Very Long Baseline Interferometry and the VLBA
ASP Conference Series, Vol. 82, 1995
J. A. Zensus, P. J. Diamond, and P. J. Napier (eds.)

Printed June 13, 1995

CHAPTER 10

Fringe-Fitting

W. D. COTTON

NATIONAL RADIO ASTRONOMY OBSERVATORY

Abstract The effects of errors in the model used by a VLBI correlator and ways to calibrate these errors are discussed. First, a relatively theoretical discussion of the causes, effects, and methods of correcting these correlator model errors are given. Later sections provide practical demonstrations of the effects and techniques for dealing with them.

Contents

10.1 Fringe Fitting Theory

10.1.1 Correlator Model Delay Errors

There are a large number of ways in which the model used by the correlator can be in error but they include: errors in the source or antenna positions used, errors in the earth model, errors in the clock epoch and rate at each antenna and errors in the atmospheric model. Correction of correlator model errors allows the data to be averaged in time and frequency, greatly reducing its volume. As shown in earlier lectures, the interferometer phase is related to the delay by the simple relationship:

$$\phi_{t,\nu} \;=\; 2\pi\nu\tau_t\,, \tag{10.1}$$

where $\phi_{t,\nu}$ is the interferometer phase, ν is the observing frequency, the subscript t indicates the time dependence and τ_t is the interferometer delay. The correlator uses this relationship to estimate the phase and remove it from the data. Differentiating equation 10.1 shows the relationship between delay and interferometer phase errors:

$$d\phi_{t,\nu} = 2\pi\nu d\tau_t\,. \tag{10.2}$$

Thus, correlator delay errors will result in time and frequency dependent errors in the interferometer phase. A first order expansion of the error in interferometer phase can be expressed as:

$$\Delta\phi_{t,\nu} = \phi_0 + \left(\frac{\partial\phi}{\partial\nu}\Delta\nu + \frac{\partial\phi}{\partial t}\Delta t\right)\,, \tag{10.3}$$

where ϕ_0 is the phase error in turns at the reference time and frequency, $\frac{\partial\phi}{\partial\nu}$ is known as the delay or delay residual and $\frac{\partial\phi}{\partial t}$ is known as the rate, delay rate or rate residual. In some cases, e.g.orbiting VLBI and very wideband observations, higher order approximations may be required. The boundary between fringe fitting and phase calibration is often fuzzy but any process that estimates the delay and/or rate residuals is referred to as *fringe fitting.*

The simplicity of the relationship expressed by equation 10.3 argues strongly for keeping the data in the time and frequency domain (as opposed to the time and delay domain). In this discussion it will be explicitly assumed that the data are kept in this form. Corrections for residual delay and rate errors can then be made as phase corrections as a function of time and frequency. After these corrections are applied, the data can be averaged in time and/or frequency. An alternative approach of estimating the delay and rate residuals and the corrected, averaged data in the delay-fringe rate domain is discussed in Moran (1976), Alef (1989a) and Walker (1989b).

10.1.2 Fringe Fitting Techniques

Most of the sources of delay error in the correlator model can be separated into contributions from each antenna. The baseline dependent errors can then be expressed as differences of antenna dependent errors. The phase errors for the baseline between antennas i and j can then be expressed as:

$$\Delta\phi_{ij} = \phi_{i0} - \phi_{j0} + \left(\left[\frac{\partial\phi_i}{\partial\nu} - \frac{\partial\phi_j}{\partial\nu}\right]\Delta\nu + \left[\frac{\partial\phi_i}{\partial t} - \frac{\partial\phi_j}{\partial t}\right]\Delta t\right). \tag{10.4}$$

Since the delay and rate residual models can be expressed in a variety of forms, there are several corresponding techniques for determining the residual delay, rate, and in some cases phase. Several of these techniques are discussed in the following.

10.1.2.1 Baseline

Equation 10.3 can be used as a model for the phase errors in the data and used to determine the baseline delay, rate and phase residual. In practice, what is usually done (Moran 1976) is to Fourier transform the data into the delay-fringe rate domain (e.g., see chapter 11). In this domain the source function is peaked at the delay and rate residual and the value of the function at that point has the correct amplitude and residual phase of the average of the data going into the Fourier transform. Baseline fringe fitting usually consists of finding this peak in the delay-rate domain and using the Fourier transform to average the data. The values of fringe amplitude and phase determined in this way are those that would also be obtained by spectrally averaging.

10.1.2.2 Baseline with Closure Constraints

There are several difficulties with the baseline oriented fringe fitting approach. The most important is that the source involved must be detectable on all baselines or data for those baselines cannot be calibrated. The second difficulty is that this technique does not maintain the antenna based, or *closure* relationships of the delay, rate and phase residuals. Both of these difficulties can be addressed by explicitly including the closure constraints (Alef 1989a; Alef and Porcas 1986). In this procedure the delay and rate residuals are estimated for each baseline and then decomposed into antenna based components using a least squares or other procedure. The antenna based delay and rate residuals are then used to predict the locations in the delay-fringe rate domain of the average data values. This procedure will result in corrupted estimates of the visibilities if the source is sufficiently large that it broadens the delay-fringe rate function. Alternatively, the delay and rate residuals can be used to correct data in the time-frequency domain; in this case proper time and frequency averaging can eliminate the distortion of large sources.

