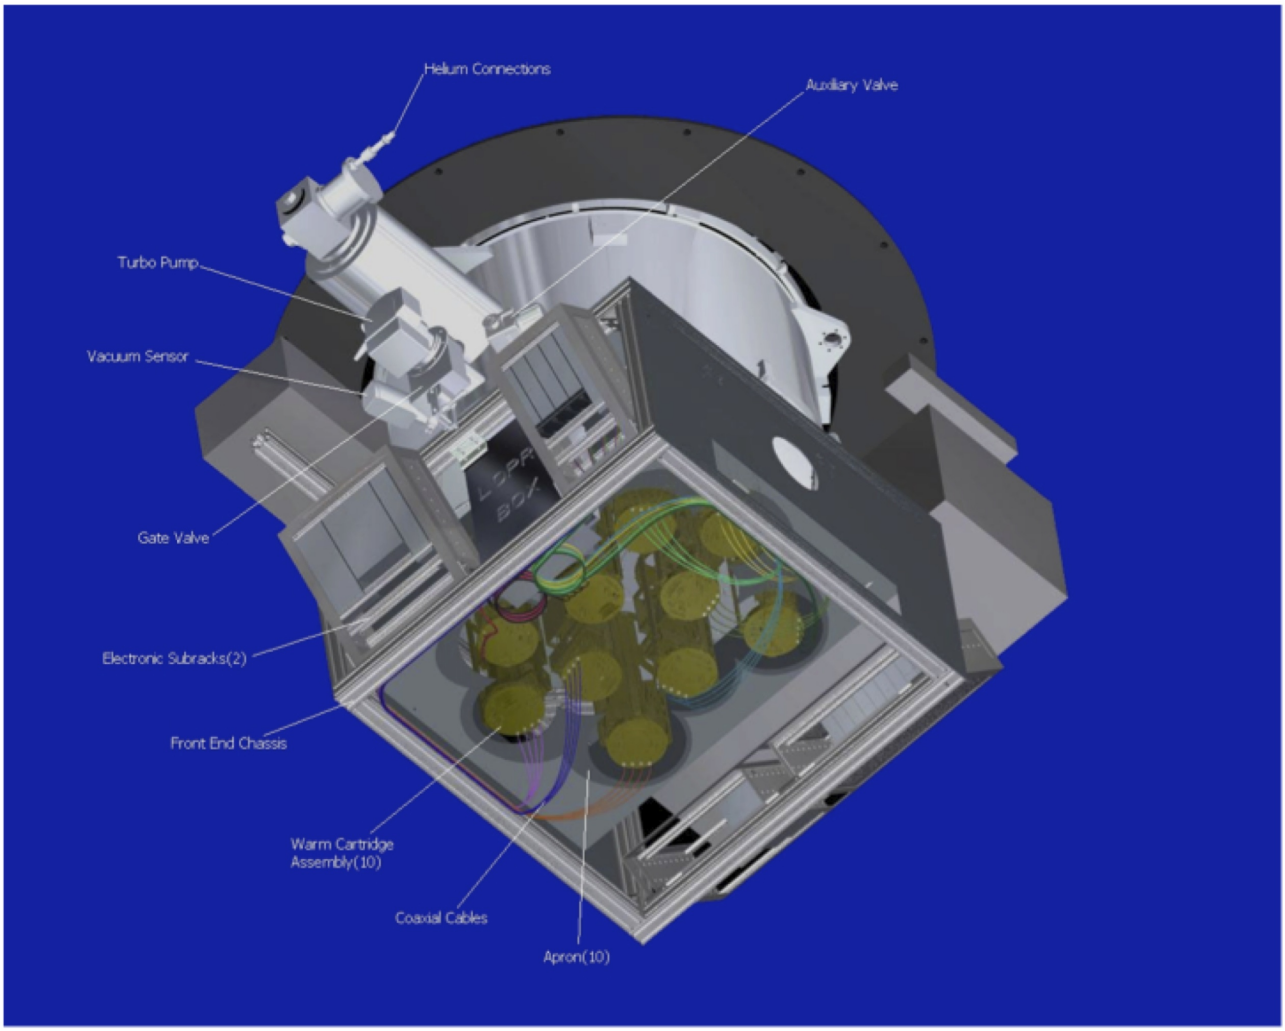

Figure A.5: Bottom view of ALMA front end, showing WCAs.

