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 2005 was about 14.4%. 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 24.4%), receiver problems (24.2%), clock problems (14.5%), and recorder problems (8.9%). The closing of the Gilcreek station is described. The current situation for the handling of correlator clock adjustments by the correlators is reviewed. The handling of these adjustments directly impact the UT1-UTC estimates from VLBI data. A change in the handling of correlator clock adjustments at the Mark IV correlators has improved the situation markedly from previous years. However some clock offset work remains to be done for S2 recordings and the VLBA correlator.

1. Network Performance

The network performance report is based on correlator reports for experiments in calendar year 2005. This report includes the 154 non-RDV 24 hour experiments that had detailed correlator reports available as of February 3, 2006. There are six RDV experiments that have been 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 36 experiments from the calendar year that had not been processed or the correlator results were not available. Most of these missing experiments 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 experiments for 2005 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 can be converted into an equivalent lost sensitivity and then equivalent lost bits. Poor recordings can be 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 154 experiments from 2005 examined here, there are 1085 station days or slightly more than 7 stations per experiment on average. Of these experiment days about 14.4% (or about 157 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
2005	14.4

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

The lost observing time for 2005 was slightly greater than for 2004, but has returned to the level of loss seen in 2003. If these observing time losses are converted into VLBI data yield losses, then 2005 had about 29% VLBI data loss, 2004 about 25% and 2003 about 29%. It is not clear whether these variations in lost data reflect real changes in the performance level or simply variations due to inaccuracies in the analysis method. It does seem however that despite the approximations in the analysis method, the calculated observing time loss has been running fairly consistently at the 12-14% level for several years. It is notable however that some unusual circumstances occurred in 2005. In particular there were some unusual and significant losses associated with clock problems at Gilcreek and Westford in particular.

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 here.

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

There are 15 stations in the 20 or more experiment category. Of these, eight successfully collected data for approximately 90% of their expected observing time. Five more stations collected 80% or more. One station in this group collected slightly less than 75%. One station in this group collected only about 66% of its data. These statistics are almost unchanged from last year, except for the station with poorest performance. That station had unusually severe problems. In general, the pattern is that among the best performing stations, performance is improving.

There are 22 stations in the 13 or fewer experiment category. The range of successful observing time for stations in this category was 0%-100%. The median success rate was about 66%. Overall the stations in this category observed successfully about 41% of the 95 station days (9% of the total analyzed) that fall in this category. Its is notable that the performance for the stations that were in the fewest experiments became somewhat worse. This may be because they have less practice and less opportunity to work out problems in their systems.

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 Hobart, Ny-Ålesund, HartRAO, and Wettzell. Given the high level of reliability of these stations had already achieved, it was quite impressive that they were able to improve their performance this much. It will be impossible for 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	24.4
Receiver	24.2
Clock	14.5
Recorder	8.9
Miscellaneous	8.0
RFI	6.2
Rack	5.1
Operations	4.7
Unknown	3.3
Software	0.5
Shipping	0.2
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, design problems that impact the sensitivity, and loss of coherence that was due to LO problems. In addition, for lack of a better choice, loss of sensitivity due to upper X band Tsys problems were assigned to this category. 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 and it was not possible to determine a cause of the loss. 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 be correlated with the rest of the experiment 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. Maser and coherence problems that could be attributed to the maser were also included in this 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 almost 50% of the losses, down slightly from last year. This was in spite of the fact that some problems were shifted to the receiver category from “Unknown”. For 2005 some stations had significant antenna problems: Gilcreek, Svetloe, Noto, TIGO, and Matera in particular. Kokee’s antenna problems of two years ago seem to have been resolved. These results discounted the inability of Matera to observe during the first third the year because of a problem with the antenna rail, i.e., the results for Matera would have been worse if this had been accounted for. However, other stations were substituted for Matera in some cases. This issue appears to have been resolved unless more extensive work is needed on the rail in the future.

Stations with significant receiver problems include TIGO, Simeiz, Seshan, Gilcreek, and Fortaleza. Most of these problems are in the category of reliability problems with the cryogenics, 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. Also the loss at the upper end of X-band, for TIGO in particular, but Gilcreek and Fortaleza as well, were included in this category. This problem represents the most significant overall loss for the network, almost 25% of the total loss. It represents almost half of TIGO’s loss as a station.

The “Clock” category represented much more of the loss than in previous years. This was due primarily to significant maser problems at Gilcreek and Westford. However, loss of coherence at other stations (Kokee and Svetloe) contributed significantly as well, although it may be that for these stations these losses had other causes.

