EVLA 'WIDAR' Correlator Description for the Preliminary Design Review

NRC-EVLA Memo# 024

Brent Carlson, June 17, 2005

ABSTRACT

A simplified description of the EVLA 'WIDAR' correlator hardware and software is presented for PDR reviewers and potential users of the system. This description is a compendium of the work of many people over several years.

Introduction

Several years of detailed design and development work on the EVLA correlator have been done. The design is now ready for fabrication and testing of first prototypes and there is considerable documentation that describes many aspects of the system. However, there is no recent document that describes the entire correlator system in a simplified and easy-to-understand way, with enough detail so that the full scope of the system can be understood. This memo aims to fill this void by describing the correlator's hardware and software in simplified, yet complete enough terms so that readers can gain a more detailed understanding of the system.

Signal Processing

WIDAR hybrid- type correlator. Sub-band correlations. Poly-phase FIR.

Detailed descriptions and investigations of the signal processing in the correlator are contained in the appendices of [1] . Essential elements of correlator signal processing are repeated here for completeness. We have named this patented[1] signal processing technique "WIDAR"—Wideband Interferometric Digital ARchitecture. A simplified block diagram of correlator signal processing is shown in Figure 1. This is a hybrid-type XF correlator whereby the wideband is split into sub-bands using digital filters, the sub-bands are separately correlated, and then the individual sub-band correlations are seamlessly[2] stitched together to yield the wideband cross-power spectrum. Although the filtering/sub-banding operation could be performed with a single poly-phase FIR/FFT filter bank, in this correlator a separate filter is provided for each sub-band to meet science requirements for targeting spectral regions of interest, while optimally using correlator resources.

Epsilon freq offset in antenna, remove in correlator. Aliasing washes out. Spurious interference washes out.

A key component of the technique is that the signal is offset in frequency at each antenna by a small, different epsilon, and this epsilon is removed in the cross-correlator. By doing so, sub-band transition band aliasing is greatly attenuated, concomitant with the

---

[1] U.S. patent application 09/936,819 "Parallel Correlator Architecture".

[2] A factor of 2 (tapering off with the steepness of the sub-band filter) sensitivity loss at the sub-band boundary is incurred.

integration time. This technique also has the effect of causing spurious interference and biases introduced after the epsilon frequency shift to be attenuated as well—an important component to generally improving data integrity, particularly in some of the more exotic integration modes of the correlator where 180^o phase switching would be impractical. Other methods are possible to reduce the effects of the aliasing, however this is believed to be the most cost-effective way to remove the aliasing **and** provide more than 60 dB of spectral dynamic range. Provided the initial sampler (A/D converter) is not saturating and generating harmonics or inter-modulation products, the spectral dynamic range is only limited by the reject-band of the digital filter.

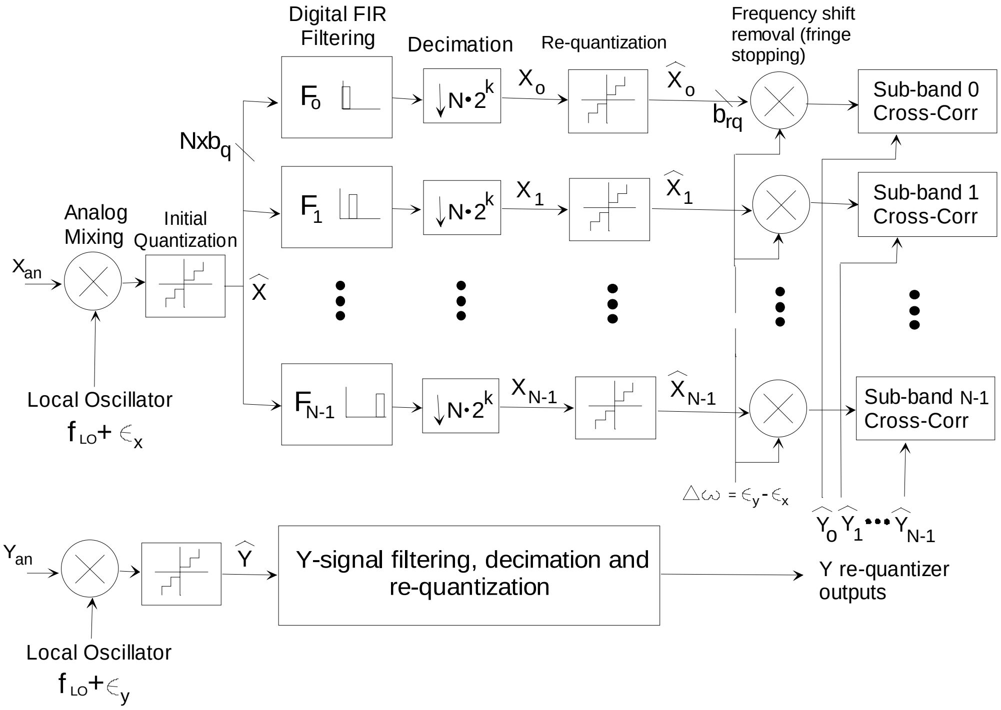

Figure 1 EVLA "WIDAR" correlator simplified signal processing.

Pre-re-quantizer power measurement for stitching.

A key measurement that allows the sub-bands to be seamlessly stitched together is the power in the sub-band after filtering, but before re-quantization. This power measurement is necessary so that the relative amplitude of each sub-band correlation can be re-established in the stitching operation. Additionally, pre-re-quantization power measurement of a sub-band that contains no interference allows the majority of the amplitude modulating effects of time-variable interference to be factored out, thus producing a wideband output ready for post-correlation interference excision or cancellation.

Digital sub-sample delay tracking.

+/-1/32nd of a sample sub-sample delay tracking is employed by controlling the phase of each sub-band's fringe stopper as shown in Figure 2, thus eliminating the need, complexity, and uncertainty of controlling the phase of the initial sampler clock to provide the same functionality.

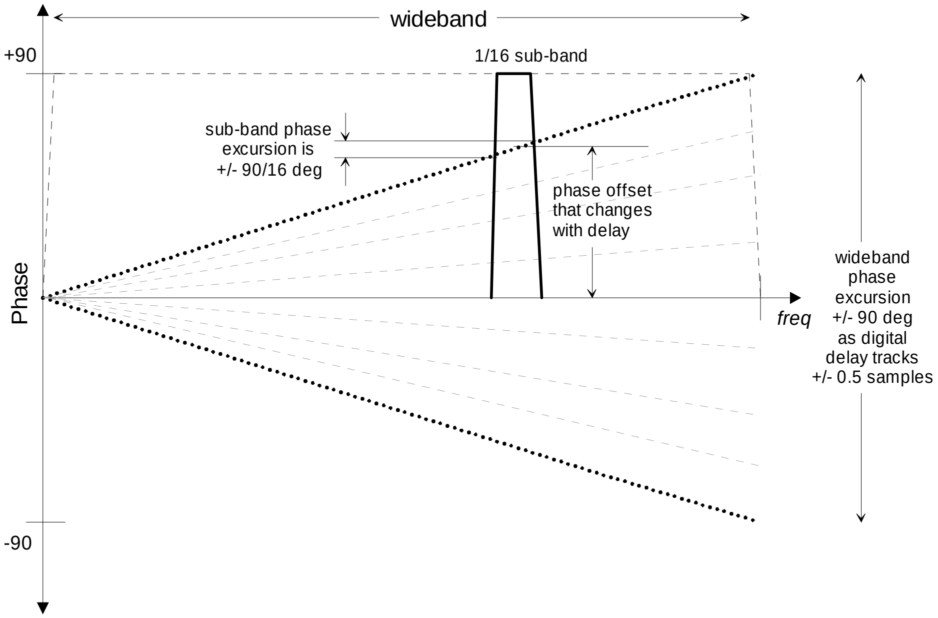

Figure 2 Sub-sample delay tracking in the correlator.

Correlator Architecture Overview

The correlator contains a number of functions to meet the requirements defined in [2]. A simplified signal flow diagram of the correlator is shown in Figure 3. Refer to this figure for the following descriptions.

Fiber-optic receiver module. Bandwidths, 16x4 data highways.

The initial samplers (A/D converters) are in the antennas, and the sampled signals arrive at the correlator via fiber. Embedded in the signaling in the fiber are 1 PPS and 0.1 PPS (10 second) time ticks that indicate antenna time epochs. Four correlator Station Boards are required to process all of the signals from each EVLA antenna, and each board processes two, 2 GHz sampled basebands. There are 3 fibers into each Station Board's **Fiber-Optic Receiver Module** (**FORM**—a.k.a. DTS Receiver), a mezzanine card developed by NRAO. The signals out of the FORM are in parallel electrical form—two sets of de-multiplexed 16x4-bit data highways[3] at 256 Ms/s/stream. These data highways exist on this interface and in the correlator, but the EVLA utilizes only 3-bit data per de-multiplexed stream. In 8-bit, 1 GHz mode, both data highways exist to service the one sampled and de-multiplexed stream. These highways go into a **cross-bar switch**. Normally, data passes through the cross-bar switch, but it is possible to feed the same data into each of the baseband paths if desired.

Delay module, 0.25 sec delay buffer.

The **Geometric Delay Tracking** block is implemented in a mezzanine card called the Delay Module. This module performs +/-0.5 sample delay tracking at 4 Gs/s on the de-multiplexed data. The modules use dynamic RAM and thus are able to have a deep delay of 1 Gsample, or 0.25 seconds of real-time delay for a reasonable price. Note that other configurations of geometric delay tracking to support other baseband bandwidths are possible with or without a different mezzanine card as discussed later.

---

[3] Previously and equivalently defined as four 32-bit data paths.

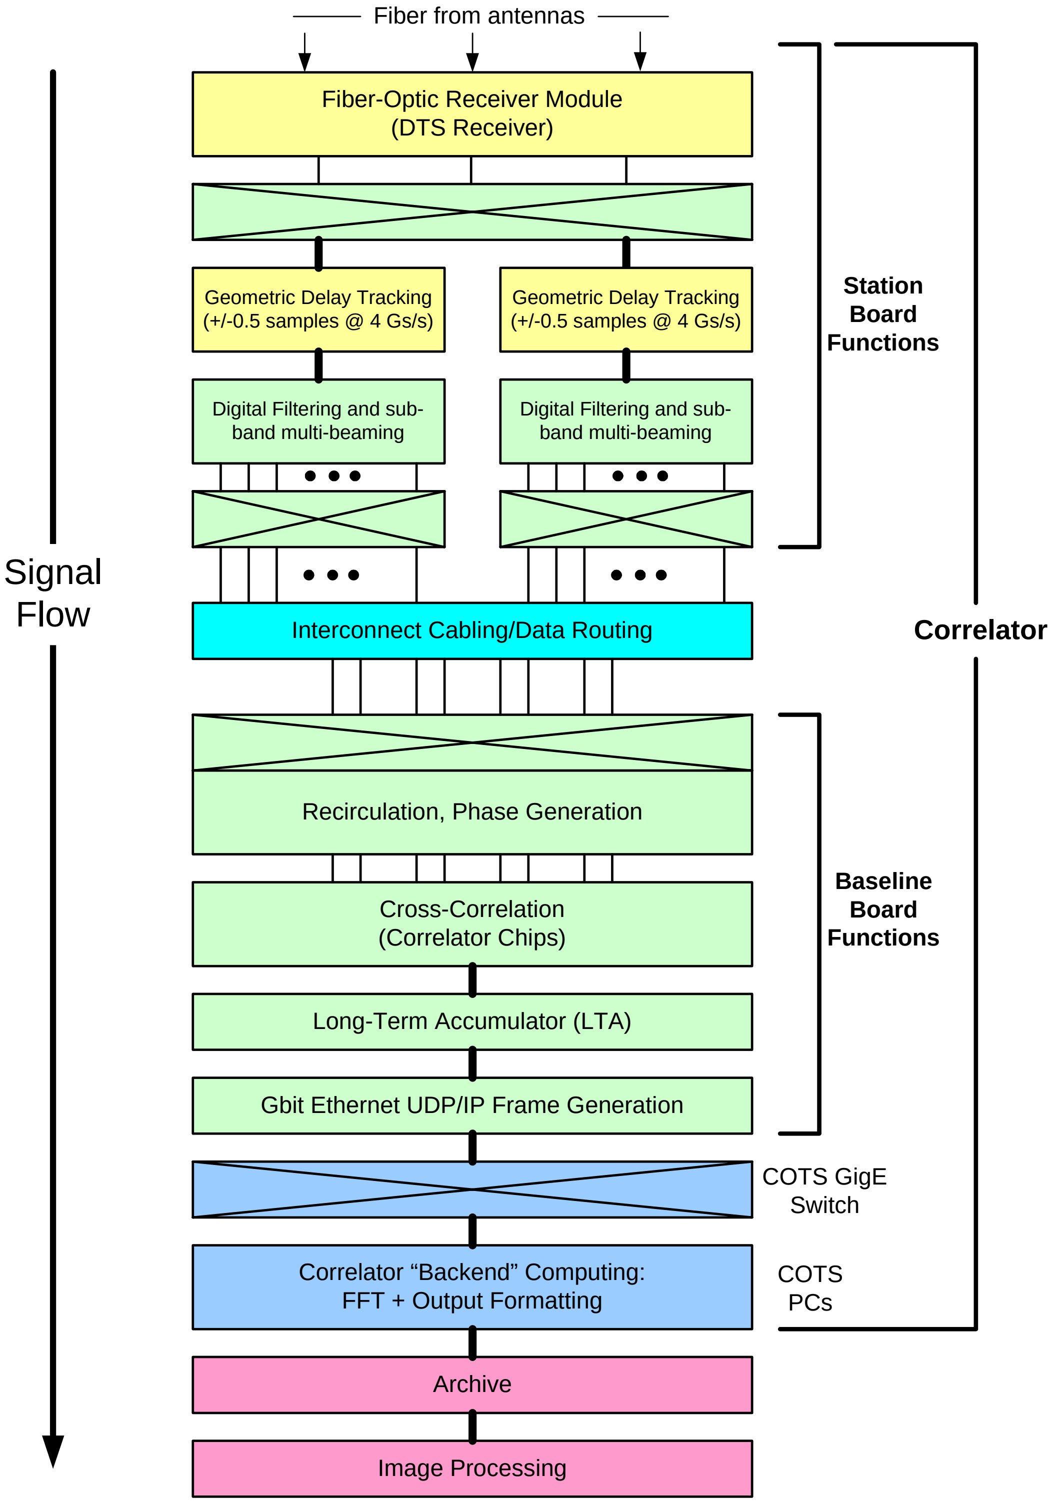

Figure 3 Simplified signal flow diagram for the correlator.

Digital filter banks, sub-ands, re-quantization.

After delay tracking, the data enters the **Digital Filtering and sub-band multi-beaming** blocks. There are two filter banks—one for each BB, and within each bank there are 18 filters. Normally, 16 filters are used for sub-band generation, one is used for noise diode power detection on a part of the band where there is no interference (i.e. part of the strategy for immunity to the effects of time-variable interference[1]), and one spare filter for the possibility of redundant correlation. If the sub-band is re-quantized to 4 bits, then the full bandwidth of 128 MHz per sub-band is available. If the sub-band is re-quantized to 7 bits, then the maximum sub-band bandwidth is 64 MHz[4]. The number of re-quantization bits, and the number of initial sampler (A/D converter) bits are independent of each other. Each digital filter contains a multi-beam delay memory that allows the interferometer beam delay center generated by the filter to be offset from the main interferometer beam delay center set by the **Geometric Delay Tracking** block (Delay Module). This can be useful for wide-field imaging and phase referencing. The allowable sub-band beam offset depends on the baseline and the available memory. Further details on digital filter capability are discussed later.

Sub-band cross-bar switch, wideband high spectral resolution.

After digital filtering, the sampled (4 or 7-bit) 18 sub-bands are fed into a full **cross-bar switch**. The cross-bar switch is needed for correlator configurations that require feeding the same data to multiple correlator modules for wideband high spectral resolution configurations, whereby different lags are acquired on different correlator boards for the same cross-correlation. The switch is also useful for system-wide off-line (and on-line) test functions.

Station-to-Baseline signal distribution at 1.024 Gbps.

The outputs of the cross-bar switch are fed to the **Baseline Board Functions** via the **Interconnect Cabling/Data Routing** block. Signals traveling on these cables have been multiplexed up by a factor of 4 to 1.024 Gbps, using source-synchronous signaling[5], and *without* 8B/10B encoding. Routing of multiplexed high-rate signals reduces the cabling requirements and cost by nearly a factor of 4, over routing low-rate signals. Fiber connections are not used for cost reasons and because the cable runs are less than ~10 m. (If fiber were required for some non-EVLA correlator configuration, it could be retro-fitted into the design with a separate pluggable module.) There is considerable flexibility for data routing via cabling in this block, and thus various non-EVLA correlator configurations can be supported, as discussed later.

1.024 Gbps de-mux, receiver lock, dynamic phase alignment.

The multiplexed 1.024 Gbps sub-band signals, along with timing, dump control, phase models, and sub-sample delay tracking models enter the **Baseline Board Functions** block, and are de-multiplexed back to their original rate of 256 Ms/s. Proper signal lock/alignment is achieved with embedded signals in the streams, with the "dynamic phase alignment" capability of the Altera receiver FPGA, and with automatic lock acquisition logic built into the FPGA. Any errors (~96% detection) at any time on these links are immediately detected by the receiver FPGA and flagged by software.

---

[4] It is possible, but not currently implemented as such, to have a 7-bit sub-band at 128 MHz. This was not done for simplicity and uniformity in data routing.

[5] A low frequency (128 MHz) clock is sent with the data and is used for clock generation at the receiver, rather than locking onto the data with a PLL to generate the clock.

Cross-bar switch, recirculation, phase generation.

Once the signals are de-multiplexed back to their original rate of 256 Ms/s[6] and properly lined up with each other, they enter a full **cross-bar** switch for routing to the correlator chips with or without recirculation. Phase models from the Station Boards are captured as well, and phase (including phase for sub-sample delay tracking) is generated and travels to the correlator chips with the data for final fringe rotation and epsilon frequency shift removal. Recirculation functions are entirely handled by the Altera receiver FPGA, and the correlator chip only passes on information to the LTA for proper binning of recirculation data, thus simplifying correlator chip design and test requirements.

Cross-correlation, correlator chip—16x128-lags.

The correlator chips are where cross-correlation and fringe stopping/epsilon frequency shift removal occurs. The chip contains input buffering that automatically and transparently aligns the data and clocks from the two different sources (“X” and “Y”), so that relative cable-length mismatch and relative clock skew is completely removed before entering the core functions of the chip. The correlator chip contains sixteen 128-complex-lag cross-correlators that can be configured in a number of different ways, including concatenation of lags for increased spectral resolution with decreased number of cross-correlations. There is some data path switching in the correlator chip and this capability, along with the cross-bar switch before recirculation (in the **Baseline Board Functions**), allows a plethora of cross-correlation configurations to be realized. For some examples, refer to [3].

Corr chip frames dump into LTA.

When dumps occur, controlled completely by hardware signaling from the **Station Board Functions**, an independent frame is generated for each 128-complex-lag correlator. The frame contains all of the information required for the **Long-Term Accumulator (LTA)**, to decide what to do with it. This can be saving the frame into the correct LTA bin along with existing data, or telling the LTA save it and then flag the bin as ready for readout. Each correlator chip has its own LTA FPGA and storage RAM—needed for high performance recirculation, pulsar phase binning, and high-speed snapshot dumping.

LTA to Gbit Ethernet chip to Backend for FFT to Archive, then Image Processing.

Once a bin is ready for readout, the LTA readout function, under command of the **Gbit Ethernet UDP/IP Frame Generation** block, transmits the data to the Gbit Ethernet block for UDP/IP frame generation and transmission on Gbit/sec Ethernet, via a **COTS GigE Switch** to the correlator **Backend** computers for further processing. The correlator **Backend** computers (computer cluster of Linux PCs) assemble the frames (“lag frames”) according to information contained in the frames and configuration information from the host computer. Finally, FFTs to the frequency domain are performed and the data is saved to the **Archive**. **Image Processing** computers access the Archive for image processing and analysis. Note that the Archive and Image Processing blocks are outside the scope of the correlator and won’t be discussed further. Initially the Backend will generate VLBI-style UVFITS, and will eventually be able to store the data in a standard format able to support the operating modes of the correlator once such a standard format is defined.

---

[6] The DATA lines are 256 Ms/s—any control lines are 128 Ms/s.

Correlator System Timing Overview

Correlator timing locked to array timing, UTC.

A simplified system timing block diagram is shown in Figure 4. Correlator timing is intimately linked to the EVLA array timing, and the source of clocks and time is the **Array Time Reference** that generates UTC time locked to a hydrogen maser time standard. There is provision for operating part of the correlator in non-real-time VLBI mode within this timing system.

Array time reference to correlator: 1 PPS tick and 128 MHz CW. Distribution to Station Boards.

A reference clock and EXTERNAL TIMECODE is provided to the correlator from the **Array Time Reference**[7]. This interface is described in detail in [4], and consists of a 128 MHz, 0 dBm sine-wave clock, and a formatted 1 PPS time tick. The **Timecode Generator Board**[8] (TGB), of which there are two parallel operating units to eliminate single point of failure, lock onto this signal and, with a time setting via Ethernet, generate the correlator TIMECODE(s) and a 64 MHz digital clock that get distributed via **Fanout Boards** to correlator **Station Boards**. One Fanout Board near the TGBs multiplies the signal to distribute it to station racks, and within each station rack there is one Fanout Board to distribute the signal to Station Boards within the rack. Cabling and fanout of these signals is virtually identical to cabling and fanout of station-to-baseline signals in the correlator, although the TIMECODE data rate is much lower.

3 TIMECODES on each TGB.

On each TGB, three TIMECODEs are generated. One TIMECODE is set to real-time UTC for main real-time correlator operation, and two TIMECODEs can be set to non-real-time values for simultaneous non-real-time VLBI correlation. Since each TGB generates three TIMECODEs, there are 6 TIMECODEs available at each Station Board.

TGB delay to match Delay Module buffer delay.

The TIMECODE generator logic on each TGB contains a delay that offsets the output TIMECODE relative to the input EXTERNAL TIMECODE by one-half the **Geometric Delay Tracking** buffer depth, so that once the time ticks from the antennas travel through the delay buffer (Delay Module), they come out at approximately the same time as the TIMECODE tick—the differences being due to fiber transport delay, and imposed geometric delay. With logic on the Station Board, it is thus possible to determine exactly what the time difference between the antenna time and the correlator reference time is, and offset it according to a delay model to achieve geometric delay compensation[9].

TIMECODE: 10 msec, 1PPS ticks, real-time UTC.

The TIMECODE signal contains 10 msec and 1 second time ticks as well as complete real-time UTC information, and this information is used for synchronization, timing, and timestamping of events within the correlator. Details of the content and protocol of the TIMECODE signal is contained in [5].

TIMECODE selection on Station Board.

Station Board logic, under CPU control, selects which of the 6 TIMECODEs are to be used for its internal timing. If the CPU detects that the TIMECODE signal has disappeared (indicating a TGB failure), it can select the TIMECODE from the other TGB

---

[7] The implementation and location of this time reference and connection to EVLA antennas is beyond the scope of this document.

[8] Defined in some documentation as the “Timecode Generator Box”.

[9] Normally the delay through the Delay Module is provided directly by reading registers on the Delay Module.

with only a momentary hick-up in system timing, thus removing the single point of failure of one TGB. Additionally, if the TIMECODE fails completely, the PLL on the Station Board fly-wheels to maintain clock distribution in the system.

