

9/15/2011

To: UVLBI Group

1

From: Geoff Crew

Subject: Current RDBE-S Status

Overview

This memo briefly describes the current (early Sept 2011) status of the effort to develop a wide-band Roach-based digital back end (RDBE-S). The initial development and testing is covered in a previous memo (number 025) in this U-VLBI Memo Series in conjunction with a higher data rate burst-mode recorder system. Briefly, the RDBE-S sacrifices the channelization and polyphase filter bank of the VLBI2010 RDBE development effort in exchange for the processing headroom to support a second iADC card and doubled bandwidth. At the same time, it is straightforward to carry out the correlation using the DiFX software correlator.

With the advent of the Mark 6 16 Gbps recorder system, it is expected that a pair of RDBE-S units feeding a Mark 6 recorder can be deployed at a station to capture 4 GHz of input I/F bandwidth in 4 512-MHz "wide" channels. Since experience with wide channels (512 MHz) is limited, an important developmental step is to test a pair of RDBE-S units in a zero-baseline test. An initial pass at this step was performed with modified RDBE "development" units at the May 2011 Technical Operations Workshop (TOW) as described in memo 005 of the Mark 6 memo series.

In this memo we report results of more extensive testing with two of the RDBE-S units such as might eventually be deployed with a more mature version of the RDBE-S firmware "personality".

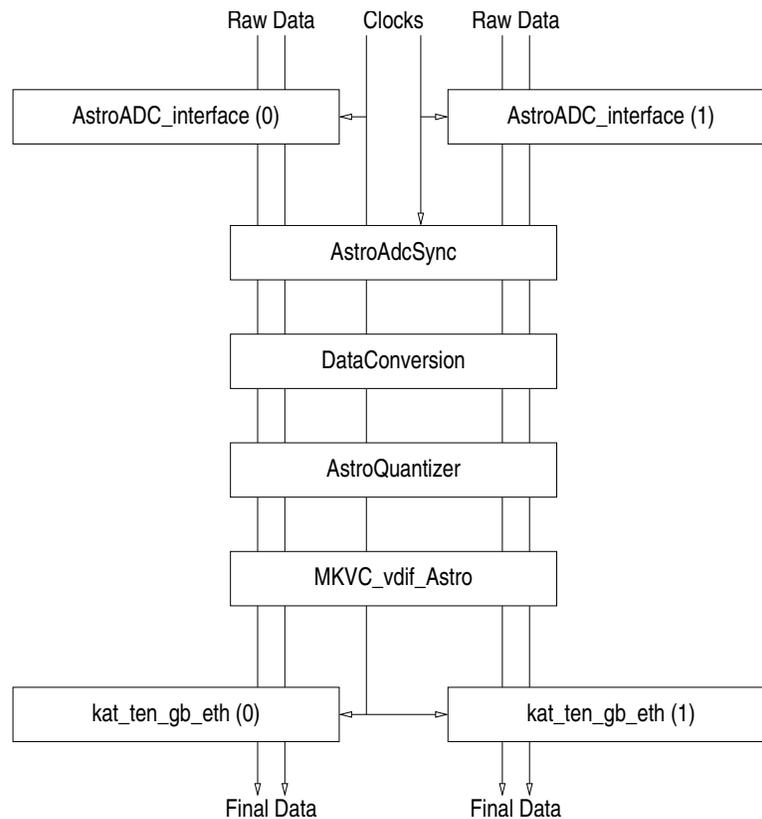


Figure 1: A simplified diagram of the Astro 8 Gbps design.

The Astro 8 Gbps Personality

A simplified diagram of the Astro design is shown in [Figure 1](#).

While the basic design of the FPGA personality remains unchanged, there have been a number of fixes to which became necessary to achieve satisfactory performance on "production" RDBE-S units. In addition, additional test code was introduced to verify various aspects of the internal modules. The testing described in this memo was carried out using the `RoachAstro8Gbps_sep02g.bin` binary which was built from a clean checkout of a branch in the SVN repository (`trunk/TopLevelBuilds/RoachAstro8Gbps_aug1` on that date. At this time, the only deficient part of the design is the 2-bit quantizer, which assumes that the distribution of ADC samples has zero mean. (As will be discussed later, the ADC provides 8-bit samples, but only effectively 6 bits, so the lowest bits are noise and not necessarily of with mean.)

Since each of the iADC cards has a separate (1024 MHz) sample clock, the raw, 8-bit samples from the iADC cards are captured in two "interface" modules as shown in [Figure 1](#). These are passed to a "synchronizer" which allows all subsequent processing to proceed on a common (256 MHz) clock derived from the sample clock of the 0-th iADC card. This data is then converted and quantized to 2 bits (as in the VLBI2010 personality) before being passed to a module that packetizes the data into VDIF packets. A number of infrastructure and test modules (*i.e.* "raw capture" have been omitted for simplicity. Each IF datastream is assigned to its own VDIF "thread", and in this testing each thread has been labelled as if it were captured at a different (co-located) station.

The packet header is as shown in [Figure 2](#).

	31		Byte 3		24 23		Byte 2		16 15		Byte 1		8 7		Byte 0		0
Word 0	0	0	Seconds from Reference Epoch														
Word 1	0	0	Reference Epoch			Data Frame # within second											
Word 2	0	0	0	0	0	0	0	0	Data Frame Length (octects; 8224 = 0x404)								
Word 3	0	0	0	0	0	0	1	Thread ID				Station ID					
Word 4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Word 5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Word 6	Packet Serial Number (PSN)																
Word 7	0xA		0xB		0xA		0xD		0xD		0xE		0xE		0xA - 0xD		

Figure 2: VDIF Header as used by the `Astro8Gbps` image. The red-shaded areas are configurable, the blue-shaded areas are dynamic (*i.e.* changing with each packet), and the green areas are fixed by the personality.

