

Pointing and Gain Measurements on the VLBA

R. Craig Walker

March 1, 2015

Abstract:

The procedures and software used for making pointing, gain, and related measurements on the VLBA are described. These measurements include startup tests, focus and rotation tests, high frequency observations aimed at improving the full pointing equation and gain curves, targeted tests to confirm operation of new or modified hardware, wideband sensitivity measurements, tracking tests, and RFI measurements. The software involved, especially that supported by the author, is described.

The target audience for this memo is the people actually performing the observations, overseeing the pointing, gain, and related maintenance activities, and anyone wishing to look at the results in some detail. Also those who manage the VLBA should be aware of the activities described here to ensure that adequate manpower is available. The level of detail and the amount of information presented is likely to overwhelm the average reader.

Introduction:

The determination of pointing equation coefficients, gain values, optimum focus and rotation values, and wideband sensitivity on the VLBA is based on single-dish observations of bright sources. This memo is an effort to gather the details of the software and procedures used for such observations in one place for future reference. Other related maintenance observations are also covered. The memo is motivated by a need to document portions of the VLBA system that I have supported since the array was built but which now must be passed to others as I retire.

Most of the observations considered here are based on a standard pointing sequence which is scheduled by SCHED within what appears to be a single scan to the user and to the VEX file. The details of the pointing sequence are specified in the *.crd* files sent to the legacy control systems at the antennas. A suite of SCHED key files or scripts that run SCHED has been provided to support various needs such as startup tests, extended pointing observations at 1cm and 7mm, focus/rotation measurements, and wideband tests. SCHED has optimization modes that are used to make it possible to use the same scripts for arbitrary time intervals.

After the observations are made, the program FIT is run to convert the monitor power data to pointing offsets, system temperatures, and antenna temperatures. Then, in most cases, program PTANAL is run to produce average or fitted pointing offsets, pointing equation coefficients, gains, gain curves, and much else. The PTANAL output is the immediate product of a pointing experiment. An important function of PTANAL is to draw attention to data that are out of expected ranges. Especially for startup

observations, this is one of the lines of defense on the VLBA against bad data, which is important because of the lack of feedback from real time correlation. PTANAL uses robust fitting and averaging methods so it is not generally necessary to edit the data.

The product of a single pointing experiment and run of PTANAL is not used to set the gains distributed to users or the pointing coefficients used on the antenna. Exceptions exist when hardware changes have just been made and more data are not available. But in most cases, long term averages, or solutions made from data taken over long periods, are used to obtain the values actually used. For pointing coefficients, the FIT output from observations made over several months are combined and used for a PTANAL solution. For gains, the PTANAL outputs from most pointing observations we have ever done are averaged (in a weighted sense) to produce the updated *vlba_gains.key* file. That file is based on a master file (*vlba_gains.master*) that must be maintained in the sense that breaks need to be inserted at times when changes in performance might be expected, such as when a receiver is changed.

For focus and rotation determinations, pointing sequences are run at several focus and rotation offsets, then a program ROTANAL is used to determine the optimum settings. For wideband observations, pointing sequences are made at fairly closely spaced frequencies across the full range of frequency supported by each receiver.

One type of support observation that does not use the pointing sequence is referred to as a PN3DB observation (for Peter Napier 3db). For such an observation, a very bright source is tracked continuously at the half power point of the beam. Separate observations are made at the azimuth and elevation half power points. Such observations can detect issues with the encoders on the antennas.

Details of everything described above are presented in the rest of this memo. Areas where changes will be required as we switch to the Executor based control system from the legacy system will be noted as will ideas about where improvements can be made to the system if someone has time and motivation.

Note that most of the tests described here are scheduled, performed, and analyzed by operations, mainly Jim Ogle, with scientific oversight, test design, scripts, and programs provided mainly by me. My replacement (I have retired) has not been identified, although Vivek Dhawan is familiar with the system and can help in urgent situations.

The Pointing Sequence:

The observations for nearly all of the following topics in this memo are based on the standard pointing sequence, which is often called the "Craig Cross" (I came up with the pattern, not the name). The pattern is illustrated in Figure 1 and described in the caption. The pattern is controlled by parameters written to the legacy control files (*.crd* files) written by SCHED for use by the legacy control computers that command the antenna motions. So far there is no way to control such patterns using the VEX file. For a pointing observation, the VEX file has a single scan covering the entire time of a pointing pattern. This scan lets the new control computer know how to set the RF switches, but that computer knows nothing about the individual pointing offsets.

