

Very Long Baseline Interferometry

R.C. Walker

National Radio Astronomy Observatory, Socorro, NM, USA, 87801

Abstract. Very Long Baseline Interferometry is the branch of radio interferometry that uses antennas that have no direct link for data and/or clocks. The baseline lengths are limited only by the size of the Earth, and not even that for space VLBI. The resolution is very high — of order a milliarcsecond. This allows imaging of galactic objects with sizes of about an astronomical unit and of extragalactic objects, even at high redshift, with sizes of about a parsec. The wide separation of antennas makes stabilization of phases difficult because of the need for a high accuracy model and because of the very different atmospheres over each antenna. But great progress has been made in the last few years with the advent of the VLBA. For many VLBI observations now, the data reduction techniques and sensitivity are similar to those of connected interferometers. In this chapter, the process of making a VLBI observation is described with emphasis on those areas that differ from connected element interferometer practice. There is special emphasis on techniques for dealing with phase.

1. Introduction

Very Long Baseline Interferometry (VLBI) has traditionally been distinguished from other forms of radio interferometry by the absence of any direct, real-time link between the stations for either local oscillators or for the received signals. Highly accurate atomic frequency standards are used to maintain coherence of local oscillators and wide bandwidth tape recorders are used to record the sampled data. Without the requirement for a direct link, the interferometer elements can be anywhere, allowing the use of baselines nearly as long as the diameter of the Earth, or longer with the use of satellites. The resolution of such an instrument is typically in the neighborhood of a milliarcsecond (mas) — the size of a dime in Los Angeles as seen from Washington D.C. This gives VLBI arrays the highest resolution of any telescopes used in astronomy. Even for sources at the edge of the observable universe, it allows images to be made of regions small enough to change structure on human time scales. Thus a VLBI imaging instrument can be thought of as a movie camera more so than most other astronomical instruments. The high resolution also allows astrometry and geodesy to be done with greater accuracy than with any other technique. The price paid for the use of such long baselines is relatively unstable phases. Also, target sources must have brightness temperatures of a million degrees or more to be detected, so observations of thermal sources are generally not possible.

For a good sampling of the kinds of science done with VLBI, see the proceedings of IAU Colloquium 164 (Zensus, Taylor, and Wrobel 1998). Good references for VLBI technique include the fundamental reference book by Thompson, Moran, and Swenson (1986: hereafter TMS) and the proceedings of two VLBI summer schools, one in Bologna in 1988 (Felli and Spencer 1989) and the other in Socorro in 1993 (Zensus, Diamond, and Napier 1995).

2. VLBI Science

One of the major areas of VLBI research is studies of jets and associated phenomena around compact objects such as those in active galactic nuclei (AGN) and collapsed stars. Collapsed objects provide the necessary conditions for the generation of extremely hot gases, often concentrated in jets with velocities approaching the speed of light. One of the best known results of early VLBI is the observation of apparent faster-than-light motion in such jets — thought to be a projection effect with a relativistic jet moving near the line-of-sight.

Another major area is studies of molecular clouds, star forming regions, old stars, and accretion disks utilizing natural maser emission from a variety of molecules. Such sources can be extremely bright, not because they are hot — typical temperatures are a few hundred degrees — but because inverted quantum level populations are allowing amplification of incident radiation rather than absorption. Such sources can have equivalent brightness temperatures of 10^{15} K or more. Perhaps one of the finest VLBA results so far was the observation of water masers in a beautiful keplerian disk in the nucleus of NGC 4258 (Miyoshi et al. 1995) — a result that provides some of the strongest evidence yet found by any method that black holes exist in the centers of active galaxies.

By being very careful about the geometric model and making accurate measurements of the relative arrival time of wavefronts at different antennas, the VLBI technique can be used to measure absolute source positions to tenths of a milliarcsecond and station locations to a few millimeters (Sovers, Fenselow, and Jacobs 1998). The source catalogs so generated form the basis for the most accurate reference frames used in astronomy (Ma et al. 1998). The station measurements are used to study plate tectonics, Earth rotation, polar motion and other phenomena of interest to geophysics. While GPS is now achieving accuracies close to VLBI and is displacing VLBI for geodesy for cost reasons, VLBI is still the only geodetic technique in use that has something approximating an inertial reference frame — the distant quasars. Therefore only VLBI can track the long term orientation of the spin axis and rate of rotation of the Earth.

3. VLBI Systems

The world of VLBI is far from monolithic. There are nearly 60 different antennas in the stations catalog used by the VLBI scheduling program SCHED, and that catalog is not claimed to be complete. There are several tape recording systems in use and at least one, and sometimes more, correlators per tape system. Some of the most common systems are sufficiently compatible to allow joint use and there are facilities for translating tapes between some of the types when necessary.

There are some important groupings of antennas involved in VLBI. The most obvious is the VLBA, which is a fully integrated instrument with 10 identical antennas, central operations, maintenance, and management, and a 20 station correlator in Socorro, NM (Napier et al. 1994). It replaced the now defunct U.S. VLBI Network. The European VLBI Network (EVN) is the most active group of separate observatories involved in astronomical VLBI. Much of the EVN data are processed at the MPIfR in Bonn, but that will change soon when

the 16 station correlator under construction at the Joint Institute for VLBI in Europe (JIVE) is completed. The geodetic community has a number of antennas devoted to VLBI that are observing nearly as much the VLBA. Their data are processed at the Washington Area Correlator, Haystack, and Bonn. The Australia Telescope has a VLBI component and correlator with the distinction that it can see the southern sky. A number of antennas in the eastern hemisphere are part of the fledgling Asia Pacific Telescope (APT). Russia is building the Quasar Network. NASA's Deep Space Network has an active VLBI program related to spacecraft navigation. Observatories involved in millimeter VLBI have formed the Coordinated MM VLBI Array (CMVA). Many of the above groups, along with unaffiliated antennas, can observe together in "global arrays". The most notable examples are the joint VLBA and EVN global sessions of 2 to 3 weeks that occur 4 times per year.

Early VLBI systems were based on computer or video tape recorders and could utilize bandwidths from a few hundred kHz to 2 MHz. The recording systems in common use today have bit rates in excess of 100 Mbps. The current VLBA and Mark IV systems used by the VLBA, the EVN, the CMVA, and most geodesy stations, are evolutionary enhancements of the older Mark III system. All are based on the same 1 inch instrumentation tape recorders. Recordings made in appropriate modes on one system can be played at the correlators for the others. The VLBA normally records at 128 Mbps, although the tape recorders can handle 256 Mbps each and there are two at each site that can be used in parallel. The Mark IV recorders have a faster maximum bit rate per track and will eventually have two heads for a maximum bit rate of 1 Gbps.

Three other recording systems have appeared. Two are the similar K4 and VSOP systems from Japan that are based on instrumentation cassette recorders. The other is the S2 system from Canada that is based on a stack of 8 commercial grade VHS recorders and is used at a variety of observatories around the world, including on the VLBI segment of the Australia Telescope. The VSOP and S2 systems were developed partly in support of the space VLBI projects that will be discussed by Ulvestad (these proceedings, p. ??).

4. VLBI Phases

VLBI observations are not fundamentally different from observations on connected interferometers like the VLA. Most of the same signal handling and data processing steps occur for both, although the mix of steps done automatically by the on-line systems and steps that must be done by the observer is somewhat different. Operationally, VLBI is distinguished by the use of tapes and delayed correlation. But by the time the observer gets the data, that should not be important.

For the observer, one of the most important differences between VLBI and typical connected interferometry is that the phases in uncalibrated VLBI data are less well behaved. This is because the lines of sight from each antenna pass through totally uncorrelated atmospheres (ionosphere plus troposphere). Also, many subtle geophysical and Earth orientation effects are too small to cause problems for the connected interferometers, but are large compared to the resolution of VLBI. With the typical correlator models in use prior to the VLBA,

the residual phase, delay, and rate variations were usually so high that phase referencing was not possible, or at least required heroic efforts. The common wisdom was that, to detect a source with VLBI, it must be detected within a coherence time, which is usually a few minutes. That detection must involve a solution for delays and rates, in addition to phases — a process called fringe fitting.

Many of the difficulties with the phases can be overcome by using accurate geometric models and by careful observing. This can greatly simplify the use of traditional VLBI calibration methods and makes possible the use of calibration techniques for many VLBI observations that are more typical of connected interferometry. In particular, phase referencing becomes possible, which allows derivation of absolute positions and allows detection of sources so weak that they can only be seen in images based on hours of integration.

In the rest of this section, the special procedures developed to deal with VLBI phases will be discussed. Then, in the following section, all the steps involved in a VLBI observation will be described with emphasis on those areas that are different for VLBI.

4.1. The Model

A correlator must cross multiply signals, from different antennas, that correspond to the same arriving wavefront. But the antennas are at different distances from the source so the wavefront arrives at different times (the delay). Also they are moving at different speeds along the direction to the source, causing different doppler shifts (the fringe rate). An estimate of these time and rate offsets is removed in the correlator hardware. The estimate is derived using the correlator model. Table 1 gives an overview of most of terms that affect VLBI delays at the level of a cm or more. Most are included in the VLBA correlator model. The observer doesn't really need to know these details, but might find it interesting to see some of the subtle effects that are involved.

