PROCEEDINGS OF SPIE

The CARMA correlator

David W. Hawkins, David P. Woody, B. Wiitala, J. Fredsti, K. P. Rauch

SPIE.

The CARMA Correlator

D. W. Hawkins[a], D. P. Woody[a], B. Wiitala[a], J. Fredsti[a], and K. P. Rauch[b]

[a]Caltech-OVRO, 100 Leighton Lane, Big Pine, CA 93513, USA.
[b]Laboratory for Millimeter-wave Astronomy, University of Maryland, College Park, MD, USA.

ABSTRACT

The Combined Array for Research in Millimeter-wave Astronomy (CARMA) requires a flexible correlator to process the data from up to 23 telescopes and up to 8GHz of receiver bandwidth. The Caltech Owens Valley Broadband Reconfigurable Array (COBRA) correlator, developed for use at the Owens Valley millimeter-wave array and being used by the Sunyaev-Zeldovich Array (SZA), will be adapted for use by CARMA. The COBRA correlator system, a hybrid analog-digital design, consisting of downconverters, digitizers and correlators will be presented in this paper.

The downconverters receive an input IF of 1-9GHz and produce a selectable output bandwidth of 62.5MHz, 125MHz, 250MHz, or 500MHz. The downconverter output is digitized at 1Gsample/s to 2-bits per sample. The digitized data is optionally digitally filtered to produce bands narrower than 62.5MHz (down to 2MHz). The digital correlator system is a lag- or XF-based system implemented using Field-Programmable Gate Arrays (FPGAs). The digital system implements delay lines, calculates the autocorrelations for each antenna, and the cross-correlations for each baseline. The number of lags, and hence spectral channels, produced by the system is a function of the input bandwidth; with the 500MHz band having the coarsest resolution, and the narrowest bandwidths having the finest resolution.

Keywords: Radio astronomy, interferometry, lag (XF) correlator

1. INTRODUCTION

Advances in millimeter interferometry including large receiver bandwidths, and larger arrays, place an increasing demand upon the capabilities of the signal processing 'back-end', i.e., the correlator system. The increasing density of electronic components, or Moore's law, continues to make the development of bigger and better correlators possible, but the long development time for correlators means that they are often behind the state-of-the-art when they are actually put to use.

The objective of the digital correlator design described in this paper was to produce a wideband cross-correlator system, with an architecture that could smoothly evolve to exploit the continuing advances in chip performance, while exploiting commercial standards and technologies as much as possible. Examples of how this objective was met are seen in both the analog and digital subsystems. The analog subsystem downconverter design implements a versatile, wide bandwidth, sideband separating downconverter, using inexpensive industry-standard RF components and an embedded microcontroller for control and communications. The digital subsystem uses FPGAs to implement the correlation logic, 6U height Compact PCI (cPCI) crate technology, DSPs for real-time signal processing, LVDS (Ultra-SCSI cabling), and used subcontractors where possible during the board design and manufacturing phases.

The technical specifications for the analog and digital components used in the correlator system are given in the following section. Sections 2.1 and 2.2 discuss the analog components of the correlator system, with Section 2.1 describing downconverter system, and Section 2.2 describing the phase-modulated correlated noise source (used for real-time system calibration). Sections 2.3 and 2.4 present the digital components of the correlator system, with Section 2.3 covering the digitizer boards which implement autocorrelation, digital delay lines, and FIR filtering and Section 2.4 covering the correlator boards which implement cross-correlation and phase demodulation.

Further author information: (Send correspondence to D. W. Hawkins)
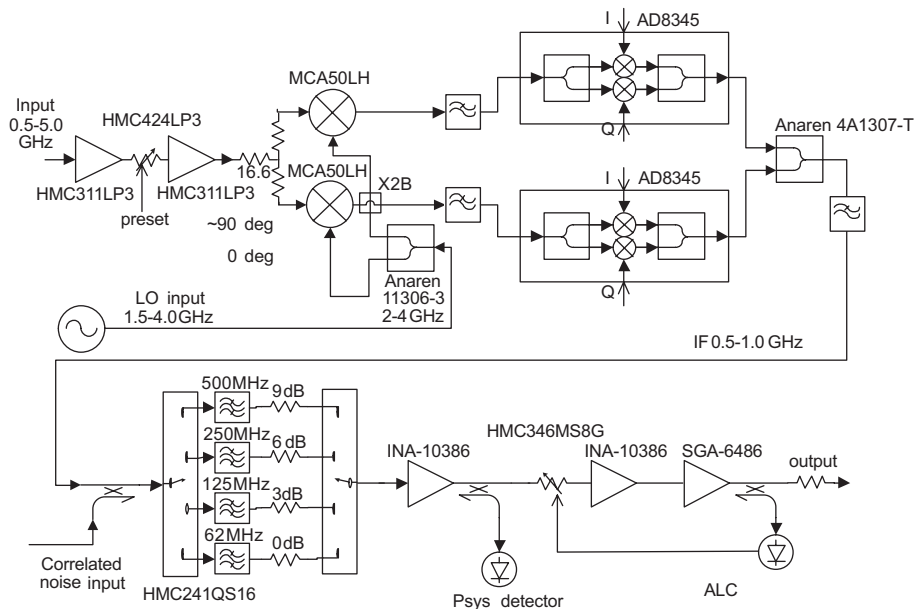E-mail: dwh@ovro.caltech.edu

Figure 1. CARMA Downconverter block diagram.

