

CARMA Memorandum Series #46

Revised CARMA Correlator FPGA Configurations

Kevin P. Rauch (UMD)

September 3, 2008

ABSTRACT

This memo provides technical information on a variety of FPGA-level design details for the revised CARMA correlator digital hardware. This includes descriptions of the board-level FPGA layout, communication paths, and baseline partitioning; specifics of fundamental VHDL design components, and their distribution by FPGA; and the FPGA memory map and control register specification.

Change Record

Revision	Date	Author	Sections/Pages Affected
			Remarks
0.1	2007-July-09	Kevin Rauch	
	Initial draft.		
0.2	2008-January-14	Kevin Rauch	
	Added descriptions of memory map, correlation logic, fractional sample delays and phase offset processing.		
0.3	2008-August-15	Kevin Rauch	
	Assigned official CARMA memo number. Updated introduction.		
1.0	2008-September-3	Kevin Rauch	
	Documented delay/phase coefficient table format.		

1. The Revised CARMA Correlator

The initial CARMA correlator system, consisting of three bands (up to 1.5 GHz bandwidth) of recycled COBRA hardware, became fully operational in 2006. At that time, development proceeded on revised and upgraded digital hardware, which will replace the original hardware and expand the correlator system to eight bands (up to 4 GHz total bandwidth). The upgraded CARMA hardware will provide $\sim 6x$ the channel resolution per baseline compared to the existing COBRA hardware when cross-correlating 2-bit samples, and will also support new capabilities, such as FPGA-based fractional sample delays, cross-correlation of 3- and 4-bit samples, and the ability to be reconfigured for dual-polarization observations. This document discusses several FPGA-specific design details of the revised hardware, such as layout of the FPGAs communication buses, FPGA memory map and location definitions, and the specifics of pertinent VHDL component implementations. When referencing specific files in the `carmacorl` CVS module, `$CCORL` is used to represent the top-level project directory.

2. Physical Configuration

In contrast to COBRA, the revised CARMA digitizer and correlator boards share a unified PCB design. Each board contains 4 correlation/signal processing FPGAs, and one system controller FPGA acting as “glue logic” between the correlation FPGAs and the board CPU. All FPGAs are Altera Stratix II devices; specifically, an EP2S60F1020C3 device for the controller FPGA, EP2S90F1020C3 devices for digitizer card data FPGAs, and EP2S130F1020C3 chips for correlator card data FPGAs. Each digitizer card receives RF input from two antennas. The correlator boards can also be used as digitizers, once the PCB is loaded with A/D conversion logic. The unified PCB design also simplifies the HDL coding; since all digitizer and correlator data FPGAs share the same signal pinout, the same high-level VHDL components can be used in both instances. All new and upgraded VHDL components were designed to be shareable in this way.

Although the physical communication bus layout is identical for the digitizer and correlator boards, the bus directions and signal processing tasks differ between the two. Figures 1 and 2 display the data pipeline and shorthand VHDL bus naming conventions for the digitizer and correlator data FPGAs, respectively, as well as the high-level data processing components present in each FPGA. All buses shown are 32-bits wide (excluding bus clocks, where applicable), except for the digitizer input bus (`dig`), which is 17-bit. The Stratix II FPGAs integrate support for a multitude of single-ended and differential I/O standards. Two such I/O standards are used in the revised correlator design: a differential standard (2.5 V LVDS) is used by the front panel I/O buses (`ext`) and rear digitizer input (`dig`); the remaining buses are single-ended, 2.5 V LVCMOS connections. Of the latter, eight buses, `1a` through `1d` and `2a` through `2d`, are used to transfer digitized samples (in various stages of processing) between FPGAs; two, `cpu1_adq` and `cpu2_adq` (not shown), are used to transfer data to the board (or host) CPU via the system controller FPGA.

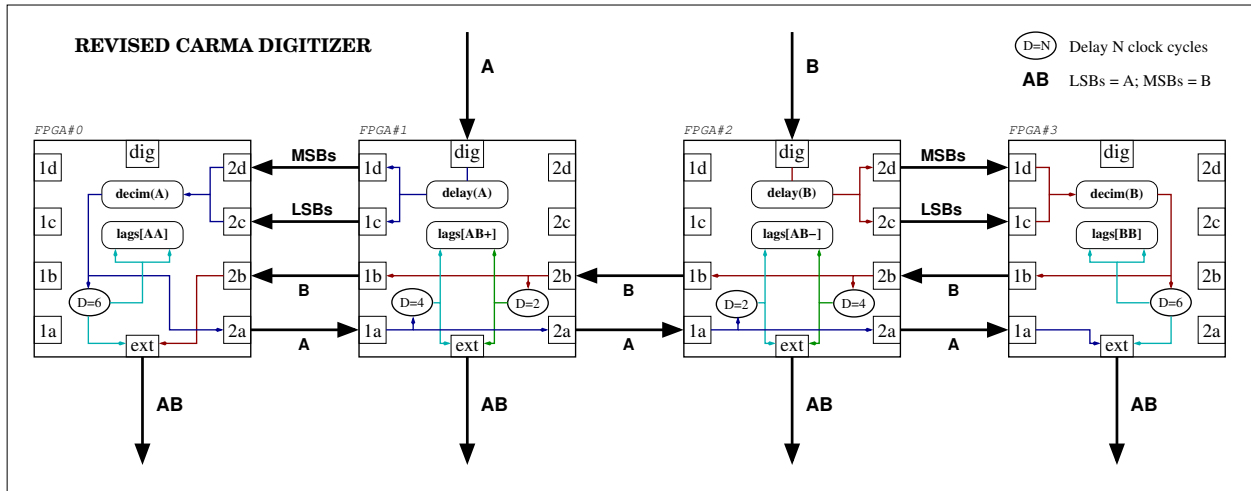


Fig. 1.— Revised CARMA digitizer board data pipelines and signal processing component layout.