These 8192-data-byte packets (31250 per second per thread) are then emitted as 4 separate "threads" on 2 physical 10 GbE interfaces.

In previous testing, a single RDBE-S unit was tested with simple tones, broadband noise, or a combination of these, as shown in [Figure 3](#).

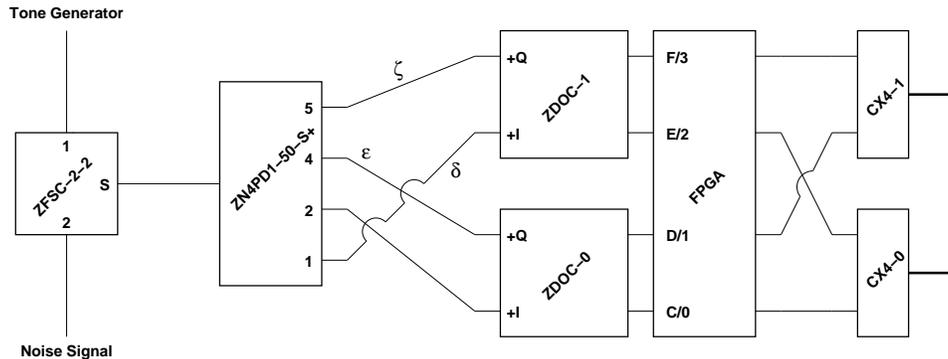


Figure 3: In single RDBE testing, 4 copies of noise or tone signals are presented to the I/F inputs on the two iADC cards which are connected to the FPGA on the Roach board by 2 zdoc connectors. The packetized VDIF data is then output onto two CX4-1 cables.

For convenience in testing in this design the "I" and "Q" inputs to the two iADC boards are crossed into the output ethernet streams. This allows only one CX4 cable to be connected for simple system checks.

In this test configuration, the four threads can be considered as four separate "stations", c , d , e and f each receiving a single 512 MHz channel from the same source (via the Minicircuit splitters ZFSC-2-2 and ZN4PD1-50-S+). Since the paths to the four inputs have different physical lengths (delta, epsilon and zeta), the four "stations" should see a common spectrum, but with different corresponding single-band delays when analyzed with `fourfit` after correlation in a zero-baseline configuration with DiFX. An example of several spectra is shown in Figure 4.

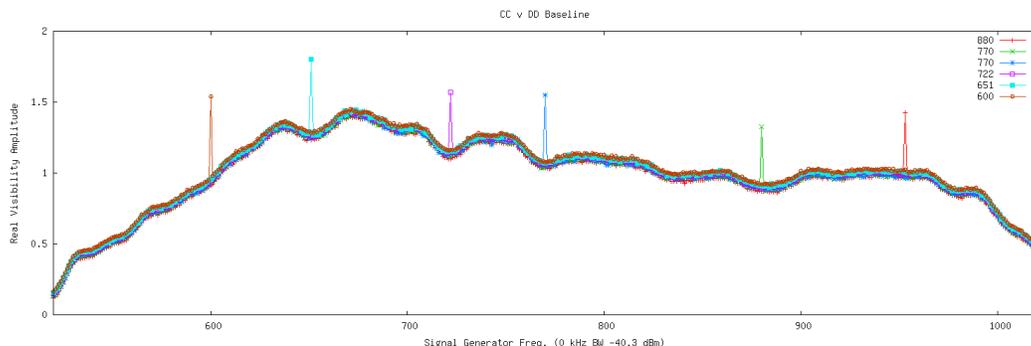


Figure 4: A number of pure tones (one per scan) were coupled with the noise source, captured by the RDBE-S and processed with DiFX. This plot shows an overplot of six such spectral from the DiFX "Swinbourne format" data files.

The data for this plot was collected on several sequential scans where the tone frequency was changed, but the noise source was left unchanged. Typically the signal is strong enough that only one or two seconds of data is required to achieve a high-SNR correlation.

