

Performance Highlights of the ALMA Correlators

Alain Baudry^{1a,b}, Richard Lacasse^c, Ray Escoffier^c, John Webber^c, Joseph Greenberg^c, Laurence Platt^c, Robert Treacy^c, Alejandro F. Saez^d, Philippe Cais^a, Giovanni Comoretto^e, Benjamin Quertier^a
&
Sachiko K. Okumura^f, Takeshi Kamazaki^d, Yoshihiro Chikada^f, Manabu Watanabe^f, Takeshi Okuda^g, Yasutake Kurono^f, Satoru Iguchi^f

^aUniv. Bordeaux/CNRS, LAB, F-33270, Floirac, France; ^bEuropean ALMA Project Office, ESO, K. Schwarzschild Str. 2, 85748 Garching, Germany; ^cNational Radio Astronomy Observatory, NRAO Technology Center, 1180 Boxwood Estate Road, Charlottesville VA 22903-4608, USA; ^dJoint ALMA Office, Alonso de Cordova 3107, Vitacura – Santiago, Chile; ^eINAF, Osservatorio Astrofisico di Arcetri, Largo Enrico Fermi 5, 50125 Firenze, Italy; ^fNational Astronomical Observatory of Japan (NAOJ), 2-21-1 Osawa, Mitaka, Tokyo 181-8588, Japan; ^gDepartment of Physics and Astrophysics, Nagoya University Furocho, Chikusa, Nagoya 464-8602, Japan

ABSTRACT

Two large correlators have been constructed to combine the signals captured by the ALMA antennas deployed on the Atacama Desert in Chile at an elevation of 5050 meters. The Baseline correlator was fabricated by a NRAO/European team to process up to 64 antennas for 16 GHz bandwidth in two polarizations and another correlator, the Atacama Compact Array (ACA) correlator, was fabricated by a Japanese team to process up to 16 antennas. Both correlators meet the same specifications except for the number of processed antennas. The main architectural differences between these two large machines will be underlined. Selected features of the Baseline and ACA correlators as well as the main technical challenges met by the designers will be briefly discussed. The Baseline correlator is the largest correlator ever built for radio astronomy. Its digital hybrid architecture provides a wide variety of observing modes including the ability to divide each input baseband into 32 frequency-mobile sub-bands for high spectral resolution and to be operated as a conventional ‘lag’ correlator for high time resolution. The various observing modes offered by the ALMA correlators to the science community for ‘Early Science’ are presented, as well as future observing modes. Coherently phasing the array to provide VLBI maps of extremely compact sources is another feature of the ALMA correlators. Finally, the status and availability of these large machines will be presented.

Keywords: ALMA radio interferometry, digital filtering, correlation, digital hybrid XF architecture, FX architecture, ALMA Early Science

1. INTRODUCTION

The Atacama Large Millimeter/submillimeter Array (ALMA) correlators are powerful and versatile digital machines in which the signals captured by all antennas of the array are combined to extract the amplitude and phase information contained in the interferometric fringe patterns associated with antenna pairs distributed within a compact area of about 100 m to an expanded configuration of 16 km. Because there is a large number of ALMA antennas, all being movable within the maximum extent of 16 km, the interferometric data after they have been properly calibrated and Fourier transformed in the spatial frequency domain provide an astronomical image of unprecedented high dynamic range. In addition, the astronomical image may rapidly change with frequency, especially when molecular or atomic spectral line sources are observed. The spectral profile of the observed source is obtained from the Fourier transform in the time

¹baudry@obs.u-bordeaux1.fr; phone +33 557 77 61 62; fax +33 557 77 61 10

domain of the auto- or cross-correlation function which can be relatively easily derived in digital correlators. With the so called XF correlator architecture, where X stands for the correlation part of the design and F is the Fourier transform, correlation as a function of time is derived by implementing digital delays (or lags) and adding many multipliers in the design. In terms of signal processing it is fully equivalent to implement an FX architecture. Both approaches, XF (or, as explained below, a digital hybrid XF or FXX) and FX have been used for ALMA and will be described here.

Two large correlators have been constructed for ALMA and both are installed in a technical building at the Array Operations Site (AOS) at an elevation of 5050 m. One correlator was built by a NRAO/European Integrated Product Team (IPT) to process data from the main array of 12-m antennas (comprising 64 antennas initially) and the other one was built in Japan to process the 16 antennas of the Atacama Compact Array (ACA), which complements the main array for wide angular imaging. These large machines are called the Baseline or 64-antenna correlator, and the ACA correlator, respectively. Both of them process 8 GHz bandwidth in each of two polarizations and provide a wide variety of observing modes. Because the current main array comprises 50 antennas only, the 64-antenna correlator can process up to 14 additional antennas of the ACA if that would be required in some science projects.

In addition to the many technical papers prepared during the construction phase of the ALMA correlators, some general papers have been published already (see e.g. [1] [2] or [3] and [4] for the Baseline and ACA correlators, respectively) but this is the first time that the main characteristics and performance of these two correlators as well as their status and availability to the community are given in the same contribution. Advanced spectral modes and VLBI which will become available in a near future will also be briefly discussed here.

2. 64-ANTENNA CORRELATOR

2.1 Main specifications and architecture

The main specifications of the 64-antenna correlator are given in Table 1. It processes up to 16 GHz bandwidth for 2016 antenna pairs. The 2 GHz ALMA baseband (BB) corresponds to the final analog intermediate frequency range (2 to 4 GHz) prior to 3-bit digitization which, at the time of the design, was a difficult goal to reach.

Table 1. Main specifications of the ALMA 64-antenna correlator.

<u>Item</u>	<u>Specification</u>
Antennas (maximum antenna pairs)	≤ 64 (2016)
Baseband (BB) inputs per antenna	8 x 2 GHz
Input sample format	3-bit, 8-level at 4 Gsample/s
Output correlation sample format	2-bit, 4-level at 125 MHz clock rate
Sensitivity options	4-bit, 16-level; twice Nyquist sampling
Baseline delay range	≤ 30 km (nominal) or more for expanded array
Spectral points per BB (FDM)	≤ 8192 for each BB pair (depending on polarization mode)
Spectral points per BB (TDM)	256, 128 or 64
Polarization products	1, 2 or 4
Temporal integration	1 ms (auto-correlation), 16 ms (auto-, cross-correlation)

Auto- and cross-correlation of the digitized incoming signals cannot be performed at 4 GHz rate, so time demultiplexing is a mandatory step in the signal processing chain. NRAO proposed ‘packetizing’ the data and processing the time-packets in 32 parallel XF correlation ‘planes’ each one working at 125 MHz [5]; the 125 MHz clock rate and the number of planes were carefully selected and thought in the late 1990’s to be a good technology compromise between reliability and complexity. Analysis in different planes is called Time Division Mode (TDM). The initial NRAO pure lag correlator design later evolved into a design implementing the core concept of the European digital hybrid correlator (also called second generation correlator or 2GC) which divides the input baseband into 32 frequency-mobile digital elements (see

e.g. [1] and Section 2.2). Each frequency element (or sub-band) can be assigned in a flexible manner to one (or more) correlator plane(s). Signal analysis is then called Frequency Division Mode (FDM). With a single correlation clock frequency, 125 MHz, both TDM and FDM are possible (see discussion on observing modes); in the FDM case the spectral resolution performance is improved by factors of up to 32 with respect to the pure XF scheme.

