VLBA TEST MEMO 69

Inappropriate EOP on the VLBA Correlator

R. C. Walker, V. Dhawan, W. Brisken, J. Benson, L. Kogan, & J. Romney
National Radio Astronomy Observatory
Socorro, NM

October 6, 2005

Abstract

A bug has been found in the job generator for the VLBA correlator that caused it to use predicted rather than measured Earth Orientation Parameters (EOP) between May 2003 and August 2005. The errors introduced often are sufficiently large to adversely affect phase referencing observations and other observations that rely on an accurate correlator model. The nature of the problem and its impact are discussed in this memo. Correcting data is straightforward using the new EOPS option in the AIPS task CLCOR. Users just wishing to make corrections without worrying about the gory details can consult Section 2, or simply follow the CLCOR help file. It has become apparent while dealing with this problem that even occasional projects from outside the time period of the bug have been correlated with poor EOP. Usually this is because their job scripts were prepared too close to the observe date. So any user depending on an accurate correlator model may wish to check their EOP.

1 Introduction

A significant issue has been found with the Earth Orientation Parameters (EOP: UT1–UTC and pole position) used in VLBA correlator job scripts since May 2003. The correlator is supposed to use rapid service values of EOP, that are based on measurements, when they are available at the time jobs are generated. Instead, because of a software bug, it was using predictions from about 2 weeks before the observe date. The Earth's rotation rate and pole position are sufficiently unpredictable that this can make a difference of tens of milliarcseconds in the observation geometry.

The problem adversely affects projects that depend on an accurate correlator model. Such projects include those that use phase referencing, especially at high frequencies or with large separations between calibrator and target or with very stringent position determination requirements. Projects that use sources scattered around the sky to solve for the atmosphere (AIPS task DELZN (Mioduszewski & Kogan 2004)) can have the atmospheric solutions degraded. Also spectral line observations that use a line feature for phase calibration but depend on a calibrator to take out phase slopes (delay) will be impacted at a low, but perhaps significant, level. Projects that are not affected include those that depend primarily on self-calibration, such as imaging of strong sources. Also unaffected are geodetic and astrometric projects that use the total delays from the correlator.

While investigating this bug, it was realized that some projects get processed with predicted rather than measured EOP even when the system is working normally. This usually happens when correlator jobs are prepared before the rapid service EOP values are available in Bulletin A. For this reason, users with projects outside the time period in which the bug was in effect might want to check the EOP used for their project and, in affected cases, make corrections.

This memo starts with a quick recipe on how to correct data. Users who are not interested in the details, but who want to make corrections can read that section only, or even just the CLCOR help file in AIPS. Later sections describe the problem in detail and present test results that show the possible effects and show that the AIPS correction works.

2 Correcting Data - The Quick Recipe

The steps required for correction of EOP data using the new EOPS option in CLCOR in AIPS are listed here with minimal explanation. If confused about something, try reading the following sections. Note that the EOP errors are almost never large enough to cause significant loss of coherence. The fringe rate and delay offsets are very small compared to any windows that are used for VLBA correlation. To make corrections, CLCOR in a version of AIPS updated since 2005 September 29 is required.

To correct for EOP errors, CLCOR obtains the EOP used for correlation from the AIPS CT table, gets new, presumably better, EOP values from an external file provided by the user, calculates total delays using both sets of EOP and using a simple model that is otherwise the same for both calculations, takes the difference, and derives CL table corrections from that difference. The EOP values are provided for UT midnight each day and are interpolated to the time of each CL table entry. CLCOR must accurately reproduce the correlator's interpolation algorithm which it does for the VLBA. The CLCOR EOP correction has not been designed with other correlators in mind nor has it been tested with data from other correlators.

The steps to correct data are:

- Download *http://gemini.gsfc.nasa.gov/solve_save/usno_finals.erp* to get good EOP values adjusted slightly from the IERS values to work best with the source and station catalogs used in VLBI.
- Choose the UV file to correct. Most multi-source files from FITLD will be OK. In order of preference:
  - A UV data file loaded with a version of FITLD more recent than mid-September 2005. It will have a new style CT table that includes a column with the time range over which that row of EOP data was used in correlation.
  - Any file whose CT table consists of multiple sets of identical rows. This will be fine, even if loaded with the older FITLD. CLCOR will use the first 5 rows, which is what the correlator did. The CT table can be examined with PRTAB. Note that the boundary between row sets is distinguished by a drop in the day of the EOP data.
  - A UV file loaded with the older FITLD and having CT table row sets that vary, usually by gaining or losing a day. For this case, CLCOR needs to relate CT rows to time ranges. It does this by reading the history records written by FITLD as it loaded the distribution file from each job.
  - A UV file with an old style CT table, that does not have times associated with each row, and that has been through DBCON, VBGLU, UVCOP or other such tasks that combine, or pare down, data sets, should probably be avoided on the grounds that the attempt to assign times to CT rows might not work properly.
- Choose the CL table to correct. It is best to select one that does not contain the results of any self-calibration-like steps. These include pulse cal corrections (PCCOR), fringe fitting (FRING), and actual self-calibration (CALIB). It can have correlator corrections (ACCOR), gain and Tsys calibration (APCAL), parallactic angle correction and other geometry changes (CLCOR), and ionospheric calibration (TECOR) already applied since they are not based on the data phases. Bandpass correction is also best done after the EOP corrections. If a CL table that contains any self-calibration results is used, that type of processing should be repeated after the EOP correction. For a fresh data set, the EOP corrections are best made any time before the pulse cal correction (PCCOR) or first fringe fit.
- Run CLCOR to correct for the EOP errors. Use OPCODE='EOPS' and the GAINVER (CL table version) chosen above. Point INFILE to the *usno_finals.erp* file that was downloaded. Be careful to process all frequency ID's — each has to be done in a separate pass of CLCOR, setting SELBAND and SELFREQ as required (it apparently is not possible to do all in one pass). An alternative is to separate the different FQID's using UVCOP before running CLCOR. Note that, if multiple runs of CLCOR with data selection are required, using GAINVER=GAINUSE after the first pass can incrementally modify the output file so that you have one final output CL table that applies to all data.
- Then redo all of the self-calibration type calibrations. Final phases for self-calibrated (including fringe fitted) sources should be the same as without CLCOR so images, bandpasses etc based on such sources do not absolutely need to be remade. The changes will be in the phases of phase referenced sources.

A script called VLBAEOPS is being added to AIPS. It optionally downloads *usno_finals.erp* and then runs CLCOR. It is one of the scripts loaded with RUN VLBAUTILS. It can use an input file already on a local disk.

Note that one way to check how much of a problem the data set has is to run CLCOR using CL version 1 as input and checking the delay and phase gains generated to see if they are large enough to be a concern. Remember that for phase referencing, the real concern is the double difference between stations and between target and calibrator.

To determine the quality of the EOP used on any particular project, see references in Section 4, especially the tables by date and project available from anonymous ftp at ftp.aoc.nrao.edu in directory pub/staff/cwalker/VLBA.

The rest of the memo gives a lot of detail for those interested.

3 What Went Wrong?

For normal VLBA correlator operations, the weekly Bulletin A published by the International Earth Rotation and Reference Systems Service (IERS) and available from the US Naval Observatory (USNO) is read by an automatic script that inserts EOP values in the database from which the job generator gathers its required information. The numbers put in the database include predicts and one week's worth of rapid service values. Once a month, the

Bulletin A also contains "final" values for times at least a month in the past. Those are also put in the data base, but are usually too late for use on the correlator. Ultimately the database includes multiple values of varying quality for each day.

The job script generator is supposed to take the most recently inserted and presumably most accurate value. Instead, starting on May 5, 2003, at the time of a change in the underlying commercial database system, a bug was introduced such that the first value inserted was used. Prior to that change, correlator jobs almost always were based on the rapid service values. Afterwards, they were mostly based on predicts which are of much lower accuracy. Until June 2004, the values for 5 days a week were 4 to 9 day predicts while the other 2 days were 4 and 5 day old rapid service values. After June 2004, the values used were 10 to 16 day predicts, accounting for a significant degradation of the values in mid 2004. The bug was fixed on August 8, 2005 so job scripts prepared after that date should have used rapid service values unless they were prepared too close to the date of observation, before the rapid service values were available (as has happened). All job scripts in the queue at that time were remade. Unfortunately, thanks to a glitch in the automatic reading of Bulletin A, some projects correlated in late September 2005 have poor EOP and will need correction if they use phase referencing.

