

UVLBI Memo #027

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
HAYSTACK OBSERVATORY**

WESTFORD, MASSACHUSETTS 01886

26 October 2011

*Telephone:* 781-981-5400

*Fax:* 781-981-0590

To: UVLBI Group

From: G. B. Crew, C. Beaudoin, D. Lapsley, and J. SooHoo

Subject: RDBE-S Testing at Westford/GGAO with Triple Noise, X-band and UDCs

## Overview

This memo briefly describes testing of the four RDBE-S units at Westford prior to shipping two units to GGAO for the fringe test, and briefly, the immediate results of that fringe test. After shipment, the Westford units were connected to the S/X receiver system to allow a GGAO-Westford fringe test in the X-band.

The move to Westford allows us to use the Up/Down converters and thereby sample a full 1 GHz band from the noise sources in two contiguous 512 MHz wide-band channels as a simulation of two VLBI stations. As a minor benefit, the Westford Maser was used for the 5 MHz and 1 PPS signals. Recabling from this test configuration to that needed for the eventual 2-GHz VLBI experiment (with the broadband receiver at Westford) is relatively straightforward. For background information and prior testing of the RDBE-S digital back end, see Memos #025 and #026 in this UVLBI series. The RDBE-Ss units were tested using the `RoachAstro8Gbs_sept14g.bin` personality, reflecting the current state of that development effort.

The main goals of the testing described in this memo are:

1. Establish that the four RDBE-S units operate equivalently
2. Verify that the DiFX/HOPS software support analysis of the two wide bands.
3. Obtain an X-band fringe GGAO-Westford while recording 16 Gbps

As described in the following sections, all units performed acceptably and two of them are now at GGAO with the Goddard Mark6 unit.

Approximately 40 noise recordings were made, referred to in this memo by run or scan number. The noise data was typically recorded to disk and then the correlation results were kept in directories within `smm-pc:/data-sc01/gbc/wftest`. Some of the more complicated analysis runs required multiple data directories. Detailed notes (`test-<date>.txt`) on these activities are to be found within the burst mode svn repository as `svn+ssh://vault/svnrepos/burst/trunk/doc/wftest`. The noise runs were made as follows:

- 01-16      Initial noise-box testing with burst mode machines and mark6 recorder to establish consistent behavior with the different software/hardware recording methods

- (notes in `test-sept27.txt`). Tests were conducted in the correlator room using RDBE-S003 and RDBE-S004.
- 17-20 Noise-box, Mark6 and RDBE-S units were moved to Westford and the preceding tests repeated with RDBE-S003 and RDBE-S004 (notes in `test-sept29.txt`).
- 21-28 The 3-noise sources were configured to use the Orca box and up/down converters. Tests with RDBE-S003 and RDBE-S004 were repeated. Initial attempts to join the two frequency bands were made (notes in `test-oct03.txt`).
- 29-33 Testing switched to RDBE-S001 and RDBE-S002, script to join bands finalized (notes in `test-oct07.txt`).
- 34-36 All four RDBE-S units recorded in parallel (but with only 8 IFs connected) (notes in `test-oct07.txt`).
- 37-38 All four RDBE-S units recorded in parallel (with only 8 IFs connected), but RDBE-S003 and RDBE-S004 were operated with the PFB personality (notes in `test-oct17.txt`).
- fringes The fringe data files copied to `sc02:/data-sc02/??data` named by station and analyzed in `sc01:/data-sc01/gbc/wfggx`.

The X-band fringe test was carried out Oct 24/25 as a tag along with the normal R1504 geodesy session.

## Noise-Box Testing at Haystack and Westford

The final incarnation of the three-noise source described in Memo #026 was used to test all four RDBE-S units at Haystack, and again at Westford. A diagram is shown in [Figure 1](#). At Westford, the house maser and 1PPS signals were used rather than the 5 MHz oscillator that had been used previously. No significant difference in performance was noted on this zero baseline test. The other difference is that the Mark6 recorder was used to capture the data packets in a single 16 Gbps recorder rather than 2 parallel 8 Gbps recorders.

In the course of these tests, it was discovered that the channel nearest the DC edge is typically discarded by `fourfit`. This behavior of `fourfit` is responsible for some of the normalization issues previously identified (*e.g.* that autocorrelations were not normalized to 100%). Discarding the DC edge was deemed appropriate to compensate for DC biases of the analog filters; however with the digital filters now in use in current backends, this is not necessarily desirable. A new parameter (`dc_block`) was introduced into `fourfit` to allow this channel to be used.

A set of runs (#02 for 2 secs, #12 and #16 for 4 seconds) were made using the burst mode software on the burst mode machines, the burst mode software on the Mark6 machine, and the Mark6 software on the Mark6 machine. Presenting these results as an 8-station matrix (stations S,T,U,V on S003 and W,X,Y,Z on S004), we find the results shown in [Table 1](#), where 12 dB of attenuation between splitters was applied. The greater disparity of the amplitudes between run #02 and runs #12 and #16 are due to different integration times and noise issues (still to be explored) at the DC edge of the wide band.

The single-band delays and phases are very consistent between the three runs as shown in [Table 2](#) and [Table 3](#). Again, there is some disparity on the baselines reflecting the

(weak) common noise source for the shorter integration time. At the longer integration times (4-sec) there is very good agreement between the burst mode and mark6 software.

## Up/Down Converter Test Configuration

In order to capture and record 1 GHz of bandwidth from these noise sources, we arranged the noise sources as shown in [Figure 2](#). The attenuators aa, bb, cc were 20, 23 and 23 dB respectively, for the three noise sources A, B, C (s/n 22772, 22771, 22269). The output from the splitters was then measured as -36.6, -36.2 and -54.2 dBm/1GHz@3072 MHz, although perhaps a band center of 2560 should have been used.

The fiber could then be connected to the Orca box which splits the two channels (in normal usage, polarizations) and feeds the 4 signals into the four up/down converters. The attenuators on the 4 up/down converters were set to provide approximately -12 dBm on the IF ports to the iADC cards within each RDBE-S unit. Thus each pair of converters (A/B or C/D) with its corresponding RDBE-S unit can then be used to simulate a station as shown in [Figure 3](#). A single Mark6 unit was then used to record the full 16 Gbps (4 Gbps on each of the four CX4 cables).

Some pictures of the Mark6 recorder are shown in [Figure 4](#). This recorder is a prototype so all of the power and data cables are hanging free. The production version will be tidier. The RDBE-S units and the fiber transmitter are shown in [Figure 5](#). The 3-spitter configuration sits atop the fiber transmitters on the right. The noise sources are still within the noise box sitting atop the RDBE-S stack on the right.

A closeup of the back of the RDBE-S is shown in [Figure 6](#). The Maser and PPS connect to jumpers J1 and J2, the four IFs are connect to J3 through J6 (in increasing thread order), J7 and J8 are no-connects (although probably the LO will be brought out here), J9 is the traditional `dotmon` connection for measuring the time offset between GPS and the VDIF frame times, and J10 contains a buffered PPS output (suitable for daisy-chaining as done in this experiment). Typically this last signal is 30 ns delayed from the PPS input.

Finally the front of the broadband rack with Orca box patch panel and up/down converter settings is shown in [Figure 7](#). The displays cycle between the Luff settings (6009 and 6137 in the noise experiment) and the attenuator settings (21 dB to 26 dB) for the various channels of the up/down converters. For reference, the frequency bands are shown schematically in [Figure 8](#). which also includes the settings for the eventual X-band test. (The future broadband settings are not yet finalized.)

The final I/F power measurements were as follows when initially connected to RDBE-S003 and RDBE-S004 (X = 4 and Y = 3 in the figure).