The VLBA correlator uses a delay model and accurate source and station positions from the geodetic community and clock offsets and rates based on long term GPS measurements. Model errors, after removal of a clock offset and rate, should normally be dominated by the atmosphere. This significant increase in model quality over previous practice has stabilized the raw phases to the extent that delays, rates, and often phases can be calibrated using calibration sources without having to improve the model in postprocessing. The traditional need to fringe fit all sources and to obtain detections within a coherence time of a few seconds or minutes still applies only at extreme frequencies where the atmospheric effects are too strong to allow referencing.

Other than atmosphere and clock offsets, it is possible for the model to be good to a few centimeters at worst for most stations. This should introduce only a very small number of turns of residual phase over many hours. Clock offsets are the largest term in the residuals from the correlator, with typical values of around 50 ns and much higher values possible. These are easy to correct to the nanosecond level because they don't depend on source position and, other than a possible linear drift, are fairly stable (although short term variations are of concern in geodetic/astrometric observations). The atmosphere is another matter. Because of the $\sec(Z)$ scaling, it can vary a lot between

Table 1-1. Terms of a VLBI Geometric Model ^a

Item	Approx max Magnitude ^b	Time scale
Zero order geometry.	6000 km	1 day
Nutation	$\sim 20''$	< 18.6 yr
Precession	~ 0.5 arcmin/yr	years
Annual aberration	$20''$	1 year
Retarded baseline	20 m	1 day
Gravitational delay	4 mas @ 90° from sun	1 year
Tectonic motion	10 cm/yr	years
Solid Earth Tide	50 cm	12 hr
Pole Tide	2 cm	~ 1 yr
Ocean Loading	2 cm	12 hr
Atmospheric Loading	2 cm	weeks
Post-glacial Rebound	several mm/yr	years
Polar motion	$0.5''$	~ 1.2 years
UT1 (Earth rotation)	Random at several mas	Various
Ionosphere	~ 2 m at 2 GHz	seconds to years
Dry Troposphere	2.3 m at zenith	hours to days
Wet Troposphere	0 – 30 cm at zenith	seconds to seasonal
Antenna structure	< 10 m. 1cm thermal	—
Parallactic angle	0.5 turn	hours
Station clocks	few microsec	hours
Source structure	5 cm	years

^aAdapted from Sovers, Fenselow, and Jacobs 1998^bFor an 8000 km baseline, 1 mas \leftrightarrow 3.9 cm. \leftrightarrow 130ps

sources, especially at low elevations. And both the ionosphere and the wet component of the troposphere are highly variable in time. See TMS for an extensive discussion of these contributions and lecture 28. Just to give an idea of the magnitude of the effects, the wet troposphere contributes between very little and about 1 ns of excess delay at the zenith, independent of frequency. The ionosphere can contribute as much as 60 ns or more at 1.4 GHz and scales with ν^{-2} . It can vary by 5% on short time scales and by an order of magnitude between day and night. The VLBA correlator only uses a seasonal and latitude dependent troposphere model and no ionosphere model. Currently there are also some problems at the 20 mas level or so, soon to be fixed, due to a poor nutation model (IAU1980) and out-of-date station coordinates.

4.2. Fringe Fitting

When there is a model error, the measured cross correlation will peak at a value of delay that is offset from the expected value. The interferometer phase, ignoring instrumental offsets, is frequency times delay, so the delay offset will show up as a phase slope as a function of frequency. For example, a delay error of 125

ns will cause a slope of a full turn across a typical VLBA baseband bandwidth of 8 MHz (Usually 4 or 8 baseband channels are used. See Section 4.2. for a definition of a baseband channel.). If the delay error changes with time, so will the phase and there will be a slope of phase with time, called a “fringe rate”. Such phase slopes can limit the sensitivity of an interferometer if they limit the bandwidth and time of the interval of data that can be used for detection of a source. Smaller slopes can limit the dynamic range of an image if allowed to degrade the amplitudes when data are averaged. This is because such degradations cannot be described by a product of antenna gains that applies to all baselines — they do not close and cannot be removed by standard self-calibration. For connected interferometers, such phase slopes are typically so small that they are not a concern. For VLBI, mainly because of clock and atmosphere model errors, they cannot be ignored.

The procedure used to find and correct phase slopes is called “fringe fitting”. A fringe fit is nothing more than a self-calibration (lecture 10) that includes not only amplitude and phase gains, but also the derivatives of the phase gains. Normally only the first derivatives are considered, although the second derivative in time is determined for some space VLBI applications. Most fringe fitting programs first transform the data, using FFTs, from amplitude and phase as a function of frequency and time to amplitude and phase as a function of delay and fringe rate. In a multi-baseline fit, the phases from various combinations of baselines connecting a reference antenna to any other can be stacked before the transform to gain sensitivity. Then the peak amplitude point in this space is found. That provides a starting guess for a least squares solution. The fit can be on a per-baseline basis or, like self-calibration, can be global in the sense that all baselines to a station contribute to finding one set of phase, delay, and rate values for the station. Details of fringe fitting are discussed in considerably more depth in Cotton (1995) and references therein.

Weak Sources: Fringe fitting on weak sources can be tricky. The FFT step must find the right peak, or the results will be meaningless. But, with open delay and rate windows, there are a large number of possible values and a detection can only be believed if it is many sigma. Both mm VLBI and space VLBI observers face this problem, the first because of the very rapid phase fluctuations at mm wavelengths and the second because the spacecraft cannot switch back and forth to a calibrator and the position of the orbiter is not as well known as that of a ground station. Both groups have developed sophisticated tools to deal with the problem (Ulvestad, these proceedings, p. ??; Rogers, Doelman, & Moran 1995; Lonsdale and Doelman 1998). It is also possible that the the ionosphere will, at times, cause such rapid phase fluctuations at low frequencies that careful fringe fitting is needed. However, with the large primary beams at low frequencies, it might be possible to find a source in the beam that can be used as a calibrator.

Accuracy Requirements: The accuracy required of fringe fit results depends on the goals of an observation. For geodetic experiments, the delays and rates from the fringe fit are the ultimate observable from which the accurate geometric results will be obtained by least squares fit. It is important to get good, high signal-to-noise ratio fit results and to fit all data. Typically after the fit is done, the actual visibility data are no longer used. For imaging experiments, the fringe

fit results are used to flatten the phase slopes. For spectral line observations, this allows phase calibration to be transferred between spectral channels and also may allow time averaging, or at least allow the use of longer self-calibration solutions. Often spectral line observations are of strong sources with high SNR, so the best possible calibration is required. For continuum observations, the purpose of the fringe fit is to allow the use of large solution intervals in time and frequency in subsequent self-calibrations and to allow data averaging in both time and frequency to reduce data set sizes. The main requirement is that the residual slopes are sufficiently small that any amplitude loss in averaging is not large enough to degrade the images.

The amplitude reduction suffered when averaging with an incorrect delay is shown in Figure 1 for a variety of final average bandwidths. The curves are simply the tops of the $\sin(x)/x$, or sinc, functions that describe averaging with phase slopes. For a project with a peak to off-source rms dynamic range target in the tens of thousands (not uncommon — 10^5 has been reached on the VLBA), such losses should be kept below about 0.3%, the level marked with the dashed line. For such projects, averaging baseband channels together is not desirable so, for typical cases, the averaged bandwidth will be 8 MHz. For that case, the delay accuracy from the fringe fitting should be kept better than about 5 ns. For very weak sources where dynamic range will not be much of an issue, it may be adequate to keep the averaging losses below 5%. But one may wish to average several baseband channels together to gain sensitivity. As can be seen from Figure 1, for a total bandwidth of 32 MHz, the 5 ns criterion again applies. For a significant portion of VLBA observations, 5 ns is greater than the full range of delays seen at any antenna. But it may be less than the scatter of delay solutions on weak sources. Therefore, it is probably not just possible, but advisable, to use the fringe fit results from one or more calibrator scans to set the delays for the whole experiment. The main exceptions are likely to be at low frequency, where the ionosphere can contribute large, variable delays, and at any frequency when some apriori parameter is poor — such as a source position with an error of more than about 30 mas.