The general architecture of the Baseline correlator is presented in Figure 1 taken from [6]. It includes some details of the digitization and transmission process at the antenna as well as long term accumulation of the correlation products and real time computing at the final processing stages. In the following Section we describe selected sub-systems of the antenna-based electronics (Tunable Filter Bank and Station cards) and the baseline-based electronics (Correlator cards). The electronics cards of the complete Baseline correlator are distributed in four quadrants and 32 racks. Each quadrant processes a 2 GHz baseband for all antennas in the array and for up to two polarizations and operates independently of the other quadrants. This architecture permits many distinct modes of operation, including observing different spectral regions with selectable frequency resolution at the same time.

Defining the ability of a single correlation machine to process many independent antenna pairs, N , and the bandwidth BW as the product $N \times BW$, the 64-antenna correlator is clearly the largest machine ever built for radio astronomy. With $N = 2096$ and $BW = 16$ GHz (see Table 1) this product is about 3.35×10^4 GHz while the current mm-wave interferometers are about 100 to 10 times less powerful. Only the future first design correlator stage of the SKA project with 250 antennas and 1 GHz bandwidth will rival ALMA when it is built.

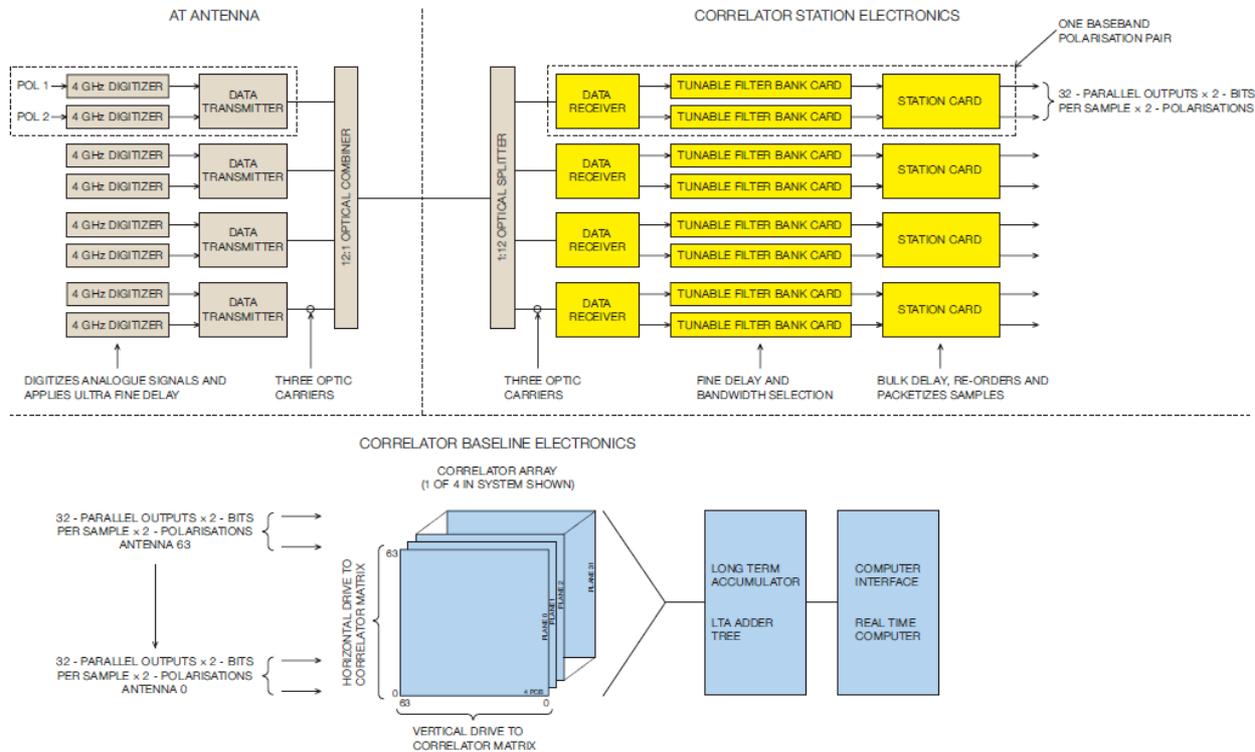


Figure 1. 64-antenna correlator block diagram (adapted from [6]).

2.2 Tunable Filter Bank (TFB), Station cards and Correlator matrix sub-systems

Tunable Filter Bank: Division of the digitized input baseband into several sub-bands is performed in the Tunable Filter Bank (TFB) (see Station Electronics in Figure 1). 32 sub-bands are extracted from the 2 GHz baseband, each with a maximum bandwidth of 62.5 MHz. The sub-bands are implemented on a single card in a matrix of 4 x 4 Altera StratixII FPGAs. The StratixII 90 nm technology was selected from considerations on power dissipation and logic resources

offered by this product at the design time. Digital filtering is performed in a 3-stage digital filter (Figure 2). The first stage is a Comb-type decimating filter which was shown to minimize power dissipation in the filter card (see [7]). It is followed by a quarter-band low-pass filter and a 64-tap half-band filter which is essential to determine the passband ripple, stop-band rejection and the final bandwidth; depending on the digital taps selected in the half-band filter, the 62.5 MHz bandwidth can be halved and the corresponding modes are called ‘half-band’ modes (see Section 2.5). One TFB card dissipates less than 60 W when all 32 filters are active. All 32 sub-bands of the TFB are frequency-mobile across 2 GHz, thus providing great spectral flexibility. This is implemented in the digital local oscillator (LO) and complex mixer which convert each sub-band frequency slice to zero frequency center. In addition, the digital LO offers new possibilities at the system level (for instance, first LO offsetting in the Front-End receiver can be compensated in the TFB LO for Front-End sideband separation) or for phasing up the antenna array (see VLBI mode in Section 2.5).