UDC	cable	dBm		RDBE IF	ADC stdev
A	A/B-0	-11.5 dBm/512@768 MHz	--	S004 IF0	29.42 33.78
A	A/B-1	-11.0 dBm/512@768 MHz	--	S004 IF1	32.72 31.08
B	A/B-2	-12.3 dBm/512@768 MHz	--	S004 IF2	28.26 29.91
B	A/B-3	-12.3 dBm/512@768 MHz	--	S004 IF3	28.54 27.20
C	C/D-2	-11.9 dBm/512@768 MHz	--	S003 IF0	30.40 28.37
C	C/D-3	-12.3 dBm/512@768 MHz	--	S003 IF1	29.76 31.57
D	C/D-0	-12.0 dBm/512@768 MHz	--	S003 IF2	28.34 27.11
D	C/D-1	-12.7 dBm/512@768 MHz	--	S003 IF3	27.20 27.26

Once connected the standard deviation of the "raw" histogram from each ADC channel could be measured (using `bpplotter.py`). Ideally, with -12 dBm power input this should be about 32 units. The last two columns above show the actual values measured at setup (Oct 3) and later (Oct 6).

After RDBE-S003 and RDBE-S004 were tested, the IF cables were swapped to the other pair (RDBE-S001 and RDBE-S002) and again measured. The UDC C values were shifted by 1 dB (from 23.24 to 22.23):

UDC	cable	dBm		RDBE	IF	ADC stdev
A	A/B-0	-12.0	dBm/512@768 MHz	--	S001 IF0	28.09
A	A/B-1	-11.6	dBm/512@768 MHz	--	S001 IF1	29.46
B	A/B-2	-12.0	dBm/512@768 MHz	--	S001 IF2	28.49
B	A/B-3	-12.0	dBm/512@768 MHz	--	S001 IF3	27.96
C	C/D-2	-11.7	dBm/512@768 MHz	--	S002 IF0	31.49
C	C/D-3	-11.9	dBm/512@768 MHz	--	S002 IF1	31.68
D	C/D-0	-11.4	dBm/512@768 MHz	--	S002 IF2	30.92
D	C/D-1	-12.1	dBm/512@768 MHz	--	S002 IF3	29.43

The noise source/up/down converter/iADC chain does not appear to be stable to better than 10-20% across the two weeks of these observations based on additional measurements made after each requantization command:

	IF0	IF1	IF2	IF3
rdbe-S001	26.78	27.91	27.99	26.31
rdbe-S001	27.11	27.77	27.42	26.72
rdbe-S001	27.45	29.28	28.14	27.65
rdbe-S002	28.62	28.61	28.27	26.97
rdbe-S002	30.59	31.62	30.77	29.05
rdbe-S002	31.38	29.40	29.95	29.20
rdbe-S003	30.01	29.02	29.54	27.31
rdbe-S003	34.00	31.95	32.48	29.76
rdbe-S003	34.52	32.71	31.25	30.08
rdbe-S003	39.86	37.20	34.43	32.43
rdbe-S004	27.53	31.35	27.37	27.75
rdbe-S004	28.31	31.16	26.73	26.71
rdbe-S004	28.94	31.26	28.05	27.70
rdbe-S004	29.86	33.86	28.74	28.83

This RDBE personality is slightly sensitive to DC bias in either the IF signal or the ADC itself. Throughout these tests, bias as measured in the histograms was seen (across all 16 IFs) to lie within the range of -1.96 to 0.52 ADC units.

An associated issue is the setting of the thresholds in the quantizer. An external "helper" script reads a sample of the raw data, computes appropriate thresholds for the quantizer, and installs them in the loaded personality. The target is for the middle two states (01 and 10) to be populated at about the 68% level. The relized values (one per thread, or 8

values per setting) are shown over the course of these tests in [Figure 10](#). The variance is somewhat greater than is desirable, but whether it is instability in the noise source itself or the procedure is still to be determined.

## Zero-baseline Fringe Results

With two bands on two stations, `fourmer` is used to join the bands after the `DiFX` correlation. Since each band at each station has different cabling, it is necessary to determine the single-band delays of each band separately prior to doing the fit for the multi-band delay. Since there are multiple copies of the same signal, similar pairs were arbitrarily labelled as R and L polarizations to simplify processing in `DiFX` and `fourfit` and to simplify bookkeeping. Thus identical results should obtain for RR as for LL correlations (each of which uses independent analog signals).

The individual per-band/per-polarization fringe plots are shown in [Figure 11](#), [Figure 12](#), [Figure 13](#), and [Figure 14](#). The common noise was not particularly flat in across these frequencies; however the common signal cross-power spectrum is representative of what is seen on the spectrum analyzer in the lab.

The procedure for doing the complete fit of the joined bands was established by the time of run #33; the full fringe plot for the station pair in use at that time are as shown in [Figure 15](#) and [Figure 16](#). The script performing these fits must currently run `fourfit` repeatedly to measure the relative delays of the RDBE-S units in the different bands, and then work out manual phase calibrations. (This is discussed below.)

The `fourmered` fringe produced frequency-weighted correlation amplitudes of 204.427 [R:  $(258.310 + 150.544)/2$ ] and 209.265 [L:  $(266.784 + 151.746)/2$ ] with a  $\sqrt{2}$  increase in SNR.

Corresponding pairs of IFs and CX4 cables were then shifted from RDBE-S001 and RDBE-S002 to RDBE-S003 and RDBE-S004 to allow a 4 station correlation. (Some care must be taken to ensure that the connected IFs correspond to the threads present on the respective CX4 cables.) The results are entirely similar and the single-band delays between the boxes were found to be consistent with the daisy-chain of the maser 1PPS signal from RDBE-S002 to RDBE-S001 to RDBE-S004 to RDBE-S003.

One quantitative issue that is currently unresolved is that while the strengths of the sources are such that the weaker, common source is about 20 dB down from the stronger sources, the correlation amplitude observed is about 2% across the pair of bands. Attenuators were added and removed to establish that the correlation amplitude varied correspondingly. Mostly likely the power measurements made with the spectrum analyzer were not over the correct band or accurately enough. Unfortunately the noise-assembly could not be network analyzed.

## MARK5B / VDIF Correlations

A special version of the VLBI2010 personality `PFBG_1_4.bin` exists that can run within the RDBE-S hardware. With all four RDBE-S units connected to the common noise source it was possible to capture zero baseline pairs with both personalities simultaneously. Analysis of this is ongoing and will be reported at a later time.

## Infrastructure at Westford and GGAO

In order to make the tests more convenient, the 1 Gbps Ethernet supporting the Haystack software correlator was extended over an existing 10 Gbps fibre to Westford as shown in [Figure 17](#). Eventually, this can be used to support a 10 Gbps connection between Haystack and Westford.

The DiFX software correlation software was installed on the Mark6 recorders to allow local (e.g. at Westford) correlation of the zero-baseline test data. However, with a 1 or 10 Gbps network it does become feasible to consider operating the Mark6 recorders as a remote `Datastream` on the DiFX correlator at Haystack. This is a work in progress.

The arrangements for a high-speed link to GGAO are currently being implemented. The current network in the MV3 trailer and connection options are shown in [Figure 19](#) and [Figure 18](#). The physical layout of the MV3 trailer is as shown in [Figure 20](#). The Mark6 and the two RDBE-S digital backends were installed in the rack labelled "Available". The `monkey` machine is not used in the test but provides passive test support and troubleshooting.

## Current RDBE-S Setup/Mark6 Recording/DiFX Commands

These tests were run using shell scripts run from the command line interface, with a number of environment variables controlling the various configuration parameters. Since there are many components to the test, it is useful to separate the commanding into separate `xterm` windows. There are three general areas: RDBE-S setup, Mark6 recording and finally the DiFX analysis scripts.

A "cheat" sheet of commands is to be found in `$BMR_ROOT/trunk/doc/wftest/simple_cmds.txt`. That directory also contains additional files with everything left out of this menu.