Correlator Averaging Losses: There is an area of concern about averaging losses that is specific to the VLBA. The correlator, for continuum projects, forms spectra of 128 or 256 channels per baseband, then averages, usually to 16. If there is a large delay error — and any clock offsets still apply at this point — there will be some amplitude loss in this average. This loss does not occur in correlators that don't average spectral channels on-line, which includes most that use the XF architecture (lecture 4). The postprocessing programs know how to correct for this effect if they know the delay with adequate accuracy. For common observing modes, this places less stringent requirements on the fringe fit results than the need to avoid losses in the full baseband average. But, if the correlator averaging produces spectral channels of 1 MHz width or greater, or if there are abnormally large residual delays during correlation, it is possible that the delay accuracy requirement will be driven by the needs of this correction. The only case where this is at all likely to occur is when the 16 MHz baseband channels are used and the on-line averaging is set to the usual default of 16 output channels. It is best to be sure that the correlator delivers output channels of no more than 500 kHz width. Note that the equivalent of the fringe fit for

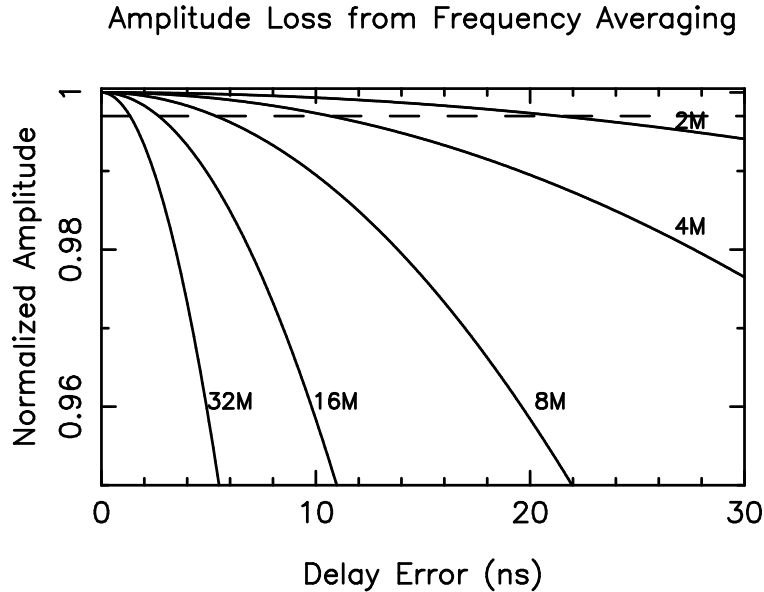


Figure 1-1. The reduction in amplitude for data averaged in frequency the presence of a delay error. The separate lines are for different averaged bandwidths.

delay could be done with a bandpass calibration, which would find and flatten any phase slopes. But this would not leave a record of the delay offset at the time of correlation and the correction for the correlator averaging loss would not be done. This might actually be acceptable if the apriori clocks are very good or if the dynamic range goals of the observation are modest.

Fringe Rates: The need for fringe fitting for the fringe rates should probably be treated separately from the need to fit for delays. At frequencies above where the ionospheric contribution can be large, the fluctuations in delay will be independent of frequency. Their impact just depends on bandwidth which is usually determined by the tape system in VLBI. Just because the phases are much less stable at the highest frequencies does not mean that the delays will be any worse than, say, at 8.4 GHz, so the effort required to determine delays should be more or less independent of frequency. The fringe rates are another story since they scale with frequency. At intermediate frequencies such as 5.0 and 8.4 GHz, with a good model and sources that can be detected in a few minutes or less, it is probably not necessary to fit for fringe rates. Just use self-calibration to solve for phases and use linear interpolation when the calibration is applied. At lower and higher frequencies, removal of rates may be needed. But it is very likely that this can be done using a calibrator near the target source. In all but the most extreme cases, it should not be necessary to fringe fit on a weak source if the observation was scheduled with adequate calibration.

Note that averaging too long in time in the presence of a fringe rate can cause non-closing errors by essentially the same mechanism as averaging in frequency in the presence of a delay error. Figure 2 shows the losses as a function of average time and residual fringe rate, assuming a linear phase slope. A rule




Figure 1-2. The reduction in amplitude for data averaged in the presence of a fringe rate error. The separate lines are for different average times in seconds. Note that losses due to a fringe rate that is known, probably as a result of a fringe fit, to the AIPS routines that apply calibration will be corrected.

of thumb would be to keep the average time below 10 seconds if the residual fringe rates can get as high as 4 mHz, but actual reasonable values will depend on details of the observation. The calibration routines do know how to correct for the amplitude loss caused by a high fringe rate as long as that fringe rate has been determined with a fringe fit.

4.3. Self-calibration

Self-calibration was covered in depth by Cornwell (these proceedings, p. ??) and will only receive cursory attention here. As a quick reminder, self-calibration is the process of using the best available source structure information to improve the antenna phase and, often, amplitude gains. The data set, calibrated with the improved gains, is then used to produce a better image. The process can and should be iterated if the starting model is poor. In fact, such an iterative loop is the standard VLBI imaging method and will be discussed more later.

Self-calibration is based on the assumption that all gain variations can be described as antenna based — they affect all baselines to an antenna equally. Since there are $N(N-1)/2$ baselines with N antennas, and each baseline gives independent information about the source, there is enough information to determine both the gains and the source structure. But as mentioned before, care must be taken to avoid doing something to the data that will cause non-closing errors — amplitude or phase offsets that are not due to source structure and cannot be described by one gain per antenna.

VLBI imaging is totally dependent on self-calibration. First of all, fringe fitting, which is required to make final model adjustments to the clocks, if nothing else, is just a variant of self-calibration. Dynamic ranges beyond about 100,

measured as the ratio of image peak to off source RMS, are extremely hard to obtain with phase referencing. If the source is strong enough to self-calibrate, then it is essentially always advantageous to do so, at least for the phases. With self-calibration, dynamic ranges of 1000 are the norm and 10^5 has been reached on the VLBA. Even in cases where phase referencing is required, it is likely that the calibrator will have to be imaged to remove structure effects, since nearly all sources are resolved to VLBI. That imaging step will rely on self-calibration.

4.4. Phase Referencing

In traditional VLBI practice, it has been necessary to detect a source on the baselines to each antenna in an integration time of less than, or equal to, the coherence time. A gain in sensitivity of between 1 and 2 orders of magnitude can be achieved if data from the whole array for the whole time of the observation can be used. This requires calibrating the phases based on something other than the target source. Connected interferometers achieve this goal by observing a nearby calibrator, self-calibrating on that calibrator (that's not usually what it's called, but really that is what it is), and then transferring those phase corrections to the target source. For this to work, it is necessary that the model errors change slowly across the sky so that they are similar on calibrator and target. It is also necessary that the errors change on time scales long compared to the switching time between calibrator and target scans. With the good model on the VLBA correlator, these conditions are often met and phase referencing is possible, even getting to be common. Phase referencing has been done on the VLBA at frequencies between 1.6 and 43 GHz using calibrators less than about 6 degrees from the target source and with intervals of between 15 seconds and a few minutes between calibrator scans.

Phase referencing is limited by whatever position or time dependent model errors remain after correlation and initial calibration. Errors in geometric parameters, such as source and station positions, will degrade phase referenced phases by up to the total error times the calibrator/target separation in radians (ie, a factor of 0.1 for a separation of 5.7 degrees). Atmospheric effects scale with the secant of the zenith angle, which, unlike the connected interferometer case, can be very different at different array elements for VLBI. Any errors in the steady ionospheric or tropospheric model induce errors scaled by the difference in $\sec(Z)$ between calibrator and target, a factor that can get rather large at low elevations even with close pairs. Atmospheric terms can also vary rapidly in both space and time which will degrade phase referencing using either too long a switching time or too widely separated sources.

The correlator model, ignoring the atmosphere and clocks, should be good to a few cm at worst, which corresponds to less than 1 to a few turns of phase across the sky (soon to be fixed nutation and station catalog problems cause the current VLBA model to have somewhat larger errors). Clocks are not much of a problem since they do not change with pointing position and can be well calibrated with phase referencing. Also, short term clock variations are generally smaller than atmospheric variations with modern maser frequency standards. But the atmosphere cannot be ignored. It is the dominant source of errors. To understand when phase referencing can be done and with what parameters, we must understand the atmosphere. Unfortunately, both important components,

the ionosphere and wet troposphere, can vary tremendously with location and time, so any estimates of how well phase referencing will work must necessarily be very approximate.

Troposphere: The variations in the wet component of the troposphere are especially troublesome for mm interferometry on any but the very shortest baselines. They will be discussed in depth by lecture 28 based on VLA and other measurements. What has been learned on the VLA about phase fluctuations can also be used to understand VLBI cases. VLBI baselines are always far beyond the outer scale length of atmospheric turbulence. But the longest baselines on the VLA are also beyond the outer scale length and so should not be much different. The characterizations of phase differences as a function of baseline length can be used to understand the differences between calibrator and target lines of sight. Those lines of sight diverge from the antenna, but are actually rather close until beyond the wet troposphere scale height. For an 0.1 radian calibrator/target separation, the lines of sight are only about 100m apart at the zenith at the scale height of about a kilometer. The separation is larger for lower elevation angles. Because of this small separation, except at low elevations, spatial effects are unlikely to be much of a contributor to the relative phase fluctuations. More important are the time fluctuations caused by the turbulence pattern blowing over each antenna. VLA data (Carilli and Holdaway, 1997) beautifully show how the fluctuations increase with baseline length as expected. But, if calibration is done on some time interval, the phase fluctuations increase with baseline in the same way until the baseline length is about equal to the distance the wind blows between calibrations. Then the rms phase fluctuations are constant for all longer baselines. There is no reason to believe that the fluctuations would not remain the same, with calibration, all the way out to VLBI baselines. Carilli and Holdaway find that calibration is good to about 5 degrees of phase with a 20 second calibrator cycle and to about 20 degrees of phase with a 300 second cycle time for observations at 43 GHz.

The above discussion would suggest that fluctuations due to the troposphere can be kept small by using a fast calibration cycle. But there is one complication that is far more serious for VLBI than for the VLA. That is that two VLBI antennas will see a source at different elevations. Any error in the total zenith atmospheric delay, not just the fluctuations, and in the mapping function that describes how the delay increases with zenith angle, will show up in the baseline phases. At low elevations, this effect will likely dominate the errors in phase referencing. Note that the same argument applies to both ionosphere and troposphere. Such an error is systematic so it does not beat down with integration.