The correlator uses EOP values from zero hours UT on each of 5 days. Those 5 values are used for a spline fit that is then used to determine the actual value to use for each delay calculation, typically at 2 minute intervals. Thus there is no one value used for each project. Any EOP correction scheme needs to be aware of this and, unless recalculating the total model, be able to duplicate the correlator's interpolation. In fact, the same algorithm should be used so that it will deal with events like steps in the data in the same way as the correlator. For example, after mid 2004, once a week there is a step when jumping from a 16 day predict to a noticably better 10 day predict. Prior to mid 2004, some rapid service values are mixed with the predicts, which also introduced steps.

While investigating the EOP problem, it became apparent that a few observations from before May 2005 and also since the fix in August 2005 also used poor EOP. The reason for this is the jobs were prepared before the rapid service values for the required days were put into the data base. Between the once-a-week updates of Bulletin A and the need for 3 days of EOP beyond the observe date, jobs prepared up to 9 days after observing can end up using predicted EOP values. A suggested change to avoid this is presented in Section 10.

4 Effect on Data

The effect of the EOP errors on an individual observation depends on the quality of the predicts used and on the geometry of the observation. The errors are dominated by the error in UT1–UTC which has an effect equivalent to an error in the RA of the source. The magnitude of the equivalent position error (UT1–UTC is measured in milliseconds of time which has to be multiplied by 15 to get milliarcseconds) is typically about 10 mas before mid 2004 and 15 mas (with more observations that are significantly worse) after that. The worst cases can be around 4 times these numbers. The X and Y pole position errors are also equivalent to a source position shift at any one instant, but the direction of that equivalent shift rotates through the day. The magnitude is typically around 20% of the UT effect. For details on exactly how the EOP values affect calculated delays, consult the references given at the end of this document.

When phase referencing, the self-calibration done on the calibrator will remove any effect of the EOP error at the position of the calibrator. The phase referencing will then apply that same phase correction to the target source. But the phase correction actually needed is a sine function of hour angle and declination so the correction from the calibrator is not quite right for the target source. The error is, to first order, the total magnitude of the effect times the source-calibrator offset in radians. The exact effect is more complicated and won't be described here.

The model errors resulting from poor EOP appear in the form of delay errors and corresponding phase offsets and slopes. These result in the position errors discussed above. The delay errors are typically a few ns and introduce phase slopes that vary between sources. Spectral line observations that depend on a calibrator to set the delay, and hence the phase relationships between channels, will be affected by such a phase slope. For an 8 MHz bandwidth, a 4 ns delay offset (in the range seen) corresponds to a phase slope of 11.5 degrees across the band. For wider bandwidths, including using multiple basebands, the effect will be larger. For some astrometric observations, such a slope may be significant.

To help users check the EOP errors for their observations, tables that contain the EOP values used and the offset from the best values known as of 2005 Sept. 7 have been placed at ftp://ftp.aoc.nrao.edu/pub/staff/cwalker/VLBA. There are separate files for each year between 1998 and 2005. There are other tables of the values used on the correlator, good values from a *usno_finals.erp* file, and differences between these two, plus plots of the differences, at http://www.aoc.nrao.edu/%7Evdhawan/eop.html. A user may also determine the EOP values used in an observation by looking at the CT table attached to the AIPS data set or by looking at the UTC Table near the end of the job scripts in the NRAO archive at http://www.vlba.nrao.edu/astro/VOBS/astronomy/. Good values for the EOP parameters may be found by following links from http://www.iers.org/iers/products/eop/. Perhaps the best values to use for VLBI are in the previously mention file http://gemini.gsfc.nasa.gov/solve_save/usno_finals.erp. That file is updated daily. It contains the latest information from the IERS site but adjusted by a small offset and drift to be a better match to the source and station catalogs used at the VLBI correlators.

One simple step that users can take to check their data, if they scheduled according to recommendations, is to see how well the phase referencing worked on their phase check source. Recall that it is recommended that a second calibrator, near the primary calibrator, be observed occasionally. That calibrator is used to check the quality of the phase referencing by checking how close a position derived from phase referencing comes to the known astrometric position. Information about the quality of the phase referencing conditions also comes from checking the dynamic range of the phase referenced image of the secondary calibrator. This is a good check, but not perfect because the impact of the EOP error depends on the orientation of the offset between calibrator and target which is likely to be different for the calibration check source and the object being studied.

It is not difficult to fix the effects of the use of incorrect EOP values. The effects on data phase are large enough to damage phase referencing, but they are nowhere near large enough to cause loss or degradation of fringes. At the time the problem was found, there was no available method in AIPS for correcting the data. Since then, an option has been added to CLCOR to allow the corrections to be made. Details are discussed in Section 7 and in outline form above in Section 2.

For an overview of the effect on VLBA projects, Figure 1 shows a plot of the offsets between the EOP values used and the values in a *usno_finals.erp* file as a function of date. The vertical bars are the year boundaries. Prior to early 2003, most observations had good values. The exceptions are probably jobs prepared sufficiently close to the observe date that rapid service values had not yet been given for all days with EOP included in the job scripts. Around May 2003, a much larger portion of observations had significant offsets. There are still some good values because the correlator still used rapid service values for 2 days a week. After June 2004, when longer range predicts started being loaded to the data base and used, the errors got significantly worse. Figure 2 shows the same data as Figure 1, but in histogram form. There are 3 histograms for the different time periods.

Since the effect on any given data set depends on details of the observation setup and geometry and on the EOP errors encountered, probably the easiest way to test the effect is to run CLCOR as described in other sections and examine the magnitude of the changes made.

5 Tests Using Science Projects

The EOP problem was first noticed as a big offset between the EOP used on the correlator and IERS final values by Mark Reid of SAO while analyzing his science data. One of his observations, processed with poor EOP shortly before the problem was found, was reprocessed with rapid service values afterward. In this case, the EOP errors were $\sim$0.7 msec, the delay corrections made by CLCOR were $\sim$0.5 ns, and the differences between recorrelated and corrected data were $\sim$2 ps.

In addition, we made observations to check the correction, and for the rest of this memo, we concentrate on the targeted test rather than Reid's observation.

Other observations from the archive were selected and processed through CLCOR to test CLCOR's ability to determine correctly the time intervals over which a each set of EOP values was used in cases when there are multiple sets and the new type of CT table was not used. The tests were successful, but not reported here in detail.

6 Targeted Test Observations of 2005 August 13

The test project (to002) was observed on 2005 August 13. It was designed to demonstrate the effects of EOP errors and to verify that we are able to correct them. One pair and one triplet of calibrators with well known positions were observed, with about 40 minutes on the pair alternating with about 62 minutes on the triplet. Three times during the observation, about 45 minutes were spent looking at calibrators all over the sky. Unfortunately one of the triplet was too week to be useful for purposes of the test. The pair was 1215+303 and 1213+350, which are separated by 4.7 degrees. The surviving pair from the triplet was 1111+149 and 1116+128, which are separated by 2.4 degrees. For the processing shown here, 1213+350 and 1116+128 were treated as phase calibrators. Less effort was put into high quality phase referencing than would be likely for a science observation. Specifically, there was no attempt to image the calibrators or to do careful editing. The weather was not great for these observations which degraded the phase referencing for the later parts of the observation. The examples shown here are from early in the observation in order to emphasize EOP effects rather than weather effects. Also, we have not yet tried an atmospheric calibration based on the segments with many calibrators. The observations were made in dual S/X band mode (13 and 4 cm bands) with a geodetic style broad frequency spread.