### RDBE-S Commands

The RDBE-S needs to have its time set, and to have a server launched:

```
to root@$host
cd /home/roach/bin
export ntpserver=192.168.1.7    # mark6 at Westford
export ntpserver=192.168.1.205 # mark6 at GGAO
ntpdate -d -s $ntpserver | grep filter.offset
ntpdate -b -s $ntpserver
ntpdate -d -s $ntpserver | grep filter.offset
./rdbe_astro 5000 6 </dev/null >/dev/null 2>&1 &
```

The `-d` option is a debugging diagnostic, the `-b` option forces `ntpdate` to immediately apply the time correction rather than the default behavior of a gradual slew. The RDBE-S personality needs configuration beyond what the normal RDBE gui can presently accomplish—a script is available for this purpose.

```
# configure on the fly:
image=RoachAstro8Gbs_sept14g.bin
side=...
make-rdbe-conf.sh $side temp $image |\
rdbe_setup.py -h $host -v 1 | bitstatus.py
```

The `side` variable has to be coordinated with the usage of the `record-vdif.sh` discussed below. The same issues are relevant to the Mark6.

After loading, two other things need to be done. First, the quantizer of the FPGA needs to properly set the thresholds between the bits states. Currently an external script set these. Next the time will need to be set for the personality. At the moment it is easiest to check that with BMR software discussed later (`record-vdif.sh`), but the Mark6 software will should eventually support this.

```
# requantize with zero handled correctly
requantize.bash $host
# set the time
make-rdbe-conf.sh $side time |\
rdbe_setup.py -h $host -v 1 | bitstatus.py
# check the time
record-vdif.sh $side eval 1 1 4
```

This last script provides feedback on the packet times and bit states relative to the current UNIX clock on the recorder. If everthing is properly NTP synched, then the time slew should be at most a few ms (allowing for packets in the backlog queue, and assuming the maser tick is near the UTC tick). *E.g.* at the end of a one-second scan starting at 2011y290d14h20m58.0000s

```
hops_time -q Vex    23@9382858.000066387 == 2011y290d14h20m58.0000s
hops_time -q Clock 23@9382858.000066387 == 1318861258

Xx_37: 1318861259.000058889  grab  23@9382858.000066387 ->active, 1000 pkts
Xx_37: 1318861259.000058889  first 23@9382857.996864000 (+0.003202386)
Xx_37: 1318861259.000058889  final 23@9382858.996672000 (-0.996605614)
```

the first packet was 32 ms early (*i.e.* one full packet) and ended one second later.

## Mark6/BMR Scan Commands

Since the Mark6 software is still in development, it has been convenient to use some of the Burst Mode software for test recordings. To make a Mark6 recording, you need to become the root user. To make a single scan and convert it:

```
su -
source /opt/mit/mark6/etc/rc.mark6
single-scan <SCAN_NAME> <DURATION>
scan-convert <SCAN>
```

It can also run a schedule:



```

su -
source /opt/mit/mark6/etc/rc.mark6
session-run -h

Usage: session-run [options]

Options:
  -h, --help            show this help message and exit
  -s SCHEDULE, --schedule=SCHEDULE
Schedule
  -t, --test            Run a test schedule based on current time.
  -d, --dryrun          Dryrun of schedule.

session-run -t

```

where the schedule file is xml in the form:

```

<experiment name="r1504" station="westford" start="2011297170000"
end="2011298170000">
  <scan experiment="r1504" source="1548+056" station_code="Wf"
    start_time="1319475600" duration="89" scan_name="r1504_Wf_297-1700"/>
  <scan experiment="r1504" source="1741-038" station_code="Wf"
    start_time="1319475775" duration="43" scan_name="r1504_Wf_297-1702"/>
  ...
</experiment>

```

The BMR driver script `record-vdif.sh` mentioned above was also used to make 1-second recordings, *e.g.* for the 4-station test:

```

( record-vdif.sh KkMm- eval 1 1 4 10 S001 &
  record-vdif.sh OoQq+ eval 1 1 4 10 S002 &
  record-vdif.sh XxZz- eval 1 1 4 10 S004 &
  record-vdif.sh TtVv+ eval 1 1 4 10 S003 )

```

There are 4 10 GbE interfaces—the Mark6 captures packets from all four of them and splits the threads out by destination UDP port (4201..4208). The `record-vdif.sh` script parses its first argument to determine the hi/lo side of the addressing pair, and the letters to label the thread stations, *e.g.* for a direct two RDBE-S comparison, these were used:

```

# - is lower address pair   if's 0..3 (192.168.5. 96/104)
# + is higher address pair  if's 4..7 (192.168.5. 112/120)

export if=- # iterm: maxwell / tigress eth[45]: disk[23]
export rdbe=S001 side=KkLlMmNn$if
export if=+ # zterm: monarth / tigress eth[23]: disk[01]
export rdbe=S002 side=OoPpQqRr$if

export if=- # iterm: maxwell / tigress eth[45]: disk[23]
export rdbe=S004 side=WwXxYyZz$if # iterm
export if=+ # zterm: monarth / tigress eth[23]: disk[01]
export rdbe=S003 side=SsTtUuVv$if # zterm

```



(The point of all this is that the 16 IF inputs to the 4 boxes needs to be tracked through to `fourfit`'s single letter station labelling scheme as K..Z.) The `record-vdif.sh` script also needs environment variables `SAVE_PATH` and `SAVE_VERS` to know how to label the files and where to store them. These things are hardwired (at present) in the Mark6 software. More details on the BMR software are to be found in Memos 025 and 026.

## DiFX Correlation Scripts

The DiFX (zero-baseline) correlation scripts currently treat each VDIF thread as a separate station and work out how to run DiFX from the file names. (This is not required, merely convenient.) Typically a setup script (`setup_m6v.sh`) is invoked to construct environment variables which get passed to the `grind_m6v.sh` script that actually runs the correlation and post-processing. The older script pair (`setup_m5v.sh` and `grind_m5v.sh`) was also used; and a modified version `grind_m6x.sh` was needed to handle polarization comparisons.

DiFX is intended to be driven by a `vex` file, which is normally created by planning software, but in this case has to be created for each test. The `vex` file, plus an (also constructed) `v2d` file, are sufficient for the correlation. Then a post-processing script is generated to run `fourfit`. For the multi-band correlations, all of this must be driven by a script that can run `fourmer` to join the bands. In this work, several scripts `wf*.sh` were used:

```
# correlate S001 and S002
ENUM=4000 ./wf12_222.sh
# correlate S003 and S004
ENUM=4000 ./wf34_222.sh
# S001, S002, S003 and S004
ENUM=4001 ./wf_412.sh
```

In each case, additional environment variables might be needed to control some of the finer points—one must consult the scripts for documentation.

## Planning for X-Fringe Test Westford & GGAO

For the X-band fringe test, the normal geodetic X-Band receiver feed will be split 8 ways to the 8 IFs of the RDBE-S units to capture the same signal 8 times (for an aggregate of 16 Gbps). For the GGAO 12-meter site, the cables should be configured as normal for a broadband experiment with each up/down converter handling both polarizations in some frequency band. This is shown in [Figure 22](#). In the X-band fringe test, all the up/down converters should get the same Luff frequency (7645 MHz), but the attenuators need to be set appropriately to achieve the normal -12 dBm/512 MHz power level suitable for RDBE-S IF input. For the eventual 2 GHz broadband experiment, a different set of Luff frequencies will be used.