2. TECHNICAL SPECIFICATIONS

In the following sections, the technical specifications of the hardware and signal processing are described.

2.1. Analog Downconversion

The multi-gigahertz (GHz) receiver bandwidth achievable from modern millimeter receivers is wider than can be conveniently sampled and processed as a single band, so most modern correlator systems are a hybrid of analog filtering and downconversion, and digital correlation. Commercial digitizers exist in the 1GS/s to 4GS/s range, with custom samplers available at higher frequencies. However, the higher the frequency the sampler operates at, the larger the number of bits-per-second that the correlator system needs to process. The commercial samplers used in the CARMA correlator operate at 1GS/s. This selection was based in part on the fact that when this data is parallelized to a lower clock frequency, the resulting digital data bus-width is compatible with the chosen FPGA and cabling technology. The choice of sampling frequency then specifies the bandwidth that the downconverters must produce. With a 1GS/s sampling frequency, the downconverter bandwidth must be 500MHz or less, located at either baseband, or at a compatible aliasing band, eg. 500MHz to 1GHz, or 1GHz to 1.5GHz. The band 500MHz to 1GHz was chosen, due to the fact this is an octave of bandwidth and simplifies the downconverter design.

Millimeter observing can be broadly classified into two frequency bandwidth regimes; wide bandwidth continuum observing, and spectral line observing. A continuum correlator needs to cover as much bandwidth as possible (ideally the full receiver output bandwidth), with coarse resolution, while spectral line observing must support a variety of bandwidths, with a selectable location within the receiver bandwidth. The CARMA downconverters provide the required bandwidth and band-location flexibility. The CARMA downconverters accept a 1-5GHz IF band and produce an output signal within the 500MHz to 1GHz band. Figure 1 is a schematic block diagram of the downconverter, while Figure 2 shows a photo of a downconverter. A single sideband mixer is used to suppress the unwanted image sideband without the necessity of multiple conversions and filters. The output bandwidth is selectable from four filters, 500MHz, 250MHz, 125MHz, and 62.5MHz bandwidth.

The single sideband mixer is implemented using a resistive in-phase divider for the input signals, two mixers with a 90-degree hybrid splitter for the LOs, two quadrature modulators and an in-phase combiner. All of the components are inexpensive commercial chips. Good image suppression requires matching both the net amplitude and phases through the two mixer chains. It is difficult to buy components with these properties over
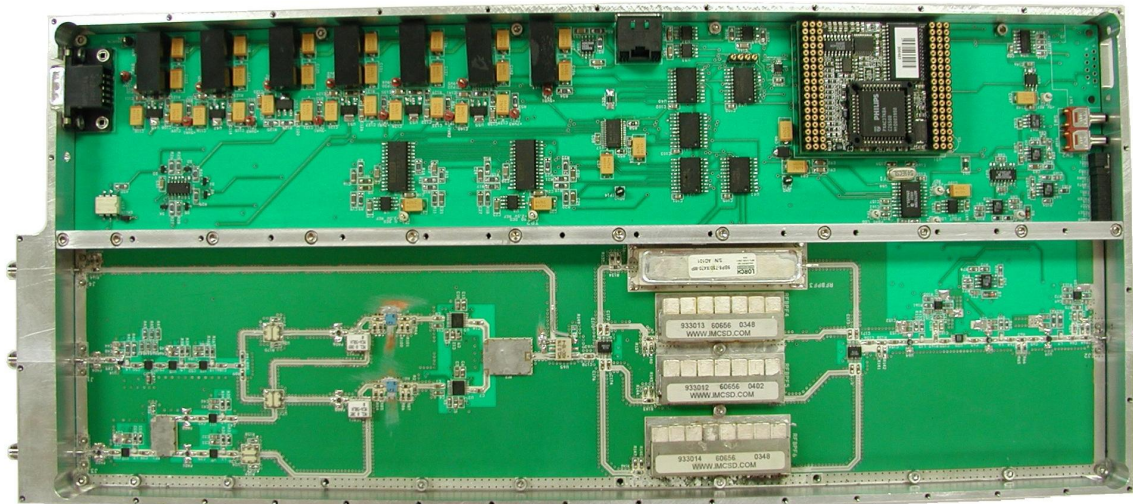
Figure 2. CARMA downconverter. The digital and RF sections of the downconverter are separated into two chambers. The top chamber of the downconverter contains the digital section, with the microcontroller near the top right. The bottom chamber contains the RF section. The four analog filters are the large components in the bottom chamber, just to the right of the center of the photo.

a wide fractional bandwidth. The solution chosen here is to use quadrature modulator chips operating in the 0.5-1.0GHz base band from each mixer to compensate for the imperfections of the various components. The broadband IF signal is fed into what would normally be the carrier input to the modulator while the modulation port is used to adjust the phase and amplitude. The modulation parameters as a function of input frequency are retrieved from a lookup table stored in EEPROM on the microcontroller module. The table is generated in a lab calibration of each downconverter. Image rejection of better than 30 dB is achieved over the full 500MHz output bandwidth for any selected input band over the 1GHz to 5GHz range.

2.2. Quadrature-modulated noise source

Data collected by the correlator system requires calibration. The signal output by the downconverter is always gain-controlled to produce an output power of nominally 6.0dBm. The digital correlator system then calculates auto-correlation and cross-correlation results that are relative to this nominal power level. Conversion of the correlator results back to radio source power requires the monitoring of various power levels throughout the RF chain. In addition to power-levels (which indicate an overall scaling of the data), individual channel scaling, or band-pass correction of the system is desired. Band-pass correction for the downconverter is performed using the quadrature modulated noise source.