The test observation was correlated twice. One pass was done with the rapid service EOP that we would normally use when the EOP bug was not active. The other pass used hand altered EOP parameters that differed from the good values by about 3 ms in UT1–UTC, 2.8 mas in X and 6.4 mas in Y for the times nearest the observations. The 5 day sequence of altered EOP were contrived to have a jump in the offset from the best values in an attempt to trigger an interpolation problem if there was one. The values used are tabulated in Table 1.

An impression of the magnitude of the phase offsets introduced by the poor EOP can be seen in Figure 3. The two test data sets only differed in the EOP used in processing, so calibration tables used for one can also be used

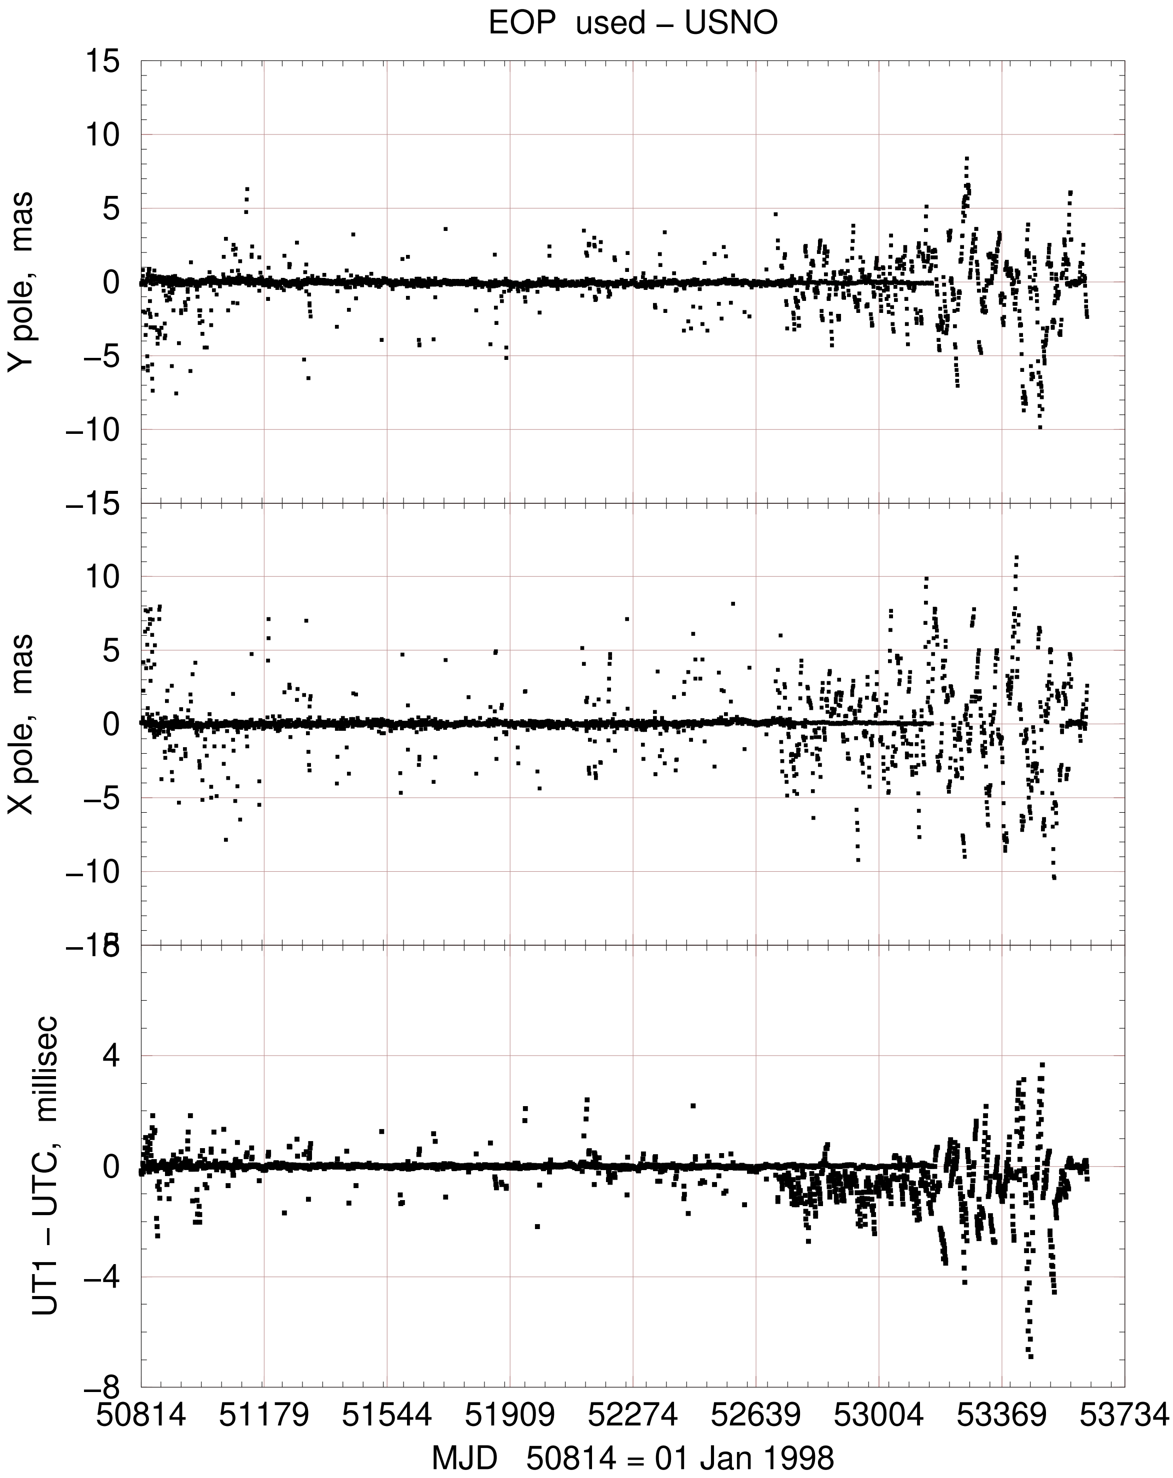

Figure 1: Plot of the offsets between the EOP parameters used in correlator jobs and reference good EOP values from a *usno_finals.erp* file. Data are shown for times since 1998. Vertical bars mark the year boundaries with the last partial year being 2005. The correlator jobs began to include poor EOP in May 2003 and were corrected in August 2005 just before the last data shown here. Longer range predicted EOP values started being used in June 2004, accounting for the worsening of the offsets at that time.