10.1.2.3 Global

Either of the baseline based fringe fitting techniques require that the source be clearly detected on at least some, if not all baselines to a given antenna. For older arrays where there were usually a couple of dominant antennas this approach is adequate. However, for homogeneous arrays like the VLBA, weaker sources can be detected using a coherent or *global* fringe fitting technique (Schwab and Cotton 1983; Walker 1989b; Walker 1989a). This technique uses equation 10.4 and all data to jointly estimate the antenna-based phase, delay and rate residuals. This procedure allows calibration of antennas which lack sufficient signal to noise ratio on any given baseline, but have a sufficient signal-to-noise ratio when used with the rest of the array.

Since an interferometer is only sensitive to differences in delay, rate and phase residuals between antennas, these values cannot be measured in an absolute sense. However, one antenna can be declared to be the *reference* antenna and the delay, rate and phase residuals for that antenna arbitrarily assigned a value of zero.

10.1.3 Solution Interval

The approximation given in equation 10.4 for the phase residual as a function of time and frequency will be adequate for a limited period of time. Random atmospheric phase fluctuations will eventually cause significant deviations from this model. The time period during which a model like equation 10.4 is adequate is referred to as the *coherence time*. Under many circumstances the *solution interval*, or time period for which a model is fitted to the data, can exceed the coherence time.

An increase in the solution interval has two opposing effects. The first is that the sensitivity is increased by increasing the amount of data in the solution. The other effect is that the decreasing coherence due to increasing deviations from the model reduce the sensitivity. In principle, the solution interval could be set so that these two effects canceled although, in practice, the timescale of the decorrelating effects are not known accurately. If self-calibration is later used in the imaging stages, it should remove much of the remaining phase residuals. Interpolation in time of a series of phase and rate solutions before application to the data has the effect of adding a higher time derivative to the solutions.

10.1.4 Calibration Sources

It is frequently possible to use strong sources near the source of interest to estimate delay and rate errors due to correlator model errors other than source position. This technique is especially useful to correct large clock errors. After corrections from a calibrator source are applied, the search window in delay and rate can be substantially narrowed reducing the probability of a false detection. Calibrator sources can also be used to calibrate a number of other instrumental problems as described later.

10.1.5 Source Structure

If the source being observed is significantly resolved, the structure of the source will introduce baseline dependent delay, rate and phase residuals. This source structure can cause problems for fully coherent fringe fitting methods such as the global fringe fitting technique. The effect of source structure is to add *noise* to delay, rate and phase solutions and may reduce the apparent signal-to-noise ratio of the solution. An improved fringe fit can be obtained if a model of the source structure (e.g., CLEAN components) is available. In this case, the model can be Fourier transformed and divided into the data. This process will, to the degree that the model is accurate, remove the effects of source structure, producing a set of visibilities with the amplitudes and phases that would have been observed for a point source of the same flux density. However, the signal-to-noise ratio will remain unchanged so the weight of the data should be reduced to reflect the effects of resolution.

The source structure is generally not known a priori. However, a fringe fit can still be performed using a point model. If this fringe fit does not detect all the antennas in all solution intervals then the data corresponding to the detections can be imaged and the resulting model used in a subsequent fringe fit.

10.1.6 Phase Referencing

It is beyond the scope of this lecture to give a detailed description of phase referencing techniques. However, the techniques discussed are related to phase referencing. If observations of a calibrator source are interleaved with observations of the target source, then fringe fit solutions for the calibrator source can be used to calibrate the delay(s), rate and phase residuals of the target source. This data can then be used to image the target source. For this technique to work, the structure of the calibrator source must be known prior to the fringe fitting so that the effects of this structure can be removed. A model of the calibrator may be determined from self-calibrating the results of a previous fringe fit of the calibrator. Additional calibration stages may be required prior to the final fringe fitting to insure the most accurate possible a priori model.

10.1.7 Multi-band Data

Most of the wideband recording systems in use for VLBI divide the bandpass up into a number of separate bands. These bands may be either contiguous or widely separated in frequency depending on the application. Since the different bands have partially independent cabling and electronics they add additional complications to the fringe fitting process.

10.1.7.1 Phase-Cals

One effect of the separate bands is that each will have a peculiar, possibly time variable, contribution to the phase and delay residuals in its data. A technique commonly used to remove this effect is to introduce a series of tones, called *phase-cals* or *pulse cals*, into the signal path. These tones are frequently introduced at or near the feeds so that they can calibrate the entire processing after that point. These tones are usually spaced every MHz and are arranged such that one or more appears in each separate band. These tones can then be used in one of several ways depending on the number of tones in each band. If there is a single tone in each band (the case for the Haystack Mark III system) then the phase of this tone can be used to correct the phase and perhaps the rate of the data in that band. If there are multiple tones in each band then the slope of phase with frequency of the tones can be used to correct the delay of the data in that band as well. The VLBA phase calibration will provide multiple tones per band.