The power leveling and locking sequence at the start of the pattern is required because the powers need

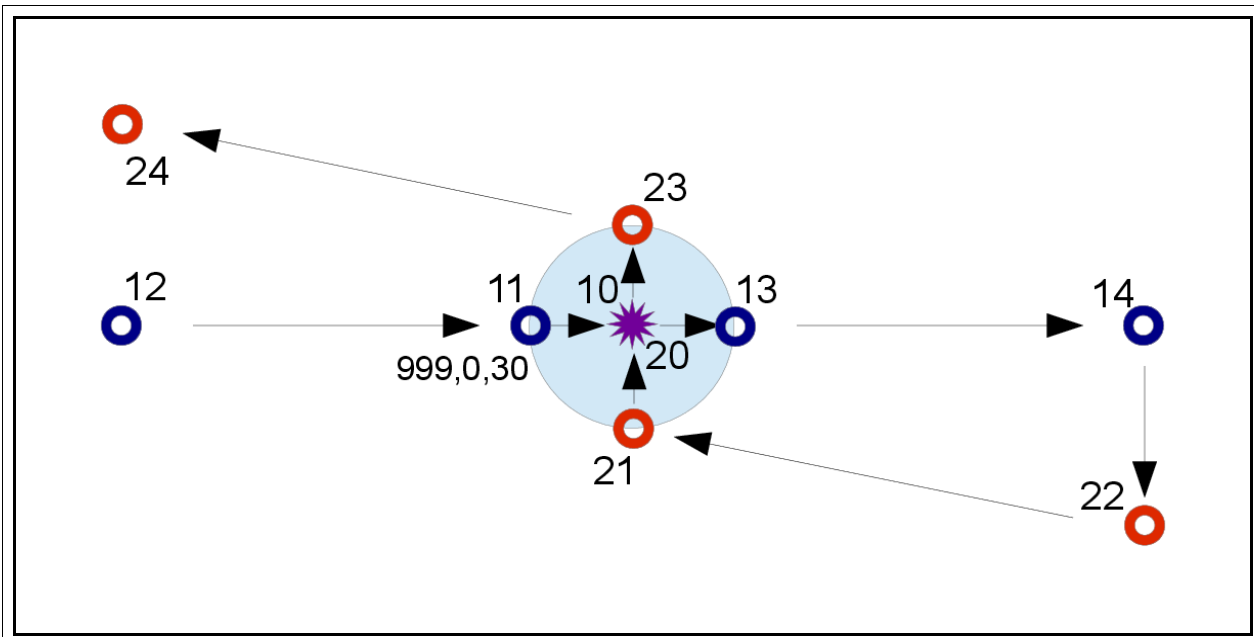


Figure 1: The pointing sequence (Craig Cross). The large shaded region is the half power primary beam. The symbols are pointing positions. The numbers are the qualifier code attached to the scan on each position. The pattern starts with scans to slew to source (999), to autolevel the power in the baseband converters (0) while at the half power point, then to lock the levels (30). Once the levels are set, the pattern is observed in the sequence 12, 11, 10, 13, 14, 22, 21, 20, 23, 24. Often that sequence is then repeated. The elevation off-source positions are offset in azimuth to allow the steps between the 5 elevation scan pointings to be evenly spaced in elevation. This improves the interpolation of the off-source power when at low elevations when there is a gradient with elevation..

reflect the impact of the source. If autoleveling were in effect throughout the pattern, the powers would stay roughly constant.. Perhaps one could use the level settings for the analysis, but that is not what was done. Three scans are used. The first (qual=999 in the *.crd* file) gets the antenna to source. The second (qual=0) is at the half power point with autoleveling requested (level=(1,-1) in the *.crd* file – with a setting for each baseband channel). This sets the levels to about 16000 counts if possible (if not possible, something is wrong with the receiver/IF converter system). The setting is done at the half power point to center the range of powers seen at the normal operating level of the baseband converters. This can be important because some of the sources used are extremely bright. The water maser in W49, for example can increase the total power by a factor of up to about 4 in a 16 MHz bandwidth! The third scan (qual=30) locks the levels at the setting seen in the previous scan. It was found long ago that, if this scan was not there and the antenna was sent straight to the off-source position while locking the levels that there was a lag in the locking and one would get a level appropriate for off-source, not what was seen at the half power position.

Once the levels are set, a five point pattern is used to measure the azimuth offset, followed by five points to measure the elevation offset. This sequence of 10 scans is enclosed in a loop in the *.crd* file. That loop keeps cycling until the overall scan end time is reached. The five points are an off-source measurement to one side of the source at 6 times the half width of the beam (distance from center to the

half power point), then a measurement at the expected half power point, then one centered on the expected source position, followed by half power and off-source measurements on the other side. The monitor system records the cal-on and cal-off power levels for each position in the pattern. The program FIT determines the pointing offset results by first subtracting an interpolated (in time) off-source power from the on-source and half-power positions and then doing an analytic Gaussian solution for peak position, peak amplitude and width for azimuth and elevation. The pattern in azimuth is straightforward for this. But there is an issue in elevation when at low elevations because the troposphere can introduce a strong gradient in the off-source power with elevation. To try to minimize the effect of this gradient on the interpolated off-source powers, the positions of the elevation off-source positions are designed such that the each of the steps in elevation going through the pattern is the same. The off-source positions are then set to the side in azimuth, which is not affected by the tropospheric gradient, sufficiently far to get away from the influence of the source.