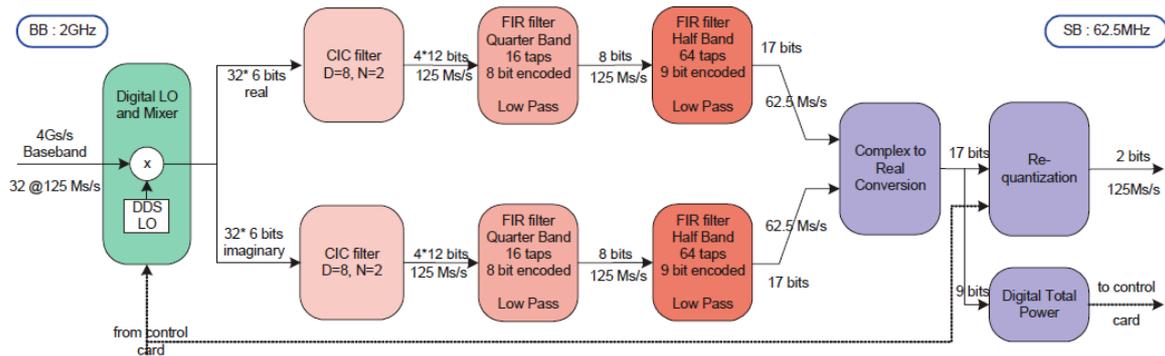


Figure 2. Functional block diagram of the Tunable Filter Bank (TFB) card.

Station cards: As with any interferometer, the geometrical delay due to the finite speed of light across the array must be compensated. In the ALMA system this correction is divided into three parts. The ‘fine’ correction, with a resolution of 16 pico-seconds (ps), is implemented as part of the digital sampler (by shifting the sampler clock phase), at the antenna. The ‘coarse’ correction, with a resolution of one sample or 250 ps, is implemented in the TFB cards. The ‘bulk’ correction is implemented in the card which follows the TFB, the Station card. Its resolution and range depend on the operating mode, but in all cases it far exceeds the 30 km maximum geometrical delay and could compensate up to a few hundred km.

In addition to bulk delay compensation, the Station card implements a mode generation functionality (FDM versus TDM where the data are organized in time-packets). This basically means that it sends the correct samples for a particular mode of operation, delayed by the correct amounts, to the correct inputs of the correlator planes. Delays are implemented in standard semiconductor memories, and allow a much-simplified design for the correlator planes.

Correlator matrix: The outputs of the Station cards drive the planes, whose function is to compute correlation functions for all of the baselines in a 64-antenna array. A custom, application specific integrated circuit (ASIC) designed in 0.25 micrometer technology and called the ALMA-1 chip, is central to the design of the correlator and is described first. This is followed by a description of the architecture of the correlator sub-system. The ALMA-1 chip includes a complete 4-antenna by 4-antenna correlation matrix (16 correlators) with 256 lags per correlator. The lags can be assigned in 1, 2 and 4 groups of 64 to a variety of inputs enabling both non-polarized and cross-polarized modes. Its operating clock frequency is 125 MHz, 32 times slower than the sample rate. Finally, it is capable of integrating the lags for up to 16 milli-seconds (ms) for cross-correlation or accumulating 1 ms samples for auto-correlation. The 4 x 4 architecture of the correlator chip naturally leads to the concept of a correlator plane, a 16 x 16 array of correlator chips able to handle the correlation of all baselines in a 64 x 64 array (see Baseline Electronics in Figure 1). The fact that the clock frequency is 32 times slower than the sample rate drives the design to have 32 planes which, as described previously, support both

TDM and FDM operation. The 16 ms integration capability of the chip results in a requirement to follow the correlator planes with a Long Term Accumulator (LTA) to provide longer integration times and limit the I/O requirements. The LTA can also place received data into bins to enable sideband separation (as required for double sideband receivers) via the classical analog LO Walsh switching scheme.

Communication between the antenna-based and correlator-based electronics cards is critical because of the large number of communication points. The XF architecture was selected mainly to maintain these communication points within a reasonable number for the 64 antennas of the Baseline correlator. The inter-rack wiring requirement is more demanding for an FX design but is well manageable for the 16 antennas of the ACA correlator (see Section 3.2). When the system was designed, the optimum trade-off between number of wires and transmission rate was judged to be 16384 high quality shielded twisted pairs and a 250 MHz bit rate. A special end-to-end communication control system has been designed ; it incorporates pattern generators and checkers and the ability to select the optimum transmission phase for each path.

Finally, it is worth mentioning that the computing power of the 64-antenna correlator is quite impressive. With a total of 4096 multipliers per chip and 32768 correlator chips distributed in the correlator racks, the number of operations performed at 125 MHz clock rate is around $1.7 \cdot 10^{16}$ per second. The number of multi-bit operations performed per second in a single TFB card is around 10^{12} . These numbers make the 64-antenna correlator the fastest computer ever used at an astronomical high site.

2.3 Power dissipation, Reliability, Single Event Upsets

Power dissipation: The ALMA correlators are located in a technical building at an elevation of 5050 m (AOS site). One of the major design considerations of the 64-antenna correlator is power dissipation because: (a) the air density at 5050-m is nearly half that at sea level; (b) dissipation is related to the component, card or sub-system failure rate; (c) energy distribution and running costs increase quickly for high elevation sites. These questions were addressed for the 64-antenna correlator at the individual component and card levels and at the entire system level in order to minimize the power dissipation and its effects as much as possible with technologies available at the time of the design.

The most critical items in terms of dissipation are the correlator and filter chips because of their large quantity; there are 32768 correlator ASICs and 8192 large FPGAs in the final design. At the time of the ASIC design and given the high dissipation expected with a large number of lags, the most appropriate technology choice seemed to be the ASIC fabrication option instead of the FPGAs available in the early 2000's. The 4096-lag ASIC designed for ALMA dissipates only about 1.6 W, thus enabling a design with a relatively large number of chips per correlator board. At the time when the digital filtering architecture was frozen the FPGA technology had made good progress and we selected the latest 90 nanometer process from Altera. Our filter firmware design has evolved toward less dissipation by making optimum use of the FPGA resources (see [7]) to achieve around 60 °C junction temperature in one FPGA and less than 60 W per TFB card when all 16 large FPGAs are at work.

When all filter and correlator cards are in use, the Baseline correlator dissipates up to 150 kW or slightly less, including power supply and computing racks; despite this good achievement for such a large machine, an efficient ventilation and air conditioning system is required. The air circulates from the correlator floor through each of the 32 racks and, in addition, a regulated fan system is installed on top of each 'Station electronics' rack where power density is the largest. Inlet temperature at the base of the racks is around 16 °C and the 5 to 10°C temperature rise through the racks results in a maximum die temperature of the TFB FPGAs of about 62 °C, well within the recommended limit.

Reliability: Power dissipation is a critical parameter for reliability, but there are several other considerations to take into account for global reliability. All cards have been fabricated with high quality industrial and production rules and special test fixtures have been developed to validate most types of cards prior to their integration in a functional rack. Many spare cards or critical components have also been fabricated or stored for long term operation. Finally, to simplify trouble-shooting, we have designed specific modules which can generate test patterns in the data path and query test functions.