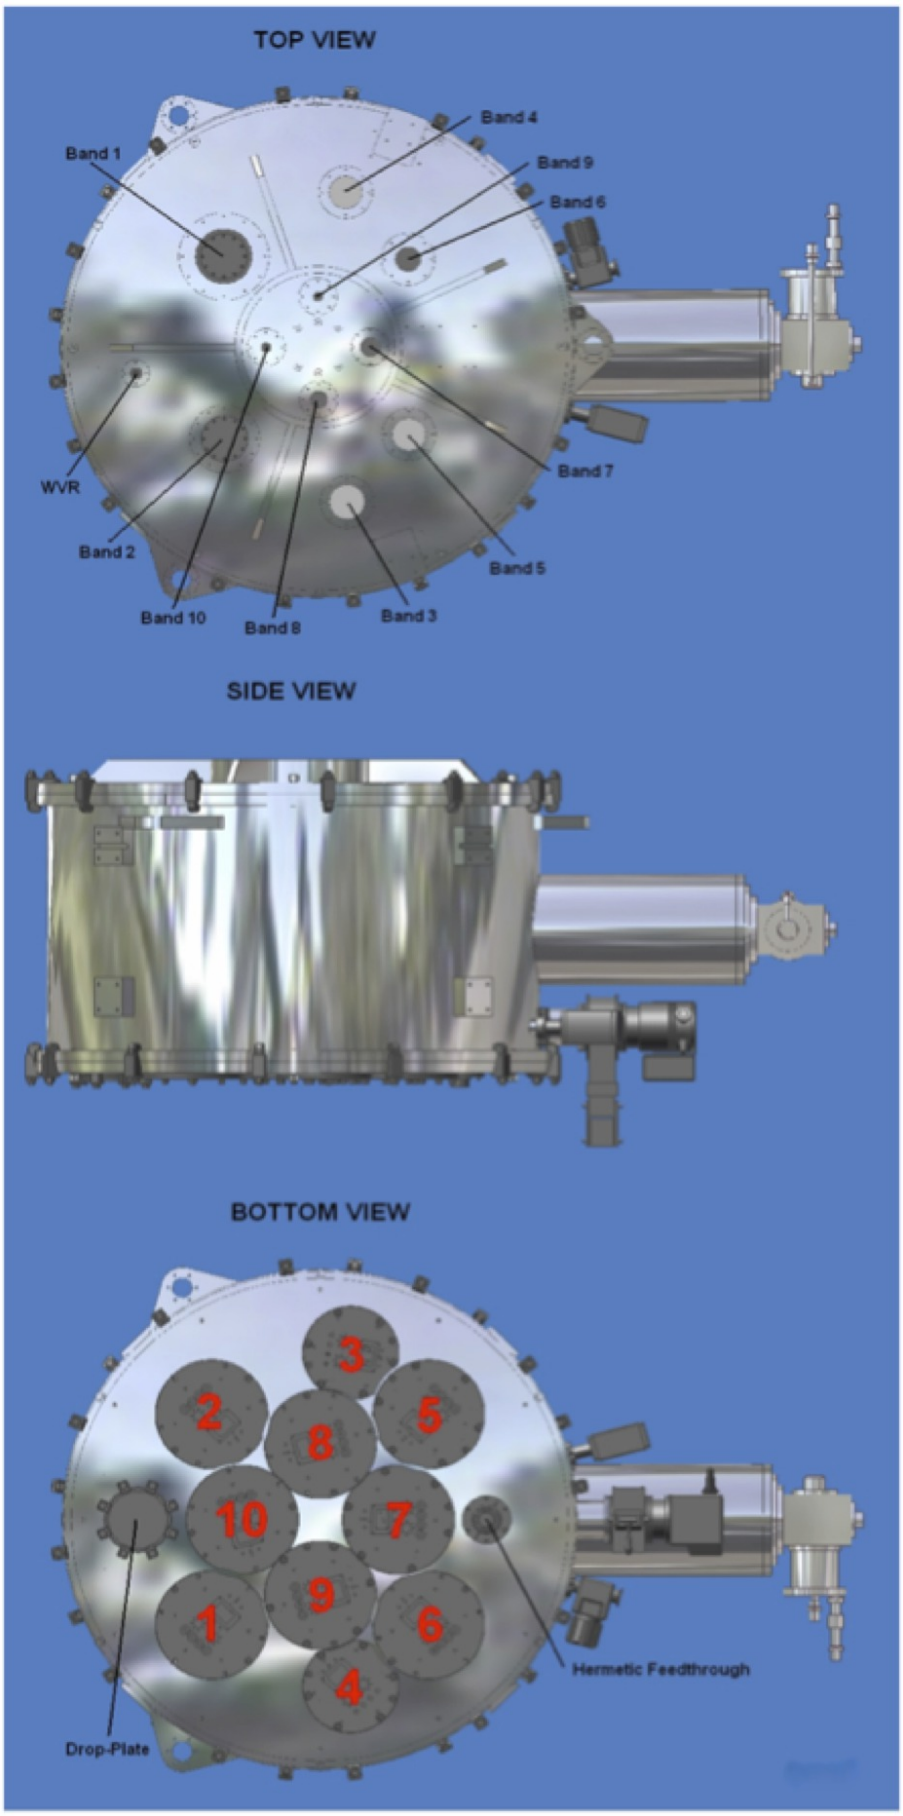

Figure A.6: Views of cryostat assembly, showing different windows (top) and the portholes for the WCAs for each band (lower view).

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 (solar observations are not available during Cycle 2). 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.

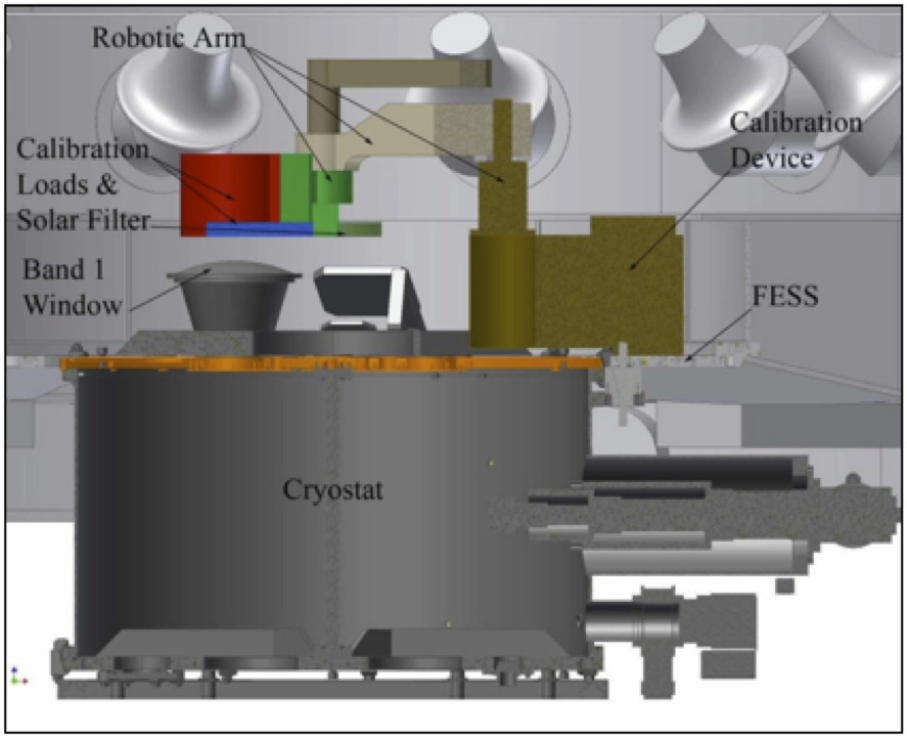

Figure A.7: Lateral view of the ACD on top of the ALMA front ends.

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.).

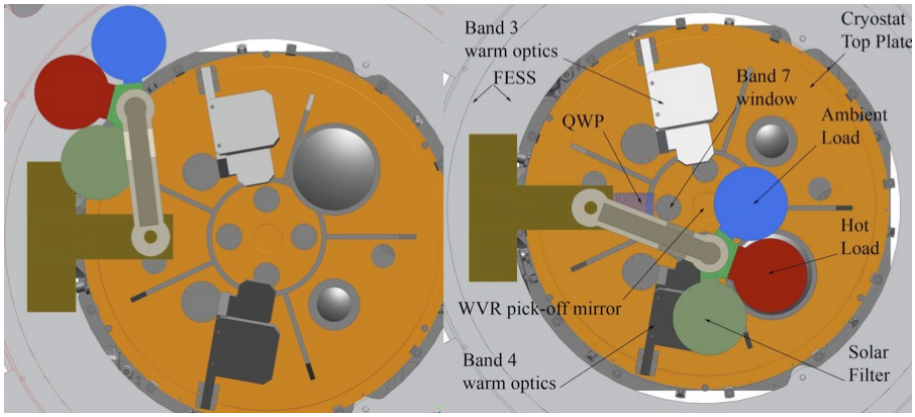

Figure A.8: Top view of an ALMA front end showing the robotic arm of the ACD retracted during normal observations or on top of one of the front end inserts for calibration. The current design has been improved by placing all the loads in a wheel.

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:

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

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

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 2.

A.6 Water Vapor Radiometers

In the mm and submm regions, variations in the water vapor distribution in the troposphere that move across an interferometer cause phase fluctuations that degrade the measurements. ALMA uses the so-called "Water Vapor Radiometry" technique to correct for these phase fluctuations. Water Vapor Radiometry involves estimating the excess propagation path amount due to water vapor along a given line-of-sight by measuring the brightness temperature of the sky at frequencies near the atmospheric water vapor resonances. These temperatures can then be transformed into a path length and the difference between any pair of antennas in the array gives the final phase fluctuations to be corrected for a given baseline. ALMA has implemented this technique by placing a Water Vapor Radiometer (WVR) on each 12 m antenna (The 7 m antennas do not have WVRs). For the WVRs to be effective, the measurements have to be taken with a cadence that is fast enough to map the actual variations in the atmosphere. The relevant shortest timescale is the antenna diameter divided by the wind speed as the path delay is averaged over the whole antenna beam and cannot therefore be corrected at any finer time resolution than that. The effective diameter is about 10 m for the ALMA antennas and the relevant windspeed is usually 10 m/s or a bit less so the fastest necessary sampling speed is 1Hz. On timescales shorter than this 1 Hz timescale, the watervapor path fluctuations are expected to lead to small apparent pointing fluctuations which are analogous to the seeing effects in single-aperture optical telescopes. ALMA selected the 183 GHz line because it is quite bright and allows a more compact design than would the 22 GHz water line. It was decided to measure the temperature of the 183 GHz line in four regions offset from the center using filters of different bandwidths. The positions of the filters are indicated as blue boxes superimposed on the profile of the water vapor line in Figure A.9. The sensitivity specification for the WVRs is 0.08–0.1 K per channel RMS.

