Delay Jumps in the VLBA's Current RDBE DDC Personality

Jonathan D Romney
NRAO and LBO, Socorro

2018 September 15

Introduction

The ROACH digital backend (RDBE) was developed in 2007-09 as part of the VLBA Sensitivity Upgrade project, as a collaborative effort involving staff of the NRAO, Haystack Observatory, and the South African KAT project. Implementation, including samplers, an FPGA to support digital signal processing and output data formatting, and a control microprocessor, was developed by NRAO and KAT staff.

A polyphase filterbank FPGA personality (PFB), with 16 fixed-bandwidth 32-MHz channels (only 15 usable), was developed as Haystack's contribution to the RDBE firmware. It became available for regular VLBA observations in the 2011B observing semester. NRAO staff developed a digital downconverter (DDC) personality that became available in the RDBE for VLBA semester 2013B. Both the PFB and the DDC personalities required challenging, then-new FPGA processing to transmit control information into the FPGA, and to format the resulting spectra for standard VLBI recording.

Early tests of the DDC's performance indicated that *very infrequent* delay jumps were occurring. I conducted several rounds of tests in 2016, attempting to characterize these jumps. There appeared to be two different types of delay jumps. Approximately equal numbers of delay transitions occurred (1) between the delays in adjoining individual schedules, and (2) between the first and all subsequent scans at a given observing bandwidth within an individual schedule. However, the test data were inadequate for a clear characterization of these effects.

A thorough, 16-month rewrite of the DDC FPGA firmware was completed earlier this year. Testing of the new version required new VLBA station control software to be loaded for each test, and subsequently removed for ongoing observations. At the same time, a large-scale ramp-up of engineering effort on designs for the ngVLA proposal limited further development, and it was decided to go back to the previous FPGA code, but to keep several important upgrades that had been developed in the control software.

A more extensive test program was initiated in June 2018, on an expanded scale aimed to yield more definitive results than in the 2016 tests. This note summarizes the results of those tests.

Test Overview

A total of 12 systematically scheduled observations were carried out, at times with low demand for scientific observations, but nevertheless interspersed among those regular runs. Each observation comprised four identical repetitions of six pairs of setups, at bandwidths of 128, 64, 32, 16, 8, and 4 MHz. Each such pair used the same setup in both sequential scans, to establish whether delay jumps occurred in multiple occurrences of the same setup. All setups observed in right-circular polarization, and the 12 observations differed only in the sequence of bandwidths observed. The 128-MHz bandwidth setup was limited by well-known instrumental constraints to a contiguous four-channel frequency range using both RDBE units at each station. Setups at all the narrower bandwidths used eight DDC channels, also requiring the use of both RDBE units. Observing frequencies were set in ascending order within each setup; either the lowest- or the highest-frequency range was set at adjacent frequencies, while the other was set with inter-channel gaps. (Overlaps instead of gaps were required for the 64-MHz setup.)

A different *sequence* of setups was scheduled within each test observation, sequenced identically in the four repetitions. Care was taken to use each bandwidth setup an equal number of times, both overall within each run, and in particular in the first setup of each observation. Among the 12 test observations, each of the six setup bandwidths occurred exactly twice as the first scan. A partially inverse-bandwidth-dependent scan duration (factor of 1.5-1.6) was used to maintain adequate sensitivity at the narrowest four bandwidths without an undue contribution to the total observing duration. The 48 scans in each observation ran for an overall duration of 42 minutes and 16 seconds.

The total of 12 observations included 576 scans over 8.45 hours of total observing time. Very strong sources were selected for each individual observation, such that elevations were above 15º in all cases. Instrumental failures limited these observations to 6-9 stations, with an average of 8.07. (The much shorter 2016 test sequence mentioned in the Introduction also used only 8 stations.) The total number of station-scans for the recent 12 observations was 4608, or 2304 repeated pairs.