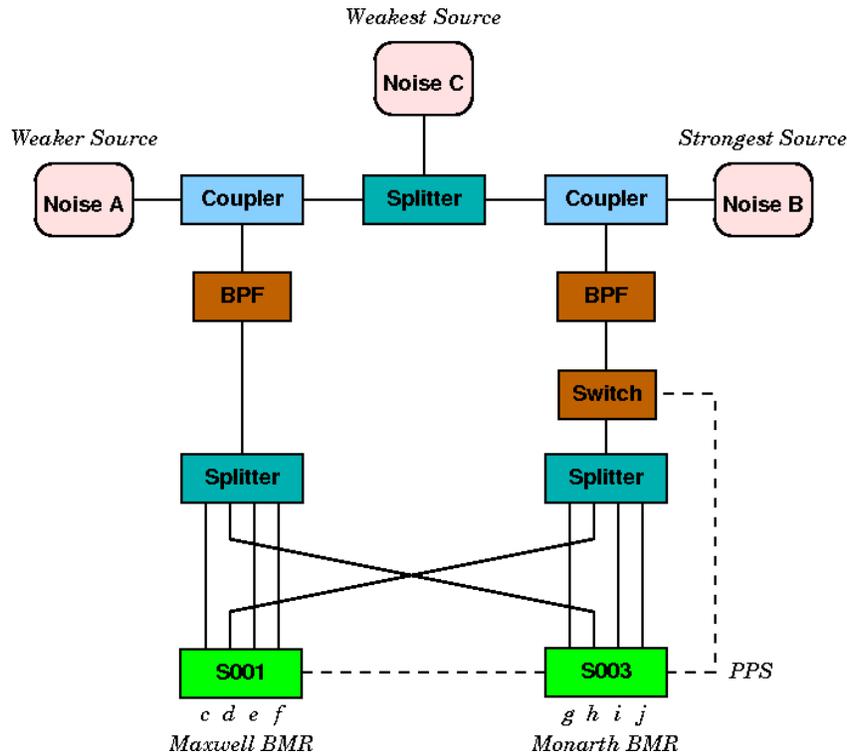


Figure 5: Three noise sources, simulating two receivers and one cosmic signal were connected to two RDBE-S units for zero-baseline testing.

Initial Two-unit Testing Setup and Results

For the testing of two units it was desirable to combine three sources into one "noise box". Two of the noise sources would simulate receiver noise of each of two stations, and the third would be an attenuated source simulating a small cosmic signal to be separated out through the zero-baseline test. Testing of all four inputs at each RDBE-S could be carried out as with the single-unit testing. Operationally, there were some difficulties setting the time in the earlier RDBE-S personalities, so it was useful to swap a pair of inputs as shown in Figure 5.

The two burst-mode recorders (*Maxwell* and *Monarth*) thus each collected 4 "stations" (c, d, e, f and g, h, i, j). A (non-maser) 5 MHz clock and PPS were distributed to the two RDBE-S units. With a logic analyzer, the "dotmon" outputs of the FPGAs were found to be about 30 ns apart—thus single-band delays on that order were to be expected from the eventual analysis.

It was desirable to observe approximately 1% correlation (about 20 dB total attenuation of the common noise signal relative to the foreground noise sources), so the weakest noise source was chosen for the common, cosmic simulator and coupled into the stronger sources with directional couplers (Minicircuits ZFDC-15-5-S) that reduced its power by approximately 15 dB. In fact this weaker noise source was substantially weaker and required an amplifier (Minicircuits ZX60-2514M) to have it begin to approach parity with the other

noise sources. Typical raw capture histograms from IF1 of each of the stations is shown in Figure 6

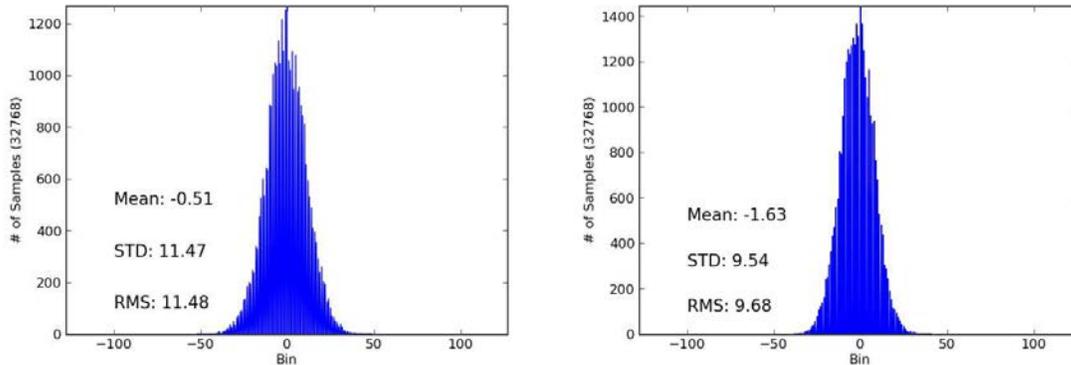


Figure 6: Raw capture histograms from IF0 on each of the two RDBE-S units are shown. They are comparable in strength, and the common, "cosmic" source is too weak to be noticeable.

As an aid to sorting out the timing issues, a switch was connected to "blank" the noise signals during the 30-ms PPS tick. Finally, second-Nyquist zone band pass filters were inserted before the final 1-4 splitters.

The results were not as expected, however. For example, a table of the amplitudes of the various baselines (in "Whitneys" == percentage times 100):

	c	d	e	f	g	h	i	j
c	10000.0	153.0	7510.0	7610.0	109.0	8510.0	38.7	40.0
d	-	10000.0	169.0	169.0	7490.0	171.0	9150.0	7580.0
e	-	-	10000.0	10200.0	119.0	8050.0	3.4	3.2
f	-	-	-	10000.0	117.0	8580.0	3.3	3.3
g	-	-	-	-	9630.0	111.0	7520.0	10300.0
h	-	-	-	-	-	9980.0	37.8	37.6
i	-	-	-	-	-	-	9910.0	7550.0
j	-	-	-	-	-	-	-	10000.0

Table 1: The amplitudes of the correlations of the baselines are not as expected.

should show comparable values (almost exactly 10000) on the common noise sets (c,h,e,f and g,d,i,j) and smaller values (about 100) on the other (intra-unit) baselines. As can be seen from the table, quite a number of the latter are of the correct strength, but others are significantly smaller. It turns out that the smaller-amplitude baselines are the ones involving the second ADC board. As this was an addition to the original VLBI2010 design, and as the clock from this design is the one that is ultimately discarded in favor of the clock from the first board, it was (erroneously, as it turned out) thought that there were issues with the design/implementation that bore investigation.

