

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. A table of relative incidence of problems with various subsystems is presented. The current situation for the handling of correlator clock adjustments by the correlators is reviewed.

1. Network Performance

The network performance report is based on correlator reports for sessions in calendar year 2003. As of the date this report was generated, 188 sessions had been processed. There are another 24 sessions from the calendar year that had not been processed yet, or the correlator results were not available. Most of the missing sessions are from the latter part of the year and/or are waiting for tapes from Antarctica before they can be processed. Roughly 85%-90% of the scheduled station days 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 operations were interrupted or missed. However in other cases, it can be more difficult to calculate. To do this a non-observing time loss is typically converted into an equivalent lost observing time. 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. In a similar fashion, poor pointing is converted into an equivalent lost sensitivity. Poor playback is expressed as the fraction of total recorded bits lost.

From the correlator reports, an attempt is made to determine how much observing time was lost at each station. This is 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 is 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 are noisy.

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 previous years' versions of this report.

For the 188 sessions with results available from 2003, there are 1040 station days or almost 6 stations per session on average. Of these session days about 14.2% (or about 148 days) of the observing time was lost. For comparison, the results for 2002 were about 12.2%, for 2001 about 11.6%, and for a subset of 1999-2000 the results were about 11.8%.

The observing time lost for 2003 is somewhat worse than previous years which were more typically around 12%. If these observing time losses are converted into VLBI data yield losses, then 2003 had about 28% VLBI data loss and previous years about 24%. Whether these results reflect a significant decline in performance is not clear. There were several significant problems that contributed large data losses to 2003's results. Typically these were receivers that either failed or had their cryogenics fail or antenna failures. Without more information about the long term performance of the network, these problems might be considered anomalous.

In previous years an assessment of each station's performance was given in this report. That practice has been discontinued. 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 not least, there seemed to be some interest in attempting to "game" the system to improve the individual results. Consequently, this year only summary results are presented.

For the purposes of this report, the stations were divided into two categories: (A) those that were included in 14 or more network sessions, and (B) those in nine or fewer. Some of the stations in the former category had been included in as many as 100 sessions. The distinction between these two groups was made on the assumption that the results would be more meaningful for the stations participating in more sessions.

There are 15 stations in the 14 or more session category. Of these 10 successfully collected data for approximately 90% of their expected observing time. Four more stations collected 70% or more. One station in this group collected about 60%. The vast majority of the commonly used stations are in the 90% or more category.

There are 22 stations in the nine or fewer session category. This category included several stations that had previously not been included in IVS sessions, including several domestic Japanese stations. The range of successful observing time for stations in this category was 35%-100%. The median success rate was 87%. Overall the stations in this category observed successfully about 78% of the time.

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 Fortaleza, HartRAO, Ny-Ålesund, and Onsala. Given the high level of reliability of these stations it will be impossible for most of them to improve by this much again this 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 for the entire network is presented in Table 1.

The categories in Table 1 are rather broad and require some explanation. 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 "Antenna" category includes all antenna problems including mis-pointing, antenna control computer failures, and mechanical break-downs of the antenna. The "Unknown" category is a special category for

Table 1. Data Lost by Sub-system

Sub-System	Percentage lost
Receiver	25.2
Antenna	17.8
Unknown	12.6
Recorder	10.9
RFI	9.3
Shipping	6.1
Miscellaneous	6.0
Rack	5.0
Operations	3.6
Clock	3.4
Software	0.1
Total	100.0

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. 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 “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 be correlated with the rest of the session data. 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. However, by far the largest contributor to this category is power failures. In retrospect power failures should have had their own 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 “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. 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 operating system problems. This category could also include errors in files generated by the DRUDG program, but none of these were noted for 2003.

From the results it can be seen that receiver, antenna, and recorder account for more than 50% of the losses. In fact for 2003 there were several unusual receiver and antenna problems. If these are not repeated in 2004, the data yield should be better. Additionally, the data losses associated with the recorder should go down significantly as more stations switch to using disk drives for media. The disk systems are much more reliable than tape recorders.