10.1.7.2 Multi- vs. Single-band Delay

In the case of multiple tones per band, the phase-cals can, in principle, be used to remove the effects of the separate electronics and the fringe fitting reverts to the single band case. However, for the Haystack Mark III system of using a single tone there are additional complications. The phase of each band will have a constant adjustment by the phase-cals but the slope of the phase with frequency within the band will be unchanged. Thus, there are two sets of delays, the *multi-band delay* as determined from the phase slope between bands and the *single-band delay* as determined from the phase slope within each band. In general these delays will not be the same. The effect of the phase-cal is to change the multi-band delay to the delay at the point of injection of the phase-cal signals. The single-band delays remain the appropriate delays at the sampler. Further complications are that the single-band delays may all be different and their relation to the multi-band delay may be time variable. Fringe fitting of multi-band data using a single tone phase cal system must estimate both single- and multi-band delays (Alef 1989a).

10.1.7.3 Sidebands

Multi-band systems frequently allow use of both sidebands from the last heterodyne stage. These sidebands may be combined into the same band as is frequently the practice with Mark III data. Since there are partially independent electronics, the phases of the two sidebands may not be aligned. In addition, the Haystack Mark III phase-cals are decoded only in the upper sideband but are applied to the lower sideband.

10.1.8 Filtering

If they are sufficiently stable, the estimates of the delay, rate and phase residuals may be filtered to reduce noise, and to remove and/or replace bad values. Any filtering operations are best done on antenna-based values in order to preserve the antenna-based nature of these corrections. However, all delay, rate and phase residual determinations are made using observations sensitive only to the differences of these values between pairs of antennas and thus are determined with respect to the reference antenna. It is important that all values in a filtering operation be referred to the same reference antenna.

10.1.8.1 Establishing a Common Reference Antenna

If all delay, rate and phase residuals are measured with respect to the same reference antenna, operations such as smoothing and interpolation are well defined. If the reference antenna varies with time, as is frequently the case, then the delay, rate and phase residuals must be re-referred to a common reference antenna. If both the primary and secondary antennas are used at a given time then it is straightforward to change from one reference antenna to another. The correction for times during which data from the primary antenna was unavailable can be estimated by smoothing and interpolating and/or extrapolating differences between the primary and any secondary reference antennas from times when they were both measured. In this procedure rates and phases must be used together as one is the time derivative of the other.

10.1.8.2 Smoothing and Interpolating Solutions

After all delay, rate and phase residuals have been referred to a common reference antenna they can be filtered. One possible operation is to smooth the solutions to reduce the effects of noise. The time sequence of solutions can be examined to detect discrepant values and mark them as invalid. Finally, in low signal-to-noise ratio cases, failed or invalid solutions can be replaced with estimated values by interpolation and/or extrapolation of valid values.

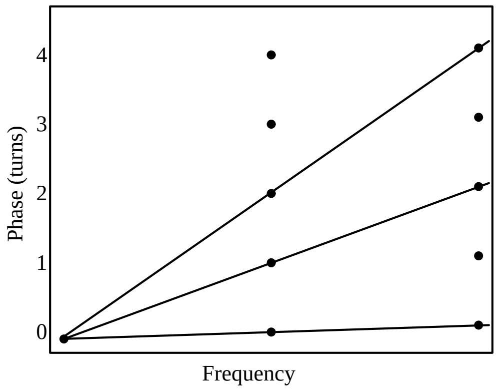

Figure 10.1: This figure illustrates the multiple possible multi-band delay solutions (lines) derived from ambiguous (2π) phase measurements (filled circles) at discrete frequencies.

10.1.9 Bandwidth Synthesis

In some circumstances it is advantageous to divide the recorded bandpass into a number of widely separated bands. This technique is referred to as bandwidth synthesis. This technique can be applied to increase the (u, v) coverage of astronomical observations (Cotton 1989) or to increase the precision of delay estimates used for astrometric and geodetic applications. In this case it may be desirable to estimate single-band delays and rates from the individual bands and then estimate the multi-band delay from the phase residual slope between the different bands. Since the multi-band delays are estimated from phase measurements at a relatively small number of discrete frequencies there will be multiple possible values of the multi-band delays. The spacing of the frequencies sampled determines the delay ambiguity. More multi-band delay solutions are possible in the presence of noise as sidelobes in the delay function may be mistaken for the main response. Single-band delays and/or knowledge of the accuracy of the correlator model may be used to choose the correct solution. Figure 10.1 illustrates the delay ambiguity.

10.1.10 Weights

As with any parameter estimation scheme the significance attached to each datum is important if it is variable. The significance of a measurement can be quantified as a statistical weight, or, the reciprocal of the variance of the amplitude for visibility data. Such a measure may be available for data which has had a proper error analysis and amplitude calibration. This is not always possible; and in practice, relative weights are sufficient.