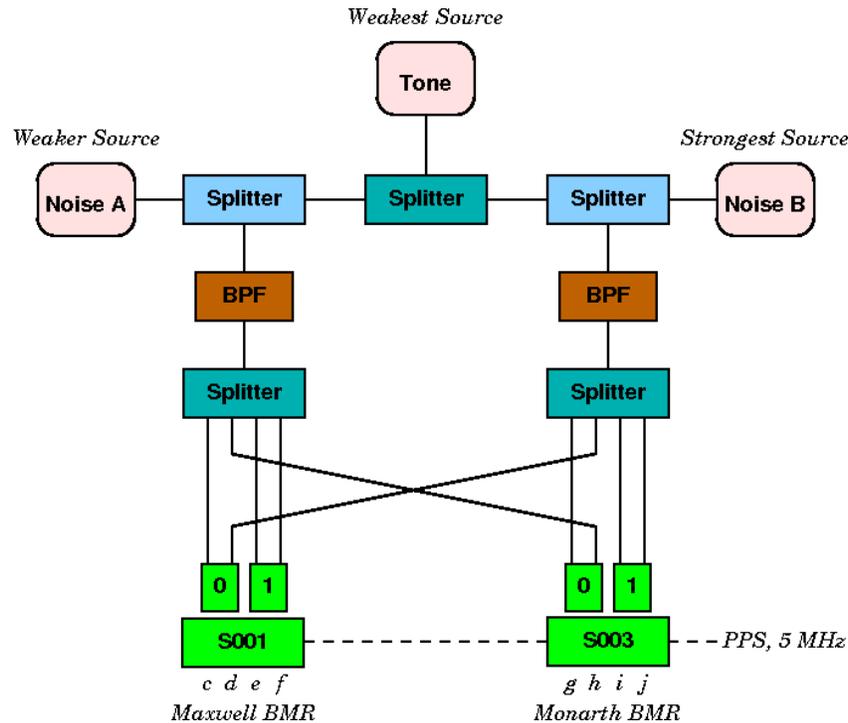


Figure 7: The previous configuration was simplified from that of [Figure 5](#) by the removal of the PPS switch and the replacement of the couplers with simple splitters to allow variation of the attenuation between "Noise C" and the other two noise sources.

Investigation of the ADC0/ADC1 asymmetry

It was at this point that the FPGA code was scrutinized and several testing capabilities added:

- The ability to blank the data using the internal PPS signal
- The ability to simplify the quantization to one-bit (*i.e.* only use 2 of the 4 states)
- The ability to load the packets in the MKVC_astro_vdif module with known data
- The ability to inject identical pseudo-random data sequences before the AstroAdcSync module

These tests turned up no surprises—the packet streams from the two units were found to be time synchronized, the one-bit data correlation results were similar (after taking loss of sensitivity into account); and the test sequences were found as expected (*i.e.* with 8 identical copies).

At this point the setup shown in [Figure 5](#) was modified as shown in [Figure 7](#) to inject simple tones or simple noise split 8 ways, and these tests likewise showed nothing obviously wrong.

At this point attenuators were inserted between the splitters so that the variation of the amplitudes of the intra-unit correlation amplitude could be measured with variation in the relative strength of Noise C. Results are shown in [Figure 8](#)

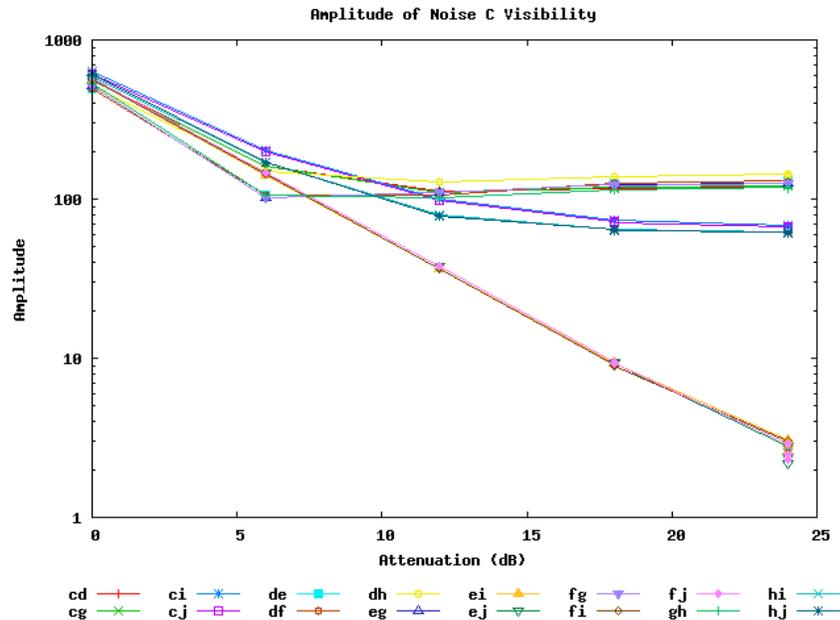


Figure 8: This plot shows the variation of 16 baselines with changes in attenuation inserted between the splitter pairs.

This plot shows the baselines which should show decreasing amplitude as the attenuation in front of Noise C is increased. In practice, only those baselines involving the second iADC in each RDBE-S show this behavior—the others saturate at about the 1% level with only a modest amount of attenuation.

The observation that the swapping of the d and h inputs also allows a coupling of the signals of Noise A and B *through* the first iADC board in each RDBE-S offers the most likely explanation. Whether this coupling occurs in the ADC chip, the iADC board proper, or the test setup has not been established. However, there is some reason to think that the iADC board does indeed suffer from the required level of cross-talk (J. Weintroub, private communication, and see also Haystack Mark5 Memo #048).