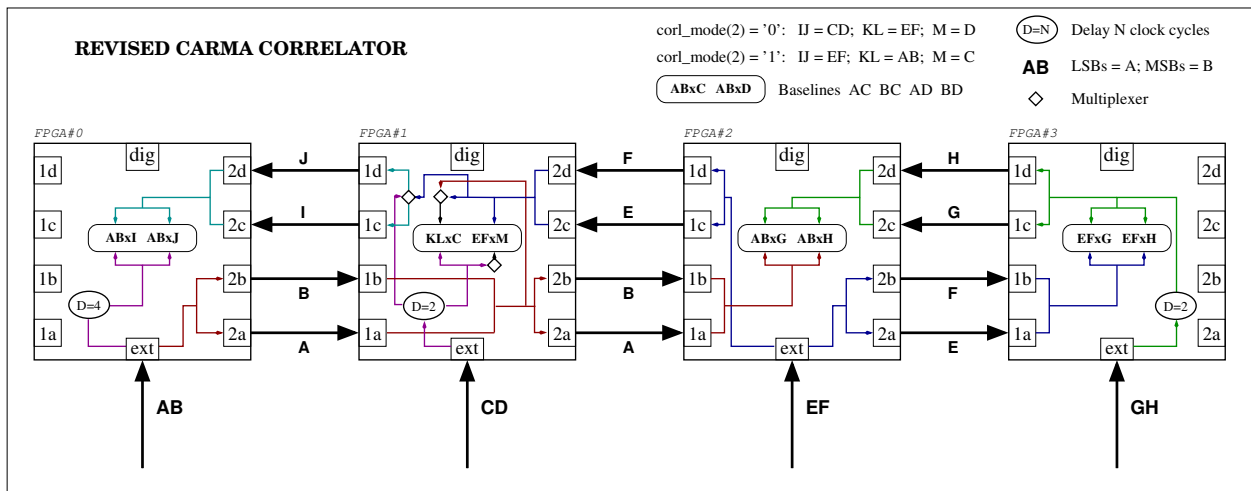


Fig. 2.— Revised CARMA digitizer board data pipelines and signal processing component layout.

3. Baseline Partitioning

The revised correlator boards contain four data FPGAs, each calculating the cross-correlation of four individual baselines, for a total of 16 baselines per board. For CARMA-15 (15 antennas), there are 105 unique baselines. Eight digitizer cards (each with 2 RF antenna inputs) are needed per band. Due to increased logic availability in the digitizers, in the revised hardware each digitizer card can be responsible for computing the cross-correlation between its own antenna input pair (cf. Figure 1). This leaves 98 baselines to be computed by the correlator cards; hence 7 correlator cards are required per band—although one of them calculates only two unique baselines. As shown in Figure 3, only two distinct baseline-to-FPGA partitioning geometries are needed to distribute the baselines among the 7 cards. The two geometries are similar enough that a single FPGA configuration can be used for both; which one to use is determined by the setting of a `cor1_mode` bit, and can be changed dynamically. The maximum fan-out any antenna is four, meaning that no fan-out boards are required for the revised hardware. In contrast, the COBRA boards required three distinct partitioning geometries, each necessitating a unique FPGA configuration, and several fan-out boards. Conceptually, the revised hardware moves card-to-card cabling fan-out into PCB (and internal FPGA) traces. Note that if digitizer cards are not used to cross-correlate their own baselines, fan-out boards would be required.

4. FPGA Memory Map

The revised FPGA memory map consists of a set of control registers and M-RAM blocks organized into a contiguous address space, as seen by the external CPU interface. Internally, each M-RAM possesses a dedicated read/write bus. Stratix II M-RAM blocks support true dual-ported, mixed-width configurations; the new memory map component (`mmap`) instantiates a 32-bit port to communicate with the external interface and a 64-bit port for use by internal logic. Memory configuration is little-endian in this regard: writing a 64-bit value to memory via the 64-bit port to an address A is equivalent to writing the 32-bit LSBs and MSBs to addresses $2A$ and $2A + 1$, respectively, via the 32-bit port.

The memory map implements a set of 32 x 32-bit control registers, beginning at address 0x0; the first 16 of these are read-only (cf. Table 1). The remainder of physical memory is provided by M-RAM blocks. Each M-RAM contains 2^{19} bits of memory, equivalent to `MMAP_BLOCK_SIZE = 0x4000` 32-bit words. The memory map reserves address space for 8 M-RAM blocks per FPGA, `MMAP_FPGA_SIZE = 0x20000` (17-bit local addresses), and allocates three additional address bits (MSBs) for the chip select, for a total of `SYS_ADDR_WIDTH = 20` bits of FPGA address space visible to the CPU. This exceeds the physical limits of 5 FPGAs and (up to) 6 M-RAM blocks per FPGA present on the revised boards. A generic (`NUM_BLOCKS`) determines the actual number of M-RAMs consumed by the memory map in any particular FPGA; reads from addresses for which no physical RAM has been allocated return a pre-defined hex string (currently 0xDEADBEEF).

The detailed memory map is listed in Table 1. Control registers occupy the first 32 locations, and shadow the corresponding region (= 0.2%) of the first M-RAM block, which physically begins at address 0x0.