A future research topic that might help VLBI phase referencing relates to the use of water vapor radiometers to measure the water vapor along the line of sight. Those measurements can be used to correct the interferometer phases. This is currently an area of active development on the mm interferometers and the next generation 22 GHz receivers for the VLA are being built so that they can be used as radiometers.

Ionosphere: The ionosphere is the other significant atmospheric contributor. It can vary by an order of magnitude between night and day. It can also vary

Table 1-2. Maximum Likely Ionospheric Contributions to Delay and Rate. ^a

Frequency GHz	Max Delay ns	Min Delay ns	Max Rate mHz	Min Rate mHz
0.327	1100	110	12	1.2
0.610	320	32	6.5	0.6
1.4	60	6.0	2.8	0.3
2.3	23	2.3	1.7	0.2
5.0	5.0	0.5	0.8	0.1
8.4	1.7	0.2	0.5	0.05
15	0.5	0.05	0.3	0.03
22	0.2	0.02	0.2	0.02
43	0.1	0.01	0.1	0.01

^aAdapted from Thompson, Moran, and Swenson 1986

with the solar cycle. Table 2 gives the maximum likely ionospheric contribution to group delay and fringe rate at 60 degrees elevation for the VLBA frequency bands for quiet (night) and active (day) times. The table is adapted from a table in TMS. Note that the delay contribution scales with frequency squared while phase effects scale with frequency. One quick conclusion from the table is that the ionosphere is a very big effect at low frequencies, and it still can contribute several turns of phase even at 43 GHz. Like the troposphere, the ionospheric contribution scales with $\sec(Z)$ above about 10 degrees elevation. Below that, it starts to level off because of the curvature of the Earth and the fact that the main ionospheric effects occur at an altitude of 300-500 km. Because the delay can be so large, if it is not modeled correctly, the $\sec(Z)$ dependence can lead to very large differences between the model errors on calibrator and target which can be a big problem for phase referencing. The ionosphere is subject to various short term variations including traveling waves. It can vary by 5% due to traveling ionospheric disturbances (TIDs) on time scales ranging from 10 minutes on up. Very short term variations also exist.

The VLBA does not use an ionospheric model of any sort. At lower frequencies, this clearly can be by far the dominant source of model errors. It is likely that phase referencing involving switching the pointing positions of the antennas between calibrator and source will not be possible at 327 and 610 MHz. The best hope of any sort of phase referencing at those frequencies may well be to find a calibrator in the beam when observing a target. Given the large primary beams at those frequencies, and the high density of sources, that might be possible if a large enough fraction of sources are compact. At 1.4–2.3 GHz, phase referencing will not be easy except under low ionospheric conditions, but it has been demonstrated to be possible. Lack of an ionospheric model could significantly degrade phase referencing at even higher frequencies, especially at low elevations.

Global ionospheric models are becoming available from the GPS community. An attempt will be made to apply such models to correct VLBA data in the hopes of improving phase referencing. This is still a research topic, but it has significant promise. GPS ionospheric data is also available directly from dual frequency receivers at some sites. If the global models prove to be of insufficient accuracy, perhaps the local data will be helpful. The ionosphere can also be measured by comparing delays, or with care, phases measured at two well separated frequencies. The VLBA is equipped with special optics that allow simultaneous observations at 8.4 and 2.3 GHz. The geodetic community uses such dual frequency observations to remove the effects of the ionosphere from their geodetic and astrometric observations.

An Example: Figure 3 shows an example of phases on both calibrator and target before and after phase referencing. Prior to referencing (top panel), both sources showed phase excursions of several turns, which is typical. This phase winding is due to imperfect models. But the model errors are clearly similar on the two sources because the phases track well. The bottom panel shows what happens when the phases of one source are used to calibrate both. The reference source phases go to zero and the target source shows steady deviations from zero of a few tens of degrees. Such deviations represent some small model error, possibly a source position error. These are test data on two strong sources so it is easy to see how well the phasing worked. The scans were 3 minutes long and the calibrator-target separation was 1.8° . The observations were at 15 GHz.

Scheduling Concerns: Just how should phase referencing observations be scheduled? First, a calibrator needs to be found as close as possible to the target source. It should have adequate flux density to provide reasonable phases during the calibration scans — see the chapter on sensitivity for how to determine this. A few hundred mJy of correlated flux density on the longest spacings is desirable. Weaker sources have been used. A major calibrator survey is in progress that will provide a total of about 3000 calibrators with positions good to a few mas or better and with known visibility functions (Peck and Beasley 1998). That should provide a calibrator that is, on average, about 2 to 3 degrees from the target. Separations up to 5 or 6 degrees can work, but expose the observations more strongly to model errors. The time between calibration scans should be between 15 to 30 seconds at 43 GHz and 2 to 3 minutes at 8.4 GHz (Beasley and Conway, 1995). At lower frequencies, longer on-target times are sometimes used, but the risk of encountering problems with turn ambiguities and with ionospheric fluctuations rises. Generally, phase referencing will be easier with shorter times between calibrators, but the total on-target integration time will be reduced if too large a fraction of the time is spent calibrating. It is highly advisable to include a few scans on a phase reference check source in the schedule. This is a source at a similar separation from the main calibrator as the target, but that is itself a potential calibrator. It can be used to check the quality of the phase referencing on that particular day.




Figure 1-3. An example of phase referencing. The top panel shows raw phases on 3 baselines. The array was cycling between 2 sources, 3C273 and 1222+037, which are about 1.8° apart, with scans of 3 minutes duration. The frequency is 15 GHz. The top panel shows the raw phases with the phase wrapping that indicates that the model is not perfect. But the phases for the two sources track each other well. The bottom plot shows what happens when the phases from both sources are calibrated using the data from one.

5. The Life History of a VLBI Observation

This section is an overview of the steps involved in a VLBI observation from scheduling to imaging. While most of what is presented applies to any VLBI observation, it applies most directly to the case of the VLBA. Special attention is paid to those areas where VLBI practice differs from that typical for connected interferometers. The VLA is used as a concrete example of a connected interferometer.

5.1. Scheduling

Scheduling a VLBI observation involves specifying the individual observation scans and enough information to set up the hardware to the desired configuration. The main VLBI scheduling programs are SKED (for geodesy), PC_SCHED (for Mark III), and SCHED (for nearly everything else including VLBA and Mark IV). Some items to keep in mind while scheduling are:

- Include two or more fringe finder scans on strong sources. Without this, if there are any problems at the correlator, debugging can be very difficult.
- Include at least one or two scans on an “amplitude check source” that is strong and has sufficiently simple structure to be easy to model or image. This source will be used to bring the amplitude calibration of all antennas and basebands to the same scale.
- If observing a weak source, include a nearby calibrator to serve as a fringe rate and delay reference and possibly as a phase reference. If phase referencing, these scans should be at most three minutes or so apart — less at low and high frequencies. For just delay and rate, they can be more widely spaced.
- If phase referencing, include a phase reference check source. This is a calibrator near the main phase calibrator that can be used to check the quality of the phasing on that day.
- If doing spectral line observations, include a bandpass calibrator. It is likely that this will also be the amplitude check source and/or the fringe finder.
- If observing polarization, obtain good parallactic angle coverage on a calibrator at every station to determine the system polarization. The fringe or phase calibrator may serve for this. It is also possible that one observation of an unpolarized source will serve for this.
- If observing polarization, observe something of known polarization to calibrate the polarization position angle.
- If not using automatic tape allocation, which is only available on the VLBA, group scans in units of the length of a tape pass for efficient tape usage. In any case, do not exceed your total tape allocation.
- If using stations with VLBA or Mark IV tape systems and only one tape drive (anything except the VLBA and VLA), allow 10–15 minutes for each tape change, following the guidance provided by the station or network.
- Allow occasional gaps of 2 minutes or more so that the tape systems can do readback checks of recorder health. One every hour or two is adequate.

5.2. Observing

VLBI observations are made based on preset schedules distributed to the antennas ahead of time. The VLBA antennas are usually in contact with the AOC in Socorro, but they do not have to be. For observations involving non-VLBA antennas, there is usually no communication between the antennas during the observations. While it might be possible to make interactive changes to the schedule, it would be difficult and not obviously productive. Therefore, unless the observer is helping with the data acquisition at one of the sites, he/she is usually not involved during the actual observations.

There is very little difference between the antennas used for VLBI and for connected interferometry. In fact, connected element antennas can be, and often are, used for VLBI. Both types focus the radiation into a feed, after which it is converted from circular to linear polarization. Then, a low level, broadband noise signal of known amplitude is injected with some duty cycle. This is used to calibrate the system temperature, which is needed in order to convert correlation coefficients into correlated power. The signals are amplified by low noise amplifiers. Then they are shifted to lower frequencies by being mixed with a reference tone (the local oscillator, or LO), which is derived from the stable frequency standard. After this mix, the signal is called an intermediate frequency, or IF. The IFs are sent to the control building on cables which are subject to bending and thermal variations.

A special feature common in VLBI systems, but not yet in connected interferometers (one is being considered for the MMA), is a second calibration signal that is injected along with the noise calibration signal. It consists of a string of very sharp pulses at, usually, one microsecond intervals that are synchronized to the reference signal from the frequency standard. This has the effect of generating tones of well established phase at one MHz frequency intervals. These tones are subject to all the same instrumental phase variations as the astronomical data. They are detected after the data are digitized and are used to calibrate variations in the delay and phase due to the instrumentation. The VLBA is sufficiently stable that most users don't actually use these tones, but they are valuable for geodesy and as a system integrity check.