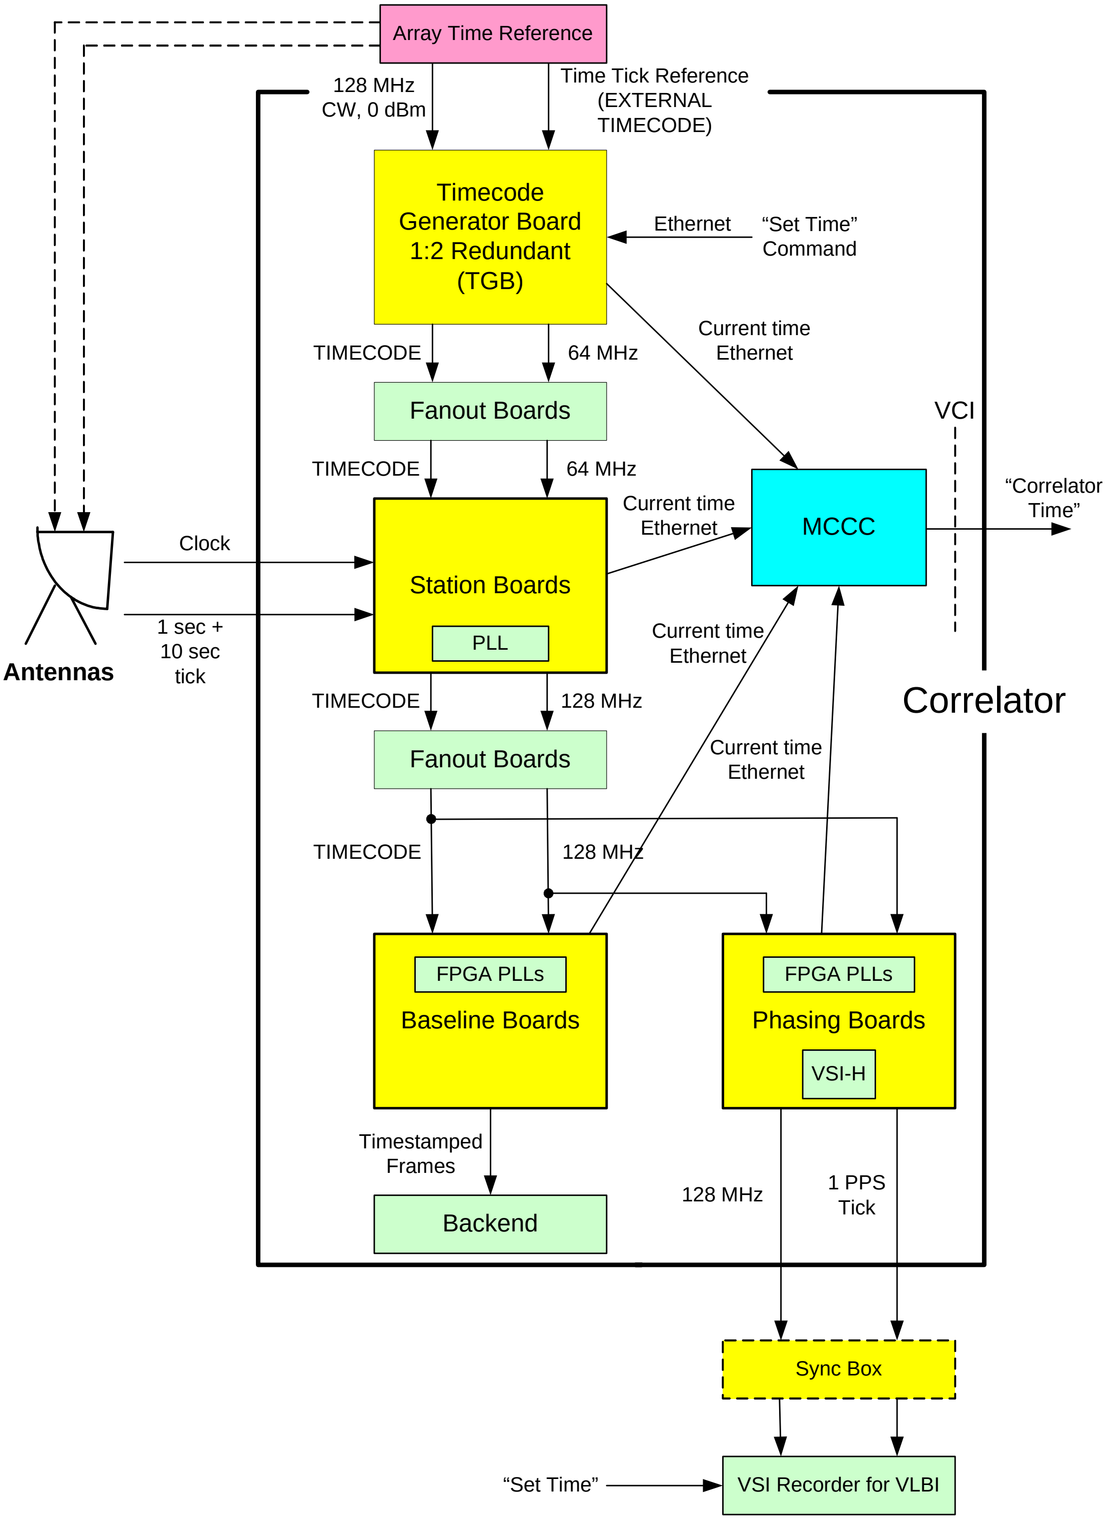

Figure 4 EVLA correlator system timing simplified block diagram.

Station Board "quad". PLLs on boards for high-rate clocks. TIMECODE regeneration on Station Board.

Each Station Board within a "quad" of four Station Boards must choose the same TIMECODE because of restrictions on the Baseline Board synchronization and timestamping logic. Each Station Board contains a high-performance crystal PLL that locks onto the 64 MHz signal to produce a pure, virtually jitter-free, 128 MHz clock for internal use and for transport to downstream Baseline Boards and Phasing Boards. High-rate clocks (256 MHz and in some cases 1.024 GHz) are only ever developed with PLLs or DLLs (Delay Locked Loops) on FPGAs. The Station Board also re-generates the selected TIMECODE for transport and use on the Baseline and Phasing Boards. If the TGB clock ever disappears, the PLL maintains generation of a clock so that a system power discontinuity will not occur.

TIMECODE use on Baseline Board. Timestamping VSI output.

On Baseline Boards, receiver FPGAs lock onto the 128 MHz clock and TIMECODE for internal use and timestamping of correlator chip data frames that eventually find their way to the correlator Backend computers and the Archive. On Phasing Boards, the same thing happens except that output data on the VSI-H interface is accompanied by the 128 MHz clock and time tick for synchronization to other Phasing Board outputs and eventually for recording by a VSI-H recording device.

Phasing data feedback to Station Board.

Not shown in Figure 4, but explained later, is the potential ability to feedback data from the Phasing Board to the Station Board via 1 Gbps links without any external hardware. This would find use when it is desired to use the correlator to autocorrelate phased-EVLA data.

Embedded CPU 10 msec interrupts.

Embedded processors on each TGB, Station Board, Baseline Board, and Phasing Board, have access on every 10 msec interrupt to the captured TIMECODE signal, and can thus report this time to the controlling **MCCC** (Master Correlator Control Computer), if so requested.

Correlator Data Routing and Connectivity Overview

Wide range of operating modes, and correlator configs.

There are many physical and software configurable settings in the correlator that permit a wide range of connections to support all of the operating modes of the EVLA correlator. Many other correlator configurations can also be built to support a wide range of antenna array and baseband configurations using the basic building blocks of the correlator, although not all of these configurations can or will be discussed here.

Figure 5 is a simplified connectivity diagram of the EVLA correlator. The cross-bar switches on the Station and Baseline Boards with the cable routing paths shown, provide significant flexibility for a plethora of configurations. Refer to [3], [6], and [7] for historical background on the evolution of this architecture.

Indep. Tunable sub-bands, cross-bar switching, output wafer.

Each **Filter Bank** on the Station Board produces 18 independently tunable[10] sub-bands "**s**" which are fed to cross-bar switches with 18 outputs "**k**". Each **k**th output can be connected to any sub-band **s**, and **k**th *pairs* are connected to a physical output **wafer**, via an output connector, along with a clock and timing and control signals.

Wafer phys/elec definition.

A wafer on a board's right angle connector is a row of 10 pins—8 inner pins for 4 differential signal pairs, and 2 outer pins for cable connector ground. The cable connector is a slim plastic HM (Hard Metric) 2.0 mm 1x8 connector with the outside ground contacts built to contact and mate with the outer 2 pins of the 10 pin board connector. Cable wafer connectors can be physically attached to each other in groups, or can be separate and form groups once plugged into the mid-plane (i.e. "Feedthru Backplane") connector that mates with the board's right angle connector when the board is inserted into the sub-rack slot. A picture of a cable wafer stack is shown below.

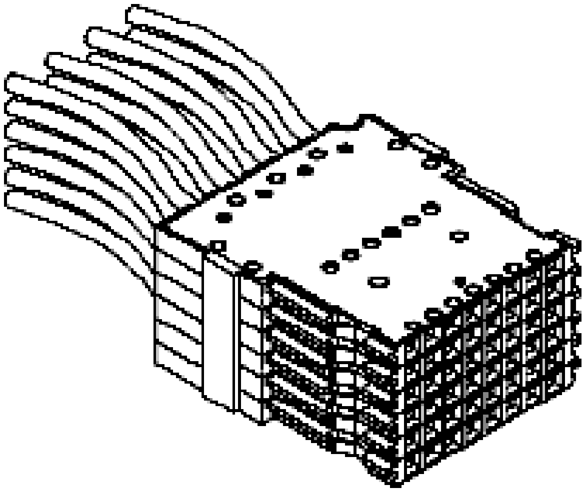

4-wafer bundle from 4 Station Boards are correlated.

Each wafer contains a pair of sub-bands, each one from one **BB** (BaseBand) input to the Station Board. Internally, and perhaps paradoxically, these sub-band pairs are referred to as BBs since the signals source from antenna BBs. A set of 4 wafers, bundled together onto one cable, are routed to one "**X**" or "**Y**" input of one or more Baseline Boards in the **k**th sub-band correlator via Fanout Boards (Fanout Boards are not shown in the figure for clarity). Thus, the set of 4 wafers contain data from the same output **k** from each of the 4 Station Boards, and each output **k** can be connected to any sub-band filter output **s**. It is the signals contained in the 4 wafers (from the **X** station) that can be correlated in the correlator chip (**CC**) with logically identical signals from

---

[10] There are some restrictions in tunability. Refer to page 22 for details.

some other antenna (the **Y** station). For correlation of the complete bandwidth, a sub-band correlator for each output **k** is required (i.e. 16, even though 18 are available). Depending on the number of stations, bandwidth, and number of spectral channels, a sub-band correlator can consist of as little as 1 Baseline Board to as many Baseline Boards as are required for the desired spectral resolution.

Cable wafer assignment provides flexibility for correlator construction.

(It is the ability to flexibly assign cable wafers to Station Board outputs and Baseline Board inputs, flexibility in the digital filters and Delay Module, cross-bar switching, and the use of Fanout Boards that provides considerable flexibility for the construction of other correlator configurations. This is why a small Feedthru Backplane that gets populated everywhere was developed, rather than a monolithic backplane fixed for a particular configuration.)

RC functions, embedded IDs, S/W learns cable config.

On the Baseline Board, a set of 4 wafers is fed into each of the 8 **X** and 8 **Y** "**RC**" inputs of the board. The signals on these wafers (data, timing, control) contain embedded IDs that exactly identify their source and that can, at all times, detect link errors. Thus, once the correlator is "cabled up" it is possible for the software to learn the cabling configuration to determine if all baselines can be processed, and to determine how observations can be processed within the physical routing constraints of the correlator. Details of the signaling protocol on these wafers is contained in [5], and details of the cable and connector physical attributes are contained in [8].

RCs operate in their own clock domain.

The **RC** is the Recirculation Controller FPGA on the Baseline Board and it contains logic to synchronize and line-up the signals within all 4 wafers for further processing. Synchronization to other RCs is not necessary at this point, and final X/Y timing synchronization is achieved on the correlator chip (**CC**). The RC contains a full cross-bar switch so that any signals from any wafers can be routed to any of the 8 inputs to the correlator chip, with or without recirculation. Refer to page 24 for more information on the functionality of the RC FPGA.

Cross-correlations: 1 to 16 baselines.

The correlator chip contains some switching circuitry that, combined with the RC cross-bar switch, allows cross-correlation capabilities ranging from correlating all pairs on one baseline, to correlating one sub-band for 16 baselines.

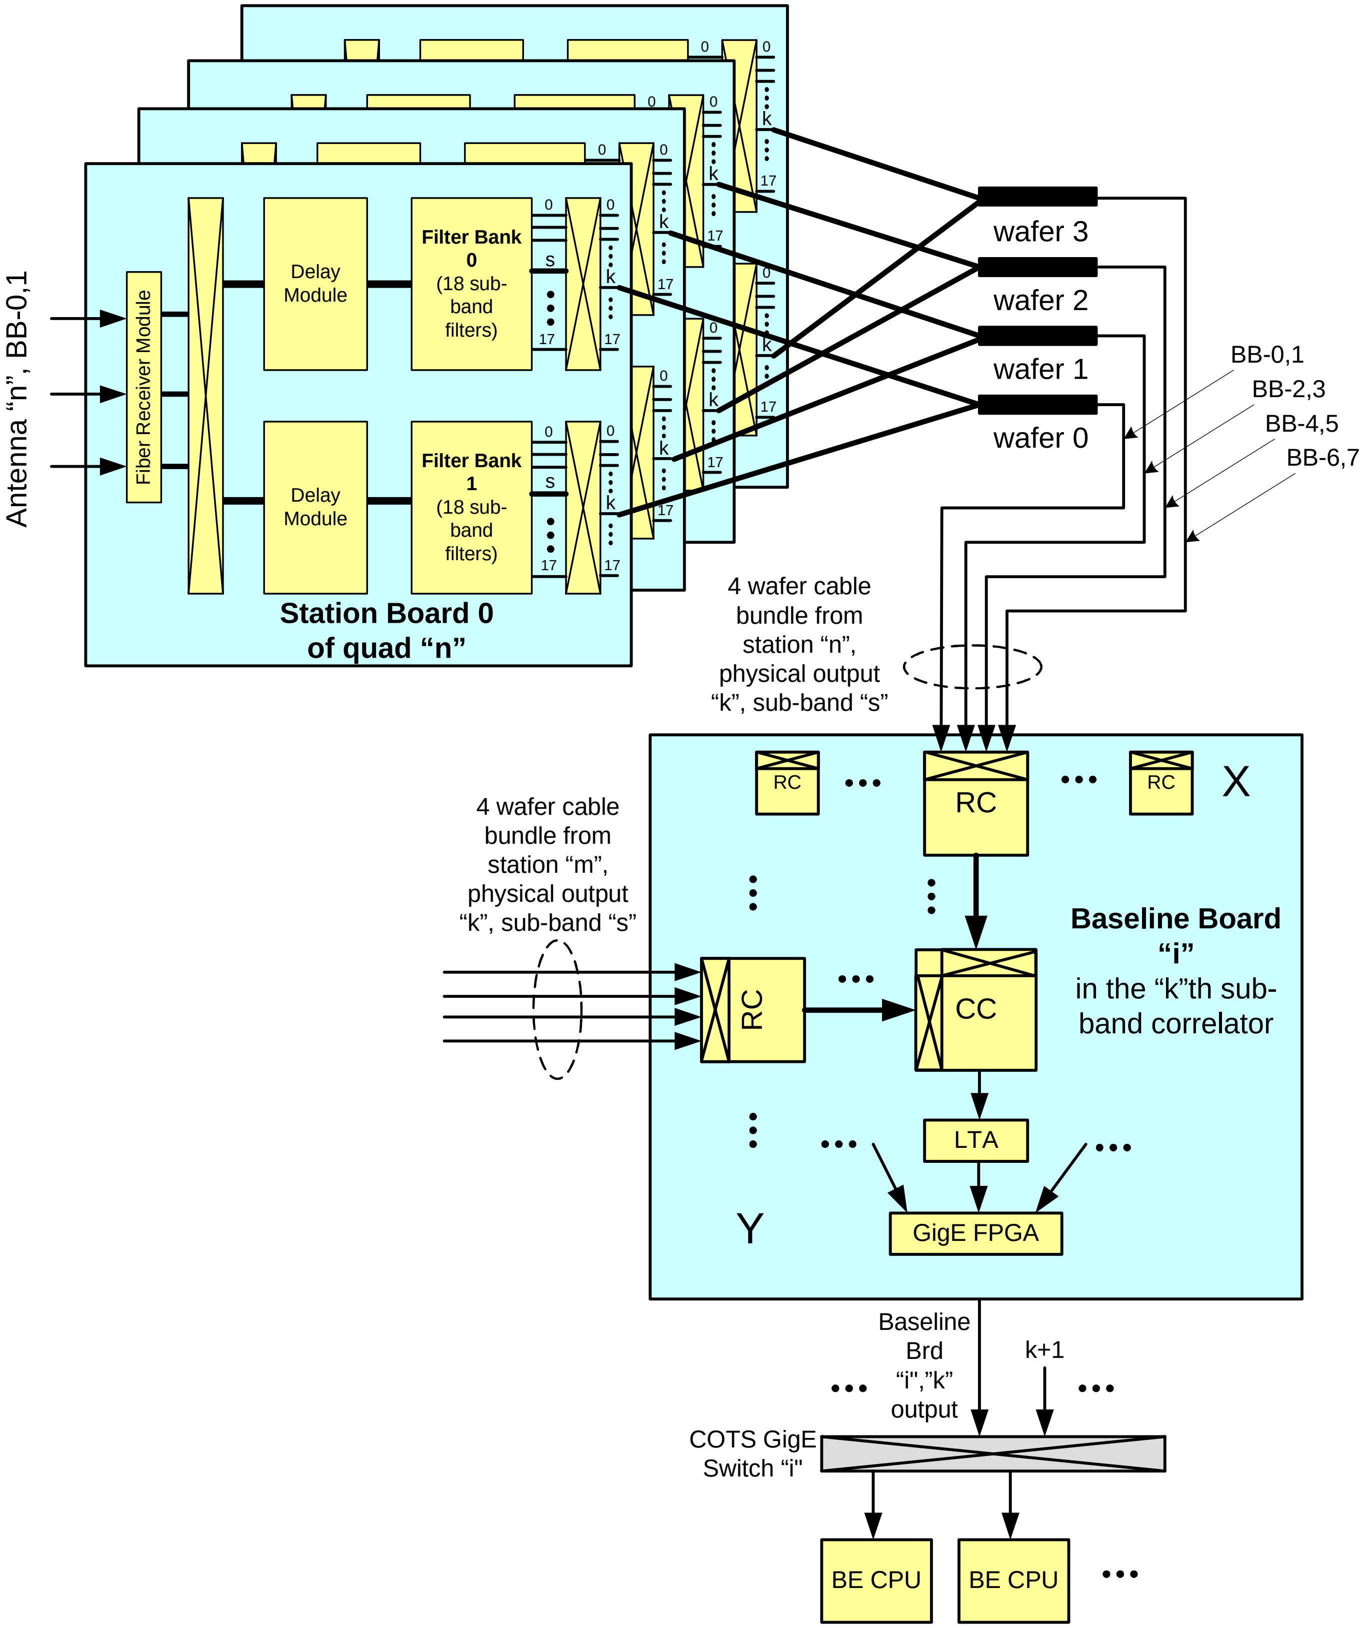

Figure 5 EVLA correlator data routing and connectivity overview.

Alternative Correlator Data Routing

2 GHz/pol'n, 8 station, 8k channels/baseline correlator.

Figure 6 is a simplified diagram of an alternative cabling plan for a non-EVLA correlator. This routing might be used where the total bandwidth for an antenna can fit on one Station Board, and only a minimum number of spectral channels are required. Using this configuration, an 8-station correlator with 8k channels/baseline can be built with 8 Station Boards, 4 Baseline Boards, and 4 Fanout Boards. Other configurations are possible.

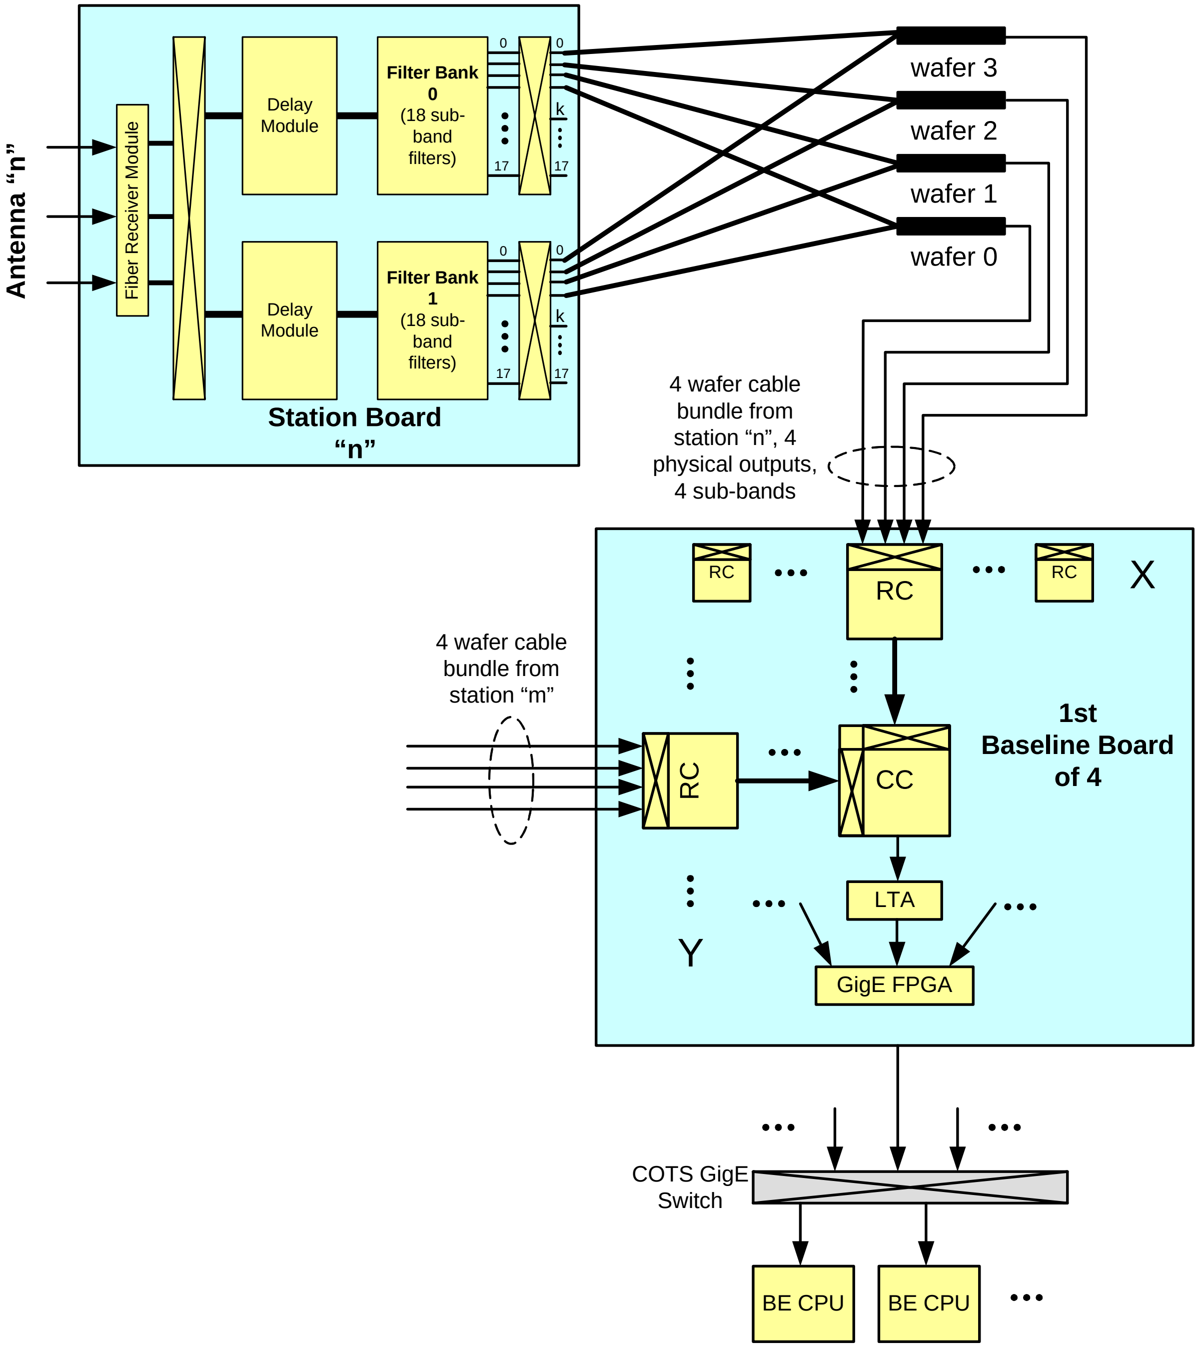

Figure 6 Alternative non-EVLA correlator data routing. Other configurations are possible.