2. Clock Offsets

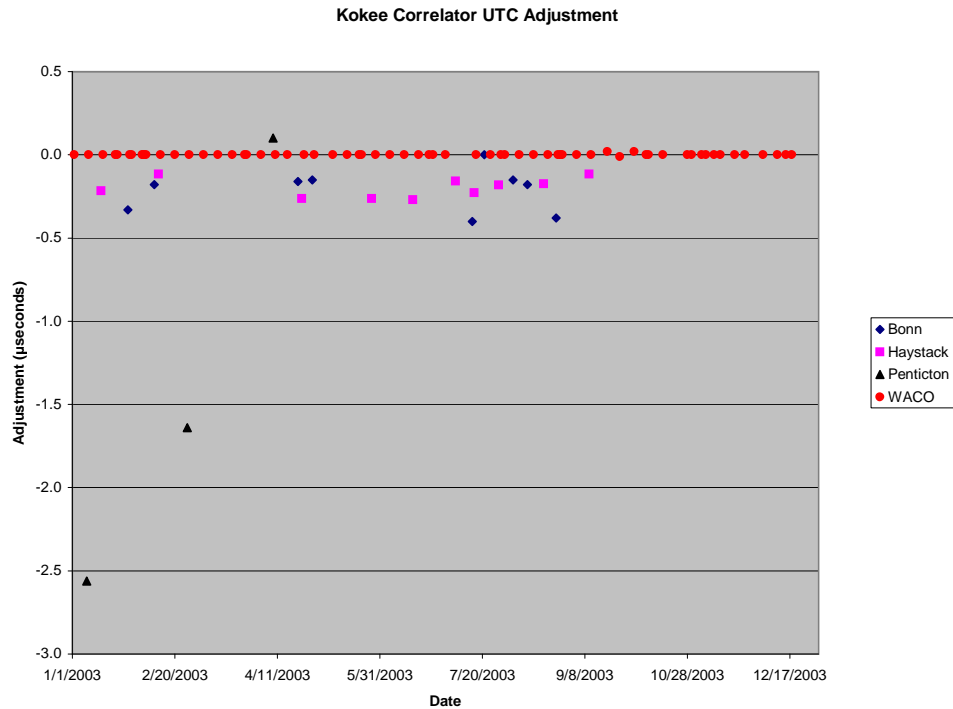


Figure 1. Kokee UTC Correlator Adjustment

As noted in the previous year's report, 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, a few μ seconds 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 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.

Last year's report 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 be compensated for when an effort is made to align the ICRF and ITRF at this level. It was also recommended in last year’s report 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 last year’s report, 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 detail in last year’s report.

At the IVS Directing Board meeting in September 2004, 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 2003. As a sample of the situation in 2003, the adjustments for Kokee are shown by correlator in Figure 1. In this figure it can be seen that the UTC adjustment applied by WACO is zero. Both Bonn and Haystack variable but small non-zero adjustments. The results for UT1-UTC will be affected at the level that the adjustments vary, less than $0.5 \mu\text{second}$. However the offsets applied by Penticton have considerably more variation. Since there were no sessions involving Kokee processed by the GSI correlator, we have no information about the adjustments for this station and correlator combination.

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, and WACO. However it is not known what the Penticton and GSI correlators do in this regard. Requests have been made to B. Petrachenko and K. Takashima, respectively, to find out more about how these correlators handle the offsets. (We also have no information about clock offsets are handled by the Miytaka correlator. However, the IVS geodetic sessions that Miytaka processes 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. A particular example is that the CRF22 session was correlated at WACO using data from two different systems: K3 formatter recorded to tape and a K5 recorder using disks. WACO found a $3.3 \mu\text{second}$ 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 be necessary to calibrate the differences between different systems. This might be undertaken by recording the same data with two or more systems and then comparing the final clock offsets that are needed to correlate them.