Network Coordinator Report

Ed Himwich

Abstract

This report includes an assessment of the network performance in terms of the yield of usable data over a 12 month period. Overall, the data loss for 2004 was about 12.5%. A table of relative incidence of problems with various subsystems is presented. The most significant causes of data loss were antenna reliability (accounting for about 33%), receiver problems (18%), recorder problems (11%), and RFI (5-10%). The current situation for the handling of correlator clock adjustments by the correlators is reviewed. The adjustments are not currently handled in a consistent way and this has an impact on the UT1-UTC estimates from VLBI data. Some work will be required to correct this problem.

1. Network Performance

The network performance report is based on correlator reports for sessions in calendar year 2004. This report is based on 153 sessions that have been processed. There are nine sessions that were processed by the VLBA, but those results were left out of this analysis since the filed reports tend to be less detailed than those from other correlators. There are another 29 sessions from the calendar year that had not been processed or the correlator results were not available. Most of these missing sessions are being processed by the Penticton correlator, are waiting for tapes from Antarctica, or are from the latter part of the year. Roughly 80% of the scheduled station days for 2004 are accounted for.

An important point to understand is that in this report the data loss is expressed in terms of lost observing time. This is straightforward in cases where the loss occurred because operations were interrupted or missed. However, in other cases, it is more complicated to calculate. To handle this a non-observing time loss is typically converted into an equivalent lost observing time by expressing it as an approximate equivalent number of recorded bits lost. As an example, a warm receiver will greatly reduce the sensitivity of a telescope. The resulting performance will be in some sense equivalent to the station having a cold receiver but observing for (typically) only one-third of the nominal time and therefore only recording one-third of the expected bits. In a similar fashion, poor pointing is converted into an equivalent lost sensitivity. Poor recordings are simply expressed as the fraction of total recorded bits lost.

Using correlator reports, an attempt was made to determine how much observing time was lost at each station. This was not always straightforward to do. Sometimes the correlator notes do not indicate that a station had a particular problem while the quality code summary will indicate a significant loss. Reconstructing which station or stations had problems and why in these circumstances does not always yield accurate results. Another problem was that it is hard to determine how much RFI affected the data unless one or more channels were removed and that eliminated the problem. Similar problems occur for intermittent poor playback. For individual station days, the results should probably not be assumed to be accurate at better than the 5% level.

The results here should not be viewed as an absolute evaluation for the quality of each station's performance. As mentioned above the results themselves are only approximate. In addition, some problems are beyond the control of the station, such as weather and power failures. Instead the results should be viewed in aggregate as an overall evaluation of how much of the data the network is collecting as a whole. Development of the overall result is organized around individual station performance, but the results for individual stations do not necessarily reflect the quality of operations at that station.

Since stations typically observe with more than one other station at a time, the lost observing time per station is not equal to the overall loss of VLBI data. Under some simplifying assumptions, the loss of VLBI data is roughly about twice the loss of observing time. The argument that supports this has been described in the Network Coordinator's section of the 2002 Annual Report.

For the 153 sessions from 2004 examined here, there are 966 station days or slightly more than 6 stations per session on average. Of these session days about 12.5% (or about 120 days) of the observing time was lost. For comparison to earlier years, see Table 1.

Table 1. Lost Observing Time

Year	**Percentage of Days Lost**
1999-2000*	11.8
2001	11.6
2002	12.2
2003	14.4
2004	12.5

* The percentage applies to a subset of the 1999-2000 sessions.

The lost observing time for 2004 was down slightly from 2003. It has returned to levels more typical of previous years, around 12%. If these observing time losses are converted into VLBI data yield losses, then 2004 had about 25% VLBI data loss and 2003 about 29%. Whether these results reflect an actual improvement in performance is not clear (e.g., there are some questions about Gilcreek's clock performance that are discussed below). It is however reassuring that the performance has returned to levels typically seen in previous years. This is in spite of significant problems with telescopes and receivers in 2004.

An assessment of each station's performance is not provided in this report. While this was done in some previous years, this practice seems to be counter-productive. Although many caveats were provided to discourage people from assigning too much significance to the results, there was feedback that suggested that the results were being over-interpreted. Some stations reported that their funding could be placed in jeopardy if their performance appeared bad even if it was for reasons beyond their control. Last and least, there seemed to be some interest in attempting to "game" the system to improve the individual results. Consequently, only summary results are presented.

For the purposes of this report, the stations were divided into two categories: (A) those that were included in 20 or more network sessions, and (B) those in 13 or fewer. Some of the stations in the former category had been included in more than 70 sessions. The distinction between these two groups was made on the assumption that the results would be more meaningful for the stations with more sessions.