Revised Correlator Baseline Partitioning

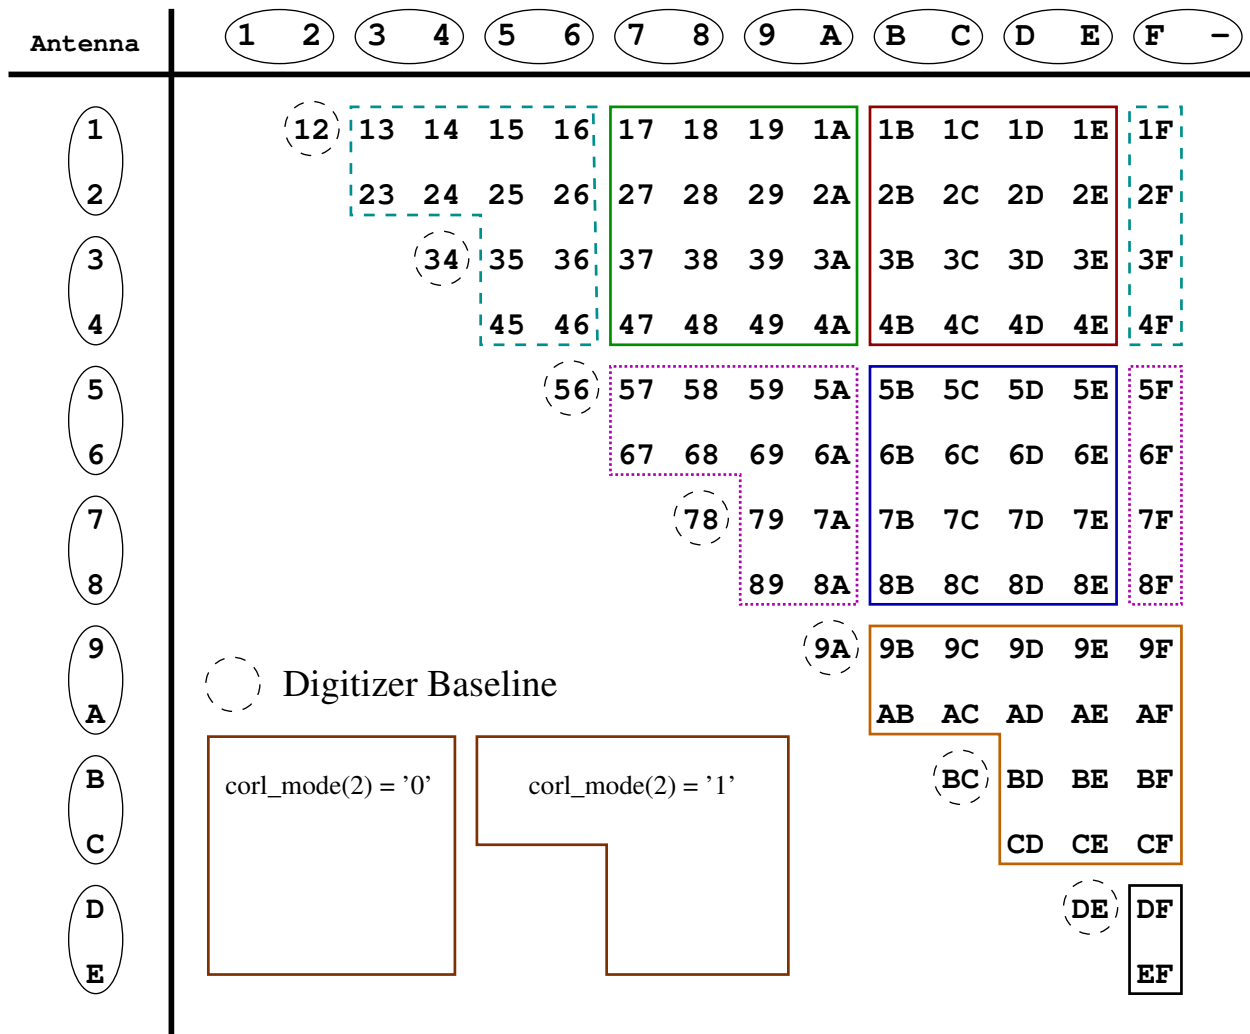


Fig. 3.— Revised correlator baseline partition map for CARMA-15. Each band requires 8 digitizer cards and 7 correlator cards.

Subsequent M-RAMs are contiguous—the starting address for block i is $A = i \times \text{MMAP_BLOCK_SIZE}$. The identical layout is used in both digitizer and correlator FPGAs. The first M-RAM block is reserved for the phase/delay coefficient tables used in digitizer FPGAs #1 and #2 (cf. Figure 1); the table contents are described in § 6.2. Subsequent blocks, one per baseline, are reserved for lag readout data; see § 5.

The control register contents are as follows (omitted bitfields are undefined):

- CTRL_REG_VERSION [Version (read-only)]
 - **bits 23-20:** FPGA type; 0xC = correlator, 0xD = digitizer
 - **bits 19-16:** hardware revision; 0xC = COBRA final, 0xD = revised prototype
 - **bits 15-7:** major version number
 - **bits 7-0:** minor version number
- CTRL_REG_COMPAT [Configuration compatibility (read-only)]
 - **bits 31-16:** *reserved* (configuration type field)
 - **bits 15-0:** bit n is set iff the configuration is compatible with FPGA # n .
- CTRL_REG_CORL_CONF1 [Correlation configuration register 1 (read-only)]
 - **bits 30-20:** Metadata elements per lag stream (NUM_META)
 - **bits 19-8:** Number of lags per stream (NUM_LAGS)
 - **bits 7-4:** Number of correlation baselines (NUM_CORL)
 - **bits 3-0:** Correlation type; 0 = auto, 1 = cross +lags only, 2 = cross -lags only, 3 = cross +lags and -lags
- CTRL_REG_CORL_CONF2 [Correlation configuration register 2 (read-only)]
 - **bits 30-27:** Written lag block address width (1+DUMP_ADDR_WIDTH)
 - **bits 26-12:** Written lag block locations ($2 \times \text{DUMP_COUNT}$)
 - **bits 11-0:** Number of quantization state counters (NUM_QCNT)
- CTRL_REG_TD_1A to CTRL_REG_TD_2E [Inter-FPGA bus readback registers (read-only)]

These control registers mirror the current (or “recent”) contents of the associated IOE registers (cf. Figs. 1-2).
- CTRL_REG_STATUS [Correlation status]
 - **bits 30-8:** Counter indicating the number of clock cycles the correlation logic was active (enabled) since the last time the status register was cleared.