As a "quick" test, independent of the current development system and test hardware, we looked for this effect in recent data taken at the SMT0 observatory in March 2011. In that observation, a single DBE1 (with one iADC board) was used for two adjacent bands (hi and low) of the down-converted LCP receiver feed; a 2nd DBE1 was used for simultaneous RCP observations. It was straightforward (for Mike Titus) to substitute the hi band LCP data in place of the lo band RCP data and recorrelate. Since the nyquist zone 2 and 3 filters used for these bands are not perfect, there is some leakage near their common edge. However at the far sides of the band, the isolation should be at least 40 dB and so the correlation amplitude for the lowest 32-MHz bands should be nearly zero. In fact, an amplitude of 2% is observed as shown in Figure 9.

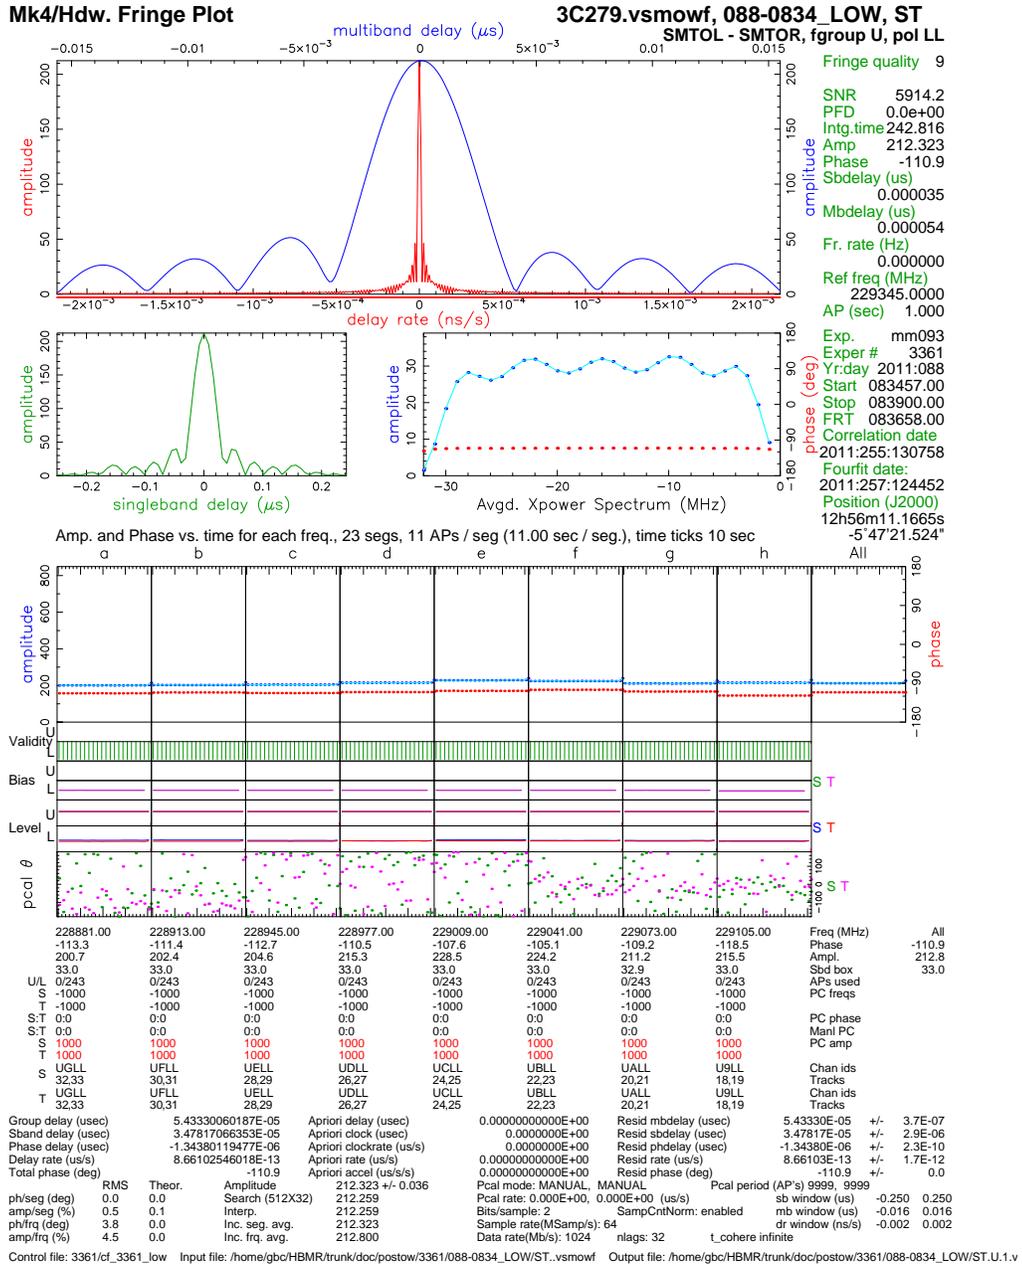


Figure 9: Fourfit plot of observed crosstalk between hi and low bands in lower channels as found in SMTO data.