There are 15 stations in the 20 or more session category. Of these, nine successfully collected data for approximately 90% of their expected observing time. Five more stations collected 84% or more. One station in this group collected slightly less than 80%.

There are 22 stations in the 13 or fewer session category. The range of successful observing time for stations in this category was 0%-100%. The median success rate was 83%. Overall the stations in this category observed successfully about 94% of the 113 station days (12% of the total analyzed) that fall in this category..

Although the results are not being reported for individual stations, a few stations deserve special recognition for how much their data collection improved from the previous year. Four stations improved the percentage of data they collected by more than 5%. These stations are Svetloe, TIGO, Seshan, and Kokee. Given the high level of reliability of these stations it will be difficult for most of them to improve by this much again next year.

The losses were also analyzed by sub-system for each station. Individual stations can contact the network coordinator (weh@ivscc.gsfc.nasa.gov) for the break-down for their individual station. A summary of the losses by sub-system (category) for the entire network is presented in Table 2.

Table 2. Data Lost by Sub-system

Sub-System	**Percentage lost**
Antenna	32.9
Receiver	18.0
Recorder	11.1
Unknown	10.1
Miscellaneous	8.0
Rack	6.8
Operations	6.1
RFI	5.0
Shipping	1.4
Clock	0.5
Software	0.1
Total	100.0

The categories in Table 2 are rather broad and require some explanation. The “Antenna” category includes all antenna problems including mis-pointing, antenna control computer failures, and mechanical break-downs of the antenna. The “Receiver” category includes all problems related to the receiver including out-right failure, loss of sensitivity because the cryogenics failed, and design problems that impact the sensitivity. The “Recorder” category includes all electrical and mechanical problems related to the recorder system (tape or disk). This includes passes that are unrecoverable because of overwriting. The “Unknown” category is a special category for cases where the correlator did not state or was unable to determine a cause of the loss, but also includes the upper X-band IF problem at TIGO which has yet to be understood. In addition, some reports of low fringe amplitude in various channels were included in this category although some of them perhaps should have been included in “RFI”. The “Miscellaneous” category includes several small problems that do not fit into other categories, including errors in the observing schedule provided by the Operation Centers. Power failures are also included in this category. The “Rack” category includes all failures that could be attributed to the rack (DAS) including the formatter and BBCs.

The "Operations" category includes all operation errors, such as DRUDG-ing the wrong schedule, starting late because of shift problems, problems changing tapes and others. The "RFI" category includes all losses directly attributable to interference. The "Shipping" category includes data that could not be correlated because the media was either lost in shipping or held up in customs long enough that it could not correlated with the rest of the session data. The "Clock" category includes situations where correlation was impossible because either the clock offset was not provided or was wrong leading to no fringes. It is difficult to be sure in some of these cases that the clock offset was the culprit, but in some it was clear. In cases where it was not, the loss was assigned to the "Unknown" category. The "Software" category includes all instances of software problems causing data to be lost. This includes crashes of the Field System, crashes of the local station software, and errors in files generated by DRUDG.

From the results it can be seen that antenna and receiver together account for more than 50% of the losses, which is an increase of almost 10% above last year. In fact for 2004 there were several unusual receiver and antenna problems. In particular, Kokee, Wettzell, and Gilcreek had antenna problems that contributed to nearly 50% or more of each station's lost data. Some of these problems are continuing. Kokee's operation continues to be hampered while one of its azimuth gear boxes is repaired. The remaining gear box will also require maintenance soon. Wettzell has recently upgraded its control system, which accounts for much of its down time. Gilcreek has both mechanical and control system problems. This antenna is about 40 years old and the control system about 30 years old. These results also discount the continuing inability of Matera to observe because of a problem with the antenna rail. For the first part of the year, Matera was still included hopefully in the scheduling process, but after about two months no further attempts were made to schedule the antenna. If it had been, it would have contributed about another 60 lost days. This would have increased the overall days lost by the network to about 17.6% with about 55% being due to antenna problems alone. However, this is an unduly pessimistic estimate since other stations were substituted for Matera for almost two-thirds of the sessions they would have missed.

Stations with significant receiver problems include Fortaleza, Westford, HartRAO, Ny-Ålesund and Seshan. Most of these problems are in the category of reliability problems with the cyrogenics, power supplies or amplifiers, but for Seshan the most significant issue is the roll-off in the X-band bandpass, which has been included in this category.