In the context of fringe fitting, weights can be considered to consist of a number of factors. First is a data sensitivity weight, ideally the proper statistical weight described above, but in any case this is a function of antenna sensitivity, integration time, bandwidth etc.

The second factor is a correction for source structure. If the data has been divided by the Fourier transform of a model then the weight of each visibility should be corrected by the square of the reciprocal of the fringe visibility of the model. This procedure takes into account the reduction of signal-to-noise ratio from the source resolution as well as correcting the observed phases.

A final possible factor is an arbitrary factor quantifying the reliability of an antenna. This is especially important if secondary reference antennas are to be chosen on the basis of the weights assigned to the data.

An alternate approach for determining weights for baseline-based fringe fitting with closure constraints has been described in (Alef and Porcas 1986). In this scheme all, or a subset, of the baselines are individually fitted in each solution interval and then the baseline weight is determined from the strength of any detected fringes. This approach eliminates the need for amplitude (and weight) calibration prior to fringe fitting. A multistage global fringe fitting procedure could also use this scheme.

10.1.11 Polarization

Fringe fitting of polarization data is similar to that for total intensity data; both parallel polarizations are fringe fitted independently. The additional calibration steps to adjust the parallel systems with respect to each other are described in chapter 15.

10.2 Fringe Fitting Practice

This section will discuss some of the more practical considerations involved with fringe fitting with illustrative examples.

10.2.1 Phase Slopes in Time and Frequency

As described earlier, errors in the correlator model will introduce residual phase slopes in time and frequency. In addition, multi-band systems will introduce band dependent phase and delay offsets. An example of this is shown in figure 10.2.

Figure 10.2 shows data as it was produced by the correlator with no calibration applied. It shows the slope of phase with time and, piecewise, the slope of phase with frequency. These data have a relatively high signal-to-noise ratio and the phases are easily seen in the 10 second integrations shown. In more typical cases the signal to noise-ratio is much lower and this sort of display shows mostly noise.

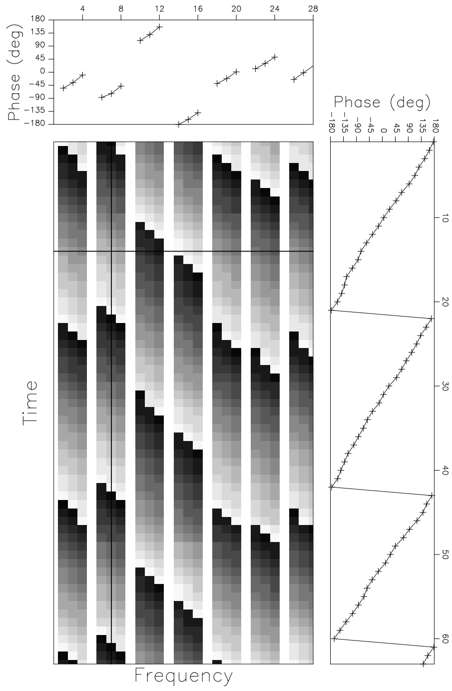

Figure 10.2: This figure shows raw phases in data from a single baseline taken using a multi-band (7 bands of 4 spectral channels each) recording. The gray scale plot shows the phase as a function of time and frequency. The frequency channels are evenly spaced and every fourth is blanked. Plots above and beside the gray scale plot show cuts along the lines throughout the gray scale giving the phase as a function of frequency at a given time and phase as a function of time at a given frequency.

10.2.2 Phase-Cals

The data shown in figure 10.2 were obtained using the Mark III system and correlated on a Haystack-style correlator and so have a single phase-cal tone per band. After the corrections derived from these phase-cals are applied the result is shown in figure 10.3. As can be seen from this figure the phase-cals roughly align the band to band phases along the slope defining the multi-band delay but this delay is very different from the single-band delays. What is not obvious from this plot is that there are residual band dependent phase errors on the order of 10 degrees. This will be demonstrated later. Phase-calibration systems using multiple tones, such as those provided for the VLBA, should align the single- and multi-band delays.

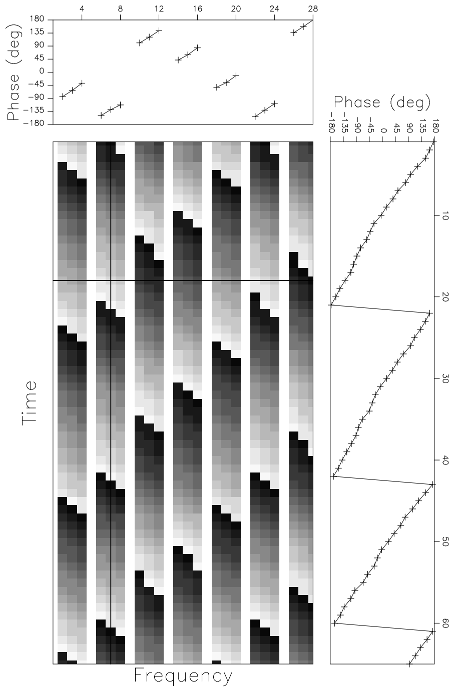

Figure 10.3: The data displayed in figure 10.2 after the application of phase-cals. Distinct single- and multi-band delays can be seen; the two delay residuals have opposite signs.