To turn a scan in SCHED into a pointing sequence, the switch PTVLBA should be set. If the previous scan requested a pointing sequence and then a scan is desired without such a sequence, NOPTVLBA can be set. The time spent on each pointing position is set by PTDUR, which is commonly set to 15 seconds (smaller numbers can result in not arriving at some positions with adequate time for pointing at low frequencies where the slews are relatively long). Scans should be scheduled with DWELL to allow for the slews. PTSLEW is added to the calculated slew time. It could probably be zero, but is typically set to the same as PTDUR. The DWELL time needs to be calculated to allow for the level setting plus however many cycles through the pattern are desired, so it is $PTSLEW + 2 * PTDUR + N * (10 * PTDUR)$. Deviation from that will cause incomplete patterns as once the DWELL time is reached, the antennas will drop out of the requested pattern loop and proceed to the next scan, which typically is a new source and/or band. The position shift between the nominal source position and the half power points, meant to approximate the distance to the half power point, is set by parameter PTINCR in the SCHED setup file used for pointing, for which there is one for each band. This does allow the user to adjust the step size if desired. Where two bands are measured at once (4cm/13cm, 50/90cm, and, potentially, 6cm/7GHz), usually a PTINCR more appropriate for the higher frequency is selected.

The total power measurement systems on the VLBA do not have any sort of Dicke switch which would allow removal of tropospheric and gain fluctuations by switching between beams. That means that the sources used for pointing must give adequate signal to dominate over such fluctuations in most cases. In practical terms, sources should be above about 5 Jy at most bands and around 10 Jy at the higher frequencies. It also means that using wider bandwidths for pointing observations may not help – the errors are dominated by tropospheric fluctuations. That flux density requirement is a fairly severe limitation on which sources can be used. The pointing observations are also used to measure the telescope gains and gain curves. Therefore, it is important that some fundamental calibrators be included. The main flux density calibrators are way too resolved to use to calibrate VLBI interferometry, but can be used for the single-dish pointing observations. Many of the observations are made on 3C123, 3C286, and 3C274 (Virgo A - for which a beam correction is needed). At 50 and 90 cm, Cygnus A and Taurus A are used along with 3C274. At high frequencies, the flux density calibrators are DR21 and the planets Jupiter, Venus, and Saturn. In addition to these sources, various strong, but variable, continuum sources plus a number of masers are used to improve the sky distribution of the observations.

For continuum sources, the power data are only gathered from one polarization pair of baseband

channels. But much VLBA pointing is done using water or SiO masers. These give an opportunity to do an improved tropospheric correction by taking advantage of the fact that the maser is confined to a small range of frequencies. For such observations two polarization pairs of baseband channels, one centered on the maser line (appropriately Doppler shifted for the source velocity and Earth motion) and one offset in frequency, usually by 100 MHz. The source will only be seen in one baseband pair while the troposphere will affect both in a very similar manner. So the power levels for on-frequency and off-frequency basebands of a given polarization are subtracted along with a constant to take out a DC offset (the value seen on the first off-source scan is used). That difference power is then fit for the pointing offsets.

Currently the antenna motions are controlled by the legacy control computer which gets its observing instructions from the *.crd* file. The *.crd* file allows this to be done in a moderately simple manner because the qualifier, the collimation offsets, and the leveling state can be specified for each scan. Also there is a loop construct. Such concepts do not fit naturally in VEX files. There is an on-going effort to move the functionality of the legacy control computer to the new control computer that was originally deployed to control the RDBE and MARK5C recorder. The control program, which is an offshoot of the VLA Executor, gets its instructions from python scripts that are constructed from the VEX file produced by SCHED. By the time antenna control is shifted to the new system, a new scheme for commanding pointing observations will be needed. It is likely that the pattern will be generated in the control computer based on inputs similar to the PTDUR and PTINCR that are passed to it somehow in the VEX file, perhaps in INTENTS.