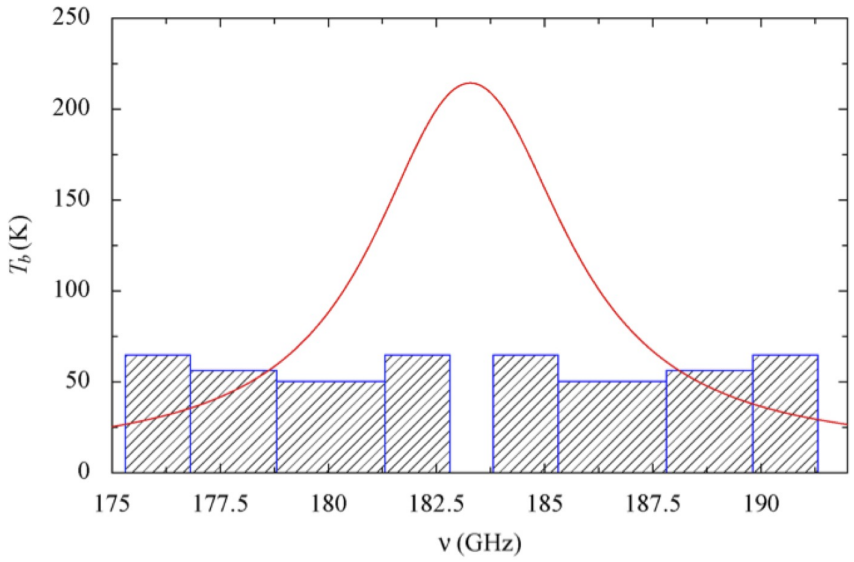

Figure A.9: WVR filters superimposed onto the 183 GHz water vapor emission line.

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

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

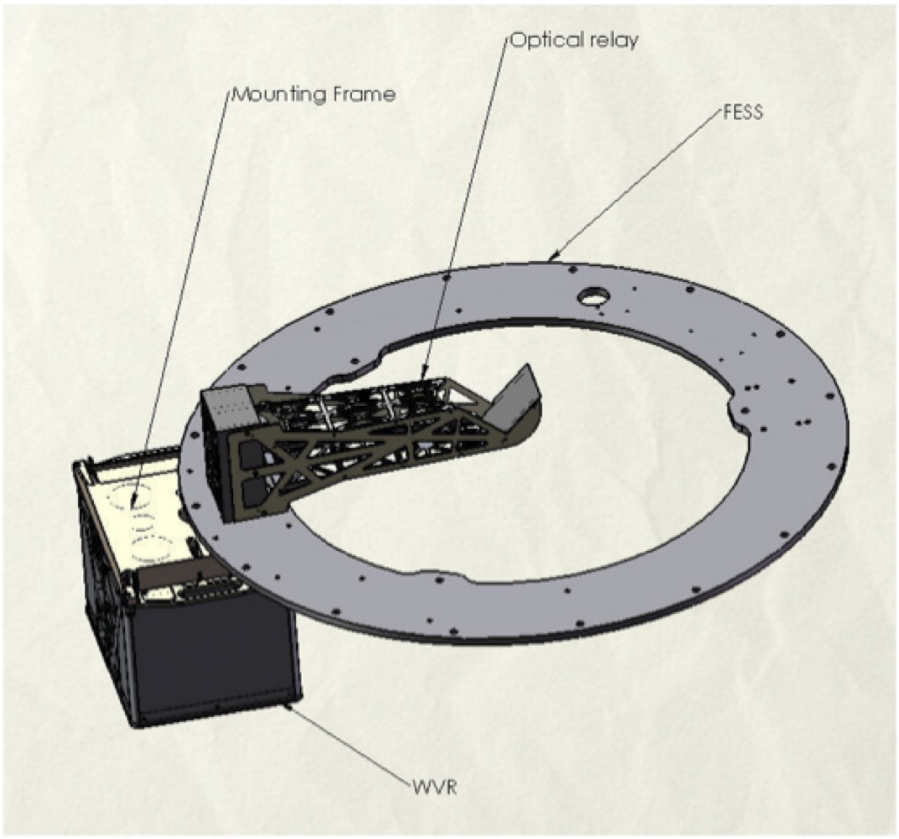

Figure A.10: Offset optics used to collect the sky emission along the optical axis of the antenna into the WVR.

atmosphere (i.e., PWV). There are also variations due to the changes in bulk ambient temperature at different heights above the observatory. It is expected that these could become significant during day time and some techniques are being currently studied to try to measure them (including thermal sounders of the atmosphere that use the profiles of the emission of the oxygen molecules). The brightness temperature variations of the sky that the WVRs have to detect are sometimes quite small, so the quality of the receiving system becomes very important. In fact, the current specification for the ALMA WVRs is that they need to allow corrections of the path fluctuations (in μm):

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

where w is the total water vapor content along the line of sight, and L_{raw} the total fluctuations observed at any given time. Therefore, this formula includes the expected error of about 2% in measuring the total fluctuations, and states the total resulting path errors after correction (L_{corr}). For a 1 mm PWV, the residual term in the formula would be 20 μm. The stability specification for the WVRs is very stringent (0.1 K peak-to-peak over 10 minutes and 10 degree tilts). To achieve this, a Dicke-switching-radiometer approach was adopted. The input into the mixer is switched periodically (5.35 Hz) between two calibrated loads (the "cold" and "hot" loads at 293 K and 351 K, respectively), and the sky using a rotating vane embedded in the light path as shown in Figure A.11.

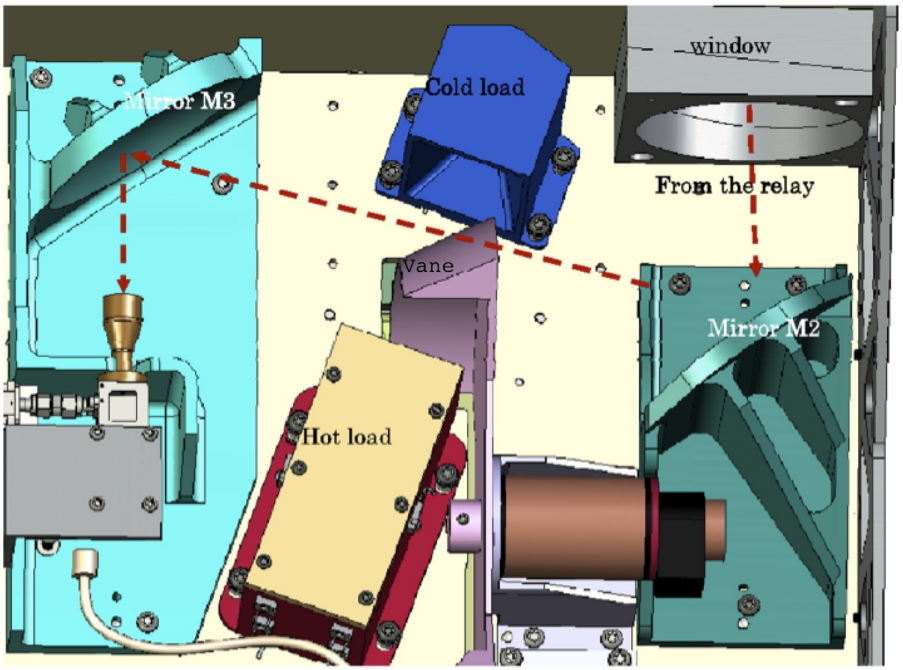

Figure A.11: Optical layout within the WVR encasing, showing the loads, the chopper vane and the input feed to the mixer.

Calibration of the measurements is done following the usual method for a 2-load system. The ratios of the output powers when observing the "hot" and "cold" loads can be used to determine the receiver temperatures. Furthermore, these output powers from the loads are also used to extrapolate to a virtual load that has a brightness temperature similar to that of the atmosphere. The specification for the absolute accuracy of the calibration is 2 K (maximum error). The mixer system is an un-cooled DSB Schottky diode pumped by an LO at 15 GHz that undergoes 2 stages of multiplication. The receiver noise temperature is about 1000 K. After amplification, the IF signal is split into four complete chains (one per filter) and a bandpass filter is applied to select the four desired sampling regions in the profile of the water vapor emission line. In each IF chain, the signal is detected with diodes and after a Voltage-to-Frequency conversion, sent to the Control section for accumulation and control. There is a possibility of LO leakage out of the WVRs that could affect

the ALMA receivers in the same antenna and others nearby. To avoid coherence, all the WVRs are tuned to a frequency slightly different (offsets by consecutive integer multiples of 10 kHz up to the total number of WVRs available). The final product sent to the ALMA Control system are time-stamped, calibrated measurements of the brightness temperatures in the 4 filter regions. The path length error due to the PWV can be calculated from these brightness temperature measurements and used to correct the data. It is envisioned that corrections at the scales of the sampling rates of the WVRs will be possible at the correlator and that refinements for longer timescales will be done offline in CASA using the wvrgcal tool.