- **bits 7-4:** Sample overflow indicator for correlation baselines 0 (bit 4) to 3 (bit 7), if present. A set bit means that at least one sample from that baseline suffered overflow during filtering/decimation (hence becoming garbage) while correlation was active. Bits stay set until cleared.
 - **bit 3:** Front input lock indicator; if set, the front-panel LVDS input PLL was unlocked for more than UNLOCK_WARNING clock cycles (currently 8) while correlation was active. Stays set until cleared.
 - **bit 2:** Digitizer input lock indicator; if set, the digitizer LVDS input PLL was unlocked for more than UNLOCK_WARNING clock cycles (currently 8) while correlation was active. Stays set until cleared.
 - **bit 1:** Correlation error indicator; if set, the correlate signal went high before lag data was successfully transferred to FPGA RAM.
 - **bit 0:** Correlation done indicator; this bit is set once lag data (including metadata and quantization counts) for the most recent integration has been successfully transferred to FPGA RAM.
- CTRL_REG_SAMP_DELAY [Decimation phasor enable control]
Writing a non-zero value n to this register causes the rotating phasor used to remodulate decimated samples to be disabled for n clock cycles. This is used by the first-light correlator to ensure phase alignment. It *should* be unnecessary for revised hardware, which can synchronize automatically to the supplied 1 pps signal, but can be used if desired. Phasors exist only in digitizer FPGAs #0 and #3; all other chips ignore the contents of this register.
 - CTRL_REG_CORL_MODE [Correlator operating mode]
 - **bit 5:** When set, the output of the decimation pipeline, which normally consists of requantized (and possibly multiplexed) output samples, is replaced by a single 32-bit, full-precision output sample. This mode enables precise decimator testing.
 - **bit 4:** When set in digitizer FPGAs, the normal prompt/delay sample input streams are replaced by the LVDS front-panel readback (prompt with the 32-bit LSBs, delay with the 32-bit MSBs). In correlator FPGAs, secondary metadata is replaced with local information.
 - **bit 3:** When set, an FPGA normally calculating multiple baselines will calculate a single correlation (correlation 0) at NUM_CORL times normal resolution.
 - **bit 2:** When set in digitizer FPGAs, the decimation and sub-ns delay components are bypassed, and the (LSBs of) the input sample bus is transmitted intact. In correlator FPGAs, this bit selects between the two basic pipeline configurations required to implement baseline partitioning (cf. Figure 2).
 - **bits 1-0:**
 - * Mode "00": normal operation of the pipeline and correlation logic.
 - * Mode "01": correlation logic replaces the prompt and delay inputs with test patterns based on CTRL_REG_TEST_PIN and CTRL_REG_TEST_DIN, respectively.

- * Mode "10": normal decimator output is replaced with CTRL_REG_TEST_PIN.
- * Mode "11": in digitizer FPGAs, sub-ns delay filter output is replaced with CTRL_REG_TEST_PIN (for DIGA input) or CTRL_REG_TEST_DIN (for DIGB input); in correlator FPGAs, the test patterns replace the front-panel input data (LSBs considered ‘prompt’).

- CTRL_REG_DEMOD [Phase-switch demodulation state]
Encodes the 180 degree phase-switch demodulation sequence for up to 16 consecutive integrations (bits 15-0, with 0 being read first). A set bit indicates that samples should be negated during the corresponding integration cycle.
- CTRL_REG_TEST_PIN and CTRL_REG_TEST_DIN [Correlation sample test patterns]
The 16-bit LSBs of each register defines the test patterns fed to the prompt (PIN) and delay (DIN) cross-correlation inputs in self-test mode (see CTRL_REG_CORL_MODE).
- CTRL_REG_OUT_ENABLE [Inter-FPGA bus output enables]
 - **bit 9:** Output enable bit for bus 2E
 - **bit 8:** Output enable bit for bus 2D
 - **bit 7:** Output enable bit for bus 2C
 - **bit 6:** Output enable bit for bus 2B
 - **bit 5:** Output enable bit for bus 2A
 - **bit 4:** Output enable bit for bus 1E
 - **bit 3:** Output enable bit for bus 1D
 - **bit 2:** Output enable bit for bus 1C
 - **bit 1:** Output enable bit for bus 1B
 - **bit 0:** Output enable bit for bus 1A

Output enable bits apply only to buses actually used for output by the configuration; they are ignored by input buses.

- CTRL_REG_SAMP_GAIN and CTRL_REG_SAMP_OFFSET [Digitizer sample gain and offset]
These registers define a linear transformation applied to the raw 8-bit digitizer input samples before they are rounded to fewer (currently 6) bits prior to processing by the sub-ns delay filter:

$$x \rightarrow x' = (\text{GAIN} \times x + \text{OFFSET})/256.$$

For example, GAIN = 282, OFFSET = 768 corresponds to $x' = 1.10 \cdot x + 3$. Both factors are 16-bit quantities (the 16-bit LSBs of each register).

Table 1. Revised CARMA correlator FPGA memory map

Symbolic Name ^a	Hex Address ^b	Description
CTRL_REG_VERSION	0x00000 ^c	FPGA configuration version.
CTRL_REG_COMPAT	0x00001 ^c	FPGA configuration compatibility.
CTRL_REG_CORL_CONF1	0x00002 ^c	Correlation logic specification.
CTRL_REG_CORL_CONF2	0x00003 ^c	Correlation logic specification.
CTRL_REG_UNUSED_04	0x00004 ^c	<i>Currently unused.</i>
CTRL_REG_TD_1A	0x00005 ^c	Inter-FPGA bus 1A readback.
CTRL_REG_TD_1B	0x00006 ^c	Inter-FPGA bus 1B readback.
CTRL_REG_TD_1C	0x00007 ^c	Inter-FPGA bus 1C readback.
CTRL_REG_TD_1D	0x00008 ^c	Inter-FPGA bus 1D readback.
CTRL_REG_TD_1E	0x00009 ^c	Inter-FPGA bus 1E readback.
CTRL_REG_TD_2A	0x0000A ^c	Inter-FPGA bus 2A readback.
CTRL_REG_TD_2B	0x0000B ^c	Inter-FPGA bus 2B readback.
CTRL_REG_TD_2C	0x0000C ^c	Inter-FPGA bus 2C readback.
CTRL_REG_TD_2D	0x0000D ^c	Inter-FPGA bus 2D readback.
CTRL_REG_TD_2E	0x0000E ^c	Inter-FPGA bus 2E readback.
CTRL_REG_UNUSED_0F	0x0000F ^c	<i>Currently unused.</i>
CTRL_REG_UNUSED_10	0x00010	<i>Currently unused.</i>
CTRL_REG_UNUSED_11	0x00011	<i>Currently unused.</i>
CTRL_REG_STATUS	0x00012	Correlation status.
CTRL_REG_SAMP_DELAY	0x00013	Decimation phasor enable control.
CTRL_REG_CORL_MODE	0x00014	Correlation/pipeline operating mode.
CTRL_REG_DEMOD	0x00015	Phase-switch demodulation state.
CTRL_REG_UNUSED_16	0x00016	<i>Currently unused.</i>
CTRL_REG_TEST_PIN	0x00017	Prompt input test pattern.
CTRL_REG_TEST_DIN	0x00018	Delay input test pattern.
CTRL_REG_OUT_ENABLE	0x00019	Inter-FPGA bus output enable.
CTRL_REG_UNUSED_1A	0x0001A	<i>Currently unused.</i>
CTRL_REG_UNUSED_1B	0x0001B	<i>Currently unused.</i>
CTRL_REG_SAMP_GAIN	0x0001C	Digitizer sample gain.
CTRL_REG_SAMP_OFFSET	0x0001D	Digitizer sample offset.
CTRL_REG_UNUSED_1E	0x0001E	<i>Currently unused.</i>
CTRL_REG_UNUSED_1F	0x0001F	<i>Currently unused.</i>
MMAP_DELAY_BEGIN	0x00080	Start of delay/phase table buffers.

