VLBA Acquisition Memo # 098

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

Operating the VLBA recorder

John C. Webber
November 12, 1987

This memo is intended to help in the initial use of the VLBA recorder. It does not deal with the formatter except to note that a formatter must be connected and set up properly. Detailed description of how the recorder accomplishes its tasks will be found elsewhere; here, we are concerned only with how to *use* the system, both for checkout and for making a recording of VLBI data.

1 Hardware configuration

1.1 Introduction

The VLBA recorder is based on a Honeywell Model 96 longitudinal tape transport capable of handling 1-inch-width tape on reels up to 16 inches in diameter. Honeywell supplies a detailed manual for the basic machine. The Honeywell parts used here are the transport power supply and voltage regulator, the capstan and reel servo boards, the reel servo driver assembly, and the motors and other hardware. Modifications have been made to the tape path to enhance its tracking characteristics and improve signal stability.

To this very basic machine we have added a VME-based controller consisting of a CPU board, a commercial A/D board, and three custom boards: a transport control module, a write module, and a monitor module. In addition, modules housed in NIM bins provide the analog functions associated with moving headstack control, write drivers, and reproduce signal amplification and equalization. A 36-head narrow-track headstack assembly includes the intimate interface electronics and preamps for reproduce. Two linear power supply assemblies provide power for the Haystack-designed electronics.

Although the recorder is intended to be connected to up to 2 VLBA formatters, it is in principle capable of recording any suitable digital signals which conform to the input specifications.

1.2 The VME Controller

The control program on the CPU board is implemented in ROM, but enough RAM is available to download revised operating systems for debugging. The ROM version will auto-start on power-up. Although some basic parameters such as the vacuum level, write voltage, tape thickness, *etc.* are set to default values at power-up, many others such as the Inchworm motor speeds, LVDT scale of distance *vs.* voltage, and head motion index positions must be downloaded before the recorder can be used. This can be done through either of two paths: the MCB interface, which is on the transport control module, or the RS-232C port on the

CPU board. A stand-alone control program called *TDC*, written for IBM-compatible pc's, will control the transport through the RS-232C port; it is described elsewhere.

A commercial board provides 16 analog voltage inputs to an A/D converter and 4 D/A outputs. These are used to monitor various voltages in the transport such as the headstack position readout, headstack temperature, and vacuum level, and to set operating levels for the vacuum motor and write voltage. The details are to be found in the description of the transport firmware.

The transport controller substitutes for the Honeywell transport controller and the Mark IIIA headstack positioner controller, and contains the MCB interface. It is a considerable improvement on the Honeywell controller. This board is connected directly to the Honeywell capstan and reel servo boards, and to other hardware within the recorder, such as the vacuum level go/nogo sensor. It is responsible for all tape movement and headstack motion.

The VME write board is responsible for switching the inputs from two formatters to two headstacks and providing a limited write enable capability. The master assignment of I.F. channels to formatter output tracks is done within the formatter. The transport write board groups the recorder heads into four groups of 9 heads each, which may be assigned groupwise to either of 2 formatter inputs. The initial implementation entails only one headstack per recorder, which is sufficient for 256 Mbit/sec recording per recorder. Ordinarily, any data channels which have been turned on by the formatter (*i.e.*, their outputs are toggling) will cause the corresponding heads to record data. However, each group within the recorder may be enabled or disabled independently, so that for some recording modes the formatter operation need not be changed when changing tape direction (as in the Mark III mode B compatibility configuration). A path is also provided from the VME write board through the analog write module and the VME monitor board whereby the integrity of the digital signal path including the write driver amplifiers may be checked in "bypass" mode. This operation requires that the recorder signal path being tested must be enabled for recording and have formatter ouputs which are active.

The VME monitor board provides the capability for an acquisition recorder to check the integrity of the write signal path and to play back a tape to check the quality of the recording made in the field. Signals from the analog read module may be chosen to be from any head in either bypass or read mode. On-board bit synchronizers provide a recovered clock from the signal, and the reclocked data along with the recovered clock may be sent to the quality analyzer board within the formatter. This decoder board will extract time and other data fields and count parity errors and sync errors. In addition, the monitor board provides one on-board mini-decoder which can accept signals from any of the 4 bit syncs. It can detect any prescribed sync pattern of up to 64 bits, then grab 32 bits of data following a programmed delay; these bits are available to the VME computer. It is thus possible to perform rudimentary functions (such as reading the time written on a tape) entirely within the transport. This is very useful during checkout and provides a back-up to the quality analyzer within the formatter.