A.7 The LO and IF System

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

1. Downconversion of the observed sky frequencies to a band in the range 2–4 GHz, which aliases down to 0–2 GHz for digitization.
2. Setting the correct power levels into the digitizers
3. Setting the center frequencies of the spws in the Correlator FDM modes (this is actually done in the correlator, but is effectively part of the LO system).
4. Application of frequency corrections for fringe rotation, and compensation for the slight differences in the Doppler shifts at each antenna due to the differential line-of-sight velocities with respect to the target.
5. Providing geometric delay corrections.
6. Suppression of the image sideband or, in the case of DSB receivers, selection of the wanted sideband(s). This is done through frequency offsets and phase modulation at each antenna.
7. Suppression of spurious signals and reduction of the effects of DC drifts in the samplers. This is done using phase modulation of the LOs.

Frequency downconversion takes place in several stages and involves 2 hardware Local Oscillators (LO1 and LO2), a 4 GHz sampler and an LO synthesised in the tunable filterbanks (TFBs) in the Correlator[3]. An overview of the operation is given in A.7.1. More detailed descriptions of the individual subsystems are given in Section A.7.2.

A.7.1 Summary of operation

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

1. The 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 all the ALMA bands, since the mixers for Bands 3 and 7 have an output IF of 4-8 GHz, Band 6 a range of 5-10 GHz and Band 9 a range of 4-12 GHz. Over most of the 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 component for all antennas, plus a small offset component generated in the FLOOG (First LO Offset Generator) which is different for each antenna (see LO1 Section A.7.2). The FLOOG is used to perform coarse fringe tracking (i.e. rough correction for the small offsets in observing frequency at each antenna), to offset the LO1

[3]Note that the 12-input correlator effectively looks like the 64-input Correlator to the end user, although it does not operate in the same way with TFBs

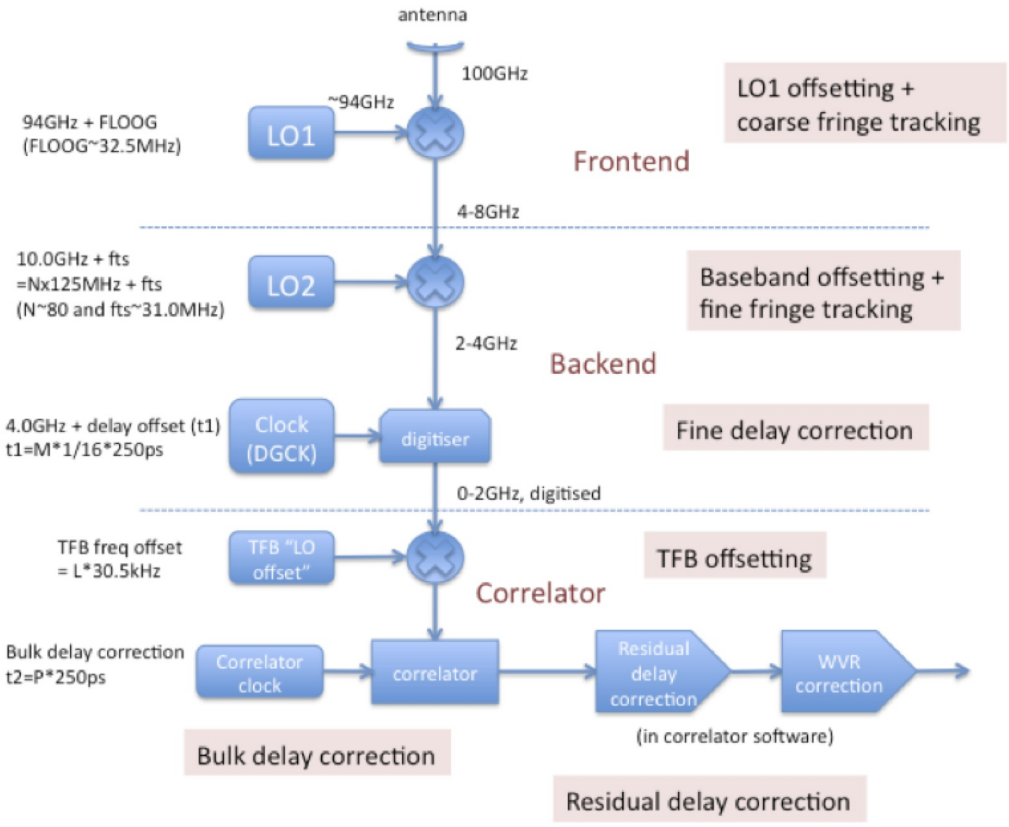

Figure A.12: Overview of ALMA frequency downconversion, LO mixing and delay corrections. This takes place in the Frontend, Backend, and Correlator. Example frequencies are given for an observation at a sky frequency of 100 GHz in the USB. Some LOs (e.g. LO1) are continuously tunable; others have quantized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

frequencies slightly to suppress internally-generated interference, and for sideband separation or selecting the sideband. It is also used to offset the LO1 phase (by 180 or 90 degrees) in conjunction with a Walsh switching pattern on the antennas to remove DC systematic errors, and will also be used for sideband suppression and, for the DSB receivers, for sideband separation.

2. In the "backend" (BE), the IF processor (IFP) splits the IF into basebands, each with frequency range of 2-4 GHz, via a set of filters and tunable second LOs (LO2) (see IFS/IFP section in A.7.2). LO2 is used to offset the individual baseband frequencies within the IF range. The LO2 and second mixer only operates in LSB, with an LO2 tuning range of 8-14 GHz. LO2 is generated by a course synthesiser which is set only in steps of 125.0 MHz, plus a second fine-tuned synthesiser (fts) which provides an offset in the range 20.0-42.5MHz (marked as "fts" in Figure A.12). The limited fts range and the 125MHz quantization means that LO2 setting is not fully contiguous; consequently there can be up to ∼30MHz difference between the desired and the set value. With a single baseband, this can be compensated by a suitable offset of LO1, but with multiple basebands this is not always possible. Without additional correction, setups with multiple basebands may have the requested lines offset from the spw center by up to 30MHz; this becomes significant in high-resolution narrowband correlator modes. However, the remaining differences in the different spws are compensated for by applying an opposite offset to the TFBLOs (LO4 - see item 4 below.)[4]. 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 frequency from the center of the spw. The finely-tunable fts is also used for fine fringe tracking. LO2 can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband.

3. The analog IF signal from the second mixer is sampled with a 4.0 GHz clock (DGCK). A fine delay (or time) offset is applied to this clock in units of 1/16 of the clock period (250 ps) (the "fine delay correction").

4. In the FDM correlator mode, digital filters (known as TFBs, or "Tunable Filterbanks") are applied to the signal, each of which can be individually offset in frequency (the TFB offsetting). This is effectively applying a digital LO (the TFB LO, or LO4), which is adjustable in steps of 30.517578125 kHz[5] and allows the spectral windows to be moved around within the basebands. At phase II of the Obsprep, the TFB is centered on the baseband if the TFB "offset" is set to the default of 3000.0 MHz; it can be moved +/-900 MHz from that frequency, the range depending on the spw bandwidth. The TFB outputs are resampled and sent to the correlator. The TFB LO can also be used to offset the frequencies in conjunction with LO1 to suppress interference and select the sideband. The correlator is used to perform the finest level of residual delay correction.