The quadrature modulated noise source consists of a single 500MHz to 1GHz noise source, the output of which is split and sent to multiple quadrature modulators. Figure 3 show photos of the noise source and a quadrature modulator. The quadrature modulators impose a phase modulation onto the noise source, such that the quadrature modulator output emulates the phase switched signal received from an antenna. The quadrature modulated noise signal is input to the downconverter, where it is coupled with the downconverted signal from the antenna receiver (prior to final analog filtering). The noise source and quadrature modulators are disabled during radio source observing. When the correlator system is commanded to perform a noise source integration, the noise source and quadrature modulators are turned on, the digital delay lines are set to zero and a noise source integration is performed. Since the delays between the noise source and antenna receivers are different, integrations on the noise source see the receiver output contributions as uncorrelated noise (i.e., the cross-correlations do not have a correlation coefficient of unity). Autocorrelation and cross-correlation noise source integrations are used to generate band-pass corrections.
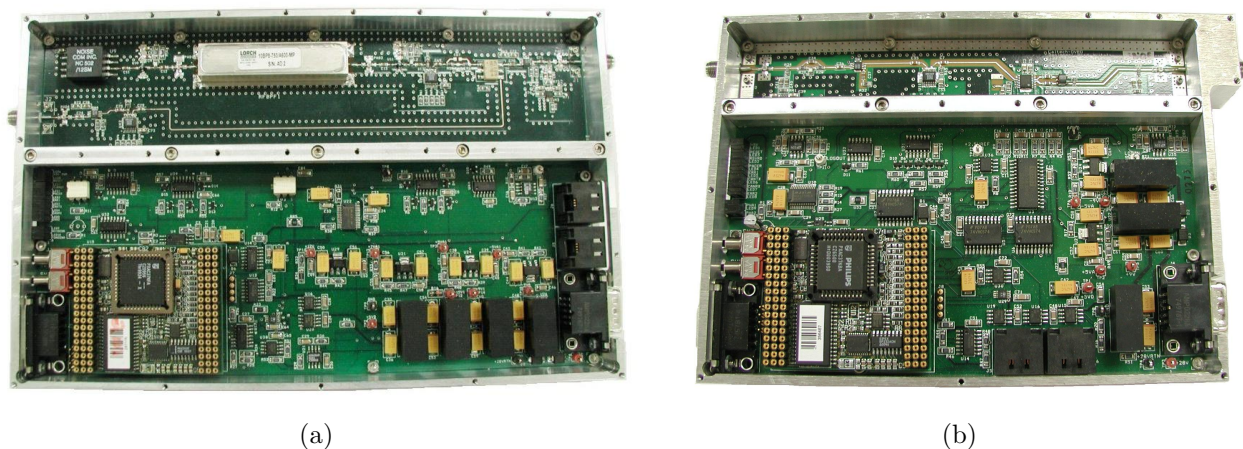
Figure 3. Quadrature modulated noise source. The quadrature modulated noise source consists of; (a) a 500MHz to 1GHz noise source feeding (b) multiple quadrature modulators; one per antenna. The quadrature modulators phase switch the noise signal so that it emulates the received antenna signals, producing a single-sideband noise source.

2.3. Digitization, FIR filtering, digital downconversion, and delay.

Figure 4 shows a COBRA digitizer board. The COBRA digitizer and correlator boards were designed around the Compact PCI (cPCI) 6U height board form factor. The COBRA digitizer boards consist of two 1GHz SPT7610 6-bit digitizers (originally manufactured by Signal Processing Technologies, now Fairchild), four Altera FLEX10K100E 2.5V Field Programmable Gate Arrays (FPGAs) (plus one used for dedicated system control), a Texas Instruments TMS320LC31 Digital Signal Processor (DSP), PLX Technologies PCI 9054 PCI interface, and associated support logic. The input to each digitizer is a 500MHz to 1GHz signal from the downconverter (or a narrower band within this range for the narrower spectral line modes). This signal is digitized at 1GHz to 2-bits of resolution. The digitizer ICs internally demultiplex this data by two producing 4-bits at 500MHz. ECL logic on the digitizer module demultiplexes this data further to 32-bits at 62.5MHz, and converts the data to CMOS /TTL logic levels. The CMOS/TTL data is then received by the digitizer board FPGAs.

The digitizer board FPGAs implement the following; quantization state monitoring, autocorrelation measurement, digital delay lines, 180-degree phase switch demodulation, FIR filtering, and digital demodulation. The quantization states, i.e., the number of times the digitizer outputs the bit patterns 00, 01, 10, 11 are monitored so that the threshold voltages on the digitizers can be adjusted until the quantization states are optimal for the 2-bit correlation scheme.[1,2] The autocorrelation for each antenna is used in the system temperature calibration. The digital delay lines are used when tracking sources. The 180-degree phase switching can removed either by synchronously inverting the sign bit output from the digitizer (the digitizer data is output in 2's compliment format), or the FPGAs can perform sign inversion when they receive the data (this operation must be synchronous with the phase switch reference).

