EVLA Online Calibration Summary

Adam Deller, 08-Feb-2012

Preamble

This document describes the calculation and application of calibration solutions in the EVLA WIDAR correlator in real time ("online calibration"). Online calibration is a prerequisite for phased EVLA operations, but is independent of the summing, framing, transmission, and the realtime processing/recording of the phased array data. Accordingly, it can be fully tested in the absence of any actual phased array operations, and in fact this is the most efficient way to do so. Inspection of the normal visibility data products of the EVLA can be used for this purpose.

System Overview

The EVLA has two independently tunable baseband polarisation pairs, each spanning 1 GHz (using 8 bit samplers, now operating) or 4 GHz (3 bit samplers, to be commissioned). A common local oscillator (LO) is used for each baseband pair, i.e., there are two LOs. Online phase corrections can be applied to the entire baseband by way of LO modulation; however, since a single LO is used for both polarisations of the pair, both polarisations are constrained to receive the same corrections. No delay corrections can be applied at the baseband level. The station boards channelize these basebands into 128 MHz chunks (8 per baseband in 8 bit mode, 32 per baseband in 3 bit mode).

The next stage of the WIDAR correlator consists of 64 baseline board pairs. Each baseline board pair can process one of the 128 MHz chunks of spectrum channelised by the station boards (both polarisations). A baseline board can optionally filter the 128 MHz chunks down by factors of 2, yielding a final bandwidth of 128/64/32/16… MHz. This final chunk of spectrum is called a subband. Since there are 64 baseline board pairs, there can ultimately be up to 64 subband pairs, each with a maximum width of 128 MHz.

Online delay calibration can be applied at the RXP chip each baseline board. Phase corrections cannot be applied at the RXP chip. Thus, phase corrections can be made at only at the baseband pair level, and delay corrections only at the subband level. These are fundamental limits imposed by the signal processing design of WIDAR. Additionally, the present realtime calibration system is only capable of propagating a single delay solution per baseband to the baseline boards. This is not a fundamental limitation, but it is one which requires some coding work to overcome, as discussed below.

The calculation of the online calibration (TelCal)

A single subband from each baseband is be designated as the reference subband for calibration purposes. Solutions from this subband can be applied to all other subbands from the baseband.

As show in Figure 1, the correct calibration of a subband generally requires the application of both a delay and phase correction. However, due to the fact that a subband has a relatively small fractional bandwidth, the phase component of the correction can be approximated with a small delay adjustment, as shown in the lower left panel of Figure 1. Both “delay-only” and “phase+delay” online calibration are implemented in TelCal; each has some drawbacks. They can be selected using the appropriate invocation of registerPhasing() in the executor script. The two methods are discussed in turn below.

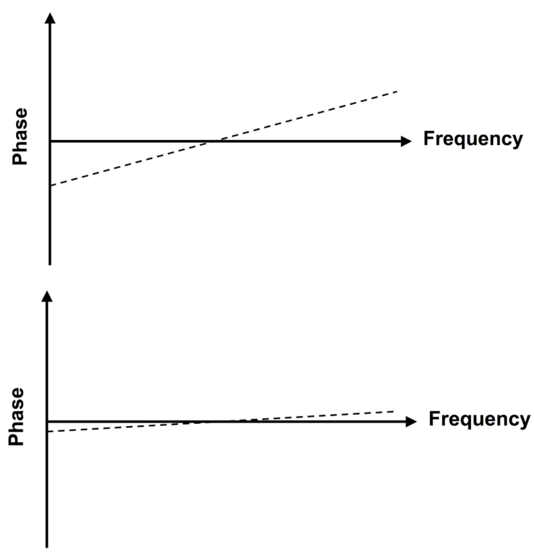

Figure 1 -- ***Top left:*** The uncorrected phase across a subband. ***Top right:*** With a perfectly corrected delay, a constant phase error remains Delay+phase calibration would adjust this phase to zero. ***Below left:*** By adjusting the delay correction, the mean phase can be set to zero without a phase correction; a residual phase slope remains.

Delay-only calibration

In this mode, TelCal solves for a delay-only correction in the reference subband. As noted above, in general the correct delay solution will leave a phase offset. Accordingly, the delay solved by TelCal in this mode differs slightly from the true delay, and leaves zero phase in the centre of the band, but a small residual slope across the band (lower left panel of Figure 1). In the worst case (low end of L band; 1000-1128 MHz, phase offset 180 degrees) the phase errors will reach +/- 24 degrees at the edges of the band – sufficient to cause moderate degradation of the sensitivity and bandpass of the subsequent summed output. Worse, the residual slope means that other subbands corrected using the same delay will be progressively more incorrect, as their center frequencies move further from the center frequency of the reference subband. Thus, only a single subband can be correctly phased per baseband. The effect is most severe at L band, but would be noticeable in all EVLA bands.

