Network Coordinator Report

Ed Himwich

Abstract

This report includes an assessment of the network performance in terms lost observing time for the 2006 calendar year. Overall, the observing time loss was about 13.6%. A table of relative incidence of problems with various subsystems is presented. The most significant causes of loss were receiver problems (accounting for about 20.8%), antenna reliability (19.0%), miscellaneous problems (18.0%), rack problems (16.3%), and RFI (11.6%). The Canadian participation in IVS observing ended in 2006. There are prospects for Korea, India, and New Zealand to start contributing to IVS. New antennas are being purchased by Australia and New Zealand. The current situation for the handling of correlator clock adjustments by the correlators is reviewed. The situation is found generally to have greatly improved compared to previous years. The handling of these adjustments directly impact the UT1-UTC estimates from VLBI data.

1. Network Performance

The network performance report is based on correlator reports for experiments in calendar year 2006. This report includes the 145 24-hour experiments that had detailed correlator reports available as of February 23, 2007. Results for 42 experiments were omitted because either they were correlated at the VLBA, have not been correlated yet, or correlation reports were not available on the IVS data centers. Experiments processed at the VLBA correlator were omitted because the information provided is not as detailed as from Mark IV correlators. The experiments that have not been correlated yet include mostly RD, JD, and T2 experiments from the second half of the year, as well as some OHIG experiments. The experiments without reports on the IVS data centers include mostly E3 experiments processed by the Penticton correlator. In summary, roughly 80% of the scheduled experiments for 2006 are included in this report.

An important point to understand is that in this report the network performance 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 and why. 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 indicates 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. 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 average lost observing time per station is not equal to the overall average loss of VLBI data. Under some simplifying assumptions, the average loss of VLBI data is roughly about twice the average loss of observing time. This approximation is described in the Network Coordinator's section of the 2001 IVS Annual Report.

For the 145 experiments from 2005 examined here, there are 935 station days or about 6.4 stations per experiment on average. Of these experiment days about 13.6% (or about 127 days) of the observing time was lost. For comparison to earlier years, see Table 1.

Table 1. Lost Observing Time

Year	**Percentage**
1999-2000*	11.8
2001	11.6
2002	12.2
2003	14.4
2004	12.5
2005	14.4
2006	13.6

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

The lost observing time for 2006 was slightly less than for 2005. It is not clear whether the year-to-year variations in lost observing time 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 should be noted that the CONT05 experiments in 2005, where a special effort was made to achieve high reliability at some of the most reliable stations in the network, an observing time loss of only 4.0% was achieved. If these observing time losses are converted into VLBI data yield losses, then 2006 had about 27% VLBI data loss, 2005 about 29% VLBI data loss, 2004 about 25%, and 2003 about 29%.

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 16 or more network experiments among those analyzed here, and (B) those in 9 or fewer. Half of the stations in the former category had been included in more than 50 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 16 stations in the 20-or-more experiment category. Of these, six successfully collected data for approximately 90% of their expected observing time. Five more stations collected 80% or more. Four more stations collected 65% or more. One station in this group collected only slightly more than 50% of its schedule data. These statistics are somewhat worse than last year. They are balanced by the fact that the stations with fewer experiments were more successful in collecting data this year. The average data loss from both groups this year were about equal, rather than the station with more experiments being significantly more successful.

There are 20 stations in the 9-or-fewer experiment category. The range of lost observing time for stations in this category was 0%-35%. The median success rate was about 6.5%. Overall the stations in this category lost about 13% of the 80 station days (about 9% of the total analyzed) that fall in this category. It is notable that the performance for the stations that were in the fewest experiments was better than last year (41%).

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 TIGO, Zelenchukskaya, Svetloe, and Medicina.

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. This table includes four years of data sorted by descreasing loss in 2006.

Table 2. Percentage Data Lost by Sub-system

Sub-System	**2006**	**2005**	**2004**	**2003**
Receiver	20.8	24.2	18.0	25.2
Antenna	19.0	24.4	32.9	17.8
Miscellaneous	18.0	8.0	8.0	6.0
Rack	16.3	5.1	6.8	5.0
RFI	11.6	6.2	5.0	9.3
Clock	4.9	14.5	0.5	3.4
Unknown	4.0	3.3	10.1	12.6
Recorder	3.3	8.9	11.1	10.9
Operations	2.0	4.7	6.1	3.6
Software	0.1	0.5	0.1	0.1
Shipping	0.0	0.2	1.4	6.1

The categories in Table 2 are rather broad and require some explanation, which is given below.