Radio source delays are tracked in the correlator system by splitting the antenna delays into coarse and fine components. A coarse, antenna-based, correction for whole nanosecond delays is performed on the digitizer boards, and a fine, baseline-based, correction for the remaining sub-nanosecond delay component is performed at the correlator boards. The coarse delays implemented in the digitizer board FPGAs provide for delays of single samples (a 1ns delay) up to the maximum RAM available in the FPGA. The Altera FPGA10K100E FPGAs contain 12 embedded memory blocks, each containing 4096 bits. Two memory blocks are used by the autocorrelation logic, leaving 10 blocks for implementing the delay line, i.e., 40960 bits. At two bits per sample,

Figure 4. COBRA Digitizer Board; Two 1GHz 2-bit digitizers, 4 Altera FLEX10K FPGAs, a TI TMS320LC31 DSP, a PLX 9054 PCI interface, and LVDS front panel data distribution.

and a sample spacing of 1ns, this corresponds to a total delay of 20.48 μs, or a little over 6km at the speed-of-light. Compensation for sub-nanosecond delay components is performed on the correlator boards by multiplying the measured cross-power spectra with a complex-valued correction function that compensates for both delay (phase slope) and lobe rotation (phase offset). The correction is performed using single-precision floating-point numbers.

To process different bandwidth signals, the correlator system performs a combination of analog and digital filtering. The system requires additional analog filtering for two reasons; only 2-bits from the digitizers are available at the FPGAs, and the FPGAs have a limited amount of resources. At input bandwidths of 500MHz, 250MHz, and 125MHz, the spectral response of the measured cross-power spectra is determined solely by the analog filter response. For the signal bandwidths of 250MHz and 125MHz, signals are digitized at 1GHz and the FPGAs perform decimation by two or four respectively. For signal bandwidths of 62.5MHz or below, an analog filter of 62.5MHz bandwidth centered on 750MHz is used, and the filtered signal is digitized at 1GHz to 2-bits, effectively oversampling the signal by a factor of 8. The 62.5MHz filter is centered on 750MHz so that after digitization the center frequency of the band aliases to 250MHz, i.e., one-quarter the sampling frequency. This center frequency allows for simple demodulation of the signal into a complex-valued baseband form. The oversampling of the original input signal allows for the use of a small FIR filter within the FPGA to remove the signal energy in the other sideband (which demodulation moved to the highest frequencies) and remove the majority of the 2-bit digitization quantization noise. The complex-baseband FIR filtered signal is then decimated by eight, to an effective clock rate of 125MHz (i.e., the Nyquist rate of the 62.5MHz band). Additional baseband FIR filter stages are used to produce the bands narrower than 62.5MHz. Prior to data being sent to the correlator boards, it is remodulated (to produce a real-valued signal) and requantized to 2-bits. The overall efficiency of the quantize-filter-quantize scheme is within a few percent of the 500MHz/250MHz/125MHz quantize-only scheme.

Digitized data is routed to the correlator boards by way of front panel low-voltage differential signaling (LVDS) cabling (which contain 32-bit data plus clock differential signals). The FPGAs drive output data to LVDS transceivers, which in turn drive cables on the front panel. The digitized data received by the FPGAs nearest the two digitizers can be routed to all four FPGAs on the digitizer board. The worst-case data volume is when the digitizers are digitizing 500MHz bandwidth signals, as each produces a 2Gbit/s (250MB/s) data
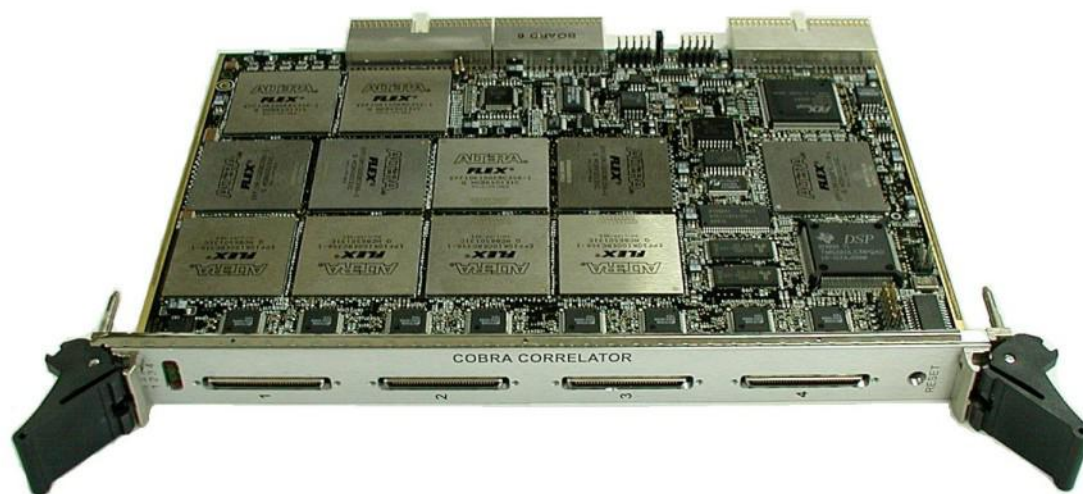
Figure 5. Correlator board.

stream. There are three main data transport modes; in the first mode, the data from each digitizer is fanned out to the two nearest LVDS connectors and is transported as 32-bit data at 62.5MHz (250MB/s per connector), in the second and third modes, the data from each digitizer is fanned out by four and the data from both digitizers is transported on all four LVDS cables as 16-bit data at 125MHz (500MB/s per connector) or 62.5MHz (for the narrower bandwidth modes).