The other most significant problem areas were Recorder, Miscellaneous, "Unknown" and RFI. Of these Recorder is by a small amount the most substantial. Most of items in this category are tape related. However, a small fraction represent problems with Mark 5 and other disk type recording systems. It is expected that Recorder issues will become less significant as more stations start using disk systems and the initial problems with these systems are resolved. However for now the results are no better than last year's recorder loss of 10.9%. The RFI category represents about a 5% loss. However, the actual number may be higher because some problems that were classified as either Receiver or "Unknown" due to lack of information might be more correctly attributed to RFI. It may be that the RFI losses should be as large of 10%, which is more in line with the previous year's result of 9.3%. The loss of TIGO due to the upper X-band IF problem represent about half of the overall network "Unknown" losses and above half of TIGO's losses as a station. No discussion of the items in the Miscellaneous category beyond the general description of these categories above is needed

A significant item that gets some attention in the correlator reports is that Gilcreek's Maser developed significant problems in the last third of the year. The problem has been difficult to troubleshoot and as of the writing of this report, was still not fully understood. The symptoms were that delay jumps of 100-300 picoseconds were being observed by the correlators and the data analysts were reporting significantly noisier than usual, by factors of several, clock variations. The noise level was bad enough that Gilcreek VLBI data were only marginally useful in the sessions where this problem occurred. In the process of trouble-shooting, it was noticed that the analysis results using Rubidium frequency standard were no worse and maybe slightly better than those using the Maser, so operations were switched the Rubidium in early 2005. This sort of problem is difficult to represent in the analysis of this report since it doesn't actually cause a loss of data, even due to a loss of coherence. However, we can to some extent work backward from the impact on the sigmas of the station's coordinates in geodetic solutions. This is difficult to accurately gauge because of differences in sessions, but it is plausible that this has roughly doubled the Gilcreek coordinate sigmas for the affected sessions. The equivalent observing time loss is about 75%. This would correspond to an additional loss of about 22 observing days for the 30 or so experiments where Gilcreek observed with this problem. This issue would probably be categorized as a "clock" problem although in the past this category has only been used for problems with the offsets. Using this analysis the overall percentage of lost days for the network would have increased to about 14.8%. The percentage of the loss due to clock problems would be about 15.8%. It should be noted however that this is a very unique problem. It has hardly ever, if not never, occurred before in the history of geodetic VLBI. We can have some hope that it is unlikely to recur.

2. Clock Offsets

As noted in the Network Coordinator's report for last year, it is important to develop consistent procedures for handling the clock offsets during the correlation process. Stations measure the offset between their formatter and the UTC time provided by GPS. The correlators typically apply a small, few μseconds or less, adjustment to the measured offsets in order to align the data to get fringes. If the adjustments are not applied in a consistent fashion by all correlators a corresponding error will be made in the UT1-UTC parameter adjustments. This will affect the quality of IVS UT1-UTC products at the level of the inconsistency in the adjustments applied for correlation. This could be corrected during the data analysis, but currently no analysis packages do this. It would require a significant amount of bookkeeping to add this feature now.

The Network Coordinator's report from two years ago recommended that the correlators develop a consistent table of adjustments to correct the local measurements of the formatter relative to GPS. This would remove a source of correlator-to-correlator and session-to-session variability in the UT1-UTC results. It was suggested that in developing this table the applied correction for Kokee should be artificially set to zero. Although not strictly correct, it is a simple approach and will maintain a level of consistency with old data, much of which was processed by WACO with an offset of zero for Kokee. However, the "true" adjustment will have to compensated for when an effort is made to align the ICRF and ITRF at this level. It was also recommended in the report from two years ago that a reference for the clock rate should be established at the same time. Although this is not as critical as the offset, it can easily be handled at the same time. Of course a good candidate station for the clock rate reference has to be found. As of the report two years ago, it seemed that Kokee's small rate relative to GPS, 1e-14 or better, would make it a good candidate. This would be a convenient choice for consistency's sake since again WACO has assumed that Kokee had a zero rate for much of the old data. This discussion is carried on in more