10.2.3 Sidebands

If data from the upper and lower sidebands are combined in a single band there is the possibility of a phase offset between the two sidebands. There are differences in the Mark III and VLBA recording systems that cause an apparent phase shift of about 130 degrees between the sidebands on baselines with mixed recorder systems. Any phase offset between sidebands must be corrected before fringe fitting.

10.2.4 Delay and Rate Fits

After the instrumental phase calibration is applied it should be possible to model variations of phase with frequency with a combination of single- and multi-band delays.

10.2.4.1 Signal-to-Noise Ratios

Fringe fits of weak sources run the risk of false detections of sources. That is, the fitted delay, rate and phase residuals may reflect the random nature of the noise rather that the instrumental response to the source. Some measure of the likelihood of a false detection is needed to reject such fits. The signal-to-noise ratio (SNR) on either an antenna or baseline basis is usually used for this. The SNR is a relatively well defined concept but the estimation of its value is not straightforward and in practice a number of schemes are in use. These mostly depend on the postfit scatter of the phases and define an SNR of 1 to correspond to an RMS phase scatter of 1 rad.

Estimates of the SNR can be used to reject solutions which do not have a reasonable probability of being true detections. The exact SNR cutoff depends on the details of the fringe fit. The larger the region of delay-rate space that is examined the larger the probability of a false detection at a given SNR limit. In many cases a limiting SNR of 5 is adequate.

10.2.4.2 Delay and Rate Windows

It is frequently both possible and desirable to restrict the region of delay-rate space to be searched. As described in the previous section, the probability of a false detection increases with the size of the region searched. If these windows are centered on zero values then systematic residuals need to be reduced as much as possible. This may involve observations of calibrator sources as discussed in a later section. In addition, when performing a multi-band fringe fit, it is often desirable to set the windows to as narrow a width as possible in order to reduce the changes of detecting an ambiguous multi-band delay.

10.2.4.3 Details of Global Fringe Fitting

Since the phase measurements are usually ambiguous by integral numbers of turns, the fitting of delays, rates and phases to the observations are non linear. Nonlinear fitting procedures such as least squares are quite sensitive to the starting solutions and frequently do not converge if the starting solution is too much in error. A multi-stage fitting procedure is suggested in (Schwab and Cotton 1983) to deal with this problem. This technique is to make use of the antenna-based phase relationships (closure conditions) and *stack* linear combinations of data from multiple baselines to produce a pseudo data set for the baseline from each antenna to the reference antenna. This pseudo data set can be Fourier transformed into the delay and fringe-rate domain and the location of the peak in this function gives the approximate delay and rate of that antenna with respect to the reference antenna. The phase of the function at the peak is an estimate of the phase residual and the amplitude is a function of the SNR. All of the information relating to a given antenna can be derived from relevant one, two and three baseline combinations. However, uncorrected source structure may cause significant baseline dependent phases which will decorrelate the *stacked* data set. In cases of sources with significant uncorrected structure, it may be useful to restrict the stacking to one and two baseline combinations,

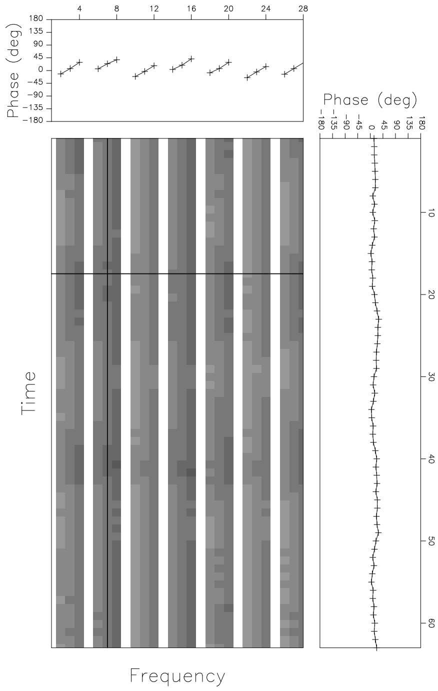

Figure 10.4: This figure shows the data displayed in figures 10.2 and 10.3 after correction for residual phase, multi- and single-band delays, and rates. No corrections for the multi-band/single-band delay differences or for errors in the phase-cals were applied prior to the fringe fitting. The resulting residual phase and single-band delay errors can be seen in the phase versus frequency plot.

or even the single baseline between the antenna in question and the reference antenna. This technique is frequently referred to as the *FFT fringe search.*

If the FFT fringe search uses only a single baseline then the choice of reference antenna becomes more critical and the use of primary and secondary reference antennas may be necessary. When a single baseline FFT search to the primary reference antenna fails to meet the minimum SNR then it is possible to use a secondary reference antenna. The choice of a secondary reference antenna can be made in a number of ways. One possibility is to pick an antenna detected on baselines to the primary reference antenna on the basis of the sum of the weights on all baselines involving that antenna. Another possibility (Alef and Porcas 1986) is to try all antennas as possible reference antennas and see which one gives the best SNR. If secondary reference antennas are used then solutions using them must be referred to the primary reference antenna before use.

