Mark V Memo #3

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

HAYSTACK OBSERVATORY

WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764
Fax: 781-981-0590

8 March 2001

TO: Distribution

FROM: Alan R. Whitney

SUBJECT: Second interim report on COTS-VLBI project

Abstract

A demonstration disc-based VLBI data system, dubbed Mark V, has been constructed and tested. The system consists of almost 100% COTS (commercial-off-the-shelf) components at a cost of ~$25K. Actual VLBI data has been recorded at rates up to 512 Mbps, and fringes have been obtained playing back into the Mark IV correlator at 256 Mbps. Data capacity of the demonstration system is 736 GB, equivalent to about 1.5 Mark 4 tapes, using an array of 16 low-priced consumer discs. In the demonstrated mode, the Mark V system acts as a direct replacement for a Mark IV or VLBA tape drive.

Introduction

Following the recommendations of the 'Interim Report on COTS-VLBI Phase I study' dated 24 January 2001, we have turned our attention to magnetic disc as the preferred medium for VLBI data storage. After a comprehensive market survey of available commercial components, a demonstration system was constructed and tested.

The Mark V Demonstration System

A simplified block diagram of the Mark V system is shown in Figure 1; a photograph of the system is shown in Figure 2.

<u>Hardware</u>

The Mark V demonstration system was built with the following commercial hardware components:

Standard PC system in commercial rack-mount PC case

- ABIT BX133 motherboard
- 800 MHz Pentium 3 processor
- 256 MB main memory

General Standards HPDI32 digital I/O interface card

- 32-bit wide input/output bus at up to 20 MHz clock rate; ECL compatible
- 1 MB on-board FIFO buffer

Boulder Instruments StreamStor 816 multiple-IDE disc interface card

- Supports up to 16 IDE discs at continuous maximum aggregate streaming rate of 800 Mbps
- Actively manages data distribution to discs to maintain streaming rate

2 - StorCase DS400-SR 9-disc-capacity chassis w/ power supply, each used to house 8 disc drives (1 slot unused)

16 - IBM Deskstar ATA-33 disc drives, 46 GB capacity each, and each mounted in a CRU Dataport V hot-swappable disc carrier

One standard 25"-high rack

The only non-commercial hardware in the system is a small crystal-replacement circuit on the digital I/O card to allow the output data rate to be driven from an external programmable synthesizer.

The General Standards and StreamStor interface cards are designed to act as either master or slave devices on the PCI bus, allowing a direct data transfer from one card to another without buffering through the main system memory. This allows nearly the full-bandwidth capability of the PCI bus to be utilized for data capture or playback. In particular, the 32-bit PCI bus operating at 33 MHz has a maximum theoretical bandwidth of 132 MB/sec.

Interfaces to VLBI equipment

Recording: The 32-bit wide I/O interface on the General Standards (GS) digital interface card was attached directly to the differential-ECL output of the Mark IV formatter for recording the equivalent of 32 Mark IV tape tracks. The Mark IV formatter clock acts as an external clock to drive the GS card at the selected tape-track data rate (4.5, 9 or 18 MHz). Because the Mark V is directly recording the Mark IV formatter output, including parity bits, the actual recording rate is 9/8 times the VLBI data rate. At a VLBI data rate of 512 Mbps (64 MB/sec), the actual aggregate recorded data rate is 72 MB/sec.

Playback: The GS 32-bit wide I/O interface was connected directly to a Mark IV Station Unit input in place of a Mark IV tape drive. Playback rate was controlled by an external programmable synthesizer in order to synchronize the playback data with the correlator.

Software

The control software for the Mark V was developed by Boulder Instruments (BI) under sub-contract to Haystack Observatory. Haystack Observatory was BI's first customer to require full-speed playback as well as recording, so BI was required to break some new ground in order to achieve the necessary capabilities. In order to quickly develop the software necessary for a feasibility demonstration, the software was developed under Windows 98, with which BI already had considerable experience, and allows only simple manual control by the user; only four user functions were implemented:

Record – start recording from beginning of media

Append – append recording to end of existing recording

Play – reproduce starting at specified byte position

Fetch data to file – starting at a specified byte position, fetch a specified amount of data to a standard Windows file-system file or files.

All 32 'tracks' are always recorded or reproduced. As indicated above, the record and playback rates are both controlled by external clocks.