An additional VME module will be designed for correlator operation; it will provide the 36 parallel bit sync and output paths from the headstack to the correlator.

1.3 NIM Bins

Several modules are mounted in NIM bins, which are used merely to support the modules. All power and signal connections are made to the front or rear panels of these modules.

One module is simply a repackaging of the Honeywell regulator board which suppplies regulated power to the capstan and reel servo boards.

The analog signal conditioning module is a repackaged version of the board used in Mark IIIA systems. There are two each of LVDT signal electronics, power detector, and temperature electronics. The LVDT section supplies the energizing oscillator for the LVDT position sensor in the head assembly and the conditioning electronics which effectively supply a d.c. voltage linearly proportional to the headstack position. This output is fed to the A/D module. Operation of the oscillator is controlled by the VME transport controller board. The power detector supplies a detected power level for the output of the head selected by the VME monitor module and may be used to peak up on a playback track. The temperature sensor supplies the physical temperature of the head block and may be used in the future for tighter control of headstack position.

The analog write module accepts the 36 outputs of the VME write module and amplifies the outputs to a level capable of supplying the necessary drive current to the heads. A signal tapped off from the write drivers is routed back to the VME monitor module and may be used as the "bypass" signal to verify the write path integrity. The write voltage necessary to produce the approximately 10 milliamps of current in each head is controlled by the variable write power supply. The actual voltage is set by one of the D/A outputs and is nominally 9 volts. Performance may be optimized by adjusting this level depending on the type of tape and longitudinal flux transition density.

The analog read module accepts the 36 outputs from the preamplifier assembly (mounted in a shielded box adjacent to the head assembly) and provides additional amplification. Two signals selected *via* the VME monitor module output address cable are fed to one of three selectable equalizers and then to comparators which in effect supply clipped signals to the VME monitor module and are the digital data from which the clocks are recovered by the bit synchronizers.

1.4 Power

There are two power supply chassis: one supplies the +5 volts for all the digital section (VME), and another provides various voltages for the analog modules. They may be individually turned on, a convenience when debugging. The write voltage is adjustable as noted above.

Two power distribution modules provide the interconnections necessary for distributing power within the rack. These are passive bus-bar assemblies, one for the analog modules and one for the VME chassis.

1.5 Miscellaneous hardware

The layout of the various subassemblies within the recorder rack is done to form a chimney so that cool air supplied from underneath the rack is circulated through the various modules and exhausted at the top.

The vacuum column of the recorder requires a working vacuum of about 9 inches of water with respect to atmospheric pressure. A calibrated vacuum sensor provides a voltage readable by the A/D converter. The vacuum is supplied by a brushless d.c. motor whose operating speed is controlled by the output of the D/A converter module. Adjustment of the vacuum is thus provided for optimizing results with different tape thicknesses. An additional go/nogo vacuum sensor provides a signal to the hardware when there is sufficient vacuum to load the tape (this takes about 5 inches of water).

2 Writing and reading data

2.1 Initial set-up

At power-up, the firmware will assume initial values to be used for such important items as vacuum motor speed. Many other parameters need to be set as well; see the documents describing the recorder control protocol and *TDC*, the checkout program written for an IBM-compatible pc. These parameters should be downloaded to define known headstack offsets, index positions, scale of distance for the LVDT's, and other pertinent information which may vary from one recorder to another or depend on station hardware configuration.

The brakes must be released and servos disabled before tape can be loaded. All parts which touch the front or back surfaces of the tape should be thoroughly cleaned before loading, using a cotton swab saturated with pure isopropyl alcohol only. *No guarantees are made concerning the integrity of the glues which hold the headstacks together if any other solvents are used!* In a pinch, drugstore rubbing alcohol may be used; however, it contains some water and occasional other ingredients which might cause problems if used regularly.