The frequency standard is usually a hydrogen maser that is located in the control building. Its output signal has to be sent to the location of the mixers and pulse cal generator, which are near the receivers high on the antenna. The cable involved can change electrical length as it changes temperature or as it is bent when the antenna pointing direction changes. The cable cal system measures the electrical length of this cable and allows corrections to be made. The VLBA, the VLA, and many other systems have cable cals. A difference is that the VLA applies the cable cal on-line while VLBI users must apply it in post-processing. As with the pulse cal, most non-geodetic VLBA users choose to trust the stability of the VLBA and do not apply the cable calibration.

One or more IFs, one for each polarization and perhaps additional ones if multiple frequencies are observed simultaneously, are sent to the VLBI racks in the control building. There each IF is split and sent to several baseband converters ("BBC", which is VLBA terminology; they are called "video converters" in Mark III/IV systems). These units contain tunable synthesizers, also locked to the maser, and filters. They mix the IF down to baseband, where one fre-

quency edge of the signal is at DC. They also limit the bandwidth to any of a number of values, which on the VLBA range in factors of two from 62.5 kHz to 16 MHz. At the VLA, the equivalent step happens in the control building after the signals are sent over the waveguides from the antennas. The VLBA has 8 such BBCs, each able to produce both upper and lower sidebands. Mark III/IV systems have 14 or 16 video converters. One output signal from a BBC is called a “baseband channel” and is the narrowest analog data signal involved in an observation. Usually be 4 to 8 baseband channels are used, often in pairs of right and left circular polarization.

The BBC outputs are sent to the samplers where they are digitized. The digitization is very coarse, either to 2 levels (1 bit) or 4 levels (2 bit). The VLA uses 3 levels. Sampling usually occurs at the Nyquist rate, which is twice the bandwidth. Despite the coarse sampling, the sensitivity loss is not great. The loss relative to an analog system is 0.64, 0.81, and 0.88 for 2, 3, and 4 level samples at the Nyquist rate (TMS). The small number of bits per sample simplifies the electronics. It is also the way to get the greatest sensitivity if the system is limited by the number of bits that can be transmitted, as it is in VLBI with its tape systems. For a given bit rate on the VLBA, 1 bit, 2 level samples give the higher sensitivity, but 2 bit, 4 level samples are only about 2% worse and do it with half the observing bandwidth. For spectral line observations, where the source is of limited bandwidth, this is a great advantage. For continuum observations, narrower bandwidths impose more relaxed requirements on finding delays before averaging and make it easier to avoid RFI. For spectral line observations, there is often excess bandwidth available, even with 2 bit samples. In such cases, faster sampling, called “oversampling”, can be used to improve the sensitivity. With factor-of-2 oversampling the sensitivity losses listed above become 0.74, 0.89, and 0.94 (TMS).

Once digitized, the formatter prepares the data for recording. It adds timing information, blocks the data into frames, adds parity bits, and, at least for some VLBA and Mark IV cases, fans out the input bit stream into multiple parallel output bit streams, one for each recording head. A VLBA or Mark IV tape recorder has 32 heads for data and another 4 for system information per head stack and the stack can be moved between passes, allowing recording of 504 tracks total. Mark IV systems will eventually be able to record on two heads. Sometimes the mapping between input data streams and output streams is rotated among several options in what is known as a barrel roll. With this, if a recording head is lost, some of several signals, rather than all of any one, will be lost. By contrast, for non-VLBI observations on the VLA, the digitization happens in the correlator and, of course, tapes are not used for pre-correlation data.

5.3. Correlation

At the correlator, the bit streams are extracted from tape. The gross geometric delay, which can be many milliseconds, is removed using both tape slewing and a buffer. If a bit stream has been fanned out, it is reassembled before correlation. The tapes from all stations in an observation are played back together, maintaining the appropriate time offset to less than one sample. The relative delay can be maintained to better accuracy than the sample interval by introducing

a “fractional bit correction”, which amounts to applying a phase slope across the frequency band either before (FX correlator) or after (XF correlator) cross correlation. In the VLA, there are delay lines in the correlator that are used to align the signals. The fractional bit problem is handled by having a sampler that can change sample epoch by a fraction of a sample interval.

The antennas are moving toward the target source at different speeds because of the rotation of the Earth (or the motion of the satellite in the space VLBI case). Thus there is a different doppler shift of the source at each antenna. To correlate without very fast phase winding, it is necessary to correct for this doppler shift. At the VLA, this is done by offsetting an LO at the antenna. For VLBI, it is considered safer to limit the complexity of the station equipment, so this correction is done in the correlator by the “fringe rotator”. Since this is done after the final filtering, the position of the filter-induced variations in gain across the band shifts with time. This shift has to be accounted for in any bandpass calibration. The “fringe rotator” is effectively a mix with a very low frequency LO. Unlike mixes at higher frequencies, it is not possible to eliminate the unwanted sideband of this mix with filters. Instead, a full complex correlator is used to generate enough information to separate the sidebands and keep only the desired one.

The internal details of various types of correlators are covered in lecture 4 and will be glossed over here. Suffice it to say that the output of the correlator is a time sequence of cross and auto correlation spectra (or lag functions that can be easily converted to spectra). Typical parameters of VLBA correlator output are 16 spectral channels per baseband channel, each averaged for 2 or 4 seconds. Many fewer channels would risk degrading the amplitudes if the apriori delays are not very good. Longer averages also risk degrading the data if the apriori model is not so good and there are high phase rates. Spectral line observations will use many more channels — between 128 and 1024 per baseband channel in the VLBA case. If very unstable phases are expected or if a wide field of view is needed, average times might be shorter, down to roughly 0.1 second. The combination of average time, spectral points, and number of baselines must not combine to give an output data rate that is too high for the hardware — currently about 500 kbps. By comparison, it is common on the VLA to average a baseband to one spectral point for continuum data and to average in time for 10 to 30 seconds. This level of averaging is possible because of the much smaller geometric and clock uncertainties. Of course, spectral line observations on the VLA can generate lots of channels.

Some correlators, especially those that do geodesy using the Mark III and Mark IV systems (but not the VLBA correlator) also detect the pulse cal tones and fringe fit the data. The exported data is then ready for the geodetic analysis packages. VLBA correlator data must be fringe fitted in AIPS prior to export to the geodetic packages.

5.4. Apriori Calibration and Flagging

All of the above steps except scheduling are usually handled by observatory staff. The astronomer is no longer expected to be involved in the actual taking and correlation of the data, at least with the VLBA. After correlation, a distribution tape is written and sent to the astronomer. Also, calibration and log files are

made available to the astronomer over the internet. All tasks from this point on are the responsibility of the astronomer, at least in current practice.

The astronomer reads the distribution tape into a postprocessing package — usually AIPS, although by the time this is published, AIPS++ may be a significant option. The processing steps are divided here into 3 stages. The first involves application of calibration and editing information provided by the monitor systems, and perhaps other external sources. There is some hope that this stage will be automated with time, as it already is to a considerable degree for the VLA. The second stage is the use of the data itself to improve the calibration and editing. This may include self-calibration of calibrators and maybe even the target source, but not the final iterative self-calibration and imaging by which final images may be produced. That is the third stage.

There is a considerable amount of monitor data associated with a VLBI observation. It is used in the editing and calibration process. Currently, these data are in ASCII files left on one of a small number of computers, one of which is at the AOC, that are designated for the distribution of files associated with VLBI. The astronomer copies them to his/her home computer by way of the Internet. Special AIPS tasks exist to read the monitor files and fill the relevant AIPS tables. Soon, the data will be provided in tables attached to the main data set at correlation time, so the somewhat tedious step of gathering, preparing, and reading in the flagging and calibration files will no longer be needed.

Flagging: Flag tables allow the user to edit data that the on-line systems knew would be bad. Common reasons include that the antenna was still slewing to source or that an LO synthesizer was not locked to the maser. After application of these flags, typical VLBA data have only a few remaining bad points, so additional editing is fairly easy. Other stations may or may not provide much useful flagging data. A typical global experiment will require considerable interactive editing, although this is improving as more stations provide useful flag information. For VLA observations, the on-line flags are applied at observe time, also providing data sets that are usually quite clean.

Pulse Calibration: The results of the pulse cal tone detection at the sites is another component of the calibration data, at least for VLBA stations. This information may be used to align the phases across the various baseband channels and to remove any delay and phase fluctuations resulting from instrumental effects, such as cable stretching. Tasks exist to load the data and to solve for the phase and delay offsets to apply to the data. The VLBA usually uses two or more tones per baseband which allows phase slopes within each baseband, not just the phase offsets between bands, to be corrected. But this process involves phase ambiguities, which are resolved by the routine that processes the tone data with the assistance of the equivalent of a fringe fit on a reference scan. Then all other scans are assumed to have delays (phase slopes) that are not too different. Cable calibration, discussed earlier, can be applied along with pulse cal corrections. Note that the pulse calibration results are not really needed for routine VLBA imaging observations because of the high stability of the electronics. Other antennas may have more problems that need correction. For the most accurate geodetic and astrometric observations, the pulse cal calibration seems to help even on the VLBA.