Station Board Functions

A simplified block diagram of the Station Board is shown in Figure 7. Station Board functionality is described in the following paragraphs, referring to the figure and data characteristics at points **(a)** through **(h)** as shown in the figure.

Fiber receiver module (FORM) cross-bar, data highways, state counts.

Sampled data and time ticks arrive at the **FORM** mezzanine card (a.k.a. DTS Receiver) from the antenna. As previously explained, four Station Boards are required to fully process the entire 16 GHz EVLA bandwidth, and this description applies to each board of the four equally. The FORM/Station Board interface is four 32-bit, "data highways" **(a)** at 256 Mbits/sec/stream, although the EVLA uses only 3 of the highways in the 2 GHz baseband mode since it uses 3-bit rather than 4-bit samplers. In 8-bit 1 GHz mode, only two of the 32-bit data highways are used. The following **cross-bar switch** [9] allows each highway to be switched to one or both Delay Modules after being converted to the required internal format from the raw A/D converter format by the "**State counts, sample bit conv**" block. Thus, the data at **(b)** is essentially the same as at **(a)**, except for this bit conversion. This block also contains wideband state counts that allow A/D converter statistics to be obtained. These state counters acquire counts of the number of occurrences of each 4-bit or 8-bit combination (state) in a time-multiplexed fashion.

Delay Module delay buffer,+/-0.5 sample tracking, baseband modes.

The **Delay Module** [10] mezzanine cards each contain a 0.25 second (1 Gsample deep) buffer, and control logic to perform real-time delay tracking to +/- 0.5 samples at 4 Gs/s. The Delay Module also supports the 8-bit mode and various modes of decreasing bandwidth and increasing number of basebands so that it is possible for the Station Board to process a range of wideband and narrowband data (although only 2 GHz and 1 GHz modes are applicable to the EVLA). The Delay Module design is such that when an integral delay change occurs, no data is missing or blanked. The 1 PPS and 0.1 PPS time ticks from the antenna, along with a data valid signal also travel through the delay buffer and experience the same delay as data.

Delay Module is mezzanine card—could be replaced with different function-specific design.

The output format of the Delay Module **(c)** is the same as the input except that geometrical delay compensation has been performed.

Since the Delay Module is a mezzanine card, it is possible that different delay functions could be included in the future by replacing the card. For example, if a larger delay buffer were required, the module could be replaced with a design with more RAM, although if this were done, *every* Delay Module in the system would have to be replaced.

Wideband autocorr, 64 lags at a time, distribution to filter banks.

The **Wideband Autocorrelator** [11] block acquires autocorrelator lags (and "cross-autocorrelator" lags between data highways) on a time multiplexed basis, 64 lags at a time. Wideband Autocorrelator data is normally only ever used for diagnostic functions and so the small loss in sensitivity incurred by time multiplexing is irrelevant. This function is implemented in an FPGA and it also prepares the data for distribution to the filter banks (output **(d)**). The autocorrelator otherwise does not directly do anything to the data going to the filter banks.

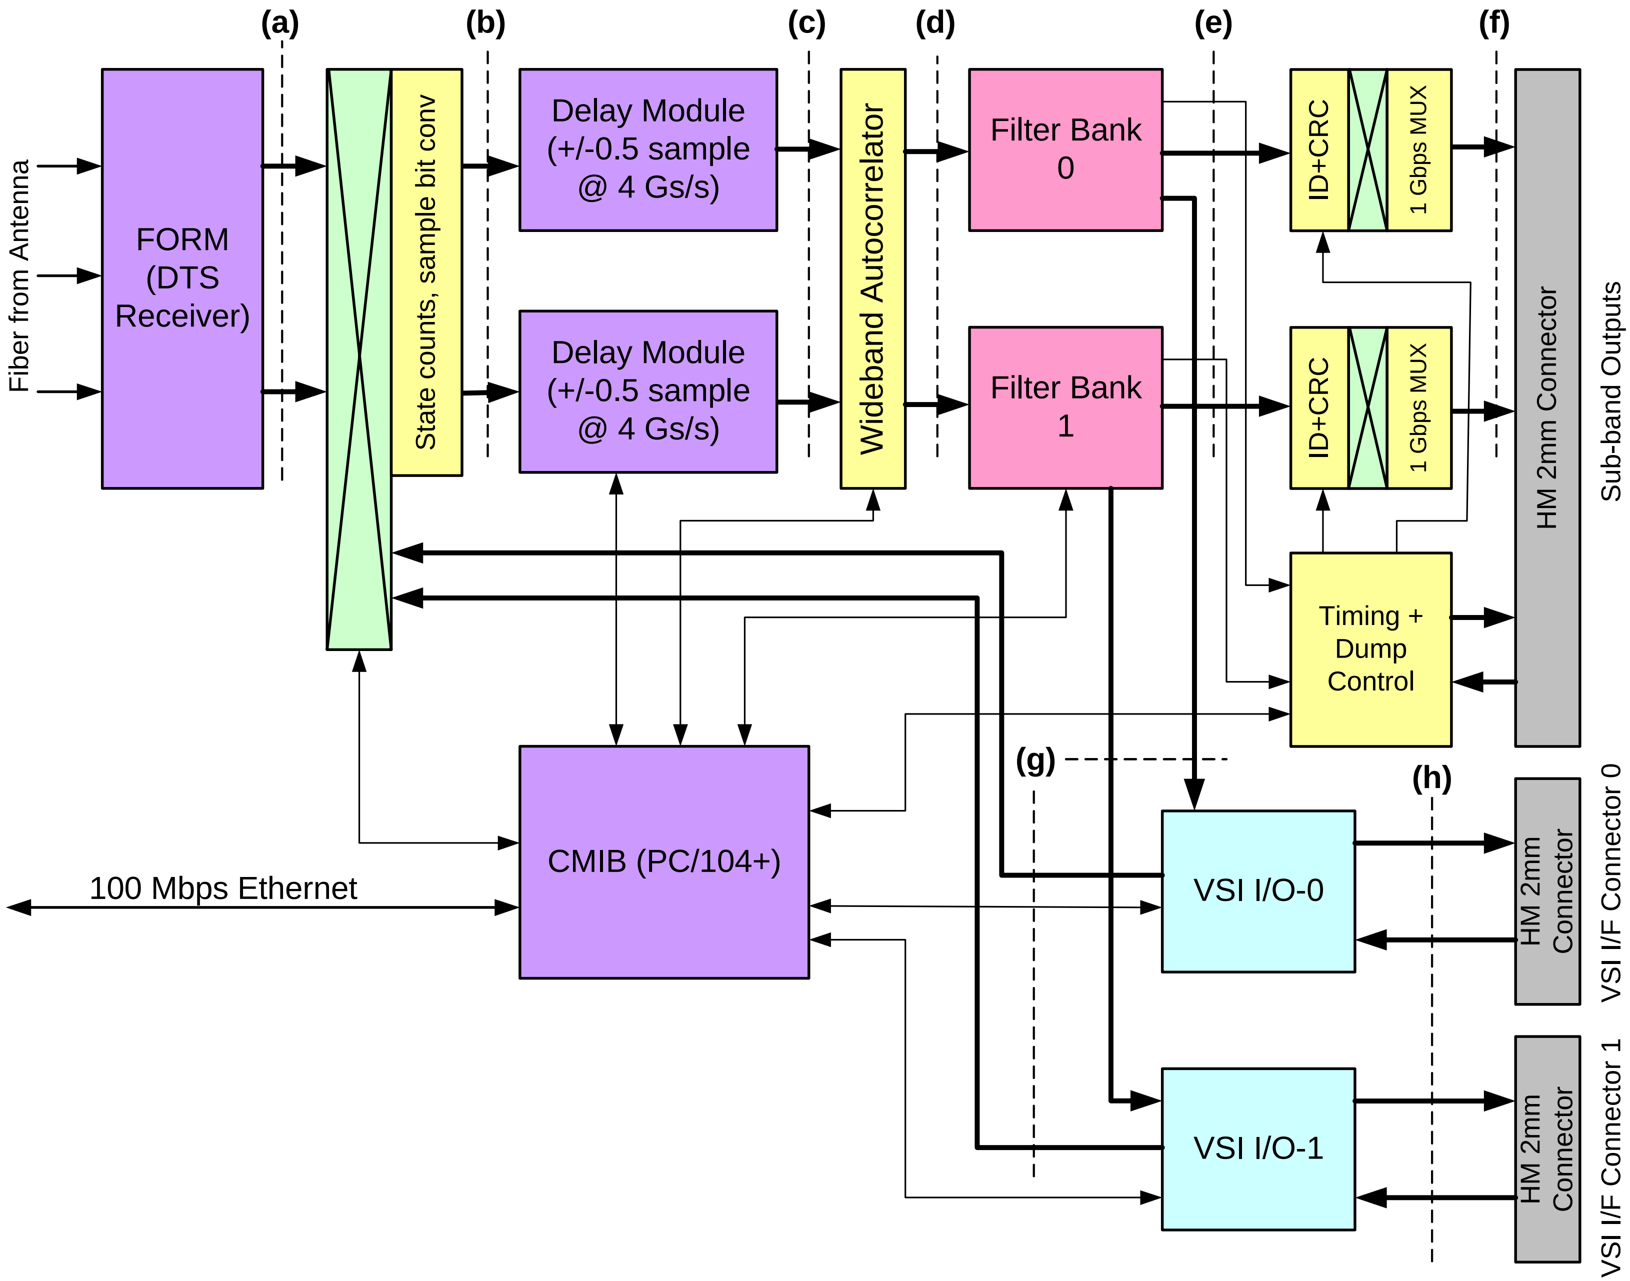

Figure 7 Station Board simplified block diagram.

Filter banks: 18 filters each

Each **Filter Bank** contains 18 individually tunable digital filters; 16 filters for selecting sub-bands to be processed by each sub-band correlator, one filter for selecting a sub-band where there is no interference for noise diode calibration and to form a reference to minimize the effects of time-variable interference, and one filter for redundancy or to provide additional sub-band correlation capability.

Each tunable digital filter is implemented in an FPGA and a complete description of the functions in the filter FPGA is given starting on page 19.

Sub-band tuning, filter output, 4 and 7-bit re-quantization.

Data exiting the Filter Banks **(e)** are independent sub-bands and each one is therefore effectively the output of a sampler. Each of these sub-bands can be a different bandwidth and placement in the original wideband signal, with some restrictions. Each sub-band stream is 4 bits wide (including the encoded data valid [5]) and one “PHASERR” stream that contains the residual phase error required for sub-sample delay tracking (refer to Figure 2) for that sub-band. When the sub-band re-quantizer is set for 7-bits, the 7-bit data is multiplexed onto the 4-bit sub-band stream by alternating least-significant and most-significant nibbles of the 7-bit word. Thus, the maximum sub-band bandwidth when 7-bit re-quantization is selected is 64 MHz.

ID and CRC code insertion, cross-bar, 1.024 Gbps mux, data capture for radar processing.

The next block [13] formats the 4-bit stream by performing a data replacement to inject the IDs and CRC check codes in the data before going to a full cross-bar switch and multiplexing up to 1.024 Gbit/sec for transmission to the Baseline and Phasing Boards **(f)**. This block contains the pulsar gate generator—a different gate can be generated for each sub-band. ID insertion is performed before the cross-bar switch so that there is no ambiguity as to the source of the data stream for downstream hardware. Refer to [5] for a complete and detailed specification of the protocol on this interface. Normally, IDs and CRC check codes replace a small amount of data every 10 msec, coincident with the time tick in TIMECODE. However, the period for replacement can be changed on the TGB (Timecode Generator Board) to once every 1 second or once every 10 seconds if there are observations that may be sensitive to this small data replacement. It is important to note that the replaced data gets flagged as invalid and thus does not correlate or produce false correlation in the correlator. This block also contains a double-buffered 32k-word data capture facility to allow acquisition of filtered data into software memory. Likely, this will normally used for narrowband modes such as for radar processing.

Timing and dump control, TIMECODE selection, phase model generation.

The “**Timing + Dump Control**” [14] block is implemented in an FPGA and takes care of timing and control housekeeping functions for the board. This includes selection of TIMECODE from one of 6 possible inputs and generation of control and timing signals on the 1 Gbps/wafer interface. These control and timing signals [5] include phase model generation (PHASEMOD), dump control (DUMPTRIG) generation under CPU control, and merging/formatting of PHASERR streams from filter chips.

Dump command generation, synchronization.

Dump generation is completely under CPU control; the CPU effectively builds a hardware dump sequence for the upcoming 10 msec period, synchronized to TIMECODE. Thus, it is possible for dump control (and the subsequent information contained in DUMPTRIG) to be synchronized to some external object, provided a model of the object relative to UTC is established. One such example is dumping using phase binning synchronized to an astronomical pulsar, given an ephemeris for the pulsar that at the very minimum contains the pulsar period. Dump control also must generate the correct sequence of dump commands to support recirculation. In all cases, synchronization of dump control across all Station Boards that are part of the same sub-array is essential for proper operation of the correlator. Synchronization is facilitated by the globally distributed TIMECODE signal and each board operates relative to its selected TIMECODE without regard to other boards. Only on the correlator chip is final X/Y clock and timing synchronization achieved and hardware facilities exist in many places to ensure that dump control or timing synchronization slip has not occurred.

Dump synchronization across all Station Boards important for proper operation.

The **VSI** (VLBI Standard Interface) **I/O** blocks **0** and **1** are each implemented in an FPGA. Each interface **(h)** supports a VSI input and a VSI output, where each input and each output contains a 32-bit data highway, a 1 PPS tick, and a clock. Currently, the VSI FPGAs are merely placeholders on the board and do not perform any functions since their design is not part of EVLA development, however all of the data routing paths exist on the board to perform VSI I/O functions. Additionally, the VSI I/O data are provided on a HM 2 mm connector, and an external plug-in module with breakout to MDR-80 VSI-H standard connectors would be required to fully realize VSI I/O functionality.

VSI FPGAs are place-holders.

VSI input data is routed **(g)** by the VSI FPGAs to the input cross-bar switch for +/-0.5 sample delay tracking and, using the digital filters, *sub*-sample delay tracking even if each VSI sampled data stream is ≤ 128 MHz bandwidth. For bandwidths ≤ 128 MHz, sub-sample delay tracking is performed by FIR interpolation using the 1st stage of the filter. Refer to page 19 for a more complete description. For bandwidths >128 MHz, sub-sample delay tracking is performed as shown in Figure 2. The VSI inputs allow the Station Board to source data from VSI data generators such as playback units for VLBI rather than from the FORM. This “hook” is provided in the design to allow the correlator to process VLBI data.

VSI I/O functionality and data routing.

Data after sub-band filtering **(g)** is routed to the VSI I/O FPGAs so that sub-band data can be sent on the VSI interface for further processing or recording. Thus, the Station Board could function as a filter bank for antenna data. This hook is provided in the design to allow for recording of NMA (New Mexico Array) antenna data for EVLA Phase-II, but could be used as a filter bank card for a VLBI antenna (for example). It is even possible to use the Station Board as a filter bank at an antenna with delay and phase rotation on record, if some application so chooses, as long as the sub-band bandwidth is ≤ 64 MHz[11] (using the Stage 2 DSSB mixer in the filter chip—refer to the filter chip description on page 22).

Station Board as a filter bank card. Delay and phase correction on record for VLBI possible.

An alternative mode of operation of the VSI I/O FPGAs is also possible. These FPGAs, in combination with the HM 2 mm connectors, can accept data from the Phasing Board[12]

Using the VSI I/O chips for feedback of phased data to correlator, 1.024 Gbps format.

---

[11] Only if 60 dB of dynamic range is required. Full 128 MHz bandwidth, with a different FPGA personality and ~30dB dynamic range is possible.

[12] Phasing Boards are not described in this memo as no design work beyond that described in [7] has been done.

outputs in the 1.024 Gbps multiplexed format, align the data streams from multiple Phasing Boards, and feed it into the input cross-bar switch. The data can then go to the correlator to perform autocorrelations or cross-correlations with other non-phased, perhaps VLBI, antennas.

Embedded processor CMIB, PC/104+, PCMC

The **CMIB** (Correlator Monitor Interface Board) consists of a COTS PC/104+ board complete with an embedded processor that communicates, via a "sandwiched" PCMC [15] (PC/104+ Mezzanine Card) with all devices on the board with a simple synchronous read/write "**MCB**" (Monitor Control Bus) interface. A PC/104+ form-factor board is a low-cost COTS embedded processor module that has a standard footprint and is about 3.5" on a side. This module runs the Linux RTOS (Real Time Operating System) and responds to 10 msec interrupts from the Timing and Dump Control FPGA.

CMIB functions. RT Linux FPGA booting, real-time, config, device drivers.

The CMIB is responsible for all FPGA bitstream/personality configuration, all observation configuration of registers on chips, real-time functions such as setting the delay tracking coefficients in the Delay Module, phase models, and dump control coefficients in the Timing + Dump Control FPGA, and all communications with the host computer (MCCC) via 100 Mbps Ethernet. Device drivers are all Linux/Unix "*/dev/deviceX <argument list>*" style and all communications including configuration information, status, etc. is via XML (eXtensible Markup Language) to the host computer. XML is an ASCII format protocol that is, as the name indicates, extensible and flexible.

PCMC card. All FPGAs in-system-program.

The PCMC card contains an FPGA that gets booted by the PC/104+ module after reset/power up. Once this FPGA is booted, all of the other FPGAs on the Station Board can be booted. Thus, all FPGAs on the board are in-system programmable. The PCMC FPGA contains a PCI bus interface that converts CPU PCI bus accesses to MCB read/write accesses. The PCMC also contains a multi-channel A/D converter that is used for monitoring voltages and temperatures on the board.

Voltage and temp monitoring, deadman thermal overload protection, remote control via CPCC.

The board design includes voltage window comparators that drive front-panel LEDs as well as drive a voltage monitor output pin on the HM 2 mm connector that eventually goes to a CPCC (Correlator Power Control Computer) for secondary monitoring. Each power supply on the board is internally thermally protected and can be remotely shutdown via a power control line, also going to the CPCC. Thus, it is possible to remotely control power to each Station Board in the system, and remotely monitor power supply voltages, independent of the health of the CMIB. If everything fails (cooling, network, CMIB, CPCC) there is a deadman thermostat switch that will shutdown the on-board power supplies in the case of a thermal overload, thus protecting the board and the system in the worst case.

JTAG for mfg test, on-chip logic analysis for debugging.

The Station Board also contains a JTAG interface and connector on the front panel that is primarily used for manufacturing testing to ensure that all connections between pins on chips are good. This interface can be used for FPGA programming, although there is no plan to do this. This interface is also available for FPGA debugging with the insertion of IP cores in FPGAs that act as on-chip logic analyzers. The results from these on-chip analyzers are read out via the JTAG interface.

Finally, the Station Board contains extensive facilities for on-line checking of the connectivity of data paths between chips using CRC check codes, but without data replacement, and with minimal chip logic overhead.

Station Board Filter Chip

The filter chip FPGA, for which there are 18 in each of the two filter banks on the Station Board, deserves special attention as it forms the heart of Station Board signal processing and contains many functions that enable flexible signal processing for the EVLA and other potential correlators. A simplified block diagram of the filter chip is shown in Figure 8.

Synopsis of filter chip functionality.

In broad-brush strokes, the filter chip contains 4 stages of digital FIR (Finite Impulse Response) filtering to allow wideband to very narrowband (31.25 kHz) operation. Narrowband operation is required for targeting spectral lines of interest, and the narrowest band of 31.25 kHz is required for radar-mode processing. The chip also contains facilities for real-time interference detection and blanking, noise diode on/off binned power detection, phase-cal extraction, state counting, sub-band beam delay tracking, sub-sample delay correction coefficient generation and, in some cases, sub-sample delay correction. All of the functions will be described in reasonable detail in the following paragraphs. For a more detailed description of the filter chip, please refer to [12].

Sub-band multi-beam delay memory, baseline length, beam offset.

The **Sub-band beam delay tracking memory** is 8k words deep and is used in those cases where it is desired to offset the sub-band interferometer beam from the main wideband interferometer beam by a small amount. The amount of offset that can be achieved depends on the sample rate and the baseline length according to the following equation [7]:

$$\frac{DL}{2 \cdot f_s} = \frac{B}{c} \cdot (\cos(\theta) - \cos(\theta + \Delta))$$

Where:

DL = Delay line length in words.
f$_s$ = de-multiplexed sample rate.
B = baseline (physical antenna separation)
c = speed of light
θ = main baseband interferometer beam angle from the azimuth.
Δ = sub-band interferometer beam offset from θ. The factor of 2 in the denominator means this can be a +/- quantity.

For example, with this buffer, the maximum sub-band beam offset is restricted to about +/-0.125^o for a B=2200 km baseline, at full bandwidth (DL=8k, f_s=256 MHz, θ=90^o). <u>This is enough buffering for all-beam imaging or in-beam phase referencing for a 25 m antenna on a 900 km baseline</u>.

Final sub-band delay error calculation. Sub-sample delay error use case#1.

The sub-band beam delay tracking logic combines the **Wideband delay tracking error** with its own residual delay error to produce a **Final residual delay error** output that is used for final sub-sample delay correction in one of three places, depending on the nature of the input data and the sub-band bandwidth. For wideband input data, and output from either Stage 1 or Stage 2 of the filter, this residual delay is used for "normal" sub-sample delay correction according to Figure 2 by offsetting the phase of the complex mixers in the correlator chip on the Baseline Board. The signal to do this is called **PHASERR** in the figure, and travels with the data to the correlator chip.

Sub-sample delay error correction case#2. Avoid delay smearing with narrowbands.

In the second case, if the sub-band bandwidth is very narrow, such as it would be coming out of Stage 3 or Stage 4 of the filter, then this residual delay correction is used to perform the necessary phase offset to track the residual delay error according to Figure 2 by using the DSSB (Digital Single-Sideband) mixer in Stage 2 of the filter. This has to be done because at low bandwidths, it takes significant time for the samples to traverse through the filter taps, and during this time, the delay has also changed significantly, resulting in delay smearing. This problem is avoided by using the DSSB mixer while the signal is still at a relatively wide bandwidth (i.e. doesn't take too long for the samples to traverse the FIR delay line) to perform the residual delay correction.

Sub-sample delay error correction case#3. Multiple basebands, use FIR for sub-sample interpolation.

In the third case, the data into the filter chip is less than 128 MHz bandwidth and can therefore occupy one 4-bit (or 8-bit) data path and is not in de-multiplexed form. Using the cross-bar switch (in front of the Stage 1 in the figure) identical data is fed to each phase of the poly-phase filter. In this case, the output of each of the 16 phases of the poly-phase FIR are used to perform the sub-sample delay interpolation. Each of the 16 phases of the poly-phase FIR is a 32-tap delay interpolator, set for a specific sub-sample delay. The most significant 4 bits of the **Final residual delay error** is used to select the output of one of the 16 phases in real-time, with no blanking between selections so that real-time sub-sample delay tracking is realized. This facility allows the correlator to process multiple (32 per Station Board) narrow baseband inputs and achieve sub-sample delay tracking without the added complications of doing fine delay tracking in the correlator chip. This mode is not used for the EVLA, but is useful for telescopes where multiple narrow baseband signals must be processed, rather than one wide baseband signal.

In principle it is possible to combine the first and last sub-sample delay interpolation methods for bandwidths like 256 MHz, and 512 MHz, but this functionality is not in the chip design because of limited logic and extra complication.