Some Results of Zero-Baseline Testing with DiFX

Unswapping the *d* and *h* inputs results in the setup shown in Figure 10:

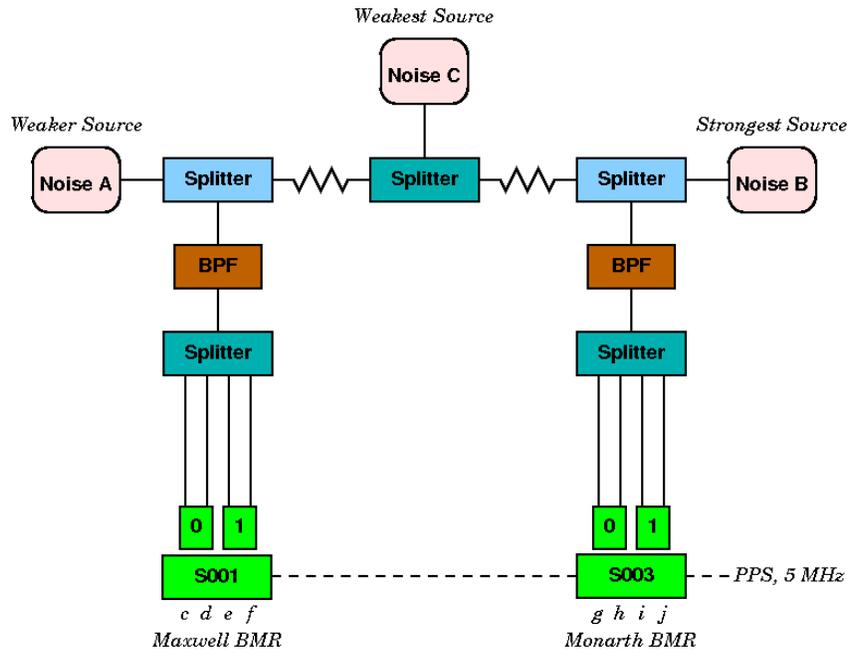


Figure 10: The previous configuration was simplified from that of Figure 7 by unswapping of the d and h inputs.

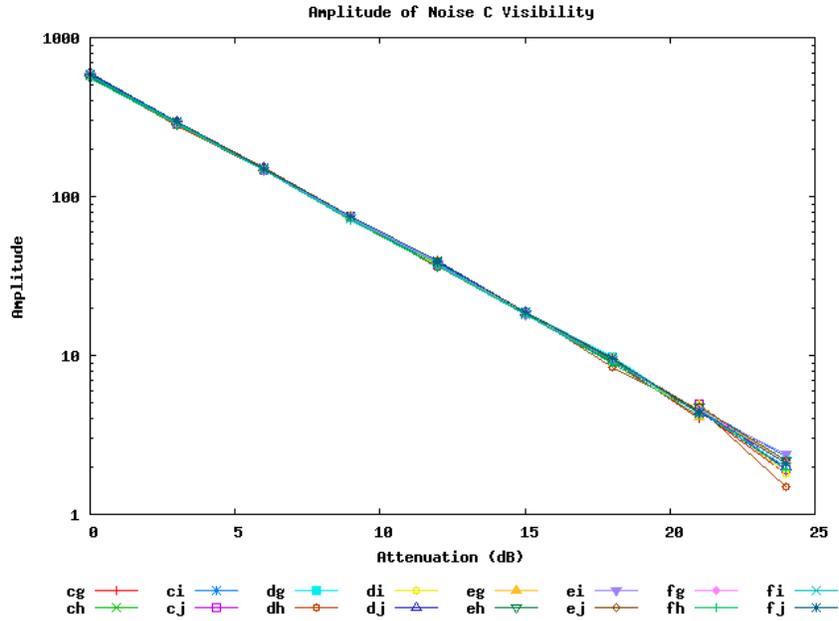


Figure 11: The amplitude of the Noise C cross-correlation visibility decreases as expected with increasing attenuation.

Now, when the attenuation between the splitters is increased from nothing to a total of 24 dB (each side) one observes the expected decrease in correlation amplitude as shown in [Figure 11](#)

At the highest attenuation levels, longer runs were needed to verify the correlation amplitude (at most 2 4-second scans). For example, with 18 dB attenuation, a 7-sec integration produces a baseline matrix:

	c	d	e	f	g	h	i	j
c	10000.0	9810.0	9790.0	9850.0	9.2	9.3	9.8	9.2
d	-	9990.0	9970.0	9930.0	9.8	8.5	9.1	9.6
e	-	-	9990.0	10600.0	9.7	9.1	9.5	9.5
f	-	-	-	9960.0	9.6	8.9	9.3	9.6
g	-	-	-	-	9930.0	9370.0	9750.0	10900.0
h	-	-	-	-	-	9690.0	9550.0	9370.0
i	-	-	-	-	-	-	9990.0	9730.0
j	-	-	-	-	-	-	-	9990.0

Table 2: The intra-unit baseline correlation amplitudes show approximately the correct value for 18 dB of additional attenuation on Noise C.

In this case, the smaller amplitude results have an SNR of about 60. The variation of the autocorrelations and the common-noise cross-correlations is not fully understood and bears investigation.

The single-band delay for this case is (in units of microseconds)

	c	d	e	f	g	h	i	j
c	+0.000	+0.001	+0.001	-0.001	+0.028	+0.029	+0.028	+0.026
d	-	+0.000	+0.000	-0.001	+0.027	+0.028	+0.028	+0.025
e	-	-	+0.000	-0.002	+0.027	+0.028	+0.027	+0.025
f	-	-	-	+0.000	+0.029	+0.030	+0.029	+0.027
g	-	-	-	-	+0.000	+0.001	+0.000	-0.002
h	-	-	-	-	-	+0.000	-0.000	-0.003
i	-	-	-	-	-	-	+0.000	-0.002
j	-	-	-	-	-	-	-	+0.000

Table 3: The delays between the baselines are approximately 30 ns as was expected from the relative PPS offsets.

which corresponds to the aforementioned 30-ns delay measured with the PPS signals into the respective units. The single-ns variations are consistent with the channel-channel timing variations that have been observed in single-unit testing. It has been suggested that the phasing of the sampling within the ADC chip itself is stable, but not fully deterministic following a reset. The variation of this delay with attenuation is available in [Figure 12](#)

Each RDBE-S was typically reset between insertion of attenuators when the attenuation was small (to give a clean re-quantization of the input signal). Once enough attenuation was applied, the RDBE-S units were not reset. Also, the sequence of attenuation was 24, 18, 12, 6, 0, 3, 9, 15, 21 (with some duplication at a few levels due to different length scans), so the observed oscillation is an artifact of the attenuation order. (*I.e.* the alternate points are stable, so the delay was in fact not changing with attenuation.)