We have now accumulated some experience on the reliability of the components in the 64-antenna correlator. Few correlator ASIC's failed at the integration center prior to moving the correlator to the high site. After an initial period of about 40 days showing a higher failure rate there seemed to be no further failures. At the 5050-m AOS site only random

and very few ASIC failures have been observed. Of the 8192 large FPGAs assembled on the TFB cards no failures were observed during a long ‘burn-in’ period at the integration center. At the AOS, only 4 FPGAs had to be replaced after several years of operation. However, the indirect effects of Cosmic Rays on these FPGAs is still poorly known and is further discussed below.

Single Event Upsets: The energetic neutrons generated by the Cosmic Rays interacting with the terrestrial atmosphere are a potential threat to the many very large scale integrated circuits used in the ALMA correlators. The cascade of particles triggered by the energetic neutrons can cause ‘bit-flips’ in the integrated circuits; these are generally called Single Event Upsets (SEUs). In some instances the SEUs can cause incorrect operation. The StratixII filter FPGAs on the TFBs can detect SEUs using cycling redundancy checks and report them to the high level computer. Such events are random by nature and require long monitoring periods.

Our SEU measurements made with 2048 FPGAs in one quadrant of the correlator at the AOS give an SEU rate of about 2.6 events/day corresponding to about 10.4 events for the full four quadrants. If there are too many SEUs, they could affect the FPGA functionality and hence the astronomical data quality. The safest and simplest approach seems then to reload the personalities and registers of all FPGAs whenever this is practical during the observing schedule. Although there is no protection possible against large SEU rates, warnings for TFB high currents have been implemented to inform the system about such extreme cases. Finally, it is interesting to mention that: (a) from a comparison of our monitoring at the integration center (130 m elevation) and at 5050 m, the measured SEU rate varies roughly in proportion with the elevation; (b) the SEU rate measured from all TFB card FPGAs provides a good indication on the fast neutrons flux at 5050 m and the potential threat for the entire ALMA equipment.

2.4 Observing modes and ALMA Early Science

The two main observing mode categories, TDM and FDM, reflect the digital hybrid architecture of the Baseline correlator. In the TDM mode, one millisecond packet of 4 GHz samples is split into 32 ‘time bins’, sent to 32 correlator planes each processing $1/32^{\text{th}}$ of the digitizer samples which ultimately are summed together to retrieve the initial data rate and the exact cross-correlation coefficients. The correlator resources used to process the individual time bins provide higher time resolution but less spectral resolution than the frequency division mode, where the input baseband is broken into several sub-bands sent to the correlator planes. In the widest bandwidth FDM mode, the Tunable Filter Bank card provides 32 sub-bands and each sub-band (62.5 MHz wide) is processed in one of the 32 correlator planes to provide across 2 GHz thirty-two times higher spectral resolution than in the TDM case. Spectral resolution is increased by processing fewer than 32 sub-bands with all of the available correlator resources, the extreme case being when one sub-band is processed in all 32 planes to reach a few kHz resolution.

A total of 21 observing modes are offered to the science community for ALMA Early Science (Table 2). There is a good variety of channel spacings offered in FDM. For a given input bandwidth, the spacing, and hence the spectral resolution, depends on the number of selected cross-correlation products (1, 2 and 4). The maximum effective bandwidth available per baseband is limited to 1.875 GHz because of the anti-aliasing bandpass filter used prior to digitization and, for smaller input bandwidths (baseband divided by 2^n), the bandwidth limitation is set by the slight sub-band channel overlap needed at the spectrum reconstruction stage. The channel spacing in Table 2 is shown for a single baseband, but as there are 4 baseband pairs available, ‘mixed resolution’ observing modes are supported for Early Science with different resolution set-ups in different basebands. Table 2 shows an apparent ‘gap’ in channel spacing from FDM to TDM (less than 1 MHz to around 8–30 kHz) and in some FDM observations there might be cases where there are more channels than actually needed. However, channel averaging options are supported by the ALMA software to control the number of output channels.

The actual spectral resolution is slightly different from the channel spacing given in Table 2. This is because the correlation function derived in the 64-antenna correlator is truncated to a maximum lag value which, after Fourier transformation, results in a sinc spectral function. The final spectral resolution depends on the weighting function adopted at the post-processing stage (see e.g. [8]) to control the sinc function ringing effect; it is, for instance, about 1.21 times the channel spacing for uniform weighting.

Table 2. Supported Frequency and Time Division modes (FDM and TDM) for ALMA Early Science. In FDM, 1, 2 or 4 cross-products are selectable at a given bandwidth. The 21 ‘Early Science’ modes shown in this Table form a sub-set of the 67 modes offered by the 64-antenna correlator.

Effective Bandwidth (MHz)	Channel Separation (kHz)								Channel Separation (MHz)		
	FDM								TDM		
	7.6	15.3	30.5	61	122	244	488	977	7.8	15.6	31.3
1875						1	2	4	1	2	4
938					1	2	4				
469				1	2	4					
234			1	2	4						
117		1	2	4							
58.6 (62.5)	1	2	4								

The channel separation shown in Table 2 is suitable for a broad variety of science goals. FDM resolutions in the range 10 to 100 kHz are adequate for line spectroscopy in protostellar discs, cold interstellar clouds or molecular cloud maser studies, while 0.5 to 1 MHz resolutions are well suited for Galactic line outflow studies. TDM modes offer for a broad bandwidth a coarser spectral resolution as is required to investigate broad band continuum sources or broad spectral line profiles in distant galaxies. For the very distant extragalactic objects it might be required to concatenate 2 or 4 basebands in order to analyze very broad velocity spreads.

2.5 Future modes and VLBI mode

The 18 frequency division modes shown in Table 2 are just a sub-set of the 63 different frequency division modes which have been validated by the Baseline correlator IPT team. All 63 modes are not supported yet for science observations, but they will become available in the future. The new modes (in addition to Table 2) correspond to double Nyquist sampling and 4-bit correlation options which all increase the correlation efficiency at the expense of reduced bandwidth or less spectral points. In addition, we have implemented ‘half-band’ modes which, thanks to appropriate numerical taps loaded in the last stage of the digital filters, increase the spectral resolution by an additional factor of two and increase the sensitivity (‘half-band’ modes employ twice Nyquist sampling).

The digital hybrid architecture allows us to use the TFB card and correlator plane resources in many different ways to produce a variety of new multi-spectral and multi-window modes. These modes have been described in [1] and [9] but they are not yet fully supported by the ALMA software. Multi-window modes are possible because the TFB frequency sub-band agility allows us to break a baseband into several ‘windows’ so that various spectral lines can be analyzed simultaneously. Multi-spectral resolution within the same baseband will be another option to allow zooming on complex spectral lines. Multi-window and multi-resolution options combined with the 67 basic FDM and TDM modes greatly enhance the overall Baseline correlator flexibility.

