

8 CORRELATOR

*Brent Carlson, Dave Fort
Tom Morgan, Bruce Rowen
Last changed 2004-Nov-15*

Revision History:

2001-June-06: Initial release.

2001-July-17: Updates from M. Rupen, J. Romney comments. Add milestone table, table of “impacts and interfaces” to the rest of the system, and risk assessment table. Add clarification text to many sub-sections. Add sub-section on sub-band stitching. Upgrade the correlator layout diagram. Revise module costs to include refined pricing and circuit boards and cables for 32 stations, with racks for 40 stations.

2001-August-14: Minor revisions based on additional Rupen and Romney comments. First draft full release.

2002-February-1: Overhaul based on new Gigabit Ethernet output and Backend configuration. Add more M&C S/W and H/W details. Many refinements to many sections based on design refinements over the last several months. Due to cost increases, NRC will now only pay for and install a 32-station correlator (i.e. not racks for 40 stations).

2002-May-22: Add Correlator Backend discussion to Summary and Introduction. Insert Correlator Backend Requirements and Design sections **8.4** and **8.5**.

2003-Aug-15: General update includes: upgraded sharp-cutoff FIR capability; explicitly state that the correlator is capable of up to 4 million spectral channels per baseline with recirculation; recirculation on 4 streams is now a requirement; specification of stream statistics capabilities; one phase-cal extractor in every FIR chip; solidification of Correlator Chip capabilities; solidification of delay capabilities including narrow baseband very fine delay tracking; VSI_H I/O on the Station Board for VLBI-ready capability; the software operating system choice is Linux at all levels; updated risk assessment table.

2003-Aug-25: Update the section on RFI mitigation to include wideband and sub-band data valid blanking requirements.

2004-Nov-15: Major upgrade to take into account refined design information. No reduction in system functionality or performance, except the correlator chip uses a 3-level rather than 5-level phase rotator, reducing sensitivity by an additional 1.75%. Enhanced functionality may be possible with additional funds (e.g. “R2” recirculation on 8 streams).

Summary

The delivered system will be a 32-station correlator, however the scaleable architecture supports up to 256 stations in 8-station increments. Each station will be capable of handling a total bandwidth of 16 GHz, arranged as 8, 2 GHz **basebands**. The correlator contains dedicated hardware (lags) for 16,384 spectral channels per baseline at the widest bandwidths and uses “**recirculation**” to provide up to 4 million spectral channels per baseline at narrow(er)-bandwidths (or wide bandwidths with sensitivity losses). The system can flexibly use and deploy spectral channel resources within internally generated and user defined **digital sub-bands**. High performance pulsar processing capabilities are an integral part of the design. The system will be delivered with a total 1 GHz phased-VLA capability—enabling up to 4 Gbits/sec of VLBI recording bandwidth at 2 bits/sample. The architecture supports phasing all 16 GHz of bandwidth and doing so is strictly a cost issue. The **Correlator Backend (CBE)** consists of a parallel cluster of commodity computers with high-speed interconnects to the correlator and the **end-to-end (e2e)** image processing and archive system. The CBE will be scalable in order to grow with increasing EVLA observational demands and correlator output data volumes, and flexible enough to handle all specified correlator operational modes. The CBE will use standard network communications hardware and software and will not rely on specialized vendor-specific implementations. In this chapter, correlator performance specifications are outlined and a reasonably complete design is presented that will meet the specifications. The principal performance specifications for the correlator are shown in Table 8-1. Development milestones are shown in Table 8-2.

Table 8-1 EVLA Correlator Principal Performance Specifications.

No. of stations (antennas) (Sec. 8.2.1)	32 (architecture supports up to 256).
Max spectral channels/baseline @ max bandwidth of 2 x 8 GHz = 16 GHz (Sec. 8.2.2, 8.2.8)	16,384 (more with “wideband recirculation” and sensitivity losses).
Max spectral channels/cross-correlation with recirculation (Sec. 8.2.2, 8.2.8)	262,144 (total 4 million channels per baseline)
Polarization products (Sec. 8.2.3)	1, 2, or 4
No. of basebands/antenna (Sec. 8.2.4)	8 x 2 GHz each (more with narrower bandwidths)
Quantization (Sec. 8.2.10)	1, 2, 3, 4, or 8-bit initial quantization; 4 or 7-bit re-quantization after sub-band filter.
Correlator efficiency (Sec. 8.2.10)	~93.2% (4-bit initial quantization, 4-bit re-quantization, 3-level fringe rotation)
No. of sub-bands per baseband (Sec. 8.2.6)	16 (provision for up to 18 for “N+1” redundancy).
Sub-band bandwidth (Sec. 8.2.8)	128 MHz, 64 MHz, 32 MHz, ..., 31.25 kHz (multi-stage filter). Each sub-band’s width and position can be set independently of any other sub-band.
Sub-band tuning (Sec. 8.2.9)	Each sub-band should remain within an appropriate integer slot to minimize band-edge SNR loss. E.g. a 128 MHz sub-band should be within 1 of 16 equally spaced slots in a 2 GHz band. Greater tuning flexibility at narrower bandwidths is possible.
Spectral dynamic range (Sec. 8.2.10)	(Initial quantization) 3-bit: ~44dB; 4-bit: ~50dB; 8-bit: ~58dB. [Test: 2 “bunches” of 4 tones/bunch, each “bunch” contained within one sub-band (128 MHz); 99% tone (interference) power; ideal samplers; dynamic range measured outside sub-bands containing interference.] With 2 tones only, results are ~10dB, ~2dB, ~4dB (respectively) worse.
Auto-correlations (Sec. 8.2.20)	Wideband (4x2 GHz pairs): 4 products of 1024 spectral channels each, SNR loss =4. Sub-band: 16,384 total spectral channels per station (widest sub-band), no SNR loss.
Pulsar processing (Sec. 8.2.14)	2 banks of 1000 time bins each/baseline. Up to 65,536 bins/baseline with software accumulation. Min. bin width: ~200 µsec (all spectral channels) ~15 µsec (64 spectral channels/sub-band/baseline). Also, pulsar gating with one timer+multi-gate generator per 2 GHz baseband.
Min. dump period (initial installed configuration) (Sec. 8.2.15)	100 milliseconds (all spectral channels). Faster with more CBE computers and/or fewer channels and/or baselines.
CBE input aggregate bandwidth	1.6 Gbytes/sec
CBE output aggregate bandwidth to e2e	25 Mbytes/sec
Max. dump period (Sec. 8.2.15)	Unlimited (within the Backend)
Maximum baseline (Sec. 8.2.16)	25,000 km with 0.5c FOTS transmission velocity (0.25 sec total delay buffer).
Sub-arrays (Sec. 8.2.18)	Cross-correlation: unlimited. Phased-VLA: 5 sub-arrays with “antenna grouping” granularity of 4 antennas.
Phased-VLA (Sec. 8.2.19)	8 digitally-phased sub-bands; architecture supports phasing all sub-bands. Simultaneous operation with interferometer modes using same array phase-center.
VLBI (Sec. 8.2.21)	VLBI-ready. Requires additional software and VSI breakout connector/PCB
Interference mitigation (Sec. 8.2.23)	Post-corr. Temporal/spectral excision—narrowband interference modulation robust. Possibly provision for post-correlation interference cancellation. Wideband and sub-band data valid blanking. Over-range detection on sub-band re-quantizers.

Table 8-2 EVLA Correlator Development Milestones

Milestone	Approximate Date	Notes
Conceptual Design Review (CoDR)	November, 2001	Architecture/features review, specifications/design freeze.
Correlator Backend (CBE) 4 node clus	Q2, 2002	Minimal configuration for initial prototyping.
CBE 8+ node test cluster	Q3, 2002	Minimal configuration for functional testing.
CBE full functionality	Q4, 2003	Ready for system test
Preliminary Design Review (PDR)	Q2/Q3, 2005	Prototype design ready, review before proto. construction.
CBE earliest connect to corr h/w	Q4, 2005	First live testing in Penticon.

On-the-sky test at VLA	Q2, 2006	Requires 4 dedicated new antennas, calibration/closure tests.
Critical design review (CDR)	Q2/Q3, 2006	Review before full production.
Begin installation at VLA (off-line)	Q2, 2006	Racks and cables: new correlator room required.
Limited production boards at VLA	Q1, 2007	16 additional Station Boards and 16 Baseline Boards. Earliest possible start of installed correlator testing. 8 station 8 GHz configuration.
Begin full installation at VLA	Q3, 2007	Begin installation of production boards
Earliest possible “shared-risk” science	Q2/Q3, 2008	Limited production boards. 8 station, 8 GHz configuration.
Correlator commissioning	Q2/Q3, 2008	Full observational mode, no apparent bugs.
Project complete	Q2, 2009	Scheduled NRC support no longer required.

8.1 Introduction

The EVLA correlator design is based on the WIDAR concept (Carlson, IEE 2000) (Carlson, Memo# 001) (Carlson, Memo# 014) where wide (2 GHz) bands are sampled, split into smaller sub-bands with digital filters, and then correlated. A key anti-aliasing technique, along with stable calculable digital filter characteristics, allow the sub-bands to be seamlessly “stitched” together to yield the wideband cross-power spectrum. Using this technique it is possible to correlate data efficiently so that about an order-of-magnitude more spectral channels can be provided compared to what other time-domain parallelization techniques can yield. A design requirement for the EVLA is to provide 16,384 spectral channels per baseline in wideband modes, with more spectral channels available using “recirculation”. Digital sub-banding has the additional benefit of increasing the flexibility of the correlator so that only those spectral regions of interest need use correlator resources. An ‘XF’ correlator has been chosen primarily to minimize the station hardware-to-baseline hardware bandwidth/cabling requirements—a significant consideration for a correlator system of this size.

The design of the Correlator Backend (CBE) is based on the requirement to not have to rely on a specialized high-speed interconnect fabric amongst Backend computers to perform required processing, and to be able to use commodity computers (PCs) and network hardware and software to meet performance goals. The data delivery network from the correlator to the Backend is designed so that each computer has all of the information (lag data) needed for processing one or more baselines—Backend computers do not have to exchange additional high-speed information. By designing the data delivery network in this fashion all processing is thus 100% parallel. This provides for linear scalability of performance with the addition of processing nodes. Use of multi-CPU processors in the nodes will provide for sufficient compute power and flexibility to handle critical real-time input demands as well as data processing, formatting and internal monitor and control activities. Inter-node communications will be limited to monitor and control messages that will be handled by message passing middleware. *Intra*-node communications will be handled by message passing middleware, and in the case of observational data, shared memory.