System Temperature and Gain: The calibration data also include the system temperatures measured during the observation. These are used, along with gains and gain curves measured at other times, to calibrate the data amplitudes. The basic equation of amplitude calibration is:

$$S_c = b\rho\sqrt{\frac{T_{s1}T_{s2}}{K_1K_2}} \quad (1-1)$$

where S_c is the correlated flux density, ρ is the correlation coefficient delivered by the correlator, T_{si} is the system temperature of the i th antenna, and K_i is the gain (degrees K per Jy) of the i th antenna at the pointing position of the observation (gain curve applied). Typically system temperatures are provided individually for each baseband channel. The scaling factor b includes both correlator specific scaling factors and corrections for the losses due to the coarse digitization. The term $(b\rho)$ equals the correlation coefficient that would be obtained with a perfect, properly scaled, analog system. For VLBA data read into AIPS with FITLD, corrections have already been applied and the b provided to the calibration tasks should be 1.0.

Gain information for VLBI antennas is generally measured in single dish calibration observations done at times different from the VLBI observations. Gains and gain curves are provided to the user in ASCII files. AIPS tasks exist to read all this information and apply it to the data. Eventually, this information will also be included with the data from the correlator. The provided gains for VLBA antennas are based on large numbers of single dish observations, often spread over years, and can, in the best cases, provide absolute calibration good to a few percent. This is comparable to what is achieved in normal VLA observations using flux density calibrators. Of course, the observer must be alert for antennas with special problems like bad weather and not use them in setting the flux scale.

When the VLA, or Westerbork, is phased up and used as an element of a VLBI array, the local correlation coefficients can be used, along with the flux density of the source, to derive the equivalent of $\sqrt{T_s/K}$ for use in calibration. The required data, and information on how to use it, are provided with the calibration data from an observation.

For non-VLBI observations with the VLA, the system temperatures, or at least something proportional to them, are measured and, usually, applied on-line (the astronomer can choose to block this). This corrects the correlator output amplitudes for many elevation and time dependent effects. Then effectively the gains are measured by observing a source of known flux density. This method could be used for VLBI, but some other instrument would have to be used to make near-simultaneous measurements of the calibrator flux density because essentially all sources that can be seen by VLBI are variable.

Absorption: At the higher frequencies, it is necessary to take into account the absorption by the atmosphere. Some stations provide gain curves that include the effect of absorption, but that is of limited usefulness because the absorption varies, especially at high frequency. Such gain curves are distinguished by a sharp drop at low elevation. The VLBA gains and gain curves do not include the effect of absorption. For the best calibration, the absorption must be dealt

with separately. The AIPS tasks involved have this ability based on a number of methods such as provided zenith opacity (from, for example, tipping scans) or using excess system temperature above a provided, or fitted, receiver temperature and spillover. For the best calibration of VLA data, something similar is also needed because the calibrator observations are generally at a different elevation from the target source observations, and the absorption varies with elevation. For the VLA, tipping scans are usually used.

Polarization: Polarization calibration is covered in lecture 25 and not much will be said about it here. But even users who have no intention of using their polarization information should be aware of the effect of parallactic angle. As any type of antenna, other than one on an equatorial mount, tracks a source across the sky, the feeds rotate with respect to the source. For observations using circular polarization, this causes a phase shift with time and that phase shift goes in opposite directions for the two hands of polarization. This can cause various problems, including compromising phase referencing. A correction gets made for this effect as a standard part of polarization calibration. But, if a polarization calibration is not being made, it is important to at least make this one correction, especially if using phase referencing. Some correlators (but not the VLBA correlator) may do it on-line, but, if not, it is a simple correction that can be computed from the source position, antenna location, and axis type. The AIPS task CLCOR can do the job.

5.5. Data Based Calibration and Editing

After the apriori calibration, there may still be stations whose amplitudes are poorly calibrated because of weather, poor calibration data, or other reasons. Also, the phases will not be calibrated, although some alignments may have been handled with the pulse cal data. Further calibration is based on the data itself, either calibrator data or data on the target sources.

AIPS Tables: A word about how AIPS deals with calibration may help some of the discussion from here. AIPS++ will be similar in concept but not in detailed implementation. While calibrating the data, the actual visibility numbers are not modified. All of the calibration and editing information is collected into tables of which there are several types. Some programs use information in one or more tables to derive incremental changes which are put in solution tables. Other programs either interpolate those solutions onto the final calibration tables, or modify the final tables directly. The final calibration tables include the cal (CL) table that includes amplitudes, phases, delays, and rates; the flag (FG) table that has the flagging information; the bandpass (BP) table that specifies the information necessary to flatten the filter responses; the baseline (BL) table which has any baseline specific corrections (not generally needed for VLBA data); and polarization information in the antenna (AN) table. Most programs which read the uv data have options to apply the information in all of these tables. Often there are several versions of some, especially the cal table, that reflect the results at progressive stages of calibration.

Interactive Editing: One of the first data based operations should be to complete the data editing. The monitor-data-based flagging, in a perfect world,

would catch everything, and all remaining data would be good. The VLA and VLBA often approach this situation fairly closely. Other systems vary. Conceptually, editing is done by displaying the data in some manner that makes bad points reasonably obvious, identifying those bad points, and adding them to the list of points to be flagged. Operationally, there are many ways to do this. One good thing to keep in mind is that with modern, large correlators, almost all bad points are the result of a station based problem of some sort (including tape playback problems), so you almost always want to flag stations, not just individual baselines. Interactive programs exist that can be used for editing and some programs have the ability to identify and edit points automatically. Also, it is possible, and sometimes more reliable, to use the more general purpose display and flagging programs separately. Editing VLBI data is not terribly different from editing VLA data except that there are usually fewer VLBI antennas, which makes it easier, and the VLBI integration times are shorter, which makes it harder. Also, because even the calibrators are resolved, some of the time-baseline image displays useful for VLA editing are more difficult to use for VLBI.

“Manual Pcal”: If real pulse cal data have not been applied, something needs to be done to remove the large phase slopes in frequency that will be there because of clock offsets, and sometimes because of inadequacies in the geometric model. Often it is useful to start with a “manual pcal”, which is simply a fringe fit of one scan (sometimes a few scans) whose results are applied to all of the data. This removes any constant clock offsets and aligns the phases and phase slopes across the various baseband channels. It is common to zero the fringe rate (phase slope with time) from this fit, on the assumption that it is likely to be dominated by scan specific effects. This process is sometimes referred to as ‘Single Band Delay’ calibration. As mentioned before, for a well behaved observation, this may be the only fringe fit that is needed. Note that applying real pulse cal data does the same thing with the added advantage that any time dependent variations in the phases through the system, including relative changes through different electronic paths of each baseband channel, will also be removed.

Example Spectra: Figure 4 shows the baseband channel spectra for a VLBA continuum observation. For each baseband channel, there are 16 frequency channels. For each baseline, two baseband channels are plotted side by side. Both the amplitudes (lower) and phases (upper) are shown. The top plots are of raw data. There are phase slopes across the baseband channels and the amplitudes are in units of correlation coefficient. The bottom plots are after amplitude calibration and after a fringe fit on this scan. The amplitudes are now in Jy and the phase slopes and offsets have been removed. Note that the phases haven’t actually gotten noisier, but rather the scale is much expanded. The knowledgeable student can probably guess, with a high probability of being right, which source this is from the very high correlated flux density¹.

¹The source is 3C 84 at 2.2 GHz.




Figure 1-4. A sample of VLBA continuum data before and after amplitude calibration and a fringe fit. For each baseline, both amplitudes (lower) and phases (upper) are shown for two baseband channels. Note the different phase scales before and after calibration — the phases have been flattened.

Fringe Fit: It may prove desirable to fringe fit all the data, or at least all of the calibrator data, to remove variable delays and rates. This can be done before or after the final amplitude calibration, although it must be done after calibration if a source model is used or if one wishes to have the data weights treated properly. When fringe fit solutions are interpolated onto the cal table, the phases are interpolated using the derived fringe rates. That interpolation can go wrong, especially with long solution intervals. When it does, the calibrated data will show a full turn of phase between the times of two solutions, something which can be quite harmful to later processing. If at all possible (source strong enough to tell), such situations should be identified and fixed, which, unfortunately, can be tricky in practice. Or one of the interpolation options that uses phases from integrated rates rather than from the fit results can be used.

Amplitude Adjustments: The apriori amplitude calibration is likely to have worked very well for some stations but not very well for others for various reasons, like bad weather or poor available calibration hardware. The “amplitude check source” can be used to fine tune the amplitude calibration for all stations and baseband channels. First an image or model needs to be made — that process is discussed later. Then the model can be used in an amplitude self-calibration

step to set the relative amplitude gains of all the antennas and channels. Either in generating the model or image of the calibrator, or after the self-calibration, the amplitudes need to be adjusted so that the average change of gain on the subset of antennas with good apriori calibration is unity.