Phasing-up the ALMA array will offer the new possibility to combine the huge ALMA collecting area with other mm/submm-wave facilities across different continents to perform very long baseline interferometry (VLBI) observations at the sub-milliarcsecond level. The original 64-antenna correlator design provides ‘hooks’ for VLBI. Hooks include adjustable phase in the TFB LOs, FPGAs dedicated to computing the sum of a mask-selectable set of antennas, and unused rack space and power capacity. An enhanced design will use the adjustable phase of the DLLs to make a fine adjustment to the delay in each of the thirty-two 62.5 MHz bands for each antenna (replicated for each baseband). The antennas will then be summed by the dedicated FPGA’s previously mentioned and the sum transmitted to new boards called Phasing Interface Cards (PICs). The primary function of the PICs will be to format the data into packets for transmission via 10 Giga-bit Ethernet. The enhanced hardware design is only one component of the entire ALMA Phasing system. Additional components include an improved frequency standard (H-maser), a fiber transmission system to get data from the AOS to the lower elevation Operations Support Facility site (OSF) economically, VLBI recorders at the OSF and a great deal of software to control the system and close the important phasing loop.

3. ACA CORRELATOR

3.1 Main specifications, requirements and architecture

ALMA consists of the main 12m-Array composed of fifty or more 12-m antennas, and the Atacama Compact Array (ACA) composed of twelve 7-m antennas (7m-Array) and four 12-m antennas, the Total Power Array (TP-Array). The ACA thus collects both the short spatial components data and the zero-baseline length data from interferometric and single-dish observations made with the 7 m-Array and the TP-Array, respectively. The spatial components obtained with the three arrays are combined to construct high-fidelity images of the target objects. The main purpose of the ACA is to provide short-baseline and total-power data of various astronomical objects with high precision, which is vital for the science goals. These data are later combined with the longer-baseline data obtained with the 12m-Array for high fidelity imaging of astronomical objects (see e.g. [10]). Since the ACA is sometimes used as a single array for observations, it is also required to have reasonable imaging capabilities. Thus, the ACA correlator has: (1) the capability to process receiver signals from up to 16 ACA antennas, (2) a high processing performance, sufficient to calculate a large quantity of correlation spectra among all of the ACA antennas, (3) high precision processing, sufficient to achieve a high spectral dynamic range, and (4) a low sensitivity loss. Concerning (3) and high spectral dynamic range, the specifications are 10000:1 for the measurement of weak spectral lines near stronger ones and 1000:1 for weak lines in the presence of strong continuum emission. Low-sensitivity loss ($< 12\%$) is also required for the ALMA digital system, including both the digitizer and the correlator. Since ALMA has adopted a 3-bit digitizer whose sensitivity loss cannot be less than 3.7%, the allowable sensitivity loss of the ACA correlator must be less than 8.3% to achieve the 12% loss goal. Additional important requirements for the ACA correlator are mentioned in [4], and we simply emphasize a few points below.

Sub-array mode and full-array mode: The TP-Array shall be independently operated as a single-dish telescope to obtain total power data, while the 7m-Array shall be operated as an interferometer to obtain short-baseline data. These two arrays shall also operate together as a single full-array to provide more sensitivity for high accuracy calibration. Hence, the ACA correlator is required to have two operational modes: the sub-array mode to perform data processing of the TP-Array and 7m-Array simultaneously, and the full-array mode to operate two arrays as a single array.

Spectral configuration: In spectroscopic observations with ALMA, it is necessary to select an appropriate IF band frequency range with an appropriate frequency resolution, because the required frequency range and resolution vary depending on the astronomical objects to be observed. Hence, ALMA requires flexible spectral configurations. Up to 32 frequency ranges with different frequency resolutions are offered to the observers. These multiple frequency ranges are called sub-bands and the way they are extracted from a 2 GHz baseband in the Baseline correlator has been described in Section 2.2 while they are more directly generated in the FX architecture of the ACA correlator (see below and following Section). The narrowest sub-band is 31.25 MHz wide to meet the highest frequency resolution ALMA specification which is 5 kHz or less. In the Baseline correlator, this is achieved with the ‘half-band’ modes (see Sections 2.2 and 2.5).

Compatibility with the 64-antenna correlator: Functional and data compatibility of the ACA correlator with the 64-antenna correlator is very important because the Front-End, Back-End and LO subsystems are all common to the ACA and the 12m-Array, and their output data are combined as previously mentioned. Therefore, the ACA correlator must be fully compatible with the 64-antenna correlator in terms of Front-End sideband separation, spectral configuration and temporal integration. The ACA correlator top level specifications are summarized in Table 3 and can be compared with specifications in Table 1.

Table 3. Main specifications of the ACA correlator.

<u>Item</u>	<u>Specification</u>
Antennas (maximum antenna pairs)	16 (120)
Baseband (BB) inputs per antenna	8 x 2 GHz
Input sample format	3-bit, 8-level at 4 Gsample/s
Output correlation sample format	4-bit, 16-level
Maximum number of correlation	120 cross-correlations, 16 auto-correlations
Baseline delay range	≤ 15 km
Sideband processing	90-degree phase switching or LO offset
Maximum number of spectral points	8192
Polarization products	1, 2 or 4
Temporal integration	1 ms (auto-correlation), 16 ms (auto-, cross-correlation)

For the ACA correlator, we selected a conventional FX method, normal FFT in the fixed-point arithmetic. It has been analyzed and simulated, and its characteristics are well known. It is adequate to achieve a sensitivity loss of 8.3% or less and a high spectral dynamic range as required for ALMA (see above). In addition, the FX correlator offers the advantage of providing a highly flexible spectral configuration. Contrary to other designs, the FX correlator always performs spectroscopy of the input signal by using the Fourier transform with the highest frequency resolution; it is thus easy to output the requested multiple frequency ranges with the necessary frequency resolutions by choosing the highest frequency resolution spectra and summing them up. Providing such flexibility is important for general-purpose telescopes such as ALMA. Another advantage of the FX processing scheme is the frequency response of the sinc squared function resulting from the calculation sequence ‘F-part’ followed by the ‘X-part’. The frequency response has a relatively steep edge and low sidelobe levels which is advantageous for precise imaging, one of the ALMA main goals.

3.2 Design

The processing flow was designed to achieve correlator functionality according to the ALMA specifications and requirements described in the previous section. The processing flow of the ACA correlator is shown in Figure 3. (This is Figure 4 in [4].) Only the designs of the major functions are described here: *FFT*, *Re-Quantization*, *LO offset*, and *Spectral and temporal integration*.