5. Revised VHDL Components

Channel resolution in each bandwidth mode for the COBRA-based hardware and revised correlator hardware is given in Tables 2-5. Each revised hardware table is for a specific (requantized) input sample bit width (COBRA hardware is limited to 2-bit samples). The numbers provided show what can be obtained with the correlation logic as currently implemented (July 2007); figures for the revised hardware may increase pending future refinements.

Each correlator card data FPGA calculates four cross-correlations, for a total of 16 baselines per card. For a typical resolution of 257 channels per sideband (512 lags per baseline), the average FPGA-to-CPU data transfer rate is 2 MB/s (16 x 512 x 32-bit words = 32 KB each 15.625 ms). Optimizing the correlation logic should permit twice this resolution in narrowband modes (62 MHz and below); an upper-limit on the required data transfer rate is therefore ≈ 4 MB/s.

Numerous improvements were made to the correlation logic to take advantage of the revised hardware's capabilities and to streamline code maintenance. A single VHDL component (`correlation`) now covers all correlation logic instantiations required; a generic (`CORL_TYPE`) controls the type of correlation performed (auto, cross: positive lags, cross: negative lags, or cross: all lags). This component also manages the processing of multiple baselines by a single chip via the `NUM_CORL` generic. Digitizer quantization state counters are also handled by the `correlation` component, controlled by the generic (`NUM_QCNT`). The values of these generics are available in new control registers (see § 4) to enable automated lag retrieval and processing without the need to hard-wire configuration information into external initialization files, which is error-prone and hard to maintain (particularly during testing).

The correlation logic used in the first-light configurations introduced the concept of a “meta-lag”, a word of arbitrary metadata written along with the primary lag and quantization state data during a lag dump. This facility has been significantly enhanced in the new `correlation` component. A generic (`NUM_META`) now controls the number of (32-bit) words of metadata read out during a RAM dump; in addition, a continuous stream of data is written to unused high memory while a correlation is active, providing integrated SignalTap-like functionality. The metadata as currently defined consists of the (prompt and delayed) input samples and—subject to bit-width constraints—the 16 bits per sample of metadata (such as the decimation phasor alignment bits) transported with the data by the pipeline.

To prevent latency problems with high resolutions, lag readout is organized for maximum throughput. The memory map allocates one M-RAM block per baseline, each connected by a dedicated 64-bit bus. Since readout data (lags, quantization counts, and metadata) is 32-bit (or less), the positive and negative lag/metadata streams are dumped in parallel, with the positive stream occupying the 32-bit LSBs of each 64-bit word. Readout for each baseline begins at the base of the corresponding M-RAM. Hence as seen from the external (32-bit) CPU interface, positive lags occupy even addresses and negative lags occupy the interleaving odd addresses, each followed by their metadata streams. The quantization state counts, if any, are appended to the positive lag/metadata channel. The readout logic (like all components) runs at 125 MHz; hence a 1024-channel spectrum (`NUM_LAGS=1024`) can be dumped to memory in $\approx 8.2\mu\text{s}$, comfortably within the phase-switch settling interval of $\approx 20\mu\text{s}$. The region of each M-RAM following the main readout

Table 1—Continued

Symbolic Name ^a	Hex Address ^b	Description
...	0x00080	Start of ring buffer 0, integration 0.
...	0x00080	Start of sub-ns coded tap table.
...	0x001CB	End of sub-ns coded tap table.
...	0x001CC	Whole-ns delay (16-bit).
...	0x001CD	Normalized phase offset (16-bit).
...	0x001CE	Start of ring buffer 0, integration 1.
...	0x0155F	End of ring buffer 0, integration 15.
...	0x01560	Start of ring buffer 1.
...	0x02A40	Start of ring buffer 2.
MMAP_DELAY_END	0x03F1F	End of delay/phase table buffers.
MMAP_LAGS	0x04000	Start of lag data for correlation 0.
...	0x04000	Lag 0.
...	0x04001	Lag -1.
...	0x04002	Lag 1.
...	0x04003	Lag -2.
...
MMAP_LAGS + 2 × NUM_LAGS	A_{meta}^d	Start of metadata for correlation 0.
...	$A_{meta} + 0$	Prompt input metadata word 0.
...	$A_{meta} + 1$	Delay input metadata word 0.
...
$A_{meta} + 2 \times \text{NUM_META}$	A_{qcnt}^d	Start of quantization state counters for correlation 0.
...	$A_{qcnt} + 0$	Sample count for quantization state 0x00.
...	$A_{qcnt} + 1$	<i>Cleared.</i>
...	$A_{qcnt} + 2$	Sample count for quantization state 0x01.
...	$A_{qcnt} + 3$	<i>Cleared.</i>
...
$A_{qcnt} + 2 \times \text{NUM_QCNT}$	A_{samp}^e	Start of continuous sample dump area for correlation 0.
...
$2^{1+\text{DUMP_ADDR_WIDTH}} - 1$	A_{end0}^e	End of continuous sample dump area for correlation 0.
$2 \times \text{MMAP_LAGS}$	0x08000	Start of lag data for correlation 1.
...
MMAP_FPGA_SIZE - 1	0x1FFFF	Top of local FPGA memory.