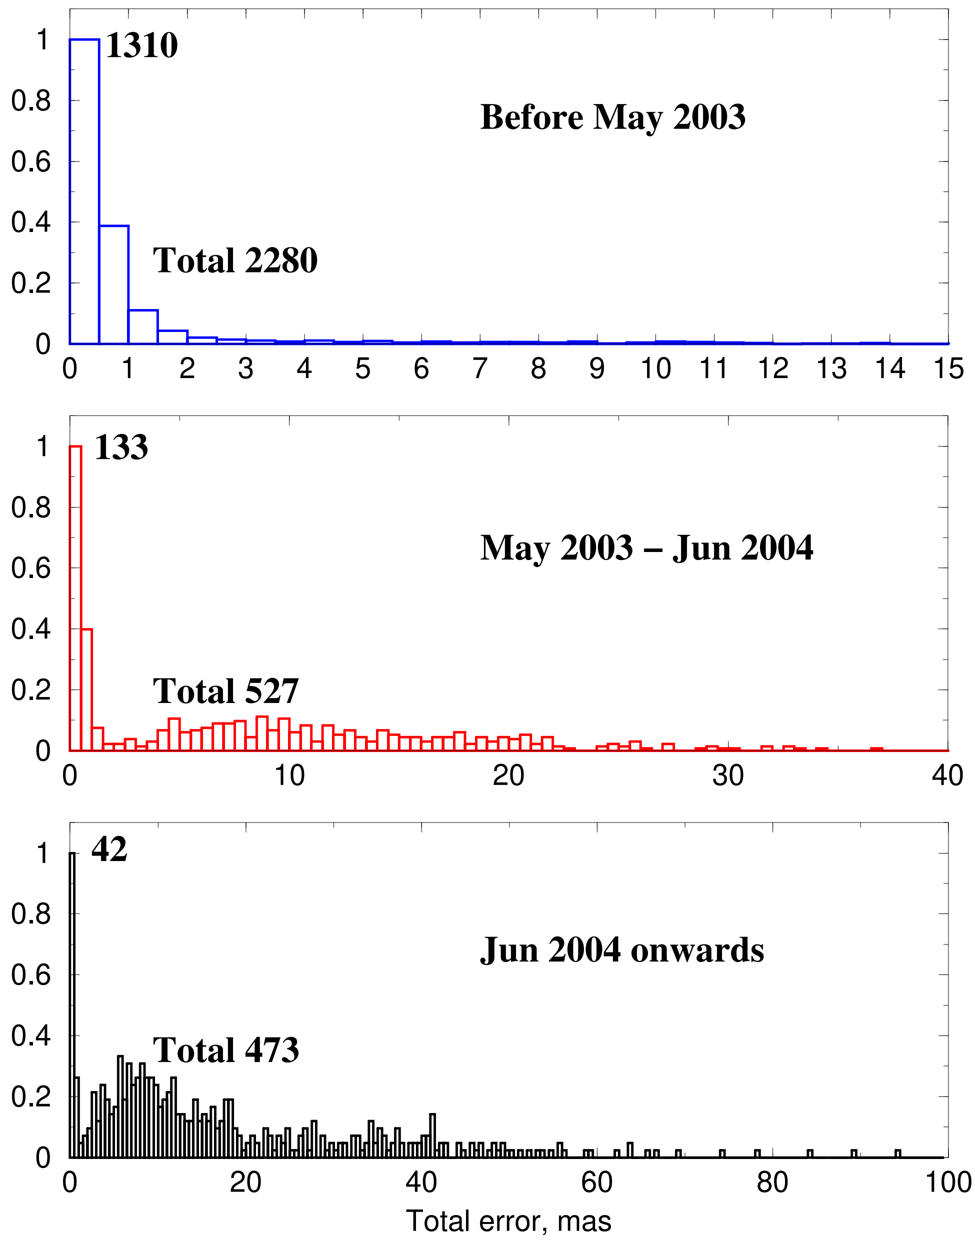

Figure 2: Histograms of the same data shown in Figure 1 for different time ranges. The horizontal axis is the square root of the sum of the squares of the different offsets, which is meant to be a measure of the total error even though UT1–UTC is not orthogonal to the other terms. Note the change of the horizontal scales between plots. The top plot shows the offsets seen while using the rapid service values. Most are better than 1 mas total error, while a few were worse. After May 2003, the number of projects with poor EOP increased dramatically, and it got even worse after June 2004.

Table 1: EOP used in Test Observations of 2005 August 13

Date	Item	UT1–UTC	X	Y
2005 Aug 12	Good EOP	–0.602573	0.02009	0.42696
	Poor EOP	–0.604434	0.01655	0.42076
	Difference	0.001861	0.00354	0.0062
2005 Aug 13	Good EOP	–0.602680	0.02162	0.42756
	Poor EOP	–0.605107	0.01845	0.42130
	Difference	0.002427	0.00317	0.00626
2005 Aug 14	Good EOP	–0.602558	0.02284	0.42816
	Poor EOP	–0.605638	0.02001	0.42178
	Difference	0.003080	0.00283	0.00638
2005 Aug 15	Good EOP	–0.602241	0.02388	0.42872
	Poor EOP	–0.602659	0.02356	0.42832
	Difference	0.000418	0.00032	0.00040
2005 Aug 16	Good EOP	–0.601808	0.02495	0.42902
	Poor EOP	–0.602220	0.02432	0.42912
	Difference	0.000412	0.00063	–0.00010

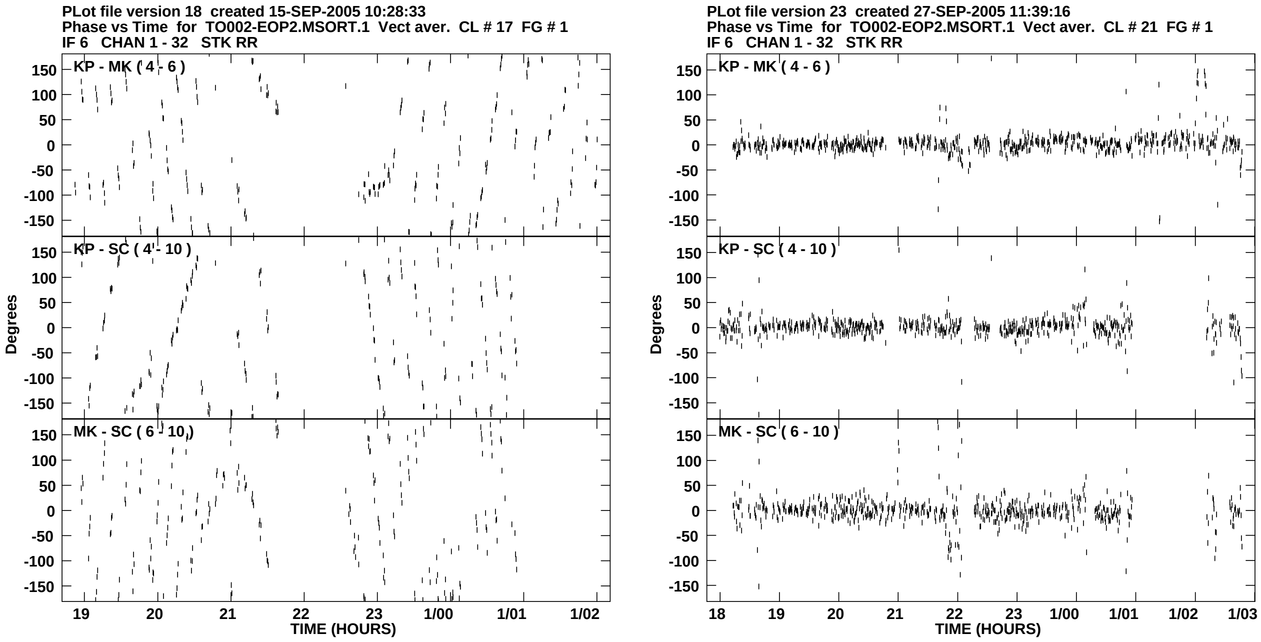

Figure 3: Plots of phases from a 4 cm IF from the test observation for baselines involving Mauna Kea, Kitt Peak, and St. Croix — a set of long baselines. These are from the processing pass in which poor EOP values were used. The left plot shows phases calibrated using a CL table copied from the data set processed using good EOP. That CL table includes results from fringe fitting so, when used on the the good EOP data set, it causes the phases to be near zero. When applied to the poor EOP data set, the residual phases will be only those resulting from the EOP errors. The right plot is the result of applying CLCOR to the CL table used in the left plot. It closely reproduces the phases in the good EOP data set calibrated with uncorrected CL table.