8.2 Specifications

8.2.1 Number of Stations (Antennas)

The installation will include a full population of 32 stations. The architecture supports up to 256 stations in 8-station increments. Final installation may be 40 stations to support EVLA “Phase-II” requirements. Thus, supporting infrastructure should plan for a 40-station correlator. Full (16 GHz) bandwidth stations can alternatively be configured for 2 stations at 4 GHz bandwidth, or 4 stations at 1 GHz bandwidth, with a factor of 4 and 16 reduction in number of spectral channels per baseline respectively. The architecture supports implementing this tradeoff dynamically, but requires additional front-end switching hardware not included in the delivered system to realize this potential.

8.2.2 Spectral Channel Capability

Dedicated correlator resources (lags) for 16,384 spectral channels/baseline at the widest bandwidths are available. Spectral channels can be flexibly deployed to desired sub-bands/basebands. “Recirculation” provides a maximum of 262,144 spectral channels per *cross-correlation* on 4 sampled data streams (sub-bands) that can be different in each sub-band correlator. Provision for recirculation on all 8 sampled data streams may be made and this option (referred to as “**R2**”) may be chosen but incurs additional production cost. Recirculation works by time-multiplexing the acquisition of correlator lags using synthesized lag delays in a memory buffer. The amount of time multiplexing is known as the **recirculation factor**. In narrow(er)-band modes where the bandwidth reduction is the same as the recirculation factor, no sensitivity degradation is realized in the cross-power spectrum. If the recirculation factor is *greater* than the bandwidth reduction, there is a *root(recirculation factor/bandwidth reduction factor)* decrease in sensitivity. Recirculation can be used at maximum sub-band bandwidth (128 MHz) with the above indicated sensitivity reduction (referred to as **wideband recirculation**). When recirculation is used, the correlator dump time and/or minimum phase-bin time is increased since it is necessary to obtain at least one pass of all lag data in each dump to produce a proper spectrum. The time increase factor is the same as the recirculation factor.

8.2.3 Polarization

Basebands can be flexibly arranged as combinations of dual-polarization **baseband pairs** and single-polarization basebands (subject to antenna system flexibility). 1, 2, or 4 polarization products can be correlated and these are selectable on a baseband/sub-band basis.

8.2.4 Sampled Baseband Capacity

Each “station input” can handle 8, 2.048 GHz basebands sampled at 4.096 Gs/s. More sampled bands—up to 128 per station input—*could* be handled if they had less bandwidth each. This could be useful if it is desired to process more (narrower) sampled basebands (for example, to avoid regions of extreme RFI), but this is currently not an EVLA requirement. The correlator can flexibly handle various combinations and numbers of sampled bands provided sample rates are properly related. For example “native” EVLA antennas with 16 GHz of bandwidth each could be correlated with “foreign” antennas with 1 GHz of total bandwidth in 8, 128 MHz chunks (or 16, 64 MHz chunks etc). For these correlations to make sense, the digital EVLA sub-bands must overlap in frequency and be the same bandwidth as the foreign antennas’ basebands. The overlap does not have to be exact, since the difference can be removed with the correlator’s fringe rotators—with the expected reduction in cross-power bandwidth. Simultaneous with EVLA-foreign antenna correlations, can be full-bandwidth EVLA-EVLA correlations.

8.2.5 Baseband Tuning

Basebands can be at any “sky” frequencies and any restrictions are governed entirely by antenna LO system flexibility.

8.2.6 Digital Sub-band Capability

The correlator has provision for up to 18 digital **FIR (Finite Impulse Response)** filters—implemented in a **Filter Chip**—for each 2 GHz baseband input. Typically, one of the sub-bands would be used for receiver switching noise diode measurements (i.e. so it can use a sub-band of the baseband with no [time-variable] interference in it for system noise temperature calibration). The Filter Chip consists of up to 4 stages of filtering and, depending on configuration, can provide sharp-cutoff sub-bands as narrow as 31.25 kHz. Thus, “radar-mode” capability is effectively built into each Filter Chip. Refer to section 8.2.12 for more detailed information.

The *delivered* correlator will be populated with 16 **sub-band correlators** and each of these can connect to any of the 18 filter outputs. Each sub-band correlator correlates all basebands of all baselines for a particular sub-band. There is provision for up to 18 sub-band correlators. Each sub-band correlator can connect to any of the 18 Filter Chip outputs (per baseband) so that “N+1” redundant capability could be achieved (“N+1” generally refers to a system’s capability of losing one module with no loss of performance or data). Additionally, provision will be made so that each *sub-band* could be on a different delay-center on the sky to support multi-beaming *within* a baseband. The maximum delay-center offset from the baseband’s delay center is currently 32 μ sec for the EVLA and 8192 samples for VLBI.

8.2.7 Sub-band Stitching

Adjacent sub-bands can be seamlessly “stitched” together with a maximum sensitivity loss of a factor of 2 at the **sub-band boundary**. The rate of reduction in sensitivity away from the boundary depends on the “steepness” of the filter transition band. (Typically, with a flat passband –6 dB cutoff filter and 511 taps, the sensitivity loss is less than 20%, 2 MHz away from the sub-band boundary for a 128 MHz passband. This includes sensitivity loss effects from re-quantization and fringe rotation.) Stitching is performed by applying the total power measurements obtained in the Filter Chips *before re-quantization* and by applying calculated digital filter bandshape corrections (Carlson Memo# 001, Carlson IEE 2000). Since the filter is applied with the LO offset in place, and this is removed in the cross-power spectrum result, baseline-based filter bandshape corrections should be applied that include the effective baseline LO offset as it affects the filter amplitudes. Depending on transition-band steepness, this special consideration is normally only required if the LO offset is greater than $\sim 1/10^{\text{th}}$ of the spectral-channel bin width. Each filter’s total power measurement (before re-quantization) can only be used properly if the total power gain of each filter is known. This gain is calculable, but also depends on tap-weight scaling (i.e. the scaling of floating-point tap weights to integer bits used in the filter) that should (effectively) be relative to some common reference value for all filters on every Station Board. Depending on sub-band roll-off and narrowband signal strength in the proximity of the sub-band boundary, stitching may require the use of adjacent sub-bands’ spectral points and careful windowing operations. Initial quantizer statistics and re-quantizer statistics are obtained in the correlator and will be required for accurate data normalization.

8.2.8 Sub-band Bandwidth

Each of the 18 general-purpose digital filters can be configured for an output bandwidth starting at 128 MHz and decreasing in powers of 2 down to 31.25 kHz. Sharp cutoff filters are possible at all bandwidths by using some or all of the 4 stages of the Filter Chip. *All* filters are independently configurable in bandwidth and placement within the baseband. Refer to section 8.2.12 for a more detailed description of the Filter Chip.

8.2.9 Sub-band Tuning Flexibility

Digital pass-bands can be placed anywhere within integer “sub-slots” corresponding to the sub-band (slot) width. For example, if the (stage 1) filter has a 1/64 bandpass, then the filter can be placed in any of the 64 evenly spaced slots in the band. More tuning flexibility is provided by stage 2 of the Filter Chip, operating on up to 128 MHz of stage 1 output. Refer to section 8.2.12.

8.2.10 Sample Word Sizes and Correlator Efficiency

The initial baseband sampled word size can be any one of 8, 4, 3, 2, or 1 bits. Each sampled baseband in each antenna could have a different word size as long as the total digital transmission bandwidth does not exceed the fiber-optic transmission system bandwidth. The correlator supports 4-bit initial quantizer word sizes, but for cost reasons, the antennas will only deliver 3 bits at 2 GHz baseband bandwidths. Refer to Table 8-1 and (Carlson, Memo# 009) for spectral dynamic range estimates. The correlator supports 8-bit initial sampling, but if used, only $\frac{1}{2}$ the baseband bandwidth will be available since the sample word width has doubled. (Each baseband is independently configurable in sample word width.) (N.B. because of frequency shifting, it is possible to use time-interleaved samplers since spectral by-products generated from amplitude mismatches do not show up in the correlator cross-power spectrum.)

After digital FIR filtering, the correlator re-samples the data to 4 bits. Alternatively, in high SNR, high dynamic range regions of the spectrum, the correlator can re-sample and correlate 7 bits (Carlson, Memo# 010). If 7-bit correlation is used, then $\frac{1}{2}$ the spectral channels and $\frac{1}{2}$ the sub-band bandwidth is available (because of internal correlator data-path routing limitations). Choice of re-sampling word size can be done on a per sub-band basis. Also, the re-sampling word size does not depend on the initial sampler word size (and vice versa).

Three-bit initial sampling and 4-bit re-sampling, along with 3-level correlator fringe rotation loss, results in a correlator efficiency of about 91% (Carlson, Memo# 011). (Four-bit sampling is $\sim 98.5\%$ efficient, 3-bit is $\sim 96.5\%$ efficient, and 3-level fringe rotation is $\sim 96\%$ efficient (Carlson, Memo# 002). Eight-bit sampling is very close to 100% efficient and thus has a negligible sensitivity loss.) For spectral dynamic range refer to Table 8-1.

8.2.11 Correlator Chip

The planned Correlator Chip contains 2048 complex-lags, arranged as 16, 128 complex-lag correlators. Adjacent internal complex-lag correlators can be concatenated together. There is no provision for directly concatenating Correlator Chips. Each accumulator is 23 bits long and is *not* truncated for high dynamic range correlation. 23-bit accumulators have a maximum integration time of 500 microseconds, however shorter readout times are needed to support recirculation and narrow phase binning. Lag-based, quantized-phase fringe stopping is performed with a planned 3-level fringe rotator (Carlson, 1999) (Carlson, Memo# 002). The maximum chip data rate is 256 Ms/s.

8.2.12 Digital Filter Chip

The digital filter chip consists of the following functional blocks (operating sequentially on the data from input to output):

1. **Sub-band delay** line of up to 32 microseconds. This delay line may be used for sub-band multi-beaming (Carlson, EVAL-NRC Memo# 014) or it can be (effectively) by-passed.
2. **Stage 1 filter.** This is a 512-tap, 16-phase poly-phase FIR filter with 32 taps per phase, and an integrated 16x16, 4-bit wide cross-bar switch. It operates on 4-bit or 8-bit input data (but only 256 taps on 8-bit data). The conversion of 1, 2 or 3-bit samples to 4-bit samples is done in another chip. For (VLBI-mode) very fine ($\pm 1/32^{\text{nd}}$ of a sample) narrowband delay tracking, each phase of the filter is loaded with sub-sample delay coefficients, and the delay tracker selects the appropriate phase in real time with no blanking as delay changes. The maximum output bandwidth from this filter is 128 MHz, and the practical minimum output bandwidth is ~ 16 MHz. The output of this filter is 16 bits, and the output can go to stage 2, or the final re-quantizer. Sub-band filtering using this stage should stay within integer slots to minimize the SNR degradation region at the edges of the filter.
3. **Stage 2 filter** with 64, 128, 256 or 512 taps, depending on the decimation factor of 2, 4, 8 or 16 respectively, so that the same *relative* filter cutoff steepness is maintained independent of bandwidth. The maximum input bandwidth of this filter is 128 MHz. This filter contains an integrated digital single-sideband mixer so that output sub-bands are finely tunable (with a 32-bit frequency synthesizer and high dynamic range mixer) anywhere within the input bandwidth. This mixer may be used or by-passed and it is also used for very fine delay tracking when the final output bandwidth of the filter is very narrow. The output of this filter is 16 bits and may go to stage 3 or the final re-quantizer.
4. **Stage 3 filter** with 64 to 512 taps operating the same as stage 2 except that there is no single-sideband mixer. The maximum input bandwidth of this stage is 8 MHz. The output of this filter is 16 bits and may go to stage 4 or the final re-quantizer.
5. **Stage 4 filter** with 64 to 512 taps operating the same as stage 3. The maximum input bandwidth of this stage is 0.5 MHz. The minimum output bandwidth from this is 31.25 kHz in order that there are an integral number of output samples in 10 milliseconds.

