

UVLBI Memo #032

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY**

WESTFORD, MASSACHUSETTS 01886

5 September 2012

Telephone: 781-981-5400*Fax:* 781-981-0590

To: UVLBI and Geodesy Groups
From: G. B. Crew
Subject: RDBE-S 16 Gbps Demo Results

Overview

This memo presents results of the recent 16 Gbps demonstration campaign. This effort began with efforts to build a burst-mode recorder system and a wide-channel digital back end; and with the introduction of the Mark6 recorder system has reached a recording milestone of 4 GHz of recorded data using the broadband GGAO/Westford stations. Many of the developmental details, tests and results are described elsewhere (see, for example memos 25, 26 and 27 in the UVLBI memo series); so we will only summarize some of that here to provide context for the observations and results described here.

This work was initiated in the summer of 2010 with an intent to quickly develop digital backend firmware and recording capability to be deployed in several months for an 8 Gbps test. In the event, the basic capability (i.e. a version of the firmware, and a burst mode recorder) was created and zero-baseline tested in the laboratory in early 2011. Field deployment further required additional RDBE units, a usable broadband system, and experiment time at two radio telescopes; and neither was readily available. In the interim, the burst mode recorder system was abandoned in favor of a continuous recording capability; and current hardware is sufficiently capable to support a doubling of the bit rate to 16 Gbps (necessitating additional digital backends to feed the beast). Prototype Mark6 hardware and software were developed later in that year and was deployed in October of 2011 as a tag-along on an R1 geodesy observation. This provided a demonstration of the capability with 8-fold recording of one X-band IF on several sources. Finally, in May of 2012 we were able to use the GGAO and Westford broadband systems to complete observations on an aggregate 4 GHz of bandwidth, and the campaign was concluded in June of 2012.

We begin with a description of the equipment, see [\[system\]](#), [page 1](#) and continue with the observations, see [\[observations\]](#), [page 9](#). There are a few issues with correlation and fringe fitting, see [\[correlation\]](#), [page 10](#). before continuing with a presentation of the results, see [\[results\]](#), [page 12](#) Finally, we acknowledge the efforts of a great many people which made this possible in a final section.

udc	f_Luff	f_L0	Ny2 band
UDC-A	7773	8592	9104-9616
UDC-B	7645	8080	8592-9104
UDC-C	7517	7568	8080-8592
UDC-D	7389	7056	7568-8080

where all frequencies are in MHz. Operationally, the UDCs are programmed through an RS-232 connection from a Mark5 (not actually otherwise used in these experiments). Typically, the attenuators are set once to provide appropriate signal strengths.

The digital backends are ROACH-based DBEs which are derived from the VLBI2010 RDBE, and given a separate designator, RDBE-S, since they differ in four significant ways:

- There is no automatic level control (ALC). This is not an issue since the UDCs have adjustable output levels.
- There are two ADC cards capable of sampling 4 IF streams per RDBE-S
- There is no internal filtering, the entire 512 MHz IF is sampled and packetized as a single channel.
- The output is produced in VLBI Data Interchange Format (VDIF), rather than Mark5B (emulation) packets.

Again, we will not dwell on the details here, as they have been described in previous memos. However, we will note that there is significant cross-talk between the two channels within each iADC cards, so the signals from each polarization pair were routed to separate cards so that polarization measurements could be made. The signals of each iADC card are routed to separate CX4 10 GbE connections (the dotted lines in [Figure 1](#)) so that the digital signals of each polarization pair ends up on a single CX4 cable. In VDIF parlance, each IF was labelled as a separate "thread" (0 through 7 from the two RDBE-S units); and each thread on the UDP packet output on these four cables was given a unique address:port assignment (ports 4201 through 4208):

```
IP 192.168.5.101.65535 > 192.168.5.102.4201: UDP, length 8224
IP 192.168.5.101.65535 > 192.168.5.102.4203: UDP, length 8224
IP 192.168.5.109.65535 > 192.168.5.110.4202: UDP, length 8224
IP 192.168.5.109.65535 > 192.168.5.110.4204: UDP, length 8224
IP 192.168.5.101.65535 > 192.168.5.118.4205: UDP, length 8224
IP 192.168.5.101.65535 > 192.168.5.118.4207: UDP, length 8224
IP 192.168.5.109.65535 > 192.168.5.126.4206: UDP, length 8224
IP 192.168.5.109.65535 > 192.168.5.126.4208: UDP, length 8224
```