A.7.2 LO/IF components

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

The First Local Oscillator (LO1)

The reference signal required to tune LO1 in the receivers is obtained as the difference of the wavelengths of two infrared lasers, the Master and Slave lasers, the Master Laser (ML) at a fixed wavelength of 1556 nm and the tunable Slave Laser (SL) which is offset from the Master Laser signal, and generated in the CLO (see Section A.7.2). The offset frequency can be anywhere in the range 27 – 122 GHz. The beat note from the two lasers constitutes the Photonic LO Reference, and the LO1 reference signal is produced from this by photomixers located in the Warm Cartridge Assembly (WCA) of each receiver. The reference signal is used to drive a YIG

[4]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

[5]The OT "adjust" button quantizes the value entered by the user to this unit

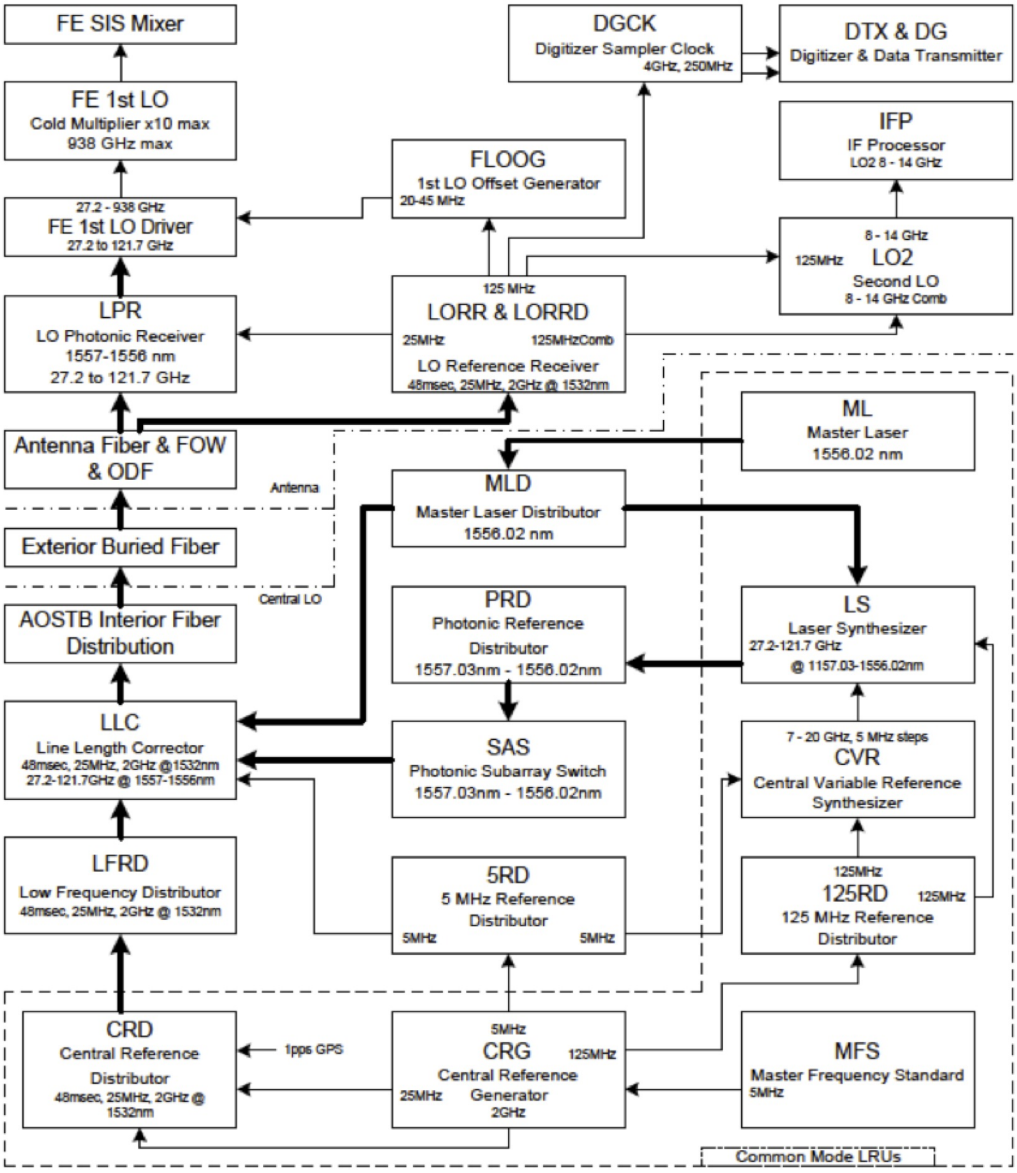

Figure A.13: Summary block diagram of the LO distribution system.

(Yttrium Iron Garnet) oscillator via a Phase Locked Loop (PLL) circuit. This produces the LO1 signal via a set of multipliers (see example Figure A.14 for Band 7). The same photonic reference signal is distributed to all antennas in the same sub-array. However, to correct for different delay rates required in different antennas, the First LO Offset Generator (FLOOG) in each antenna generates a variable offset frequency of 20-45 MHz which is also fed into each PLL.

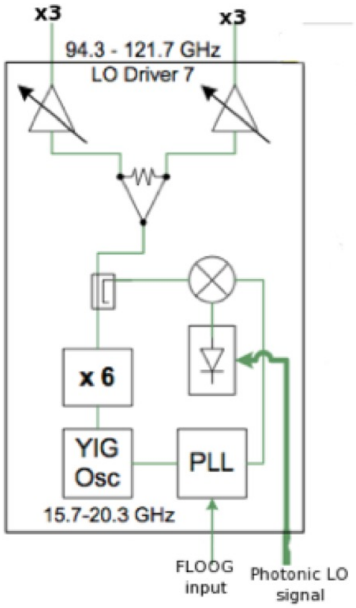

Figure A.14: Block diagram showing generation of LO1 in a WCA - in this case Band 7 (diagrams for the other bands are shown in the description of the individual bands). Note that an additional multiplier (in this case, x3) is used to generate the LO1 frequency, at 282.9 – 365.1 GHz. The photonic LO signal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

The IF switch and IF Processor units

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

The output of each front-end cartridge is connected to a IF Switch unit (IFS) situated in the Frontend, which selects between bands, provides some amplification, and has variable attenuators to set the output levels.

The four (or two) outputs from the IF switch unit are fed into two IF Processor units (IFP), one per orthogonal polarization (Figure A.15). The IF processors divide the incoming 4-8 GHz IF bands from both sidebands into four 2 GHz basebands and downconvert them to the 2-4 GHz range using the second LO (LO2). Since each baseband is fed by a separate LO2, it is possible to locate them at different frequencies within the output bandwidth of the receiver (see Section 5 and Table 5.2 for limitations). The LO2s are common to both mixer polarizations which means that both polarizations will have the same spectral setups.

The IF processor unit has Total Power detectors for tuning/optimization of the IF power levels into the digital samplers, although they are are not used for observations. Figure A.15 shows a basic block diagram

of the IF Processor. It is important to note that the switch network means it is NOT possible to select IF configurations with one baseband in one sideband and three in the other (except for DSB receivers, where this is done using sideband selection).

The LO2s can be used for sideband separation when combined with the first LO (Section A.7).