^aDefined in \$CCORL/revised/share/fpga/src/revised_components.vhd.

^bAs seen by the 32-bit external CPU memory interface.

^cThese locations are read-only.

^dCalculable using the contents of CTRL_REG_CORL_CONF1.

^dCalculable using the contents of CTRL_REG_CORL_CONF2.

Table 2. Spectral resolutions for COBRA-based correlator bands

Bandwidth (MHz)	Channels (per sideband)	δV [3 mm] (km/s)	V_{tot} [3 mm] (km/s)	δV [1 mm] (km/s)	V_{tot} [1 mm] (km/s)
500	17	94	1500	31	500
62	61	3.4	188	1.1	62.5
31	65	1.7	93.8	0.56	31.2
8	65	0.42	23.4	0.14	7.81
2	65	0.10	5.86	0.03	1.95

Table 3. Revised CARMA correlator spectral resolution [2-bit samples]

Bandwidth (MHz)	Channels (per sideband)	δV [3 mm] (km/s)	V_{tot} [3 mm] (km/s)	δV [1 mm] (km/s)	V_{tot} [1 mm] (km/s)
500	129	12	1500	3.9	500
250	193	3.9	750	1.3	250
125	289	1.3	375	0.43	125
62	385	0.49	188	0.16	62.5
31	385	0.24	93.8	0.081	31.2
8	385	0.061	23.4	0.020	7.81
2	385	0.015	5.86	0.005	1.95

Table 4. Revised CARMA correlator spectral resolution [3-bit samples]

Bandwidth (MHz)	Channels (per sideband)	δV [3 mm] (km/s)	V_{tot} [3 mm] (km/s)	δV [1 mm] (km/s)	V_{tot} [1 mm] (km/s)
500	81	19	1500	6.3	500
250	129	5.9	750	2.0	250
125	225	1.7	375	0.56	125
62	321	0.59	188	0.20	62.5
31	321	0.29	93.8	0.10	31.2
8	321	0.073	23.4	0.024	7.81
2	321	0.018	5.86	0.006	1.95

buffer is used to hold a continuous record of the incoming samples/metadata, which is useful for diagnostic purposes. The region is written (filled) only once per integration—i.e., it is *not* a circular buffer—and (depending on the lag count) can hold hundreds to thousands of samples. During normal operation its contents will not be read by the CPU, but it is always available for inspection should problems arise. The correlation logic accepts a `DUMP_ADDR_WIDTH` generic determining the number of (64-bit) locations per block it can write. Currently, the configurations write to only the first half of each block; the remaining half is unused.

6. FIR Filter Design

6.1. Digital Decimation Filters

The revised correlator hardware implements all sub-500 MHz bandwidth modes via digital decimation and band-shaping. Decimation is a multi-stage process involving a series of simple half-band filters, each followed by decimation by a factor of two. Successive filters gradually increase in complexity as the final sample rate is approached. Following decimation, a single high-precision edge-defining filter is used to produce a sharp final bandpass. For additional discussion of the method, see Rauch (2003).

All revised correlator bandwidth modes were designed to meet the following criteria:

1. Peak-to-peak bandpass ripple of 0.1 dB.
2. Peak sidelobe power below -40 dB.
3. Two edge channels aliased in excess of -40 dB (@ nominal resolution).
4. Three edge channels aliased in excess of -40 dB (@ maximum resolution).

The anti-aliasing specification implicitly depends on the channel resolution. Nominal resolution refers to the channel resolutions achieved as of July 2007; maximum resolution was defined to be 50% higher than this. The general recommendation for observers is to clip the outer three edge channels to avoid aliasing artifacts (one of these being the phaseless, half-width edge channel). The preceding figures refer to the digital decimation filters only; in the wider bandwidth modes, the analog filter in the spectral downconverter limits band flatness to 1-3 dB peak-to-peak and, for 500 MHz, edge channel anti-aliasing (-20 dB at 475 MHz and -15 dB at 480 MHz).

Figures 4-7 plot the net decimation filtering response for the revised 500, 250, 125, and 62.5 MHz bands, respectively. The spectral downconverter analog filter response is *not* included. In each case the band is centered on zero frequency (DC) when the filter is applied, and subsequently re-modulated into the positive frequency band; hence only half of each band (and residual out-of-band artifacts) is shown. Filter response is symmetric about zero frequency. The heavy (green) line is the signal level and the thinner (blue) line is the aliased noise level. Vertical lines denote the location of spectral channel boundaries for the mode. The number of filter taps used increases as the bandwidth decreases (and number of channels increases; cf.

Tables 3-5), in such a way that only ~ 2 edge channels suffer noticeable out-of-band aliasing, regardless of the absolute channel resolution. The corresponding plots for the narrowband modes (31, 8, and 2 MHz) are very similar to Figure 7 as the same edge-defining filter is used in all four modes.

6.2. Fractional Sample Delay Filters

Figure 8 displays the amplitude and phase performance of the set of sub-ns (fractional sample) delay filters used in the digitizer cards to align antenna input signals. The dashed vertical lines indicate the limiting “usable” analog bandwidth, defined here as the frequency beyond which out-of-band aliasing exceeds -18 dB. Filter performance depends on the delay (each delay corresponds to a unique set of filter coefficients). The dashed lines denote the worst-case amplitude and phase variations over all possible delays; the solid curve plots the response for one particular delay. Amplitude response is flat to within 0.1 dB over the lower 94% of the 500 MHz band, and phase delay accuracy meets the nominal CARMA specification over the lower 97% of the 500 MHz band. All remaining bandwidth modes—carved from the center of the 500 MHz band—exceed the desired accuracy over 100% of the band. The width of the region near Nyquist exhibiting degraded delay performance scales inversely with the filter length; logic usage increases linearly with the filter size. The implemented filter (discussed in detail below) balances the two competing constraints.