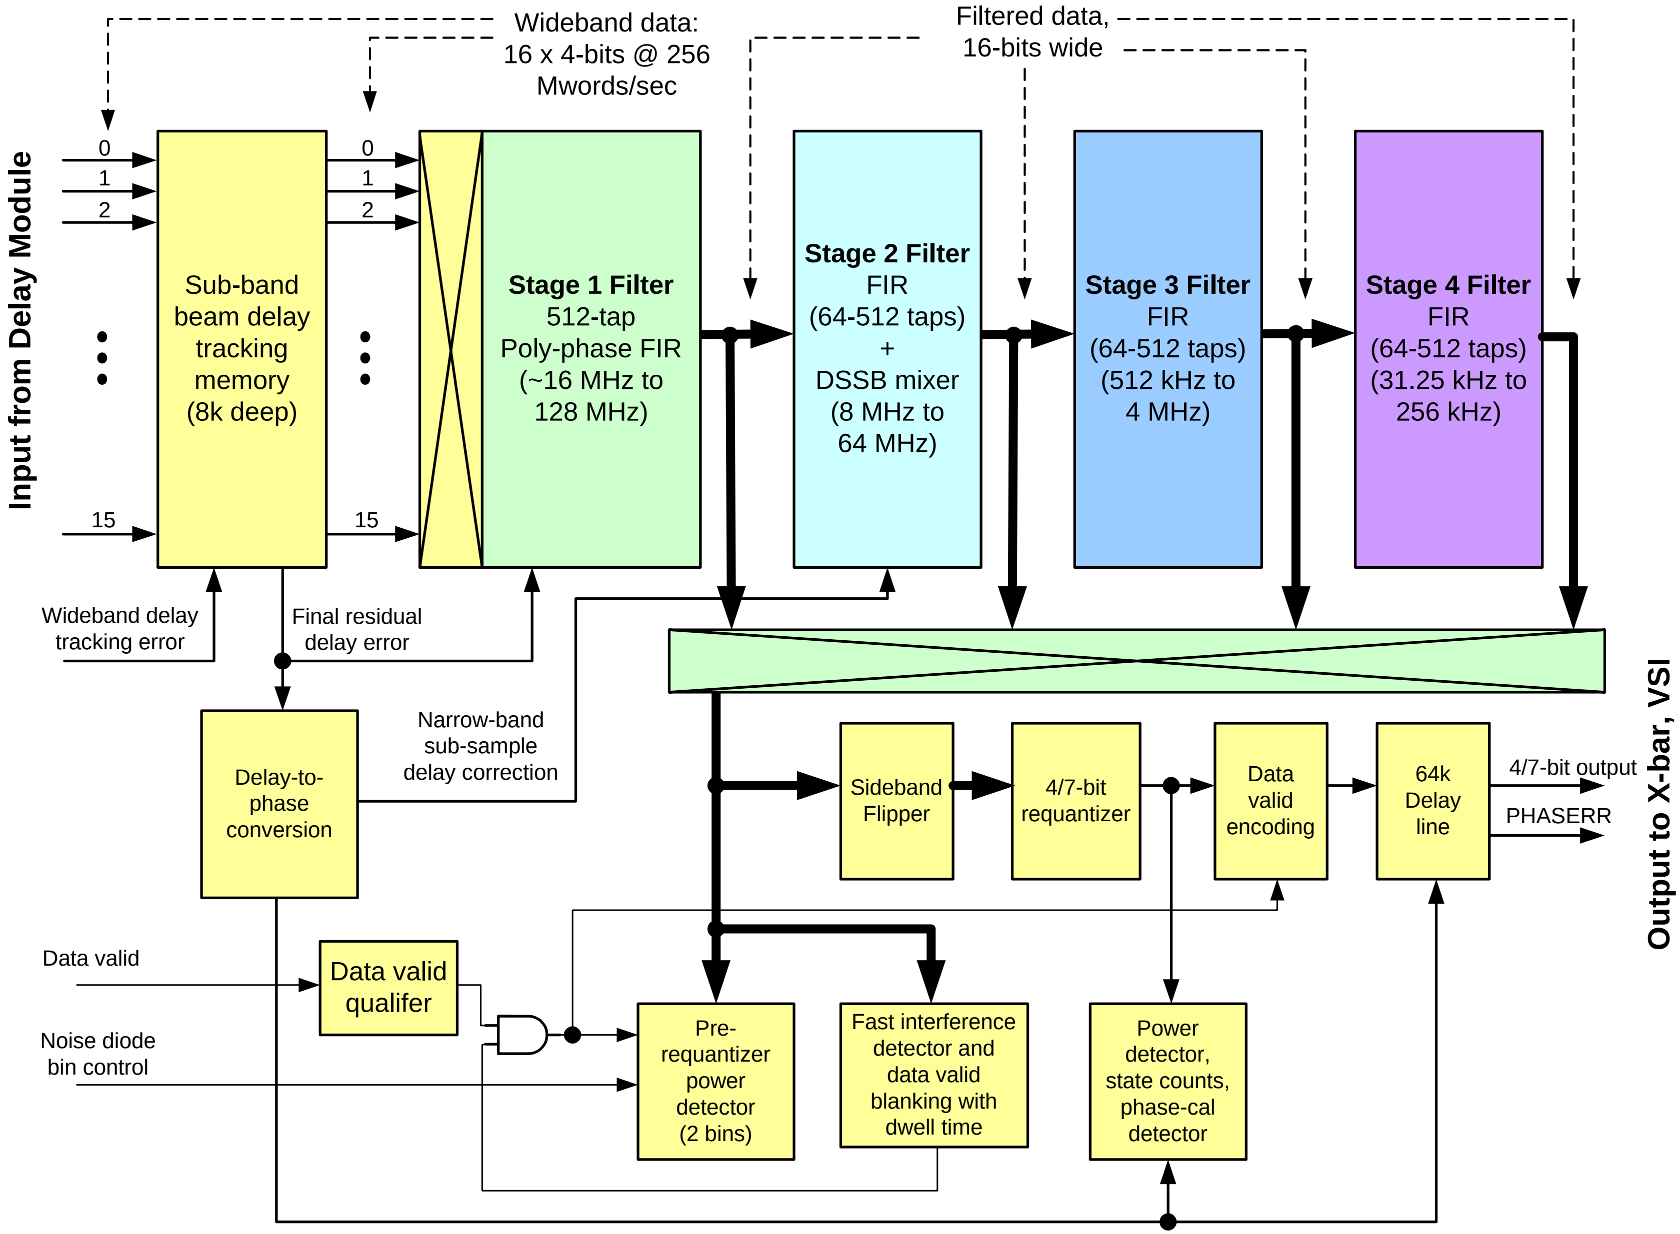

Figure 8 Filter chip simplified block diagram.

Stage 1 filter description. 512-tap poly-phase FIR, 1-8 bit operation. Tap RAM LUT with 12-bit outputs, 60 dB reject-band attenuation.

Selection switch for case#3 sub-sample delay correction.

Stage 1 bandwidths.

Stage 1 is a 16 phase, 512-tap poly-phase FIR filter that consumes the bulk of the logic and power dissipation in the FPGA, produces an output with a bandwidth from ~16 MHz to 128 MHz[13], and must always decimate by at least a factor of 16. For dynamic range, this stage implements a simple filter-and-decimate operation with no explicit mixer except for the mixing operation of decimation, and thus the tunability of this stage is restricted to be one of N integer slots in the wideband, where N is 16, 32, 64, or 128 [1] . This filter can operate on data from 1-bit wide to 8-bits wide. 8-bit wide data is processed by distributing the math operations across phases [6], so that for 8-bit data (or anything greater than 4 bits), only 256 taps are available. At each tap, the product of every data value and the coefficient is calculated and stored in a small RAM LUT (Look-Up Table) as part of configuration. Thus, in real-time, the data forms an address into this LUT, and the **12-bit** output—the product of the data times the tap coefficient—is fed to an adder tree that eventually produces the output of the filter. 12 bits is enough for 60 dB reject-band attenuation [16]. This stage contains an output selection switch, and an input cross-bar switch for third-case sub-sample delay interpolation as described above. In addition, the sub-band beam delay tracking memory, and this stage can handle input baseband combinations including: 1 x 2 GHz; 2 x 1 GHz; 4 x 512 MHz; 8 x 256 MHz; 16 x (128 MHz…16 MHz). Bandwidths less than 128 MHz to as low as 16 MHz can be handled with clock-enable circuitry built into the design. Anything less than 16 MHz requires external logic to perform a zero-insertion interpolation prior to entering the filter chip to be processed. Note that at all times, each filter chip only ever produces one sampled band output so that, for example, to handle 16 x 128 MHz basebands requires the use of all 16 filter chips.

Stage 2 operation, bandwidths, DSSB mixer, 16-bit output.

Stage 2 of the filter runs at an effective clock rate of 128 MHz and produces output sub-bands from 8 MHz to 64 MHz wide, from the input 16-bit data[14]. The result out of this stage's (and stages 3 and 4) adder tree is 40 bits wide, and 16 bits out of the upper 32 bits are selected under software control for passing onto the next stage, or for transmission to the re-quantizer. Thus, there is never a cumulative word growth problem when stages are concatenated, and the most that happens is that the output of a stage is clipped (and if this occurs, counted and reported to the CPU, allowing the CPU to adjust the selection of the 16-bit window). Depending on the output bandwidth, there are 64 to 512 taps in the filter—arranged such that the *relative* transition band steepness is always maintained, independent of bandwidth. This stage must decimate by at least a factor of 2, and at most decimate by a factor of 16, hence the output bandwidths. This stage contains a high dynamic range (60 dB) DSSB (Digital Single SideBand) mixer and can therefore place the output sub-band anywhere within the input sub-band of 128 MHz. As previously mentioned, this mixer is used for sub-sample delay tracking according to Figure 2, when the sub-band output bandwidth is very narrow to avoid delay smearing.

---

[13] Anything less than 128 MHz would start to see increasingly worse relative transition band roll-offs, and so to maintain consistent transition band steepness, other stages of the filter chip should be used with the maximum output bandwidths indicated.

[14] The number of bits passing from one stage to the next is 16 bits and thus requantization losses are not incurred by using multiple stages.

Stages 3 and 4, same as 2 w/o DSSB mixer.

Stages 3 and **4** of the filter are identical in function to Stage 2, except that there is no DSSB mixer in them, and the output bandwidths are as shown in Figure 8. Since these stages run at lower and lower clock rates, they consume less and less power and chip resources compared to Stage 1 and 2, and essentially “come for free” as they use built-in DSP resources in the chip that would otherwise be unused.

Stage output selection switch, sideband flipper, re-quantization.

As shown in the figure, a selection switch is used to select the 16-bit output from one of the filter stages for further processing and re-quantization. The **Sideband Flipper** can be set to flip the sideband back to the correct sense, if desired, after decimation (Figure 8-3 of [1]). This is accomplished by changing the sign of every other sample, synchronized to the 10 msec time tick in the chip. The **4/7-bit re-quantizer** consists of a scaler/multiplier and followed by a 4 or 7-bit truncater (i.e. by selecting a fixed 4 or 7-bit window after scaling, and clipping to ensure the output is symmetric—15 or 127 levels). The scaler can be set and left for the observation or periodically adjusted (thereby forming a software AGC) under CPU control so that optimum quantization occurs.

Data valid qualifier, programmable operation.

Delay line to align data out of different filter chips, at different bandwidths.

After re-quantization, the data valid signal, developed by the **Data valid qualifier** and the fast interference detector, is encoded in the data [5]. The **Data valid qualifier** senses data on the input and, according to a duration set by the CPU, blanks data valid for that duration—normally the entire time that the associated sample is in any one of the filter stages’ shift registers so that the filter does not produce output with bad data—but could be set to zero time to turn off this feature. Finally, the data, along with PHASERR and the time tick, is fed into a **64k Delay line**. This delay line is used to line-up data from all filter chips on the board to the same time, independent of sub-band output bandwidth. This is because, depending on how many stages have been used and the resulting sub-band bandwidth, the “wavefront” will have experienced different relative delays traveling through the filter chip. This delay line compensates for this effect.

Other blocks in the filter chip that perform calculations on the data but do not directly affect it are as follows.

Pre-re-quantizer power detector, noise diode switching, synchronization, and binning.

The **Pre-re-quantizer power detector** calculates the power in the filtered sub-band before the re-quantizer, and saves it in one of two bins under control of the input **Noise diode bin control** signal (this signal is developed with the use of a timer on the input cross-bar switch after the FORM on the Station Board—Figure 7). This power detector accomplishes two functions. The first is to make the critical measurement needed, along with filter gain and scaling coefficients, to properly scale sub-band cross-power output amplitudes so that sub-bands may be stitched together [1] . The second is to measure the sub-band power in a sub-band free of interference (i.e. by using one of the spare filter chips) when the antenna noise diode is on or off, thereby providing information required for absolute T_{sys} calibration [17].

Fast interference detection, blanking, programmable dwell time.

The **Fast interference detector and data valid blanking with dwell time** block simply looks at the data before re-quantization, and if it hits a threshold set by the controlling CPU (CMIB), pulls data valid low and holds it low for a configurable dwell time from 0 to 64k samples at the sample rate of the sub-band output. There is a counter that provides feedback to the CPU to indicate how many times this circuit has been tripped. This counter can be used in a number of ways to fine-tune the setting of the interference detector for the desired tradeoff between burst interference suppression and loss of data. Since this functionality is provided on a sub-band-by-sub-band basis, sub-band specific optimizations can be performed to deal with burst interference such as from radar, and passing interference sources such as satellites etc.

Post-re-quantization power detection, state counts, phase-cal detection.

Finally, the **Power detector, state counts, and phase-cal detector** block is used to acquire the full sensitivity lag-0 autocorrelation, the reduced sensitivity[15] state counts, and phase-cal *after* re-quantization. The lag-0 autocorrelation and state counts are used for sub-band cross-correlation Van Vleck correction, and the phase-cal extractor is used to acquire the amplitude and phase of an injected tone (normally injected in the antenna's RF electronics) in the data. Since this injected tone does not experience geometric delay, every 10 msec the CPU sets the initial phase of the internal phase-cal extractor's tone generator by taking into account the total geometric delay compensation incorporated into the signal, and then until the next 10 msec, integer delay sample changes are automatically compensated for in the phase of the internal tone generator.

Final data, PHASERR, time tick formatting and switching before Baseline Board.

The re-quantized sub-band data, PHASERR, and time tick are formatted and ready for further processing by the external cross-bar switch chip and VSI I/O chip (the cross-bar switch and the VSI I/O chip both get the same data). The cross-bar switch chip inserts IDs and CRC check codes, switches the data to the desired output, and multiplexes it up to 1.024 Gbps for transmission to the Baseline or Phasing Board. The VSI chip converts the data to 1 or 2 bit quantization, lines all of the data streams up in time, and transmits it on the output connector.

Baseline Board Functions

8x8 array of correlator chips, 1.024 Gbps demux, Recirculation Controller FPGAs.

A simplified block diagram of the Baseline Board is shown in Figure 9. This board is where final alignment of data and cross-correlation in an 8x8 array of correlator chips occurs. Multiplexed 1.024 Gbps data, control, timing, and 128 MHz clocks enter the board via the HM 2 mm connectors and are routed to the X and Y **Recirculation Controller FPGA**s.

Recirculation Controller FPGA functions: switching, recirculation, phase generation, 7-bit data mapping.

In this Recirculation Controller chip [18], the 1.024 Gbps signals are de-multiplexed and the data and control streams that source from potentially different Station Boards (Figure 5) are locked and aligned before traveling to a cross-bar switch to route them to specific output streams traveling to the row or column of correlator chips. In the chip, streams can be selected for recirculation and eventual routing to the correlator chip for correlation. The correlator chip is virtually unaware that recirculation is occurring and merely passes control data from this chip to the LTA via the correlator chip data frame. Recirculation is controlled by the "DUMPTRIG" control stream from the Station Board, and memory offsets and burst control for recirculation that depends on whether the data is an "X" or "Y" station is determined by this FPGA. This chip also loads the 32-bit point-

---

[15] Due to the time-multiplexed acquisition of state counts, with 4-bit data there is a 4X sensitivity loss, and with 7-bit data there is an 11X sensitivity loss, compared to the situation where there is a state counter for each state.

slope [5] phase models ("PHASEMOD") into 32-bit linear frequency synthesizers to generate the real-time fringe stopping phase that is passed onto the correlator chip after being truncated to 4 bits. The PHASERR signal, that contains the real-time phase offset required for sub-sample delay tracking, is added to the 32-bit fringe stopping phase before truncation. This chip also decodes data valid, generating a data valid signal for each stream going to the correlator chip, and maps 7-bit re-quantized data streams into dual 4-bit LSN and MSN output streams for 7-bit correlation. Thus, the correlator chip does not know about 7-bit correlation—it is up to controlling software to configure this chip and the correlator chip to perform the necessary distributed arithmetic to effect 7-bit correlation.

Recirculation external dual-port RAMs, 4 or 8 stream recirculation options.

Not shown in Figure 9 are external dual-port RAMs connected to each Recirculation Controller FPGA. These RAMs are the deep and fast memories that allow the FPGA to implement recirculation functions. The board design includes the copper traces for 4 external RAMs that allow recirculation to occur on all 8 streams with an appropriately sized FPGA. However, these RAMs are expensive, and thus the default configuration for the EVLA is 2 RAMs and a smaller FPGA (footprint compatible with the larger FPGA) to perform recirculation on a maximum of 4 streams (two 7-bit streams).

Recirculation Controller FPGA clock domain, alignment, synchronization.

Each Recirculation Controller FPGA operates in its own clock domain without regard to other chips. Data transmitted to a row or column of FPGAs is "self-aligned" and includes a 32 MHz clock[16] for use by the **correlator chip**. The correlator chip contains input FIFO buffers and automatic alignment circuitry so that the X and Y data are finally aligned before going to the core of the chip for cross-correlation.

The correlator chip performs all of the configured cross-correlations and, for each dump of each cross-correlation, transmits a data frame to the LTA for further processing. A more complete description of the correlator chip is given starting on page 27.

---

[16] This low-frequency clock is used to reduce power dissipation in the correlator chip.

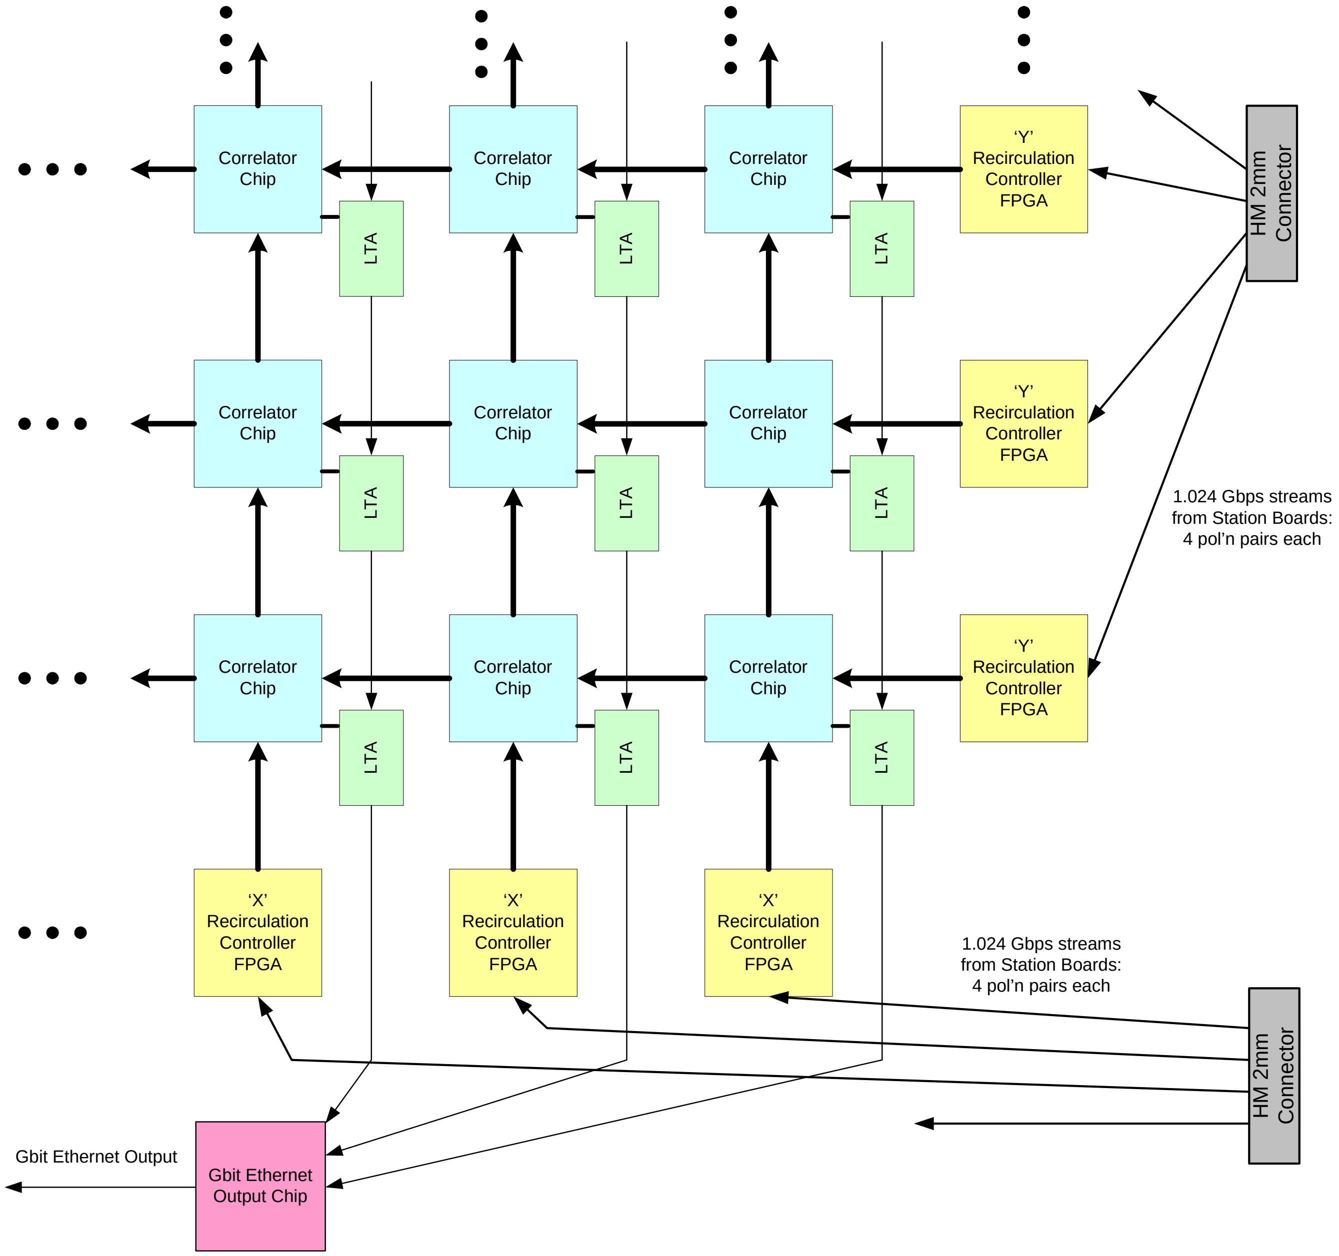

Figure 9 Baseline Board simplified block diagram.

LTA operation, dump time, # of bins, Tx to Gbit Ethernet chip.

