Supplementary Report on the EVLA/WIDAR Correlator Critical Design Review

24 December, 2008

- CDR Review Committee -

Roger Cappallo, Dave Hawkins,
Barry Clark, Mike Revnell

During the course of the EVLA WIDAR Correlator Critical Design Review many fruitful discussions were held, points were raised, and many questions were answered. Although the (top level) summary review of the committee has already been submitted to and accepted by Greg Fahlman at HIA, it was felt by the committee that some less formal, but more detailed comments may prove helpful to the design team.

On-the-sky (OTS) testing

OTS demonstrated many features of the Station and Baseline boards, and the control software (which is still in development).

OTS tests produced a high-dynamic range image (73000:1) of a calibrator with the image center offset from the source, with only 4 stations and a 2 hr track. The source was correctly imaged, and no artifacts were visible at the image center (which can be indicative of signal processing errors). OTS tests produced a blank-sky image that was clear of any false images. Multiple source tracking showed that phases behaved as expected (with the exception of a residual term).

Issues arising from the OTS tests

- the lack of functionality on some boards (due to known hardware issues)
- Station Board Stage 3 filter issues
- different LO frequency offsets did not have any effect
- phase-residuals
- phase-glitches at scan boundaries
- data dropouts (low rate of occurrence)
- frequency spikes (low rate of occurrence, perhaps 10^-6 duty cycle) in correlator "magic channels"

The phase residuals and glitches are hypothesized as being due to errors in the software control of the delays. It is also possible that there are epoch uncertainties in the model's alignment to the data. The frequency spikes are likely due to an error in the Station Board filter software control. The data dropouts were low priority and their causes have not been analyzed (other than to hypothesize that the issue is CBE software related).

Brent and Mike Rupen were going to discuss these issues further on Dec 4th. None of the issues were considered to have a high-probability of being hardware related, so did not impact board sign-off.

Mike's presentation contains a number of OTS tests that still need to be performed. However, those tests are not considered a prerequisite to board production.

Station Board

An interesting note from the Station Board discussion was that the digital filtering logic on the Station Board is considered to have higher complexity than the WIDAR correlator ASIC.

OTS testing had problems configuring the Stage 3 filters. Dave Fort has tested each of the filtering stages of the boards in lab testing, so the OTS issues are considered to be related to control software.

Several I/O interfaces in the correlator system use a novel scheme of capturing single-rate data on both rising and falling clock edges (as would be used for double data-rate data capture), and then using the edge that yields error-free results. The scheme was questioned as an unnecessary complication, given that timing analysis should be able to yield the ideal location for single data-rate capture. Furthermore, empirical determination of the best edge can lead one to choose an edge near the transition, with the result that at some point in the future, under different thermal or aging conditions, a race-condition will ensue. Arguments in favor of the double-clocking scheme (put forth by Dave Fort) was that it was a more automated way of obtaining valid timing across the I/O interface, and that the use of the CRC ensured that if there was an issue with this approach that it would be detected as a CRC failure.

The delay module connectors were identified as being an issue. The connector pins were found to be too short, and reflow soldering did not adequately attach the connectors to the boards. A compatible AMP connector, with longer pins, will replace the problem connectors on the production boards.

The mechanical attachment of the delay module to the Station Board uses two connectors. Hawkins commented that there was no mechanical attachment associated with one of the connectors, i.e., no screws plus standoffs either side of the connector. Apparently attempts were made during later PCB routing stages to add screw holes, however, it was too difficult, and so the screws were not added.

The FPGA’s are cooled via a monolithic heat-sink, machined to lie across the tops of all of the chips. It was pointed out that during operation one might expect thermal gradients from top to bottom, as well as radially outward from the chips, and that this might cause warping of the plate, and a loss of good thermal contact if the thermally conductive pads are unable to expand.

During testing of multiple Station Boards, there were multiple occurrences of FPGA 1.2V core supply module failure. The cause of this failure is possibly related to a 'bad lot' of modules from the manufacturer. There have also been multiple instances of a block of FPGAs de-configuring. The common aspect of the block of FPGAs is their 1.2V core supply power segment. These two power supply issues could be related, and further investigation of the issues is recommended.

Station Board production can commence once the power supply issue is understood and resolved.

Baseline Board

Extensive lab testing of the v2.1 boards is complete. One board was used for OTS testing. No OTS issues were associated with the baseline board.

The phased-array mode of the Baseline Board has not been implemented. The risk associated with not implementing this mode before production is not considered high enough to delay production. If there are issues with the scheme currently envisioned, there are several alternatives to implement phased-array mode, enabled by the daisy-chained data output coming from the baseline board.