Delay filter coefficients are computed using a windowed sinc function method, which has a closed-form solution for each coefficient. Given a fractional delay δ , $0 \leq \delta < 1$ and a filter containing N coefficients, the formula for coefficient c_i , $i \in \{0, \dots, N-1\}$ is

$$c_i = W(i - \delta, N) \operatorname{sinc}[\pi(1 - \varepsilon)(i - D)],$$

where W is the windowing function, ε is a fixed parameter chosen to minimize delay errors, and $D = \delta + N/2$ ($\delta < 1/2$) or $D = \delta + N/2 - 1$ ($\delta \geq 1/2$) is the effective filter delay. The filters developed for the CARMA digitizers employ a Hann window function, $W(x, N) = [1 - \cos\{2\pi x/(N-1)\}]/2$, with $N = 86$ and $\varepsilon = 0.01165$. The coefficients are quantized to 15-bit precision for use by the digitizer FPGA sub-ns delay filter component (`frac_delay`). Quantization reduces the net number of coefficients to a maximum of $N = 80$, as several edge coefficients underflow to zero. An optimized subroutine to compute CARMA digitizer delay coefficients, `fd_carma_subns_coef()`, can be found in `$$CORL/share/fpga/test/frac_delay.c`. See Laakso, Valimaki, & Karjalainen (1996) for a review of fractional delay filter design methods.

The Altera FIR filter component requires that reloadable filter coefficients be stored in (private) RAM blocks. The contents of this tap RAM must be loaded serially into the component using a hard-wired interface; the RAM is not directly accessible. Updating delay coefficients in the digitizer FPGAs is done by sequentially reading an area in the main (PPC-visible) FPGA memory map, treated as a circular buffer. The address limits for this buffer are `MMAP_DELAY_BEGIN` and `MMAP_DELAY_END` (cf. Table 1). The output of this buffer is then fed into the appropriate filter components. The Altera FIR components can only receive one input sample and produce one output sample per clock. The digitization logic outputs demux-by-8 @ 125 MHz samples (say x_0 to x_7); hence in practice the sub-ns delay component needs to instantiate a vector of 8 identical filters, each accepting 8 input samples and producing one output sample per clock (y_0 to y_7).




Fig. 4.— Revised CARMA correlator decimation filter performance for the 500 MHz band.




Fig. 5.— Revised CARMA correlator decimation filter performance for the 250 MHz band.




Fig. 6.— Revised CARMA correlator decimation filter performance for the 125 MHz band.




Fig. 7.— Revised CARMA correlator decimation filter performance for the 62 MHz band.

The input samples fed to each filter are skewed by one clock cycle between filters, so that the phase of each output sample matches that of its youngest input sample. Specifically, filter $i \in \{0, \dots, 7\}$ receives input samples x_{i-7} to x_i ; implementing this requires caching the 7 youngest input samples from the previous clock cycle.

Since each filter must consume 8 input samples per clock, these filters themselves need to be split into 8 sub-filters, each containing 1/8 of the filter coefficients, whose individual outputs are summed to produce a single output sample. To fully utilize logic resources, the number of filter taps should therefore be a multiple of 8. For the CARMA digitizers $N = 80$ and each sub-filter contains 10 taps. In terms of the taps $c_i, i \in \{0, \dots, N - 1\}$ of the original filter, the taps $a_j^k, j \in \{0, \dots, (N/8) - 1\}$ of sub-filter $k, k \in \{0, \dots, 7\}$, are $a_j^k = c_{8j+k}$; the input to filter k is x_k .

Altera FIR filters with reloadable taps do not accept the actual coefficients as input, but rather a coded data stream produced by their `coef_seq` command line utility. A copy of the C++ source code for `coef_seq` can be found in `$$CCORL/share/fpga/test`. The format of the coded data is not specified by Altera; empirically, for a filter with $N \times m$ -bit taps, it consists of $2^N \times (m + 3)$ -bit values (the first of which is always zero) representing some kind of pre-computed multiplication table. The critical thing to note is that the amount of coded data increases exponentially with the number of taps and linearly with the bit-width of the input (for a fully parallel filter, as required here). A Stratix II M512 block has a capacity of 32×18 -bit elements, precisely enough to hold the coded data for a filter with 5×15 -bit taps. The EP2S90 FPGAs loaded into digitizer cards contain 488 M512 blocks and 408 M4K blocks (with 256×18 -bit capacity). To reduce RAM usage to an acceptable level, the 10-tap sub-filters are further sub-divided into two 5-tap filters. In total the sub-ns delay filter instantiates 128 divided sub-filters, requiring 128 RAM blocks per bit of input. The implemented component accepts 6-bit input samples—in practice, the number of meaningful bits expected from the digitizers—and consumes 480 M512 blocks and 288 M4K blocks (as the latter are poorly utilized in this context, their use should be minimized). Output samples are rounded to 8-bit precision.

The coded coefficient stream for a single delay filter therefore consists of $(N/5) \times (2^5 - 1) = 496 \times 18$ -bit values (the leading zero in the coded streams are inserted on-the-fly by the filter reload state machine). Since the FPGA memory map is 64-bit wide internally, it is convenient to group values into triplets occupying the lower $3 \times 18 = 54$ bits of each quadword of the circular buffer in FPGA RAM. This nominally amounts to 166 quadwords of FPGA RAM per delay coefficient set; however, an extra quadword containing the whole-ns delay value (32-bit LSBs) and normalized phase-offset (32-bit MSBs) is appended to each coefficient set. Hence each phase/delay update occupies 167 quadwords (334 longwords) of FPGA RAM. The phase/delay update rate is 64 Hz (1/15.625 ms), which implies an average data transmission bandwidth of ~ 83.5 KB/s to each of the two digitizer FPGAs receiving this information (FPGA #1 and #2); since the CPU bus is shared, the effective load is ≈ 167 KB/s, a very small fraction of the available bandwidth. The memory map allocates one Stratix II M-RAM block (8192 quadwords) for the delay buffer, enough to hold 750 ms of phase/delay update data ($48 \text{ sets} \times 167 = 8016 \text{ quads}$).