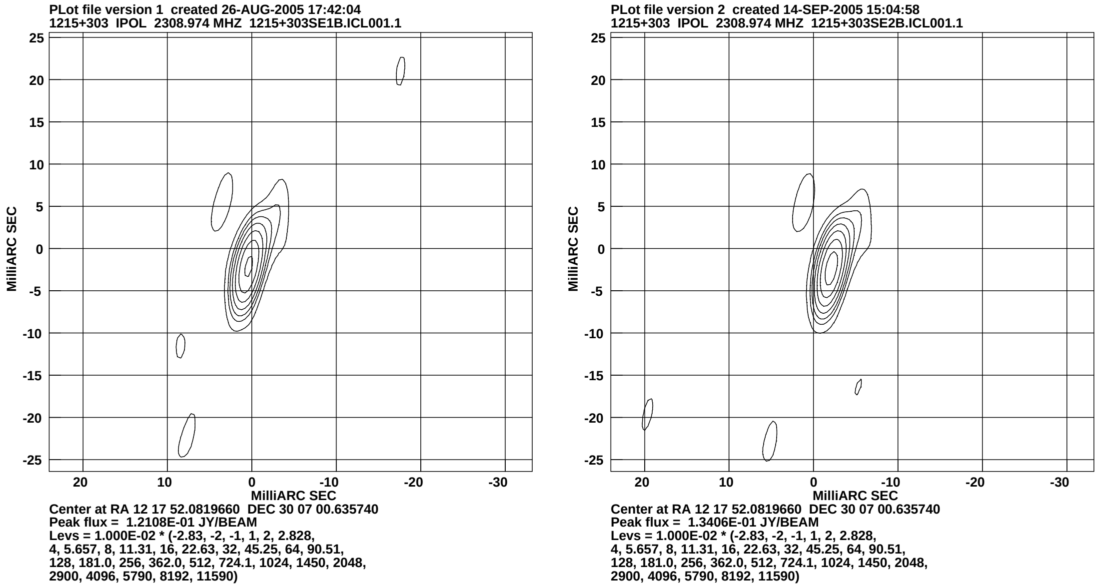

Figure 4: Phase referenced images of 1215+303 at 13 cm from the test observation. The left image was made from data processed with good EOP. The right image was made with EOP offset by amounts typical of what was seen while the EOP bug was in effect. There is a clear position shift, but at this relatively low frequency, not much change in image quality.

for the other. The data set processed with good EOP was fully calibrated, including fringe fitted, to the point where the residual phases of the calibrators are all near zero. The CL table with that calibration was copied to the data set correlated with poor EOP and applied in making the plot on the left. The plot shows the phases on 3 long baselines. There is rapid phase winding, due entirely to the EOP offsets. The plot on the right is made with a CL table produced by using CLCOR to add corrections for the EOP offset to the CL table used in the left plot. The phases are now flat and near zero so the correction has been very effective. This also demonstrates that essentially all of the phase offsets in the left figure are actually due to EOP offsets.

Figure 4 shows the 13 cm phase referenced images of 1215+303 from the good EOP correlation pass (left) and from the poor EOP pass (right). For this low frequency, where the rotations caused by the poor EOP are not much larger than the beam, the image quality is not much affected. But notice the shift of about 2 mas in position. Figure 5 shows a similar pair of images of the same source, but at 4 cm. Now there is a noticable degradation of the image quality in addition to the position shift.

Figure 6 shows the 13 cm phase referenced images of 1116+128. Figure 7 shows the 4 cm images of the same source. Again the left images are made from data with good EOP while the right image is made using poor EOP. As with the other source, the effect at 13 cm was primarily a position shift. But at 4cm, despite the smaller calibrator/target separation than the other pair, the phase referencing for this pair did not seem to work as well. The position offset scales about with the separation (about 1 mas in this case with about half the separation.), but the image quality was more affected by the poor EOP.

Note that these tests show a significant impact on image quality for observations at 4 cm. Phase referencing observations at higher frequencies is common and will be even more affected, although they typically use calibrators that are closer to the target.

CLCOR was run on the data set correlated with poor EOP to test the CLCOR corrections. The *usno_finals.erp* file used for the corrections was edited so that, over the days of interest, the values matched those used for the "good" EOP correlation pass. The most direct comparison was made by differencing the total delays from the good EOP pass and the corrected poor EOP pass. The total delay here is the geodelay column in the CL table plus the delay column for a CL table that has only had CLCOR affect the delay (no FRING etc). Figure 8 shows the results of that comparison for the station delay for 2 representative stations, Kitt Peak and Mauna Kea. What will actually matter is the difference between these station delays. The differences, after some debugging, are now down to the level of about 1.4 ps max in the test data. Note that 1 ps is equivalent to about 3 degrees of phase at 8.4 GHz. One slight concern is that the difference is increasing with time through the observation as is the scatter between sources during the 3 time periods (beginning, 0.9-0.93 day, and end) when sources scattered around the sky were observed. That effect is not yet understood. It is not clear how a long observation would be affected. A project is in progress to correlate a 24 hour run twice with different EOP to explore this effect, but the results are not ready for this memo.

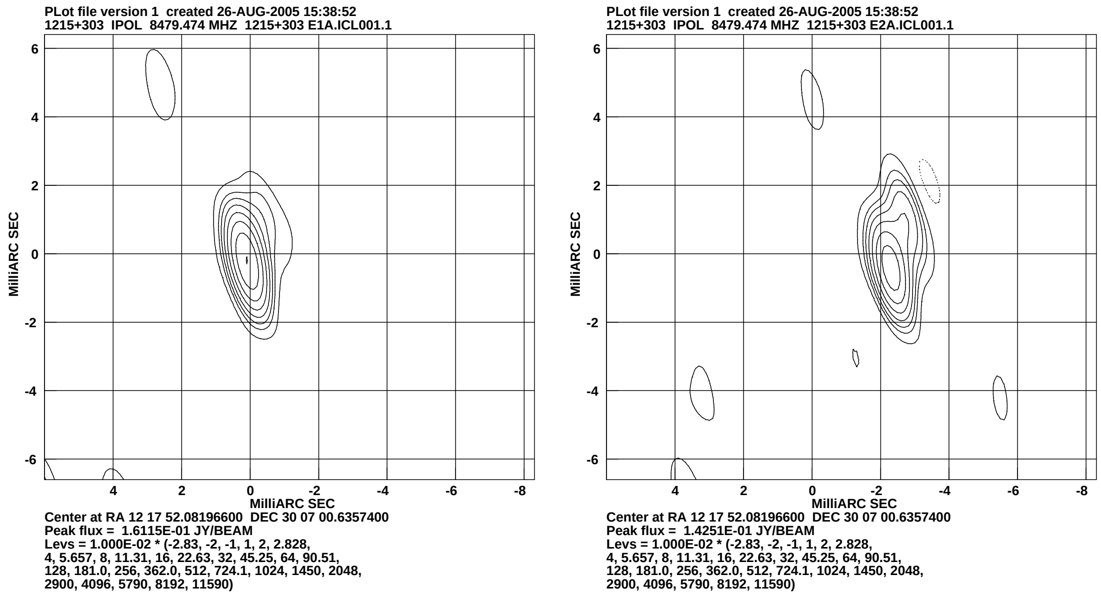

Figure 5: Phase referenced images of 1215+303 at 4 cm from the test observation. The left image was made from data processed with good EOP. The right image was made with EOP offset by amounts typical of what was seen while the EOP bug was in effect. The position offset seen as 2.3 GHz is seen. At this higher frequency, the image quality is also somewhat degraded