Since the Fourier transform in this procedure is done using an FFT, all data must be gridded into equal time bins which may not accurately reflect the true time average for each data sample. For this and other reasons it is possible

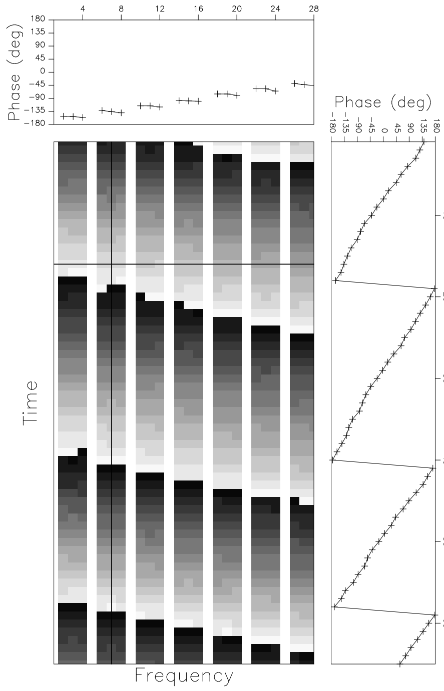

Figure 10.5: This figure shows the data displayed in figure 10.3 after applying a correction for phase cal errors and the difference between single- and multi-band delays determined from a calibrator source. Note there is still a difference between the single and multi-band delays.

to improve on the solution derived from the FFT fringe search by using a full nonlinear least squares solution using all the data. A least squares solution starting from the result of an FFT fringe search should converge fairly quickly and reliably. An estimate of the antenna SNR can be derived either from the least squares estimate of the uncertainty of the phase residual or the scatter of the data phases after correction by the fitted model.

10.2.4.4 Multi- and Single-band Delays

If the data being fitted have, after application of the phase-cals, distinct multi- and single-band delays, the fringe fitting must take this into account. All of the procedures described above can be modified to do a full 3-D fringe search. In practice, it is usually possible to correct for the difference between single- and multi-band delays using calibrator source(s) or the target source at a single time and then determining the, hopefully, smaller time dependent differences between residual multi-band and single-band delay. Calibration of these differences is discussed further below.

The results of fringe fitting done without correction for multi-band/single-

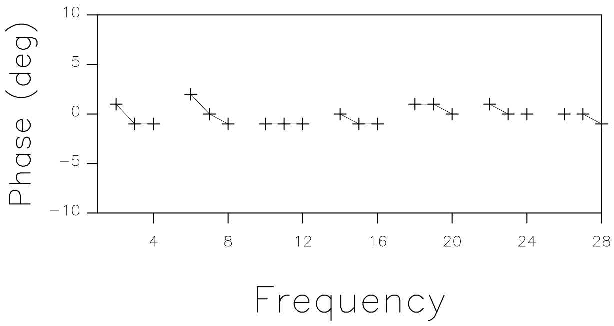

Figure 10.6: This figure shows the data displayed in figure 10.5 after it is fringe fitted, corrected, and time averaged. Note the expanded phase scale.

band delay differences and errors on the phase-cals is shown in figure 10.4. This procedure leaves residual phase and single-band delay errors.

10.2.4.5 Phase-Cal Errors

The phase-cal system used for Mark III observations frequently suffers residual phase errors on the order of 10–20 degrees. If uncorrected, this can degrade the dynamic range of images derived from this data. This effect can be seen in figure 10.4. Calibrator or other high SNR observations can be used to make corrections to the phase-cal values.

10.2.5 Calibrator Sources

Observations of a source which gives strong fringes to all antennas can be used to make delay and rate residual corrections. In particular, corrections to the phase-cals and differences between the multi- and single-band delays can be corrected as well as large clock errors. Phase-cal errors and multi- versus single-band delay differences can usually be determined from observations at a single time and then applied to all data. These calibrator solutions should fit each band independently. Unless there is a large rate offset from a clock (rare these days), the rates from the calibrator fits should be zeroed before applying them to the other data. Small rate residuals are usually caused by atmospheric fluctuations or errors in the assumed source position and are usually not relevant for other sources or times. Figure 10.5 shows the result of applying calibrator solutions to the data shown in figure 10.3.

