VLBA TEST MEMO NO. 21

LESSONS FROM THE MARCH 27 1989 CDP OBSERVATIONS

R. C. Walker

This memo outlines many of the unsolved problems that were encountered during the March 27 CDP observations. These observations were far from smooth. The CDP observations stress the system far more than typical Network observations and therefore teach us much more about overall system performance. Special efforts to make them run smoothly should pay off in almost all that we do. At the end of this memo, I suggest that we institute occasional dry runs of CDP observations to shake out problems. Paul Rhodes also suggests that we have a debriefing after CDP runs to review any problems encountered. This is probably a good idea. In any case, this memo should serve as a record of problems already encountered.

Many of the problems of the Mar. 27 run were related to the need to frequency switch. For the long R&D run in October, the CDP project intends to use spanned bandwidths that are approximately twice as wide as those normally used. For this, they feel that 8 channels is just not enough, especially in view of the very active solar conditions which cause highly variable ionospheric delays. Therefore, we are using a recording mode with 4 channels at each of S and X band. Of the 4, the 2 inner ones are being switched in frequency so that Pie Town provides some data on a total of 6 frequencies from each band or 12 frequencies overall. Recall that the Mark III stations are recording 14 basebands. These observations are also using the Mark III double speed mode, recording 4 MHz per band and running the tape at 270 ips. One pass through a tape lasts 6.5 minutes.

Problems not related to frequency switching:

We are still having problems with OPTIMOVE picking unreasonable move directions and causing lost data. I will try to document this from the monitor data. The CDP schedulers care much more about predictability than optimization. They assume simply that the telescope will go the shortest available route to the next scan regardless of what comes later. If they see an unreasonable situation coming, they simply reorder their scans to force reasonable behavior. We should have a mode that mimics this behavior. For other observations, some optimization is appropriate. However is the current algorithm too complicated? It seems to be taking a long time to get right.

An error message indicating failure to reset the formatter clock was seen a few times per hour on average. Adding the WEAther screen to the OBS, ACU, and RECorder screens caused the message to appear more often – as often as every few minutes. When the message appeared during setup scans when there was time to check, the formatter clock seemed to be ok and there was no message on the FMTERR screen. I understand that perhaps NEWD is taking too long to run and is not checking the formatter clock in a reasonable time, causing the message to appear.

The VLBA round trip cable length measurement system seemed to be broken. At least the monitor points are all -32753. The module has been replaced since.

The prepass takes almost exactly the full time scheduled if we don't run INIASTR and LDSCHED at the same time. Running those programs slows it down. If we then have to run the programs after the prepass is done (and especially if anything goes wrong) we miss the next scan or two. We need to alleviate this situation. Some things to

investigate: a) make it reasonable to run MASTR during the prepass, b) don't break schedules at tape changes, c) schedule more time for tape changes, and d) get rid of the need for the prepass.

The weather station pressure gives highlighted values maybe a third of the time. This probably means that the MCB request was not answered fast enough. Usually the number on the screen is just the previous non-highlighted value. Occasionally it is 1419 mbar.

The wind gauge went for several hours without reading above zero. The peak wind speed also indicated zero. At the same time, there was a breeze in Socorro. Could it have been frozen? It had been raining during the day and this zero wind was seen as the temperature fell below 0 C. However, it may have in fact been calm according to Eric.

The fringe test showed significant problems with data playback quality at Haystack. This needs to be fixed.

We blew 3 fusses in the elevation drive.

It would be nice to have a scan number that keeps incrementing all the way through the schedule. For schedules longer than 60 scans, the observation number on the OBS screen looses any relationship to the scan number, especially if the system is restarted in the middle of ths schedule. If one wants to jump to a specific scan, it is necessary to search for the number of the scan with the corresponding stop time first.

We need to get better at getting schedules ready ahead of time. I suppose this is mainly my problem.

The following are from Susan Koski:

LDSCHED should check for typographical errors - at least obvious ones. For example, it passed positions of the form: ra = 19h04m32.*s.

The message, LOGER: Could not resume program USEDISK popped up once. This usually indicates that monitor data will no longer be logged until the system is restarted. In this case, it appears that USEDISK took longer to write to disk than LOGER expected, but there was no loss of monitor data. It suggests that whatever is taking NEWD so long to run is constipating USEDISK as well.