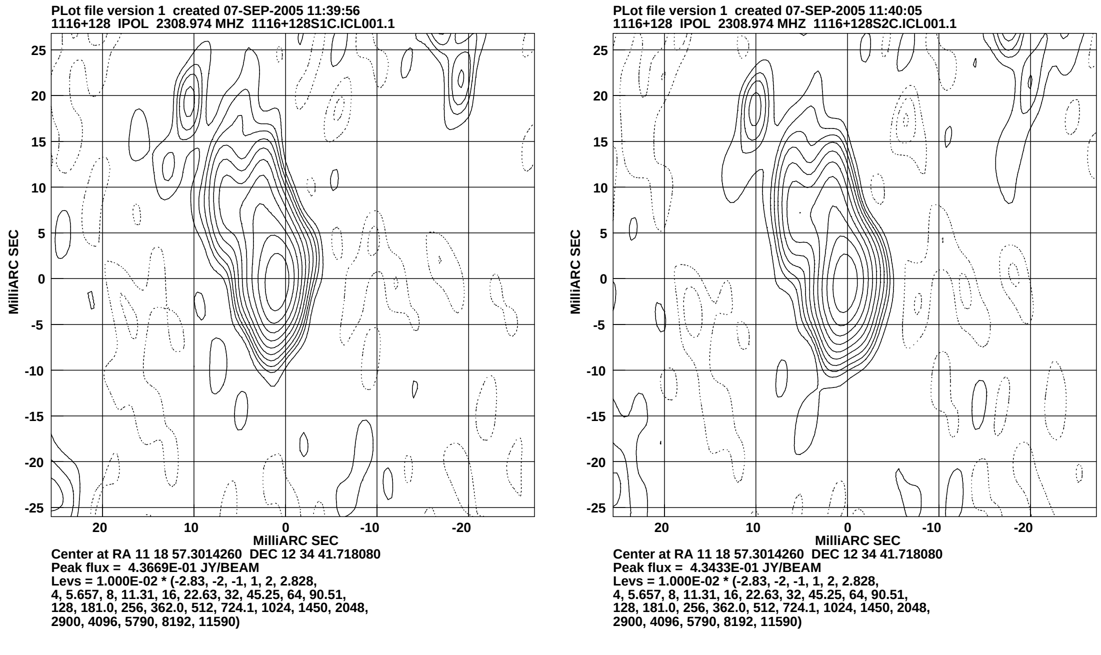

Figure 6: Phase referenced images of 1116+128 at 13 cm from the test observation. The left image was made from data processed with good EOP. The right image was made with EOP offset by amounts typical of what was seen while the EOP bug was in effect. The offset between calibrator and target is smaller than for the source in Figure 4 so the image shift is smaller. Again there is little change of image quality at this frequency.

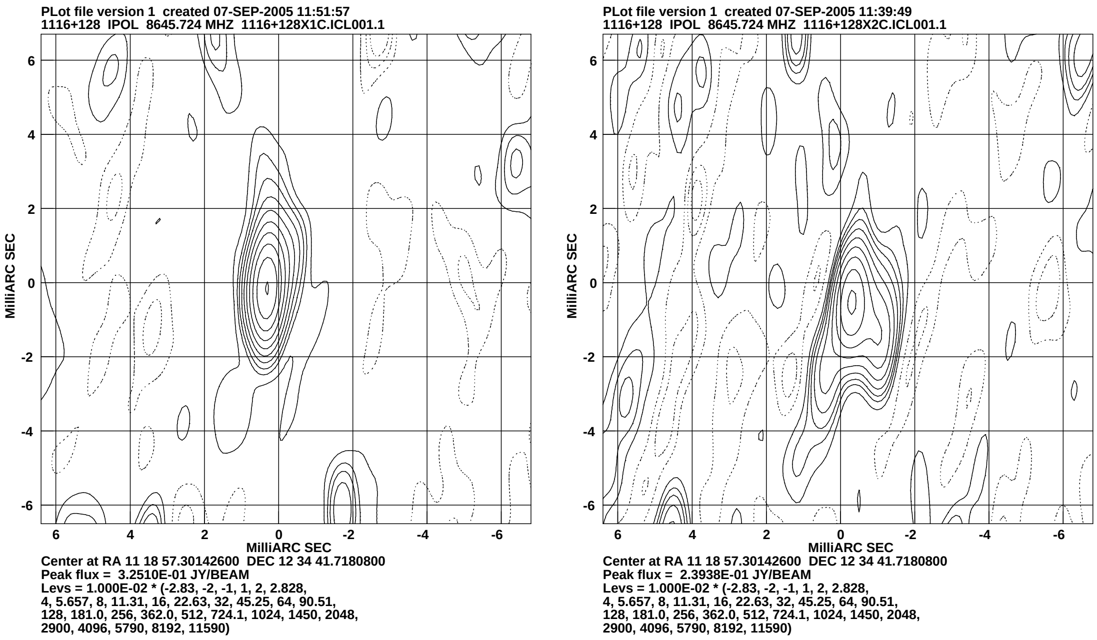

Figure 7: Phase referenced images of 1116+128 at X band (4 cm) from the test observation. The left image was made from data processed with good EOP. The right image was made with EOP offset by amounts typical of what was seen while the EOP bug was in effect. A position shift is seen and degradation of the image quality is pronounced.

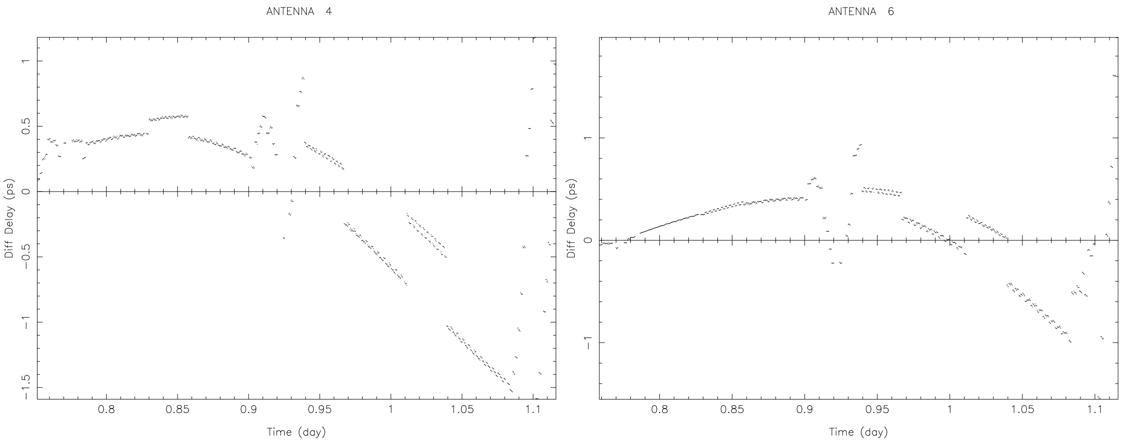

Figure 8: Difference between the CL table geodelay values from the processing pass done with good EOP and the geodelay plus the CLCOR correction for the pass done with poor EOP. The left plot is for Kitt Peak and the right plot is for Mauna Kea. If the CLCOR corrections were perfect, the difference would be zero. Note that 1 ps at 8.4 GHz is 0.0084 turn of phase or about 3 degrees of phase, so the differences are small.

Another test was done by comparing the final calibrated phases for reference and target sources for data correlated with good EOP and data correlated with poor EOP and corrected with CLCOR. For this test, the data were calibrated and fringe fitted separately on each dataset (A CL table was not copied from one to the other as in the test described earlier). The differences are a few degrees and of variable sign so they seem to be within the fluctuations due to the different correlation passes and different fringe fits. Without the CLCOR corrections, the calibrated target source phases are different by significant fractions of a turn. Note that, with this type of processing, the phase calibrator phases should be the same, as seen, because of the post-correction fringe fit.

An image equivalent to that of Figure 5, made with poor EOP data corrected by CLCOR looks identical to that in the left panel of Figure 5, the one made with good EOP. A simple fitted position for the peak of the source gives a position different from the one made with good EOP by 27 μas, which is probably limited by noise. The offset of the peak in the image made with poor EOP is 2.4 mas, or nearly 100 times worse. The image off-source noise, from an histogram fit, for the corrected data is 6% higher than for the good EOP data set, but, since no effort has been made at editing or other cleanup, that may just be differences in the processing. It is unlikely to be due to the EOP correction.

The conclusion of the tests so far is that the corrections available with CLCOR are good to the levels required by most users. Users with extreme requirements (order 10 μas is possible with care), should consider redoing the model if possible if they have data processed with poor EOP.

7 Correcting Data - The Long Version with Full Explanations

There are two methods, in principle, to correct existing data. VLBA data sets contain a record of the total delay model used on the correlator. That record is in the form of polynomials in the IM (Interferometer Model) and CL tables. A very good way in principle to correct any issues with the correlator model is to simply calculate a new total delay model and adjust the data by the difference between that result and the correlator model. That is the scheme used for geodesy and that is why geodesy observations have not been affected by the EOP error. Currently there is no capability to make this style of correction in AIPS. This incident has added to the motivation to add one.

The second method is to calculate the difference between a delay model made using the same EOP that were used in correlation and one that uses improved EOP. The difference is then applied to the data. As long as, other than the EOP, the two models are identical, the difference is not very sensitive to the quality of the overall delay model so a fairly crude model can be used. This is the scheme implemented in CLCOR.