Figure 10.6 shows the time averaged phase versus frequency plot for the data shown in figure 10.5 after global fringe fitting solutions have been applied. Even after correction by a calibrator, the fringe fitting found a 20 ns offset between the single- and multi-band delays for the baseline shown. Other baselines included in the same global fringe fit had differences between multi- and single-band delays

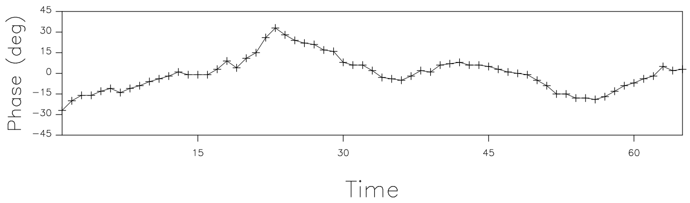

Figure 10.7: This figure shows the effects of using a solution interval in excess of the coherence time. The phase as a function of time for a single baseline is shown after the removal of a single rate and phase residual fitted to multiple baselines using a global technique. The time axis is labeled in arbitrary units but this figure shows about 12 minutes of data at 6 cm. wavelength.

several times this value. The time interval shown was divided into 3 solution intervals. Figure 10.6 demonstrates that the combination of calibration using calibrator source observations and fitting both multi- and single-band delays removes most of the systematic delay problems.

If there are time variable systematic effects left over after the calibration using the measurements from the single time interval, then observations of calibrator sources at other times can be used to remove them. Removing these systematic calibration errors will allow using a smaller window in delay and rate. These fringe fits of calibrator data should be done in the same mode as is to be used on the target source.

10.2.6 Solution Interval

As long as the phase deviations from the model are relatively small, the solution interval may exceed the coherence time. Figure 10.7 shows the effect. The phase fluctuations remaining are due to variable atmospheric delays and are large enough to affect the image derived from this data but small enough for the fringe fitting procedure to work. These phase fluctuations will be removed at a later stage if the self-calibration process has a sufficiently short solution interval.

10.2.7 Weights

If the visibility amplitudes have been calibrated and the associated weights calibrated as well, then most of the determination of the weights used in fringe fitting is done. If the array being used is homogeneous (e.g., VLBA) then the relative weights are comparable and may not need adjustment. However, for heterogeneous arrays relative weighting may be needed. If the visibility weights have not been calibrated then it is possible to determine relative antenna weights based on their sensitivities. The relative antenna weight should be the inverse of the square of its sensitivity expressed in Jy. The weights should also be modified to account for resolution if a source model is available.

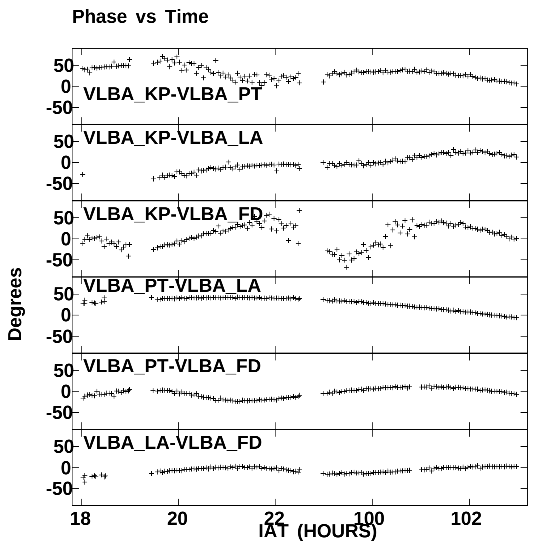

Figure 10.8: This figure shows the residual phases on a number of baselines due to the structure of the source being observed.

10.2.8 Source Model

Structure in a source will cause baseline, frequency and time variations of the interferometer phase. As global fringe fitting tries to use these baselines coherently, source structure can cause problems. If a reasonable model of the source is available then the Fourier transform of the model can be divided into the data to correct for the source structure. Figure 10.8 shows the phase on a number of baselines without correction for the source model and figure 10.9 show the same data after division by a source model.

One of the effects of using a source model is to reduce the weights on baselines on which the source is resolved so that this data will have less influence on the fitting. If a source is known to be resolved (usually apparent from direct inspection of the data) it may be desirable to downweight the long baselines by using a gaussian model. As such a model has no effect on phases it can be safely used to reduce the weights on long baselines.

10.2.9 Suggested Procedure

This section describes a procedure for fringe fitting data with Mark III type phase-cals, one tone per band. Some of these steps can be eliminated if the the VLBA phase-cals are used successfully. The following procedure is based on a global fringe fitting technique in an antenna based calibration system.

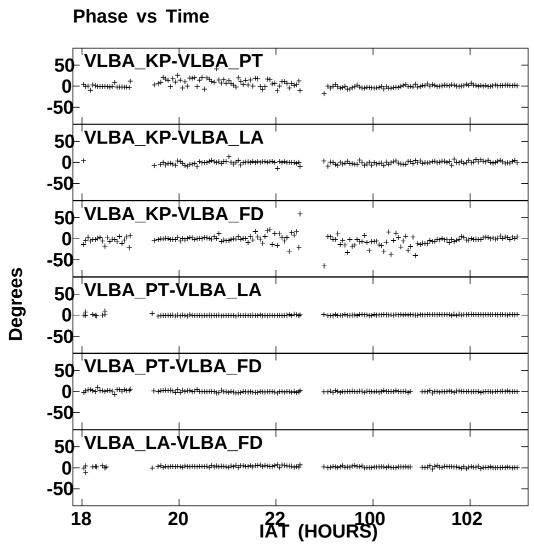

Figure 10.9: This figure shows the phases in figure 10.8 after being divided by a model of the source structure.