Note that since stages 2, 3, and 4 operate on 16 bits, even when all stages are in use, there is only one re-quantization loss in the Filter Chip—that of the final re-quantizer.

6. **Pre-re-quantizer power meter**, with two accumulation bins for noise diode calibration. This power meter operates on the 16 bits out of one of the filter stages before re-quantization.
7. **Final re-quantizer.** Sixteen bits from one of the four filter stages is selected to be re-quantized to 4, 5, 6, 7 or 8 bits by this block. Note that the Baseline Board can only except 4 or 7-bit samples. For VLBI recording, 1 and 2-bit samples are produced in another (VSI) chip from the same samples that are sent to the Baseline Board. Over-range detection on input to the re-quantizer is used to flag data as invalid for burst interference suppression.
8. **Sideband flipper.** This flips the sign of every other sample to change the frequency sense of the output sub-band. This block can be enabled or disabled.
9. **Phase-cal extractor.** This is a tone extractor with a 32-bit frequency synthesizer that operates on the re-quantized data.
10. **State counters and power meter.** This block acquires re-quantizer output statistics. Refer to section 8.2.20.

11. **Delay line.** This block delays the output of the re-quantizer so that the outputs of all Filter Chips (that may have different output sub-band bandwidths) on the Station Board can be properly lined-up in time before correlation.

The final output of the Filter Chip is 4 bits wide, and contains the 4-bit or 7-bit (time-multiplexed) re-quantized data.

8.2.13 Radar Mode

Each Filter Chip will have the ability to output sharp-cutoff, narrow sub-bands required by radar mode and so “radar mode” is effectively no longer a separate mode of the correlator. However, radar processing requires 1 Hz resolution on ~30 kHz bandwidth while simultaneously obtaining reasonable spectral resolution on the wide 2 GHz baseband the 30 kHz is part of. One way this can be accomplished is by using two, 128-lag correlator cells within one sub-band correlator and recirculation x256 to get 1 Hz resolution on the 31.25 kHz, while using the rest of the correlator to correlate all of the required sub-bands that make up the 2 GHz baseband. Additionally, it will be possible for the CMIB to capture and readout filtered and re-quantized data for software processing. The maximum filter output bandwidth for which all data can be captured is TBD, but the minimum specification is to be able to capture all data (8-bit samples) at a 31.25 kHz bandwidth for one sub-band of each baseband. Currently, two 31.25 KHz 8-bit sub-bands can be captured and wider bandwidths with correspondingly fewer bits are also possible. This functionality replaces the front-panel radar-mode filter output connector previously specified.

8.2.14 Pulsar Processing

There are 2 banks of **1000** time bins each per baseline. One bank is active while the other bank is being downloaded to back-end computers. Alternatively, 1 bank of 2000 time bins can be used if correlator dead time while downloading data is acceptable. If all spectral channels are dumped, then the minimum bin width is ~200 μ sec; if only 64 spectral channels/sub-band/baseline are dumped, then the time bin can be as narrow as ~15 μ sec. Up to 65,536 bins/baseline can be accommodated with back-end computer software accumulation. Pulsar gating with one timer and multi-gate generator per 2 GHz baseband is available. The multi-gate generator can produce 16 pulsar gates with configurable delays relative to the timer epoch so that each sub-band can be gated “on” at different times to track different pulse arrival times at different frequencies.

8.2.15 Real-Time Data Output Performance

The real-time data output performance is governed by several factors. The correlator hardware itself has a very wideband data output pipeline so it is most likely that any performance limitations will be determined by the performance and configuration of the correlator’s back-end computers. The minimum dump period for all spectral channels if the extreme, highest-performance output pipeline is used (4, 1 Gbit/sec links) is ~2.6 msec. This is a dump rate of over 12 Gvis/sec in a 40-station correlator. The delivered system will have a pipeline—*out of the Baseline Boards*—capable of dumping all spectral channels every ~11 milliseconds. With the *planned* number of back-end computers (Figure 8-3) all spectral channels should be able to be dumped about every 100 milliseconds (~315 Mvis/sec in a 40-station correlator). If fewer spectral channels are dumped, then shorter dump times could be obtained. The maximum correlator hardware **LTA (Long Term Accumulator)** integration time is signal-characteristic dependent but is about 16 seconds for low SNR cross-correlations and 2 seconds for auto-correlations. Backend computers can integrate data for an arbitrarily long period of time.

8.2.16 Delay

The delivered correlator contains enough delay buffering for 0.25 seconds of delay and this translates into 25,000 km baselines if there is a $0.5c$ FOTS data transmission velocity over the same distance. The delay may increase in the production module depending on cost and availability of SDRAM memory chips. The delay may be increased in the future by replacing the Delay Module mezzanine card (on all Station Boards), if desired. The delay rate that the correlator can handle is essentially unlimited. Each baseband can have its own independent delay model and hence independent delay center on the sky. Precision, fully digital $\pm 1/32^{\text{nd}}$ of a sample delay tracking on 2 GHz basebands is a feature of the WIDAR architecture (Carlson, Memo# 007). There is no associated data blanking as the correlator tracks delay. WIDAR sub-sample delay tracking eliminates the need and uncertainty associated with sampler clock

phase modification. In addition, the Filter Chip provides the ability to finely track delay on narrower basebands (≤ 128 MHz) by loading each “phase” of the stage 1 poly-phase FIR filter with delay-interpolation coefficients, and seamlessly selecting the correct phase of the filter to implement the correct sub-sample delay in real time. This method eliminates the need to perform special delay functions on a baseline basis in the Correlator Chip, and provides $\pm 1/16^{\text{th}}$ sample of baseline delay tracking with virtually no restrictions on delay rate.

8.2.17 Doppler/Frequency Shift

The Correlator Chip contains digital complex phase-rotators with effectively no limitations in Doppler phase rate or artificial frequency shift. The rotators are driven by linear digital frequency synthesizers that can be updated every 10 msec. The fundamental limitation is the sub-band bandwidth, but it is suggested that the maximum phase rate not exceed $\frac{1}{2}$ the widest sub-band so that phase does not contribute to Correlator Chip heating (through fast toggling of CMOS transistors). Digital filter anti-aliasing requires offsetting each antenna’s Local Oscillator by a small amount. It is suggested that this be about 10 kHz, but tunable in 100 Hz steps for narrowband radar mode (Carlson, Memo# 005). There should be an adequate frequency shift between signals/antennas being correlated so that digital mixer edge effects are not apparent and so that sufficient anti-aliasing attenuation occurs. A minimum of 100 cycles of differential phase rotation within an *incoherent* integration period is recommended. If desired, frequency shifts could be dynamic so that anti-aliasing occurs even over arbitrarily long *coherent* integration times.

8.2.18 Sub-arrays

There is no limit to the number of interferometer sub-arrays. Additionally, separate sub-arrays can have mutual antennas as long as the configuration within a sub-array is consistent within the constraints of correlator data routing (and as long as the configuration software is capable of doing this!). Phased-VLA sub-arrays *may* be defined differently on different sub-bands since phasing hardware sums data independently on each sub-band. There is a maximum of 5 phased sub-arrays per sub-band, and antennas must be included in a given sub-array in “fixed” groups of 4. (The actual “fixing” of the groups of 4 is determined by correlator cable connections, and could be modified by changing connections in front of the correlator.) In this case, if an antenna is in a particular group of 4 and in a particular sub-array, it can be disabled but it can’t be assigned to another sub-array.

8.2.19 Phased-VLA

The delivered system will include 8 digitally phased sub-bands, on up to 48 antennas, for a total bandwidth of 1 GHz. Exactly which sub-bands of which basebands are phased is a free parameter and is determined and fixed by rack wiring (i.e. the wiring can be changed, but not dynamically and not under program control). The architecture supports phasing all (18) sub-bands of all basebands—phasing more than the 8 planned sub-bands is strictly a cost/configuration consideration and no additional design effort (or installed cable replacement) is required for post-installation upgrades. The EVLA will operate using VLBI-standard frequencies, and so all-digital phasing will be done. One wide-word (~ 8 -bit) output that contains one phased sub-array’s data stream before re-quantization will be provided on a front-panel connector. This allows expansion for phasing more than 48 stations with an external synchronizer and digital adder. The delivered system will not include the hardware required to feed the phased outputs back into the correlator for auto-correlation or cross-correlation processing. Also, additional hardware that is not part of the delivered system will be required to connect phased outputs to VLBI data recorders.

8.2.20 Auto-correlations, Data Statistics, and Phase-Cal

Four wideband auto-correlation products will be provided for every baseband pair. Each product will have 1024 spectral channels, but with a factor of 4 sensitivity loss over an ideal auto-correlation. This loss of sensitivity comes from only using $1/16^{\text{th}}$ of the data or acquiring the auto-correlations in 64-lag chunks every 10 milliseconds due to hardware limitations. *Sub-band* auto-correlations will be performed with cross-correlator hardware and 16,384 spectral channels per antenna will be possible. Sub-band auto-correlation results may contain transition-band aliasing so it is not possible to seamlessly stitch sub-band auto-correlation spectra together (except where a “cross-auto-correlation” is performed—if the antenna LO system is sufficiently flexible). Sixteen wideband state counters are provided per baseband (64-bit data highway) that, in ≤ 4 -bit mode are time-multiplexed across the input data streams.

Time-multiplexing is under CMIB control and parameters can be modified every 10 milliseconds. In 8-bit input mode, there is one accumulator and it can be set to count occurrences of any of the 256 states in a similar time-multiplexed fashion. After filtering and re-quantization, one accumulator (state counter) is used to time-multiplex the acquisition of state counts across 16 (4-bit) or 128 (7-bit) possible states. Also, full sensitivity total power accumulators are provided for both 4-bit and 7-bit requantized data. Finally, each Filter Chip contains a dedicated phase-cal tone extractor with a linear frequency synthesizer and full delay-tracking compensation that operates on the filtered and requantized data stream.