The RDBE-S firmware uses the same quantization scheme as its VLBI2010 parent, viz, a histogram of the raw 8-bit data should be approximately Gaussian, and a digital threshold is determined to partition it into 4 (2-bit) states such that the middle pair of states (surrounding the zero point) should comprise about 68% of the output. Because of the lack of the polyphase filtering which ensures that there is no DC signal, a slight modification was made to remove any residual DC bias that might have been introduced by the ADC. The RDBE-S has a "raw" capture mode which provides this histogram as a diagnostic. These are shown in [Figure 2](#) and [Figure 3](#). These examples are from the June 19th observations.




Figure 2: Raw ADC data from the two RDBE-S units (RDBE-S001 and RDBE-S002) at Westford for its 4 IFs.




Figure 3: Raw ADC data from the two RDBE-S units (RDBE-S003 and RDBE-S004) at GGAO for its 4 IFs.

The Python graphic which produced the plots has a nasty artifact in that (depending on the histogram data) it either does or doesn't introduce return-to-zero lines from the actual data points, which is why some of the histograms are filled in, and some are not.

Once initialized (network/thread configuration) the RDBE-S can be commanded to quantize the data stream at any time. Typically this is done on a per-scan basis. An external script is at present used to select the proper thresholds between the 4 states (--, -, + and ++); but this could be added to the firmware. The script provides output such as follows (at the start of the Jun 19th observations):

```

rdbe-S001 KkLlMmNn-
-----IF0-%----- | -----IF1-%----- | -----IF2-%----- | -----IF3-%----- |
0-- 0- 0+ 0++ | 1-- 1- 1+ 1++ | 2-- 2- 2+ 2++ | 3-- 3- 3+ 3++ |
 16 32 36 14 | 15 35 34 14 | 16 32 34 16 | 16 33 33 16 |
-----
IF0 mean 0.24 std 21.22
IF1 mean -0.38 std 20.23
IF2 mean -0.10 std 20.44
IF3 mean -0.51 std 21.24

rdbe-S002 OoPpQqRr+
-----IF0-%----- | -----IF1-%----- | -----IF2-%----- | -----IF3-%----- |
0-- 0- 0+ 0++ | 1-- 1- 1+ 1++ | 2-- 2- 2+ 2++ | 3-- 3- 3+ 3++ |
 16 33 34 16 | 15 34 33 16 | 15 34 34 15 | 16 33 34 15 |
-----
IF0 mean -0.84 std 21.36
IF1 mean -0.83 std 20.52
IF2 mean -0.73 std 21.15
IF3 mean -0.53 std 21.07

rdbe-S003 SsTtUuVv+
-----IF0-%----- | -----IF1-%----- | -----IF2-%----- | -----IF3-%----- |
0-- 0- 0+ 0++ | 1-- 1- 1+ 1++ | 2-- 2- 2+ 2++ | 3-- 3- 3+ 3++ |
 16 35 31 16 | 16 36 31 15 | 15 34 34 15 | 15 34 35 14 |
-----
IF0 mean -1.15 std 21.26
IF1 mean -1.66 std 19.48
IF2 mean -0.31 std 19.05
IF3 mean -0.65 std 20.28

rdbe-S004 WwXxYyZz-
-----IF0-%----- | -----IF1-%----- | -----IF2-%----- | -----IF3-%----- |
0-- 0- 0+ 0++ | 1-- 1- 1+ 1++ | 2-- 2- 2+ 2++ | 3-- 3- 3+ 3++ |
 16 33 32 17 | 16 33 33 16 | 16 33 33 15 | 15 34 34 16 |
-----
IF0 mean -0.93 std 18.98
IF1 mean -0.28 std 19.43
IF2 mean -0.31 std 19.15
IF3 mean -0.77 std 20.43

```

For each IF on each RDBE-S, a 2-letter station designator is used (Kk through Rr at Westford, and Ss through Zz at GGAO) as an identifier which was useful during development. The station labels for `fourfit` are derived from these by dropping the second letter. These letter assignments are also shown in [Figure 1](#) immediately above or below the RDBE-S boxes.