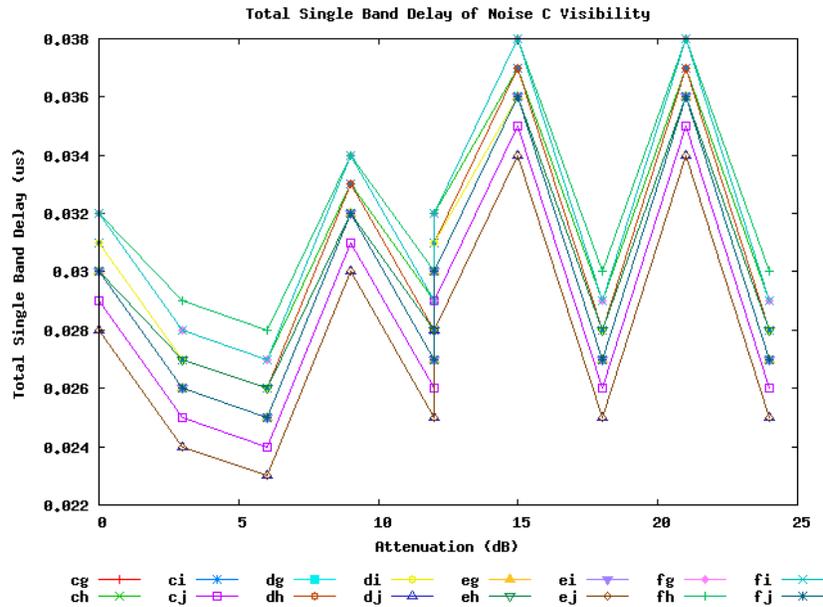


Figure 12: The single band delay of selected baselines is shown. The apparent oscillation with attenuation is a result of the order in which the data was gathered, with changes in delay following from resets.

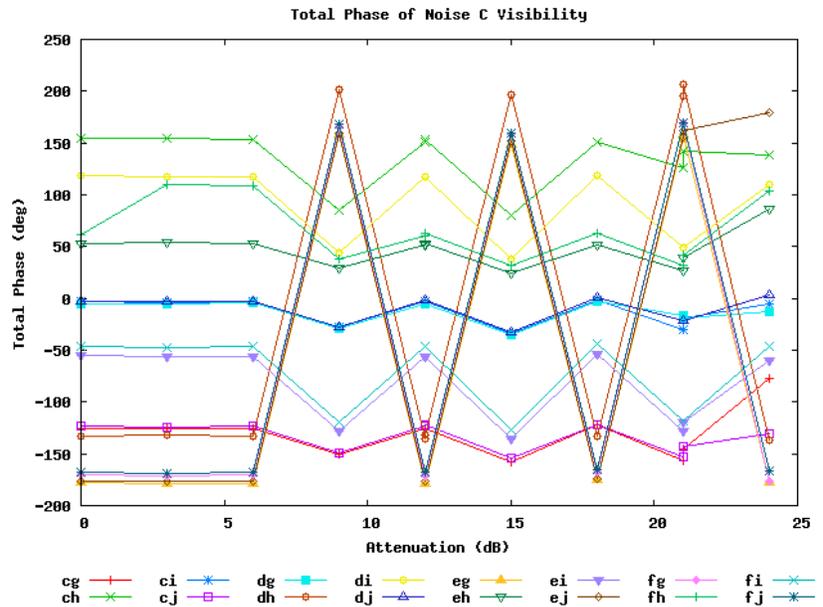


Figure 13: The total baseline phase for selected baselines is shown. As with the single band delays of Figure 12, the oscillation is not real, but results in a change of phase following a reset.

Finally, a similar behavior was observed with the relative phases of the baselines, shown in Figure 13:

Some Additional Testing

The departure of the fully-correlated signals from 100% needs to be (thoroughly) investigated. A few simple numerical experiments suggests that errors in the quantization can indeed produce several percent variations in the correlation amplitudes. This should be revisited once the quantizer is performing correctly. (I.e. once the apparently DC bias in the ADC histograms is removed prior to quantization.)

Although 8-bit values are returned by the ADC, only 6 of the bits are significant. In Figure 14 we show the some numerical results from limiting the ADC to only 6 bits. The numerically simulated data was constructed to have exactly 0.1 % correlation (i.e. a correlation amplitude of 10.0 in Fourfit's units). The three pairs of plots in the figure correspond to no truncation of the input data, and truncation into 64 states between 1.5 and 3.0 sigma prior to the normal 2-bit quantization. It is clear that this truncation injects noise at the DC edge of the spectrum and reduces the correlation amplitude.