Analysis

I analyzed the results from the full set of 12 observations in the simplest manner I thought would be sufficiently complete. I used AIPS task FRING to solve for delays within each scan at each bandwidth, with fairly short integration times to validate the presence of fringes by multiple solutions within each scan, and task SNPLT to display the delay solutions. The delay shifts among the four scans at each bandwidth, in each cycle of the schedule, clearly showed at which stations delay jumps had occurred, in which channels (IFs) and scans, and of course the magnitudes of the jumps.

Since the reference station for such fringe fits can also undergo a jump itself, I did multiple fringe fits, with different references, in cases where a large number of jumps were seen. There were no surprising results in those cases. Delay jumps appeared simultaneously at as many as five stations.

Primary Results

The new test results substantially confirm most of what was suggested in the more limited 2016 tests, and delineate the effects occurring in the current system much more definitively.

No delay jumps occurred in the transition from the first to the second sequential scan at the same bandwidth, within any pair, at any bandwidth.

At bandwidth changes, two different types of delay jumps occurred. The primary type of jump is similar to those seen in the previous tests. It occurs in a predictable manner: A delay jump occurs **universally** in the transition from any setup with a bandwidth of 4, 8, or 16 MHz that observes in the **first scan** of the schedule, to a **second scan** at *any other* bandwidth. No such jumps occurred with a first-scan bandwidth of 32, 64, or 128 MHz. There were also two cases where a delay jump occurred in the transition from a **second** setup at 4 or 8 MHz to a **third** setup, but only when the **first** setup observed at 128-MHz bandwidth. In all cases, the magnitude of the jump is (256 ns) / (pre-jump scan BW in MHz). These two types of delay jumps occurred in 1.22% of the total number of station-bandwidth changes.

In comparison, the occurrence level of this type of jump in the 2016 tests was about 4%. The pattern of jumps from the first to the second setup in the schedule is also seen in a re-analysis of the bandwidths observed in common between those earlier tests and the current series, but was not evident within the small number of cases available at that time. The apparent special case with an initial 128-MHz setup could not be detected because that bandwidth failed in those tests.

The incomplete understanding of those results led to a misleading informal terminology that should now be abandoned: the delay jump does *not* begin with the "second use" of a given setup. It occurs immediately, at the beginning of the **second scan** in the schedule (with an exception for initial 128-MHz scans as described above). Although indeed present in the second use of the initial setup, it is also in place for any other setups that precede a second use of that first setup.

This type of jump generally occurred in subsets of the 8 channels. Occurrences were often limited to channel groups 1-4 or 5-8 (or both), which are processed in different RDBE units. There were even two cases where these two groups jumped with opposite signs at a single station. More general cases, limited to channels 7-8 alone, or to 5-6 and 7-8 simultaneously but separately, occurred in three of the 12 tests.

While these tests observed in a single polarization, it is expected that similar effects will occur, with independent jump amounts, in R/L polarization setups.

Finally, and unfortunately, a single delay jump occurred in a second category which does not fit the model just described. This occurred in the second use of the 4-MHz setup, in the **eighth** setup scheduled overall. This one event was 0.04% of all cases. And, for completeness: no delay jumps occurred at all in four of the 12 observations — all beginning with scans at 32 or 64 MHz bandwidth.

Details of the delay-jump occurrences described above are shown in the following table.

RDBE / DDC Delay Jump Occurrences in Test Observations TD108A-L

	Bandwidths [MHz]						BW Sequence						Delay Jump Details					
td108	4	8	16	32	64	128	1	2	3	4	5	6	IFs	Stations	IFs	Stations	Total Jumps	
a		1					8	64	16	32	4	128	5-8	BR/FD/KP			3	
b	8		1				16	4	128	8	64	32	1-4	MK	5-8	BR/FD/HN/OV	4	1
c							64	8	4	16	128	32						
d	1						4	64	8	128	32	16	1-4	FD/OV			2	
e			1				16	32	128	8	64	4	1-4	BR/KP/OV/SC	5-8	HN/SC	6	
f		2					128	8	32	4	16	64	5-8	BR/KP/MK/OV			4	
g	1						4	16	64	128	32	8	1-4	MK			1	
h		1					8	32	64	4	16	128	5-6	FD/KP/SC	7-8	HN/MK/SC	6	
i	2						128	4	32	16	8	64	7-8	HN/SC			2	
j							64	16	4	32	128	8						
k							32	128	8	64	4	16						
l							32	128	16	64	8	4						
																	28	1
																2304	1.22%	0.04%

KEY to BANDWIDTH table flags above

Blank: No delay jumps occur at this bandwidth in this observation.