It is essential that the transparent windows in the vacuum column be cleaned at every tape change! A dirty vacuum window will cause the reel servo to operate in a different regime from that for which it has been adjusted, with possibly disastrous results such as destruction of a tape.

Once the tape has been threaded in the correct path and the vacuum door closed, the tape may be loaded *via* the LOAD command. This should cause the tape to be sucked into the vacuum column and equilibrium established so that the reel servos are holding the tape stationary. Moving the reels by hand should produce resistance in each direction of motion; holding the reel so that the servo is drawing considerable current is also a severe test of the driver transistor assembly.

For best results, a high-speed conditioning pass is needed before data are read from or written on a freshly-mounted tape. This accomplishes three principal things:

1. The tape is thoroughly cleaned because the headstack acts as a knife-edge scraper to remove loose oxide and backing material from the active surface of the tape. This reduces the likelihood of a clogged head and insures good head-to-tape contact, which is essential for recording the shortest wavelengths.

2. The tape is warmed to its operating temperature and humidity. This helps the repeatability of the tracking signature. Without a prepass, a tape which has been stored in a colder place may show deviations of up to 10 μm on the first pass, compared to subsequent passes. This is unacceptable with the 5 μm guard bands used in 16-pass mode,

and undesirable with the 12-pass mode used currently in Mark III operation, in which the guard bands are 15 μm.

3. There is also a physical accommodation of the tape which we believe to be due to a relaxation effect. A tape which has been stored for a long time ajusts itself to a different mechanical state than if it has been recently run and subjected to stress. We are not sure of the time scale for this effect, but it contributes to making the first pass of a tape have a different signature from subsequent passes.

After the prepass in the forward direction to the end of the tape (make sure that low-tape is on to keep from running off the end!), the vacuum door should be opened and the headstack cleaned to remove the debris accumulated during the prepass. This insures that loose material will not be re-distributed when the tape is rewound. No further cleaning should be necessary for as long as the tape is mounted (up to 24 hours).

2.2 Reading a pre-recorded tape

One should always have on hand a pre-recorded tape which is known to have good data in order to check the reproduce system. Several selectors within the VME controller should be set in order to get the desired reproduced signal. For a system with only one formatter, all groups should be connected to Formatter 1 (MCB address 80; SELECT command in *TDC*). Current firmware provides for the use of only headstack #1. The desired tracks for monitor outputs A and B, equalized by one of three equalizers, are routed to formatter outputs 1 and 2 *via* MCB addresses 90-9B (the OS1 command in *TDC*). Recorder #1 as shipped should have equalizer 2 selected for 270 ips reproduction, or equalizer 1 for 135 ips playback. The synthesizer for the X21 clock to the bit synchronizers must be set to 189 MHz for 270 ips reproduction of Mark III format. This results in the on-board bit syncs getting the correct clock, which is also routed through the tape recorder to the decoder board in the formatter. Synthesizer, equalizer, and tape speed must all be appropriate for the format used to record the data.

The record function should be disabled for all groups in order to read data from tape; otherwise, any head which receives a signal from the formatter which is toggling will record over the data you are trying to read. It is *very* useful to attach an oscilloscope (anything that goes to 20 MHz will do) to the analog output of the analog monitor module (NIM bin) when trying to see what is on a tape. The oscilloscope input should be teed off the output; note that a cable already is connected from the analog output to the total power detector input of the analog signal conditioning module in the upper NIM bin. This cable should remain attached in order to allow for peaking up on a track. The oscilloscope should be triggered from the CLOCK output on the VME monitor module; this is the clock recovered from the reproduced data by the selected on-board bit synchronizer.

The headstack should be positioned at the desired index using the offset appropriate to the direction of tape motion (see the MOVEF and MOVER commands in *TDC*). If record is disabled on all heads, the trace on the oscilloscope should be a slightly fuzzy straight line at a vertical scale of 0.1 V/div. If it is a wild mess, you probably have heads enabled for recording. Starting the tape forward should show an "eye pattern" on the oscilloscope. Peaking up (try PEAK 15 command to *TDC*) will place the headstack at the maximum of the signal power. It should now be possible to recover data from the tape.