8.2.21 VLBI

The correlator is fundamentally a VLBI correlator and the system will be delivered with all of the “hooks” in place for VLBI. Each Station Board has two VSI (VLBI Standard Interface) inputs and two VSI outputs—one of each per baseband (data highway) to allow data to be piped into the Filter Chips from some source (such as a VSI playback device) and out of the Filter Chips to some destination (such as a VSI recording device). Each VSI input or output can handle sixteen 2-bit sampled data streams at rates up to 256 Msamples/sec. VSI signals on the Station Board are broken-out to two connectors, each of which contains a VSI input and a VSI output. These connectors plug into Common Backplanes. To convert from the backplane connectors to standard VSI-H MDR-80 connectors requires special purpose breakout modules which are also not part of the delivered system. Refer to Figure 8-2.

8.2.22 Maintenance

All (semi-conductor-populated) modules and module-to-module communications will be designed for hot-swap capability. Additionally, the design is such that swapping out one module will have the minimum possible impact on other modules and their data products. The estimated MTTR is about 10 minutes (with maintenance personnel on-site). The estimated MTBF is not known, but state-of-the-art commercial devices, design, and production techniques will be employed for maximum benefit. It is not anticipated that regular semiconductor failures will occur. All hardware modules will have active (via computer) and dead-man (thermal switch) temperature monitoring and shutdown. Separate cooling fan monitors will be employed so that fan failures can be detected immediately, rather than waiting for components to heat up. It will be possible to remotely power-cycle individual modules using a power control computer that is not part of normal correlator processing (for increased reliability). While the correlator is on-line, embedded synchronization codes will allow for constant monitoring of module health and module-to-module communication integrity. When off-line (for example, when slewing antennas) it will be possible to enable internal test vectors for complete correlator system testing. The intent is that a test is treated like a normal observation, except that, instead of processing data “from the sky”, test vectors are processed instead. *The degree to which antenna and antenna transmission systems are included in this kind of testing is currently undefined.*

8.2.23 Interference Mitigation

The correlator contains some special real-time burst interference nulling hardware. This includes accepting data valid flagging from the antenna and not correlating when it is flagged bad, and detecting saturation before the re-quantizer in the Filter Chip to flag and not correlate invalid sub-band data. Additionally, high-speed dumping (with scaleable performance back-end computing), and high spectral dynamic range provided with many-bit samples will enable post-correlation, temporal/spectral excision. The WIDAR design strongly attenuates the modulating effects of time-variable narrowband interference on normalized correlation coefficients (Carlson, Memo# 009), so post-correlation excision of non-saturating burst-like interference should be quite effective. Post-correlation interference cancellation, should it be found to be effective, can easily be handled since the interference detection antenna is just another antenna to the correlator. The correlator also has the capability of processing 8-bit sampled data for high-spectral dynamic range even in the presence of overwhelmingly powerful narrowband interference.

8.2.24 System Timing

All actions in the correlator are synchronized to distributed “TIMECODES” and a 64 MHz clock. Correlator delivery will include a “TIMECODE Generator Box” (TGB) that generates the TIMECODES. The TGB requires an externally-provided 128 MHz clock and 1 PPS time tick synchronized to the UTC 1 PPS. To support simultaneous

VLA/VLBI operation, **three** TIMECODEs will be distributed. One TIMECODE is the current real-time UTC for the VLA, and two TIMECODEs are programmable to any UTC epoch for VLBI. Each TIMECODE that is generated can be delayed a programmable amount (from the input reference 1 PPS, with some TBD maximum), to take into account delays through the FOTS and the large baseband delay buffer on the Station Board. One 64 MHz clock and three TIMECODEs are distributed and available to every Station Board. Station Boards, in turn, synchronize and generate timing for downstream Baseline Boards and Phasing Boards.

8.2.25 Computing and Data Highways

The Correlator installation will include three classes of computers. The top-level monitor and control and power control computers will be high-availability units, either CompactPCI or VME based rack mount systems. Each Station, Baseline, and Phasing Board will be directly controlled by an embedded PC/104+ format computer called a CMIB (Correlator Module Interface Board) installed as a mezzanine card. Finally, back end data processing computers will be COTS (Commercial Off-The-Shelf) desktop or rack mount PCs arranged in an N+1 redundant configuration and tied together as (Beowulf) clusters. Communication between the top-level monitor and control computer and the embedded computers will be via 100 Mbps Ethernet over a twisted pair switched network. The top-level computer will also share a network connection with the Back End computing cluster manager. Lag data from the Baseline Boards will be transmitted to the back end computers using 1 Gbps Ethernet through Gigabit Ethernet switches (Rowen 2001). These Gigabit switches will provide for a more scalable and run-time configurable Back End computing design and help isolate the Correlator from future technology upgrades.

One possible correlator network topology is shown in Figure 8-3—other topologies for monitor and control are possible, and indeed likely. A standard Linux OS distribution will be used in all these computers with the CMIBs running a pre-emptable kernel version to address the tighter timing requirements of this computing system layer. All inter-system communications will rely on standard Unix/Linux facilities to minimize development effort and maximize scalability and compatibility with external systems. CMIB hardware control will be abstracted behind standard Linux device drivers to allow higher level control processes to be developed with a greater amount of portability should computing platforms be changed/upgraded. All monitor and control communications will be message based to best isolate failing systems from impacting overall system performance and to provide for a highly modular distributed system. Messages will be XML (eXtensible Markup Language) based. Use of XML provides correlator system designers and users an industry standard protocol that is both human readable and includes a wealth of industry standard applications and tools to manipulate XML messages. The intended design will provide the system engineer easy access to correlator hardware internals as well as provide system users a polished orthogonal interface where the tedious details of hardware setup are hidden to avoid confusion, but still available if desired.

8.2.26 Environment

The correlator will be designed for a “benign office environment” with an ambient temperature of $\sim 25^{\circ}\text{C}$ at the altitude of the VLA. Board and rack design will be such that its operating temperature range will be 0°C to $+35^{\circ}\text{C}$. However, for reliability the ambient temperature should be kept at about $+15^{\circ}\text{C}$. The (32-station) correlator is expected to require somewhere between 120 kW and 160 kW of power (a 40-station correlator requires somewhere between 180 kW and 220 kW). Correlator delivery will include installation and configuration of the system power supply (by the manufacturer). The system power supply will be a 48 VDC plant with battery backup as used in central-office telephony systems. This plant should have very high reliability, be very efficient (compared to an AC-AC UPS), and will be on-line serviceable and hot-swappable. This plant requires three-phase 208 VAC, has a power factor of 0.9, is approximately 90% efficient, and meets FCC Part 15 Class ‘A’ EMI specifications. Overload protection and remote monitoring and control will be as supplied by the manufacturer. This power plant could have a back-up time of ~ 1 hr 15 min at full load (although the requirement is for only 15 minutes of backup). It may also be necessary to install one or more AC-AC UPSs for supplying power to correlator COTS PCs and CompactPCI computers. These additional UPSs will be included in the installation. NRAO will be responsible for providing the (sufficiently clean) office space, HVAC, and mains power service for the correlator. The correlator (with enough room for 40 stations and up to 16 phased sub-bands) will require a floor surface approximately 32 ft x 32 ft (~ 1000

ft²—refer to Figure 8-4). This is based on 10 station racks each with a footprint of 2.7 x 3 ft, and 16 baseline racks each with a footprint of 2.7 x 3 ft. It is expected that racks are to be at most 7 ft high.

8.3 Correlator Architecture

8.3.1 System Overview

A simplified block-diagram of key correlator systems is shown in Figure 8-1. Sampled data arrive from the antennas into the station-based processing subsystem where coarse delay compensation occurs, the data are digitally filtered into sub-bands, and then data along with models and synchronization information are transmitted to the 16 downstream sub-band correlators. *Each sub-band correlator correlates (or, is capable of correlating) one sub-band from all 8 basebands on all baselines.* Each sub-band can be any width and placement (within slot constraints mentioned previously) within its associated baseband.

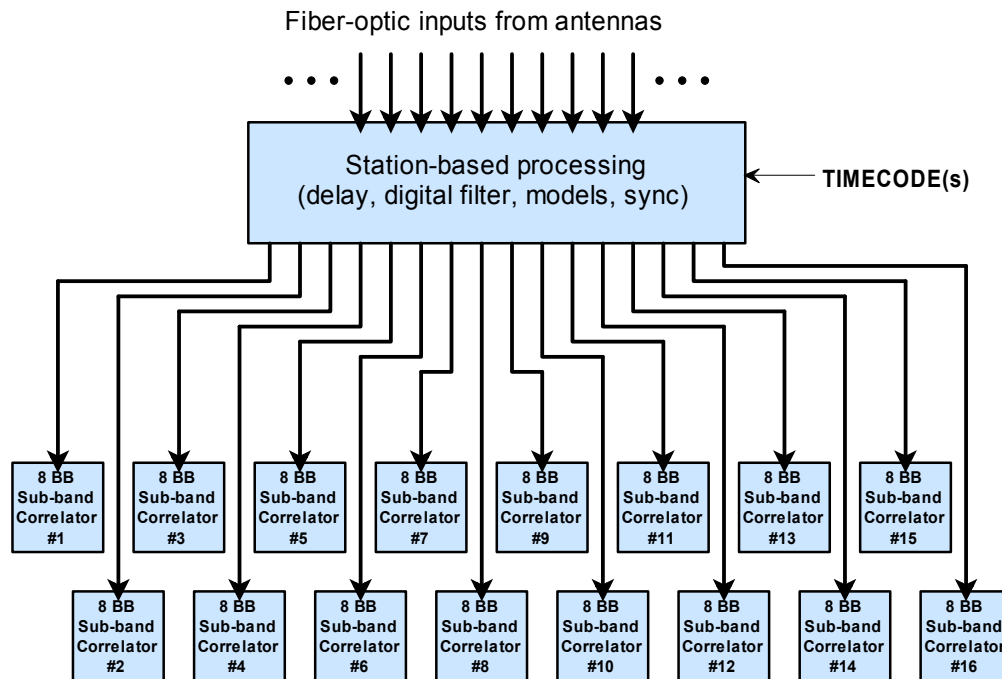


Figure 8-1 Simplified correlator system block diagram.

The width and placement of a sub-band is *entirely* governed by the Filter Chip tap weights and chosen decimation factor. In a particular Correlator Chip (bearing in mind that each Correlator Chip in a row or column is fed the same X or Y data), it is perfectly acceptable to be correlating one or more sub-bands with different bandwidths and placement within their respective basebands. For example, within a Correlator Chip, one correlation could be slot 1 of 16 slots within baseband 1 with a 128 MHz bandwidth, while another correlation could be using recirculation on slot 63 of 128 slots within baseband 5 with a 16 MHz bandwidth. Switching circuitry on the output of the station-based processing block effectively allows any sub-band filter output to be routed to any sub-band correlator. Thus, redundant correlations across sub-band correlators could be performed for test or on-line redundancy purposes.