The following needs to be added to MCBID.H to allow room temperatures to be monitored:

ID	#41	#42
Beg Addr.	#4100	#4200
Block Size	#80	#80

These are for the ped. room utility module and the station bldg utility module. Power glitches currently clear these addresses and the necessary monitor data will then not be logged.

Problems related to frequency switching:

If NEWD is started before the previous NEWD is done, the system crashes. If there are screens running, NEWD can take more than 10 seconds to run. We were trying 10 second switching so we kept crashing the system until the problem was understood. We went to 20 second switching which seemed to run ok. However, Nancy Vandenberg thinks that 20 seconds is too long, especially for the next run in which some scans will be only 98 seconds long (this time they were 196 seconds). Can we shorten the time that it takes NEWD to run?

The length of the last scan of a frequency switching loop is not necessarily well determined, at least if the loop is entered

directly rather than by allowing the preceeding scan to reach its stop time. It would be easy to have that last scan be short and crash the system. Also, the scheduler must know that short scans are not allowed and schedule the switching rate and scan length such that short last scans are avoided. Since the scan lengths are being determined by the need to be a submultiple of the tape length (only about 6 minutes at double speed) and the switching time must be an integral number of correlator dump times (1, 2.5, 5, and 10 seconds with the longer ones preferred), the problem is rather constrained and it would be easy for someone somewhere to forget a constraint and set up something that won't work. Can something be done to prevent a scan change within a loop from being initiated too close to the loop end time? Perhaps such problems are a job for a schedule checker.

How well determined are the exact start times within a loop? If we use frequency switching at two or more VLBA sites, the switches must be kept aligned. I suppose we need a policy of only entering a loop from the stop time of the preceeding scan, in which case the switch times seem well constrained. This is ok for short scans but may be a problem if anyone wants to do this for long scans (I'm not sure who would).

The actual change of frequency was determined at Pie Town to occur about 2 to 3 seconds after the beep indicating the start of the scan change. Can this be shortened, perhaps by changing the order in which NEWD does things?

In order to separate the switch cycles at the correlator or in post-processing, the times of the switches must be known. Ignoring this just causes a SNR loss for the data but wipes out the phase cal detection. We need to develop an appropriate method of passing switch cycle information to the correlator, including taking into account the lag mentioned in the last point.

Note that our monitor data (including TSYS) logging intervals must be powers of two in seconds. This is unlikely to be commensurate with the switching interval. I don't think that this has any serious effects, but I'm not sure yet.

In the schedule for a frequency switching run, it is necessary to set the scan duration to some value like 20 seconds during the loop. Then, for the next setup scan, the stop time should be the overriding parameter. However the schedule uses the duration if it is not zero. Therefore, the duration must be set to zero. This causes an error message to appear on the operator's screen any time LDSCHED loads such a scan. This message should be removed or in some way bypassed for these observations. Or perhaps the stop time should be made to override the duration in all cases where it is reached first, as I originally expected, so that a large duration could be used.

During frequency switching, the operator's terminal is beeping at each switch. This causes one not to notice the beep of an error message. Perhaps the error messages should cause a quick double beep.

The frequency switching schedule is very bulky – 5244 lines for the 24 hour run on March 27. It takes a long time to move from one computer to another (about 30 minutes) and to make any changes. I have suggested some bulk reducing measures to Nancy Vandenberg but don't expect to gain more than about 20 percent. I suppose not much can be done for now. If frequency switching becomes common, perhaps a special observing mode could be provided.

Dry run recommendation:

I recommend that we add dry runs of CDP experiments to our test proceedures. Most of the problems encountered on Mar. 27 did not

require correlated data to be obvious. If we do dry runs, we should be able to find these problems without wasting the time of other observatories. We could also take time to explore the nature of any problem in detail, and perhaps make fixes, without worrying about lost data. Where in a live experiment, we avoid anything that we are suspicious can cause problems, in a dry run we can purposefully probe the limits of the system. I propose the use of CDP schedules for the dry runs because they stress the system more than most other types of observations.