Antenna This category includes all antenna problems including mis-pointing, antenna control computer failures, and mechanical break-downs of the antenna.

Clock This 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. Phase instabilities reported for Kokee were included in this category.

Miscellaneous This category includes several small problems that do not fit into other categories, mostly problems beyond the control of the stations such as power, weather, and errors in the observing schedule provided by the Operation Centers. As of the 2006 report, this category also includes errors due to tape operations at the stations that were forced to use tape because either they didn't have a disk recording system or did not have enough media. This category is dominated by power and weather issues.

Operations This category includes all operational errors, such as DRUDG-ing the wrong schedule, starting late because of shift problems, operator (as opposed to equipment) problems changing recording media and others.

Rack This category includes all failures that could be attributed to the rack (DAS) including the formatter and BBCs.

Receiver This 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 clearly more accurate choice, loss of sensitivity due to upper X band Tsys and roll-off problems were assigned to this category.

Recorder This category includes problems associated with disk recording systems and network transfer of data. As of the 2006 report, no problems associated with tape operations are included in this category.

RFI This category includes all losses directly attributable to interference including all cases or amplitude variations in individual channels, particularly at S-band.

Shipping This category includes data that could not be correlated because the media were either lost in shipping or held up in customs long enough that it could not be correlated with the rest of the experiment data. There were no such instances this year.

Software This 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.

Unknown This category is a special category for cases where the correlator did not state the cause of the loss and it was not possible to determine the cause with a reasonable amount of effort.

From the results it can be seen that antenna and receiver together account for almost 40% of the losses. This is down slightly from previous years. For 2006 the stations with significant antenna problems include Zelenchukskaya, Svetloe, Algonquin, and Westford. In particular, the antenna problems at Algonquin were so severe that the need for expensive repairs appears to have helped precipitate the decision by the Canadian government to terminate their country's VLBI operations.

Stations with significant receiver problems include Matera, Zelenchukskaya, Ny-Ålesund, Fortaleza, and Medicina. Most of these problems are in the category of reliability problems with the cryogenics, power supplies, and amplifiers. Losses due to receiver cryogenic problems were significantly reduced this year compared to previous years. For some stations, Fortaleza in particular, the most significant issue is the roll-off in the X-band bandpass, which has been included in this category.

The "Miscellaneous" category is significantly larger than for previous years. This is primarily due to power outages, primarily at Algonquin, weather problems, at HartRAO and Hobart primarily, and a security issue at Tsukuba. There was also a bureaucratic problem that prevented Matera from observing for the first few weeks of the year.

The "Rack" category was larger this year primarily due to the fact that Seshan observed with only eight BBCs. However, there were some Mark IV formatter failures, as well as a lack of availability of the correct filters for R1s at Fortaleza.

The "RFI" category is larger this year than last, possibly due to more uniformity in attributing amplitude fluctuations to this cause as well as the fact that Matera, the station with the most severe RFI problems, was able to observe more this year because their antenna was working more reliably.

The "Clock" category represents less loss than last year. This is primarily because the long term failures of the Masers that occurred in 2005 were not repeated. However, loss of coherence at other stations (Kokee in particular) were included in this category and contributed to the increase over 2003 and 2004.

The "Recorder" category is down significantly from previous years primarily because almost all recording is done on disk and the problems with the few tape recordings that were made are reported now in the "Miscellaneous" category. The decrease in data loss due to recorder operations of about 11% to about 3% probably represents the "disk dividend" we have been hoping to get as tape use is curtailed.

2. Closure of Canadian Stations

The Canadian government stopped all VLBI operations. This was apparently a cost-cutting move. This appears to have been precipitated by the need for a costly repair to the antenna at Algonquin.

3. New Stations

The station at Badary, Russia should start observing in 2007. There are prospects for new stations on several fronts. Both Australia and New Zealand are in the process of obtaining new antennas, two and one, respectively. Korea is planning to build one antenna primarily for geodesy. There is also interest in using the Korean VLBI Network (KVN), which will consist of three stations intended primarily for astronomy, for geodesy. There is interest in India in building a network of four telescopes that would be useful for geodesy. Many of these antennas may become available for use in the next few years. Efforts are being made so that these antennas will be compatible with VLBI2010.