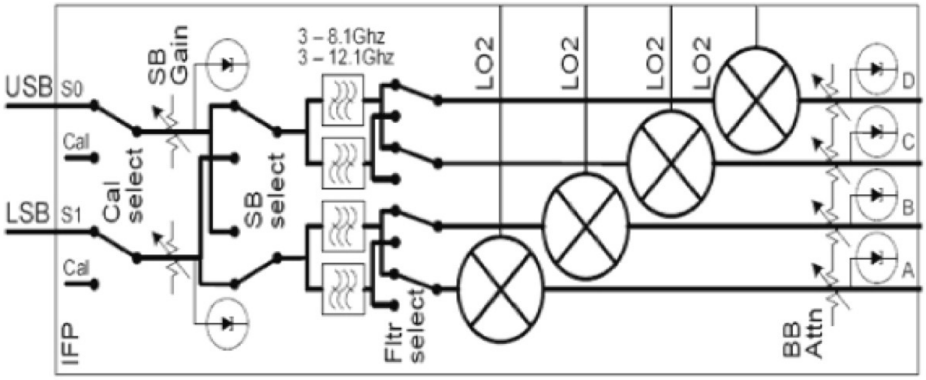

Figure A.15: Block diagram of IF Processor.

The IF processor also has anti-aliasing filters, which define the 2 GHz baseband width and remove out-of-band signals (Section A.7.2). This results in the higher noise levels on the upper and lower 50-100 MHz of channels in the TDM correlator mode (see Section 5.5). These filters cause a decrease in the effective IF range to approximately 1.875 GHz.

Digitization and Transmission

The outputs of the IF Processor units are fed into the Data Transmission System modules (DTS), that include digitizers and formatters to convert the signals to optical wavelengths for transmission via optical fibers. There are four DTS units per antenna, each one handling data for a given baseband pair (i.e., the same 2 GHz baseband from each of the two orthogonal polarizations). Each baseband is digitized by a separate digitizer at 4 GHz (i.e., Nyquist sampling for a 2 GHz bandwidth), quantizing each sample into 3 bits (8 levels) per polarization, so that a total of 6 bits must be transferred per baseband pair. The digitized signal is then transferred to the formatter part that packages the data in frames of equal size. The output of each DTS module is fed to three optical fibers, each transporting 2 bits, and the signal leaves the antenna after passing through a Fiber Optic Multiplexer (FOM). All DTS modules are fed with reference/timing signals from an associated Digital Clock (DGCK), which is also used to do the fine delay tracking.

The outputs of the DTS are sent, via the optical fibers, to the AOS Technical Building where the process is inverted (conversion from optical to digital signal) at the DRXs (Data Receiver units), before the signals are sent to the correlator. Delay corrections due to changes in the length of the optical fibers are done using metadata information to realign the frames sent from the transmitting side at the antenna (DTX) and the receiving side at the Technical Building (DRX). Figure A.16 shows a block diagram of a single DTS module.

Reference and LO Signal Generation and Distribution

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

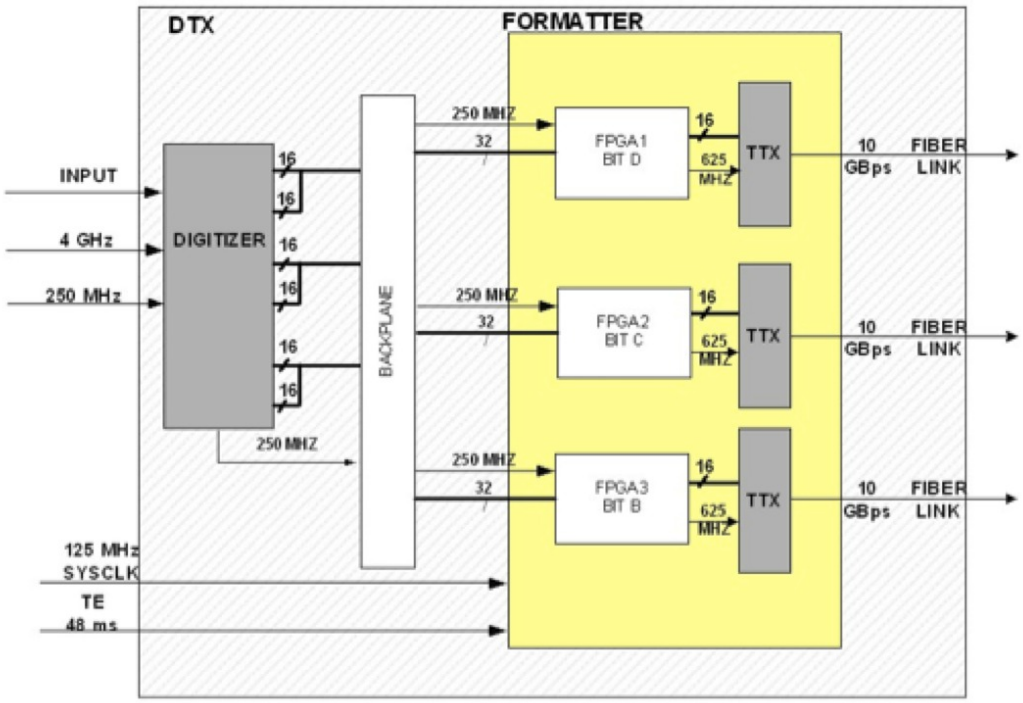

Figure A.16: The DTX and DRX Signal Digitization and Transmission system.

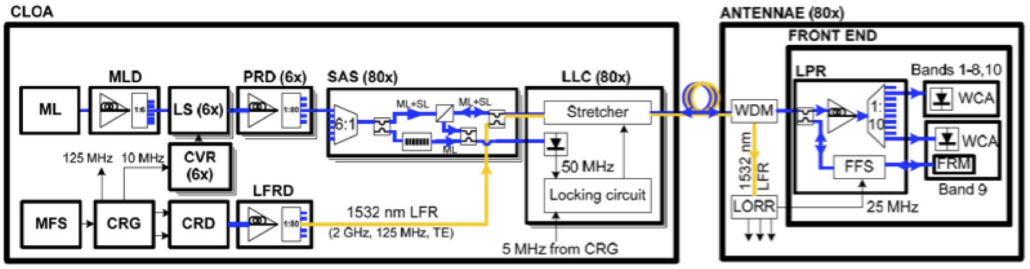

Figure A.17: LO block diagram, showing the Central LO (CLOA) and the LO section in the WCA in each front end. For description of acronyms, see text.

Reference Signal generation

The ALMA frequency and phase standard is a Rubidium atomic clock, the Master Frequency Standard (MFS), which produces a signal at a frequency of 5 MHz. This signal is fed into the Central Reference Generator (CRG) module, which produces several signals as multiples of the 5 MHz signal. The 5 MHz signal is fed into the Line Length Corrector (LLC, See Section A.7.2). The 125 MHz signal becomes one of the standards used by many components in the ALMA system. At the AOS Technical building it is used by the Slave Lasers in the Laser Synthesizer modules (Section A.7.2). At the antennas, it is fed into the FLOOG, the Digital Clock (DGCK, see Section A.7.2) and the LO2 Synthesizers. The 2 GHz signal goes into the DTS cards at the antennas. All these reference signals are modulated into a 1532 nm IR laser in the Central Reference Distributor (CRD) module. The CRD has an internal 48 ms (TE) clock that is also modulated into the same signal. The modulated 1532 nm signal is sent to an optical distributor (with 80 outputs), the Low Frequency Reference Distributor (LFRD), that feeds it into the Sub Array Switch (SAS) modules, where it is merged with the signals from the Master and Slave lasers (see Section A.7.2).

LO signal generation

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

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

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

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

Optical Signal Distribution

Figure A.18 shows the three laser signals after combination in the Sub Array Switches (SAS). The Master and Slave laser signals have wavelengths of about 1556 nm and the laser carrier signal for the reference signals a wavelength of 1532 nm. The signals are distributed via a single-mode fiber optic line to each of the antennas. The fibres are distributed in buried trenches, and fed into the Cassegrain cabin on each antenna through Az and El fibre wraps.