1 Delay is discrepant at some stations, in some channels at this BW, in the FIRST SETUP of the entire observation.
Delay changes to match the other channels at those stations, for SUBSEQUENT SETUPS.

2 Delay is discrepant at some stations, in some channels at this BW, in the SECOND SETUP of the entire observation, following an INITIAL 128-MHz SETUP.
Delay changes to match the other channels at those stations, for SUBSEQUENT SETUPS.

8 Delay is discrepant at one bandwidth/station in the EIGHTH SETUP in the entire observation; the SECOND occurrence of this setup.
Delay then RETURNS to the value seen in the first occurrence of this setup.

18/8/17 -jdr

Impact on and Advice to Users

First, the delay jumps seen in the recent tests are in a range similar to the typical timing differences of ~100 ns among VLBA stations, so they should have no significant effect on fringe detectability, except possibly at the very narrowest 1- or 2-MHz bandwidths.

This limits the scientific impact of the jumps primarily to observations which depend on carrying over delay/phase solutions from calibrators to target sources *observed at different bandwidths*. Many experts consider this to be bad practice in general, but sensitivity considerations may make it unavoidable in some circumstances. Interpolation of delays among occasional atmospheric calibration blocks may fail between an initial and subsequent blocks, but the atmospheric results obtained in each block should still be valid.

The restriction of almost all the jumps to the transition out of the first- or second-used setups makes an appropriate single initial dummy scan a viable method to eliminate these jumps in future observations. This scan should use a bandwidth different from any other within the intended observation, and is only necessary when the intended first bandwidth observed is 16 MHz or narrower (presumably extending down to 1 MHz), or when an initial scan at 128-MHz bandwidth is to be followed by such a narrowband scan. Such cases should be easily avoidable in scheduling.

The preceding statements apply to an individual observation, as specified in a single schedule file. However, some observations may be interrupted by one-hour daily observations at two (or occasionally more) stations for USNO's Earth Orientation Parameter measurements. VLBA operations software treats resumption after such an interruption as an entirely new observation, with the scans preceding the current time simply being skipped. The references herein to the "first or second scans" will re-apply at that time at the affected stations, so that careful checks for delay jumps are appropriate in data taken immediately after such an interruption.

As a simpler alternative to all the above, we recommend that delay-sensitive DDC observing be limited to bandwidths of 32, 64, and 128 MHz. This range includes the most commonly used DDC configurations, and we believe the additional load on the VLBA media pool will not be significant. Narrower bandwidth observations could also be accommodated by observing at 32 MHz, and requesting "zoom mode" correlation to achieve the desired bandwidth. However, this may not be satisfactory if the narrow bands serve to avoid strong interference.

The preceding relatively good news stands in contrast to the situation for the one remaining delay jump that was seen. Such a jump — which occurred in 0.04% of bandwidth changes in the tests reported here — could not be avoided by either of the approaches suggested above. In observations of relatively strong sources, they could be detected by examination of the delays measured in fringe fits.

Future Correction via Phase Cal

Detection and automatic correction of delay jumps using phase-cal data has been developed in AIPS, and is described in AIPS Memo #123 by Eric Greisen. This could become a major, wide-scale approach to elimination of concern about delay jumps in continuum observations. However, the required, substantially expanded PC table is not yet included within the FITS-IDI files currently available from the VLBA archive. Although most of the phase cal measurements in these files are of excellent quality, there are often a number of solutions that failed for reasons currently unknown. AIPS includes a variety of tasks, also described in Memo #123, for visualizing and editing the files. Availability of the expanded PC table is expected by about the beginning of calendar year 2019.