Phase-switching within the CARMA system will be performed on a $\sim$ 1ms timescale (500ms/512-phase states to be exact). The phase-switch sequency used contains nested 180-degree sequences 16-states long. Data from within a 16-state 180-degree phase-switch sequence is demodulated (by sign-inversion) by the digitizer board FPGAs, so that the correlator board FPGAs do not need to demodulate the 180-degree phase switch cross-correlation sequencies. Since phase-switching of the telescope local-oscillators (LOs) is not 'instantaneous', the digitizer and correlator boards are synchronized to the phase-switch reference, and only integrate data when the phase-states have 'settled'. At the beginning of a phase switch cycle, both the digitizer and correlator boards stop integrating for $30\mu s$ after a phase switch edge (this off-time is ultimately determined by the settling time of the IF system). The digitizer boards calculate autocorrelation estimates, and accumulate the estimates internal to the FPGAs for the duration of a 16ms 180-degree phase switch sequence. At the end of a 16-state phase-switch sequence, the data is transferred to the on-board DSP where it is normalized, converted to floating-point format, Fourier transformed to power spectra (using the FFT), and accumulated further. Autocorrelation data is transferred to the host CPU every 500ms. Monitor point data is also transfer to the host every 500ms (board temperatures, supply voltages, processing time statistics, etc).

2.4. Correlator

Figure 5 shows a COBRA correlator board. The COBRA correlator boards consist of ten Altera FLEX10K100E 2.5V Field Programmable Gate Arrays (FPGAs) (plus one used for dedicated system control), a Texas Instruments TMS320LC31 Digital Signal Processor (DSP), PLX Technologies PCI 9054 PCI interface, and associated support logic. Each correlator board receives data from multiple digitizer boards via front panel LVDS receivers. Each cable can be used to deliver the digitized signal from a single-telescope (at up to 32-bits at 62.5MHz) or two-telescopes (2×16-bits at 125MHz).

The 10 FPGAs on the correlator boards are used to calculate either 5 baselines or 10 baselines of cross-correlations. Section 3 discusses how these options map onto the OVRO, SZA, and CARMA correlator systems.

The correlator board FPGAs calculate and accumulate cross-correlation lag estimates for the duration of the 16-state 180-degree phase switch sequence (the 180-degree phase switch is removed by the digitizers, so the FPGAs simply need to accumulate the data over the valid period within a phase-switch period). At the end of a 16-state phase-switch sequence, the data is transferred to the on-board DSP where it is normalized, converted to floating-point format, Fourier transformed to power spectra (using the FFT), sideband separated, phase compensated, and accumulated further. Crosscorrelation and monitor point data is transferred to the host CPU every 500ms.

The interferometer systems in which the COBRA hardware has been used all use double-sideband mixers in the antennas and employ 90-degree phase switching to separate the sidebands. In the SZA and CARMA systems the phase of the LO reference to each antenna is generated using a lower frequency direct-digital-synthesis (DDS) output, and up-conversion to generate the frequency required at the antenna. Each DDS generates phase as the sum of the per-antenna phase switch state, and a continuously-changing offset determined by the radio source position (which is used to calculate a per-antenna delay). The phase offset is referred to as the first LO lobe rotation, and its correction is referred to as lobe rotation correction, or 'fringe-stopping'. The lobe rotation phase rotates rapidly, and if the phase term was not corrected (tracked) at the first LO, the correlator response would 'wash-out'. Demodulation of the antenna receiver output results in a secondary lobe rotation term that must be corrected for on the time-scale of 20ms (for CARMA's longest baseline). This time-scale is short enough that the correlator boards can perform the correction. The (more expensive) alternative is to use DDSs on the downconverter LOs.

A more detailed description of the processing performed on the digitizer and correlator boards is now appropriate. The correlation logic implements a 2-bit deleted inner product multiplication scheme.[1,2] The multiplication of 2-bits from telescope X with 2-bits from telescope Y can be described as a 4-input, 3-bit 16-output lookup table (LUT). Each bit in the 3-bit output LUT maps exactly to the 4-input LUT architecture of the Altera FLEX10K FPGAs. The 3-bit output of a multiplier LUT is accumulated by a 3-bit adder, producing a multiplier-adder component (requiring a total of 7 logic cells). Wide bandwidth input signals are correlated using parallel processing. For example, for the widest bandwidth mode of 500MHz, 2-bit digitized data is packed into a 32-bit word at 62.5MHz (16-samples in parallel) and distributed to the correlator boards at either 32-bits at 62.5MHz or 16-bits at 125MHz. Data transported as 32-bits at 62.5MHz is received by the correlator boards and remultiplexed back to 16-bits at 125MHz internal to the FPGA (FPGA-to-FPGA data transmission is performed at 62.5MHz where possible to maximize the available setup and hold-time between FPGAs, and internal data processing is performed, for the widest bandwidth modes, at 125MHz to maximize logic usage). The correlator logic is parallelized so that each of the 8-samples packed into the 125MHz data streams coming from the digitizer boards is cross-correlated. This produces 8-independent samples that need to be accumulated into the same lag. Rather than using 8-independent counters and summing the contents of those counters when the data is read-out, the FPGAs use full-adders and add pairs of carry bits from the multiplier-adders. The full-adders are cascaded until there is a single carry output. This single carry output is then prescaled (to reduce its toggling rate) and then the prescaled carry is accumulated. The carry-cascade and prescaling logic effectively truncates the least-significant-bits of the correlation result. This truncation relies on the assumption that the input signal is random gaussian noise, and that the signal has been accumulated for some minimum integration time (so that the final counter has a significant count value along with noise bits in the counter least-significant-bits). Prescaling of the lag outputs allows for the counters in the FPGAs to be operated at a lower-frequency, eg., in the case of the 500MHz band, the lag logic operates at 125MHz, yet the counters can operate at around 16MHz. This scheme results in lower power consumption, and eases the timing requirements within the FPGAs (the lag logic is 'critical', while the counter logic is not as critical). At the end of a 16-state 180-degree phase-switch sequence, the correlator lag counters are copied to RAM internal to the FPGAs and the counters are cleared. This allows the DSP time to read-out the lag results for the previous 16ms correlation, while the FPGAs are busy calculating the current 16ms correlation.