The precise contents of the delay/phase coefficient table are as follows. The `frac_delay` component receives two coded coefficient streams in parallel, one for each sub-channel (5-tap) filter of the pair which to-

gether implement a particular 10-tap sub-filter, as described above. To support parallel readout, the delay table therefore interleaves coefficients from two individual filters. Let $a_l^{k,0} = a_l^k$ and $a_l^{k,1} = a_{l+5}^k$, $l \in \{0, \dots, 4\}$, be the coefficients for the two sub-channel filters implementing sub-filter k , and denote the corresponding `coef_seq`-coded tap streams by $s_m^{k,0}$ and $s_m^{k,1}$, $m \in \{0, \dots, 31\}$, where $s_0^{k,0} = s_0^{k,1} = 0$. One delay/phase table consists of 167 quadwords; except for the final quad (containing the whole-ns delay and phase offset), each quadword contains three 18-bit coded taps, packed contiguously into the LSBs. Let $T_i = \{s_2, s_1, s_0\}$ represent table quadword i , where s_m are the coded tap values (in MSB to LSB order), and unused MSBs are cleared (though ignored by the FPGA). FPGA RAM addresses increase with i , and quadwords are stored in little-endian order (RAM addresses refer to 32-bit words). Then

$$\begin{aligned}
 T_0 &= \{s_2^{0,0}, s_1^{0,1}, s_1^{0,0}\} \\
 T_1 &= \{s_3^{0,1}, s_3^{0,0}, s_2^{0,1}\} \\
 &\vdots \\
 T_{20} &= \{s_1^{1,0}, s_{31}^{0,1}, s_{31}^{0,0}\} \\
 T_{21} &= \{s_2^{1,1}, s_2^{1,0}, s_1^{1,1}\} \\
 &\vdots \\
 T_{164} &= \{s_{31}^{7,0}, s_{30}^{7,1}, s_{30}^{7,0}\} \\
 T_{165} &= \{0, 0, s_{31}^{7,1}\} \\
 T_{166} &= \{\Phi, \Delta_{\text{ns}}\},
 \end{aligned}$$

where $\Phi = \lfloor 2^{16} \cdot [\phi / (2\pi)] \rfloor$ is the 16-bit normalized phase offset for absolute phase offset ϕ , Δ_{ns} is the 16-bit whole-ns delay value, and both Φ and Δ_{ns} are zero-extended to 32-bit values. Note that the null ($s_0^{k,x} = 0$) coded taps are not stored in the table, but inserted on-the-fly during the reload process.

7. Phase Offset Correction

In contrast to the first-light CARMA correlator, the revised correlator applies phase corrections continuously to incoming samples as part of the signal decimation process; the corrections remove antenna-based phase offsets arising from bulk downconversion and lobe-rotation differentials. As described in Rauch (2003), downconversion entails shifting the band center from half the Nyquist frequency (250 MHz) to DC by multiplying the input samples $\{x_k\}$ by $e^{-i\pi f_0 k}$, where $f_0 = -1/2$. The resulting frequency modulated samples are $\{x_0, +ix_1, -x_2, -ix_3, x_4, \dots\}$. Phase correction in the revised correlator is implemented by further multiplying the modulated samples by $e^{i\phi}$, where ϕ is the phase offset. The corresponding phase-corrected, frequency modulated samples are $\{x_0 \cos \phi + ix_0 \sin \phi, -x_1 \sin \phi + ix_1 \cos \phi, -x_2 \cos \phi - ix_2 \sin \phi, x_3 \sin \phi - ix_3 \cos \phi, x_4 \cos \phi + ix_4 \sin \phi, \dots\}$.

A new offset ϕ is applied each integration by the decimation logic in digitizer FPGA #0; the value of ϕ to use is driven by FGPA #1 onto the "E" data bus connecting FPGA #1 to #0 (and similarly by FPGA #2 for FPGA #3) and remains static for the duration of each integration. The values are stored in FPGA RAM as




Fig. 8.— Revised CARMA correlator fractional sample delay filter performance. The top and bottom windows display frequency and delay responses, respectively. CARMA-defined per-antenna and per-baseline limits are also indicated.

part of the delay/phase tables (see Table 1); each time the phase/delay reload occurs between integrations, a new value of ϕ appears on the bus. Applying the offset requires calculation of the $\cos \phi$ and $\sin \phi$ factors. An Altera numerically controlled oscillator (NCO) megacore component is instantiated for this purpose by the decimation logic. The NCO component accepts scaled output frequency and phase parameters, ϕ_{INC} and ϕ_{PM} respectively, where the physical output frequency $f_o = \phi_{\text{INC}} f_{\text{clk}} / 2^M$ and phase $\phi = 2\pi \phi_{\text{PM}} / 2^P$. Here f_{clk} is the actual frequency of the NCO input clock, M is the internal accumulator precision, and P is the angular precision. The NCO component outputs two waveforms in two's-complement format, $2^{N-1} \sin(2\pi f_o t + \phi)$ and $2^{N-1} \cos(2\pi f_o t + \phi)$, where N is the magnitude precision. The parameters M , N , and P are fixed for each NCO instantiation; the ϕ_{INC} and ϕ_{PM} inputs are fully dynamic (can change each clock cycle). In the present case the NCO does not actually oscillate ($\phi_{\text{INC}} = 0$) as we are only interested in computing the (quasi-)static values $\sin \phi$ and $\cos \phi$. The CARMA implementation uses $N = 18$, $M = 20$, $P = 16$, and rounds the resulting phase-corrected, frequency modulated samples to 12-bit precision. The choice $P = 16$ corresponds to a phase resolution of $360/2^{16} = 5.5 \times 10^{-3}$ deg.

REFERENCES

- Laakso, T.I., Valimaki, V., and Karjalainen, M. 1996, IEEE Signal Processing Magazine, vol. 13, no. 1, pp. 30-60.
- Rauch, K.P. 2003, CARMA Memo 12.

Table 5. Revised CARMA correlator spectral resolution [4-bit samples]

Bandwidth (MHz)	Channels (per sideband)	δV [3 mm] (km/s)	V_{tot} [3 mm] (km/s)	δV [1 mm] (km/s)	V_{tot} [1 mm] (km/s)
500	33	47	1500	16	500
250	49	16	750	5.2	250
125	97	3.9	375	1.3	125
62	161	1.2	188	0.39	62.5
31	161	0.59	93.8	0.20	31.2
8	161	0.15	23.4	0.049	7.81
2	161	0.037	5.86	0.012	1.95