There is another pointing sequence that follows the same pattern as the one above, but which is generated on the fly by the control computer rather than by separate scans in the *.crd* file. That is the pattern used for reference pointing. Since the control computer has access to the flags, it is able to set a target integration time at each position. Because of variations in slew times, this leads to somewhat variable total times for a reference pointing measurement. The control computer also does the reduction of the data to new collimation offsets as soon as the pattern is done. If the results pass some quality tests, they are added to the results from the a priori pointing equation to generate the offsets sent to the antenna. Once a set of offsets is established, it is used until either a new set of offsets is determined, or the project ends. There is no way to revert to the defaults. The reference pointing does not use the frequency difference method on masers as described above. That is a desirable enhancement for the future. Reference pointing is usually done at 7mm for use in 3mm observations. At other bands including 7mm, the VLBA does not normally need reference pointing.

Scheduling pointing observations:

Some of the details of how SCHED constructs pointing patterns have already been covered. This section is more about how do build a pointing schedule using SCHED. The simplest scheme is to specify the desired scans in a manner very much like specifying a normal observing schedule. The DWELL times need to be chosen as described earlier, but otherwise one can observe whatever sources and bands are desired. But this could get very tedious given that pointing observations are done frequently, on short notice, and at a wide range of sidereal times. A new sequence would need to be designed for each observation. SCHED has some optimization modes that are helpful here. Command files can be made that only require that the observing date, start time, total duration, and experiment code be specified. This is very useful for operations personnel who can build such a schedule easily

when needed.

For pointing schedules, one of SCHED's simpler optimization modes is generally used. It is invoked with `OPTMODE=SCANS` and a specification of a minimum elevation using `OPMINEL`. Then a loop of scans is given that covers the desired bands and is spread over the full range of right ascension so that at least one scan is up at any time. SCHED simply cycles through the loop, keeping only those scans for which all of the antennas are above the specified elevation. When a separate schedule is built for each antenna, only that antenna needs to be above the minimum elevation. For most of the standard pointing files, the SCHED input files are built by a script that constructs a separate schedule for each of a list of antennas, then runs them, and renames the SCHED output files that are usually not station specific to add the station name. A variety of such scripts, covering a variety of bands and combinations of bands are in `/users/cwalker/bin` and are now under version control at:

<https://svn.aoc.nrao.edu/repos/VLBA/Pointing/trunk/Scripts/Pointing>

The pointing scripts usually have an extension of `“.com”`. Also many such files, that have actually been used can be found by looking for files with names that end in `.com` at `/home/vlbiobs/pointing/mmmmyy` where `mmmyy` is the month and year. That is the area where the SCHED outputs and the analysis results can be found. There are also a couple of pointing command files among the SCHED examples (currently `doptg.com` and `doptg2.com`).

The most frequent type of pointing observation is the startup test. The SCHED examples mentioned above are for startup tests. They are based on a schedule that cycles three times through all of the VLBA bands, including two frequency settings for the 6cm receiver, one near the old standard and one near the methanol line, and two frequency settings for the 1cm receiver, one on the water line and one just below 24 GHz. The 50 and 90 cm bands are observed simultaneously as are standard frequencies near the hydrogen and OH lines (21 and 18cm) for the 20cm receiver. To the extent possible, the sources observed are primary flux calibrators to facilitate measurement of gain. Typically each calibrator is observed in many successive scans at different bands. The minimum elevation is usually set at 20 degrees to avoid low elevation scans. For 3mm, at the stations that have it, reference pointing is done to center the pattern better before the pointing pattern is done. Startup tests, as suggested by their name, are meant to be run at the end of maintenance days to check that all is working, an important exercise on an interferometer with no real-time feedback from fringes. Startup tests can last anywhere from about 2 hours on up, with 4 to 6 being typical. The actual duration is likely to depend on when the first dynamically scheduled project starts after maintenance. Startup tests are sometimes run at other times when there are gaps in the science schedule

Another common type of observation is the “kq” test. That is a schedule that uses significantly more sources than the startup tests, but only observes using the 1cm and 7mm receivers. Those are the highest frequency bands for which the system works well and is available at all antennas. With the smallest primary beam size, they are the most sensitive to pointing offsets. They are also high enough frequencies that gain curves (gain vs elevation) are needed. The primary goal of these observations is to provide data for solutions for the coefficients of the pointing equation and the gain curves. For that, it is desirable to have a good distribution of sources in azimuth and elevation, hence the larger number of sources.

Most other pointing schedules address either the needs for 3mm observing, where reference pointing (generally done at 7mm) is used, or individual bands in cases where some work had been done at an antenna. Scripts are available for such observations at most of the bands. There are some special

types of observing which take special schedules. One type is designed to measure subreflector focus and rotation offsets. The current scheme is to do a 5 point pattern with a full pointing sequence at each point. The points are at the expected focus and rotation position and at offset positions on both sides of the expected position in both focus and rotation. The pattern is established by the SCHED variables FOCOFF and ROTOFF. This pattern is very slow and it might be possible to predict the pointing at each position and just measure on-source powers. But that would require various software modifications that have not been done. Another type of observation is referred to as “wideband” observations. For wideband observations, a moderately dense sampling of frequencies across the full receiver band is made. The object is to check the performance across the band, and to explore the performance at the band edges where some deranged users like to look for spectral lines.