Each correlator board DSP reads the 16ms lag counters from the correlator FPGAs, converts them to floating-point format, normalizes them, then performs an FFT to produce cross-power spectra. The data is then phase corrected. Phase correction consists of applying the baseline-based sub-nanosecond delay correction (the digitizer boards remove whole nanosecond delays on an antenna basis), and applying a downconversion band specific offset correction. To compensate for these two phase components, the correlator needs to be sent delay information

(continually updated) and frequency information (which is setup at the start of observing). For double sideband data, the phase slope due to delay is identical in each sideband, however, the phase offset is not. Within a 16ms phase switch data set, there is insufficient data to sideband separate, so the data must be corrected for each sideband, and accumulated into a lower sideband sum and an upper sideband sum. When the 90-degree phase switch cycle completes (every 500ms) the two sideband sums contain sideband separated and phase corrected data. In addition to delay and frequency information, the correlator needs to be told when a noise integration is being performed. For noise integrations, the delay lines on the digitizers should be zeroed, and no delay correction be performed. The 90-degree phase switch sequence completes every 500ms, at which time the accumulated and sideband separated data is transferred over the PCI bus to the Linux host CPU, where the data is transferred over ethernet into the data processing pipeline.

The COBRA Correlator boards are reconfigured for each input bandwidth setting. The digitizer board FPGAs and correlator board FPGAs are downloaded with bandwidth specific configurations. This reconfiguration allows for optimal reuse of the FPGA resources, i.e., narrower bandwidth settings increase their spectral resolution due to the decrease in bandwidth *and* increase in the number of lags produced by the autocorrelation and crosscorrelation logic. The DSP and Host CPU code is general purpose and reused for each band setting.

3. CORRELATOR SYSTEMS

The following sections describe the application of the COBRA digitizer and correlator boards for the OVRO 6-telescope, 4GHz bandwidth, COBRA Wideband correlator system, the SZA 8-telescope, 8GHz bandwidth, correlator system, and the CARMA 15-telescope, 4GHz bandwidth system.

3.1. COBRA Wideband System

The COBRA Correlator design was initiated at the Owens Valley Radio Observatory (OVRO) to replace their aging analog continuum correlator system, and to potentially replace their spectral line system. The main requirement for the wideband system (in reality a coarse resolution spectral line system) was that it produce a channel spacing on the order of 20MHz so that the system could be used in the search for high-redshift galaxies. The OVRO interferometer contained 6-telescopes, each producing an output bandwidth of 4GHz located at an intermediate frequency (IF) of 0.5GHz to 4.5GHz (the telescopes are now part of the CARMA array). The COBRA Wideband System processed the 4GHz IF as 8 500MHz bandwidth bands, with the downconverter LOs at fixed frequencies across the IF (there was no tunability). Each 500MHz band in the correlator system produced 6 autocorrelations and 15 cross-correlations using 3 digitizer boards (digitizing 2 antennas per board) and 3 correlator boards (processing 5 baselines per board). Data was distributed between digitizer and correlator boards on the LVDS front panel cables as 32-bit data at 62.5MHz. The correlator boards produced 64-lags per cross-correlation, which Fourier transformed to produce a 33-channel spectrum (positive frequencies) with 15.625MHz channel spacing for each sideband. The digitizer board autocorrelations were calculated across 33-lags (positive lags) to produce a matching 33-channel spectrum. This scheme resulted in the combined 4GHz spectra for each sideband being covered by 257-channels with 15.625MHz spacing. If the lag data was Hanning weighted, then the spectrum would be decimated by two to produce 17-channels per band, 129-channels across 4GHz, with 31.25MHz channel spacing.

Figure 6 shows a photo of the COBRA Wideband System. On the left of the image are the downconverter modules, (analog) RF cables link the downconverter outputs to the digitizer inputs, and the digitizer boards are linked to the correlator boards via the (digital) LVDS cabling. Section 4 contains examples of the measurements obtained using the COBRA Wideband System.

3.2. SZA Correlator System

The SZ Array (SZA) primary objective is the detection and analysis of the Sunyaev-Zeldovich (SZ) Effect. The requirements for the correlator were that it process an 8GHz wide bandwidth, be able to track source delays (since the antennas were not mounted on a common platform), and be cost-effective. The SZA Array consists of 8-telescopes and a receiver IF bandwidth of 8GHz (over the frequency range 1GHz to 9GHz). The correlator system requirements were met using the COBRA analog and digital hardware, and an additional downconversion stage. The additional downconversion stage takes the 1GHz to 9GHz antenna IF and produces two 1GHz to 5GHz

Figure 6. The 4GHz bandwidth COBRA Wideband Correlator System. The analog downconverter system is shown on the left-hand-side of the photo, while the digital correlator is shown on the right-hand-side. The black cabling is for the RF signals, while the lighter cabling is the digital LVDS cabling.

output signals; one from the lower 4GHz of the input band, the other from the upper 4GHz downconverted using a 10GHz lower-sideband downconversion. The two 1GHz to 5GHz bands are then processed by two identical blocks of downconverters. The downconverter modules were upgraded relative to the COBRA versions to improve their spectral response over the 1GHz to 5GHz band. The initial SZA downconverters are fixed frequency models as used for the COBRA Wideband system. The downconverters will eventually be upgraded to the flexible frequency spectral line downconverters used in the CARMA system.