Bandpass Calibration: The data can also be used to do a bandpass calibration. Without such a calibration, there are variations in both amplitude and phase across the individual baseband channels as a result of the imperfect analog filter shapes, as can be seen in Figure 4. These can be removed by measuring them on calibration sources using algorithms similar to self-calibration. Bandpass correction is necessary to calibrate spectral line data where you want all channels calibrated individually as well as possible. For continuum observations, in which the spectral channels are averaged together once any phase slopes are removed, a bandpass calibration might improve the accuracy of the closure parameters. But the value will depend on the quality of the filters used in the experiment. For an observation using the VLBA and equivalently good systems, the improvement is probably not great, although it has not yet been quantified. Bandpass calibration can usually be based on a single calibrator scan, although some antennas (such as the VLA) show low-level time variability and it may be necessary to calibrate frequently for the best possible results.

There is a subtlety to bandpass calibration in VLBI that can be important to a decision about whether to use it on continuum data. Because the doppler shifts of the antennas, which are variable, are removed after the final analog filters, the bandpass shape tends to shift around within each baseband. A bandpass shape measured in one scan can only be used to calibrate another scan if it is shifted in frequency by an appropriate amount. Doing so is easy enough, but one is left not knowing how to calibrate one or more of the edge channels, because that region was not measured in the calibrator scan. The current bandpass programs deal with this problem by throwing out the edge channels and for consistency, the same number of channels is thrown out at both ends of the spectrum. Thus a continuum observation will typically lose 2 of its 16 channels per baseband when a bandpass calibration is done. The resulting loss of bandwidth might hurt more than the improvement in the closure characteristics helps. For weak source observations, bandpass calibration is not worthwhile. For very high dynamic range observations on strong sources, it might be.

Spectral Line Amplitude Calibration: For observations of strong maser sources, it is possible to do the relative amplitude calibration using the autocorrelation data. The ratio of the source power to the total system power, including system noise, is what you need to calibrate and is what the autocorrelation provides. The individual autocorrelation spectra, as a function of time, can be fit to a template spectrum to obtain calibration gains. This can be a very effective way to calibrate, especially in the presence of time variable pointing or absorption effects. This method is covered more completely in lecture 24.

5.6. VLBI Imaging

After the above steps, the data set consists of phases that are reasonably flat with time thanks to good models or fringe fitting. The phases across individual basebands are flat and aligned between basebands. And the amplitudes are in

some good approximation of Jy. The data can typically be averaged at this stage to make a file of much more manageable size. For continuum data, usually all spectral channels in a baseband channel are averaged, and for very weak sources, baseband channels are averaged together as well. The time average may be limited by atmospheric fluctuations, but if not, will usually be for something like 10 to 30 seconds. Any longer does not help the data volume all that much and risks various problems such as having a lot of partial records. Of course, spectral line data will usually not be averaged in frequency. The averaged data set is then ready for imaging.

There are two major aspects to imaging - final calibration and actually producing a deconvolved image from the calibrated data. The calibration steps outlined above do not produce calibrated phases for the target source, which are required in order to make a good image. In VLBI, this has traditionally been done almost exclusively through self-calibration, as opposed to the usual connected-element practice of transferring phase from a nearby calibrator. But, as noted earlier, the high quality models now available are changing this. Phase referencing has become reasonable over a wide range of VLBI observations as long as a nearby calibrator can be observed often.

Phase Referencing: The process of transferring phase from a calibrator is the same as for the VLA (lecture 7). One complication is that it is much more likely that the calibrator will have to be imaged before being used to derive the calibration phases for the target source. However, for observations of a very weak target source with limited dynamic range possibilities, the structure phases on mildly resolved calibrators are not likely to be a big problem and can often be ignored. Another difference from VLA practice is that, especially at extreme frequencies, it is possible that the phase change between calibrator scans will be greater than 180 degrees. This requires utilization of fringe rate information from fringe fits on the calibrator to resolve turn ambiguities.

Figure 3 of the Sensitivity chapter in these proceedings (Wrobel & Walker, p. ??) shows an example of a phase referenced image from the VLBA at 8.4 GHz. It is of the radio source in the nucleus of the Seyfert galaxy NGC 5548, a source of about 2 mJy. The image has an off-source rms of 90 μ Jy.

Phase referencing has the advantage that it can be used on sources much too weak to detect in a coherence time - the maximum time over which self-calibration (including fringe fitting) can be done. It also provides an absolute position, or at least a position relative to the reference source. Such a position can be used for aligning different frequency images or images made at different times. But, the accuracy of the phase transfer will be limited and hence the quality of image that can be made will also be limited. If the target source is strong enough for self-calibration, it will almost always be possible to significantly improve the image quality. It is also possible to make images with self-calibration without the first step of the phase referencing — in fact, this is the normal procedure for strong sources.

Self-Calibration: The process of making an image of a source that is strong enough to detect in a coherence time, but that has very poorly calibrated initial phases, involves an iterative procedure in which both the source structure and the antenna calibrations are determined. This procedure has sometimes

been called “hybrid mapping”, but that term is going out of favor. It is now usually just called self-calibration. There is some room for confusion, because that term is also used for one of the steps of the procedure, the one in which the antenna calibrations are determined. The usage in any given instance will have to be determined from context. Self-calibration is covered in more depth than here in lecture 10.

The procedure is fairly simple in concept. You start with some sort of initial model or image — a point source works fine, but if something better is available, the process will take less time to converge. That model is used to self-calibrate the data. Don’t require very high quality fits at first - you won’t get them. The self-calibrated data are then used to make a new image, which hopefully is better than the original. That image is then used for the next round of self-calibration. For the first several iterations (typically 2 to 10), the amplitudes should be held fixed in the self-calibration. The apriori amplitude calibration is almost certainly better than what is suggested by the initial poor model. For weak sources and well calibrated data, amplitude self-calibration may not be justified. But for many cases, once progress with phase only self-calibration stops, the amplitude calibration can be added. For sources of modest complexity and with decent data, 10 to 30 iterations might be required to produce a final image and the process is straightforward.

The images from a sample self-calibration sequence are shown in Figure 6. This was a simple source with reasonable uv coverage, so the convergence was quick.

Some users have developed automated imaging procedures which cycle through the self-calibration loop and produce a final image without human guidance. Such procedures have produced something on the order of 90% of the images for some of the larger VLBI surveys. But they don’t yet work well for very complicated sources at high dynamic range. Such cases require considerable hands-on guidance. Developing a reliable automatic imaging procedure for VLBI would be a good thesis project for someone interested in algorithms and would be a great boon to VLBI observers.

Self-calibration of large, complicated sources can take time and patience. Large numbers of iterations may be required and every few iterations, one or more parameters of the process needs to be changed or progress stops. The basic problem is probably that, for a source covering a very large number of beam areas, the limited uv coverage of a VLBI array (compared to the VLA for example) just doesn’t have enough sampled uv points to specify robustly the source structure and all antenna gains. The self-calibration loop is basically a fitting procedure, and in these circumstances, it keeps getting stuck in local minima and needs to be pushed out of them to resume progress toward the global minimum. Parameters that can be tweaked between iterations to do this include the clean window, the robustness, the taper, the uv range for self-calibration, and the parameters of any clean component editing that is being done. One should calculate the expected noise level, or determine it from something like a stokes V map, and not be satisfied until the off-source RMS is reasonably close. This process is not particularly difficult, it just takes time.

Figure 7 is an example of what can be done with persistence using just VLBA data on a low declination source. It shows 3C 120 at 1.6 GHz based on

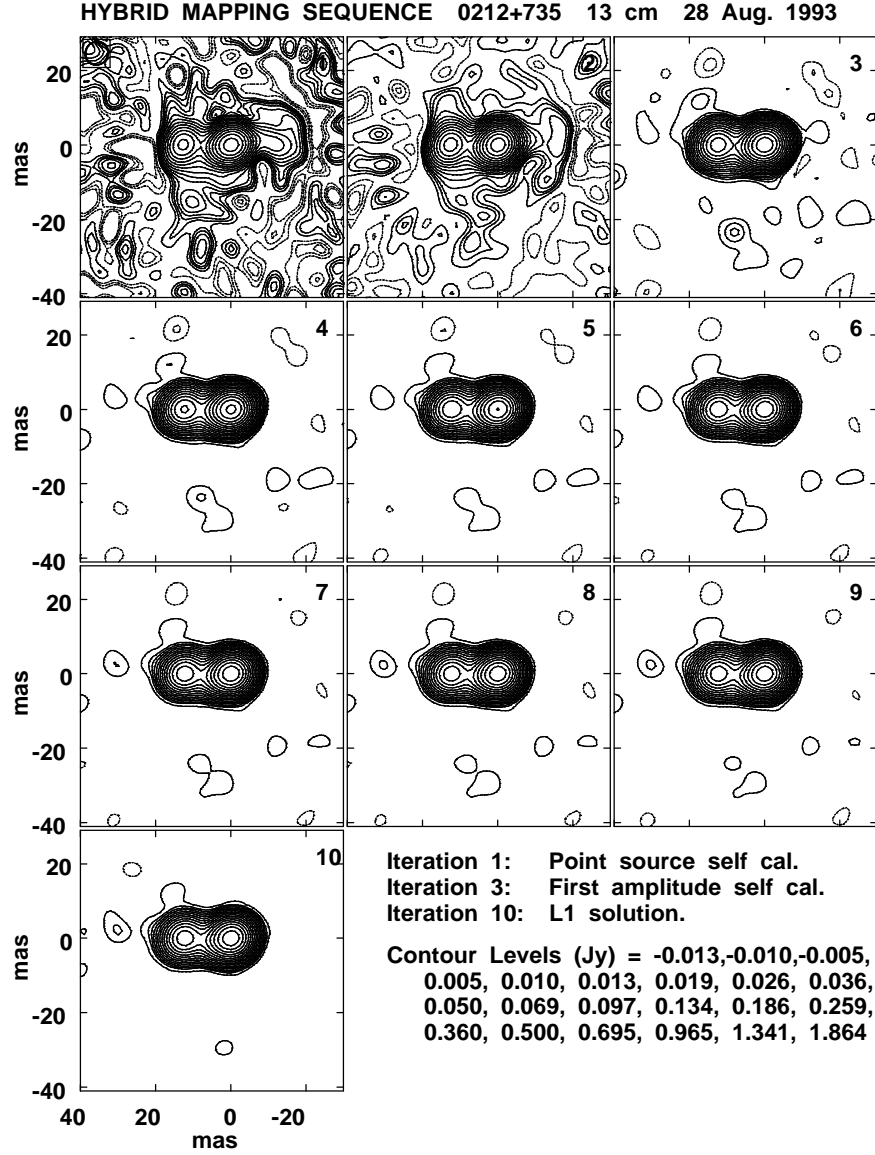


Figure 1-5. The images from each step of a self-calibration sequence. The data are hourly snapshots on a circumpolar source at 2.3 GHz, taken during a geodesy observation. The first two iterations used phase only self-calibration, starting from a point source. The rest used amplitude and phase self-calibration.