FFT: The FFT is required to perform 1048576 ($= 2^{20}$; hereafter called 1 Mega.) or more points in order to achieve the highest frequency resolution of 5 kHz in a 2 GHz bandwidth. The 1 Mega spectral points FFT provides the highest frequency resolution of 2 GHz/512K frequency channels (where $K=2^{10}$). This corresponds to 3.815 kHz resolution, which is finer than the 5 kHz required by ALMA and matches the highest frequency resolution of the 64-antenna correlator (see ‘half-band’ modes). The bit length of the FFT calculations is driven by the required spectral dynamic range and sensitivity loss described in Section 3.1. The spectral dynamic range of 10000:1 requires that the calculation bias within the ACA correlator is less than -60 dB relative to T_{sys} , assuming strong astronomical signals of $1/100 T_{sys}$ at the highest frequency resolution. This suggests that 10-bit or more accuracy is necessary for the FFT calculations, while a longer bit length is preferable to decrease the calculation noise caused by number rounding. For these reasons, a 16-bit fixed-point FFT calculation is adopted. In the 1 Mega-point FFT operation, 4-Gsps time series data of finite length is used as input data and each data segment is called an FFT segment. Since 4-Gsps data is used to obtain correlation data of 1 or 16 ms integration corresponding to 4×10^6 or 64×10^6 samples, 1 Mega-point FFT calculation is initiated every 250 microsec or every 10^6 samples. This indicates that 1 Mega-point FFT lacks 48576 ($= 2^{20} - 10^6$) samples. To account for this difference, adjacent FFT segments are overlapped by about 5% of a segment length.

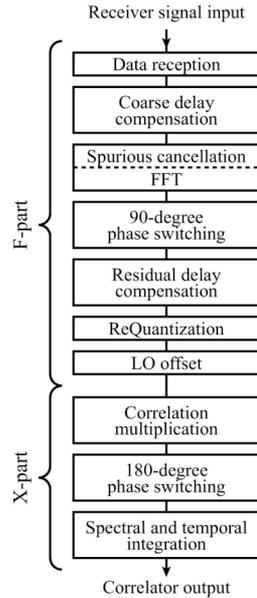


Figure 3. Processing flow of the ACA Correlator (taken from [4]).

Re-Quantization: The 1Mega-point FFT outputs 512K 16-bit complex values (voltage spectra) for each baseband of all 16 antennas. All of these values have to be transferred to the X-part of the design in order to calculate all correlation spectra over 2-GHz bandwidth in the correlation multiplication block. The total data rate for one baseband pair is then $16\text{-bit} \times 2$ (real and imaginary parts of the complex values) $\times 512\text{K}$ frequency channels $\times 2$ polarizations $\times 16$ antennas/ $250\mu\text{s}$ (effective time length of 1 FFT segment) = 2150 Tbps. In order to reduce the data rate from the F-part to the X-part and thus the hardware size of the X-part, the complex values are re-quantized to 4-bit real and 4-bit imaginary numbers at the output of the F-part. The word length of the re-quantization is determined by a trade-off between possible data deterioration and the impact on the hardware size. The ALMA signals are already degraded by the 3-bit digitizer (3.7% loss) and a 3-, 4-, or 5-bit re-quantizer adds another 3.7%, 1.1% or 0.34% loss, respectively; this additional loss was estimated by re-quantization simulations using a Monte Carlo method. If there were no other noise sources in the correlator, selection of 3-bit, 4-bit, or 5-bit re-quantization would correspond to a total sensitivity loss of 7.3%, 4.8%, and 4.0%, corresponding to an increase of the observing time of 16%, 10%, and 9%, respectively, with respect to the case without the re-quantization. The hardware cost of a correlation multiplier is approximately proportional to the square of the word length at its input and we expect a relatively large cost increase of the ACA correlator from the adoption of 5-bit instead of 4-bit re-quantization. Since the 3-bit choice would reduce the sensitivity too much and a 5-bit solution would not be affordable, we decided to choose the 4-bit solution.

LO Offset: Another way to suppress spurious signals and the unwanted receiver Front-End sideband is to adopt the LO offset method in addition to the classical 180-degree and 90-degree phase switching. An LO offset adds a slightly shifted frequency offset to the LO frequency at the antennas (modulation) and subtracts this offset before correlation (demodulation). The spurious signals generated between the first LO where the offset is applied and the point where it is removed (with different frequency offsets for different antennas) are thus suppressed after correlation. In addition, because the offset removal is correctly performed for either the upper or the lower sideband of the Front-End receivers, the LO offset method efficiently removes the undesired sideband of the ALMA two-sideband receivers. On the other hand, for the ALMA double sideband receivers, the first LO 90° phase switching method needs to be applied to separate the upper and lower sidebands. The ACA correlator realizes the LO offset demodulation by shifting the spectra by the appropriate number of frequency channels at the output of the 1 Mega-point FFT, before the correlation multiplication stage (see Figure 4). In the 64-antenna correlator, the LO offset is demodulated in the digital LO of the TFB card (see first block in Figure 2). The minimum step of the frequency channel shift in the ACA correlator must match the

frequency resolution of the TFB digital LO, which is 30.5 kHz; this is exactly 8 frequency channels of the ACA correlator ($30.5 \text{ kHz} = 3.815 \text{ kHz} \times 8 \text{ channels}$). The maximum frequency shift in the ACA correlator is about 39.1 MHz.

Spectral and Temporal Integration: After correlation, the sub-band spectra of the requested frequency ranges are integrated up to the requested frequency resolutions. The maximum number of sub-bands and bandwidth are 32 and 2 GHz, respectively. Spectral integration is available for 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, and 1024 frequency channels \times 3.815 kHz. If integrations larger than 1024 frequency channels \times 3.815 kHz are needed, further integration is performed by post-processing computers. Two types of spectral integration methods are available. One is non-weighted spectral integration, which just adds up spectral data without weighting in the frequency domain. The other one is frequency profile synthesis, which performs convolution of the weighting function in the frequency domain. The latter method is used for frequency profile compatibility with the 64-antenna correlator. Some details are given in Section 4.3. Sixteen and one milli-second temporal integrations are performed in the ACA correlator for the interferometric observations (auto- and cross-correlation spectra for all 16 antennas and all baselines) and the single-dish observations (auto-correlation and cross-polarization spectra for all 16 antennas), respectively. The spectral and temporal integrations are fully compatible with the 64-antenna auto- and cross-correlation products.