8.3.2 System Module Connectivity

Figure 8-2 is a diagram showing the interconnectivity of all correlator modules. A brief description is provided here and more details for each module can be found in system design documentation. All connectivity in the correlator is facilitated by 36 “wafer” connectors and Common Backplanes. A “wafer” contains 4 signals: 1 clock, control/timing signals, and 2 sampled data streams. Wafers in the connector and the ensuing cable can be grouped together or used separately in various ways. Signals on wafers run over cost-effective, high-capacity differential pair copper wires.

Data from the antennas arrive via fiber-optic links where they are wavelength demodulated before being presented to Fiber-Optic Receiver Modules (FORMs) mounted as a mezzanine card on the Station Board. On these cards, the

fiber-optic signal is demodulated into electrical signals for use by Station Board electronics. Each “station input” in the correlator consists of four Station Boards—also known as a “Quad”. Each Station Board handles two, 2 GHz sampled basebands—also referred to as a baseband pair. The Station Board Delay Module mezzanine card compensates for wavefront geometrical delay as well as delay through the fiber-optic system. Data then go to the Digital Filter Banks, the output of which are 16 (up to 18) sampled data streams no longer in de-multiplexed parallel form as they were going into the filters. These data go through crossbar switches before going to the “SB Output Connector”.

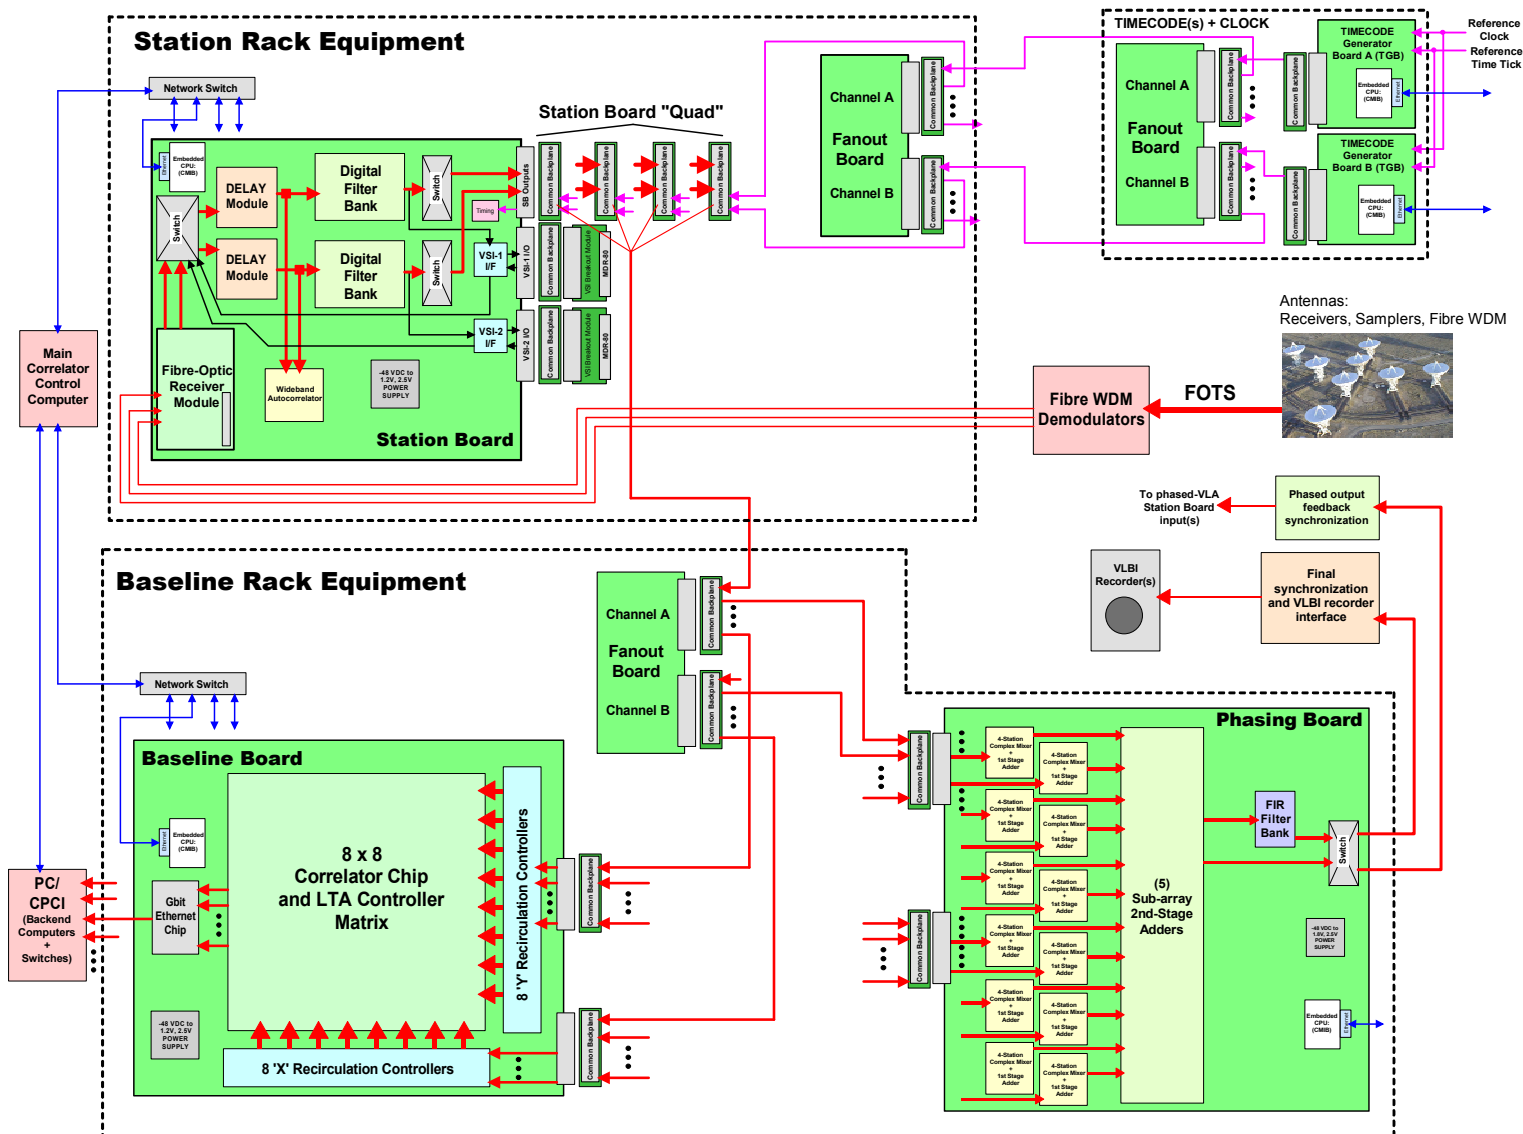


Figure 8-2 Correlator module connectivity diagram.

Each cable from one station Quad that is routed to a Baseline Rack contains data, timing, model, and synchronization information for one sub-band of all 8 basebands from one station (quad). All real-time information required for the down-stream Baseline Boards (recirculation, phase-binning, dumping, phase models, delay-error induced phase error models) is generated on the Station Boards and flows with the data on each cable. Within a Baseline Rack, data get distributed and fanned-out to all of the Baseline Boards and the Phasing Boards via the Fanout Board.

On the Baseline Board, there are 8 ‘X-station’ and 8 ‘Y-station’ inputs—each input being data from one sub-band cable from one Quad of Station Boards. The input data are resynchronized and formatted for transmission to a row or column of Correlator Chips by the 8 ‘X’ and 8 ‘Y’ Recirculation Controllers. The Correlator Chips correlate data and respond to commands coming from the Recirculation Controllers. After integration, and on command from Recirculation Controllers, the data are read out of the Correlator Chip by its own dedicated LTA (Long-Term

Accumulator) Controller and saved in 256 Mbit LTA DDR SDRAM. Although having one LTA Controller for each Correlator Chip seems extreme, it offers significant performance advantages and is cost-effective since a relatively small (and inexpensive) FPGA can be used. When enabled by an on-board scheduler, LTA data are transmitted via Gigabit Ethernet (and a switch—not shown—see Figure 8-3) to a Backend computer (PC) for further processing. The data on the Baseline Board are not handled by a microprocessor so there are virtually no bottlenecks to data flow off the board.

On the Phasing Board, data for one sub-band pair (possibly more—the final functionality of the Phasing Board is not yet defined) from all antennas enters via the Common Backplane. These are the same data that go to the Baseline Boards. Each Phasing Board likely sums antennas for one sub-band pair of one baseband pair. Data are summed in two stages to keep on-board data path widths within device capabilities. In the first-stage, data from antennas are summed in groups of 4—hence the “group of 4” restriction mentioned in the specifications section. Each antenna’s data are complex mixed before complex addition to remove the Doppler shift and the frequency shift required by the WIDAR technique. There are 5, second-stage adders—each one being the output of one sub-array. After second-stage addition, the complex data are combined using the Hilbert transform filter, the second part of the digital single-sideband mixer. Details and test results are found in (Carlson, Memo# 008). The final summed output is available in normal sub-band “wide” mode, or it can be filtered with on-board FIRs to generate more, smaller, sub-bands for VLBI recording.

8.3.3 System Network Topology

The correlator system is designed for scaleable performance: there are virtually no bottlenecks to output data flow and the system’s real-time data handling performance is largely governed by Backend COTS computing performance. The proposed network configuration for the correlator is shown in Figure 8-3. (All network connections are isolated with either fiber or transformer-coupled/isolated Ethernet. Isolation is required to minimize ground-loop noise between sub-systems operating on different UPSs.) In the figure the MCCC is the Main Correlator Control Computer and the CPCC is the Correlator Power Control Computer. Each Station, Baseline, and Phasing Board has an embedded CMIB. Finally, Baseline Board data are transmitted to Backend computers on Gigabit Ethernet through Gigabit Ethernet switches. Switch connections are such that all data required for a particular FFT arrives at one computer and that processing can be dynamically load-shared across computers, by changing Backend destination IP addresses on the Baseline Boards. This eliminates the need for an additional wideband network fabric that would be required if a distributed FFT is performed. More straw-man details of network topology and Backend processing are in (Rowen, 2001), (Morgan, 2003), and section 8.4.

8.3.4 System Installation

The correlator will be a large system. For cost and performance reasons, it is desirable to minimize the correlator installation footprint. A smaller footprint requires shorter and less expensive cables and results in better signal performance—particularly at the clock and data rates under consideration. A preliminary floor plan for a 40-station correlator installation is shown in Figure 8-4.