There were a small number of revision 2.1 Baseline boards available for testing. Brent's debugging of these boards indicated several minor changes were required to the design. The revision 2.2 boards are ready for pre-production (an 8 board build).

Baseline Board production can commence once multiple-board tests have been performed using the pre-production boards.

System-level Comments

During review of the WIDAR documentation, it was noted that there was no electronic support to limit the 48V supply in-rush current during hot-swap insertion, or power-on of the boards. When any of the WIDAR boards are powered on, there will be high current on the 48V rail as the input capacitance on each of the power supply modules used in the system are charged. The only limit to this current is the wiring and connector resistance between the 48V supply and the module capacitance. The peak of the current could exceed the maximum expected current during normal board operation. The current transient is most damaging during hot-swapping of a board, since the current flows as soon as power supply contacts touch, and thus when the contacts have highest resistance (since they are not fully seated).

Fortunately, the MTBF analysis indicates that boards will be replaced infrequently (72K hrs for the SB; 56K hrs for the BLB), and the intended minimal number of power cycles will reduce the impact of this design issue. To further mitigate the potential for board or connector damage it is recommended that Baseline and Station Board per-board breakers are thrown (power turned off) before boards are removed, and that the Cross-bar board

backplane breaker be thrown before any Cross-bar board is removed. High inrush currents will still be experienced by the system whenever the power is cycled, but since each board power-entry connector will be fully seated connector damage is unlikely.

Prior to attending the review, the thermal design of the system seemed to us like a potential for risk. However, after review of the results of extensive thermal testing, and viewing the EVLA correlator room with its massive water-cooled under-floor A/C and powerful rack fans indicates that the thermal aspects of the system are well understood and under control.

NRC will deliver loaded boards. NRAO is responsible for ordering any replacement parts. NRAO should procure lifetime buys for the FPGA's, DDR RAM's, and other IC's that are expected to be needed for repairs/rework over the projected 20 year lifetime of the correlator.

Miscellaneous Questions and Their Answers

- Has slaved mode of 4 BLB's ever been tested? – *master/slave mode is no longer part of the design*
- Have the gigE PHY chip and complete serial I/O chain been tested? - *yes*
- Are 5% spares really enough? Is this in addition to an initially working system, which will require some extra boards to be built? – *the board manufacturer is building an additional 3% to cover unusable boards. NRAO have ordered additional spares, to approach 10% total.*
- Has a general analysis of the effects on dynamic range of limited valid sample counting been done? (only 2 lags are counted in the correlator chip – I worry about propagation of blanked regions, both interference and time and channel codes, giving complex "shapes" of valid sample counts) – *the correlator chip uses biased arithmetic, but a '0' is added in for invalid samples, so an exact count of valid samples is not necessary*
- Can 3 bit mode be tested somehow with a truncated 8 bit data source? – *no, but all data are treated the same within the correlator regardless of original sample size, due to resampling*
- Has there been any OTS test at a high delay rate? -- *No, not in the OTS tests. Test vectors have been run at high bit shift rates in Penticton, but with inconclusive results. Brent plans to revisit this.*
- Has the system been tested with large fringe rates, or large phase accelerations? – *same answer as above. Tests were inconclusive on the SB. Independent tests performed on the BLB conformed with software predictions.*
- In the station board, are the FIR filter coefficients dynamically loaded? (they would need to be, if one wants to track the residual delay & remove phase slope) – *in principle they can be, as the station board has a delay generator, but that HDL code is not currently written*

- Is there a delay vernier bit? (if not, the delay errors on a baseline basis are twice those in the text, i.e. +/- 1/16 sample across sub-band) – *no, and this is in fact the error on a baseline basis*
- Can negative delays (e.g. for observations from spacecraft beyond the Earth's center of mass) be accommodated? – *yes, by offsetting the 0 point.*
- Autocorrelations are time shared – is the duty cycle 50%? How is the switching implemented? – *due to control software issues involved in reconfiguring the correlator blocks, there is some dead time when switching the antennas to be auto-correlated. The exact details are TBD.*
- Are only 8 of the 16 baseline racks used? – *yes, with the new resized array, and new connectivity scheme, only 8 racks are needed for up to 32 antennas*
- Has phase switching been tested? What mechanism ensures synchronicity of phase flip and its removal? – *there are no plans to use phase switching in the array*
- Is phased array mode still a requirement? If so, how do complex gains per antenna get applied to the sample data? – *yes, though that part of the development has been postponed, and is considered primarily a firmware issue.*