The StreamStor interface allows recording/playback on 1, 2, 3, 4, 5,6, 7, 8, 10, 12, 14 or 16 discs, depending on user data rates and storage-capacity requirements.

As indicated above, recording utilizes a direct PCI-to-PCI data transfer to achieve a maximum recording rate of 72 MB/sec (32 'tracks' at 18 Mb/sec/track). Due to a firmware 'bug' uncovered in the GS I/O card, playback could not utilize a direct PCI-to-PCI transfer in the first software release. As a result, playback data is first read to main memory and then transferred to the GS card, which doubles the data rate on the PCI bus and currently limits the reproduce data rate to 36 MB/sec.

Testing

The testing strategy was a simple 3-step procedure:

1. Goal: Test record/playback

   Procedure: Attach the Mark V to a Mark 4 formatter and record 32 tracks of data at various data rates up to 512 Mbps (unformatted data rate); playback through a Mark IV decoder to verify.

   Results: This test was achieved without incident. Zero errors were observed on playback.

2. Goal: Record real VLBI data

   Procedure: The Mark V system was operated in parallel with the Westford Mark IV system for a short test experiment with the GGAO antenna at NASA GSFC. Data were collected at 256 Mbps using only four disc drives with a standard S/X geodetic configuration. The Mark IV and Mark V systems took advantage of the Mark IV formatter's capability to provide two independent, but identical, sets of 32-'track' outputs so that the Mark V operation had no impact whatsoever on the Mark IV recorder.

   Results: Approximately one-hour of data was recorded successfully without incident.

3. Goal: Correlate Mark V data

   Procedure: The Mark V was attached to the input of a Station Unit on the Haystack Mark IV correlator. Data synchronization with the correlator was achieved by manually adjusting a frequency synthesizer controlling the Mark V playback rate until proper synchronization was achieved. Though a bit cumbersome, this procedure actually worked quite well.

   Results: Normal fringes were achieved with no difficulty. Figure 3 shows the X-band 'fringe' plot of the correlation of the Mark V at Westford and the Mark IV tape at GGAO. Figure 4 shows the X-band 'fringe' plot from the same time segment of data, but using Mark

IV tapes at both stations. Due to non-reproducible playback errors on the tapes (primarily due to the high error-rate on the GGAO tape), the results are not identical but are within the expected limits. The S-band fringe plots are not included but are similar. The playback from the Mark V is believed to be entirely error free; the low level of errors observed on the Mark V fringe plot is believed to be due to short periods of bit-synchronizer unlocking when the Mark V reproduce data rate was changed. Note that the multi-band group delay results are identical to the picosecond (this is better than expected!)

Future Work

Replacement for Mark IV or VLBA tape drive

The Mark V at this stage is very much a demonstration system. However, only two significant improvements need to be made to the existing system in order to make it a viable low-cost replacement for either a Mark IV or VLBA tape drive operating at a maximum data rate of 512 Mbps:

1. The 'bug' in the General Standards I/O card which prevents playback using direct PCI-to-PCI data transfer needs to be fixed.
2. A modest amount of additional software must be developed to support remote control of the Mark V system.

Field Trials

Due to the low cost of implementation, several Mark V systems can easily be deployed to stations and correlators as Mark IV/VLBA tape drive replacements for actual field. For a modest investment, this will demonstrate the efficacy of the Mark V concept and prove Mark V usability and reliability in real-world day-to-day VLBI.

Increase Data Rate to 1 Gbps

Discussions with Boulder Instruments indicate the possibility of increasing the maximum date rate of the StreamStor card to at least 1 Gbps with some update and re-design. Certainly, the existing suite of disc drives are able to support this rate.

Mark V and VSI

Making Mark V fully compatible with the international VLBI Standard Interface (VSI) is clearly a high priority and will require the development of two interface cards, a VSI-formatter card and a VSI-deformatter card. This will allow the Mark V system to be interchanged with other VLBI data systems in the world and allow its use in a truly global sense. In addition, the use of VSI interfaces will allow the inefficiencies of recording parity bits to be eliminated, increasing effective disc capacity by ~15%. Haystack intends to design and build these important VSI interface cards for the Mark V system.

Mark V and e-VLBI

The StreamStor disc data are stored in a special format with its own integral directory. However, the ability to fetch data to a standard file allows one to do local processing or to transmit the data via network connection using a standard ftp file transfer. Thus, with suitable software, the Mark V can easily serve as a low-cost data buffer for near-real-time e-VLBI at either/both the station and correlator.