The test was planned to be conducted as a tag-along (for the 12m at GGAO) following Westford's schedule on the R1504 session (Oct 24/25). The minimum scan duration at Westford (43 seconds) was applied to the strongest sources

scan	line	source	az	el	UTCstart	UTCstop	dur	flux
-----	----	-----	---	---	-----	-----	----	----
297-1702	24	1741-038	124	26	17:02:55	17:03:38	0:43	4.66
297-1741	177	OJ287	288	11	17:41:23	17:42:06	0:43	8.96
297-1746	211	1124-186	233	9	17:46:41	17:47:24	0:43	1.92
297-1807b	279	1741-038	139	35	18:07:28	18:08:11	0:43	4.66
297-1827	381	1156+295	275	42	18:27:55	18:28:38	0:43	2.13

so presumably these will give us the best shot. Unfortunately, these are all at low elevation (best for geodesy) and may be problematic with the trees at GGAO (due to be trimmed soon) and radar mask. (The mask is an arc of 60 degrees around azimuth 192, elevation 0. The elevation of the mask as function of azimuth is (trig functions in degrees, azimuth within 60 degrees of 192),  $El = \arcsin(\sin(60) * \cos(\arcsin(\tan(Az - 192) / \tan(60))))$ .) Of these, 1741-038 is the best candidate. Prior to the R1, 4C39.25 (10 Jy) is available at both sites for an initial testing. For simplicity, the GGAO 12m will point only at 4C39.25 and then 1741-038:

```
Westford:
  10:00a EDT - 12:30p Ken points at 4C39.25
  12:30p EDT - 13:00p Ken switches over to running R1

GGAO-12m:
  10:00a EDT - 12:30p Ed points at 4C39.25
  12:30p EDT - 2:30p Ed points at 1741-038 for
                297-1702 and 297-1807b
```

Looking more closely at the factors affecting SNR:

- We lose because a smaller dish: 12m/18m dishes
- We lose because of CP v LP:  $\sqrt{1/2}$
- We gain because of wider BW:  $\sqrt{400/80}$
- We may lose due to RFI issues: probably not an issue
- We may lose due to spectral shape: several dB of slope

Curiously, these factors balance out (the product is 1.05) so that we should expect SNRs of at least 20 (the minimum SNR driving the Westford scheduling) on these sources. The reduction script will be a variant of the `wf*sh` scripts used for the noise testing, except that the correct site positions, source coordinates and (projected) EOP parameters will be used:

```
* GGAO
  site_position = 1130735.400 m : -4831268.100 m : 3994243.700 m;
* Westford
  site_position = 1492206.600 m : -4458130.507 m : 4296015.532 m;
```

```

def 4C39.25;
    source_type = star;
    source_name = 4C39.25;
    IAU_name = 0923+392;
    ra = 09h27m03.01391s;
    dec = 39d02'20.8520";
    ref_coord_frame = J2000;
enddef;
def 1741-038;
    source_type = star;
    source_name = 1741-038;
    IAU_name = 1741-038;
    ra = 17h43m58.85613s;
    dec = -03d50'04.6167";
    ref_coord_frame = J2000;
enddef;
def EOP297;
    TAI-UTC = 34 sec;
    A1-TAI = 0.03439 sec;
    eop_ref_epoch = 2011y295d;
    num_eop_points = 5;
    eop_interval = 24 hr;
    ut1-utc = -0.342420 sec : -0.343780 sec :
              -0.345310 sec : -0.346890 sec : -0.348390 sec ;
    x_wobble = 0.18690 asec : 0.18700 asec :
              0.18700 asec : 0.18680 asec : 0.18670 asec ;
    y_wobble = 0.34470 asec : 0.34350 asec :
              0.34220 asec : 0.34090 asec : 0.33960 asec ;
enddef;

```

The 12m receiver is linearly polarized (H and V). It's not clear that the analysis chain fully supports linear polarizations with circular polarizations, so we will process all threads as RCP and reconcile the results after `fourfit` is run.

## X-Fringe Test Results

Prior to the start of the R1 session, we were able to point the GGAO 12m dish at the 10 Jy source 4C39.25. The first fringe was obtained from a 23 second scan in thread 1 from both stations, as shown in See [Figure 23](#). The Westford X-band receiver only had 400 MHz of useful bandwidth, so this `fourfit` fringe was made with a passband limiting the analysis to between 8592.0 and 9032.0 MHz. The diagonal pattern near the DC edge of the band in the cross-power plot is a plotting artifact. The nominal GPS-dotmon values of 10.9 and 3.2 us were used for the correlation, so the single-band delay residual is plausible.

Rather than follow the Westford schedule, it was deemed easiest to begin with a pointing at one of the early, bright sources (1741-038). The fringe from the first such scan is as shown in See [Figure 24](#).

Unfortunately, one of the RDBE-S units (RDBE-S003) was damaged in shipment and we were unable to repair it and complete its setup prior to the loss of the 1741-038 source below the radar-avoidance mask at GGAO. So the scans on 4C39.25 and 1741-038 were only made with one RDBE-S unit (RDBE-S004).

Equally unfortunate—as we were preoccupied with resurrecting the RDBE-S we did not notice some remaining tuning issues with the Mark6 and so these scans were made with

some small packet loss. *E.g.* the dip in the amplitude at the beginning of the 1741-038 is due to such losses.

Thus we shifted to a different, bright source (0552+398)

```
def 0552+398;
  source_type = star;
  source_name = 0552+398;
  IAU_name = 0552+398;
  ra = 05h55m30.80561s;
  dec = 39d48'49.1650";
  ref_coord_frame = J2000;
enddef;
```

which was recorded (unattended) for the following scans.

scan	line	source	az	el	UTCstart	UTCstop	dur	flux
-----	----	-----	---	--	-----	-----	----	----
298-0238	2506	0552+398	60	28	02:38:30	02:39:13	0:43	4.06
298-0304b	2642	0552+398	63	32	03:04:41	03:05:24	0:43	4.06
298-0357a	2863	0552+398	69	41	03:57:09	03:57:52	0:43	4.06
298-0425	2999	0552+398	72	46	04:25:02	04:25:45	0:43	4.06
298-0451a	3135	0552+398	75	50	04:51:08	04:51:51	0:43	4.06
298-0517	3271	0552+398	78	55	05:17:50	05:18:33	0:43	4.06
298-0547a	3407	0552+398	82	61	05:47:21	05:48:04	0:43	4.06
298-0710	3781	0552+398	95	76	07:10:01	07:10:44	0:43	4.06
298-0756	4019	0552+398	116	84	07:56:57	07:57:40	0:43	4.06
298-0819	4155	0552+398	162	87	08:19:50	08:20:33	0:43	4.06

A representative scan is 298-0547 for this scan, all 8 threads were found to fringe (GGAO to Westford). The corresponding thread fringe plot (for the same thread 0) is as shown in See [Figure 25](#).

Unfortunately, the compressor for the receiver at Westford failed shortly into the R1, so the receiver warmed up (Tsys about 300K) with some loss of sensitivity relative to the earlier scans.

Very few packets were dropped (and DiFX is not sensitive to such small losses).

In all of these, plots, the curvature of the phase is most likely to be instrumental (*e.g.* dispersion in cables), and the decline in amplitude of the cross-power spectrum is dominated by the decline in the X-band receiver output at the highest frequencies.

A complete analysis is ongoing.