The other most significant problem areas were Recorder, Miscellaneous, RFI, Rack, and “Unknown". Of these Recorder is by a small amount the most substantial. Losses in this category are down slightly from last year. Most of the items in this category are tape related. However, a small fraction represent problems with Mark 5 and other disk type recording systems. The improvement in this category may finally be reflecting the "disk dividend" we have been hoping to get as tape use is curtailed. If so, we would expect to see the losses in this category to continue to decrease. This should be easily accomplished since more than one-third of 2005's Recorder losses were due to tape problems at Fortaleza, which is now using Mark 5 exclusively for recording. The RFI category represents about a 6% 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 result from two years ago of 9.3%.

A significant item that gets some attention in the correlator reports is that Gilcreek's Maser developed significant problems in the last third of 2004. This problem continued into 2005 and was not resolved until about a third of the way into the year.

The performance of the stations in CONT05 was excellent. The overall loss for these experiments was about 4%. This compares to an overall loss of about 16% for the other experiments of 2005.

2. Gilcreek Closure

NASA decided at the end of 2005 to cease operations at the Gilmore Creek station. This decision was forced by administrative issues, but was probably inevitable. The performance of the station had sagged considerably over recent years due largely to the aging of the equipment and the lack of continuity of personnel knowledgeable enough to maintain it. Although the station's performance was starting to recover at the end of the year due to new personnel working there, administrative problems forced the issue. Despite the long history and central role of this station in the geodetic VLBI network, the closing of the station could not be avoided. The positive note that comes out of this is that NASA reports that the money that was being used for the operation of this station will be reprogrammed to support development of the next generation, "VLBI210", system.

3. Clock Offsets

As noted in the Network Coordinator's reports for the last few years, 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 three years ago recommended that the correlators develop a consistent table of adjustments for correcting the local measurements of the formatter offsets relative to GPS. This would remove a source of correlator-to-correlator and experiment-