The **LTA** is built using a small FPGA and external 16M x 16 DDR SDRAM. There is one LTA for each correlator chip to provide the fast dump performance for recirculation and pulsar phase binning. Additionally, to minimize the number of gates in the custom correlator chip, the maximum on-chip integration time is 512 µsec, and this in itself requires high-speed data transfer to the LTA. The LTA chip queries the correlator chip for data and, when ready, accepts the frame, sends it to the output OR stores and integrates it in the LTA, depending on the command contained in the data frame itself. Thus, the LTA and correlator chip act in-concert without any real-time interaction by the CPU (CMIB—not shown in the figure). When the command in the frame indicates that this is the last dump, the LTA chip saves the frame and flags it as ready for readout. Output frame transfer logic in the LTA chip traverses the memory and, when requested, transmits ready LTA data frames to the **Gbit Ethernet Output Chip** for UDP/IP frame construction and transmission to the Backend as shown in Figure 5. The LTA is large enough to provide 2000 bins (2 banks of 1000 bins each) for every data product generated by the correlator chip. Normally, frame transfer from the correlator chip to the LTA has priority so that correlator chip data will only start to be discarded when the LTA bin the frame is destined for is full. In addition, there is burst-mode control settings in the LTA chip that force priority so that the controlling CPU can decide how the LTA is filled and emptied, depending on observation requirements.

Gbit Ethernet Output Chip, SFP transceivers, 4x1 Gbps or 1x10 Gbps outputs.

The **Gbit Ethernet Output Chip** has four Gbit output data stream channels, and on the board design, each of these is routed to an SFP (Small Form factor Pluggable) cage although in the current design of this FPGA, only one output is active. Copper or fiber[17] SFP transceiver modules can be plugged into these cages to provide between 1 and 4 Gbps of output data. Additionally, copper is provided on the board to permit the installation of a 10 Gbps "XENPAK" transceiver module and could, with the correct FPGA design and a bigger FPGA (footprint compatible with the default design) allow transmission of 10 Gbps Ethernet.

For more detailed information on the operation of the LTA and Gbit Ethernet Output Chip, refer to [19] and [20].

Board overhead functions same as Station Board.

The Baseline Board contains similar overhead functions as the Station Board as described in a previous section. This includes JTAG interface, voltage and temperature monitoring via the CMIB, deadman thermal overload protection with a thermostat, and remote power supply monitoring and control with the CPCC.

Correlator Chip

4 Mgates, 130 nm, 3.5 W.

The correlator chip is where final X/Y data alignment and cross-correlation occurs. The chip is a ~4 million gate standard-cell custom chip fabricated in 0.13 µm CMOS. Estimated power dissipation at full speed (256 MHz) is approximately 3.5 W. The

---

[17] Current switch and cabling pricing and performance indicates that copper is the most cost-effective solution.

correlator chip design has undergone extensive functional testing and a test plan [21] is in place for validation and verification of the design with an extensive suite of tests [22].

Corr chip partitioning, CCQs, CCCs, switching.

A simplified block diagram of the correlator chip is shown in Figure 10. Correlation resources in the chip are partitioned in a hierarchy, of which only the top-level hierarchy is shown in the figure. A **CCQ** (Correlator Chip Quad), contains four 128-lag **CCC**s (Correlator Chip Cells). Thus, in total there are 16 CCCs in the chip. Since there are four pairs of data streams into the chip, the 16 CCCs can produce all four polarization products for all pairs. Data routing, limited two-stage switching, and the full cross-bar switch in the Recirculation Controller permit formation of a wide variety of cross-correlation products in the chip [3]. For example, each CCC can be configured to cross-correlate data from unique inputs, or can be concatenated with adjacent CCCs to produce longer-lag cross-correlations. As another example, if the streams of data entering the chip source from different antennas (in some non-EVLA configurations), it is possible to correlate all 16 baselines for one product in the chip.

Dump control formatting, routing of dump control to CCCs automatic, embedded ID extraction for output frame.

Dump control signaling, originating from Station Boards, is formatted for use by the correlator chip by the Recirculation Controller. Each CCC is dumped and generates a data frame independently, however, when CCCs are concatenated for longer-lag correlations, logic in the chip automatically routes the correct dump control signals to all of the concatenated CCCs so that all the CCCs are always dumped at exactly the same time. This logic removes the burden, and potential error, of doing this with the controlling CMIB. Additionally, the correlator chip extracts embedded IDs in the X and Y input data and includes these IDs in the correlator chip data frame so that at the final destination (Backend), it is possible to ensure that the frame contains the expected/correct cross-correlation data product.

Input error detection, off-line test vectors, signal regeneration for next chip.

The correlator chip contains synchronization detection circuitry that can alert the controlling CMIB (through polling status registers on the chip) that the X and Y input data streams are not aligned, that X and Y dump control signals are not aligned (may not be a problem except during recirculation), or that any one of the X or Y input signals has gone dead. In addition, there is a pseudo-random test vector receiver on the chip that can receive test vectors from the Recirculation Controller to test correlator chip connectivity off-line. The chip regenerates and re-times X and Y data and clocks for transmission to the next correlator chip in the row or column so that only short (<1.5”), point-to-point connections are required on the board, as indicated in Figure 9.

Refer to [23] for a more detailed description of the correlator chip.

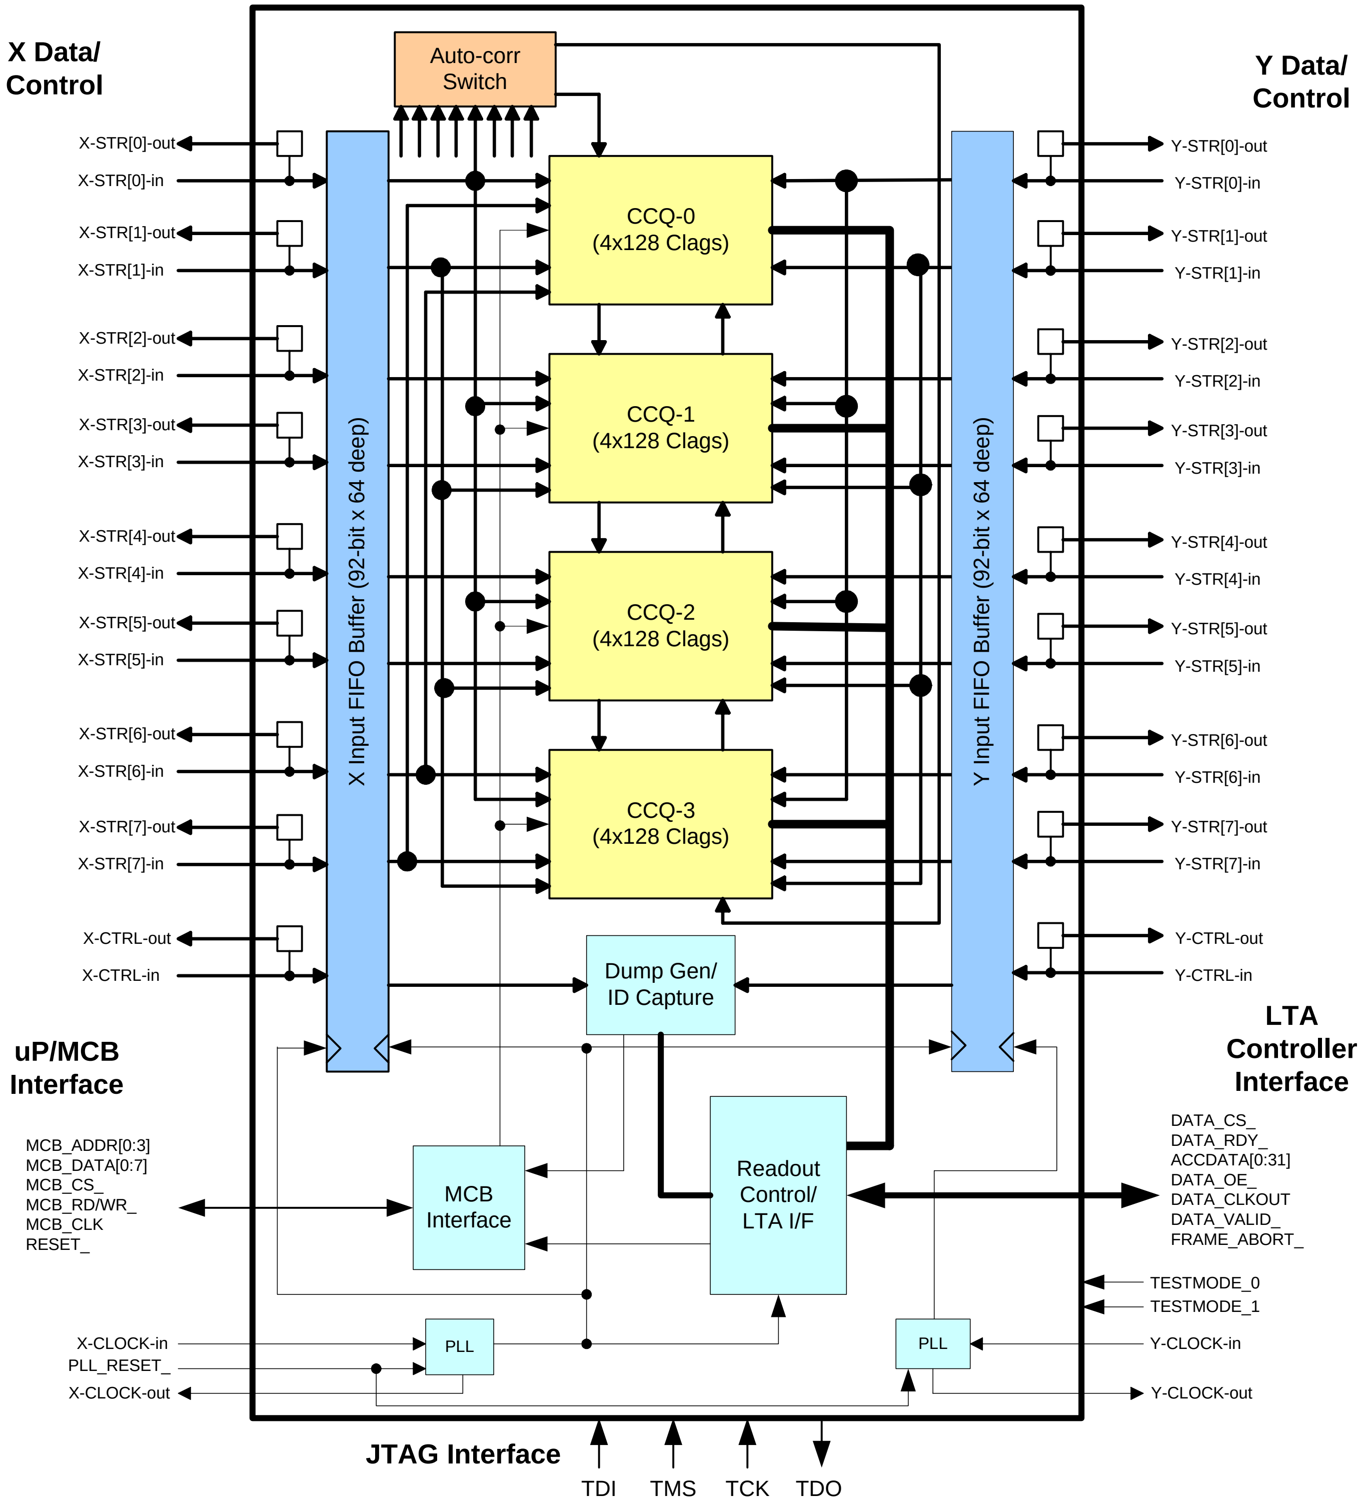

Figure 10 Simplified correlator chip block diagram

Software

Overview: MCCC MAPPER, VCI, CMIBs, CALC

A simplified function and data flow diagram of the correlator software is shown in Figure 11. For a complete detailed description of the correlator software architecture, refer to [24]. The basic software structure is reasonably straightforward. The EVLA monitor and control system talks to the correlator's **MCCC** (Master Correlator Control Computer) across the VCI (Virtual Correlator Interface) to configure the correlator in a high-level fashion. This host computer then maps these high-level configuration requests into low-level configurations (MAPPER function), based on knowledge of the connectivity of the system, and the resources available in the system. The Station and Baseline Board CMIBs (embedded processor modules) configure the hardware and then provide monitor data such as chip status registers, voltage and temperature monitor points etc. to the host MCCC for processing and potentially to notify the EVLA monitor and control system of problems. The Station Board CMIB must also perform real-time delay tracking functions and the models to do this come from the MCCC. The source of the delay models is CALC running outside of the correlator in the EVLA monitor and control system[18].

Data flow to Backend computers, UVFITS output.

Station-based data products, correlator data frames, and configuration and auxiliary data from the MCCC flow to the **Backend** computers for final processing into a standard format for the archive. The definition of this format is not yet available, however, in the meantime the Backend will save data in "VLBI-style" UVFITS format. The Backend is defined to be within the correlator, and it consists of a cluster of computers, each one processing all lags for a subset of baseline products.

MCCC MAPPER functions.

The primary function of the MCCC is the MAPPER—mapping high-level requests for correlation to low-level configuration data for the CMIB drivers. This MAPPER has knowledge of the resources in the system, their cable connectivity, and their functionality. The intent is that requests from the VCI can occur in a very high-level, general way and that the MAPPER discovers the best way to implement these requests in the correlator. The MAPPER is developed such that it is portable—it could be used in other high-level software outside of the correlator to allow for "what-if" analysis of observation feasibility and capability. The MCCC also contains real-time monitor and error handling functions that act as a gateway/filter between the EVLA monitor and control functions and the embedded CMIB processors. For a complete detailed description of the VCI and MCCC functionality, refer to [25], [26], and [27].

Embedded CMIB communications, XML, GUIs.

Communication between computers occurs over a 100 Mbps switched Ethernet[19] network using TCP/IP and XML (eXtensible Markup Language). XML is an ASCII language that is flexible and powerful. The CMIBs run real-time Linux and all device drivers are Unix/Linux style "*/dev/device_type <arg list>*". Invocation of one or more of these device drivers complete with their arguments can be made on the command line of the CMIB shell, or by executing an XML script. XML scripts can be built by hand with a text editor, built by the MCCC MAPPER, or built by GUIs.

---

[18] Although, as of this writing, the location of CALC for the correlator is not decided.
[19] Except that correlator data frames are transmitted to the Backend on 1 Gbps Ethernet UDP/IP.

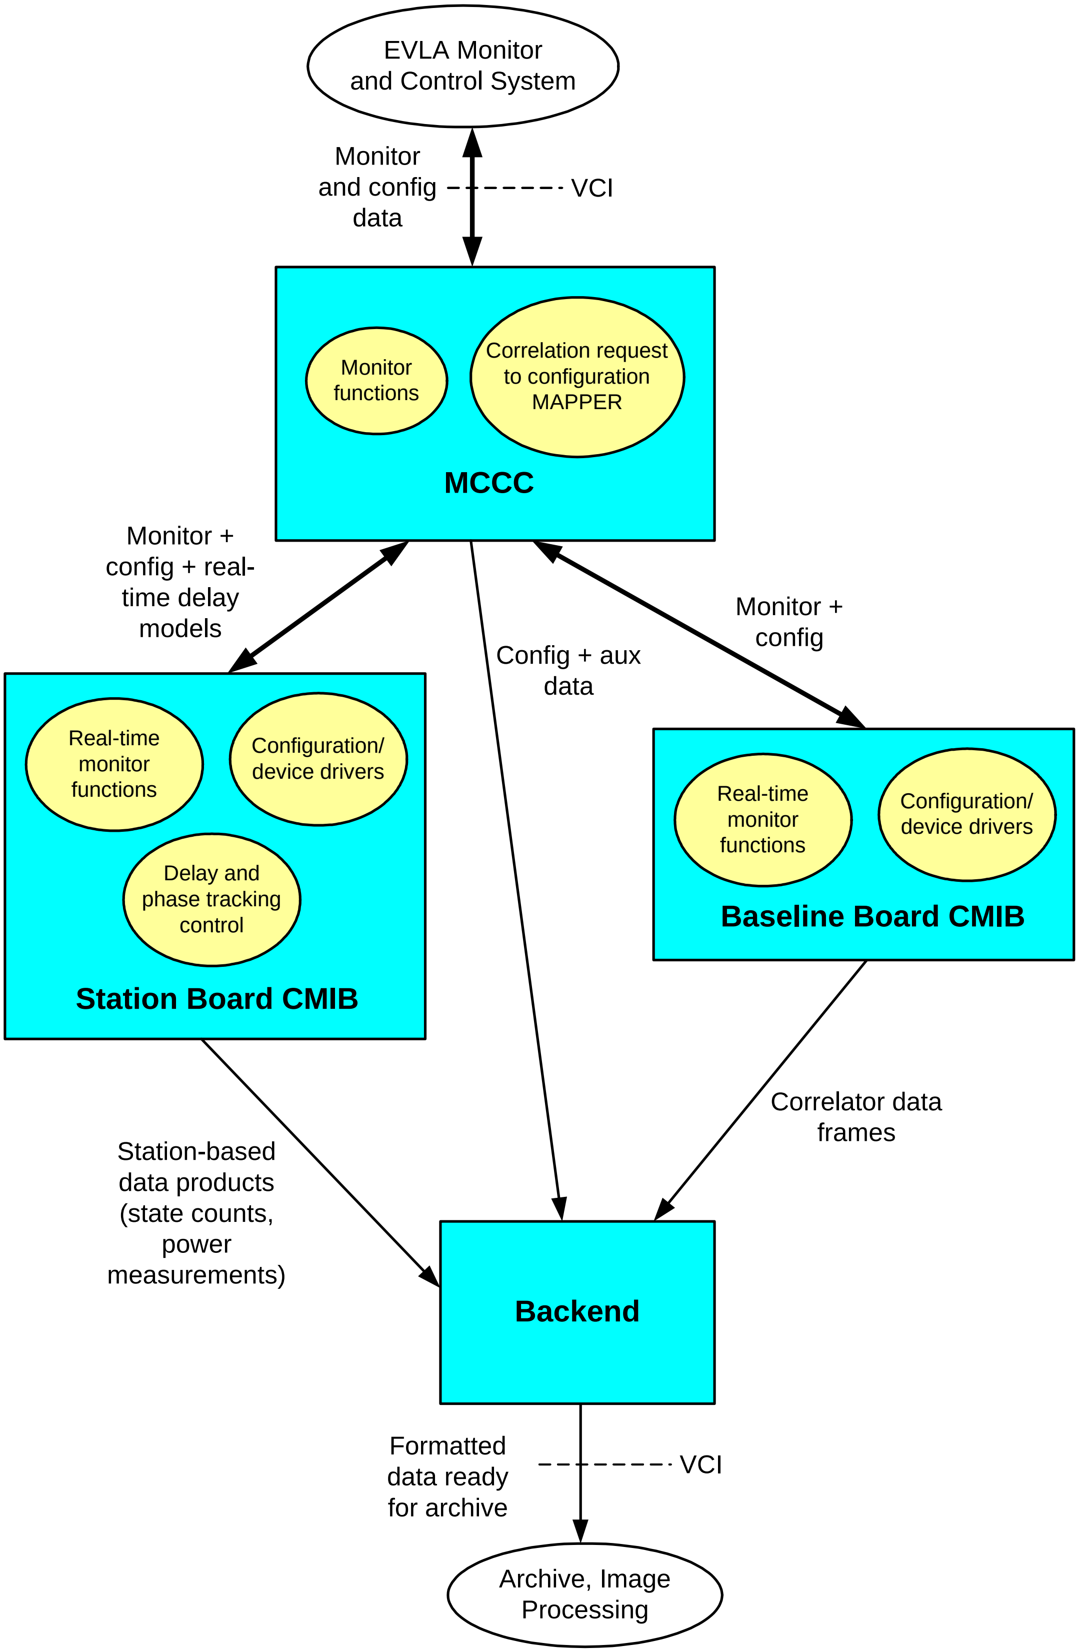

Figure 11 Simplified correlator software functions and data flows.

Thus, complex XML scripts, built using one or more methods, are executed the same way by the CPUs, independent of their build source. Real-time “deamon” processes in the MCCC and CMIBs are communicated with and requested to perform functions in a similar fashion.

Example XML file.

An example of an XML file passing across the VCI to the MCCC to configure the Station Board is as follows:

```
<VciStbConfigReq msgId="10" cfgId="15">
  <TimeStamp dateTime="2005-05-13 16:17:25-0700" millis="1116026245592"/>
  <BoardId rack="2" crate="1" slot="0"/>
  <MinHwIntegTime microSec="15"/>
  <NoiseDiode status="Off" />
  <Baseband id="0" dataPath="0" input="FO_0" bw="2048000000" bitsInInpStream="3">
    <Subarray id="2"/>
    <Station id="2" antenna="Antenna_2" antType="EVLA"/>
    <Polarization pair="1" id="Right"/>
    <Models lo="1024000000" freqOffset="0" singlePhaseCentar="Yes" fringeRotation="On" />
    <Subband id="0" >
      <Filter id="0" bw="128000000" centralFreq="64000000" useMixer="Yes" flipSideband="No"/>
      <Products quantity="4" bitsToCorrelate="4" spectralChann="64" recFactor="0"
sensLossAllowed="No"/>
      <IntegTime stbIndex="0" hw="1" lta="1" cbe="10"/>
    </Subband>
  </Baseband>
  <Baseband id="1" dataPath="1" input="FO_1" bw="2048000000" bitsInInpStream="3">
    <Subarray id="2"/>
    <Station id="2" antenna="Antenna_2" antType="EVLA"/>
    <Polarization pair="2" id="Right"/>
    <Models lo="3072000000" freqOffset="2048000000" singlePhaseCentar="No" fringeRotation="Off" />
    <Subband id="0" >
      <Filter id="0" bw="128000000" centralFreq="64000000" useMixer="No" flipSideband="Yes"/>
      <Products quantity="4" bitsToCorrelate="4" spectralBins="64" recFactor="0"
sensLossAllowed="No"/>
      <IntegTime stbIndex="0" hw="1" lta="1" cbe="10"/>
    </Subband>
  </Baseband>
</VciStbConfigReq>
```

Web-based Java GUIs for testing, system integrate+test and operational system use.

During initial testing of correlator hardware, web-accessible Java GUIs that can run on virtually any computer on the Internet will build XML scripts for configuration of devices on boards, under point-and-click control. Real-time status information about the board also passed as XML to the GUI, will be displayed by the GUI to assist in understanding the behaviour of the device under consideration. These GUIs will live on into system-level integration and test and eventually normal operations so that it is always possible to obtain a low-level, graphical view of the configuration of the correlator for debugging purposes, or simply to act as an aid to trace system behaviour anomalies. An example of a GUI screen for the Recirculation Controller FPGA on the Baseline Board is shown in Figure 12.