The 4 numbers within the boxes of the script output are the relative population of the 4 states. (I.e. a third of the signal in each of the two low states, and the remaining third split between the two high states.) The actual mean and standard deviation of the raw

Gaussian is also shown. Typically these values fluctuate as they are drawn from a small sample (32768 samples).

The UDC attenuators were tuned to provide a standard deviation of approximately 20 units for these Gaussians. Ideally, 32 units would provide the maximum sensitivity, but this was not attainable on all channels, so a uniform choice of 20 was made instead. So, with Nyquist frequency, 2-bit sampling (1024 MHz) on eight 512 MHz IFs, we have 2 Gbps for each IF, 4 Gbps per CX4 cable, for a total of 16 Gbps into the two recorders. Note that in principal 2 of the 4 Gbps packet streams could be combined onto a signal 10 GbE cable. This was not done as it would have required additional development of the RDBE-S firmware. (I.e. it was simpler to use an additional CX4 output for the extra iADC in the RDBE-S than to add the additional logic to buffer and combine the two packet streams.)

In these experiments, prototype Mark6 recorders (shown in [Figure 4](#)) were used. In this configuration, each Mark6 had two dual-ported network interface cards.

Prototype Mark 6



Figure 4: The Mark6 prototype hardware used at Westford. The host computer and 2 disk packs are in the upper 4U module; the 2 additional disk packs are supported by the lower 4U module. A 1U separator provides a cable paths to the host RAID cards.

These cards are designed to support an aggregate 10 Gbps packet flow, and present two software ethernet devices to the Linux kernel (i.e. eth2 and eth3 from one card, and eth4 and eth5 from the other card). Symmetrically, a pair of hardware RAID cards were used. Each of these has "slots" to support 16 disks RAID0-ed in groups of 8 to form four multi TB ext4 filesystems (i.e. disk0, disk1, disk2 and disk3). The disks were packaged in

these same groups of 8 and connected to the raid cards via external SAS cables which group the signals of four disks into one cable. (I.e. a total of 8 cables per Mark6 recorder.)

The point of this arrangement is that the RAID0 of 8 disks provides approximately 8 times the write speed of a single disk, and that the RAIDing is performed in the hardware means that the kernel is relieved of the required RAID computations. During every scan, system performance was monitored; after every scan diagnostic plots were produced. Samples are shown in Figure 5 and Figure 6 which show the network packet performance and disk write performance, respectively. At 16 Gbps, about 2 GB of data enters the system every second. The machines at Westford and GGAO had 12 and 6 GB of system memory, respectively, so during the first few seconds data can be captured and buffered to memory. After that, a steady-state flow of data to disk is required. The first plot shows a more steady rate, as the packet buffers are smaller and thus must be emptied more regularly; the file buffers are larger and therefore the instantaneous write rate fluctuates more. In both cases, the average rates are in excess of 4 Gbps per network/disk device.

The actual rate is in excess of 4 Gbps, because while each packet holds 8192 B of data, there are in addition VDIF and network header data also being captured. Thus the aggregate rate per thread-pair is greater than 4 Gbps. In this developmental setup, the entire network packets were captured and conversion to VDIF-packet-only files was carried out as a secondary step. In the Mark6 product that will eventually be deployed, the files will contain only VDIF packets.




Figure 5: A plot of the instantaneous network packet processing rate for a 2-minute recording.




Figure 6: A plot of the instantaneous file write rate for a 2-minute recording.