In this plan, maximum 9 m-length cables are required for data distribution from the Station Racks to the Baseline Racks. In a 40-station correlator, each sub-band correlator uses 15 Baseline Boards—leaving one slot open for a Phasing Board. Signal arrival time mismatch at the Baseline Boards will be completely compensated for by buffers in the Correlator Chip. Racks need only front and rear access and can be installed side-by-side. Each rack holds 16 boards plus a 6U sub-rack for FORMs and Fanout Boards. Provided there is floor space, more Phasing Boards and racks can easily be added at a later date without requiring replacement of existing cabling. Each rack is about 7 feet high. The fronts of the racks are shown with bold lines in the figure. All high-speed cabling is within the racks and under the (raised) floor. Any other cabling (e.g. network cables shown in Figure 8-3), will be run in cable trays hanging from the ceiling. The MCCC, CPCC, and Backend computers can be located in the room, or located in a completely separate room.

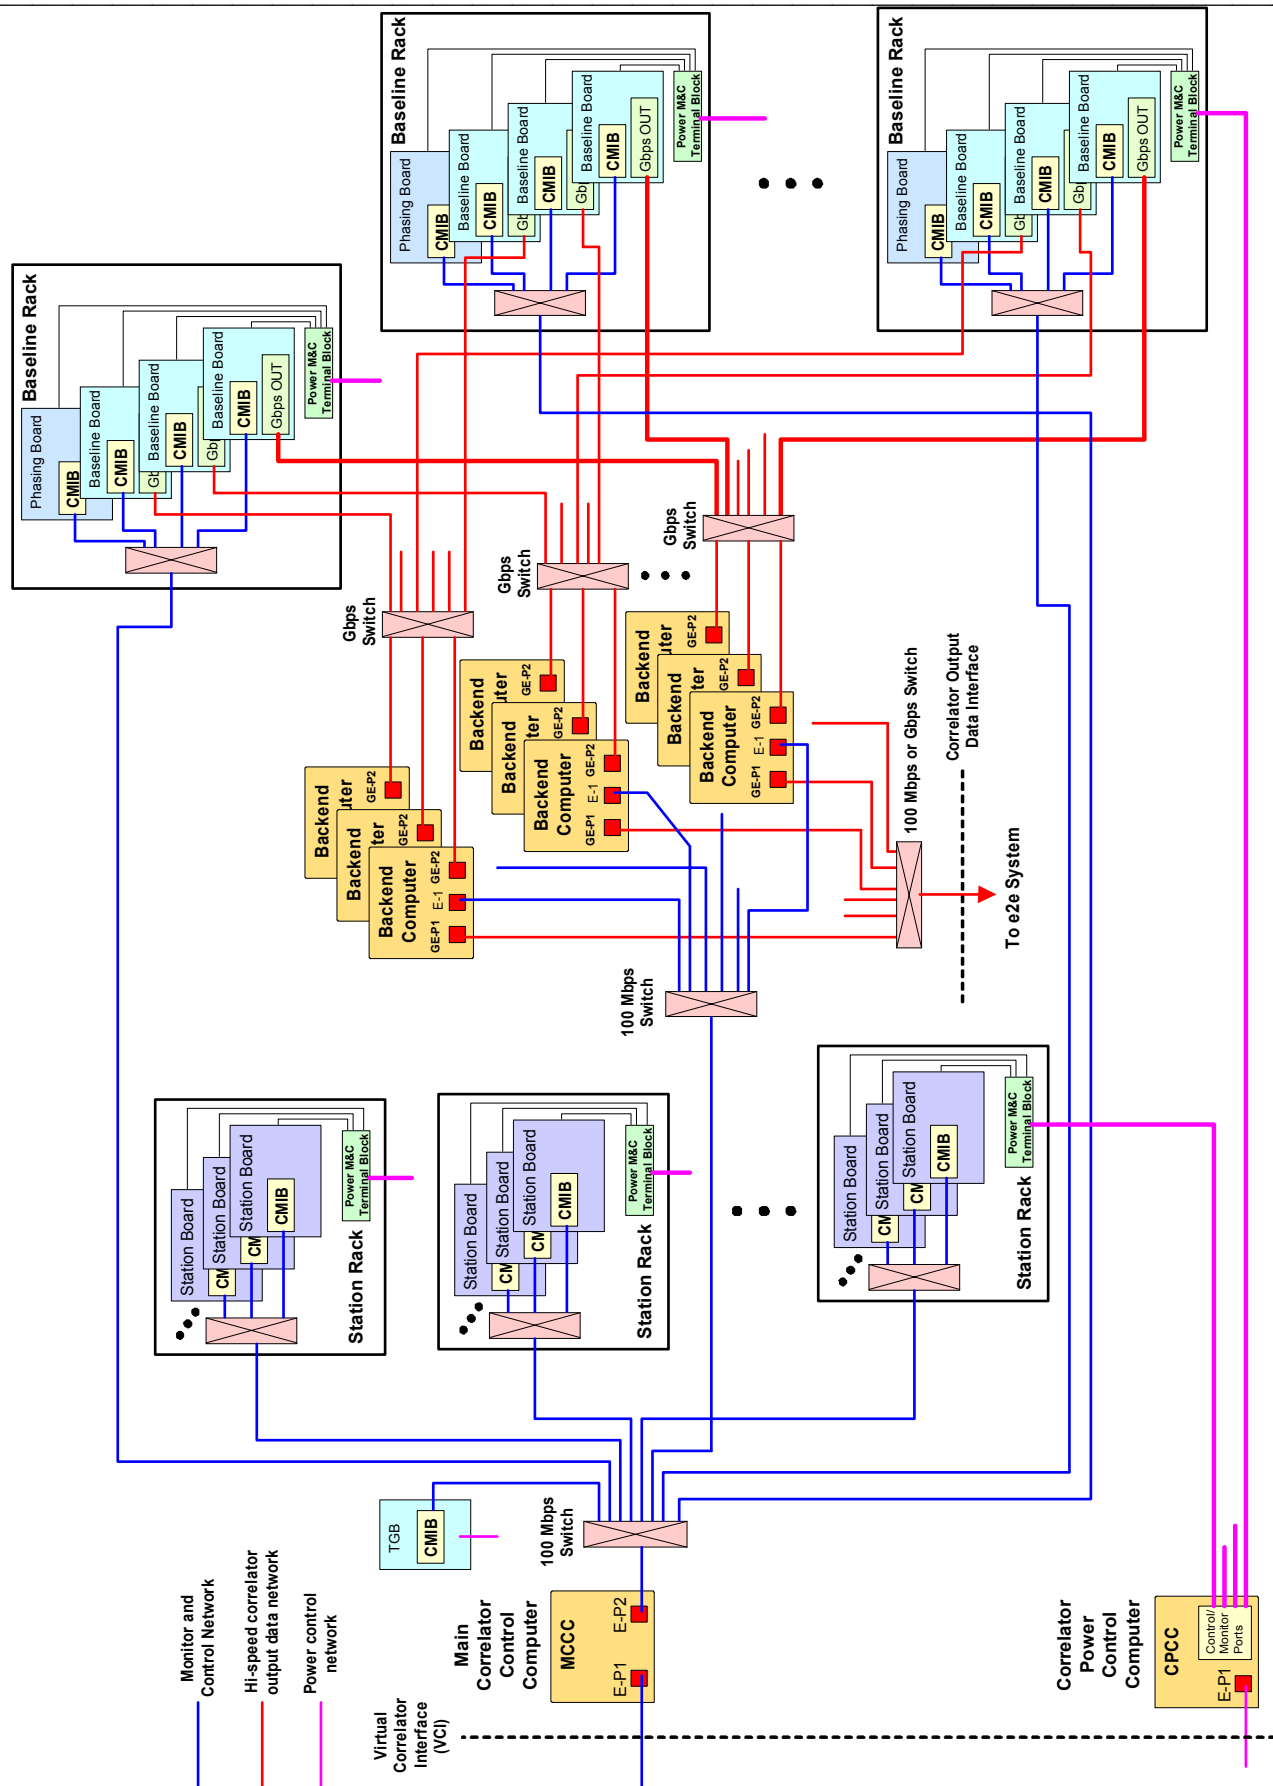


Figure 8-3 Possible correlator computing/network topology.