The pointing schedules usually use a different source catalog from the one used for interferometric observing. The catalog is in the usual *\$\$SCHED/catalogs* directory and is called *sources.pointing*. This catalog should not be used for interferometry as the source positions are typically good at the arcsecond level, not the milliarcsecond level. Also many are epoch 1950 coordinates. The catalog starts with a series of water masers, then SiO masers, mostly with *CALCODE='L'* so that the frequency difference mode will be use by FIT. That is followed by a series of continuum sources, most rather strong. Many will overlap sources in the main source catalog, but the positions here, while good enough for pointing, are not to the same standard as the main catalogs. A fraction of the continuum sources will not be in the main catalogs because they are far too big. Many of the catalog entries have comments, mostly from early days of the VLBA, about the strength and suitability of the source. One issue with many of the masers is that they have features spread over several arc seconds. Variations in the relative flux density of the features can cause shifts in the centroid position that are significant for pointing. The pointing source catalog has not had much maintenance for a very long time so it should be revisited.

FIT:

FIT is the program mentioned above that reads switched power data, plus other bookkeeping data from the monitor data, and produces the pointing offsets, antenna temperatures, and system temperatures. There are two output files. The first, traditionally given an extension *.tsm*, is the main data output used by later programs. As we shift to using power data from the RDBE/DDC, I will encourage a shift to the extension *.fit*. That should avoid confusion with output from the program TSM used to extract calibration data from the monitor data on normal observing projects, and make it clear that the file is from the RDBE, not the legacy Baseband Converters (BBCs). The second output file type is given the extension *.tsm.XX*, (or *.fit.XX*) where XX is the station code. There is one file per station. This file contains detailed power data from each pattern. If all goes well, it will not be used. But as soon as there is something wrong with the results, as seen in the *.tsm* (*.fit*) file, this much more detailed file is the place to go for debugging. The results from FIT, and from PTANAL which is discussed later, are generally kept in subdirectories of */home/vlbiobs/pointing* and can be examined there. It would be useful to be looking at some of these files while reading the FIT and PTANAL sections of this memo.

There are two versions of FIT. The first, written long ago, reads the legacy system monitor data. It derives the pointing results from the power records (total power and switched power) derived from the old BBCs. The second version is much more recent and was written by Walter Brisken as a plug replacement for the first version, except that it reads power data derived in the RDBE using the DDC

personality (PFB may be added, perhaps with all 32 channels). Both programs are controlled by the standard commands package and have essentially identical inputs. They have limited inputs: the antennas to use, the time range, whether to use total power or system temperature or a variant on total power more useful for satellites, which pattern to use (Craig Cross, off-on-off patterns, or both), whether or not to use the frequency difference method on sources identified by an “L” in the calcode; the path to the monitor data, and the output file name. Here is a sample of inputs using the typical choices to use total power, Craig crosses only, and use the frequency differencing on line sources:

Inputs for FIT

```
=====
      (DATASELECT)=====
antennas . . . . FD,LA,PT,KP,OV,MK
uttimerange. . . 2014MAY15 at 01:45:00 to 2014MAY15 at 11:30:00
findpeak . . . . TPWR
pattern. . . . . 0
lineoverride . . OFF
      (FILES)=====
mondatapath. . . /home/jansky/mdata/
outfile. . . . . tmay15qw.tsm
```

One key thing to remember when running FIT is to give a time range that starts before the project was loaded to the control computers. Typical practice is to give a start time 15 minutes before the start, but it might even need to be more. The problem is that the equipment records, that record, among other things, the pointing parameters in use, only get written when the schedule file is loaded. If that is missed and the equipment records are not there, the analysis is compromised.

Note that it is possible to use much wider bandwidth baseband channels with the DDC, and it is even conceivably possible to use the full 512 MHz of the IF. In principle, this would allow use of weaker sources, or produce better results on the ones normally used because of the higher SNR. But experience during the original VLBA commissioning suggests that this would not be of much benefit. Over the time scales of the pointing patterns, the tropospheric fluctuations dominate over the system noise. To take advantage of the wider bandwidth, it would likely be necessary to develop another observing strategy, such as slewing through the source so the time over which the off-source power level interpolations are made is much shorter.