Observations

Observations with the complete system were carried out on 3 separate days: May 17 and May 18 (days 138 and 139) of 2012, and June 19 (day 171) of 2012. The observational plan was to visit several sources of varying strengths for scans long enough to achieve significant SNR. One of the brighter sources (3C84 or 0316+413) was selected as it was used on previous days for other broadband observations. A weaker, but still bright source (0552+398) was selected because it had been used in the X-band observations of the previous year. Additional sources were selected to be nearby and so that the full set spans 2 orders of magnitude in intensity. A final consideration was to have relatively small parallactic angles at the time of observation (to simplify the eventual polarization analysis).

These primary targets were:

Source	R.A. (hours)	Decl. (deg)	400km	El (Wf)	PA	PA dif
0316+308	+3.323336456	+31.029012897	X 0.09	58.8	-57.9	7.2
0550+356	+5.902647028	+35.692055733	X 0.18	33.1	-58.2	2.2
0554+580	+5.987053952	+58.067624161	X 0.29	41.7	-73.4	-0.5
0059+581	+1.046045106	+58.403093503	X 1.75	74.1	-173.3	-10.5
0552+398	+5.925223781	+39.813656939	C 6.96	35.0	-60.5	1.8
0316+413	+3.330044469	+41.511695578	C 25.41	62.9	-75.2	5.2

The 4th and 5th columns above indicate the source strength in Jy for a 400 km baseline from the VLBA calibrator list. The elevation at Wf is given in the 6th column, the parallactic angle in the 7th, and the difference in the 8th. The PA calculations are for 1400 UT, and 0316+413 (last row) is more commonly known as 3C84. For future reference, the backup targets were:

Source	R.A. (hours)	Decl. (deg)	400km	El (Wf)	PA	PAdif
0058+498	+1.021388205	+50.079163753	X 0.12	82.4	-168.1	-15.3
0057+334	+1.010636310	+33.751700467	X 0.14	81.0	-9.3	35.5
0318+438	+3.360241213	+43.989578272	X 0.17	63.1	-79.8	4.4
0554+343	+5.966691281	+34.313441028	X 0.20	31.7	-57.2	2.2
0057+678	+1.014350998	+68.139037531	X 0.38	64.4	-176.5	-7.7
0554+242	+5.951309328	+24.232027406	X 0.74	26.1	-52.7	3.0

In all cases, time with the system was for a variety of reasons limited, so it was not possible to obtain a full sequence of scans on a variety of sources with the complete system as originally planned. We did manage to obtain a variety of results that collectively serve the basic demonstration purpose of the campaign, so we here accept the deficiencies as part of life and move on to the analysis.

On the first two days of observation, May 17 and May 18 (days 138 and 139), there were issues with one of the up/down converters at GGAO, with the result that one band (both polarizations) was not operated with a sensitivity comparable to its peers. The problem was ultimately traced to bad (lossy) cables. A correlation on the affected band was (surprisingly) possible, but interpretation of the results was problematic, so for those two days, we will omit that band and work with the remaining 3 GHz. On the first day, the requantization with each scan was not part of the initial plan; and unfortunately the RDBE-S units were initialized with no signal present, so while the data was recorded as 2-bit samples, the digital backend was effectively making 1-bit samples. The impact of this is calculable (Van Vleck factors of 0.68 and 0.88 respectively), so we do include results from both days.

On the final day, June 19 (day 171), the observing plan was cut short by an emergency need to put Westford back into S-X geodesy service; so complete, no-issue observations were only carried out on the first source of the day, 3C84.

Amusingly, name confusion (0316+413 v 0316+308) caused 3C84 (0316+413) to not be observed at one station on May 18. (Note to self: choose the names of pointing procedures with care to avoid ambiguity.) A further operational consideration was that it was easy to bring scans captured at Westford to Haystack for correlation, but we did not have the bandwidth to pull all of the valid GGAO scans up to Haystack. Unfortunately, this meant that we were not able to recover all the data at Haystack before it was necessary to rebuild some of the RAID systems due corruption resulting from power and cabling issues at each of the sites. So, to summarize: no one source was observed on all 3 days, and not all of the scans captured were fully analyzed. Nevertheless in [\[results\]](#), [page 12](#) we show promising results from what we have.

In the following table, we summarize the scans that were successfully recorded with durations in seconds.