4. 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. This could be corrected during the data analysis, but currently no analysis package does this. It would require a significant amount of bookkeeping to add this feature now.

The Network Coordinator's report for 2002 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-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 2002 that a reference for the clock rate should be established at the same time.

Significant progress was made in 2005 in implementing these recommendations. The calendar year 2006 is the first one in which this approach was applied to all experiments. The offset adjustment for Wettzell in 2006 is shown in Figure 1. The scatter of these measurements is greatly improved for Mark IV correlators (Bonn, Haystack, and WACO). Generally the variation is at the 0.1 μsecond level, but there are few larger outliers for the Bonn correlator. It is not known if these outliers are due to typographic errors or if the adjustments for these experiments were larger than expected. They should be investigated. In any event, the results are greatly improved compared to previous years where the adjustments were scattered over a range at least 10 times as large.

It is interesting to note that the adjustments used by the GSI (K-5) 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 for the "K" type Intensives compared to a few μseconds for R1 and R4 experiments. Consequently, this bias does not significantly impact the results. It has recently been realized that the K-5 correlator does not include the effect of the clock offset and rate in its calculation of the time-tags and delays that it produces [Y. Koyama (NICT), personal communication]. Thus the quality of the correlator clock model is not currently impacting the UT1 results for the K-5 correlator. This will be corrected in the future.

A further issue is how stable the UT1 rate measurements are. This depends on the accuracy of the correlator models for the Maser (and associated station electronics) rates for the observing stations. The desired accuracy of the IVS Working Group 2 report was for 0.3-0.5 μseconds/day.

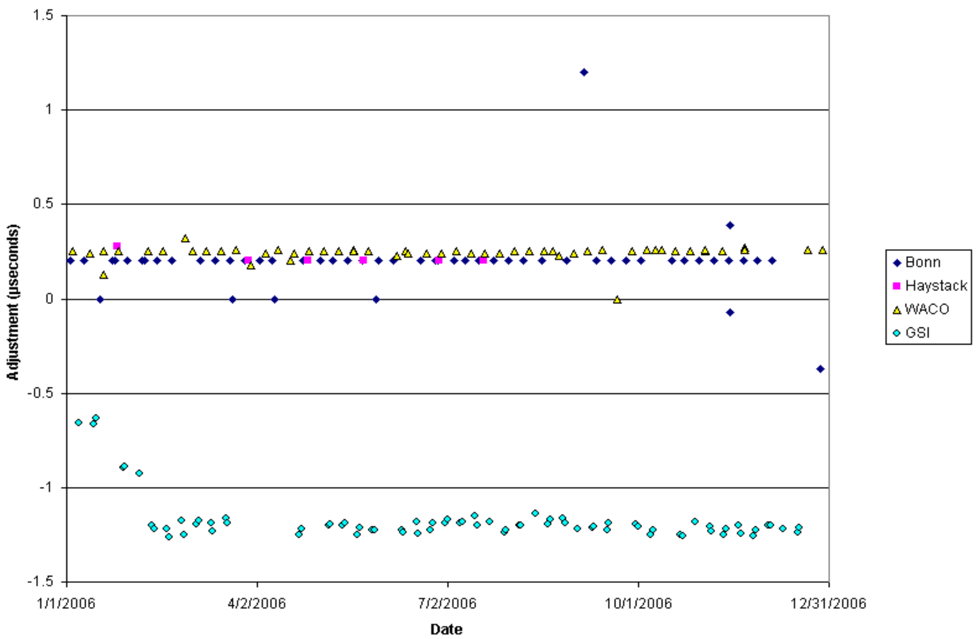

Figure 1. Wettzell Correlator UTC Clock Adjustments

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. It turns out that it is possible to determine the station rates at about the level required from the formatter offset values recorded during 24-hour experiments.