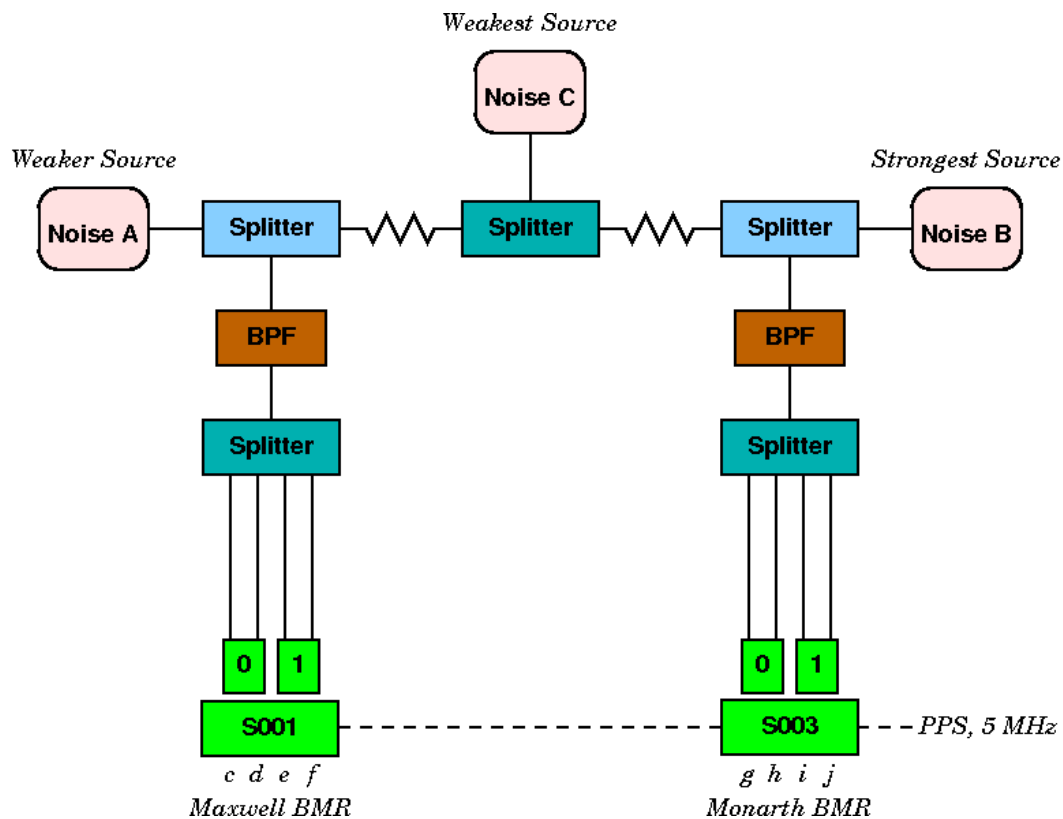


Figure 1: The final state of the noise box as used to check-out the four RDBE-S units at Haystack, and then again at Westford. However, the four network interfaces of the Mark6 was used to capture all the packets, rather than the two per machine of the two burst mode machines *maxwell* and *monarth*.

amplitude t02-3000 -----								
	S	T	U	V	W	X	Y	Z
S	10000.0	9760.0	10000.0	10000.0	36.6	61.8	96.5	57.4
T	-	10000.0	9990.0	9800.0	39.1	35.7	46.3	49.6
U	-	-	10000.0	10300.0	39.5	53.4	71.8	42.6
V	-	-	-	10000.0	39.2	44.1	49.9	39.6
W	-	-	-	-	10000.0	9780.0	9880.0	9910.0
X	-	-	-	-	-	10000.0	9920.0	9830.0
Y	-	-	-	-	-	-	10000.0	9730.0
Z	-	-	-	-	-	-	-	10000.0
amplitude t12-3000 -----								
	S	T	U	V	W	X	Y	Z
S	10000.0	9600.0	9800.0	9820.0	45.1	55.0	65.8	36.6
T	-	10000.0	9850.0	9640.0	37.3	47.5	58.2	50.1
U	-	-	10000.0	10100.0	42.4	43.4	42.7	39.6
V	-	-	-	10000.0	40.5	39.7	41.6	45.8
W	-	-	-	-	10000.0	9810.0	9920.0	9900.0
X	-	-	-	-	-	10000.0	9900.0	9840.0
Y	-	-	-	-	-	-	10000.0	9760.0
Z	-	-	-	-	-	-	-	10000.0
amplitude t16-3000 -----								
	S	T	U	V	W	X	Y	Z
S	10000.0	9770.0	10000.0	9980.0	46.1	58.6	62.1	37.7
T	-	10000.0	10000.0	9820.0	37.3	49.6	55.4	50.7
U	-	-	10000.0	10300.0	43.3	45.0	43.4	39.8
V	-	-	-	10000.0	40.7	40.1	40.6	46.1
W	-	-	-	-	10000.0	9810.0	9930.0	9960.0
X	-	-	-	-	-	10000.0	9910.0	9850.0
Y	-	-	-	-	-	-	10000.0	9760.0
Z	-	-	-	-	-	-	-	10000.0

Table 1: The total correlation amplitude is shown for the 28 baselines between the 8 IFs on two RDBE-S units with various hardware/software combinations.

tot_sbd t02-3000 -----									
	S	T	U	V	W	X	Y	Z	
S	+0.000	+0.005	+0.001	+0.002	-0.031	-0.026	-0.031	-0.029	
T	-	-0.000	-0.004	-0.002	-0.035	-0.031	-0.037	-0.034	
U	-	-	+0.000	+0.002	-0.031	-0.026	-0.032	-0.029	
V	-	-	-	+0.000	-0.033	-0.028	-0.034	-0.032	
W	-	-	-	-	+0.000	+0.005	-0.001	+0.001	
X	-	-	-	-	-	+0.000	-0.005	-0.004	
Y	-	-	-	-	-	-	+0.000	+0.002	
Z	-	-	-	-	-	-	-	+0.000	
tot_sbd t12-3000 -----									
	S	T	U	V	W	X	Y	Z	
S	+0.000	+0.005	+0.001	+0.002	-0.033	-0.028	-0.033	-0.031	
T	-	+0.000	-0.004	-0.002	-0.037	-0.033	-0.038	-0.036	
U	-	-	+0.000	+0.002	-0.033	-0.028	-0.034	-0.032	
V	-	-	-	+0.000	-0.035	-0.031	-0.036	-0.034	
W	-	-	-	-	-0.000	+0.005	-0.001	+0.001	
X	-	-	-	-	-	+0.000	-0.005	-0.004	
Y	-	-	-	-	-	-	+0.000	+0.002	
Z	-	-	-	-	-	-	-	+0.000	
tot_sbd t16-3000 -----									
	S	T	U	V	W	X	Y	Z	
S	+0.000	+0.005	+0.001	+0.002	-0.033	-0.028	-0.033	-0.031	
T	-	+0.000	-0.004	-0.002	-0.037	-0.033	-0.038	-0.036	
U	-	-	+0.000	+0.002	-0.033	-0.028	-0.034	-0.032	
V	-	-	-	+0.000	-0.035	-0.031	-0.036	-0.034	
W	-	-	-	-	+0.000	+0.005	-0.001	+0.001	
X	-	-	-	-	-	+0.000	-0.005	-0.004	
Y	-	-	-	-	-	-	-0.000	+0.002	
Z	-	-	-	-	-	-	-	-0.000	

Table 2: The total single-band delay is shown for the 28 baselines between the 8 IFs on two RDBE-S units with various hardware/software combinations.



totphase t02-3000 -----								
	S	T	U	V	W	X	Y	Z
S	0.0	-130.0	-157.0	-173.0	-92.2	-165.0	2.8	-10.5
T	-	-0.0	-26.4	-42.9	104.0	-83.6	26.1	11.5
U	-	-	0.0	-15.0	72.9	1.2	172.0	150.0
V	-	-	-	0.0	144.0	14.9	-180.0	112.0
W	-	-	-	-	0.0	-128.0	35.2	-36.0
X	-	-	-	-	-	0.0	164.0	39.3
Y	-	-	-	-	-	-	0.0	-123.0
Z	-	-	-	-	-	-	-	0.0
totphase t12-3000 -----								
	S	T	U	V	W	X	Y	Z
S	-0.0	-130.0	-157.0	-173.0	-132.0	162.0	-25.5	-49.2
T	-	0.0	-26.2	-42.6	3.6	-132.0	27.9	-27.3
U	-	-	0.0	-15.0	30.2	-44.7	121.0	-7.0
V	-	-	-	0.0	48.3	-82.1	74.4	8.0
W	-	-	-	-	-0.0	-128.0	34.6	-36.1
X	-	-	-	-	-	0.0	163.0	39.3
Y	-	-	-	-	-	-	-0.0	-123.0
Z	-	-	-	-	-	-	-	-0.0
totphase t16-3000 -----								
	S	T	U	V	W	X	Y	Z
S	0.0	-130.0	-157.0	-173.0	-133.0	163.0	-26.6	-47.9
T	-	-0.0	-26.2	-42.6	4.7	-134.0	30.0	-27.0
U	-	-	-0.0	-15.0	30.4	-43.1	122.0	-6.5
V	-	-	-	-0.0	48.6	-81.4	79.3	8.7
W	-	-	-	-	0.0	-128.0	34.6	-36.2
X	-	-	-	-	-	0.0	163.0	39.3
Y	-	-	-	-	-	-	0.0	-123.0
Z	-	-	-	-	-	-	-	0.0

Table 3: The total phase is shown for the 28 baselines between the 8 IFs on two RDBE-S units with various hardware/software combinations.