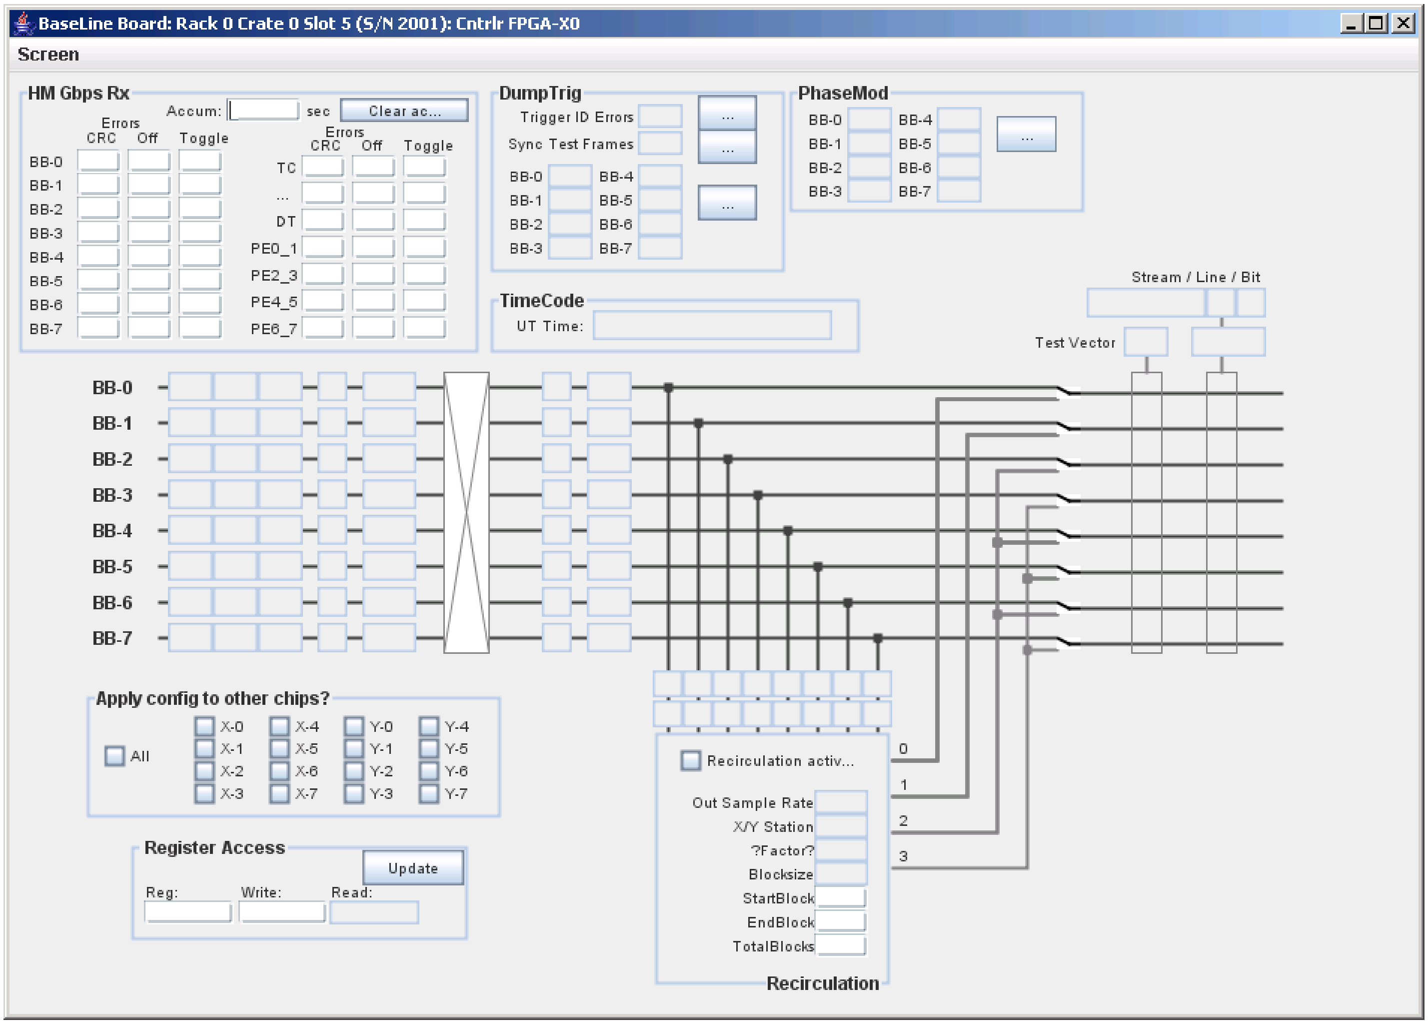

Figure 12 Example GUI screen for Recirculation Controller FPGA.

Backend cluster data flow.

The correlator Backend software is developed so that it can run on the order of a 100-node cluster of computers as a separate sub-system of the correlator. The data flow from the correlator Baseline Boards on 1 Gbps Ethernet using UDP/IP, through COTS Gbit Ethernet switches is wired and arranged such that all correlator lags required for an FFT end up on the same Backend computer. Thus, there should be little, if any, high-speed traffic between Backend computers.

Backend processing, FFT, meta data, Archive, sub-band stitching operations.

A block diagram of the correlator Backend software is shown in Figure 13. Lag data frames from the correlator are accepted, sorted, and concatenated with frames that are part of the same cross-correlation before FFT. Frequency domain data is then merged with auxiliary “meta” data, and data acquired from Station Boards before being sent for storage to the Archive. Note that many operations required to stitch sub-bands together and properly scale the data are not applied in the Backend, rather information required to do such operations off-line are sent with it.

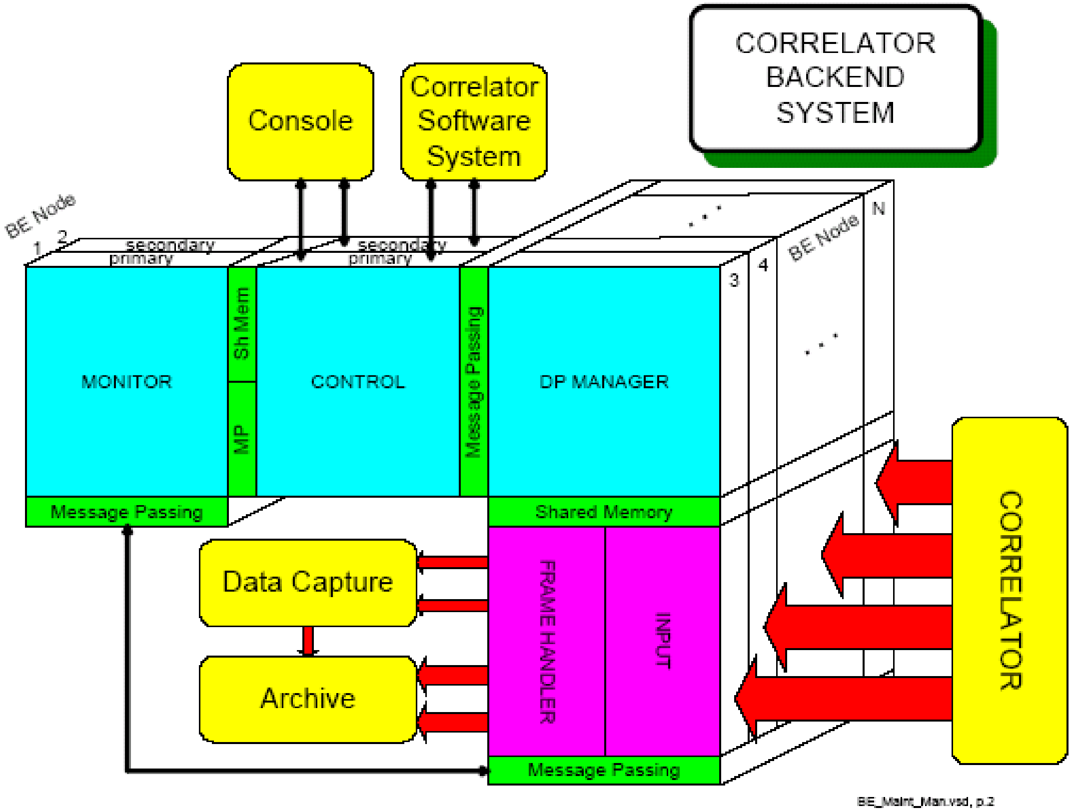

Figure 13 Correlator Backend software block diagram.

Backend M&C, cluster computer management, lag frame packet re-routing.

The Backend includes the data processing computers and a monitor and control computer that manage the data processing computers. If a data processing computer crashes or becomes inoperable, the M&C computer can restart it or communicate with the MCCC to instruct the correlator to route lag frame packets to other computers for processing. Once a new/replacement computer becomes available, the M&C computer brings it on-line and instructs the correlator to route packets to it again.

Backend output format, utilities, interference mitigation.

Initially (V1.0) the Backend will output VLBI-style UVFITS files, and will (V2.0) eventually write data to the Archive in a more flexible format, once that format is defined. The Backend also includes facilities for capturing raw data for analysis, displaying plots of Real and Imaginary vs lags, and amplitude and phase vs lags, as well as routing data for real-time display to other computers on the network. Although not currently implemented, the Backend will be able to apply interference mitigation/excision techniques in real time, once such algorithms have been developed, and once the need to employ such algorithms is recognized.

For a complete, detailed description of the operation of the Backend, refer to [28].

System Infrastructure

Synpopsis: boards, cabling, computers, power.

The EVLA correlator consists of 24 racks[20] of Station Boards and Baseline Boards, interconnect cabling, Fanout Boards for signal repetition and fanout, a couple of Timecode Generator Boards (TGBs), the MCCC, the CPCC, the Backend computer cluster, and the system COTS 48 VDC power supply. All boards in the system operate off –48 VDC, and so a COTS 48 VDC power plant supplies power to all racks. This method was chosen to minimize downtime when a power supply fails—if a power supply on a board fails, the board is replaced. If a rectifier unit in the power plant fails, it is replaced on-line without power interruption since the plant is configured for N+1 redundancy.

Design and test for cooling, minimal complexity.

Extensive testing [29] has been performed to optimize cooling, and minimize complexity in the correlator racks. A system requirement is that all boards that contain active electronics must hot swappable; it must not be necessary to power down the system or the rack to replace a defective circuit board.

Basic rack design, airflow, fans.

The basic rack design (Figure 14, [30]) for Station Boards and Baseline Boards is the same. There are two 12 U[21] x 400 mm deep x 24” wide crates that hold up to 8 boards each, and between them is a 6 U crate to hold Fanout Boards, and in one instance, the 2 Timecode Generator Boards. At the top of the rack is an airflow duct, and mounted at the very top are 4 “squirrel-cage” impeller fans, that suck air through the rack from the floor and exhaust it out the top.

Fan capability, board power dissipation.

The fans chosen for the job have a high static pressure of 0.5” of water. The linear airflow rate is about 16 LFS (linear feet per second) or 1200 CFM (cubic feet per minute), when the rack is fully loaded with boards[22]. The maximum power dissipation of each Station Board is approximately 400 W, the Baseline Board is approximately 550 W, and the Fanout Board has a relatively negligible power dissipation of 20 W.

Thermal test rack; results.

A thermal test rack [29] was built to develop this cooling strategy (i.e. rack layout), measure actual airflows, and determine actual temperature rises so that board and system design could proceed with confidence. In the test rack, the chip temperatures on the main power dissipating components, attached to a monolithic heatsink, showed *maximum* temperature rises above ambient of 25 to 30 ^{o}C. The specified air inlet temperature is 15 ^{o}C (although 10 ^{o}C air may be required at the 7000 ft altitude of the VLA), which means that the maximum operating chip temperature is 40-45 ^{o}C. Maximum chip operating temperatures in this range are a requirement for system reliability [31].

Pictures of the thermal test rack are shown in Figure 15.

---

[20] For a 32-station correlator; for a 40-station correlator, there are 26 racks.

[21] 12 U is approximately 19”.

[22] The airflow rate may change if we choose to go to a bonded-fin heatsink rather than an extrusion heatsink. Nevertheless, the same or better thermal performance will need to be obtained.

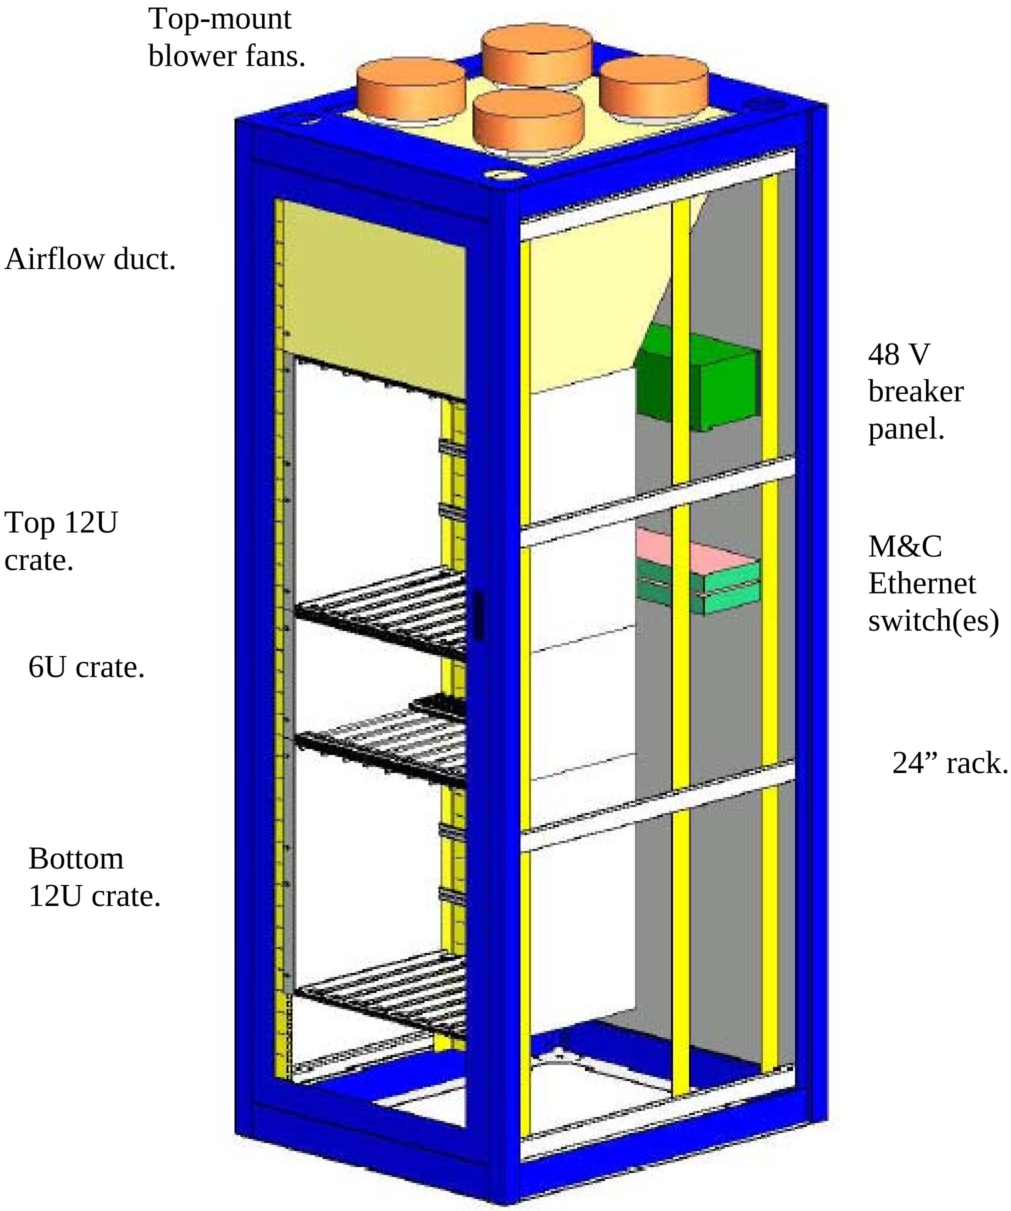

Figure 14 Basic correlator rack, side panel removed.

Hot-swapping, fan failure, rack height.

The fans and boards in all racks are all hot-swappable. Fans can be swapped-out, one at a time from the top of the rack, and tests show that if a fan fails in the worst case the temperature in the rack rises 8 ^{o}C, since the rest of the fans continue to operate, and the failed fan does not auto-rotate. The total usable rack height is 44U, and the total rack height is 7’ 3”.

Figure 15 (Right) Close-up of heat-producing boards, heatsinks, and fans in the thermal test rack. (Below) Thermal test rack.

Rack cable routing plan.

A side view of the baseline rack with the basic cable routing plan and airflow indication is shown in Figure 16. The station rack is similar (i.e. the rack designs and basic cable layout is similar), except there is only one Fanout Board, and no major cable distribution for the Fanout Boards to the Station Boards. Note that all 1.024 Gbps cable connections are to the Feedthru Backplane at the rear of the rack, and that all boards plug into the Feedthru Backplane from the front. Airflow is not affected by cable routing since it all occurs within a duct where there are only boards.

The following points are of note in the figure.

Rack plenum, raised floor.

- The rack is bolted to a plenum that is mounted on the floor of the screened room installation. The raised floor tiles and hangars fit around the rack and thus the rack load is not directly on the floor.

Adjacent racks bolted together, inter-rack airflow.

- Adjacent racks in a row do not have (recessed) side panels installed, and are bolted together for increased stability. With the bottom 12U sub-rack as shown, there will be an additional air channel between racks that are adjacent and this helps to reduce the airflow impedance between racks. When a door is opened, the racks will draw in room air for a short period of time until it is closed. It is not anticipated that this will be a problem, since it is a transient condition.
- Airflow within sub-racks is not affected by cable runs.
- Cables from the rear-mount M&C Ethernet switch are routed in the space beside the sub-racks. This cable routing is not shown in the figure.

Remote power M&C, fan speed M&C; CPCC.

- Two wires for remote power supply monitor and control from each Feedthru Backplane, and for fan speed monitor and control connect to terminal blocks in the rack. A cable from each terminal block in each rack routes to the CPCC. This is not shown in the figure.

48 VDC cable routing.

- -48 VDC power and return cables (from the 48 VDC plant distribution panel[23]) enter from the top rear of the rack into a breaker panel mounted inside the rack on the rear vertical rails. There is one breaker in this panel for each board, and a pair of wires from each breaker in the panel runs to each board's Feedthru Backplane power-entry mounting screws. Since each board requires two power connectors, a small backplane-to-backplane power routing PCB is to be installed as shown in Figure 17. Additionally, there will be one breaker for all Fanout Boards in a rack.

Overhead cable trays.

- Two cable trays will be required above the racks. One for –48 VDC power and one for signaling. These are nicely separated into a cable tray running along the rear of the racks, and a cable tray running along the front of the racks.

Cable strain relief.

- There is a large mass of 4 wafer 1.024 Gbps cables and these will need strain relief. Commercially available or custom strain relief mechanisms are being considered.

---

[23] The 48 VDC return is grounded at the plant distribution panel, but is isolated at each rack.

- The front door is smoked plexi-glass, and the rear door is solid metal. Side-panels are only installed on end racks of a row.

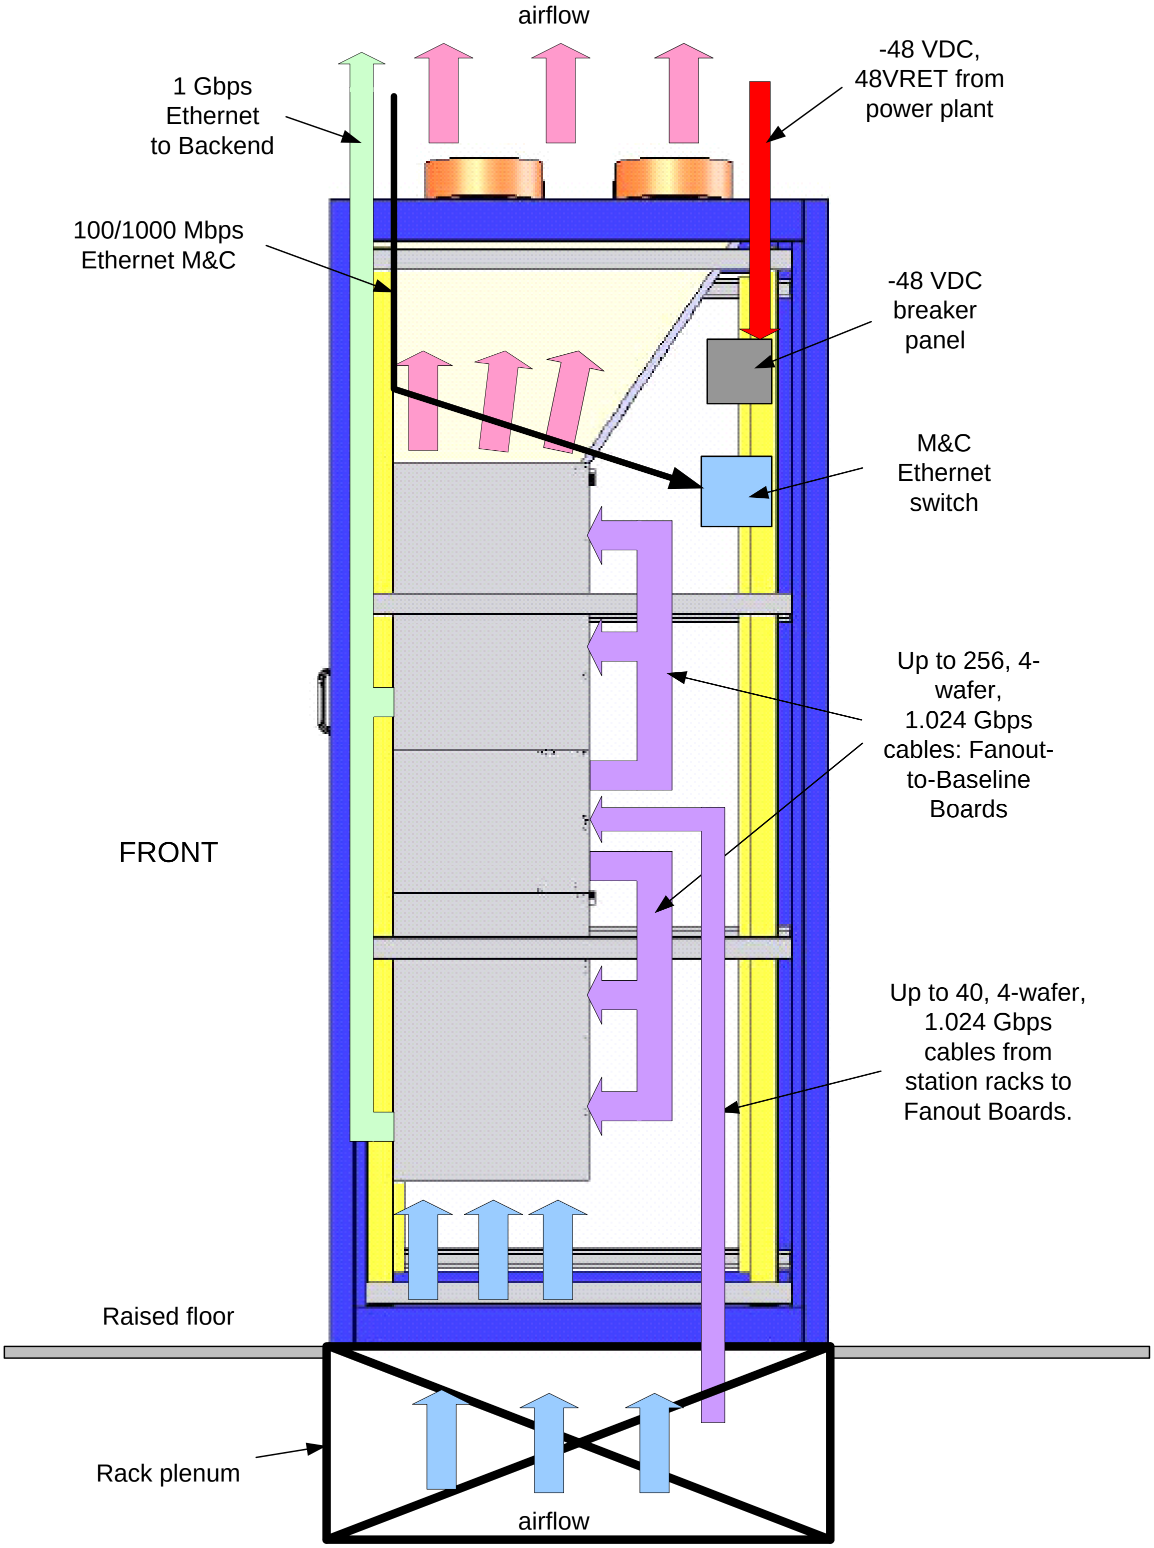

Figure 16 Baseline rack side view and basic cable routing plan. Airflow paths are shown.