A plot of the Correlator Clock Rate Model for Wettzell for 2006 using this approach is 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 (Bonn, Haystack, and WACO) vary at about the desired level, about 2.5e-13 (5e-13 peak-to-peak, i.e., simplifying by saying the peak-to-peak is about twice the RMS). This is about good enough and is an improvement of a factor two over last year. The rates for Wettzell used by the GSI correlator vary by about 1.4e-11 (peak-to-peak). These exceed the desired level by a factor of about 30. 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. The reason for the scatter in the rates used by the GSI correlator is that they are determined from the clock offsets measured during the short span of data collected in the K Intensives [Y. Koyama (NICT), personal communication]. This explains the larger range of the values. The origin of the systematic nature of the rates is not clear, but presumably is related to some systematic effect that is being sampled by the clock offset measurements. As mentioned previously the correlator clock model is not being taken into account in the results of the K-5 correlator. But this is planned for the future.

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. A value of 0.26 μseconds was measured between the K-5 and the Mark IV formatters, K5 later than Mark IV [Y. Koyama *et al.*, Timing Offset of the K5/VSSP System,

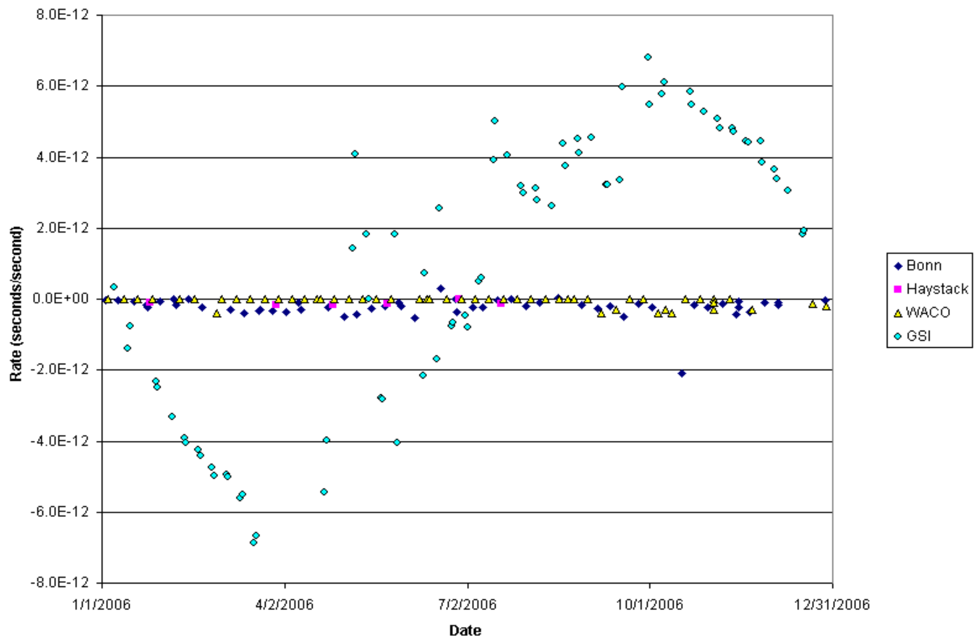

Figure 2. Wettzell Correlator Clock Rate Model

IVS NICT TDC News No. 26, p. 6-8]. This should be compared to the value that has been empirically determined from processing of K-5 data at Mark IV correlators [K. Kingham (USNO), private communication]. Unfortunately, while in the past there was a need to consider how S2 recorders (and correlator) would be integrated into this scheme, due to the cessation of Canadian VLBI operations, this is no longer an issue.

It is also important to consider whether there are offsets between different recording rates within a given recording system. It was recently discovered that there appears to be a 8 μsecond offset for the Mark IV formatter when the tape data rate is 18 MHz as compared to other lower data rates [D. Smythe, Timing Offset of the Mark IV Formatter, Haystack Observatory Mark 5 Memo #047]. This offset essentially affects only 1024 Mb/s recordings. Otherwise the Mark IV formatter is known to not have an offset between different data rates [D. Smythe (Haystack Observatory), private communication].

The difference between correlators must also be considered. The VLBA correlator has moved to the same relative offset used by the Mark IV correlators, i.e. Kokee with a VLBA formatter is 0 μseconds [C. Walker (NRAO), private communication]. However, we don't know if the VLBA correlator has an offset relative to the Mark IV correlators. We must also consider that an offset of the K5 correlator relative to the Mark IV. This will have to be investigated before experiments that yield more precise estimates of UT1, typically 24-hour with long baselines, are processed with the K5 correlator.