Other COTS implementations of Mark V

In the rapidly developing computer world, there are other potential COTS implementations of a Mark V look-alike system besides those from General Standards and Boulder Instruments. This combination just happened to be viable at the moment. We expect that other vendors will offer competing solutions, which should also be examined.

Summary and Conclusions

The fact that the entire Mark V demo system was conceived, built and tested within two months is a powerful argument for the efficacy and advantages of building with COTS components. The ~$25,000 cost of the system is very modest compared to any other system with similar capabilities and should decrease (perhaps significantly) with even modest quantities. Furthermore, there are no consumables, such as headstacks, within the system and the disc-media cost is expected to continue to drop to levels below tape-media costs over the next few years.

Once the 'improvements' discussed above are completed under Phase II of the program, we propose deployment in two stages. The first stage would be the phased replacement of Mark IV and VLBA tape drives at stations and correlators. The second stage would be the addition of VSI interfaces for global compatibility.

Further work on the Mark V system will depend on the availability of additional funding for Phase II. As indicated in the first Phase I interim report, a specific proposal for Phase II will be prepared and submitted in early summer 2001 to allow sponsors to program the necessary funds in FY02 (beginning October 2001).

Acknowledgements

The author wishes to acknowledge Will Aldrich for his expert assistance in helping to create the necessary interfaces to existing Mark IV equipment and development of the crystal-replacement electronics to steer the Mark V for correlator synchronization, to John Ball, Roger Cappallo and Mike Titus for Mark IV correlator support for the Mark V system, to Peter Bolis for his expert assistance in tracking down various necessary pieces of the system, and to Mike Poirier and Dan Smythe for assistance in Mark V observations at Westford.

Distribution:

J. Bosworth, NASA
T. Clark, NASA
J. Romney, NRAO
R. Schilizzi, JIVE
J. Webber, NRAO
W. Wildes, NASA

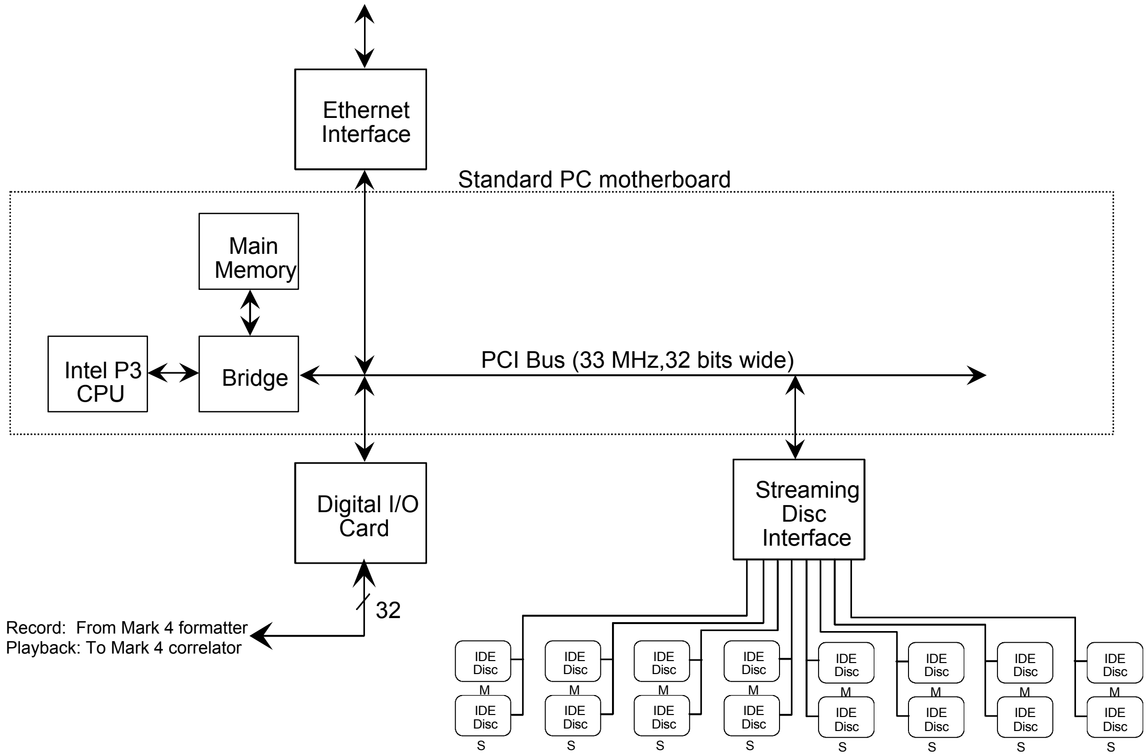

Figure 1: Simplified Block diagram of Mark V Demonstation System

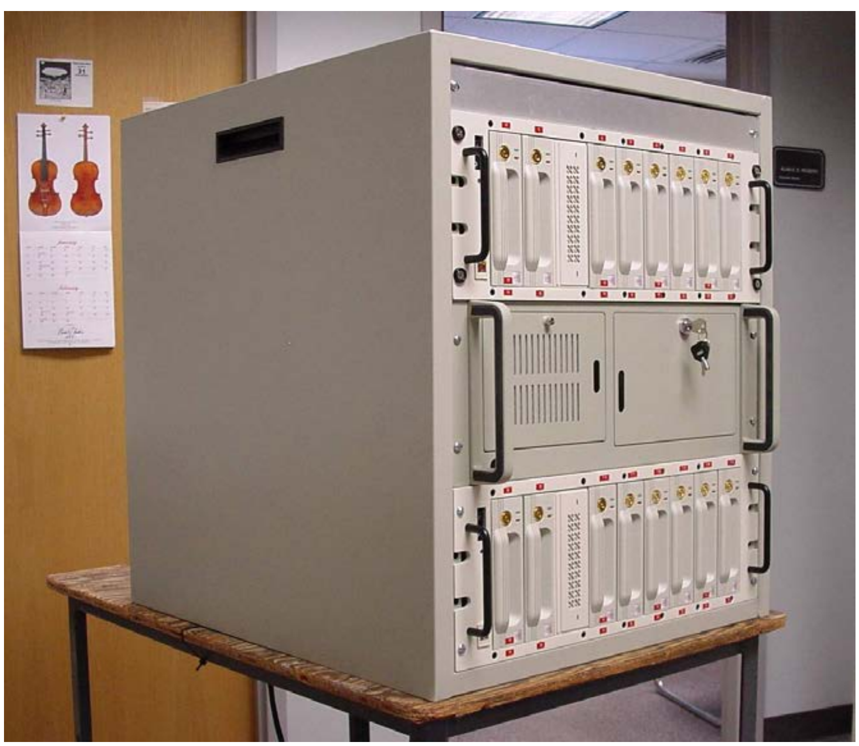

Figure 2: Mark V Demonstration System

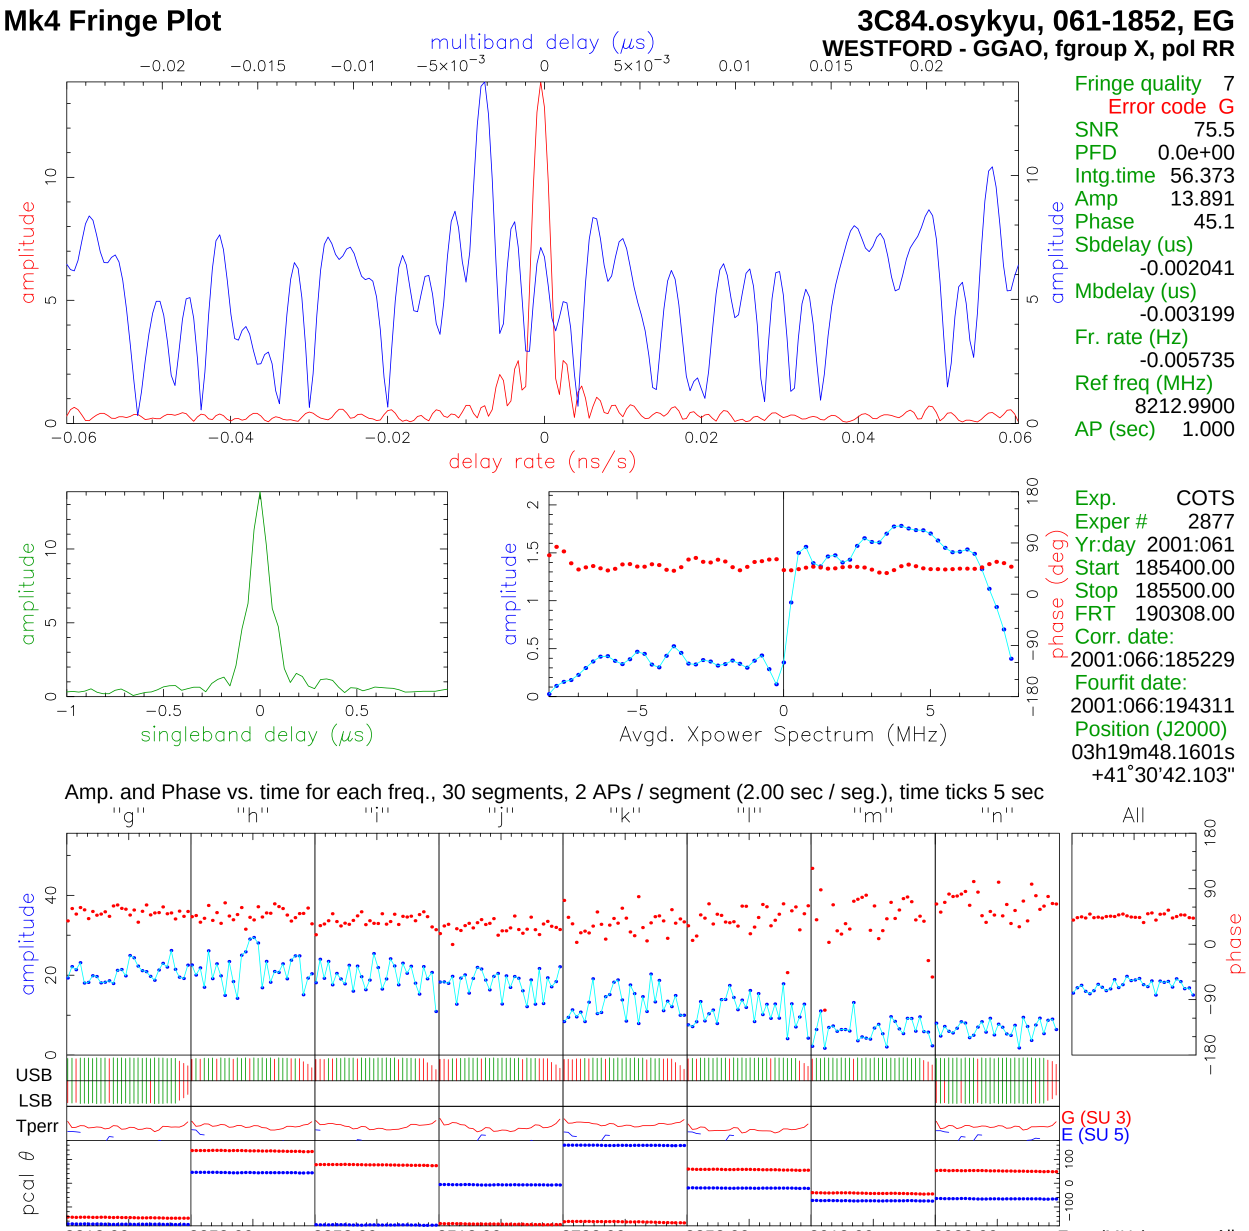

	8212.99	8252.99	8352.99	8512.99	8732.99	8852.99	8912.99	8932.99	Freq (MHz)	All
	51.6	47.6	38.5	27.7	32.8	42.4	48.8	66.1	Phase	45.1
	20.7	21.3	19.8	17.7	12.6	10.7	5.5	6.3	Ampl.	14.1
	33.1	33.1	33.0	32.8	32.8	32.7	32.7	32.7	Sbd box	32.9
U/L	60/60	60/0	60/0	60/0	60/0	60/0	60/0	60/60	APs used	
E:G	-173:-145	45:136	-176:77	-7:-173	160:-163	-21:57	-74:-43	-66:52	PC phase	
E:G	0:0	0:0	0:0	0:0	0:0	0:0	0:0	0:0	Manl PC	
E:G	44:30	41:38	42:34	44:29	46:26	51:24	43:21	47:21	PC amp	
E	X1U,X1L	X2U	X3U	X4U	X5U	X6U	X7U	X8U,X8L	Chan ids	
	2,4 6 8	10,12	14,16	18,20	22,24	26,28	30,32	3,5 7 9	Tracks	
G	X1U,X1L	X2U	X3U	X4U	X5U	X6U	X7U	X8U,X8L	Chan ids	
	2,4 6 8	10,12	14,16	18,20	22,24	26,28	30,32	3,5 7 9	Tracks	