An example of the first lines of a FIT output file are given below. Please excuse the small font – that is needed in order to fit the full line width within the page size used for this memo. The file starts with a reflection of the antennas processed and the time range. The line of column headers corresponds to the actual data lines, which come after various bookkeeping records. In this case, but not always, the information starts with some weather records and a record of where the FRM was positioned in focus and rotation. The weather record fields are, in order after “WEATHER”: 1) analog temperature (°C - not usually recorded as indicated by the -99.0), 2) analog dew point (°C - also not recorded), 3) wind velocity (m/s), 4) wind direction (deg), 5) pressure (mbar), 6) rain, 7) snow, 8) digital temperature (°C), 9) digital dew point (°C), minimum temperature (°C – not recorded), 10) maximum temperature (not recorded), 11) dummy logical to keep parser happy, 12) digital wind (not recorded), and 13) maximum gust (m/s).

The PTEQN and PTEQN2 lines give the parameters of the pointing equation in use. The items in the

! Produced by: FIT Jun 27 12 17:00

! For antenna(s): SC,HN,NL,FD,LA,PT,KP,OV,BR,MK
! For UT timerange: 2014NOV26/330 at 03:00:00 to 2014NOV26/330 at 10:30:00

St	Band	Date	Time	Source	P	El	Az	El Off	El Wid	Az Off	Az Wid	Tant	Tsys	Freq	Beam	Tcal	sx	El Col	Az Col
SC	14 330	03 03	WEATHER			-99.0	-99.0	1.2	77	-99.00	0.00	-99.00	27.90	27.00	-99.0	-99.0	T	-99.0	8.2
SC	14 330	03 07	FRM			36.25	213.33										T		
SC	14 330	03 12	WEATHER			-99.0	-99.0	7.9	78	-99.00	0.00	-99.00	27.70	25.70	-99.0	-99.0	T	-99.0	9.1
SC	14 330	03 20	WEATHER			-99.0	-99.0	6.3	77	-99.00	0.00	-99.00	27.60	25.40	-99.0	-99.0	T	-99.0	10.2
SC	14 330	03 20	PTEQN			17.7841	64.602	1.02	1.66	-1.59	0.00			-10800.48			T		
SC	14 330	03 20	PTEQN2			-0.0500	0.140	0.11	0.14	-0.11	-0.02	-0.06					T		
SC	14 330	03 20	FEPARM	0	4cm	13cm	0001	83.4	348.8	0.0	0.0	-1.69	2.85						
SC	14 330	03 20	FEPARM	1	3mm	-	0000	82.0	195.6	0.0	0.0	1.81	0.23						
SC	14 330	03 20	FEPARM	2	7mm	-	8000	81.1	226.0	6.5	0.0	-0.06	5.95						
SC	14 330	03 20	FEPARM	3	1cm	-	8000	80.0	313.0	6.0	0.0	2.40	7.24						
SC	14 330	03 20	FEPARM	4	2cm	-	8000	82.9	250.6	6.5	0.0	0.53	5.70						
SC	14 330	03 20	FEPARM	5	3cm	-	8000	80.1	12432.9	0.0	0.0	2.30	3.40						
SC	14 330	03 20	FEPARM	6	4cm	-	8000	83.0	66.9	0.0	0.0	-1.49	8.40						
SC	14 330	03 20	FEPARM	7	6cm	-	8000	77.5	279.3	0.0	0.0	-0.10	6.37						
SC	14 330	03 20	FEPARM	8	13cm	-	0001	71.3	358.6	0.0	0.0	1.49	8.93						
SC	14 330	03 20	FEPARM	9	13cm	-	0001	62.4	358.6	0.0	0.0	1.06	8.78						