Figure 2: The three noise sources used in the noise box were recombined with broader-band splitters suitable for input to the updown converters.

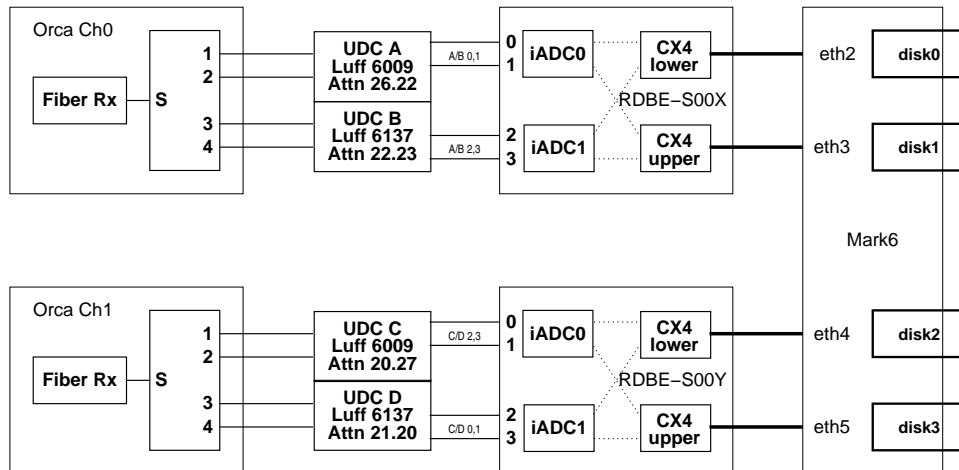


Figure 3: The fibers carrying the noise signal are received by the two channels of the Orca box, up/down converted and converted by two RDBE-S units, and recorded by the Mark6 recorder.




Figure 4: Various views of the Mark6 recorder.




Figure 5: The RDBE-S units and noise infrastructure was laid out on a tea-cart alongside the broadband rack at Westford.



Figure 6: A closeup of the back of the RDBE-S unit showing the connections.




Figure 7: The front of the broadband rack, showing the up/down converter connections to the Orca box. Note that Orca channel 0 is patched to converters A and B, while channel 1 is patched to converters C and D. The displays cycle between Luff frequency setting and attenuator settings.




Figure 8: Frequency bands in use in the 3-noise Orca/up/down configuration and the proposed X-band fringe experiment. The Up/Down converter outputs  $f_{\text{output}} = 22500 - 4 \times \text{Luff} + f_{\text{input}}$  MHz. Thus the 3 Luff tunings of 6009, 6137 and 7645 produce sky coverage 2048 - 3072 MHz. The 2nd Nyquist zone filter (center frequency 768) aliases these into an image LO at 2560, 3072 and 9104 and is presented in the fringe plots as -512 - 0 MHz.



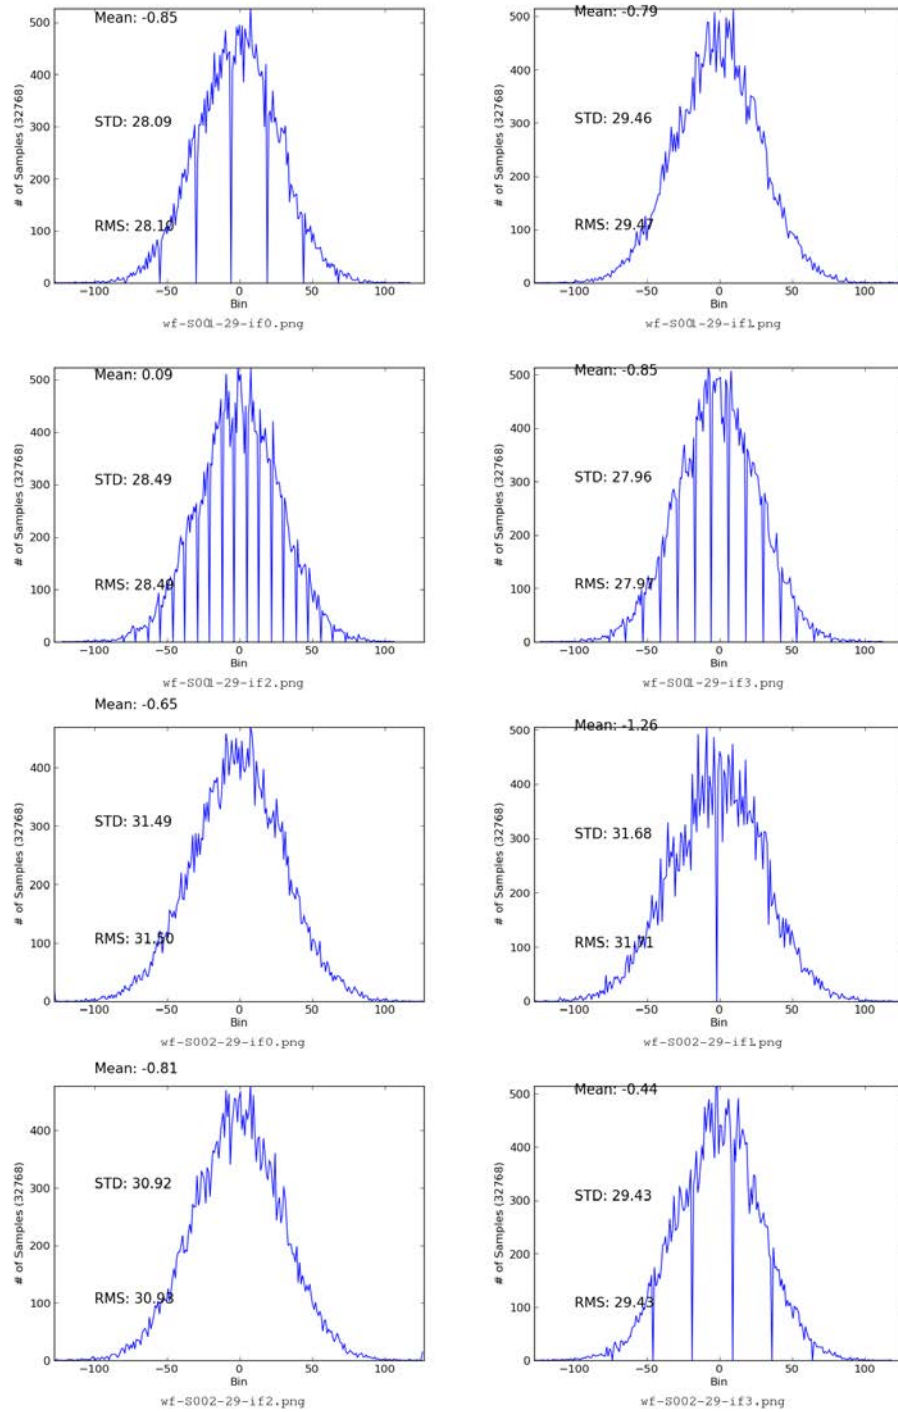


Figure 9: Histograms of the eight IFs captured midway through this testing with RDBE-S001 and RDBE-S002. The python plotting widget generates vertical blue line plotting artifacts that should be disregarded.