1. Sideband phase offset

   If the bands used contain both sidebands then there may be a phase offset between the sidebands. There is a 130 degree offset for baselines with mixed Mark III and VLBA record systems. Data for a calibrator or other strong source can be inspected for offsets. Any such offset needs to be corrected before fringe fitting.

2. Clean up phase-cals (Mark III)

   Check that all phase-cals are present and have reasonable values. If valid phase-cals are not available make sure that the phase-cals are set to zero phase.

3. Amplitude Calibration

   The amplitudes and weights should be calibrated prior to fringe fitting. This simplifies determining the relative weights of the data.

4. Fringe fit on calibrator.

   Select a short segment of data on a calibrator source or other source that gives good fringes to all antennas, bands etc. This data should be fringe fitted independently in each band and polarization. The fitted rates should be set to zero and the solutions applied to all data. This will remove the differences between multi-band and single-band delays as well as the errors in the phase-cals at that time. Since the calibrator source should be fairly strong use a relatively high SNR limit of 10 or more.

5. Fringe fit on target source.

   There are numerous options for fringe fitting sources. For sources with strong signals to all antennas the details of the fringe fitting are not terribly critical. For weak or resolved sources the details are more critical. If a reasonable model of the source is available, either from another frequency or time or from a self-calibrated model from a previous fringe fit it should be Fourier transformed and divided into the data. If a detailed model is not available, an appropriate circular gaussian can be used to crudely modify the weights to account for resolution.

   For sources with strong detections to each antenna or widely spaced frequency bands make separate solutions for each band. The residual fringe rate should be averaged in time per time units (not Hz) before being applied to the data.

   Solutions for sources with weak detections to at least some antennas should combine the data from all bands and determine both multi- and single-band delays. Some experimentation with the possible options may be required to find the best combination for a given source. Minimum SNRs of 5 to 7 depending on the search window will reduce the number of false detections.

   Strongly resolved sources should use only a single baseline to each antenna in the FFT search stage. In this case, additional antenna weight factors may be needed to force the choice of secondary reference antennas.

6. Filter solutions

   If the solutions from the fringe fitting are excessively noisy, contain bad values or occasionally failed due to marginal SNR then it may be desirable to do some filtering. Solutions should all be referred to a common reference antenna before any filtering. Box car smoothing can reduce noise in solutions and median window filtering can remove occasional bad solutions.

7. Iterating fringe fit.

   If the fringe fitting of a resolved source using a point (or gaussian) model failed to get solutions for a significant amount of the data due to low SNR then is possible to take the data successfully calibrated and image it using the usual techniques. The resulting source model can then be used in a subsequent fringe fit with, hopefully, better success.

10.3 Bandwidth Synthesis

There are special problems involved with fringe fitting data with bands widely spaced in frequency. As described in a previous section, the value of the multi-band delay may be ambiguous, especially if the single-band delays are distinct from multi-band delays. Low SNR makes the problem more serious as the delay function will have many large sidelobes which could lead to spurious solutions. In addition, the FFT fringe search can become very inefficient if the frequency space is poorly sampled. For many astronomical applications the value of the multi-band delay is not important but in some applications, such as astrometric or geodetic measurements, the multi-band delay is the fundamental measurement. Accurate multi-band delays may be needed in phase referencing measurements to resolve phase ambiguities.

In many of the cases where widely spaced frequency bands are used, the SNR is relatively high and the sources observed can be detected in each band individually. In these cases, the fringe fitting can be done separately in each band and then the multi-band delays estimated later from the phase residual in each band, possibly subject to constraints from the single-band delays.

References

Alef, W. 1989. "Introduction to phase reference mapping". See Felli and Spencer (1989), p. 261.

Alef, W., & Porcas, R. W. 1986. "VLBI fringe-fitting with antenna-based residuals". *Astron. Astrophys.* **168**, 356.

Cotton, W. D. 1989. "Polarimetry". See Felli and Spencer (1989), p. 275.

Felli, M., & Spencer, R. (Eds.) 1989. *Very Long Baseline Interferometry*, Volume 283 of *NATO ASI Series C*. Dordrecht: Kluwer.

Moran, J. M. 1976. "Very Long Baseline Interferometric observations and data reduction". In M. L. Meeks (Ed.), *Astrophysics—Part C: Radio Observations*. Volume 12C of *Methods of Experimental Physics*. New York: Academic Press, p. 228.

Schwab, F. R., & Cotton, W. D. 1983. "Global fringe search techniques for VLBI". *Astron. J.* **88**, 688.

Walker, R. C. 1989a. "Calibration methods". See Felli and Spencer (1989), p. 141.

Walker, R. C. 1989b. "Principles and practice of VLBI". In R. A. Perley, F. R. Schwab, and A. H. Bridle (Eds.), *Synthesis Imaging in Radio Astronomy*. San Francisco: Astronomical Society of the Pacific, p. 355.