In order to correct for poor EOP, CLCOR uses information from internal AIPS tables, from user input, and from an external file that the user needs to download. First, the EOP actually used on the correlator must be determined and getting that right has revealed some issues with the interaction between the correlator software and the correlator job generator. It has also revealed a weakness, recently corrected, in the correlator model parameters (CT) table in AIPS. The vast majority of users will not be affected by these issues but CLCOR needs to deal with them so they are described here.

The correlator obtains its Earth Orientation Parameters from the "UTC Table" in a correlator job script. There are typically a few to a few tens of job scripts needed to cover the whole time range of a project. The UTC Table has values for UT1–UTC, Xpole, and Ypole for midnight UT of several days. The correlator reads the first 5 such days, making an assumption that there are only 5 present and that those 5 are the ones it actually wants. The values of the parameters can change significantly from day to day (up to about 1 ms = 15 mas for UT1–UTC) so the correlator interpolates between the provided values to get a value for each individual delay calculation. That interpolation is done with a spline fit. An example of such a spline fit is shown in Figure 9. These are the "poor" EOP points from the August 13 test and include a step of the sort that can be induced by the presence of 10-16 day predicted values.

Complications for CLCOR occur when a project crosses midnight. For projects that cross midnight, the job generator provides EOP for more than 5 days for some or all jobs, including those that apply to times before midnight. Usually a project just crosses midnight once and most or all jobs have EOP from 6 days. But there are very occasional projects with more days that have more entries (One late 90's case with 12 has been seen). The correlator always takes the first 5 days of EOP data in the job, regardless of how many days data are there or when scans times are relative to the EOP rows. So to determine what values the correlator used, CLCOR must determine, for any given time, what day's EOP were in the job script covering that time. CLCOR is able to do that for essentially all cases without user help.

In VLBA correlator distribution data, there is usually (maybe always) one file per correlator job. The values of EOP used are in the CT table in that file. When the data are loaded with FITLD, and the distribution files (jobs) are concatenated, only one CT table is written. The rows of EOP data from each of the input jobs are concatenated. In most cases, where the rows from each of the jobs are identical, that means that the CT table contains many rows that repeat the same information. But there are exceptions. It appears that in jobs for times well past midnight (details still a bit unclear), the row for the first day covered in earlier jobs can be dropped. The correlator model for that job then will not be based on the same EOP values used in the earlier jobs. So the routine doing corrections needs to know when this happens.

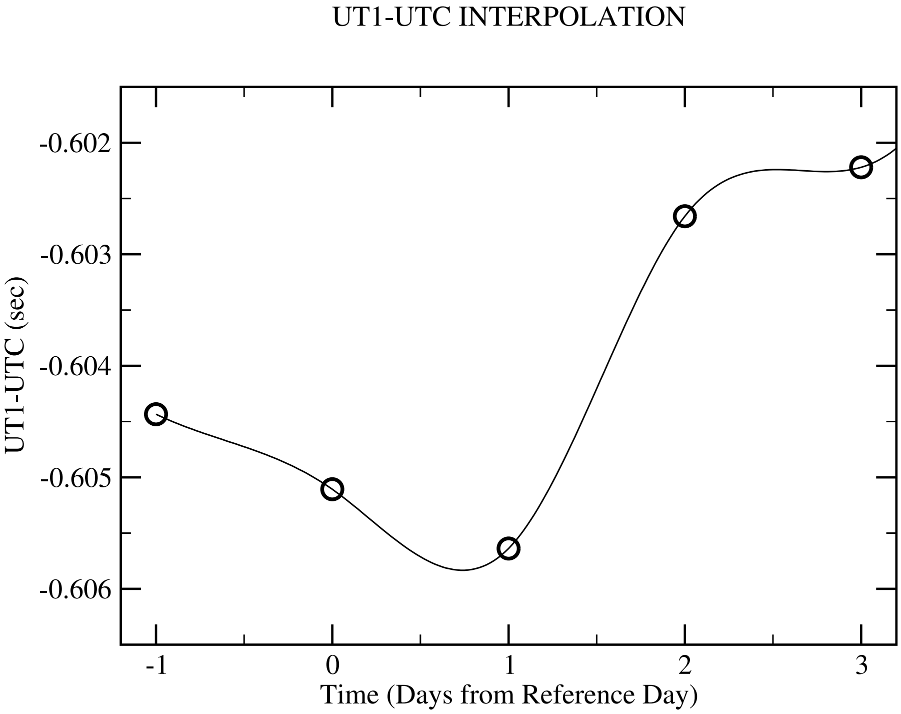

Figure 9: The correlator gets one set of EOP numbers each day. For each delay calculation, it interpolates to the actual time using a spline fit. This plot shows the spline fit for the UT1–UTC values used in the August 13 test observations for the correlator pass with "poor" EOP. The actual observations spanned a several hour period near 1 day after the reference day.

Unfortunately, the original version of the CT table (before 15 Sept 2005), which simply copies the information from the distribution files, does not include a time range for which each row applies. That time range is obvious from other information for the individual jobs, but was lost once FITLD concatenates the files. The only record of the time ranges of each job that AIPS had was in the history records FITLD writes, which are not really meant for other AIPS tasks to read. As of mid September, 2005, FITLD has been modified to write a new varient of the CT table that includes the time range to which each row applies. Thus for newly loaded data, it will be easy for CLCOR to tell what rows should be used in corrections.

When run CLCOR will attempt to figure out which CT rows to use for the interpolation to each CL table row time. For most observations, all sets of CT table rows are the same. Sets are rows from a single correlator job. They are easily distinguished from other sets by a drop in time between successive CT table rows when the new set starts. CLCOR checks if all sets are the same and, if so, simply uses the first 5 rows for correcting the entire data set. If they are not the same, CLCOR uses the time range column for new type CT tables or, for the old type, reads the history records from FITLD and uses the time information there to assign times to each CT table row. Note that, in the case of non-identical sets, getting the time information right depends on the structure of the CT table relating nicely to the history records from FITLD. If a post-FITLD concatenation operation such as DBCON or VBGLU has been done, this might be problematic and the user should be very alert as to what is going on.

Once CLCOR has the proper 5 days of EOP data, it interpolates to the time of the CL table entry using the same code that is used by the correlator. Thus the interpolation should be accurate to high precision.

There are a small number of projects from years ago that extended over several days. The job scripts contained EOP data from many days — up to 12 in a standout case. Since the correlator uses the first 5 days EOP regardless, it is likely that later days in the observation were correlated with EOP values actually extrapolated off the end of the time range of the values provided. Spline fits are notoriously poor at extrapolation. After only 3 days or so from the last point used, the EOP would be very poor. Any user with data from many days should look very closely at their situation and it might be best to use a full model recalculation to make corrections.

CLCOR also needs to know the desired "good" values of EOP. It is set up to read those from an external file, mentioned earlier, that can be downloaded from http://gemini.gsfc.nasa.gov/solve_save/usno_finals.erp. It is easiest to just feed this file to CLCOR (parameter INFILE) without modification, although it can be edited if you desire. If it is edited, beware of the somewhat odd units of tenths of arcseconds for X and Y and microseconds for UT1–IAT (not the UT1–UTC normally given). Those units must not be changed. The first column is the Julian Day number.

Alternate inputs for desired EOP are not yet implemented.

Data correlated with poor EOP has delay and phase offsets that vary with position on the sky. Any calibration that uses the measured phases and delays of a source to derive a correction for that source will remove the effect of the EOP errors. Such calibrations include fringe fitting and, for phases, self-calibration. This is why projects for which all sources are fringe-fitted and/or self-calibrated are not affected by the EOP error. It is also why some calibration steps, like bandpass calibration and pulse cal calibration, which are typically based on fringe fitted sources, are not affected. Where the corrections are needed is in situations where the corrections for one source are applied to another source at a different location. Then the relative delay in the model must be accurate and the EOP corrections need to be applied if poor EOP were used.