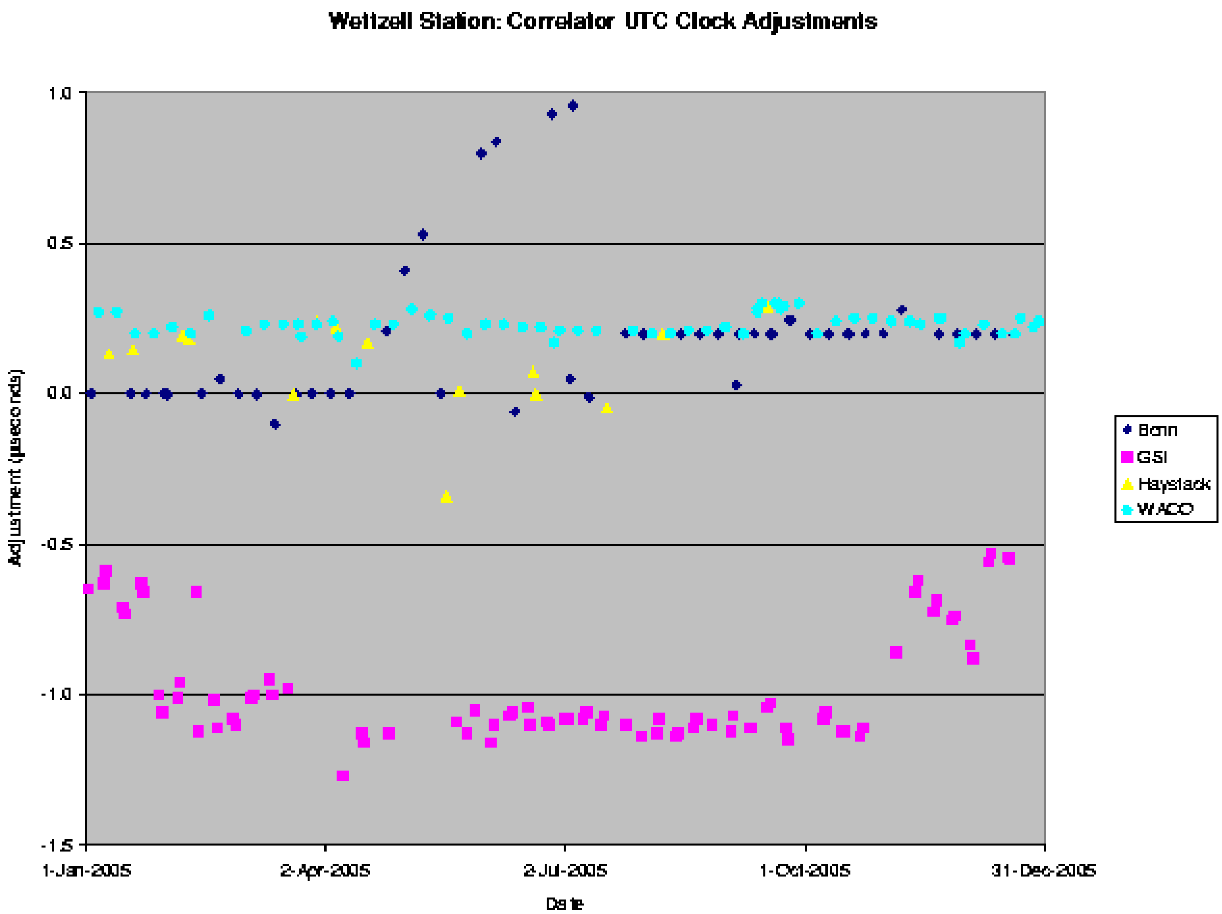

Figure 1. Wettzell Correlator UTC Clock Adjustments

to-experiment 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 (when Kokee uses the VLBA formatter,when they use a Mark IV, the correction should be increased to about 0.4 μseconds). 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 a correction 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 the report from two years ago that a reference for the clock rate should be established at the same time.

Significant progress was made in 2005 in implementing this recommendation. As of about August, the three Mark IV correlators started using consistent correction values [K. Kingham (USNO), B. Corey (Haystack Observatory), and A, Mueskens (University of Bonn), private communication]. An example of the improvement can be seen in Figure 1, which shows the clock corrections used for Wettzell. Until approximately August, the adjustments used by the three Mark IV geodesy correlators (Bonn, Haystack, and WACO) can be seen to vary by as much as a μsecond. After August, the variations are much smaller, on the order of 0.1 μseconds. This