SC	14 330	03 20	FEPARM	10	20cm	-	8000	28.4	139.7	0.0	0.0	-2.54	5.90						
SC	14 330	03 20	FEPARM	11	90cm	-	8000	679.7	139.7	0.0	0.0	-1.72	7.65						
SC	14 330	03 20	FEPARM	12	4m	-	8000	0.0	0.0	-0.0	0.0	0.00	0.00						
SC	14 330	03 20	RAILPRM			0.22	0.10	-0.40	0.32										
SC	14 330	03 20	RAIL1			0.16	0.27	0.29	0.21	0.17	0.08	0.34	0.36	0.32	0.30	0.22	-0.60	0.02	
SC	14 330	03 20	RAIL2			0.40	0.27	0.12	0.24	0.25	0.32	0.40	0.40	0.36	0.34	0.27	0.12	0.05	0.03
SC	14 330	03 20	RAIL3			-0.04	-0.09	-0.19	-0.39	-0.83	-0.42	-0.42	-0.39	-0.31	-0.23	-0.31	-0.30	-0.23	-0.20
SC	14 330	03 20	RAIL4			-0.25	-0.33	-0.31	-0.23	-0.10	0.04	0.16	0.12	0.14	0.06	-0.69	-0.23	0.17	0.32
SC	14 330	03 20	RAIL5			0.31	0.21	0.27	0.38	0.57	0.69	0.50	0.46	0.27	0.26	0.33	0.42	0.29	0.02
SC	14 330	03 20	RAIL6			-0.15	-0.43	-1.14	-0.67	-0.44	-0.58	-0.52	-0.36	-0.33	-0.29	-0.22	0.07	0.30	0.45
SC	14 330	03 20	RAIL7			0.41	0.43	0.52	0.60	0.65	0.67	0.59	0.44	-0.28	0.23	0.40	0.11	-0.21	-0.27
SC	14 330	03 20	RAIL8			-0.38	-0.43	-0.45	-0.37	-0.32	-0.36	-0.28	-0.33	-0.29	-0.21	-0.15	-0.01	0.07	0.10
SC	14 330	03 20	RAIL9			-0.85	-0.43	-0.32	-0.29	-0.20	0.00	0.00	0.05	0.16					
SC	14 330	03 25	FRM			8.30	66.88												
SC	14 330	03 29	WEATHER			-99.0	-99.0	4.1	78	-99.00	0.00	-99.00	27.60	25.20	-99.0	-99.0	T	-99.0	8.6
SC 4cm	14 330	03 32	3C454.3	R		23.89	-79.69	-0.06	5.11	-0.20	4.93	1.521	49.1	8418.99	4.90	3.61	F	8.40	-1.49 1 U
SC 4cm	14 330	03 32	3C454.3	L		23.89	-79.69	-0.27	5.02	-0.02	4.92	1.496	46.5	8418.99	4.90	3.32	F	8.40	-1.49 2 U
SC 4cm	14 330	03 34	3C454.3	R		23.30	-79.54	0.01	4.79	-0.17	4.98	1.534	49.4	8418.99	4.90	3.61	F	8.40	-1.49 1 U
SC 4cm	14 330	03 34	3C454.3	L		23.30	-79.54	-0.18	4.85	0.00	4.97	1.504	46.8	8418.99	4.90	3.32	F	8.40	-1.49 2 U
SC	14 330	03 35	FRM			2.85	139.80												
SC 4cm	14 330	03 35	3C454.3	R		22.89	-79.44	"Power data incomplete"						8418.99	4.90	3.61	F	8.40	-1.49 1 U
SC 4cm	14 330	03 35	3C454.3	L		22.89	-79.44	"Power data incomplete"						8418.99	4.90	3.32	F	8.40	-1.49 2 U
SC	14 330	03 36	FRM			2.85	139.69												
SC	14 330	03 37	WEATHER			-99.0	-99.0	3.3	77	-99.00	0.00	-99.00	27.70	25.60	-99.0	-99.0	T	-99.0	7.7
SC 20cm	14 330	03 38	3C454.3	R		22.49	-79.34	-1.14	40.39	5.16	36.34	1.201	38.1	1438.49	28.66	1.91	F	5.90	-2.54 1 L
SC 20cm	14 330	03 38	3C454.3	L		22.49	-79.34	-1.76	30.71	1.26	31.70	1.429	35.6	1438.49	28.66	2.01	F	5.90	-2.54 2 L
SC 20cm	14 330	03 38	3C454.3	R		22.49	-79.34	0.03	33.03	0.13	30.27	0.933	34.1	1658.49	24.86	1.61	F	5.90	-2.54 3 L
SC 20cm	14 330	03 38	3C454.3	L		22.49	-79.34	-0.55	26.52	1.35	27.48	1.616	43.3	1658.49	24.86	1.56	F	5.90	-2.54 4 L