When the EOP correction is applied, it will change delays and phases on all sources. If it is applied to a phase calibrator whose phases have been forced to zero already by fringe fitting or self calibration, it will push those phases away from zero. Another round of fringe fitting or self calibration would need to be done to bring them back to zero to preserve the calibator position. All such fringe fitting and self calibration results, from before and after the EOP corrections, need to be applied to the target source(s) too.

It is probably best to apply the EOP corrections to a CL table that does not contain any data based delay or phase calibrations — that has not been fringe fitted or self calibrated. Such a CL table can contain, for example, the correlator offset corrections (ACCOR), parallactic angle corrections (also CLCOR), absolute gain calibration (APCAL), and ionospheric calibration (TECOR). Flagging will not be affected. Fringe fit and/or self calibration will need to be redone after EOP corrections regardless so it is best to do them only after the corrections. It is also probably best to do pulse cal calibrations (PCCOR) and bandpass calibrations (BPASS), both of which depend on data phases, after the EOP corrections, although that is not strictly necessary because they almost certainly will use fringe fitted sources whose final calibrations are not changed by the EOP corrections.

8 Improving the EOP Used on the VLBA Correlator

As described earlier, when the VLBA correlator job generator is working properly, the EOP values used are usually rapid service values from Bulletin A. Those values are far better than the predicts that got used while the EOP bug was in effect, but they can be based on data as little as 18 hours old, so they should be considered preliminary. Between 1 and 2 months after a given date, the "final" values are also published in Bulletin A and put in the data base. These are better values, but the recording media really need to be recycled before the final values are available.

A quick scan of some old Bulletin A issues on the archive at the USNO site show that the changes from the rapid service values to later, nominally better, values are on the order of a several tens of microseconds for UT1–UTC and a tenth of a milliarcsecond for the pole positions. These offsets are probably well below the errors in the data from other effects such as atmosphere, so the use of rapid service values should not be limiting most observations.

The standard scheme of acquiring EOP from the Bulletin A will start to give poor values (predicts) more often as we try to shorten turnaround times to support Mark 5 recording system deployment. If an observation is at the right time of the week, it can end up with at least one day with predicted, rather than measured, EOP if the jobs are made as much as 9 days after the observations. Recall that the jobs contain 3 days of EOP data from after the observations. Now it is possible to correct such projects after correlation, but it is still undesirable to use poor EOP data.

An enhancement in the process of being implemented is to switch to using a "daily" file as the source of EOP for the correlator. The *usno_finals.erp* file recommended above for corrections will be used. Any changes to the EOP values are very small after a few days after real time, so the daily file, for which all days' data are updated to the latest available each day, would give very close to final values for most correlations. If we switch to a daily file, which has decent, if not final, measurements for the day before the file is made, jobs made more than 4 days after the observations will be based entirely on measured values. This could be significant for fast turn-around. In recent days, the correlator has been running closer to real time than typical in the past and the need for better EOP has been apparent.

A related change that should be made is to rationalize the determination of which EOP values are used on the correlator. The current situation where the job scripts contain a variable number of EOP rows and the correlator uses the first 5, no matter what days those are for, is not very satisfactory. Either the correlator should select rows based on date, not just the first 5, or the job script writer should only put out 5 rows.

If we update the EOP values in the correlator every day, the probability that we could end up with different job scripts (time ranges) for a project having different EOP are increased. We will need to be careful to avoid such situations. Typically now all job scripts are made at the same time, and that will protect us against variable EOP in the future. This is actually an old issue related not just to EOP, but source and station positions etc. You don't want any to vary between jobs for a given project. Ideally the job generator would have a mechanism to guarantee that this doesn't happen, but realistically the software support is not available to do this and we will need to continue to rely on the vigilance of the analysts.

9 Additional Issues

In the process of debugging the CLCOR corrections, it was found that there is a misunderstanding about the Greenwich Sidereal Time that is recorded in the AN table in AIPS. That value is passed from the correlator. It is used in many places in CLCOR and probably elsewhere in AIPS, not just for EOP corrections, for calculating hour angles. For that purpose, the value needed is the Greenwich Apparent Sidereal Time. But in files from the VLBA correlator, the AN table contains the Greenwich Mean Sidereal Time. The differences are on the order of of a few tenths of a second of time. For most purposes, this is not significant. Even for calculating the EOP corrections, it is not significant because the same value is used for calculating both the before and after delays and the impact on the difference is small. But this trap for the unwary should be fixed.

Also while working on the EOP corrections, it was noted that CLCOR modified the geometric delay column in the CL table when doing some types of correction (source and station positions and axis offsets) and not others. In the interests of having the geometric delay reflect what has actually modified the phases, we decided not only to not add the EOP correction to that column, but to make CLCOR not add any other correction either since it does not actually modify data phases. Later, while thinking about the impact of the EOP error on the spacecraft tracking project, we were reminded that the data path to the totals that we delivered to JPL depended on CLCOR making modifications to the geometric delay. So we have broken the spacecraft tracking data path. We are now considering what to do about this, with the likely outcome being to add a user option to decide whether or not to alter the geometric delay column.

Users should be warned that the EOP correction as implemented in CLCOR is specific to the VLBA correlator. This is because CLCOR must be able to reproduce the values of the EOP used on the correlator. Since the parameters change by significant amounts every day — on the the order of a millisecond of time for UT1–UTC in recent years — some sort of interpolation must be done between the daily values. CLCOR has code adapted from the CALC model calculator used on the VLBA correlator to do this interpolation. Other correlators may use different code and so may get somewhat different results. Also the choice of which EOP values to use for the interpolation, especially after crossing a day boundary may be different. And finally, other correlators may not provide the required CT table data. So CLCOR probably won't even run for data from other correlators, and, if it does, may get results that are not quite right. But if it does run and there are no oddities like steps in the data, the results may not be too bad.

Note that the need for EOP corrections will rise in the future if eVLBI becomes popular. Fundamentally, an eVLBI correlator will be using predicted EOP values so making corrections after the fact will probably have to become a routine processing step. On the other hand, Non-VLBI instruments, like the VLA, have less stringent accuracy requirements and can use predicted EOP without degrading the vast majority of observations.

10 Preventing Future Problems

This whole episode raises the question of how we can prevent anything similar from happening in the future. It is easy to keep tabs on the EOP in the future, but surely the next problem will be something else so a scheme is needed to detect anything that degrades the VLBA's astrometric capabilities. The difficulty is that anything like the EOP problem is subtle and easily lost in other large, but normal, error sources such as the atmosphere. The raw correlator output from any test involving new observations cannot be directly compared with a standard run because of variations in clocks, atmosphere, station positions, EOP etc. A standard observation could be recorrelated periodically to be sure the correlator had not changed, but that would miss something like the EOP error because it would not be sensitive to failure to keep up with changing parameters.

A possible test, perhaps to be added to the weekly "MT" test observations, would be to do something like an hour of phase referencing between 2 or 3 very bright calibrators with well known position. Such calibrators should be much brighter than those used in the obserations reported here. They would need to be processed in AIPS, probably with the pipeline script, and relative positions and image RMS's should be measured. If there is any significant change, that would be cause for a staff scientist to investigate.

Another thing to do is just to keep an eye on the raw phases coming off the correlator. This is easiest for observations that are a long track on a bright source. If those sources seem to be subject to excessive phase winding, flags should be raised. At frequencies such as 4 and 2 cm, the number of turns of phase across a day should only be a few — we supposedly know all of the geometry except the atmosphere and clocks to better than a wavelength.

References

[1] Mioduszewski, A. J. & Kogan, L. 2004, Strategy for Removing Tropospheric and Clock Errors using DELZN, AIPS Memo 110, NRAO.

[2] Sovers, Ojars J., Fanslow, J. L., & Jacobs, C. S. 1998, Reviews of Modern Physics, 70, 1393.

[3] Sovers, Ojars J., & Jacobs, C. S. 1996, Observation Model and Parameter Partials for the JPL VLBI Parameter Estimation Software "MODEST" — 1996 Jet Propulsion Laboratory Publ. 83-39, Rev. 6, Pasadena CA.