Group delay (usec)	8.41862276727E+02	Apriori delay (usec)	8.41865475930E+02	Resid mbdelay (usec)	-3.19920E-03	+/-	7.5E-06
Sband delay (usec)	8.41863434918E+02	Apriori clock (usec)	2.2250004E+00	Resid sbdelay (usec)	-2.04101E-03	+/-	4.6E-04
Phase delay (usec)	8.41865491177E+02	Apriori clockrate (us/s)	0.0000000E+00	Resid phdelay (usec)	1.52469E-05	+/-	2.6E-07
Delay rate (us/s)	-9.37687155016E-02	Apriori rate (us/s)	-9.37680001654E-02	Resid rate (us/s)	-7.15336E-07	+/-	1.5E-08
Total phase (deg)	309.7	Apriori accel (us/s/s)	-9.16614319451E-07	Resid phase (deg)	45.1	+/-	0.8

	RMS	Theor.		
			Amplitude	13.891 +/- 0.184
ph/seg (deg)	3.9	4.1	Search (128X256)	13.324
amp/seg (%)	7.5	7.1	Interp.	13.425
ph/frq (deg)	11.0	2.0	Inc. seg. avg.	13.887
amp/frq (%)	43.3	3.5	Inc. frq. avg.	14.119

Pcal mode: NORMAL, NORMAL
Pcal rate: -9.422E-09, -2.645E-08 (us/s)

Control file: /correlator/data/2877/cf_2877 Output file: Suppressed by test mode

Figure 3: Correlation results – Westford (Mark V) to GGAO (Mark IV tape)