40-Station Rack Layout

28 ft

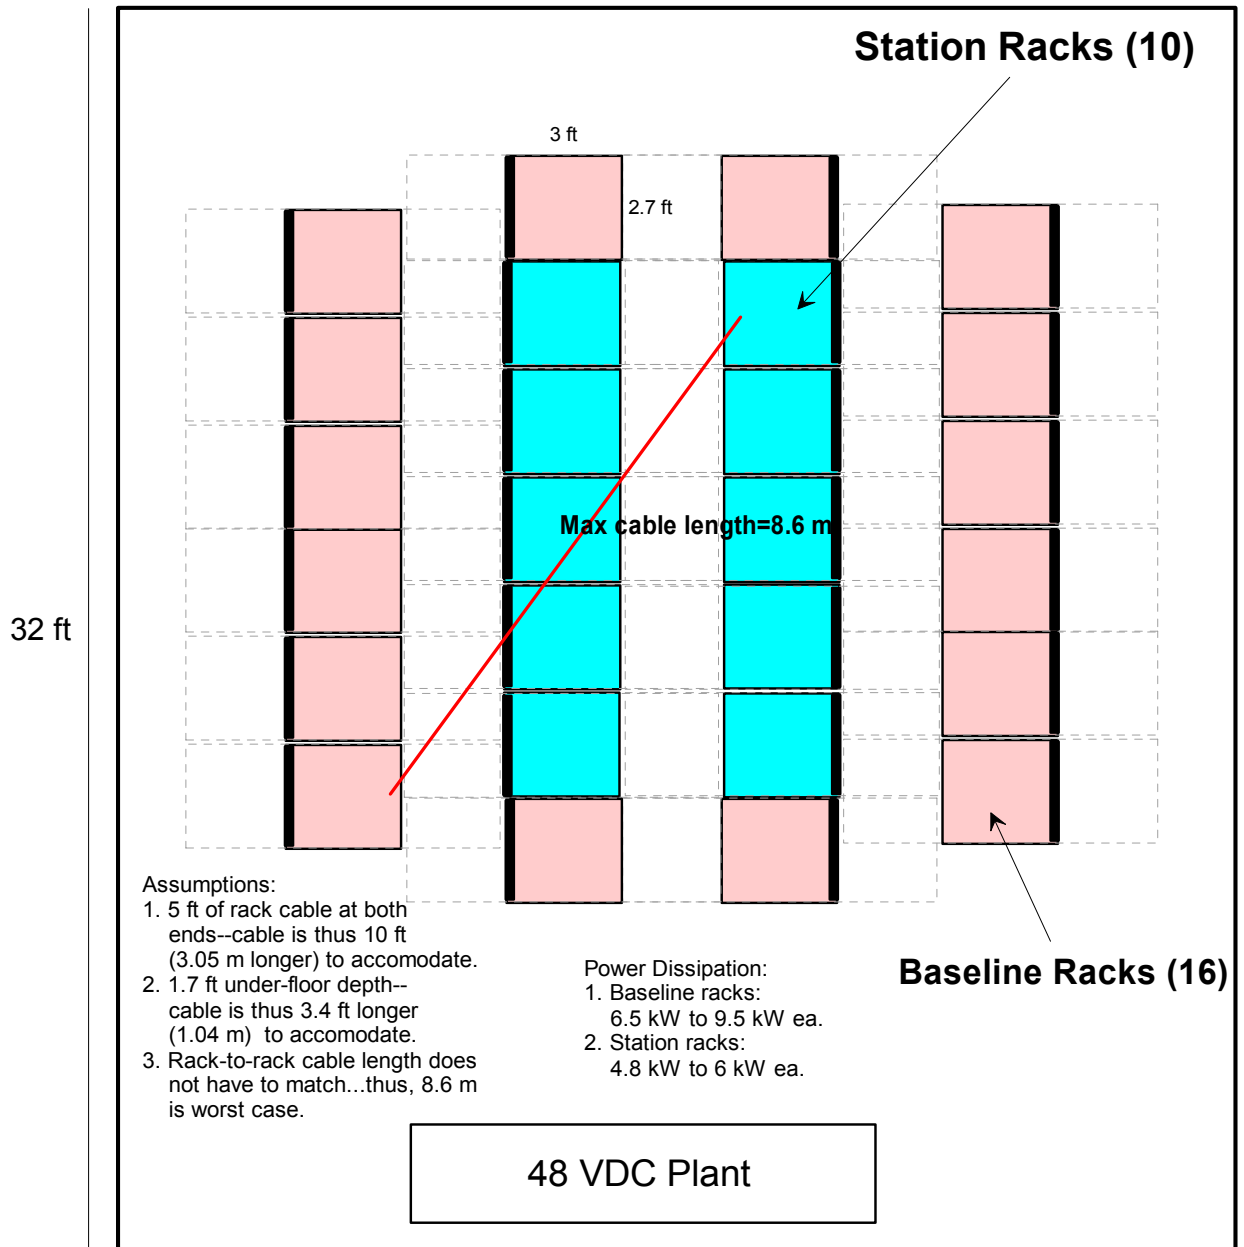


Figure 8-4 Preliminary 40-station correlator system floor-plan.

8.4 Correlator Backend (CBE) Requirements

A complete requirements specification is given in: “System Requirements Specification: EVLA Correlator Backend“, project document A25251N0000, revision 2.0, May 10, 2002. It can be found on the Computing Working Documents web page at <http://www.aoc.nrao.edu/evla/techdocs/computer/workdocs/index.shtml>. The following is an overview of the key requirements.

8.4.1 Assumptions

Packetization of lag frames (a lag frame consists of one lag section of up to 128 complex lags and identifier information from one Correlator Chip (section 8.2.11)) including setting of the correct Backend node destination IP address (section 8.3.3) for the given baseline will be handled by the correlator. The CBE will provide a mapping of baselines to node IP addresses for the current correlator mode. The lag frame packets will not necessarily arrive in a set order and their delivery will be a one-time event. Resends of missed or bad packets will not be possible due to Baseline Board hardware limitations and performance requirements.

Indirect (i.e., non-correlator lag frame) data will arrive in a timely fashion. That is, with no significant delay prior to its being needed at a particular point in Backend data manipulation, processing, or formatting.

The e2e System will be designed to handle output rates and volumes delivered by the Backend during times of peak production. The Backend will provide results to the e2e that will be complete and ready to be ingested by the archive.

8.4.2 CBE Input

The Backend shall be able to receive the following correlator outputs: lag frames, quantizer power measurements, Filter Chip parameters (power measurements etc), frequency shift parameters, windowing parameters, and quantizer and re-quantizer state counts. It shall also be able to receive observational mode, meta-data (“sky frequencies”, polarizations etc.), status requests, and other EVLA data from the Monitor and Control System.

8.4.3 CBE Output

The CBE shall be able to deliver formatted observational output to the e2e and status, warning, error, and system component failure and recovery reports to M&C.

8.4.4 Correlator Interface

All lag frame data shall be sent directly across the correlator to Backend interface using Gbit Ethernet and UDP/IP. All Backend cluster nodes shall have a switched path to each correlator output point (the Baseline Boards)—although some non-direct paths will not have the same performance of more direct paths. The interface shall have sufficient bandwidth to meet the initial maximum aggregate data transfer rate of 1.6 Gbytes/sec. All lag frames from the same baseline (that could be distributed across Baseline Boards) shall be routed to the same Backend cluster node.

8.4.5 M&C Interface

All non-lag frame correlator data along with other EVLA data, M&C requests, Backend responses and Backend-generated reports shall pass to and from the M&C System via the Virtual Correlator Interface (VCI). If sufficient bandwidth is not available to handle all traffic, critical auxiliary correlator data may have to be routed directly from the Main Correlator Control Computer.

8.4.6 e2e Interface

All final, formatted astronomical results shall be sent directly to the e2e across this interface. It shall have sufficient bandwidth to meet the initial maximum aggregate data transfer rate of 25 Mbytes/sec. All Backend cluster nodes shall have a path to the e2e system.

8.4.7 User Interface

The Backend shall be capable of presenting a command line interface on any and all nodes for use by software development and test personnel. It shall also have selectable internal test modes that produce printed values for key variables at critical locations in the code. It is not a requirement that the CBE have an external GUI of its own.

8.4.8 Data Processing

Backend cluster nodes shall be able to perform the following data manipulation and processing tasks: lag set assembly (a “lag set” is required for an FFT), data valid normalization, coarse quantization correction, time stamp adjustment, residual phase rotation correction, FFT, interference removal/reduction, windowing, integration, and output formatting. These are capabilities that shall be available. A subset will normally be used on any given data stream. The Backend will *not* perform sub-band stitching operations.

8.4.9 Internal Monitor and Control

The Backend shall be self-monitoring. Input, output and data processing rates shall be measured and error and warning statistics shall be maintained in order to continually monitor system health and anticipate problems. All internally generated reports and status information shall be passed to the M&C for presentation to the outside world.

8.4.10 Reliability

The Backend shall be capable of attempting recovery from a number of failure modes. The failure of a single node, including the node running the Backend monitor and control functions, shall not affect any other node. The loss of an external network connection shall not affect internal operations until all on-board storage resources are filled, in the case of loss of the e2e connection, or until necessary auxiliary data is needed, in the case of loss of the M&C connection. The system shall be able to kill and restart corrupted processes, and reboot failed processors and network connections. It shall report all problems, recovery attempts and outcomes to the M&C. The goal is to avoid total system reboots for a period of time greater than or equal to the normal EVLA maintenance interval.

8.4.11 Scalability

The total system shall be scalable to higher rates of input, output and data processing with an ultimate objective of meeting the full data generating capability of the correlator (16 Gbytes/sec). Hardware shall be extensible in a manner that is transparent to software and vice versa. Upgrades shall meet seamlessly with unchanged components.

8.5 Correlator Backend Design

The CBE will be a distributed cluster based system with the nodes logically linked via message passing middleware. High-speed external switched networks will be used to connect to the correlator and e2e systems. (See Figure 8-3 for a schematic of the possible topology of these networks.) The node hardware will be multi-CPU Intel or Intel-clone processors configured with large amounts of memory and disk storage. The current operating system of choice is Linux, and the message passing middleware will be either PVM or MPI. All are open source, and based on widely accepted industry standards.

There will be two main software subsystems running on the nodes. One node, with one or more shadow nodes, will run a subsystem consisting of the Backend internal monitor and control functions. The remaining nodes will run the processing pipeline subsystem that will consist of input, sorting, data processing and output functions. Several of the processing pipeline nodes will not actually be receiving data from the correlator, rather, they will serve as standby capability in the event of a node failure elsewhere.

8.5.1 Backend Control Function

Backend Control is the gateway to the CBE. All inbound and outbound non-correlator frame data will pass through it. Backend Control will also maintain a statistical model of the CBE system state. It will incorporate measurement, error and warning data from the processing nodes along with periodic status checks performed by the Monitor Function. There will be three classes of messages. The most basic are messages that are simply routed to another

destination. A second class is messages that will also be routed, but in the process data for the statistical model will be extracted. The third type has the Control Function as a destination and is also used to update the statistical model. Backend Control will generate messages, based on the state of the statistical model, to request check and repair, and offload services from the Monitor Function.

8.5.2 Backend Monitor Function

The Monitor Function performs system component monitor and recovery operations based on directions received from the Backend Control Function. It will do status checks of networks, processors and processes. It will attempt network and processor restarts, and be able to kill and restart damaged processes. It will also perform off-loads of data processing from malfunctioning to standby nodes.

8.5.3 Input Function

The Input Function will receive data packets from the correlator/Backend network interface and deposit the lag frames contained in them directly into large blocks of memory reserved for its use. It will signal the Sorting function when a memory block has been filled. In the case of all memory being filled, it will discard new input until more memory becomes available. The primary objective of the design of this function is to maximize throughput by minimizing overhead.

8.5.4 Sorting Function

Once a full memory block is released by the input function, the Sorter will sort the lag frames into lag frame sets; a lag frame set being all lag frames needed to form a complete, properly ordered series of correlator lags required for FFT. The lag frame data itself will not be moved during sorting, instead, the memory location of each frame will be sorted into tables for later reference by the Data Processing Function. The Sorter will keep track of the lag frames that have been sorted and provide a mechanism whereby the Data Processing Function can identify completed lag sets. It will also provide the mechanism whereby the Input Function can identify the next available memory block.

8.5.5 Data Processing Function

The Data Processing Function will access the sort tables generated by the Sorter and identify completely sorted lag sets. If all auxiliary data needed during processing is available, it will access the specified memory locations, retrieve the lag frame data and assemble the lag set. The lag set will then be passed through the selected processing steps and on to an output storage location. During data processing, floating point exceptions and other computational errors will be trapped. The function will periodically recheck completed lag sets that were not processed due to lack of auxiliary data and process them if the needed data has become available. Occasionally a few lag frames may be discarded due to age and total unavailability of auxiliary data.

8.5.6 Output Function

Prior to formatting and transmission to the e2e system, processed data will be stored in an output data memory area reserved for its use. The Output function will have access to auxiliary disk storage to prevent memory overflow. Data will be written to local disk when output memory use reaches a preset limit. Previously written data will be retrieved from disk prior to its being needed for formatting. Once all meta-data needed to format a particular set of processed results has been received from the M&C System and internal sources in the Backend, the pertinent observational data will be drawn from output storage and organized according to the specified format. It will then be sent to e2e via the Backend/e2e interface network. The send will be checked for successful completion, and in the event of failure, resends will be done until success is achieved. Only then will the output storage area be released for reuse.

8.6 Deliverables

Table 8-3 summarizes the modules that will be developed and delivered by the NRC. This table includes items and costs for a 32-station correlator. Costs are estimated and do not include NRE (Non-Recurring Engineering) costs or technician test, burn-in, and handling costs. Spares are not included. All figures are in 2004 U.S. dollars.