Source	May17	May18	Jun19
3C84	10,10,30	none	5,5,10
0552+398	10,30,30,60	10,30,30,30,30,30,30	none
0059+581	10,30,30,60	10,10,20,20,40	none
0554+580	30,30,60	10,30,40	none
0316+308	60,60,60	60,60,60,60	none
0550+356	60	10,30,30,60,60,60	none

Correlation and Fourfitting

The data was correlated with DiFX using files of the VDIF packets stored on RAIDS available to four of the correlation nodes. One of the challenges of this process is that DiFX is meant to be driven by a VEX (VLBI EXperiment) file, but the current version of VEX (1.5) is inadequate to fully describe the experiment. Furthermore, the separation of the different threads into separate files is not currently handled by DiFX. Note that the different bands are independent, so there is no pressing reason to correlate them simulataneously. (That's a plus, since it's not currently practical to do this anyway.) A side effect of this is that the Mark4 format data files need to be combined as a post-processing step in order to present the complete observation to `fourfit`. And finally, the program that does this merger (`fourmer`) cannot relate the different "stations" that are needed to set up the intermediate processing steps in order to allow proper analysis of the polarization data.

A final problem is that while both stations have phase cal data (and both systems appeared to be working); the phase cal data reported by DiFX does not appear to be useful within `fourfit`. Brief investigation of the problem does not make it clear whether the problem is with `fourfit` or DiFX. In fairness to DiFX, this processing procedure is unusual insofar as each "station" receives only one polarization (i.e. either H or V) but both are active in the experment. Also, the bands are rather wider than normal (512 MHz compared with the more usual 32 MHz) so phase cal tones at 5 MHz produce a bit more information than the system is currently equipped to handle. These problems can all be investigated and ultimately solved; but here we have taken the prudent shortcut of using manual phases (on the brightest sources, and for the weaker sources, those from a "recent" bright one).

So: the procedure carried out was as follows:

1. a script `gang_of_four.sh` was invoked for each scan to carry out 4 correlations, one per IF band:
2. locate the data files and gather the information necessary to set up the correlation (using a script `driv3_m6v.sh`),
3. construct `vex` and `v2c` files suitable for invoking `vex2difx` which prepares the correlation and write a script to run the correlation,
4. construct post processing scripts (to run `difx2mark4` and `fourfit` to immediately characterize the results,
5. actually run the correlation and post-processing scripts (these steps accomplished with a script `comb3_m6v.sh`),
6. once all 4 bands have been correlated run a script (`gang_fourmer.sh`) that assembles the 4 correlation results into the equivalent of the all-four-at-once correlation,
7. construct manual phase calcs suitable for use with `fourfit`,
8. and finally `fourfit` the final data.

The final set of labels from [Figure 1](#) shows the stations names and `fourfit` labels used through these steps. Some of the tools are not yet happy with the use of X/Y or H/V for polarization labels, so we used X/H/L or Y/V/R interchangeable. The "stations" Aw and Bw refer to two halves of the Westford system; `fourfit` refers to these stations as A and B. The "stations" Cg and Dg refer to two halves of the GGAO system; `fourfit` refers to these stations as C and D. Thus thread 0 (port 4201) at Westford carries Aw, Band A, L

(which is really H) polarization (upper right-hand corner of the figure). The complete H polarization at Westford is carried by the 4 threads(ports) 0,1,4,5(4201,4202,4205,4206) for bands A,B,C and D.

The correlated data and some fourfit results are archived for future reference in the Haystack correlation library as experiment 3407.

Results

In the following table, we list all the targets and scans successfully recorded and recovered at Haystack. For each we present 3 representative amplitude correlation values. The columns marked **b:c** are the effective number of bits and the number of channels. The column marked **dur** is the duration in seconds of the multi-channel fits. The operational software used to capture the scans was not overly concerned with starting/stopping/capturing data on the requested boundaries. As a result the correlated scan duration are somewhat less than the requested scan durations resulted, since the correlation is carried out on the overlap of the data streams. The parenthesized number is the actual correlated duration for the multiband scans. The correlation intervals were requested to be 0.4s (and therefore the deficit is some multiple of this amount).