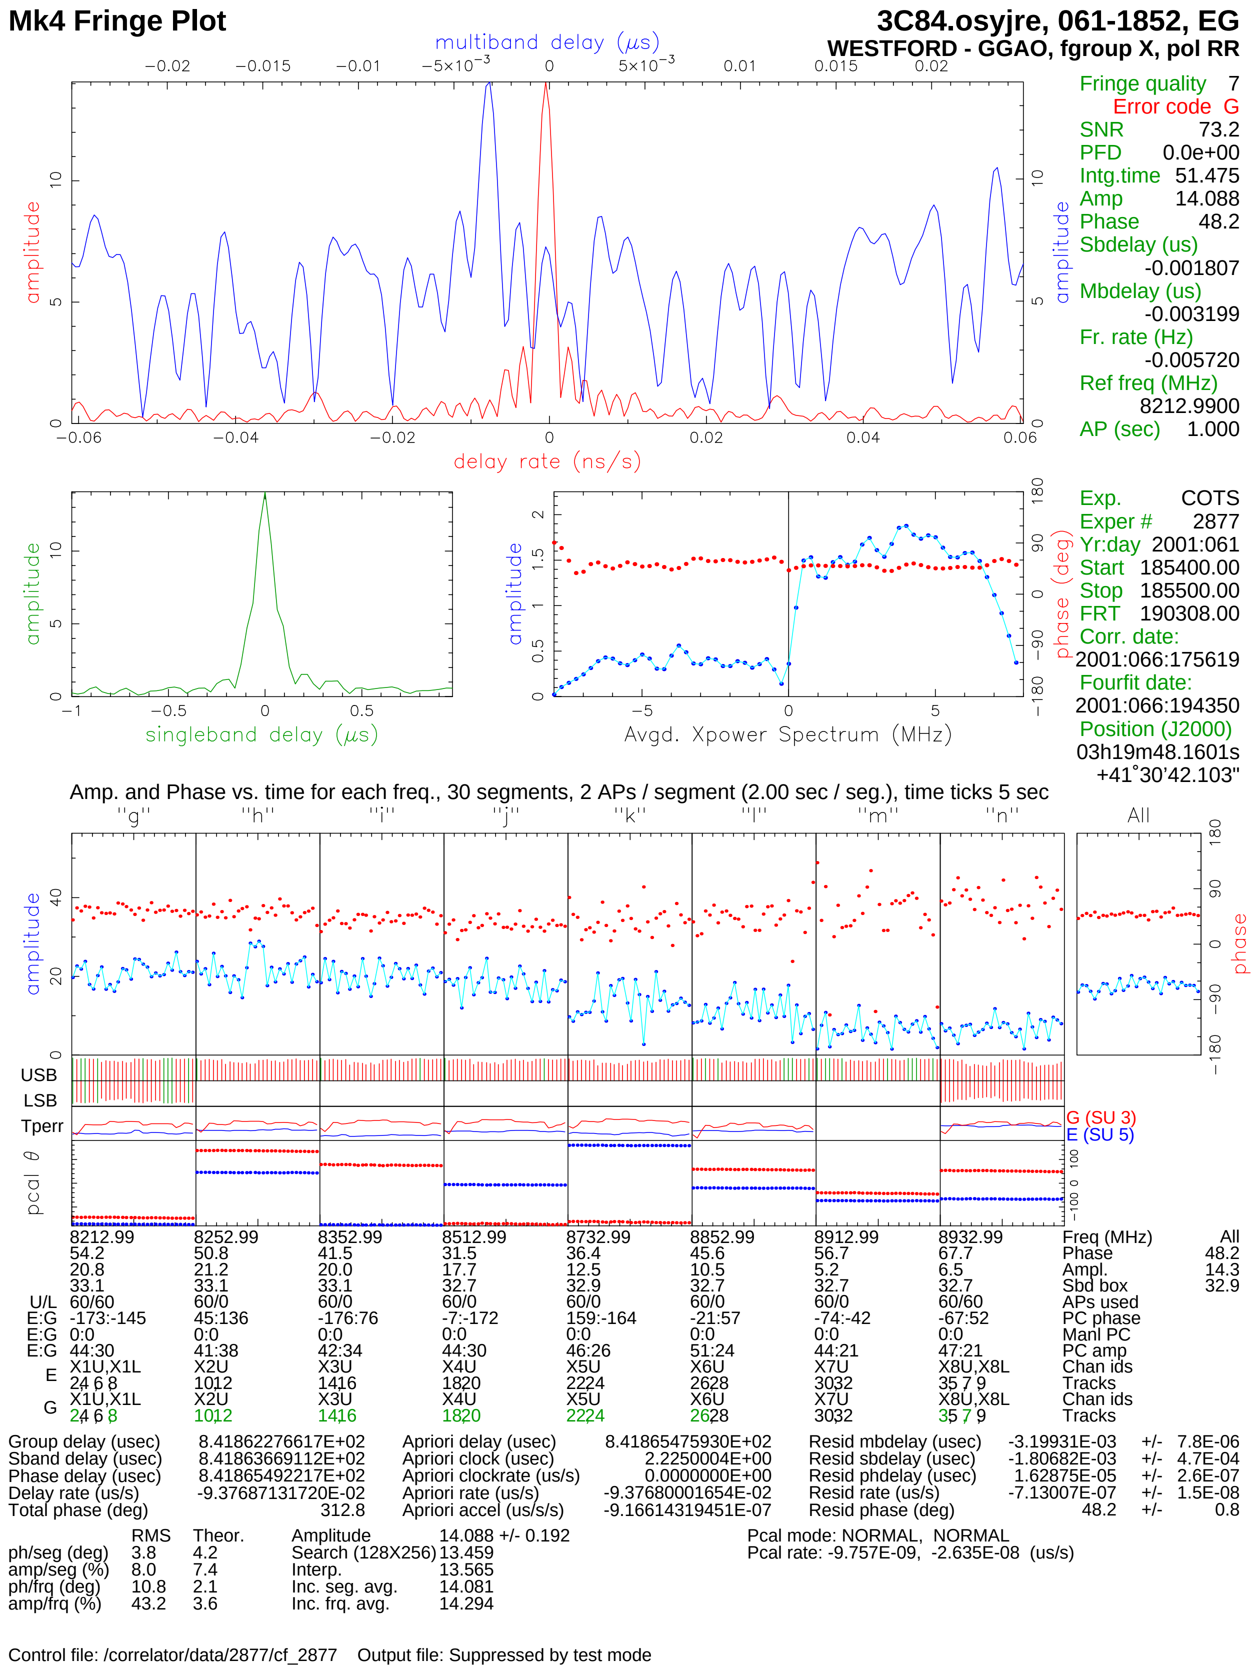

Figure 4: Correlation results – Westford to GGAO (both Mark IV tape)