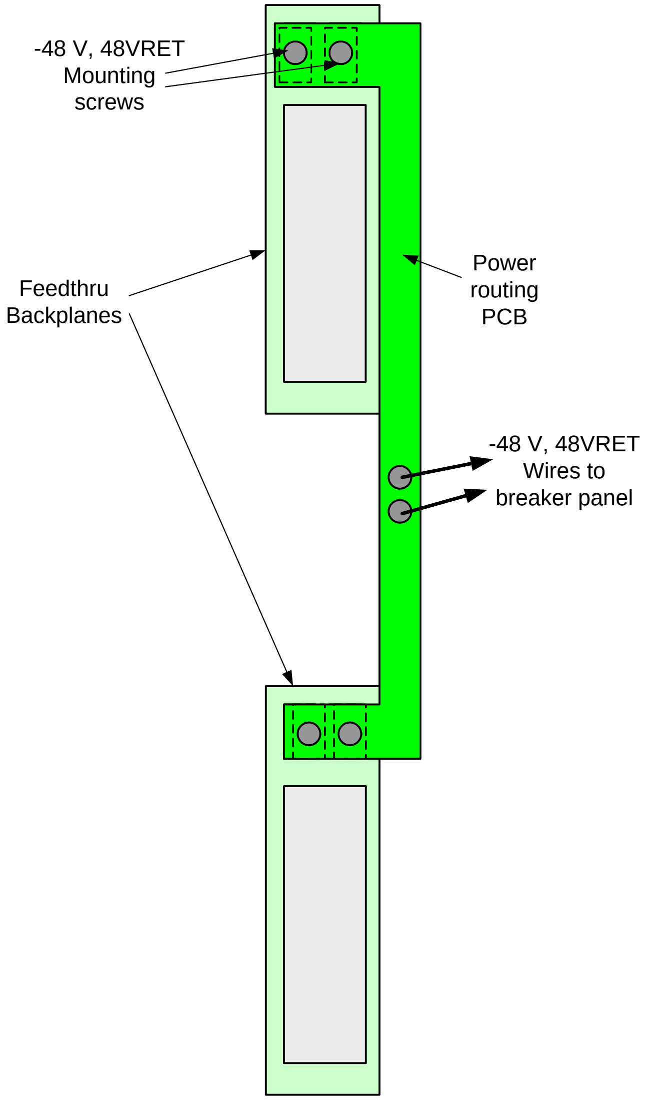

Figure 17 Feedthru Backplanes with installed Power routing PCB. This PCB eliminates the need for extra wires to connect power to each backplane.

Sub-rack duct front covers.

- The front of the 12U "sub-rack ducts" will be covered by virtue of the front panels on the large PCBs. For unused 12U slots, filler panels with partial air blockage will be required. For the 6U sub-rack ducts, the Fanout Boards are recessed and thus a hinged, transparent front cover will be required.

- If necessary, a separate rack-mount duct can be installed from the bottom 12U sub-rack to the floor if the gap under the bottom 12U sub-rack is problematic for airflow.

Fan speed monitor and control.

- Fan speeds are individually controlled and thus it is possible, by monitoring the temperature of boards within a rack, to equalize cooling across the system under computer control.

Power failure battery backup, sleep mode, minimize thermal cycles.

- The power plant has 5 minute backup under full power. The current plan is to put the correlator into “sleep mode” by asserting the reset lines of the power consuming chips (except the embedded processor which will remain active) on the boards if there is a power failure greater than ~2 minutes. In this case, the fans will be slowed down to try to maintain temperature for as long as possible in an effort to eliminate a thermal cycle. It is believed that eliminating power-outage induced thermal cycles will improve the reliability of the system.

The current floor plan for the EVLA correlator installation is shown in Figure 18.

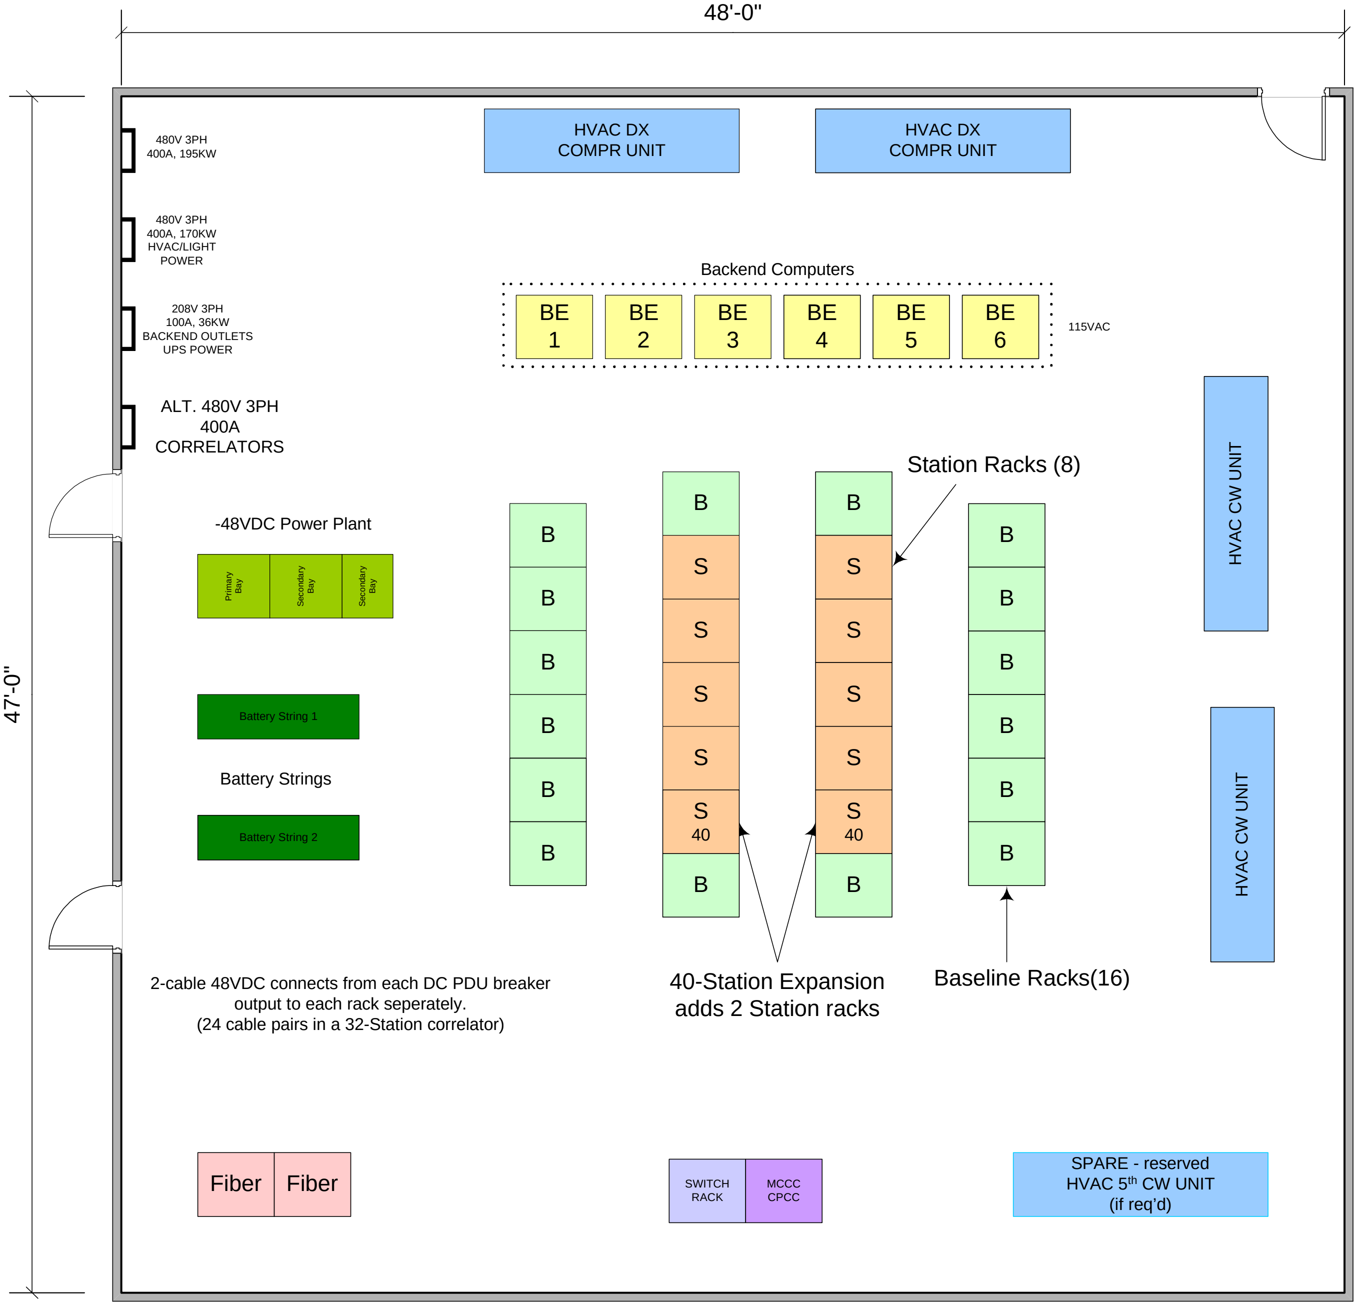

Figure 18 EVLA correlator floor plan.

For the floor plan of Figure 18, the following points are of note:

- All components shown in the figure are in a screened room.
- There is a minimum 2.5' of clearance to the front or back of each rack.
- The Backend (BE) computers and GigE switches are in separate racks from the correlator.
- The Fiber racks contain the WDM de-multiplexers to produce fiber outputs to go to the correlator, 12 for each antenna.

Correlator room attributes, including total power reqs.

- The –48 VDC power plant and its batteries are installed in the room. The plant and batteries are safety rated for an office environment and no particular safety precautions are required. The input to the 48 VDC plant is 480 VAC, 3 phase.
- The total system power dissipation, not including the Fiber racks and the HVAC systems is **188 kW** for a 32-station correlator and **255 kW** for a 40-station correlator.
- For the Backend computers, GigE switches, MCCC, and CPCC a battery-backed 110 VAC source is required[24]. For a 64-computer Backend, the power requirements are approximately 35 kW. A 3-phase, 480 VAC, battery-backed supply will be provided to the correlator room for this purpose, thus likely a 50 kW power transformer will be required in the room to develop the necessary 110 VAC voltage for this purpose.
- As previously mentioned, each correlator rack requires ~1200 CFM airflow at ~15 °C. It may be necessary to drop this to 10 °C at the altitude of the VLA, to achieve desired maximum chip temperatures.

For more detailed information on correlator room requirements and plans, refer to [32].

MCCC config, packaging.

The MCCC (Master Correlator Control Computer[25]) is a COTS PC, rack-mount or desk-top, likely identical to the Backend PCs. The MCCC will consist of two redundant on-line computers so that if one fails, the other one continues to operate. Hot-swapping of this computer and the Backend computers is a consideration when choosing the package/form factor for these computers.

CPCC functions.

The CPCC (Correlator Power Control Computer) is a high reliability/high availability compact-PCI computer, with a number of digital and analog I/O boards that facilitate the monitor and control of power supplies on each board and monitor and control of the rack fans. Monitor and control is independent of the network and any the health of any CMIB or MCCC computer.

---

[24] The CPCC *may* be powered off –48 VDC, although there is some desire to control the –48 VDC plant with the CPCC. The CPCC may require its own, dedicated AC UPS.

[25] It is named as such to differentiate it from the 'CCC' of the correlator chip.

Board power supply M&C.

Each board has one power supply control line that must be TTL high for the power supply on the board to be on. The control line drives an LED in an opto-coupler on the board being controlled. Thus, if the CPCC dies, or the connection gets broken, the power supplies on the board all shut down to protect the circuit board. The monitor line from the board is an indication of the health of every power supply on the board, and whether the power supplies are all producing voltages within specification. If this line is high, everything is fine, if it goes low, then at least one power supply voltage is out of spec.

Ground-loop problems and solutions.

The power supply monitor and control signals on each board are referenced to signal/chassis/earth ground, since there are only 2 pins[26] available for each board to perform this function. Thus, routing all of these lines (as well as the fan speed monitor and control lines) and at least one signal ground from each rack to one central computer could provide a high frequency ground loop path for switching power supply-induced common-mode noise that is undesirable.

There are a couple of ways of handling this problem. The first way is to isolate all of these lines with opto-couplers by using a control board in the CPCC with built-in opto-couplers. The problem then is that a reference voltage must be developed at the rack, and fed back to the control board, not to mention the cost of the large number of COTS optically-isolated boards required. This option requires a separate circuit board, -48 VDC power supply feed, power supply etc. to do this. The second way is to use TTL I/O, referenced to ground through a standard connector (e.g. DB 37-pin x 2 per rack), and then build a filter board with male/female connectors on opposite ends to filter any high-frequency common-mode noise. This would still require a small PCB design (or perhaps an acceptable unit can be purchased off-the-shelf), but would be a completely passive solution and could merely be installed in-line with the cable running back to the CPCC. Then, only a COTS terminal block would be required in each rack to connect the pair of wires from each board to the DB 37-pin connector. Yet a third way is to use COTS DC-DC isolators for each line—likely cost prohibitive.

Fan speed monitor and control reqs.

Fan speed monitor and control requires a control board with the capability of detecting and counting pulses for monitoring the speed (3 pulses per revolution), and for producing a variable amplitude DC output of 0-10 V for control.

---

[26] One of the three power pins on the blind-mate power connector is used for monitor or control; with two connectors, this provides two pins. There are no HM 2.0 mm pins available for this purpose.

System Circuit Boards

A layout diagram is shown for each of the circuit boards in the system in the following figures. On each board with a CMIB, a 16-bit readable serial number is attached and fixed at manufacturing time.

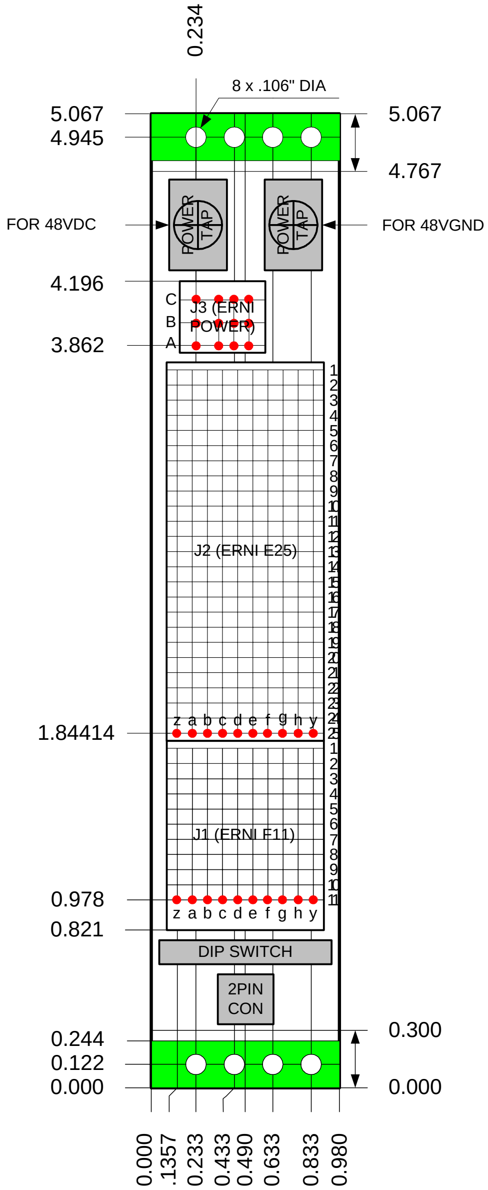

Figure 19 Feedthru Backplane layout diagram. All boards and 1.024 Gbps cable wafers plug into this backplane. The DIP Switch sets an 8-bit ID that can be read by the plugged in motherboard's CPU, and with two boards, a 16-bit ID can be formed [33]. A minimum of two of these backplanes are required for each (large) board—for power delivery, and for IDs.

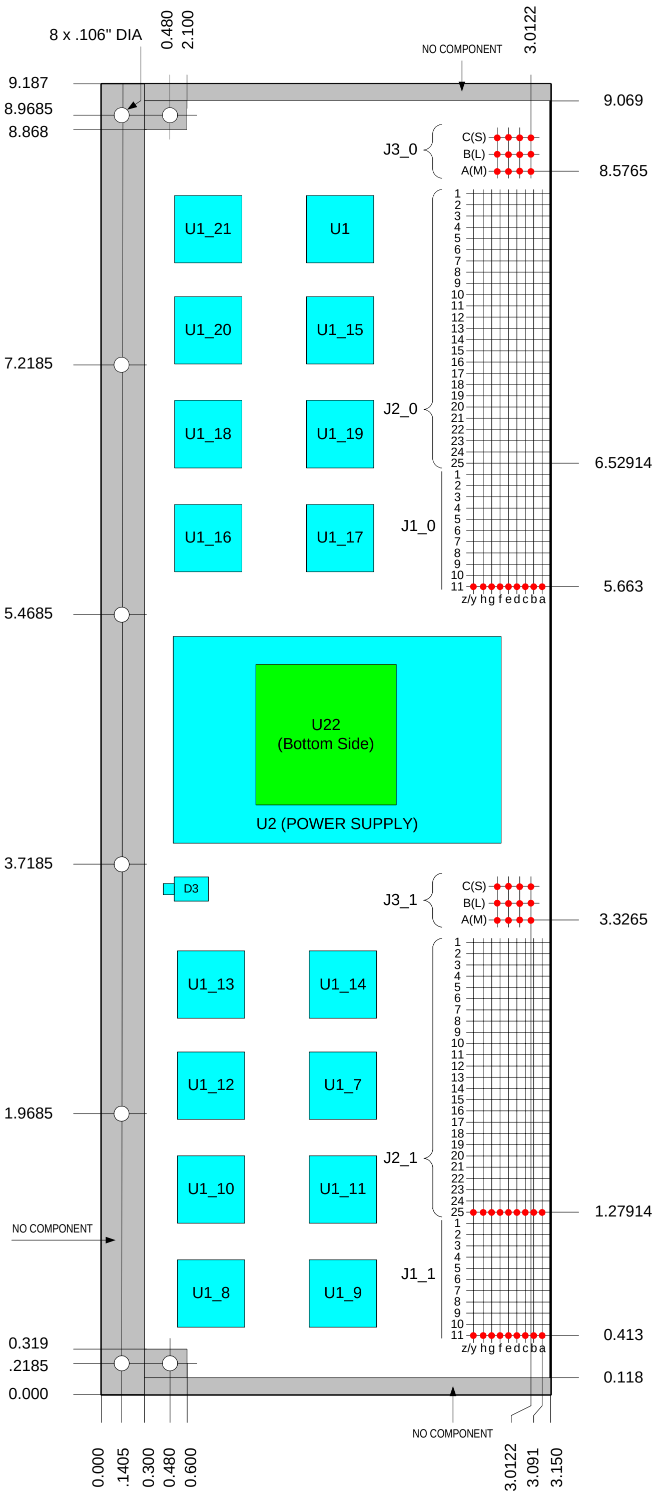

Figure 20 Fanout Board layout diagram. This is a short, non-standard depth board. The board fans-out two sets of 4 wafers a factor of 8 each. Dual fanout capability was chosen to save cost (shared power supply and board).

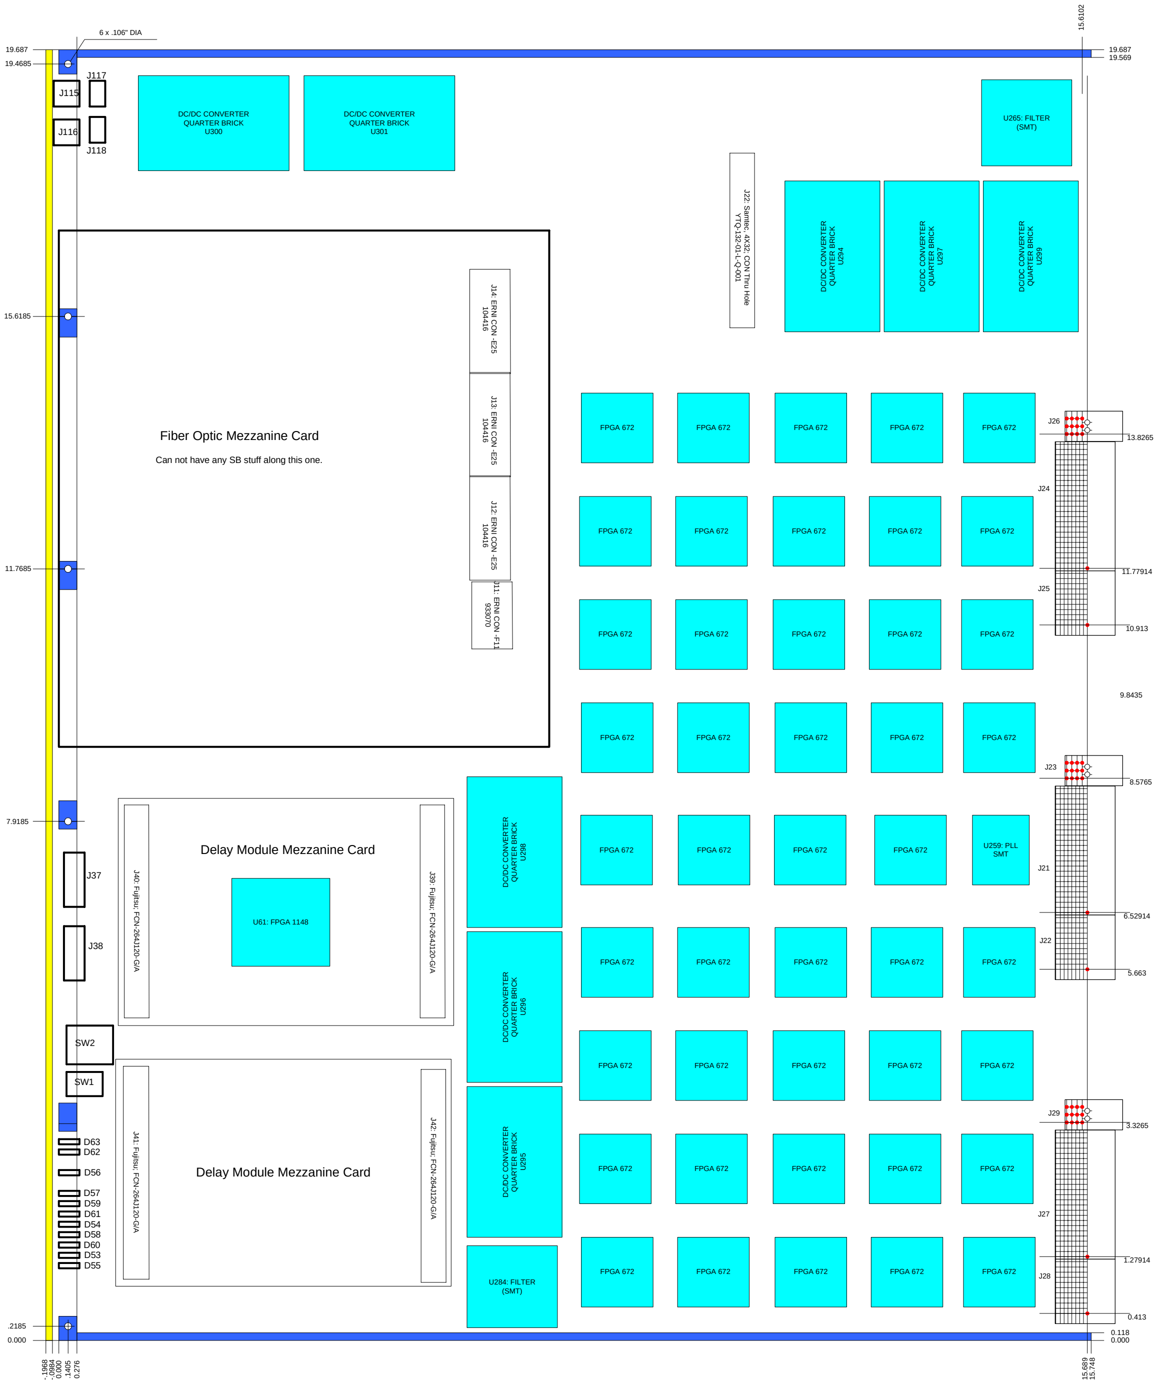

Figure 21 Station Board layout diagram. NRAO's FORM/DTS receiver module is a mezzanine card, as are the Delay Modules, and the CMIB (PC/104+ and PCMC). One output connector is for 1.024 Gbps wafer connections, and 2 connectors are for VSI I/O. A minimum of 2 connectors is required for power delivery and IDs.