Figure 10: Across all scans and all RDBE-S units, the percentage of the middle states (01 and 10) is plotted. (The horizontal axis is sequential for each pairing.) The target is about 68%.



## Mk4/DiFX Fringe Plot

mark6wf\_33.vtywnf, 285-2028, ab  
S001\_aa - S002\_bb, fgroup S, pol RR






Figure 11: S001-S002 zero-baseline fringe plot for the R polarization in the low frequency band.

## Mk4/DiFX Fringe Plot

mark6wf\_33.vtywnj, 285-2028, ab  
S001\_aa - S002\_bb, fgroup S, pol RR



Amp. and Phase vs. time for each freq., 11 segs, 1 APs / seg (0.10 sec / seg.), time ticks 1 sec  
All





Figure 12: S001-S002 zero-baseline fringe plot for the R polarization in the high frequency band.

## Mk4/DiFX Fringe Plot

mark6wf\_33.vtywnn, 285-2028, ab  
S001\_aa - S002\_bb, fgroup S, pol LL



Amp. and Phase vs. time for each freq., 11 segs, 1 APs / seg (0.10 sec / seg.), time ticks 1 sec  
All




Figure 13: S001-S002 zero-baseline fringe plot for the L polarization in the low frequency band.

## Mk4/DiFX Fringe Plot

mark6wf\_33.vtywnq, 285-2028, ab  
S001\_aa - S002\_bb, fgroup S, pol LL




Figure 14: S001-S002 zero-baseline fringe plot for the L polarization in the high frequency band. Autoscaling affects the appearance of the cross-power spectrum due to a large value at the DC edge.




Figure 15: S001-S002 zero-baseline R polarization.




Figure 16: S001-S002 zero-baseline L polarization.

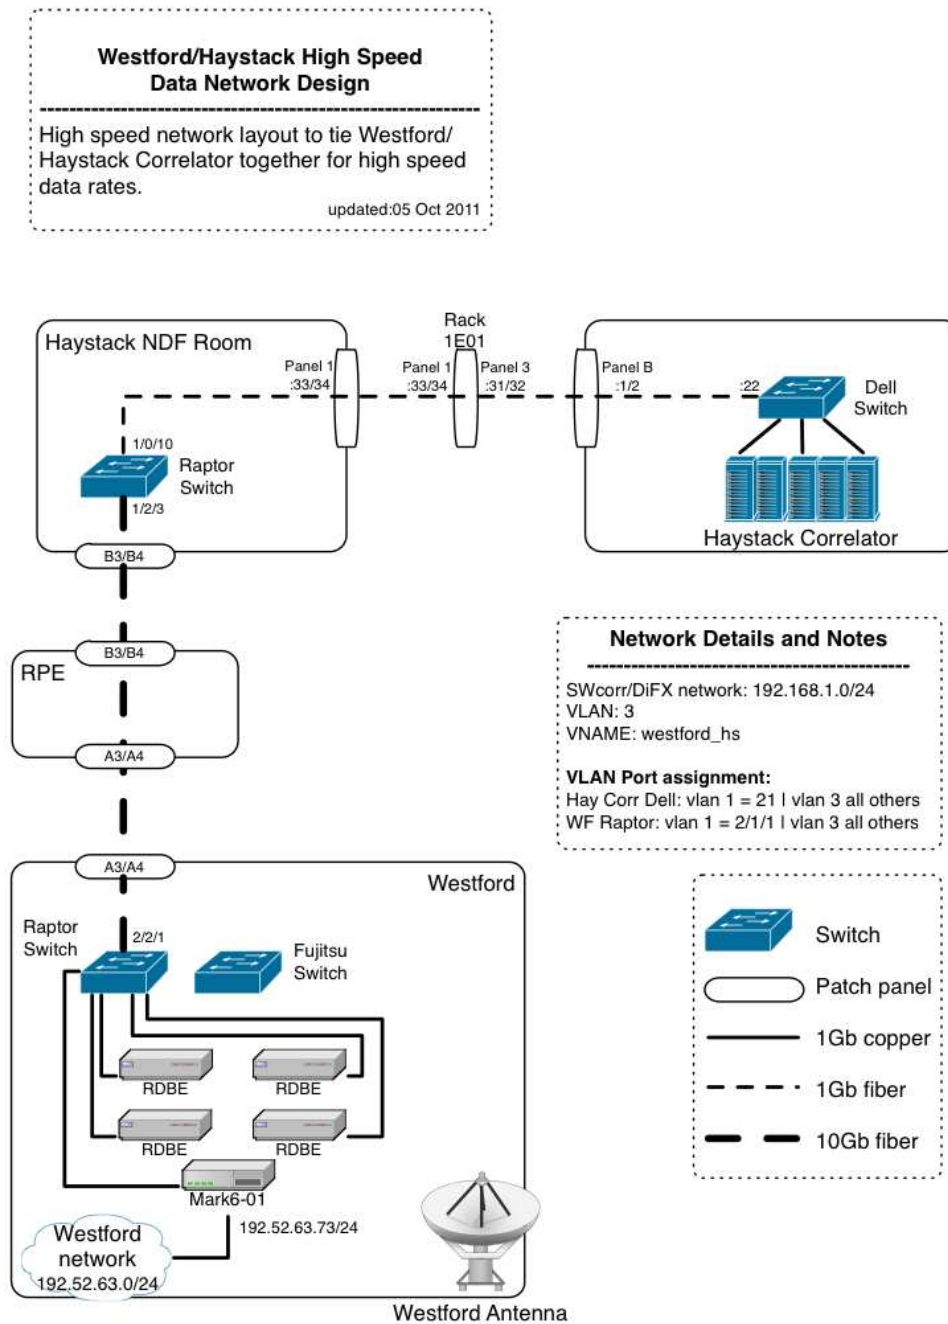


Figure 17: The network infrastructure at Haystack and Westford which will eventually allow 1 Gbps DiFX correlation.



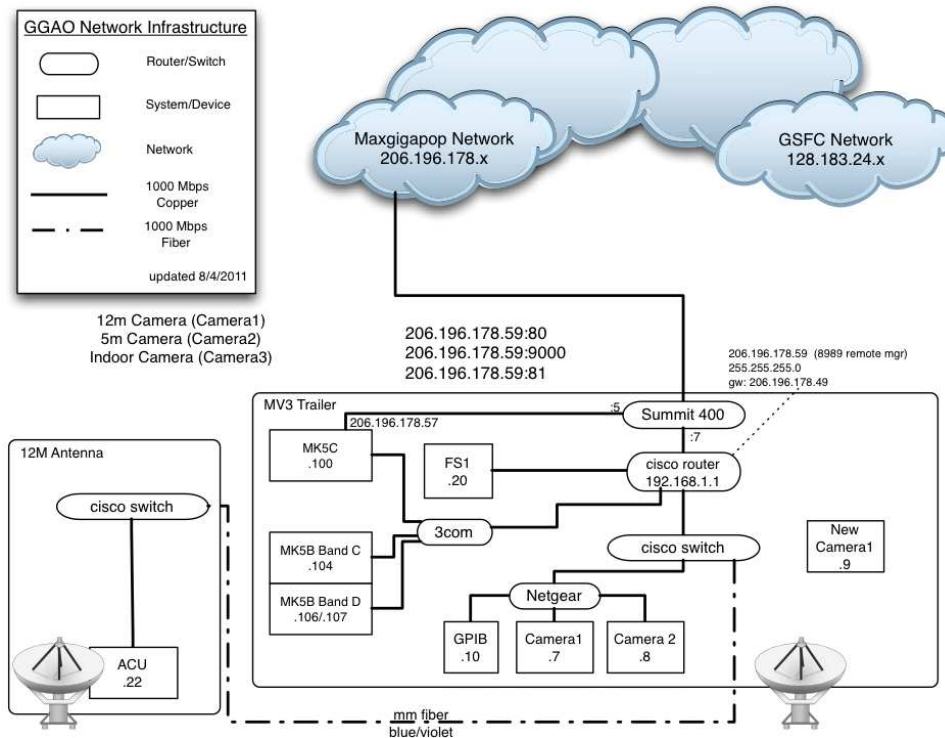


Figure 18: The network infrastructure at Haystack and GGAO which may eventually allow 1 Gbps DiFX correlation.