Table 8-3 Cost estimates for NRC-supplied correlator deliverables

Qty	Item/Description	Cost (ea) USD
128	Station Board (c/w mezzanine cards)	\$15,055
266	Fanout Board	\$666
1258	Common Backplane	\$61
160	Baseline Board (c/w mezzanine cards)	\$18,212
8	Phasing Board (c/w mezzanine cards)	\$10000 (est)
2	TIMECODE Generator Board (c/w mezzanine cards)	\$2400
1	High-speed cables for 32 stations	\$404,450
1	Sub-racks and racks for 32-station correlator	\$225,600
1	COTS computers (MCCC, CPCC, Backend (50), Ethernet switches)	\$241,970
1	48VDC, 4000A plant including batteries, shipping, installation, cables. Can be field-upgraded to 6000A. AC-AC UPS for back-end COTS PCs (est. 30 kVA)	\$344,000
n/a	Correlator software	

The total installed system cost is about \$6.4 million. The total estimated cost of the correlator including NRE, labour, and all of the above deliverables is \$14.8 million dollars, *not* including a contingency of \$1.5 million (based on a 0.818 exchange rate).

Table 8-4 summarizes additional modules and components that NRC does not develop or supply as part of the correlator installation. Not all modules are required, depending on desired configuration (VLBI correlation, phased-VLA correlation etc). Quantities are for a 32-station correlator. Higher-level on-line, interface, data processing, and VLBI control software is not included in this table. Quantities in **boldface** are considered essential for the system to perform its basic functions.

Table 8-4 Additional (NRAO-supplied) correlator deliverables.

Qty	Item/Description
135	Dual-input (2 x 2 GHz bandwidth) Fiber-Optic Receiver Module (FORM) card to plug into the Station Board. C/w test vector receivers and test vector transmitters (to Station Board receivers). Each baseband output is 4 Gs/s arranged as 16 demultiplexed streams, 4 (or 3) bits/stream, @ 250/256 Mbps each.
1+	Final Phasing Board output synchronization and VLBI recorder interface. This includes a VSI-H output to feed back into the correlator (i.e. into the Station Board's VSI-H inputs).
12+	VSI-H connector breakout module to breakout VSI-H I/O capability to the Station Board.

8.7 Interfaces and Impacts on Other Systems

Table 8-5 summarizes correlator interfaces and/or associations to the external world, and a description of possible impacts on other parts of the system.

Table 8-5 Table of correlator interfaces and potential impacts on other systems.

Interface/Location	Description	Impacts
Fiber-Optic Receiver Module (FORM). Sec. 8.2.4, 8.2.10	3 fibers into each module. Two, 48-bit data highways out of each module. Each data highway contains a 2 GHz BB, 3-bit samples at 4 Gs/s. Provision for 4-bit samples. Provision for one 8-bit sampled stream at 2 Gs/s.	BERT transmitter in the antenna's fiber-optic transmitter and in the (correlator) receiver module allows transmission system and interface testing. Real-time, non-invasive CRC checks allow on-line connectivity testing between the FORM and the Station Board. Supports 1, 2, 3, 4, or 8-bit sampling with flexible baseband widths. 8-bit sampling reduces the sampled bandwidth by a factor of 2.
Delay Module (Station Board). Sec. 8.2.16	This modules inserts wavefront delay in the station data path. The depth of this delay determines the maximum baseline.	Design will be for total 0.25 seconds of delay with +/-122 ps resolution. Delay can be increased with new module. Designed to tradeoff bandwidth for number of BBs.
Correlator clock/timing interface (TIMECODE) Sec. 8.2.24	Reference clock (128 MHz), and reference time tick (1 PPS). Required for correlator TIMECODE generation.	Requires clock and time epoch (1 PPS) from array maser/timing master.
LO system (antenna). Sec. 8.2.6, 8.2.7, 8.2.12, 8.2.13, 8.2.17	LO offsets for anti-aliasing, sub-sample delay tracking, and narrowband harmonic/inter-modulation product reduction.	Requires 100 Hz LO tuning resolution for LO offset capability. An antenna can have the same LO offset in every one of its basebands. Optionally, different LO offsets in the same antenna allow sub-band "cross auto-correlation". System control should ensure that minimum acceptable net phase rotation rate is ensured on all baselines. Time-variant LO offsets could be employed for more aliasing attenuation on long <i>coherent</i> integration times. LO offsets could be turned off, and correlator would lose sub-sample delay tracking and anti-aliasing capability.
Noise diode switching (antenna). Sec. 8.2.6	Noise diode switching in the antenna receivers for system noise calibrations. A reference Filter Chip will synchronously switch with the noise diode to acquire power data with the diode on and with the diode off.	Switching/binning in the correlator will be synchronized to switching in the antenna using a timer and apriori knowledge of the switching period and phase. It is not defined what this switching rate will be.
VLBI recorder interface (Phasing Board) and output for feedback into the Station Board. Sec. 8.2.19	Interface box to synchronize outputs of multiple Phasing Boards for transmission to VLBI recorder, and feedback into the correlator.	NRAO-developed. Rack space and physical location is currently undefined.
VSI-H I/O (Sec. 8.2.21)	Single antenna VLBI I/O to/from Station Boards	Correlator will be able to connect to VSI-H record/playback device with VSI-H connector breakout module (not part of delivered system).
Internal correlator monitor and control bus (Station, Baseline, Phasing Boards). Sec. 8.3.2, 8.3.3	Interface to Station, Baseline, and Phasing boards. Will use 100 Mbps Ethernet. Will use embedded PC/104+ "CMIB".	Station Board data products output from this interface (auto-corr, sub-band power, quantizer statistics, phase-cal) internally, but through the Backend computers externally.
Internal data output interface (Baseline Board). Sec. 8.2.25, 8.3.2	Baseline Board data output pipeline on Gigabit Ethernet.	Wideband output with delivered output data rate of ~100 Mbytes/sec from each Baseline Board. Potential for upgrade to ~800 Mbytes/sec (10 Gigabit Ethernet) output capacity from each Baseline Board.
Correlator system monitor and control interface. Sec. 8.3.3	Network interface to higher-level control computers. 100 Mbps Ethernet. Refer to (Vrcic)	Virtual correlator interface to allow high-level configuration, control, and monitoring.
Correlator system data output interface. Sec. 8.3.3	Straw-man concept is Gigabit or 100 Mbps network connections. See Figure 8-3.	In the straw-man concept, correlator Backend computers perform FFTs, excise interference, and integrate.

8.8 Risk Assessment

Table 8-6 Areas of risk, and planned risk mitigation strategies in descending order of importance.

Risk	Risk Mitigation
Personnel	All engineers have been hired and appear to be working well on development. It is possible that if the schedule slips, more engineers will have to be hired or contracted to do development work, at some extra cost to the project.
Speed (256 MHz clock rates)	Mentor Graphics high-performance development tools are being used and it is now demonstrated that FPGAs are able to perform functions at the required speeds. System signaling and synchronization are designed to accommodate skewed clocks and data, and only the high-speed data paths run at the full rate of 256 MHz.
Correlator Chip	The first prototype Baseline Board will now be tested with a full-function and speed ASIC prototype. Preliminary power dissipation and cost estimates indicate that the 2048-lag chip is a 4 million gate design, will dissipate about 3.5-6.0 W, and will cost about \$2.3 million USD including NRE.
Filter Chip	The plan is to prototype in an FPGA and then do a conversion to a cheaper and lower-power gate array, if necessary. Currently, the design appears to fit in a Xilinx Virtex-4 SX35 device which may be cheap enough and cool enough for production. The FPGA would eliminate the design risk leaving only the reliability of the SX35 as a risk item.
Personnel turn-over	Define and enforce documentation standards to minimize single person dependencies.
Disruptive ground loop noise	Use differential signaling from rack-to-rack. Use common-mode noise filters on all module 48 VDC lines. Use large, low-impedance shunts between racks. Use signal-cable common-mode chokes if necessary.
Major supplier insolvent	Altera, has become a major player in the high-speed FPGA market and we are using both Altera and Xilinx devices as required. FPGA design is largely device/manufacture independent, thus the loss of a major supplier is no longer seen to be any risk. Cabling is no longer thought to be a risk area as well since a solid, industry-standard product, has been identified.

8.9 References

- Carlson, B.R., Dewdney, P.E., Efficient wideband digital correlation, Electronics Letters, IEE, Vol. 36 No. 11, p987, 25 May, 2000.
- Carlson, Brent, 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.
- Carlson, Brent, WIDAR Correlator Sensitivity Losses, NRC-EVLA Memo# 011, January 30, 2001.
- Carlson, Brent, A Closer Look at 2-Stage Digital Filtering in the Proposed WIDAR Correlator for the EVLA, NRC-EVLA Memo# 003, June 29, 2000.
- Carlson, Brent, Simulation Tests to Quantify the Spectral Dynamic Range and Narrowband Interference Robustness of the WIDAR Correlator for the EVLA, NRC-EVLA Memo# 009, Nov. 1, 2000.
- Carlson, Brent, Refined WIDAR EVLA Correlator Architecture, NRC-EVLA Memo #014, October 2, 2001.
- Carlson, Brent, Simulation Tests of Phasing Subsystem Signal Processing in the WIDAR Correlator for the EVLA, NRC-EVLA Memo# 008, Nov. 7, 2000.
- Carlson, B.R., Dewdney, P.E., Burgess, T.A., Casorso, R.V., Petrachenko, W.T., Cannon, W.H., The S2 VLBI Correlator: A Correlator for Space VLBI and Geodetic Signal Processing, Publications of the Astronomical Society of the Pacific, 1999, 111, 1025-1047.
- Carlson, Brent, An Analysis of the Effects of Phase Dithering in a Lag-based Fringe-Stopping XF Correlator, NRC-EVLA Memo# 002, May 26, 2000.
- Carlson, Brent, Requirements for 8-bit Processing in the Proposed WIDAR Correlator for the EVLA, NRC-EVLA Memo# 010, January 29, 2001.

Carlson, Brent, Simulation Tests of Sub-Sample Delay Tracking in the Proposed WIDAR Correlator for the Expanded Very Large Array, NRC-EVLA Memo# 007, October 3, 2000.

Crochiere, R.E., and Rabiner, T.R., Multirate Digital Signal Processing, Prentice-Hall, New Jersey, 1983.

Morgan, Thomas R., Requirements and Functional Specification, EVLA Correlator Backend, A25252N0000, Rev. 1.0, September 23, 2003.

Rowen, B., WIDAR Correlator Backend Processing Options, NRAO, Socorro, 2001, November 19.

Vrcic, S., Protocol Specification, EVLA Correlator Monitor and Control—Virtual Correlator Interface (VCI), A25201N0000, October 28, 2004.