The single channel value is for the lowest frequency band (**fourfit** channel **a** H pol (DC edge at sky frequency 8080 MHz, UDC-D, typically the strongest signal of the four bands), and the multi-band values are for 3 channels on the May scans, and 4 bands on the June scans for the two aligned polarizations (HH and VV). The SNRs (as calculated by **fourfit**) are in parentheses.

Manual phase cals **fourfit** control files were constructed for scans 138-1438 (on 0059+581, May 17th), 139-1256 (on 0552+398, May 18th), and for scans 171-1917, 171-1919, and 171-1927 (on 3C84, used on each scan, June 19th). In the first session, we have the following results:

Source	Scan	b:c	dur(act)	HH-a(SNR)	HH amp(SNR)	VV amp(SNR)
3C84	138-1357	1:3	10(9.60)	44.232(368.0)	39.520(573.1)	37.661(547.5)
3C84	138-1400	1:3	30(28.80)	43.814(641.8)	40.784(1035.3)	38.973(989.0)
3C84	138-1407	1:3	10(9.60)	43.310(361.1)	39.953(577.6)	38.411(556.6)
0552+398	138-1411	1:3	10(9.60)	8.992(84.6)	8.334(112.3)	8.013(117.6)
0552+398	138-1414	1:3	30(28.80)	9.853(142.0)	7.712(193.6)	7.454(187.8)
0552+398	138-1420	1:3	30(29.10)	10.017(147.9)	8.344(211.5)	8.041(203.8)
0059+581	138-1438	1:3	10(9.60)	8.638(73.2)	7.700(113.0)	7.297(107.1)
0059+581	138-1442	1:3	30(28.80)	8.927(131.0)	7.705(195.8)	7.279(185.0)
0059+581	138-1448	1:3	30(28.80)	8.929(130.1)	7.936(200.5)	7.427(187.6)
0554+580	138-1506	1:3	30(28.80)	0.836(12.1)	0.657(16.7)	0.618(15.7)
0554+580	138-1511	1:3	30(28.80)	0.731(10.6)	0.677(17.1)	0.641(16.0)
0554+580	138-1517	1:3	60(59.04)	0.833(17.4)	0.662(23.9)	0.593(21.5)
0316+308	138-1525	1:3	60(59.04)	0.215(4.5)	0.149(5.4)	0.165(6.0)
0316+308	138-1534	1:3	60(58.72)	0.199(4.2)	0.166(6.0)	0.205(7.4)
0316+308	138-1537	1:3	60(59.04)	0.172(6.2)	0.172(6.2)	0.165(6.0)
0550+356	138-1541	1:3	60(59.04)	0.370(7.7)	0.336(12.0)	0.267(9.5)

Since this is 1-bit data, there is some loss of sensitivity relative to the normal 2-bit data. However, this is the only data set (at present) with multiple scans on multiple sources, so it is adequate for making some general observations about the data. For example, the HH amplitude is in all cases stronger than the VV amplitude. Likewise, the single-channel HH result has a slightly higher correlation amplitude than the corresponding 3-channel HH

result. Among scans on the same source, the longer ones have comparable amplitudes, but higher SNRs that scale roughly with time as expected. Unfortunately, from the May 18th session with 2-bit data, only 4 scans on 4 sources were eventually recovered to Haystack and available for correlation.

Source	Scan	b:c	dur(act)	HH-a(SNR)	HH amp(SNR)	VV amp(SNR)
0552+398	139-1256	2:3	10(10.08)	11.732(101.9)	9.576(144.0)	9.238(138.9)
0554+580	139-1353	2:3	30(28.80)	1.007(14.5)	0.875(21.9)	0.807(20.3)
0550+356	139-1422	2:3	60(59.04)	0.507(10.5)	0.397(14.4)	0.393(14.3)
0316+308	139-1439	2:3	60(59.04)	0.223(4.6)	0.194(7.0)	0.231(8.4)