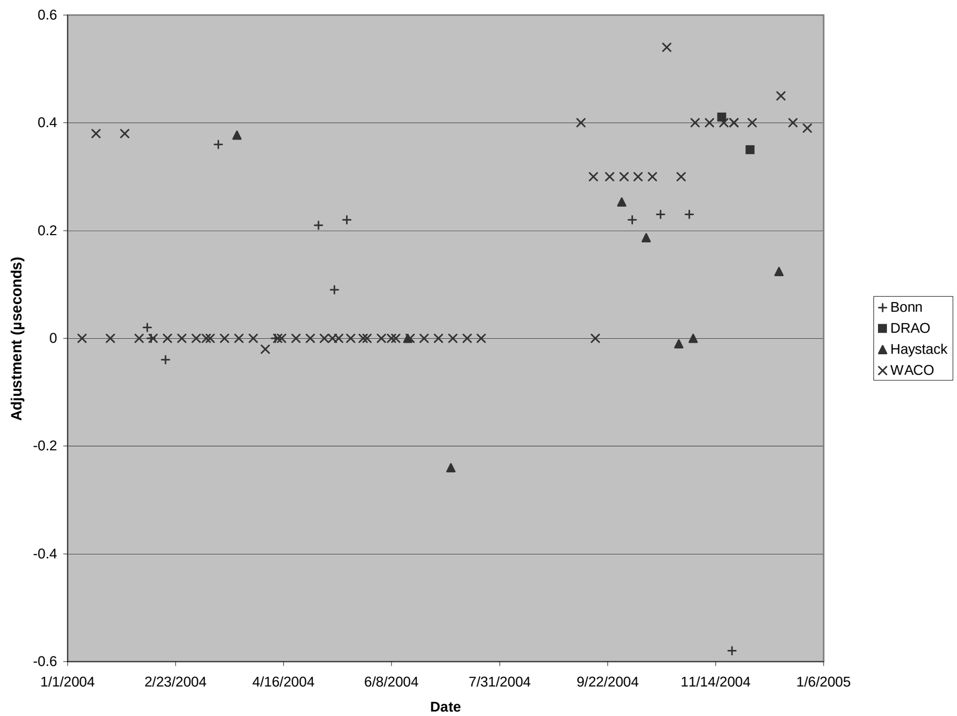

Figure 1. **Kokee UTC Correlator Adjustment**

detail in the report from two years ago.

Some progress has been made in implementing this recommendation, but this effort has not been completed. At the IVS Directing Board meeting in September 2003, the Correlator Representative to the Directing Board had offered to develop a consistent set of adjustments for the correlators to use. Currently this set of adjustments is under development [K. Kingham, USNO, private communication]. Consequently it would be premature to expect the offsets to be consistent for 2004. As a sample of the situation in 2004, the adjustments for Kokee are shown by correlator in Figure 1. In this figure it can be seen that the UTC offsets applied by WACO are small and in some cases zero. The non-zero values are probably related to use of Mark IV formatter which has an intrinsic offset of about 0.4 μseconds relative to the VLBA formatter that was also in use at that station. This is an example of equipment dependent offsets. WACO appears to have compensated appropriately for this case. Both Bonn and Haystack used variable but small non-zero adjustments. The variations however seem unrelated to the equipment in use. In the case of the large offset for Bonn in November this probably introduced a nearly one μsecond offset in the UT1-UTC for this experiment relative to other nearby sessions. The offsets applied by Penticton (DRAO) appear to be more consistent with the other correlators than had been the case in previous years. However, we have no information about the relative offset of the S2 equipment used for the sessions that Penticton correlates relative to the equipment used for other sessions. In general, the results for UT1-UTC will be affected at the level that the adjustments vary, in this case less than about 0.5 μsecond.

It is not only important that the UTC adjustments applied by the correlators are all consistent, but also the final clock value must be applied in the generation of the time-tags for the observations. It is known that this is done at the three Haystack Observatory developed correlators: Bonn, Haystack, WACO, and Penticton (DRAO). However it is not known what the GSI correlator does in this regard. A request has been made to K. Takashima to find out more about how the GSI correlator handles the offsets. (We also have no information about clock offsets are handled by the Miytaka correlator. However, the IVS geodetic sessions that Miytaka process are not primarily intended to measure UT1-UTC.) If the final clock value is not applied in the generation of the time tags, the session-to-session variation in the locally measured formatter to GPS will be included in the UT1-UTC parameter estimates.

Another area of concern is that different recording systems may require different adjustments. The Mark IV to VLBA formatter offset difference mentioned previously is an example of this. A more extreme example is from the CRF22 session from 2003. It was correlated at WACO using data from two different systems: K3 formatter recorded to tape and a K5 recorder using disks. WACO found a 2.3 μsecond clock offset between the two systems. This difference corresponds to almost a kilometer of cable difference. This does not seem realistic. If no cause for this is found, and probably even if it is, it will necessary to calibrate the differences between different systems. This might be undertaken by recording the same data with two or more systems possibly in heterogeneous networks and then comparing the final clock offsets that are needed to correlate them. Measuring and using these system specific offsets is a necessary step in implementing consistent correlator clock offset handling.