Each 500MHz band in the SZA Correlator system produces 8 autocorrelations and 28 cross-correlations using 4 digitizer boards (digitizing 2 antennas per board) and 3 correlator boards (processing 10 baselines per board). The number of baselines calculated by the correlator boards was boosted to 10 per board relative to the COBRA Wideband System, since the SZA does not require high resolution and this reduces cost. The increase in the number of baselines requires that each correlator board receive a larger number of antennas. By increasing the data rate on the LVDS cables, each LVDS cable can route the pair of antennas from a digitizer board on the same cable, hence, a correlator board can receive up to 8 antenna inputs. The cabling of the SZA correlator is such that each correlator board receives 3 cables of LVDS data, containing 6 antennas worth of digitized data, and calculates 10 cross-correlations (from the 15 possible baselines for those antennas). Each auto- and cross-power spectrum for a 500MHz band contains 17-channels, with 31.25MHz channel spacing. The combined spectra for each of the 8GHz bandwidth sidebands is then covered by 257-channels of 31.25MHz channel spacing (with Hanning weighting and spectral decimation by two doubling this channel spacing).

3.3. CARMA Correlator System

The Combined Array for Research in Millimeter-wave Astronomy (CARMA) (`www.mmarray.org`) was formed by merging Caltech's Owens Valley Radio Observatory (OVRO) millimeter array and the Berkeley-Illinois-Maryland Association (BIMA) millimeter array. CARMA contains a total of 15-telescopes made up of the six 10m diameter telescopes from OVRO and nine 6m diameter telescopes from BIMA. The antenna IFs on both systems will be upgraded so that both systems initially support receiver IFs of 4GHz, with wider bandwidths in the future. The requirement of the 'First Light Correlator' is to process 15-telescopes, 105-baselines, and 4GHz bandwidth. The downconverters in the CARMA system are the spectral line downconverters discussed in Section 2.1. The CARMA First Light Correlator will contain 8 flexible frequency bands, where each band can be independently tuned within the 4GHz IF, with an independent bandwidth. Practical use of the system will configure some bands for continuum measurement, while others will be used for spectral line work. A band in the CARMA correlator consists of 8 digitizers boards (to process the 15 telescopes), 5 fanout boards, and 11 correlator boards (each calculating 10 baselines per board). The required fanout of digitized data from the digitizer outputs to the correlator inputs exceeds the number of front panel connectors on the digitizer boards (a fanout of six is needed from some boards), so in addition to the digitizer and correlator boards the CARMA crates contain fanout boards that receive the data from a single digitizer board and fan it out using LVDS transceivers to 6 output connectors. Each CARMA band will be housed in two compact PCI crates, with the two crates controlled by a single CPU and the crates linked via a PCI-to-PCI bridge (this makes the use of two crates transparent to the control CPU). Each band produces 15 auto-correlations and 105 cross-correlations. The 4GHz COBRA Wideband System consists of 24 digitizer boards and 24 correlator boards (plus spares of each). This system and the spares will be recycled into the first 1.5GHz of the CARMA First Light Correlator (requiring 24 digitizers and 33 correlator boards). The remaining 2.5GHz will be generated as part of a digitizer and correlator board upgrade. The upgraded boards will implement higher resolution digitization and will contain higher density FPGAs. The new boards will allow for increased flexibility in the FIR filtering on the digitizer board, and higher resolution auto- and cross-correlations. Table 1 shows the spectral channel spacings for each of the input bandwidths for the COBRA hardware based bands. The digitizer and correlator board upgrades are not complete, so their characteristics are not yet available.

The SZA and CARMA Correlator systems will initially operate as fairly independent systems; the SZA 8GHz wideband correlator, and the 4GHz CARMA spectral line correlator. Cross-correlation between antennas in the two arrays will be achieved through the use of a manual RF switch-yard. Ultimately it is desired to have a more automated method for supporting a 23-element interferometer, and for supporting multiple sub-arrays within these 23-telescopes. Recent advances in the telecommunications industry now support the distribution and routing of multi-gigabit data streams through network switching hardware (this high bandwidth equipment forms

Table 1. CARMA correlator band lag and spectral characteristics.

Bandwidth nominal	(actual)	Number of Lags	Number of spectral channels	Spectral channel spacing
500MHz		32	17	31.25MHz
250MHz		64	33	7.81MHz
125MHz		80	41	3.13MHz
62MHz	(62.5MHz)	96	49	1.29MHz
31MHz	(31.25MHz)	112	57	558kHz
8MHz	(7.81MHz)	128	65	122kHz
2MHz	(1.95MHz)	128	65	31kHz

the back-bone of the internet). Future correlator research at CARMA will look at exploiting this technology for the transport of digitized data to the correlator boards. The use of a configurable data routing network between the digitizer outputs and correlator inputs will produce a correlator system that can arbitrarily route antenna (digitized) data to cross-correlation sub-arrays.