Delay-only calibration appears to work stably (~4 scans to stabilise; 10 second scans provide sufficient S/N for a typical calibrator source). However, since only one subband per baseband can be correctly calibrated using this method at the present time, this mode is currently unusable for wideband phased array applications. At most, 2x2x128 MHz can be correctly calibrated online using this method; sufficient for VLBI at 2 Gbps, but insufficient for pulsar searches. Additionally, the residual phase errors mean that this method is generally undesirable at L band and S band.

Delay + phase calibration

In this mode, TelCal derives a delay + phase correction for each reference subband. This can then be applied to all subbands from the same baseband. So long as the phase response is linear across the entire baseband (i.e., the bandpass is linear and there are no dispersive delays) this approach should exactly correct all of the subbands. Thus, delay+phase calibration is a more accurate approach than delay-only calibration.

However, there are presently two drawbacks with this approach. The first is that the algorithm has proven to be unstable in testing, usually taking a long time to settle (~8 scans) and sometimes returning to an unstable state after settling. A recent code change implemented by Barry may have partially or fully corrected this shortcoming – Amy will undertake testing to determine if this is the case. The second drawback is linked to the method of phase correction. Since this is applied using LO modulation, the same correction must necessarily be applied to both polarizations of a baseband pair. Since they will in general require different corrections, this effectively means that only one baseband from each baseband pair will be usable, halving the overall bandwidth which can be phased.

Since only 2 basebands can be phased using this approach, the total bandwidth which can be phased using the 8 bit samplers is 1 GHz (dual polarization). For VLBI, this offers the advantage of potentially selecting multiple narrower subbands, and avoiding the decorrelation at low frequencies inherent in the delay-only calibration approach. However, 1 GHz of bandwidth is insufficient for the requirements of the pulsar projects.

Readiness of online calibration for pending projects

Case 1 – VLBI: The present delay-only calibration should suffice for most VLBI utilisation. However, delay+phase calibration would be preferable, as it is more accurate at lower frequencies where the fractional bandwidth of a subband is large, and it allows more flexibility in the selection of subbands. In particular, 32 MHz subbands could be used to match those produced by the new RDBE systems on the VLBA, easing the complexity of the setup of the VLBA DiFX correlator, and allowing non-contiguous subbands.

Case 2 – Pulsar searching: Neither the current delay-only calibration nor a delay+phase calibration (even if it is proved to now be robust) can calibrate the full 2 GHz of bandwidth available with the 8 bit samplers. Only an upgraded delay-only calibration, in which different delay corrections can be propagated to the different subbands, will meet the needs of these projects.

Requirements for full-bandwidth calibration (delay-only)

The calculation of the correct delay for each subband based on the reference subband is trivial. However, propagating these calculated delays to the baseline boards is not currently possible – only a single value of delay can be propagated for each baseband, and is sent to all baseline boards which process subbands derived from that baseband. According to Barry, modifying this behaviour such that a different delay can be sent to each subband (baseline board) is relatively straightforward. He believes that the more difficult task is to provide the executor with the information associating each baseline board with a given subband. Once the EVLA executor can obtain this information, it can send the correct delays to the correct baseline boards.

Barry believes that the simplest case occurs when the subband bandwidth is 128 MHz (and so there is a 1:1 match between baseline boards and the original 128 MHz chunks from the station boards). He intends to attempt an implementation for this specific case. Happily, this is a case that is immediately useful for wideband phased array work, such as the pulsar search projects. If this implementation proves successful, the pulsar search projects could begin.

Actions and conclusions

Amy: Test Barry's latest delay+phase calibration to see if the stability has improved. Once working correctly, this mode should be used for any phased array projects unless wide bandwidths are required (i.e., those in which all four basebands must be phased).

Barry: Attempt an implementation of subband-based delay-only correction for the case of 128 MHz subbands. Success here would enable the pending wideband phased array projects.

??: Investigate how to provide the necessary information to Barry/the executor to allow association between subbands and baseline boards, such that the general case of delay-only correction can be solved. This could be seen as a longer-term aspiration, since the cases above should cover 99% of the planned phased array science.