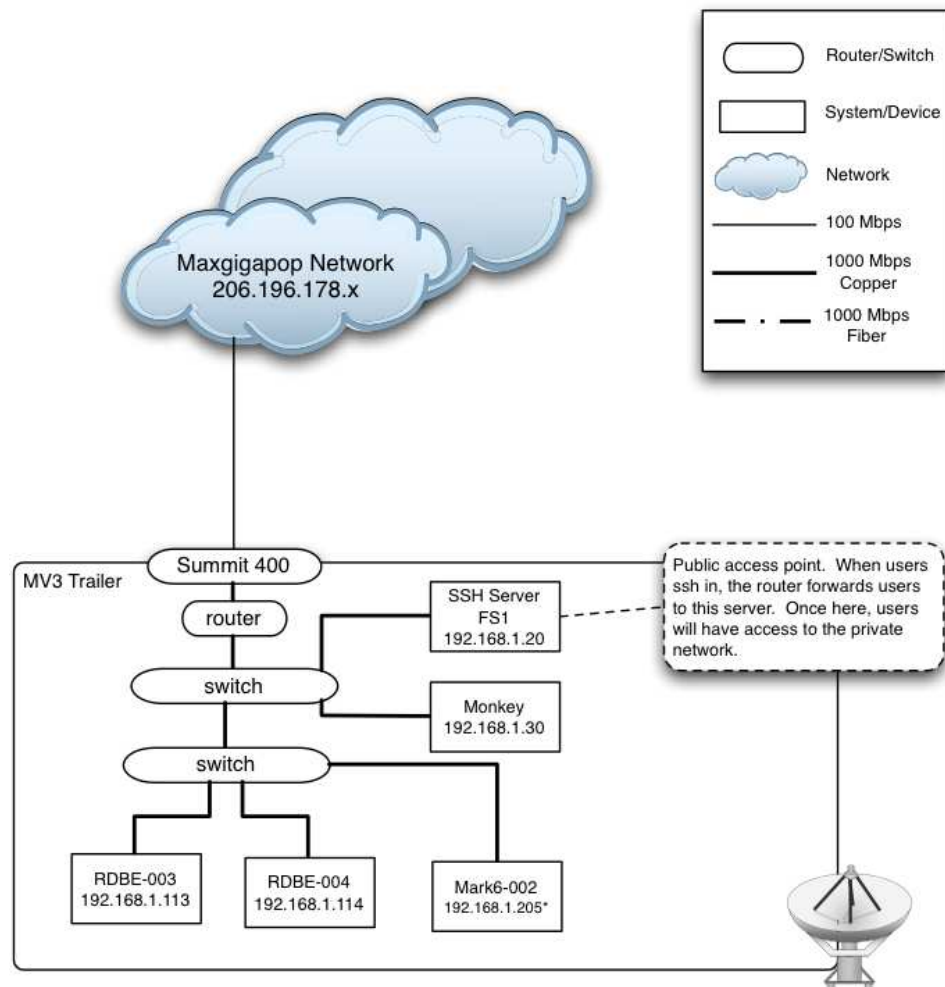


Figure 19: The network infrastructure within the trailer showing the RDBE-S units and Mark6 recorder.

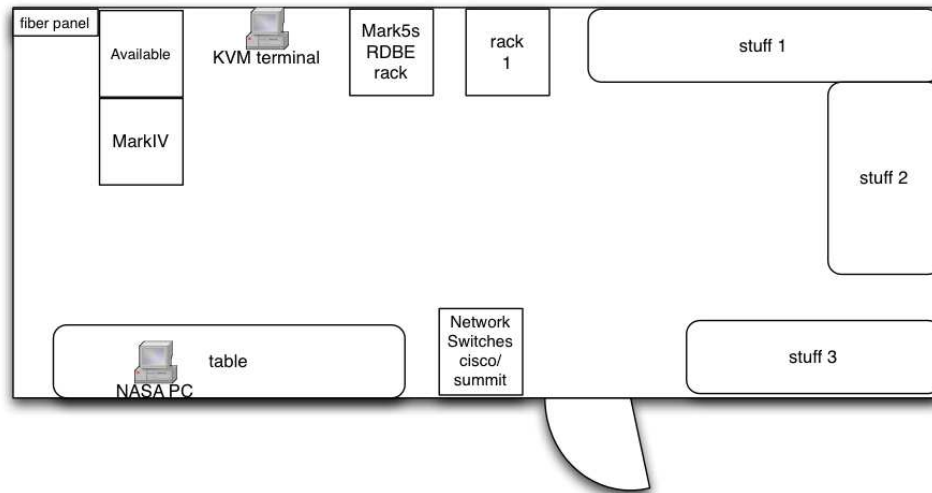


Figure 20: The layout of the MV3 trailer at GGAO is shown. The RDBE-S units and the Mark6 recorder are placed in the "Available" rack.



Figure 21: The two RDBE-S units were installed at the bottom of the available rack at GGAO. The Mark6 recorder is above the network switch (the disk bays are empty). Above the Mark6 are 4 prototype disk packs containing 32 disks.

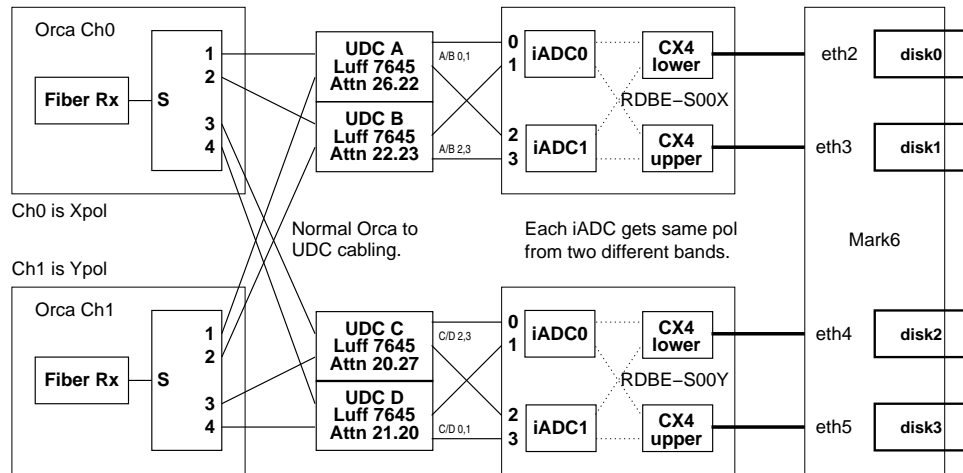


Figure 22: For the 2 GHz broadband experiment the Orca/UDC/RDSE-S units are cabled as shown. Because of the frequency spread, each up/down converter must handle both polarizations, and because of the iADC cross-talk issue, each iADC must not get the two polarizations from the same frequency.




Figure 23: GGAO-Westford fringe plot for 4C39.25 on the thread 0 IF at each station obtained with a 23-second scan 297-1556.




Figure 24: GGAO-Westford fringe plot for 1742-038 on the thread 0 IF at each station obtained with a 43-second scan 297-1702.



0552+398.vunolm, 298-0547, KW  
S001\_Kk - S004\_Ww, fgroupp X, pol RR




Figure 25: GGAO-Westford fringe plot for 0552+398 on the thread 0 IF at each station obtained with a 43-second scan 298-0547.