Finally, the session on June 19th was prematurely terminated after only a few scans on 3C84. (The Westford station was needed to replace a failed station in normal S-X operations.) These were, however, full 4 GHz (2 polarizations, 4 channel, 512 MHz/channel) scans:

Source	Scan	b:c	dur(act)	HH-a(SNR)	HH amp(SNR)	VV amp(SNR)
3C84	171-1917	2:4	5(4.32)	60.571(342.9)	55.739(633.4)	51.596(586.3)
3C84	171-1919	2:4	5(4.32)	60.624(343.3)	55.748(633.5)	51.393(584.1)
3C84	171-1927	2:4	10(9.72)	60.158(520.1)	55.201(940.9)	50.673(863.7)

For comparison, the amplitude and SNR results from the 1-bit May 17th can be trivially scaled by the 2-bit/1-bit factor, viz, $0.88/0.64 = 1.375$, and by the ratio of the actual duration to that obtained below. Doing so, we obtain the following table:

Source	Scan	b:c	dur(act)	HH-a(SNR)	HH amp(SNR)	VV amp(SNR)
3C84	138-1357	*2:3	10(4.32)	60.819(339.4)	54.340(528.6)	51.784(505.0)
3C84	138-1400	*2:3	30(4.32)	60.244(341.8)	56.078(551.3)	53.588(526.7)
3C84	138-1407	*2:3	10(4.32)	59.551(333.1)	54.935(532.8)	52.815(513.4)
0552+398	138-1411	*2:3	10(10.08)	12.364(119.2)	11.459(158.2)	11.018(165.7)
0552+398	138-1414	*2:3	30(10.08)	13.548(115.5)	10.604(157.5)	10.249(152.8)
0552+398	138-1420	*2:3	30(10.08)	13.773(119.7)	11.473(171.2)	11.056(164.9)
0059+581	138-1438	*2:3	10(28.80)	11.877(174.3)	10.588(269.1)	10.033(255.1)
0059+581	138-1442	*2:3	30(28.80)	12.275(180.1)	10.594(269.2)	10.009(254.4)
0059+581	138-1448	*2:3	30(28.80)	12.277(178.9)	10.912(275.7)	10.212(257.9)
0554+580	138-1506	*2:3	30(28.80)	1.149(16.6)	0.903(23.0)	0.850(21.6)
0554+580	138-1511	*2:3	30(28.80)	1.005(14.6)	0.931(23.5)	0.881(22.0)
0554+580	138-1517	*2:3	60(28.80)	1.145(16.7)	0.910(23.0)	0.815(20.6)
0316+308	138-1525	*2:3	60(59.04)	0.296(6.2)	0.205(7.4)	0.227(8.2)
0316+308	138-1534	*2:3	60(59.04)	0.274(5.8)	0.228(8.3)	0.282(10.2)
0316+308	138-1537	*2:3	60(59.04)	0.236(8.5)	0.236(8.5)	0.227(8.2)
0550+356	138-1541	*2:3	60(59.04)	0.509(10.6)	0.462(16.5)	0.367(13.1)

For the 3C84 scans, these scaled SNRs are not as large since there are only 3 (rather than 4 bands) contributing to the result (by a factor of $\sqrt{3/4} = 0.866$). The amplitudes are also less consistent for 3C84 as well as all the other sources as expected (lower SNR 1-bit data compared with higher SNR 2-bit data). The scaled 1-bit data has slightly higher amplitudes than the unscaled 2-bit data (at the 10% level, perhaps). It is not immediately clear why this is the case. `Fourfit` fringe plots for the 171-1927 scan are shown in Figures [Figure 7](#) and [Figure 8](#). For comparison, HH fringe plots from each of the May 18th targets are also shown.




Figure 7: This is the fourfit fringe plot for the HH polarization of one of the the 4 GHz scans (171-1927) on 3C84. It is a 25.4 Jy source (C Band).




Figure 8: This is the fourfit fringe plot for the VV polarization of one of the the 4 GHz scans (171-1927) on 3C84. It is a 25.4 Jy source (C Band).




Figure 9: This is the fourfit fringe plot for the HH polarization of the May 18th observation of 0552+398. It is a 6.96 Jy source (C Band).