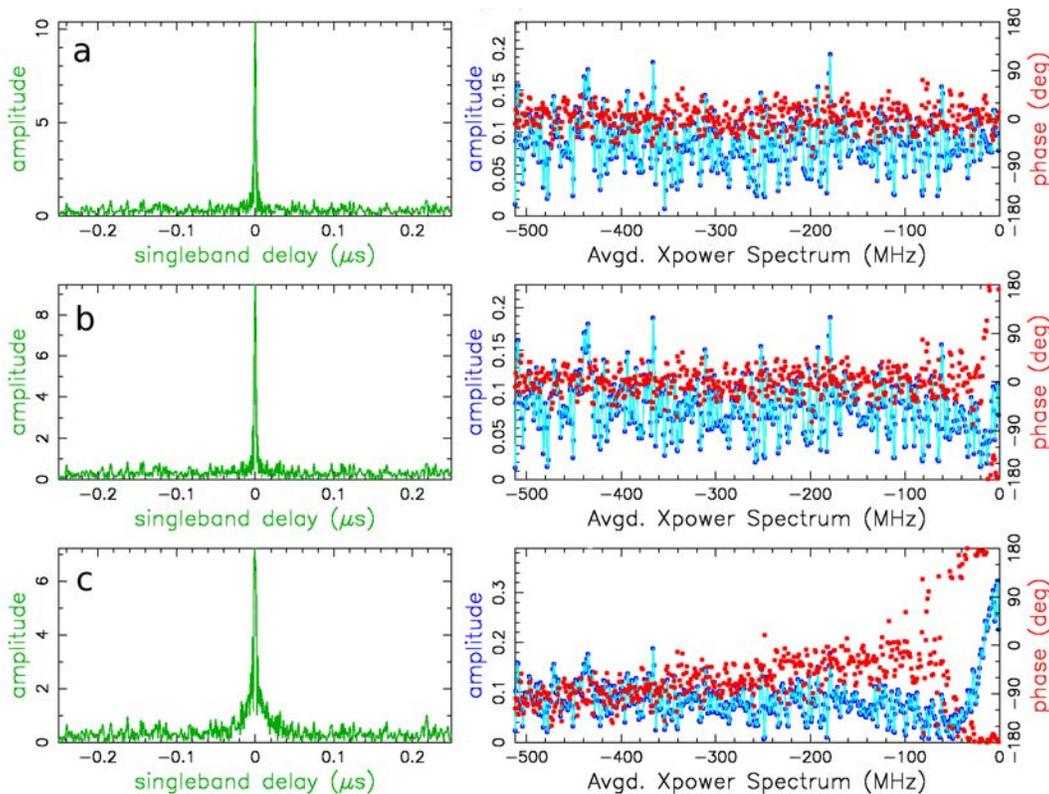


Figure 14: Single-band delay and cross-power spectra from fourfit plots are shown for a numerical experiments for a baseline pair that (a) has full precision, (b) where the ADC effectively has only 6 bits spanning -1.5 sigma to 1.5 sigma, and (c) where the ADC effectively has only 6 bits spanning -3 sigma to 3 sigma.

Finally, a version of the personality `RoachAstro8Gbs_sept14g.bin` was built which allows specification of the 3 thresholds separating the 4 bit states (rather than all the personality to set them). This version was run twice (runs 74 and 75), first with the default quantizer, and then with the thresholds as set by an auxiliary script

station_run	time	state:	00	01	10	11	
cc_74:	1316027525.000127077	BS[0]	0.170	0.318	0.360	0.152	(67.8% 32.768Ms)
cc_75:	1316027957.000092983	BS[0]	0.167	0.325	0.353	0.155	(67.8% 32.768Ms)
dd_74:	1316027525.000117064	BS[1]	0.164	0.344	0.351	0.141	(69.5% 32.768Ms)
dd_75:	1316027957.000121117	BS[1]	0.161	0.341	0.339	0.160	(68.0% 32.768Ms)
ee_74:	1316027525.000127077	BS[2]	0.160	0.322	0.355	0.163	(67.7% 32.768Ms)
ee_75:	1316027957.000121117	BS[2]	0.161	0.326	0.355	0.159	(68.1% 32.768Ms)
ff_74:	1316027525.000117064	BS[3]	0.184	0.328	0.338	0.149	(66.7% 32.768Ms)
ff_75:	1316027957.000078917	BS[3]	0.162	0.356	0.335	0.147	(69.1% 32.768Ms)
gg_74:	1316027525.000097036	BS[0]	0.169	0.350	0.351	0.130	(70.1% 32.768Ms)
gg_75:	1316027957.000642061	BS[0]	0.143	0.374	0.326	0.158	(69.9% 32.768Ms)
hh_74:	1316027525.000053883	BS[1]	0.192	0.346	0.344	0.118	(69.0% 32.768Ms)
hh_75:	1316027957.005271912	BS[1]	0.160	0.381	0.291	0.168	(67.2% 32.768Ms)
ii_74:	1316027525.000988007	BS[2]	0.171	0.329	0.363	0.138	(69.1% 32.768Ms)
ii_75:	1316027957.000051022	BS[2]	0.167	0.331	0.335	0.167	(66.5% 32.768Ms)
jj_74:	1316027525.000057936	BS[3]	0.183	0.316	0.343	0.157	(66.0% 32.768Ms)
jj_75:	1316027957.000051022	BS[3]	0.161	0.342	0.343	0.155	(68.5% 32.768Ms)

which produced a middle state fraction that is slightly better, but arguably still not ideal. And in any case, this produced no improvement in the DiFX correlation amplitudes. (*I.e.* that some of the autocorrelation amplitudes are reported as something other than 10000.0 is a bit of a mystery.)