3.3 Hardware implementation

For the hardware design, we adopted the following policy: (1) distributed processing; (2) use of common modules and card structures; (3) use of optical transmission in long-distance data transmission; (4) efficient cooling; and (5) distributed power supplies. Distributed processing allows us to select processors with less demanding performance and enables us to use simple, but general-purpose, processing chips, such as low frequency FPGAs. This also enables having distributed heat sources, which is favorable for thermal dissipation. The use of common modules and card structures reduces the number of case/card types and enables simple assembling and, accordingly, simple maintenance operations. Optical transmission is more reliable than electrical cables in the bulk/long distance data transmission within the ACA correlator (e.g., 500 Gbps data in total are transferred from the F-part to the X-part) and between the ACA correlator and post-processing computers. The ACA correlator is installed in the Array Operations Site (AOS) technical building at an altitude of about 5050 m in a room next to the 64-antenna correlator room. The cooling system was designed to maintain high efficiency in an environment where the air density is half that at sea level. Stable electric power is supplied to each electrical component by ‘DC power generation and voltage conversion’ placed close to the components. The ACA correlator consists of four quadrants, each one processing one baseband pair. One quadrant is mainly divided into three parts (see Figure 4 without M&C module): Data Transmission System Receiver and FFT Processor (DFP) modules; Correlation and Integration Processor (CIP) modules; and Monitor & Control Interface (MCI) module. The DFP module, corresponding to the F-part of the FX correlator, receives optical signals transmitted from each antenna, restores observation data, and performs FFT processing of the restored data. The CIP module, corresponding to the X-part of the FX correlator, performs the correlation multiplication of the voltage spectra (FFT outputs), and outputs auto- and cross-correlation spectra for various temporal and spectral resolutions. The MCI module is responsible for monitoring and control of the DFP and CIP modules in a quadrant; it is also used as the interface with a control computer.

These three types of modules adopt a common module/card structure, which consists of a motherboard, processing cards, such as FFT and correlation multiplication, a power-supply unit and a cooling fan unit. The processing cards are aligned in parallel within a module so that the air entering from the front side of a module goes straight through a module, and is exhausted from the back side. Each module is equipped with a switching power supply unit for AC–DC conversion from AC230 V to DC12 V. The generated DC12 V is converted to an appropriate voltage (e.g., DC1.2 V for FPGAs) by DC–DC converters, which are distributed to a motherboard and cards in order to supply the necessary electric power to each electric component on a short distance. Optical transmission is used to realize reliable transmission of bulk data (3.125 Gbps per optical fiber) between the modules over a long distance (up to about 3 m). For communications between the MCI module and the other modules, 100 Base-T Ethernet is used through a network switch. The 48-ms timing signal and 125 MHz reference signal sent from the Back-End sub-system are received at the Reference Signal Distributor (RSD) panel, and are distributed to each quadrant through the panel. As a result, a quadrant of the ACA Correlator is composed of eight DFP modules, four CIP modules, one MCI module, one RSD panel (installed only in the first quadrant), one Ethernet controlled outlet and one network switch. They are installed in two racks.

Further details of the hardware implementation are given in [4] and we just give below additional information on one essential element of the ACA correlator, the FFT card. The FFT card is equipped with two different FPGA types, one for

FFT and another one for control. The FFT FPGA (Xilinx Virtex-4 LX60) performs 1 Mega-point FFT and related processing including re-quantization and LO offset demodulation (see Section 3.2). The control FPGAs (Xilinx Virtex-4 LX80) allocate the segmented data transferred from the Data Transmission System Receiver (DTS-Rx) cards to the FFT FPGAs, and transmit the FFT outputs collected from the FFT FPGAs to the CIP modules. One FFT card includes 8 FFT FPGAs. An FPGA pair processes 1 Mega-point FFT operation of 16 FFT segments corresponding to a unit of data transfer from the DTS-Rx card every 32 ms. Hence, one FFT card processes 64 FFT segments and, in total, 512 FFT operations are performed by 8 FFT cards of a DFP module every 32 ms (see Figure 4). After the FFT operations, the voltage spectra are compensated for residual delay and re-quantized to 4 bits. Finally, the 4-bit re-quantized voltage spectra are shifted in the frequency channel domain for LO offset demodulation (if required) and transferred to the CIP modules.

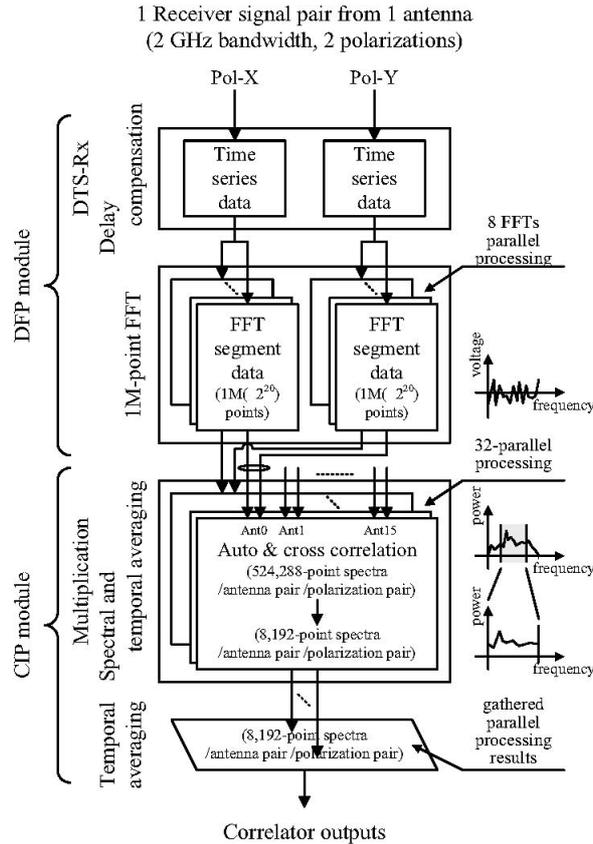


Figure 4. Detailed Processing of the ACA correlator (taken from [4]).

4. STATUS & FORWARD LOOK TO THE FUTURE

4.1 Baseline correlator status

The 64-antenna correlator is organized by quadrants. The first and second quadrants were delivered to the AOS in 2008 and 2009 respectively and they are routinely used for ALMA Early Science and commissioning and science verification tasks; these two quadrants are able to process up to 32 antennas for all four 2 GHz baseband pairs. The third quadrant was delivered to the AOS in 2010 and is being used primarily as a remote software test bed and for engineering testing, firmware development or SEU rate testing. The fourth quadrant has been extensively tested at the integration center and

is now at the AOS. Migration to the full 4-quadrant operation will be completed in September 2012 in order to double the processing capacity (up to 64 antennas). In this configuration all modes described in this paper will become available to the science community.