data taken in June 1994. The resolution is 7 by 15 mas, extended north-south. The jet is followed out past 0.5 arcseconds.

It is useful, early in the self-calibration process, to explore the beams obtained with a few values of robustness and taper (Briggs 1995). Sidelobes in VLBI tend to be high because of the limited number of baselines. With uniform weighting, the innermost sidelobes tend to be especially large, which can get confused with near-in structural features. With natural weighting, the central




Figure 1-6. A VLBA image of the superluminal radio source, 3C 120, at 1.66 GHz. This demonstrates what can be done with VLBA data on a complicated source.

condensation of baselines, especially with the VLBA, can give a large platform extending well away from the central peak of the beam. This can inhibit proper imaging of large resolved features. It is often possible, with a robustness between -1 and 1 , to get rid of the big platform without seriously enhancing the close sidelobes.

One should not use too large a box when using the CLEAN algorithm for the imaging. With limited uv coverage, CLEAN can manage to describe the noise in the data with a modest number of points scattered around a large image and the computed rms will be artificially low. One easy way to tell if this is a problem is to plot a histogram of the values in an image (see Wrobel & Walker, these proceedings, p. ??). The main peak should be centered on zero and should look Gaussian. If it is a sharp peak with wide wings, you have used too large a CLEAN box. This is mentioned here because it is a common error in VLBI.

Remember that good self-calibration requires that all of the source be described by the model, which is typically a list of CLEAN components. For this to be true, the CLEAN needs to be deep. A shallow CLEAN would leave a significant amount of source flux density out of the model and so it would not fit the data well. That said, in the early iterations, it is typically best to only do a shallow CLEAN, or at least only use the first few points after a merging of components at each position, since any lower level points are likely to be spurious. It is probably better to have an incomplete model containing only real points than a more complete one containing a lot of spurious points.

In AIPS, the self-calibration procedure can be done using a RUN file that strings together separate CLEAN component editing, self-calibration, imaging, display and statistic gathering tasks. An alternative is to use the task, SCMAP,

which integrates most of these steps and includes some editing capabilities. Outside of AIPS, DIFMAP provides a highly integrated procedure that is especially useful with smaller data sets where the computations happen quickly enough that the process can be very interactive. This chapter is based mainly on the AIPS approach because of the author's experience. But for some classes of observations, DIFMAP is almost certainly a better approach, although it does not yet have robust weighting. The strongest convergence of the self-calibration procedure that the author has experienced was obtained using Dan Briggs's NNLS algorithm (Briggs 1995) which is an imaging/deconvolution algorithm based on a least squares fit. That algorithm is not available in AIPS or DIFMAP, but it is in AIPS++ and SDE. It does have some limitations on very large sources. It is especially powerful when the main emission region is partially, but not completely, resolved on the longest baselines.

Self-calibration is also done for data from linked interferometers like the VLA. With larger numbers of baselines and better starting models than is typical for VLBI, the convergence is usually faster, often 1 to a few iterations. But the very highest dynamic range images, especially on sources of some complexity, can require 20 or more iterations. The highest dynamic range images on the VLA also require special efforts to reduce and calibrate closure offsets. One important way to reduce such offsets is to set the delays more accurately than normal, either before observing or in postprocessing using a fringe fit. There is really no fundamental difference between linked interferometer and VLBI data - just differences in the degree of difficulty of certain processing steps.

6. Conclusions

VLBI has reached a new level of maturity in the last few years, largely because of the advent of the VLBA. The instrument is producing data of quality undreamed of not long ago. And by setting an example, it is pushing the rest of the astronomical VLBI world to be satisfied with nothing less. The postprocessing software has not fully kept up as much of the effort has shifted to the AIPS++ project, but it is still only somewhat harder to use for VLBI data than for connected interferometers like the VLA. This situation should continue to improve.

The high quality model used on the VLBA correlator has brought a couple of paradigm shifts to VLBI processing. There once was much concern about all the subtleties of weak source fringe fitting. Now fringe fitting can be confined to calibrators at all but the most extreme frequencies or for space VLBI. Most users should only need to face the much easier task of fringe fitting strong sources. But perhaps the most exciting development is the near routine use of phase referencing. This has meant that VLBI observations can address the science on sources of a millijansky or less — a regime only reached in the past with heroic effort. This greatly increases the number and types of sources that can be observed.

It will not be possible to observe typical thermal sources with VLBI without increases in sensitivity of several orders of magnitude, not something that is likely to happen soon. Therefore the range of source types that can be studied is inherently narrower than for the linked interferometers. But the sources that

can be detected can be studied with resolution far beyond what is available with any other direct imaging technique. A particular strength of VLBI, and one that will require a lot of observing time to exploit, is that the observed sources are small enough to exhibit structural changes on human time scales, even for sources at the edge of the observable universe. Perhaps more so than any other type of telescope, a VLBI array can be a movie camera.

References

- Beasley, A. J. & Conway, J. E. 1995, in *Very Long Baseline Interferometry and the VLBA*, Astronomical Society of the Pacific Conference Series, Volume 82, eds. J. A. Zensus, P. J. Diamond, & P. J. Napier (San Francisco: ASP), 327–341.
- Briggs, D. S. 1995, *High Fidelity Deconvolution of Moderately Resolved Sources* Ph. D. thesis, New Mexico Institute of Mining and Technology.
- Carilli, C. & Holdaway, M. 1997, MMA Memo 173.
- Cotton, W. D. 1995, in *Very Long Baseline Interferometry and the VLBA*, Astronomical Society of the Pacific Conference Series, Volume 82, eds. J. A. Zensus, P. J. Diamond, & P. J. Napier (San Francisco: ASP), 189–207.
- Felli, M. & Spencer R. E. 1989, *Very Long Baseline Interferometry, Techniques and Applications* (Dordrecht: Kluwer Academic Publishers).
- Lonsdale, C. J. & Doelman, S. S. 1998, in *IAU Colloquium 164: Radio Emission from Galactic and Extragalactic Compact Sources*, Astronomical Society of the Pacific Conference Series, Volume 144, eds J. A. Zensus, G. B. Taylor, & J. M. Wrobel, (San Francisco: ASP).
- Ma, C., Arias, E. F., Eubanks, T. M., Fey, A. L., Gontier, A.-M., Jacobs, C. S., Sovers, O. J., Archinal, B. A., and Charlot, P. 1998, A. J. 116, 516.
- Miyoshi, M., Moran, J., Herrnstein, J., Greenhill, L., Nakai, N., Diamond, P., and Inoue, M. 1995, Nature, 373, 127.
- Napier, P. J., Bagri, D. S., Clark, B. G., Rogers, A. E. E., Romney, J. D., Thompson, A. R., & Walker, R. C. 1994, Proceedings of the IEEE, 82, No. 5, 658.
- Peck, A. B. & Beasley, A. J. 1998, in *IAU Colloquium 164: Radio Emission from Galactic and Extragalactic Compact Sources*, Astronomical Society of the Pacific Conference Series, Volume 144, eds J. A. Zensus, G. B. Taylor, & J. M. Wrobel, (San Francisco: ASP).
- Rogers, A. E. E., Doelman, S. S., & Moran, J. M. 1995, A.J., 109, 1391.
- Sovers, O. J., Fanelow, J. L. & Jacobs, C. S. 1998, Reviews of Modern Physics, 70 (In Press for October 1998)
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 1986, *Interferometry and Synthesis in Radio Astronomy* (New York: John Wiley & Sons). (There is a newer edition).
- Zensus, A., Diamond, P. J., & Napier, P. J. 1995, *Very Long Baseline Interferometry and the VLBA* Astronomical Society of the Pacific Conference Series, Volume 82 (San Francisco: ASP).
- Zensus, A., Taylor, G. B., & Wrobel, J. M. 1998, *IAU Colloquium 164: Radio Emission from Galactic and Extragalactic Compact Sources*, Astronomical Society of the Pacific Conference Series, Volume 144. (San Francisco: ASP).