One of the on-board bit synchronizers (selected *via* MCB address 9C) may be routed to the on-board decoder and such data as time and AUX code read back, 32 bits at a time. The GET command to *TDC* provides display of these quantities. The operation of the decoder board in the formatter to provide error counts and other information is described elsewhere.

2.3 Writing data on a blank tape

Since up to 16 passes are possible, the same tape may be used for many write tests without degaussing. A high-flux degausser is desirable for erasure to the tape limit and should be used if available. In a pinch, a Nortronics hand-held degausser may be used. Depress the power switch while the unit is a few feet from the tape, which should lie flat on a table. Lower the unit into contact with the flange and *slowly* move it in a circular pattern, spiraling outward from the center. After every portion has been wiped, lift the unit about 5 mm from the flange and repeat the wiping (this isn't easy–it's attracted to the tape). Continue raising and wiping slowly until you have reached about 3 inches height; then accelerate the wiping and lifting speeds, finally turning off the degausser at least 3 feet from the tape. Flip the tape over and repeat. The entire process should take about 1 minute per side. The inner portion of the tape will retain a residual signal 3-4 dB above the ultimate tape noise level, so some tracks will not perform as well as possible. However, it is definitely better than nothing. You can look at the tape to see how effective the degaussing has been.

With the prepassed tape loaded, position the headstack at the desired index #. Make sure the synthesizer is set to 21 times the frequency appropriate for the format (189 MHz for Mark III format with 8 MHz sampling) since the divided clock is used by the formatter. The formatter should be set to the desired mode (beyond the scope of this discussion) and the appropriate groups of tracks enabled (see VLBA Acquisition Memo #97). The correct write voltage should be set (it is determined empirically and defaults to 9.0 volts). Starting the tape at the appropriate speed will cause a recording to be made. Note that if the speed is changed, the longitudinal bit density on the tape will be changed. With FUJI H621 tape, the minimum usable speed for recording at 8MHz sample rate is about 180 ips. This requires reducing the write voltage by a factor equal to the ratio of the standard speed to the newly-chosen speed. To make a standard Mark III tape, use 270 ips and 8 MHz sample rate.

When the desired length of tape has been recorded, merely stopping the tape is not enough to end the recording. Be sure also to disable all groups before moving the tape back to the starting point! It is convenient to use the tape positioning feature (MCB address B7) to move back to the starting point, since the firmware uses the "offset" rate generator for this purpose, leaving the nominal recorder speed untouched.

A final word on recording: it is possible to write a tape which has a strange-looking eye pattern but which decodes perfectly. This usually occurs if actual sampled data from the baseband converters are recorded with no I.F. signal present on the data acquisition rack. All-ones or all-zeroes are perfectly acceptable data patterns, but they will produce odd eye patterns. For test purposes, it may be better to use the pseudo-random noise generator within the formatter in order to get normal-looking eye patterns. If an I.F. signal is available (even from a noise generator), using it will also test the baseband converters and the formatter's A/D converters.

2.4 Recording a real experiment

The first fringes using the VLBA DAR and REC were obtained in the "Mark IIIA Compatibility Mode" described in Acquisition Memo #97. This is one of the ROM'd modes in the formatter and may be set up with a minimum of hassle (the aforementioned hassle is described elsewhere). For forward passes, VLBA recorder groups 0 and 2 (even tracks) are enabled, and for reverse, groups 1 and 3 are enabled. For each of 12 passes, the headstack must first be positioned to the desired index number at the offset appropriate for the impending direction of tape motion. Once the formatter mode is set, it need not be changed because the enabling of the tracks may be changed within the recorder.

As an example, here is a sequence of commands to *TDC* which will result in recording a Mark IIIA-compatible tape (assuming various initialization parameters have already been set):

```
SPEED 270
ENABLE 02
MOVEF 1
WAIT 5
FOR
WAIT 400
STOP
MOVER 1
WAIT 10
ENABLE 13
REV
WAIT 400
STOP
ENABLE
```

This sequence leaves the tape stopped at the beginning, all heads disabled, after recording two 400-second passes at index position #1.

This brief discussion is intended only to give a good idea of what is involved in using the VLBA recorder system and provide a basis for experimentation. Details have been omitted for the sake of brevity. For further information, see the detailed documentation.