LO Path Length Corrections

The LO Reference signals are generated at the AOS Technical Building and distributed via optical fibers to all the antennas. To guarantee that the phase of the LO signals is stable during the observations for fibers

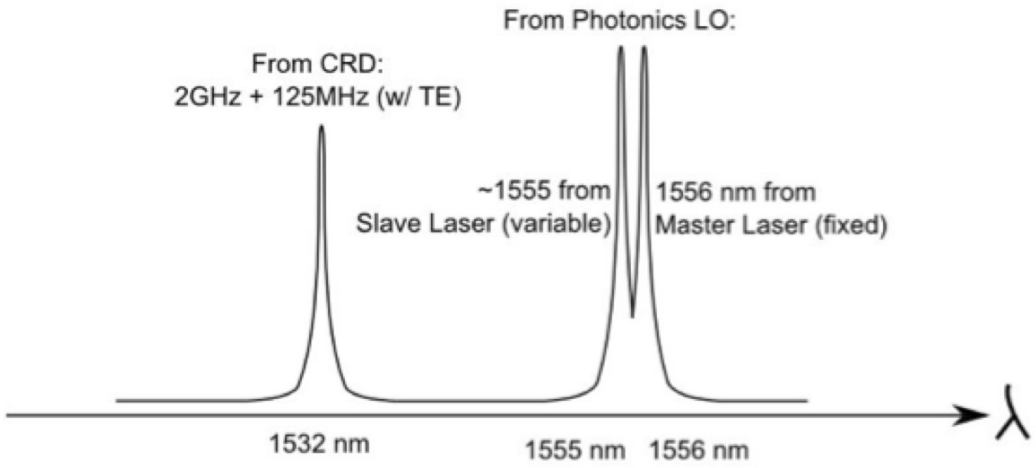

Figure A.18: The ALMA reference signals. Within each antenna, the optical fibers are split and fed to both the LO Reference Receiver (LORR) for the demodulation of the reference/timing signals, and the LO Photonic Receiver for the LO Reference signals.

of up to 15 km in length, compensation for the changes has to be done in real time. The method adopted by ALMA is based on a round-trip optical interferometer. Phase fluctuations for an optical fiber transmission system are mainly caused by thermal expansion of the fiber and mechanical stresses, which produce birefringent effects and changes in the absolute polarization of the signals. These changes, in turn, cause differential group propagation delays (PDM) that show up as LO phase jitter. A block diagram of the Line Length Correction system implemented by ALMA is shown in Figure A.19.

The two-wavelength laser synthesizer signal (master and slave lasers) is adjusted in polarization and mixed at the SubArray Switch (SAS) and then passed through a 3-port polarizing beam splitter assembly (PBS). The polarization is aligned so that all the light passes through the beamsplitter. It then passes through a piezo-driven fiber stretcher assembly and the fiber to the antenna. At the antenna end there is a 3-dB coupler, so that half of the light goes to the turnaround assembly and half to the photomixer in each WCA. The turnaround assembly consists of a fiber frequency shifter (located at the LO Photonic Receiver module) and a Faraday Rotator mirror located within the WCA of the Band 9 cartridge in each 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 1kHz) copes with the small range variations (about 0.1 mm).

A.7.3 Delay corrections, sideband suppression and interference rejection

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

Delay corrections

ALMA handles delay corrections via the “Delay Server” software package. It computes the corrections for all the different components involved with a cadence of one minute and distributes them buffered. The three main

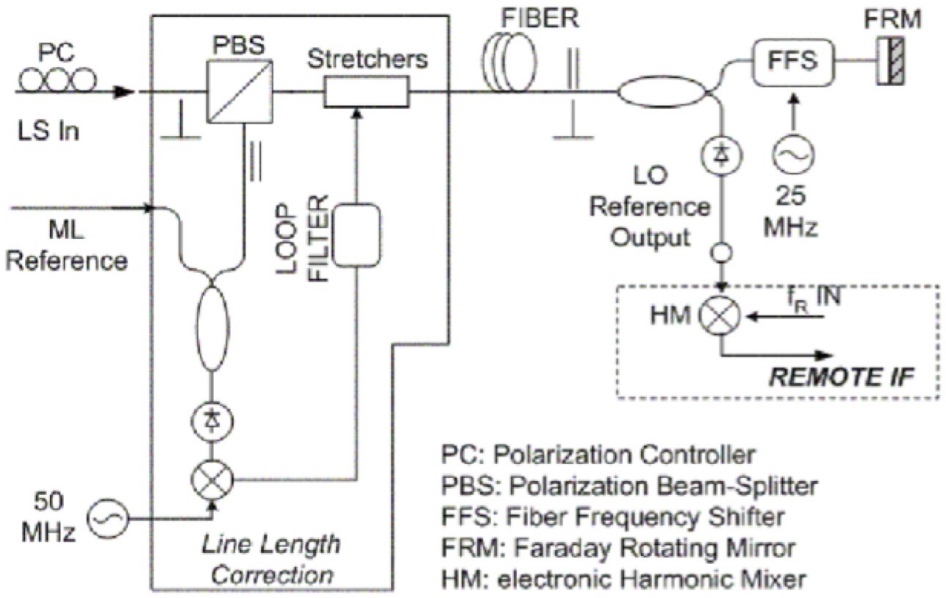

Figure A.19: Block diagram of the Line Length Corrector system for ALMA.

components along the data flow chain where the corrections are applied are: the First LO Offset Generator (FLOOG), the Digital Clock (DGCK) and the correlator (see Figure A.12). Fringe tracking is done at the FLOOG by slightly offsetting the frequency of the LO1 signal. Currently, the delay handled by the FLOOG is in steps of 250 ps. The FLOOG is also used for phase and frequency switching for suppression and separation of sidebands, and for rejection of internally-generated interference, described in the next subsections.

Fine delay corrections are handled by the DGCK that feeds the corrections into the four DTS modules in each antenna. The delay correction resolution of these is 1/16 of the FLOOGs (i.e., 1/16th of 250 ps). The bulk delay correction is handled by the Correlator in integer multiples of the 250 ps units. On top of these corrections, the correlator also handles the "residual" delay corrections at much higher temporal resolution (<250 ps/16) by applying a linear phase gradient across the passband after correlation. Also, the correlator applies relative delay corrections between all the basebands and polarizations of a given ALMA band receiver. Currently, the first baseband of the X polarization is used as reference.

Sideband suppression - LO offsetting

Some of the ALMA receivers (e.g. Band 3, 6 & 7) are inherently single sideband (SSB), either through having a mixer or quasioptic design which rejects the unwanted sideband. But their intrinsic sideband rejection is typically only about 10-15 dB - adequate for rejection of the unwanted sky noise, but not enough to remove strong lines from the other sideband. Others receivers (e.g. Band 9) are double sideband (DSB), and the relative response of the two sidebands may not be equal, significantly affecting calibration. Accordingly, additional schemes are necessary for more effective removal of the unwanted sideband (known as sideband suppression), and for correlation of both sidebands independently (or sideband separation - see next section). Sideband suppression in ALMA is done using the FLOOG, and either LO2 (2LO offsetting) or a combination of LO2 and LO4 (3LO offsetting). A small frequency offset F_o is added to LO1 and subtracted from the other LOs, so that while the signal sideband remains at the same frequency, the image sideband is shifted $2F_o$ away from it's nominal value. A different value of F_o is applied at every antenna, so that all signal sidebands are at the same frequency, but all image sidebands are at slightly different frequencies and no longer correlate.

Note that each of the basebands has an independent LO2 and LO4. So by setting the sign of the offset in (LO2+LO4) differently, each baseband can be set up to observe in a different sideband.

Sideband separation - 90 degree Walsh switching

For future Cycles, it will be possible to apply a 90 deg phase switch in the FLOOG and in the correlator processing, allowing correlation of the upper and lower sidebands separately. For Band 9 (DSB) it will effectively double the bandwidth from 8 GHz to 16 GHz per polarization. This is under development and will not be available for Cycle 2.