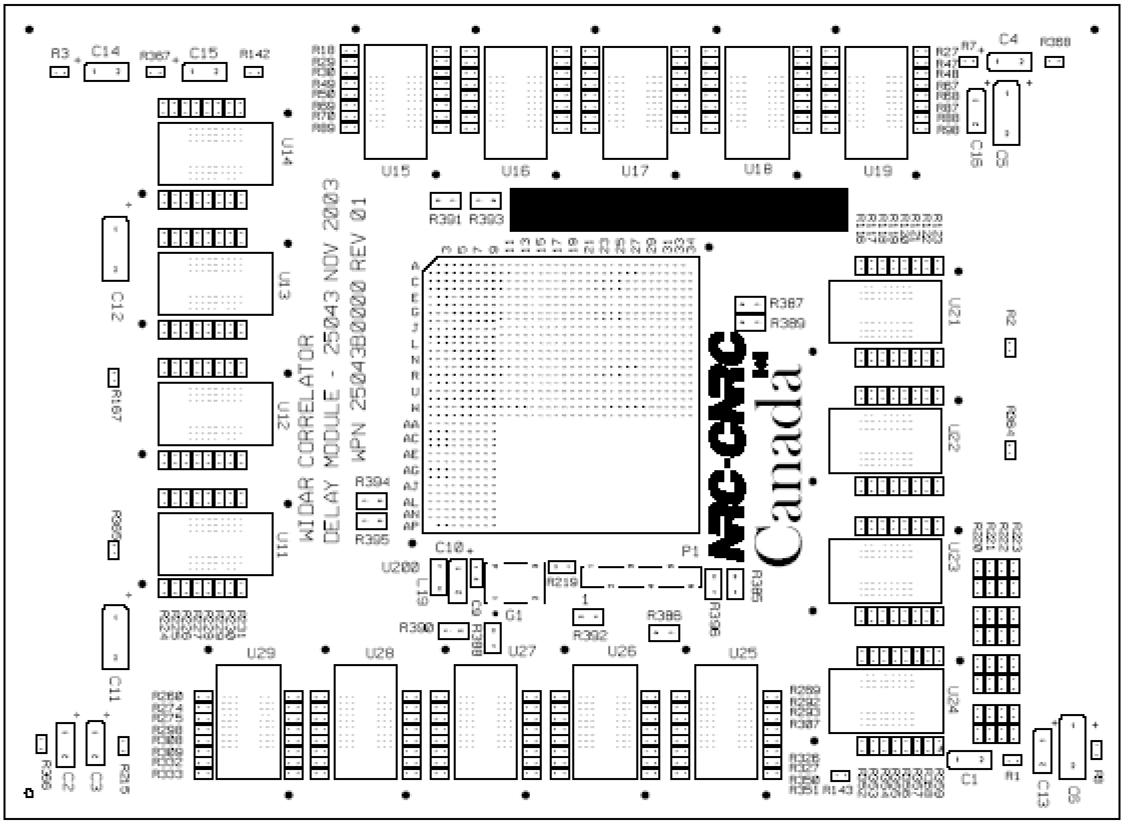

Figure 22 Delay Module mezzanine card. Two of these modules plug into the Station Board.

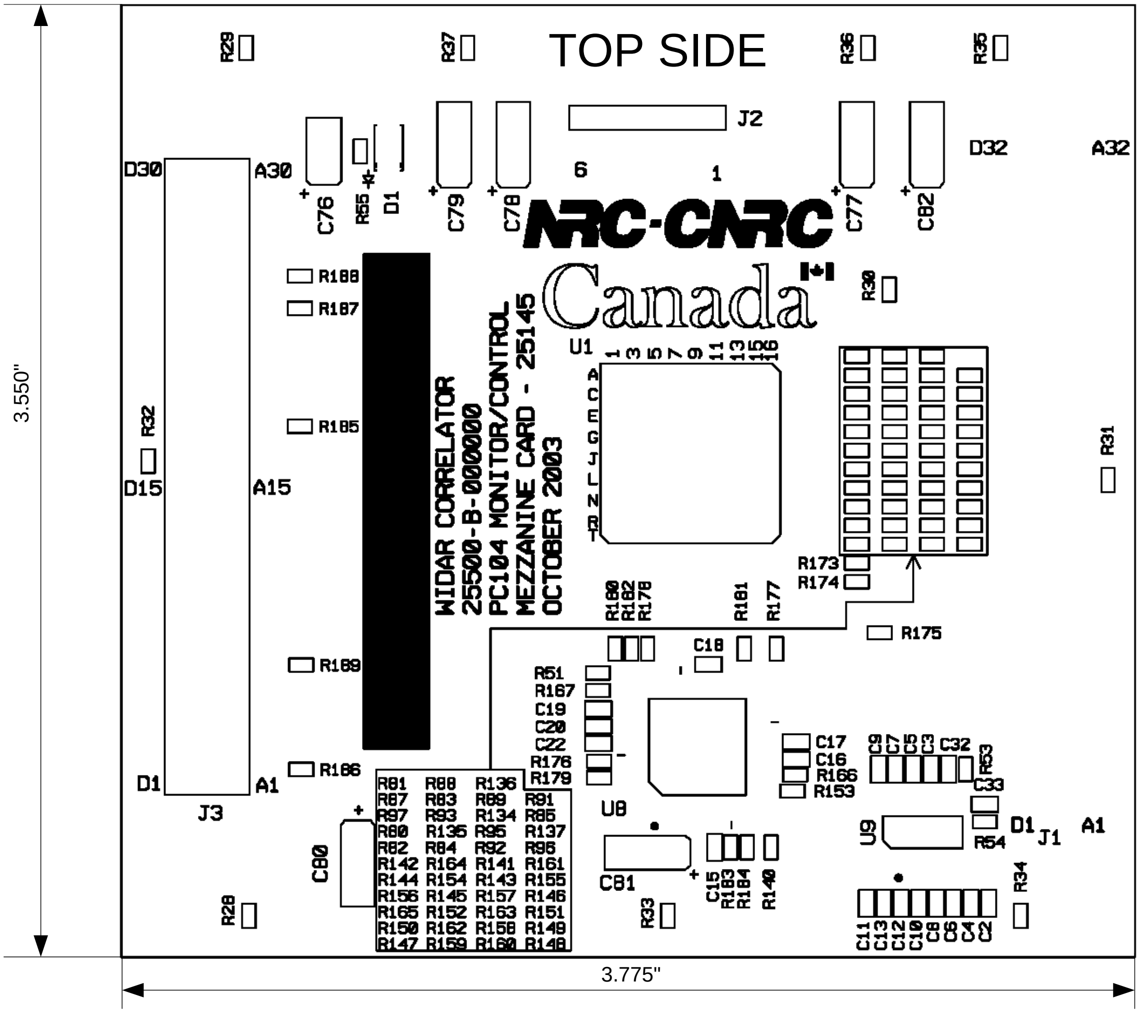

Figure 23 PCMC mezzanine card. This card contains the PCI bus-to-synchronous "MCB" bus interface FPGA and an A/D converter for temperature and voltage monitoring. It is sandwiched between the CPU PC/104+ COTS board and the motherboard.

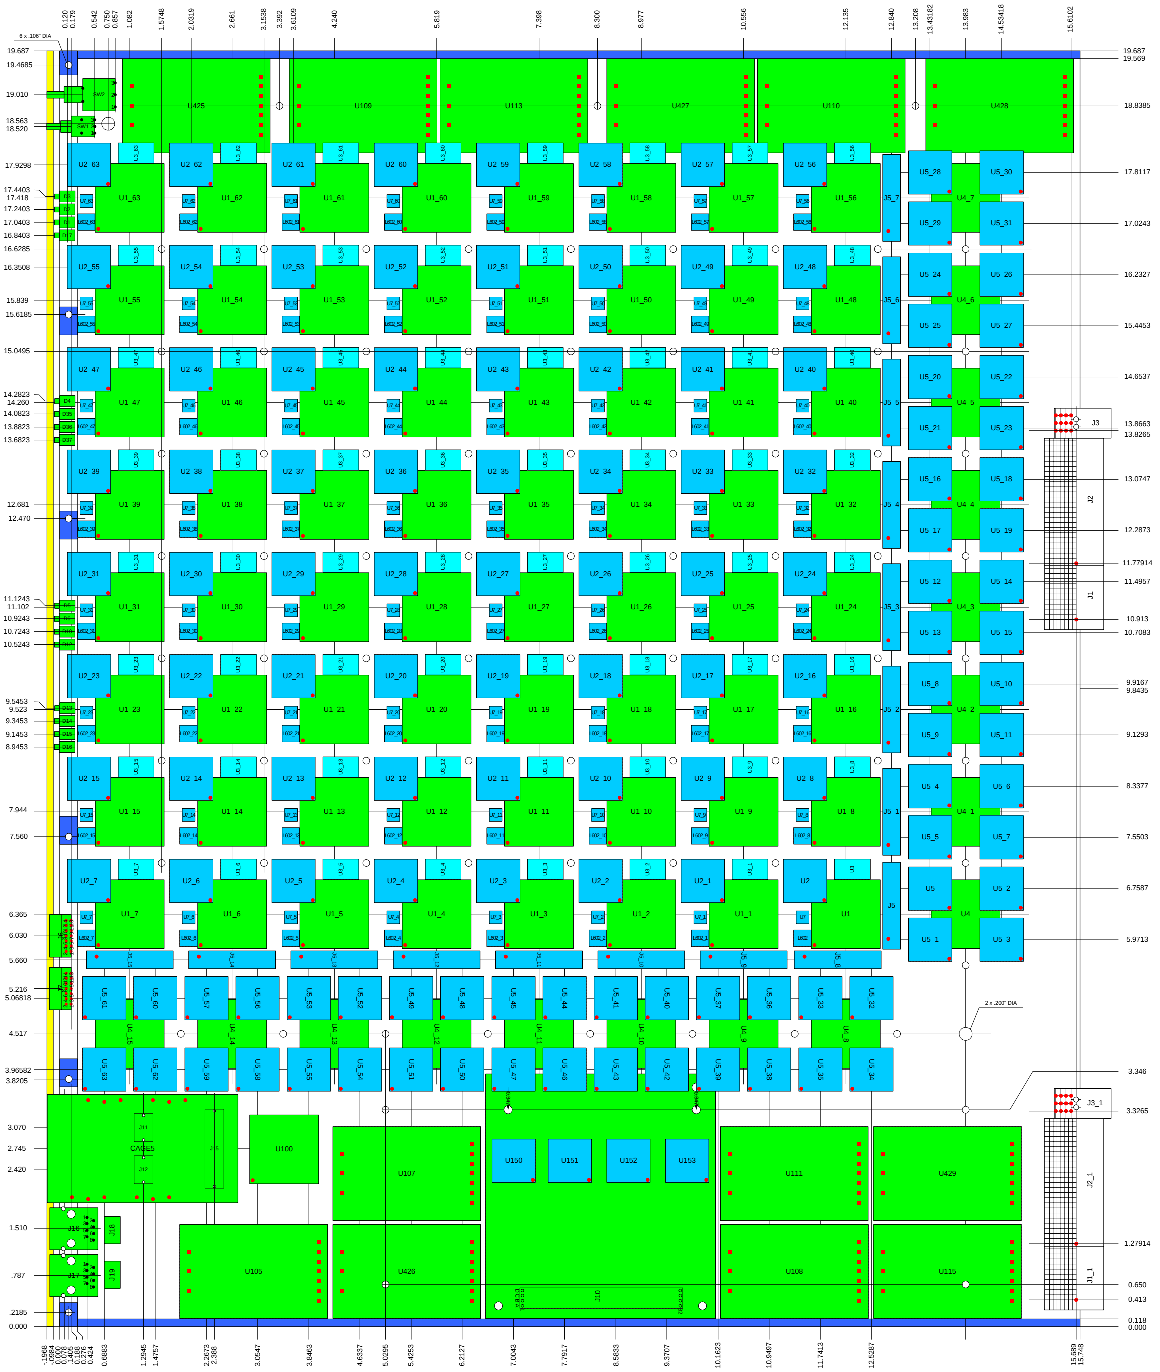

Figure 24 Baseline Board preliminary layout diagram. This board is the most complex board in the system because of number of devices, and because it requires double-sided BGA surface mounting.

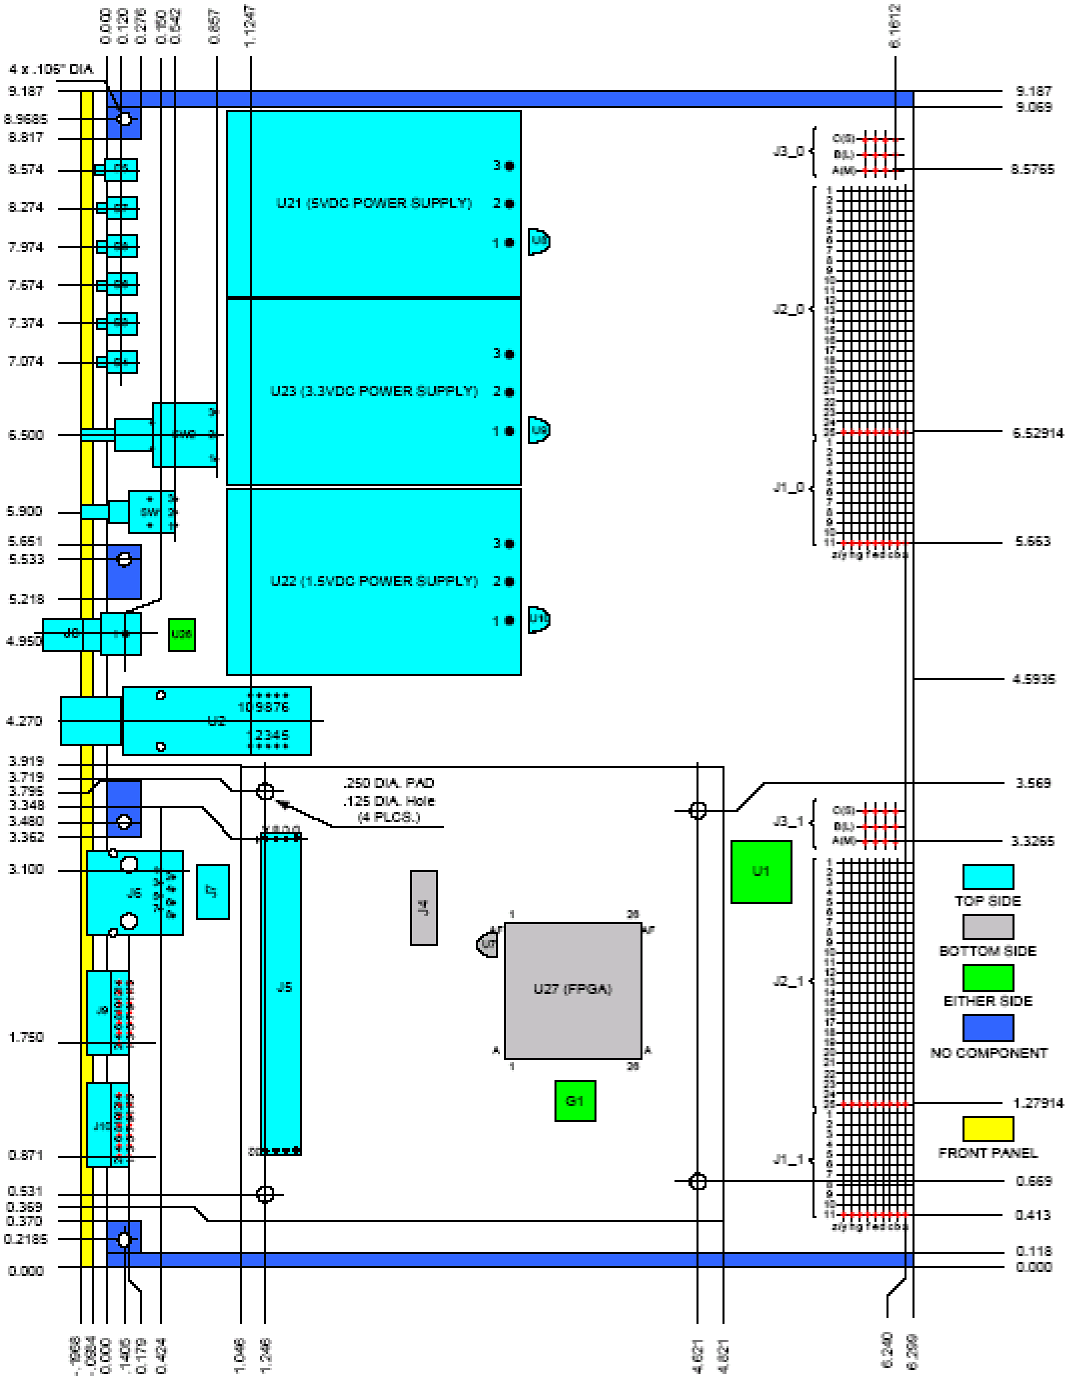

Figure 25 Timecode Generator Board layout diagram. There are two of these in the system to eliminate single point of failure. Both boards are fed the same signals, and downstream Station Boards can select one of three TIMECODEs from each board.

References

[1] Carlson, B., A Proposed WIDAR Correlator for the Expansion Very Large Array Project: Discussion of Capabilities, Implementation, and Signal Processing, NRC-EVLA Memo #001, May 18, 2000.

[2] Carlson, B., Fort, D., Morgan, T., Rowen, B., EVLA Project Book, Chapter 8: CORRELATOR (http://www.aoc.nrao.edu/evla/admin/projbook//chap8.pdf)

[3] Carlson, B., Two Correlators for the Price of One: How a VLBA Correlator Could Fit Within the Proposed 40-Station WIDAR EVLA Correlator, NRC-EVLA Memo #006, September 28, 2000.

[4] Heng, Z., INTERFACE CONTROL DOCUMENT: TIMECODE and Clock External Interface Specification, ICD Document: A25022N0090, Revision 1.0, February 22, 2005.

[5] Carlson, B., PROTOCOL SPECIFICATION: HM Gbps Cable Signaling Specification, Protocol Document: A25022N0041, Revision 1.1, February 24, 2005.

[6] Carlson, B., Requirements for 8-bit Processing in the Proposed WIDAR Correlator for the EVLA, NRC-EVLA Memo #010, January 29, 2001.

[7] Carlson, B., Refined EVLA WIDAR Correlator Architecture, NRC-EVLA Memo #014, October 2, 2001.

[8] Heng, Z., INTERFACE CONTROL DOCUMENT: HM Gbps Cable Physical Specification, ICD Document: A25022N0040, Revision Draft, April 13, 2004.

[9] Fort, D., Ljusic, Z., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Station Board Input Chip, RFS Document: A25054N0000, Revision 1.0, March 2005.

[10] Ljusic, Z., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Delay Module, RFS Document: A25043N0000, Revision 1.0, March 1, 2005.

[11] Fort, D., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Station Board Wideband Correlator FPGA, RFS Document: A25051N0000, Revision 1.0, June 2005.

[12] Fort, D., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Station Board Filter Chip, RFS Document: A25044N0000, Revision 1.0, January 15, 2005.

[13] Ljusic, Z., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Station Board Output Chip, RFS Document: A25055N0000, Revision 1.0, May 2005.

[14] Heng, Z., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Station Board Timing Chip, RFS Document: A25052N0000, Revision DRAFT, February 2004.

[15] Heng, Z., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: PC/104 Monitor/Control Mezzanine Card, RFS Document: A25145N0000, Revision 1.1, May 22, 2003.

[16] EVLA Correlator Conceptual Design Review, Detailed Technical Points, B. Carlson (http://www.aoc.nrao.edu/evla/admin/reviews/correlatorcodr/detechpoints_bcarlson.pdf)

[17] Carlson, B., Summary of Discussions Held During the July 10-14, 2000 Workweek in Socorro Regarding the EVLA-WIDAR Correlator, NRC-EVLA Memo #005, August 22, 2000.

[18] Carlson, B., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: Recirculation Controller FPGA, RFS Document: A25090N0000, Revision 2.2, September 10, 2004.

[19] Carlson, B., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: LTA Controller FPGA, RFS Document: A25091N0000, Revision 1.1, February 16, 2005.

[20] Carlson, B., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: GBit Ethernet Chip, RFS Document: A25092N0000, Revision 1.0, January 20, 2005.

[21] Carlson, B., TEST AND VERIFICATION PLAN: EVLA Correlator Chip, TVP Document: A25082N0001, Revision DRAFT, August 26, 2004.

[22] Smegal, R., TEST AND VERIFICATION PLAN: EVLA Correlator Chip Functional Test Cases, TVP Document: A25082N0002, Revision DRAFT 2.0, January 14, 2005.

[23] Carlson, B., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: EVLA Correlator Chip, RFS Document: A25082N0000, Revision 2.4, February 15, 2005.

[24] Vrcic, S., Rowen, B., The EVLA Correlator Software Architecture, RFS Document: A25200N0001, Revision 1.1, February 17, 2004.

[25] Vrcic, S., PROTOCOL SPECIFICATION: EVLA Correlator Monitor & Control Virtual Correlator Interface (VCI), RFS Document: A25201N0000, Revision 2.6, March 18, 2005.

[26] Vrcic, S., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: The EVLA Correlator Software; Master Correlator Control Computer (MCCC), RFS Document: A25202N0001, Revision 2.1, December 6, 2004.

[27] Vrcic, S., A WIDAR Correlator for the Expansion Very Large Array: Technical, Functional, and Performance Discussion Document; MCCC Software; Generating Baseline Board Configuration Based on the Configuration of Station Boards, a Proposal, NRC-EVLA Memo #018, February 2004.

[28] Morgan, T., USER’S GUIDE AND MAINTENANCE MANUAL; EVLA WIDAR Correlator Backend System, A25250N0001, Revision 1.1, May 6, 2005.

[29] Webber, R., Halman, M., TEST AND VERIFICATION RESULTS: EVLA Correlator Rack: Thermal Mock-up Tests, TVR Document: A25031N0003, Revision 1.0, April 29, 2005.

[30] Halman, M., EVLA Correlator Sub-rack and Duct Assembly Drawing, D25015A0020, June 13, 2005.

[31] Carlson, B., TEST AND VERIFICATION REPORT: EVLA Correlator Reliability Report #1, TVR Document: A25010N0003, Revision Report#1, June 6, 2005.

[32] Webber, R., REQUIREMENTS AND FUNCTIONAL SPECIFICATION: EVLA Correlator Room Requirements, RFS Document: A25012N000, Revision DRAFT, May 26, 2005.

[33] Carlson, B., INTERFACE CONTROL DOCUMENT: EVLA Correlator System Numbering Plan, ICD Document: A25010N0002, Revision 1.1, February 24, 2005.

Acknowledgements

The system documented in this memo is the result of the work of many people over a period of several years. The following contributions are of note:

Dave Fort: Signal processing, Station Board functionality, Filter Chip design, FPGA design, system design, project management.
Zoran Ljusic: Delay Module, Station Board FPGA design, Station Board design, system design.
Zhang Heng: TGB, PCMC, Fanout Board, Feedthru Backplane, Baseline Board designs. FPGA design, system design.
Sonja Vrcic: Software architecture design, MCCC software/MAPPER design.
Ralph Webber: Rack design and development, system infrastructure specification and development, power plant specifications, room requirements, laboratory setup.
Mark Halman: Rack/sub-rack design and development. Thermal test rack construction and test.
Peter Dewdney: Project management, signal processing development, system development.
Brent Carlson: Signal processing, architecture development, system design, Baseline Board FPGA/chip design, project management.
Amy Fink: Project management, engineering assistance, prototype component acquisition.
Donna Morgan: Project management assistance.
Rick Smegal: Correlator chip validation and verification test suite development and testing.

And, from NRAO:

Bruce Rowen: CMIB specification, device driver design, RTOS/Linux testing, specification, development, XML protocol.
Tom Morgan: Backend specification, design, and development.
Kevin Ryan: Jave GUI design and development.

I would also like to thank the following individuals for their ideas and contributions to the project:

Rick Perley, Bryan Anderson, Craig Walker, Michael Rupen, Barry Clark, Ken Sowinski, John Wu, Don Matson.

…and others who have provided direct or indirect input and contributed ideas to the development of this system.

Index