4.2 ACA correlator status

The design and actual implementation of the ACA correlator were verified by conducting functionality, stability and performance tests. A known digital input signal was sent to the ACA correlator with typical configuration parameters and its outputs were checked to confirm that they were, consistent with 64-bit floating-point arithmetic results for design verification, and/or bit-accurately coincident with the ACA correlator simulator results for implementation verification. In the performance test, a 1-bit digitized thermal noise signal was also used to evaluate the upper limits of the calculation bias and sensitivity loss. Functionality and stability tests were executed both in Japan and at AOS to confirm proper implementation and functioning of each quadrant. The performance test was performed in Japan only. The four quadrants of the ACA correlator were delivered to the AOS in May 2012 and they are being used for ALMA commissioning and science verification tasks. Because the ACA will be available at the next round of Early Science observations the supported ACA correlator spectral configurations will be identical with those of the 64-antenna correlator (see Table 2).

4.3 Merging both correlator outputs

The ACA correlator has adopted an FX architecture while the 64-antenna digital hybrid FFX correlator is basically an XF design. This difference introduces different frequency responses between these two correlators. However, to combine the ACA and 12m-Array data images, frequency channel profile compatibility is required. Therefore, we have implemented in the ACA correlator a ‘frequency profile synthesis’ function based on the convolution function in the spectral domain to achieve frequency profile compatibility. We found that the frequency response of an XF correlator convolved with a Hanning window function is well synthesized by the linear combination of twelve FX frequency responses when one frequency channel width of the 64-antenna correlator is twice that of the ACA correlator. More generally, it was found that the adequately weighted sum of the ACA correlator channels can match the 64-antenna correlator spectral profiles. The maximum profile differences between the two correlators are estimated to be as small as 0.6 % or less except when a δ function is used as an ‘object profile’. Details of the calculation are described in [11].

4.4 Forward look to the future

Some key components in the ALMA correlators are now superseded by more advanced products in terms of technology process or fast serial links. For instance, the FPGAs assembled on the TFB cards use a 90 nm process available when we started our design while the most recent products available nowadays use a 28 nm technology which offers huge memory capacity and many more and faster interface links. The availability of new FPGA generations should clearly enhance the performance of future correlators, although the SEU susceptibility of smaller-feature devices at high elevation sites needs to be weighed carefully. The situation is entirely similar with the ACA correlator. Despite these continuing technology advances, we consider that the 64-antenna and the ACA correlators are large machines of the future for many years to come because of their remarkable performance and their extreme configuration flexibility offering many observing modes, and simply because there is still a large fraction of validated modes which have not yet been commissioned by the ALMA observatory beyond the Early Science phase or because very high data rate extraction from the correlator hardware is not yet possible and will offer new perspectives in the future.

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

Acknowledgements: Many people in addition to the authors should be thanked for their contribution during the construction, assembly and testing phases of the ALMA correlators. The 64-antenna Correlator IPT thanks continuous support from the ALMA Computing

IPT correlator sub-group. R. Amestica and J. Perez played a key role in integrated testing of the Baseline correlator at both the integration site in Charlottesville and the 5050 m site in Chile. The ACA correlator team thanks continuous support from the ALMA Computing IPT ACA correlator sub-group. T. Matsui plays an important role in operational and integrated testing of the ACA correlator at both Mitaka in Japan and the AOS/OSF sites in Chile. The Commissioning Science and Verification team also contributed in offering to the science community new observing modes initially designed within the correlator teams.

REFERENCES

- [1] Escoffier, R., Comoretto, G., Webber, J., Baudry, A., Broadwell, C.M., Greenberg, J., Treacy, R., Cais, P., Quertier, B., Camino, P., Bos, A., and Gunst, A., “The ALMA correlator”, *Astron.&Astrophys.* 462, 801-810 (2007)
- [2] Baudry, A., Webber, J., “The ALMA 64-antenna correlator: main technical features and science modes”, URSI XXX General Assembly and Scientific Symposium (Istanbul), Session J10 on Mm and Sub-mm Science and Technology (August 2011)
- [3] Baudry, A., “The ALMA correlators”, ALMA Newsletter No. 7, 18-31 (January 2011)
<http://www.almaobservatory.org/en/outreach/newsletter/252-newsletter-no-7>
- [4] Kamazaki, T., Okumura, S., Chikada, Y., Okuda, T., Kurono, Y., Iguchi, S., Mitsuishi, S., Murakami, Y., Nishimuta, N., Mita, H., and Sano, R., “Digital spectro-correlator system for the Atacama Compact Array of the Atacama Large Millimeter/submillimeter Array”, *PASJ*, 64, 29 (2012)
<http://pasj.asj.or.jp/v64/n2/640029/640029.pdf>
- [5] Escoffier, R., “The MMA correlator”, ALMA Memo Series No. 166 (1997)
<http://www.alma.cl/almamemos/100358/memo166.pdf>
- [6] Escoffier, R., Webber, J. and Baudry, A., “64 Antenna Correlator Specifications and Requirements”, ALMA-60.00.00.00-001-C-SPE document (2008)
<http://edm.alma.cl/forums/alma/dispatch.cgi/documents/showFile/100591/d20090218174917/No/2008-08-07-ALMA-60.00.00.00-001-C-SPE.pdf>
- [7] Camino, P., Quertier, B., Baudry, A., Comoretto, G. and Dallet, D., “The new 3-stage, low dissipation digital filter of the ALMA correlator”, ALMA Memo Series No. 579 (2008)
<http://www.alma.cl/almamemos/100612/ALMA%20Memo%20579.pdf>
- [8] Comoretto, G., “Algorithms and formulae for hybrid correlator correction”, ALMA Memo Series No. 583 (2008)
<http://www.alma.cl/almamemos/100615/ALMA%20Memo%20583.pdf>
- [9] Escoffier, R., Comoretto, G., Broadwell, C., Lacasse, R., Webber, J., Baudry, A., “Observational modes supported by the ALMA correlator”, ALMA Memo Series No. 556 (2006)
<http://www.alma.cl/almamemos/100591/ALMA%20Memo%20556.pdf>
- [10] Iguchi, S., Morita, K-I., Sugimoto, M., VILA VILAR´O, B., Saito, M., Hasegawa, T., Kawabe, R., Tatematsu, K., Sakamoto, S., Kiuchi, H., Okumura, S., Kosugi, G., Inatani, J., Takakuwa, S., Iono, D., Kamazaki, T., Ogasawara, R., and Ishiguro, M., “The Atacama Compact Array (ACA)”, *PASJ*, 61, 1 (2009)
<http://pasj.asj.or.jp/v61/n1/610101/610101.pdf>
- [11] Kamazaki, T., Okumura, S. K., Chikada, Y., “Frequency profile difference between ACA Correlator and 64 antenna Correlator”, ALMA Memo Series No. 580 (2008)
<http://edm.alma.cl/forums/alma/dispatch.cgi/almamemos/showFile/100611/d20081029184302/No/ALMA%20Memo%20580.pdf>