Interference rejection - 180 degree Walsh switching

The FLOOG is additionally used to reject spurious signals prior to digitization by applying 180 deg phase switching according to orthogonal Walsh function patterns, with pattern cycle time of 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

A.8 Calibration source selection

A.8.1 Antenna-based Sensitivity Determinations

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

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

ALMA Sensitivities

The present ALMA sensitivity must be known in order to choose the calibrators. These sensitivities are presently determine from an ascii table that contains the nominal ALMA sensitivity as a function of frequency. It was obtained from the ALMA Sensitivity Calculator by running it many times. A partial list of the table is show in Table A.2.

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

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

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

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

Table A.2: Table showing ALMA examples of sensitivities as function of frequency.

Calibration Parameters and Derived Antenna-based RMS

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

- Phase calibration minimum SNR (over entire BW): 15.0
- Maximum phase calibration scan length: 2.0 minutes
- Maximum phase calibration separation from target: 15.0°
- Bandpass Cal/chan minimum SNR: 50.0
- Maximum bandpass calibration scan length: 10.0 minutes
- Maximum bandpass calibration separation from target 50.0°

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

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

A.8.2 Choosing the Calibrators

The image sensitivity for the experimental parameters can be obtained from the above sensitivity table, scaled by the square-root of the appropriate bandwidth and integration time. The conversion of the image sensitivity S_{image} to *antenna-based* sensitivities S_{ant} also depends on the number of antenna, N_{ant}, and is given by Equation A.6

$$S_{ant} = S_{image}(N_{ant} - 3)^{0.5} \tag{A.6}$$

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

Table A.3: The minimum phase calibrator flux density in mJy need for an antenna-based phase determination (see Eq. A.6) with 4 deg rms (15 SNR) with a 2 min integration time, assuming 37 antennas. The total bandwidth in both polarizations of 7.5 GHz is assumed.

The Phase Calibrator

From the experiment frequency, ALMA sensitivities, the number of antennas, the total bandwidth, and the maximum scan length, the minimum flux density for a phase calibrator can be determined. Two minimum flux densities are determined: (1) The entire observing bandwidth and both polarizations can often be combined coherently using the bandpass calibrator scan, and one amplitude and phase solution per antenna per scan can be obtained. (2) An amplitude and phase solution is determined for each spw and polarization per scan. With four spw's and two polarizations, this minimum flux density is $\sqrt{8}$ larger than the full bandwidth solution. The full bandwidth solutions are often made at the highest frequencies in order to find calibrators that are relatively close to the target.

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

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

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

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

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

The Bandpass Calibrator

Although the same strategy in picking the phase calibrator was originally used to determine the bandpass calibrator, a more simpler scheme is now used. Forty-four bright, compact quasars are now monitored periodically and are tied to solar system objects with accurately known absolute flux density models, with ALMA at bands 3 and 7. All are sufficiently bright to be used as bandpass calibrators (with a caveat given below). In addition, many have sufficiently stable flux densities than can be followed to about 5 to 10% between frequencies 90 GHz to 400 GHz, and can be used as secondary flux density standards for an ALMA experiment. Nevertheless, some sources can vary by a factor of three in a month or less (egs. 3c454.3 in 2012), but the monitoring will be sufficiently frequent to remove such variables from consideration as soon as possible.

The present list of bandpass sources is shown in Table A.4. Some changes will be made for extremely variable sources or by the addition of sources that have become strong at bands 7 and 9.

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

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

Other Calibrator Types

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

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

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

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

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

Table A.4: ALMA bandpass calibrator sources and their representative fluxes. Data downloaded on Sept 19, 2013.

Appendix B

Acronym List

ACA	Atacama Compact Array
ACD	Amplitude Calibration Device
ACS	ALMA Common Software
ALMA	Atacama Large Millimeter/Submillimeter Array
AoD	Astronomer on Duty
AOS	Array Operation Site
APDM	ALMA Project Data Model
AQUA	ALMA Quality Assurance software
ARC	ALMA Regional Center
ASC	ALMA Sensitivity Calculator
ASDM	ALMA Science Data Model
AZ	Azimuth
BB	Baseband
BE	Backend
BL	Baseline
BLC	BaseLine Correlator
BWFN	Beam Width between First Nulls
CASA	Common Astronomy Software Applications package
CCA	Cold Cartridge Assemblies
CCC	Correlator Control Computer
CDP	Correlator Data Processor
CFRP	Carbon Fiber Reinforced Plastic
CLO	Central Local Oscillator
CLT	Chilean Local Time
CORBA	Common Object Request Broker Architecture
CRD	CentralReference Distributor
CRG	Central Reference Generator
CSV	Commissioning and Science Verification
CW	Continuous Wave
DC	Direct Current
DEC	Declination
DGCK	Digital Clock
DMG	Data Management Group within DSO
DRX	Data Receiver module
DSB	Double Sideband
DSO	Division of Science Operations
DTS	Data Transmission System
DTX	Data Transmitter module
EL	Elevation
EPO	Education and Public Outreach
ES	Early Science
ESO	European Southern Observatory
FDM	Frequency Division Mode
FE	Frontend
FITS	Flexible Image Transport System
FLOOG	First LO Offset Generator
FOM	Fiber Optic Multiplexer
FOV	Field of View
FPGA	Field-Programmable Gate Array
FT	Fourier Transform
FWHM	Full Width Half Maximum
FWHP	Full Width to Half Power
GPS	Global Positioning System

HA	Hour Angle
HEMT	High Electron Mobility Transistor
HPBW	Half Power Beam Width
IF	Intermediate Frequency
IFP	Intermediate Frequency Processor
IRAM	Institut de Radioastronomie Millimetrique
LFRD	Low Frequency Reference Distributor
LLC	Line Length Corrector
LO	Local Oscillator
LO1	First LO
LO2	Second LO
LO3	Digitizer Clock Third LO
LO4	Tunable Filterbank LO
LORR	LO Reference Receiver
LS	Laser Synthesizer
LSB	Lower Sideband
LTA	Long Term Accumulator
MFS	Master Frequency Standard
ML	Master Laser
MLD	Master Laser Distributor
NGAS	New Generation Archive System
NRAO	National Radio Astronomy Observatory
OMC	Operator Monitoring and Control
OMT	Ortho-mode Transducer
OSF	Operations Support Facility
OST	Observation Support Tool
OT	Observing Tool
OUS	Observing Unit Set
PBS	Polarization Beam Splitter
PDM	Propagation Delay Measure
PI	Principal Investigator
PLL	Phase Lock Loop
PMG	Program Management Group within DSO
PRD	Photonic Reference Distributor
PWV	Precipitable Water Vapor
QA	Quality Assurance
QA0	Quality Assurance Level 0
QA1	Quality Assurance Level 1
QA2	Quality Assurance Level 2
QA3	Quality Assurance Level 3
QL	QuickLook pipeline
RA	Right Ascension
RF	Radio Frequency
RMS	Root Mean Square
SAS	Sub Array Switch
SB	Scheduling Block
SCO	Santiago Central Office
SD	Single Dish
SED	Spectral Energy Distribution
SIS	Superconductor-Insulator-Superconductor Mixer
SL	Slave Laser

SNR	Signal-to-Noise Ratio
SPW	Spectral Window
SRON	Netherlands Instutute for Space Research
SSB	Single Sideband
2SB	Sideband separating Mixer
STE	Standard Test Environment
STI	Site Testing Interferometer
TA	Technical Assessment
TDM	Time Division Mode
TE	Time Event
TelCal	Telescope Calibration subsystem
TFB	Tunable Filterbanks
TFB LO	Local Oscillator at the Tunable Filterbanks
TMCDB	Telescope Monitor and Configuration DataBase
Tsys	System Temperature
$T_{\rm rx}$	Receiver Temperature
USB	Upper Sideband
VLA	Very Large Array
WCA	Warm Cartridge Assembly
WVR	Water Vapor Radiometer
XF	Correlation-Fourier Transform Type Correlator
YIG	Yttrium-Iron Garnet Oscillator

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