Figure 10: This is the fourfit fringe plot for the HH polarization of the May 18th observation of 0554+580. It is a 0.29 Jy source (X Band).




Figure 11: This is the fourfit fringe plot for the HH polarization of the May 18th observation of 0550+356. It is a 0.18 Jy source (X Band).




Figure 12: This is the fourfit fringe plot for the HH polarization of the May 18th observation of 0316+308. It is a 0.09 Jy source (X Band).

The most obvious difference between the two fringe plots is the shape of the cross-power spectrum, which appears to be more rounded for the VV (labelled RR by `fourfit`) polarization than that of the HH (labelled LL by `fourfit`). This is most likely due to front-end circuit characteristics (i.e. filters and fiber transmission), rather than back-end electronics (i.e. properties of the ADCs). Across the 4 bands, the correlation amplitude appears to be slightly steeper in VV (than HH). The VV and HH amplitudes `a` through `d` were: 58.391, 54.446, 53.161, 46.373 for VV compared with 60.158, 56.111, 55.428, 49.687 for HH. Again, this is most probably an electronic artifact.

Finally, the parallactic angles were 60.3 (Westford) and 66.3 (GGAO) for 171-1927; $\sin(\text{PA})$ is about 10%, and that is the relative strength of the amplitude of the cross-hand correlations:

3C84 Scan	HH amp(SNR)	HV amp(SNR)	VH amp(SNR)	VV amp(SNR)
171-1917	55.74(633.4)	5.09(57.2)	6.51(73.5)	51.60(586.3)
171-1919	55.75(633.5)	4.81(54.3)	6.39(72.0)	51.39(584.0)
171-1927	55.20(940.9)	4.80(80.9)	6.40(108.3)	50.67(863.7)

More detailed analysis should await completion and characterization of the broadband electronics. In any case, we do not have sufficient data here to untangle the gain and D terms (and we did not intend to be so able).

Finally, contemporaneous SEFD measurements were planned for the June 19th session, but could not be carried out before the broadband receiver was removed to be replaced with the normal S-X receiver. Based on the catalog strengths of these sources, we can say that the correlation amplitudes are in the proper relationships and approximately as expected, but that a detailed comparison must be a future project.

Summary

In this memo we have present the observational details and some of the immediate correlation results from 3 VLBI sessions with the broadband electronics plus single-channel-flavor digital backend plus prototype Mark6 recording system. The main purpose here was to demonstrate that

- A total bandwidth of 4 GHz is possible with current technology
- That such a system can be deployed

Acknowledgements

This work would not have been possible without a great many contributions by a large number of individuals. Attempting to mention everyone:

- Shep Doeleman who motivated and supported the RDBE-S hardware and burst mode recorder system development,
- Alan Whitney who coordinated the Mark6 effort,
- Arthur Niell who provided moral and planning support for the observations, as well as negotiating with NASA for use of the antennas,
- Alan Hinton who did the initial development work on the RDBE-S firmware,
- Russ Mewhirter who completed the firmware implementation,

- Chester Ruzczyk who provided hardware and software support throughout the development, testing and use of the system,
- Jason SooHoo who provided essential network, hardware and logistical support throughout the project,
- Chris Beaudoin who provided significant technical support on the Westford and GGAO broadband systems,
- David Lapsley who wrote the Mark6 application used in these tests,
- Roger Cappallo, Brian Corey and Alan Rogers who have offered guidance on a variety of technical issues,
- Mike Titus who provided correlation assistance,
- Mike Poirier, Dave Fields and the rest of the Westford team for operational support at the Westford station,
- Wendy Avelar, Ed Himwich and the rest of the GGAO team for operational support at GGAO,
- Adam Deller, Walter Brisken and the other DiFX developers for the software to correlate the observations,

This work made use of the Swinburne University of Technology software correlator, developed as part of the Australian Major National Research Facilities Programme and operated under licence. (See Deller, et al. 2011, PASP, 123, 275 for a description of DiFX.) This work was sponsored under NASA ... and NSF