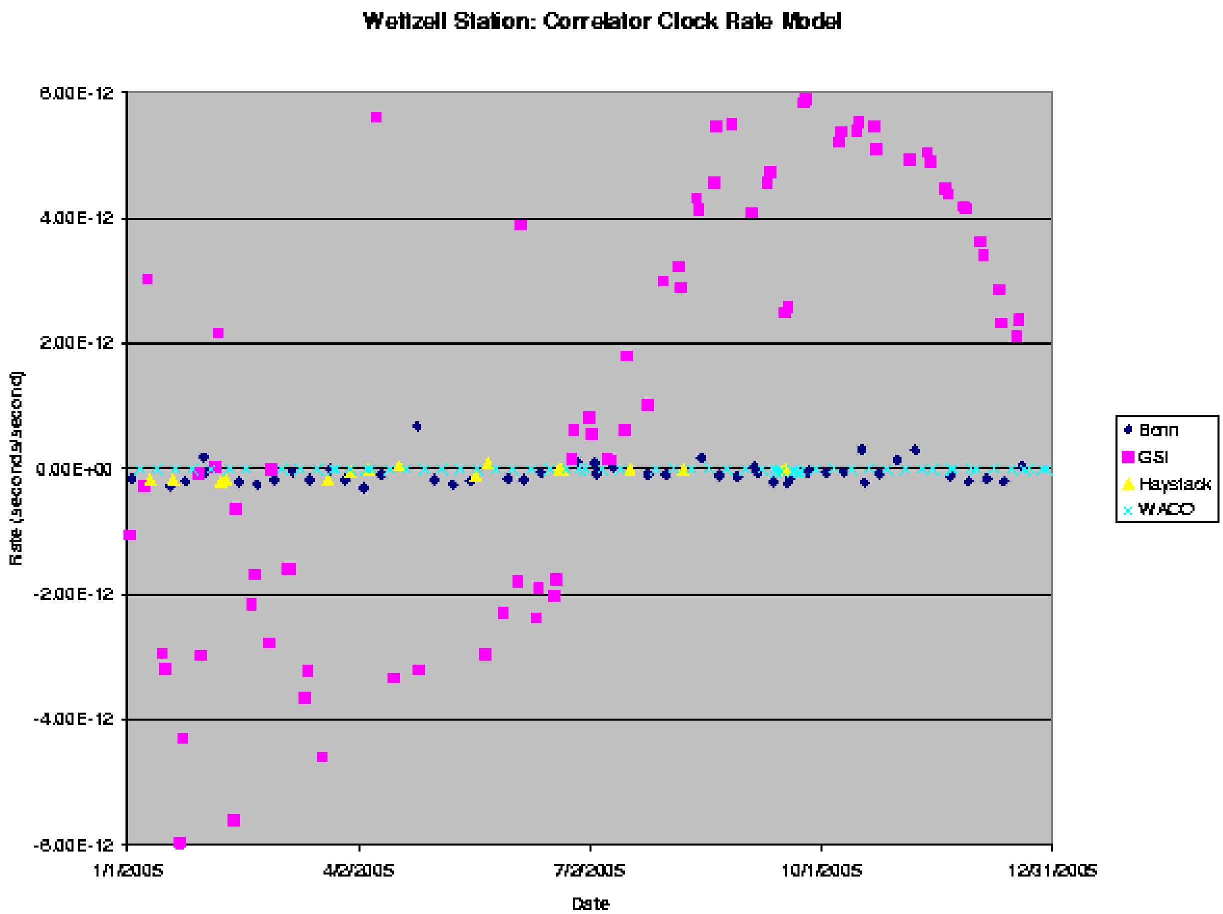

Figure 2. Wettzell Correlator Clock Rate Model

is a significant improvement and will lead to better accuracy of the UT1-UTC estimates. It is interesting to note that the offsets used by the GSI correlator are biased by about 1 μsecond compared to those used by the Mark IV correlators. If this is correct it would bias UT1-UTC estimates from GSI correlated experiments compared to Mark IV correlated experiments. As it happens, the GSI correlator is currently used only for "K" type Intensives and domestic Japanese 24 hour experiments. The UT1-UTC precision of these experiments is much less than that of the 24 hour international experiments, around 10-20 μseconds compared about 1-2 μseconds. Consequently, this bias does not significantly impact the results. However, it should be corrected. Discussions about how to handle this have begun [K. Takashima (GSI) and Y. Koyama (NICT), private communication]

A further issue is how stable the UT1 rate measurements are. This depends on the accuracy of the correlator models for the Maser rates for the observing stations and the accuracy of the Maser rates themselves. The desired accuracy of the IVS Working Group 2 report was for 0.3-0.5 μseconds/day. This translates into a clock rate of about 3.5e-12 to 5.8e-12. It would be desirable to have the correlator clock models consistent at a level 10% of that, 3e-13, or better.

A plot of the Correlator Clock Rates for Wettzell are shown in Figure 2. It is difficult to see on the scale of the plot, but the clock rates used by the Mark IV correlators vary at about the desired level, about 6e-13 (peak-to-peak, i.e., simplifying by saying the peak-to-peak is about twice the RMS). This is about good enough, but perhaps could be improved. The rates for Wettzell used by the GSI correlator vary by about 1.2e-11 (peak-to-peak). These exceeds the desired level by a factor of about 20. As with UT1-UTC estimates for GSI, this is probably not an issue given the lower precision of the experiments processed by this correlator. However, this should probably be improved.

A further issue is how close the rates of the underlying Maser are to UTC. The $de facto$ standard for the Maser rates is that Kokee's Maser is assumed to be "zero" by the correlators. Contrary to past Network Coordinator reports it is not known what the rate of the Kokee Maser relative to UTC is (there was a misunderstanding because there are two Masers at Kokee). The actual rate of the Kokee Maser and how to handle the rates in the future is being investigated. It may be a better solution for rates to use the actual measured rates at each station rather than working to the assumption that one station has zero rate.

Another area of concern is that different recording systems may require different adjustments. There is a difference between Mark IV and VLBA formatters of about 0.4 μseconds [K. Kingham (USNO), private communication]. This was accounted for when Kokee changed from using VLBA to Mark IV formatters. There is also a measured value between the K-5 and the Mark IV formatters [Y. Koyama, NICT, private communication]. However, this is not of much concern since Tsukuba and Kashima which use K-5 formatters have had their offsets determined relative to the Mark IV formatter stations by the Mark IV correlators already. There should be no problem as long as they do not change the formatters that they use, or if they do, it is done in a experiment with a population of Mark IV formatters. This is true in general for any station. A different issue is how stations that use S2 formatters will be aligned relative to stations that are already aligned. This question is being investigated [B. Petrachenko (NRCan) and A. Whitney (Haystack Observatory), private communication].

One remaining issue besides the K-5 (GSI) and S2 correlator data is aligning the clock offsets used by the VLBA correlator with those used by the Mark IV correlators. The VLBA is planning to do this [C. Walker (NRAO), private communication].