4. SCIENCE

The COBRA Correlator system was installed onto the Owens Valley millimeter array during 2002. Figure 7 shows one of the first astronomical tests using a single 500MHz bandwidth band. Continuous observing with the system occurred in early 2003. The system was fully operational during the complete 2003 to 2004 observing season up until the OVRO shutdown, due to the change-over to CARMA, in June 2004. Scientific data obtained during the 2003 to 2004 observing is now being used in publications.

One of the original motivations for the development of the wide bandwidth COBRA correlator system, was the search for CO emission in high-redshift objects. These searches have traditionally suffered due to the lack of accuracy of optically-derived redshifts and the limited correlator bandwidth available for the search. Figure 8 shows a 3GHz spectrum containing a QSO line detection from the first COBRA publication[3]. Hainline *et al*[3] contains maps and analysis for several radio sources. There are several other publications currently in submission and review.

ACKNOWLEDGMENTS

REFERENCES

1. B. F. C. Cooper, "Correlators with two-bit quantization," *Australian Journal of Physics* **23**, pp. 521–527, 1970.
2. J. B. Hagen and D. T. Farley, "Digital correlation techniques in radio astronomy," *Radio Science* **8**, pp. 775–784, Aug. 1973.
3. L. J. Hainline, N. Z. Scoville, M. S. Yun, D. W. Hawkins, D. T. Frayer, and K. G. Isaak, "A Study of CO Emission in High Redshift QSOs Using the Owens Valley Millimeter Array," *Astrophysical Journal* **609**, July 2004.
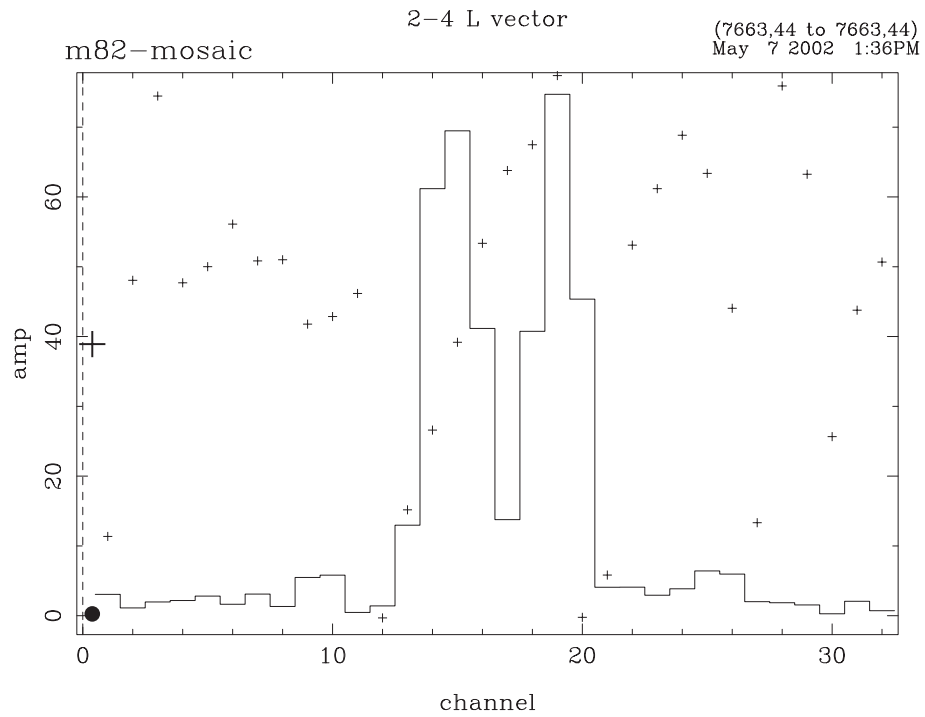
Figure 7. M82 Spectra (500MHz bandwidth, 32-channels)–one of the first astronomical measurements obtained using COBRA.
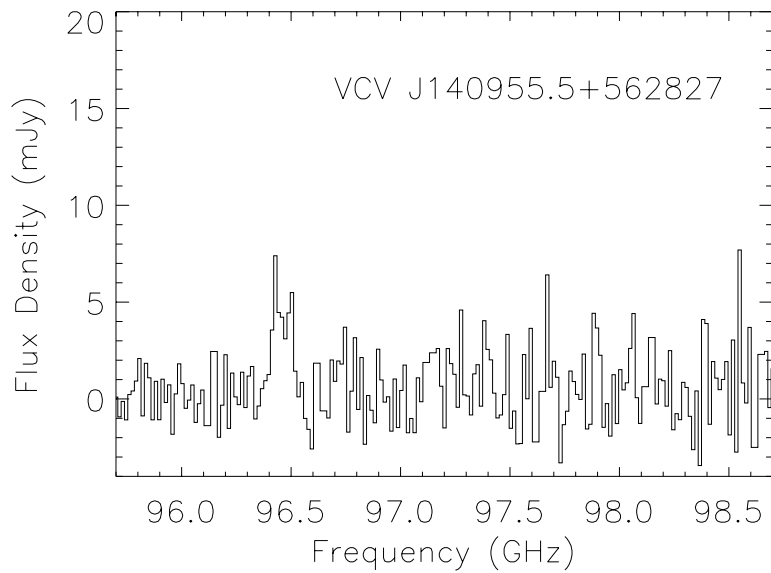

Figure 8. High redshift QSO detection using COBRA. Hainline *et al*[3] contains the analysis of, and maps obtained from, several QSO detection experiments.