PTEQN line after "PTEQN" are (units are arc minutes unless otherwise noted): 1) station latitude (deg), 2) station longitude (deg), 3) west tilt, 4) north tilt, 5) sag, 6) axis non-perpendicularity, 7) azimuth encoder offset. Note that the zero of the azimuth encoder is to the south while the zero of normal azimuth is to the north, so the offset is close to 180 deg (10800 arc min). The items after "PTEQN2" are: 1) azimuth encoder centering cosine term, 2) azimuth encoder centering sine term, 3) elevation encoder centering sine term (the elevation cosine term is the same as sag), 4) azimuth 2 theta term cosine, 5) azimuth 2 theta term sine, 6) elevation 2 theta term cosine, 7) elevation 2 theta term sine. The pointing equation will be described in more detail later.

The FEPARM lines give parameters for each receiver, with one line per receiver plus one extra line for the S/X system. The numbers after the receiver name(s) are: 1) a code whose meaning it not clear to me, but which is not used in the analysis, 2) the focus position (mm), 3) the rotation position, 4) the variation of focus with elevation, 5) the variation of rotation with elevation, 6) the azimuth collimation offset, and 7) the elevation collimation offset.

The RAIL lines give the parameters and data needed to correct for non-flatness of the rail. The RAILPRM line gives the four rail parameters. More will be said about the what they mean later when the pointing equation is discussed in detail. In order they are the FP, BP, AZ and EL terms. The following lines are the rail height data with a point at each 3 degrees of azimuth starting at zero. The data are based on survey measurements taken of the height of each azimuth rail bolt. The surveys are done during major maintenance visits. The bolts are at 3 degree intervals, but the first bolt is at an arbitrary azimuth not necessarily near zero nor on an even 3 degrees from zero. The survey measurements are also in thousandths of an inch with an opposite sign from what is wanted. The

engineering division delivers a file of rail heights, which is kept in the directory /home/jansky3/vlbaops/TCAL, the same area where the tcals files are kept, which will be discussed shortly. In those files, the actual bolt azimuths are given and the height is given in mm. Program PTANAL can be used to read these files, interpolate to the required azimuths, and write text in the format needed to install them at the antennas. The use of rail heights is the subject of VLBA Test Memo 61 by R. Walker (1999).

The next line in the example output is an FRM line. That indicates the focus and rotation positions to which the FRM has been commanded. That will change every band, and as the elevation varies. For measurements of focus and rotation, those lines will change nearly every scan.

There is one other type of bookkeeping line that does not appear in the example line shown. That is one with the words "REFERENCE POINTING" after the elevation and azimuth. After that text, the fields are: 1) frequency, 2) expected beam size, 3) Tcal, 4) flag for S/X, 5) derived elevation collimation offset relative to the a priori position, 6) derived azimuth collimation offset relative to the a priori position, 7) channel number, 8) sideband.

Finally we get to the actual data lines. There are two forms of these lines. The one you don't want to see has an error indication in a text string after the elevation and azimuth. The text describes why FIT did not produce a result from a scheduled pointing pattern. Common reasons include that the baseline is too steep (The off-source powers measured at the start and stop of the azimuth or elevation sequence differ by too much), there is a high negative value (The interpolated off-source value was significantly higher than the power measured at the on-source or half-power position), and the beam is too large (The width of the derived Gaussian is interpreted as the beam width, whose expected value is pretty well known from frequency and antenna size. If it is far from expectations, the data are assumed to be bad.). There is one other possible string here that is not actually an error. That is "Off-line data", which means that these channels are assumed to be the ones off-source in frequency to use with the immediately preceding lines for the frequency difference method used on spectral line sources.

The other form of the line is the actual pointing result. The columns correspond to the column headers at the top of the file. The first 10 items are the station code, receiver, year (last 2 digits), day of year, hour, minute, source, polarization, elevation, and azimuth. Then comes the elevation position of the peak of the derived Gaussian in arc minutes from the a priori offset plus the reference pointing offset. Next is the full width, half maximum (FWHM) of the derived Gaussian in elevation. This should be close to the expected beam size. The next two numbers are the same, but for azimuth. The following 2 numbers (fields 15 and 16) are the antenna temperature and system temperature. Then comes the observing frequency in MHz, the expected beam size (which can be compared with the fitted sizes), and the Tcal at the observing frequency. The logical flag is positive for S/X – the ellipsoid reflector is deployed. Other fields have a dummy in that logical to avoid some parsing issue with old code. The last 4 numbers are a priori elevation collimation offset, the a priori azimuth offset, the baseband channel number, and the BBC sideband (usually the channels have upper RF sideband, but the BBC sideband is of interest when checking the BBC electronics).

A FIT file can, and usually does, include several stations. FIT files can be combined (with the Unix command "cat", for example) to run a pointing analysis on data from several pointing sessions.

The station specific file from FIT contains the detailed power measurements for a pointing pattern.

Example data for one pointing pattern at 6cm at LA is:

```

LA 6cm 14 324 03 47 58 AZ= 78.93 EL= 40.03 FQ= 4993
PTVLBA: 3C123 R AZpk= -0.02 wid= 8.25 ELpk= 0.19 wid= 8.24 amp=1106.68
TP 15751 16009 16860 16000 15744 15750 15969 16825 15993 15705
SP 276 277 276 276 276 278 281 274 275 277
TS 32.08 32.53 34.39 32.59 32.06 31.84 31.98 34.56 32.75 31.93
LA 6cm 14 324 03 47 58 AZ= 78.93 EL= 40.03 FQ= 4993
PTVLBA: 3C123 L AZpk= -0.72 wid= 8.18 ELpk= 0.37 wid= 8.24 amp=1124.27
TP 15516 15868 16610 15684 15511 15523 15717 16582 15775 15459
SP 320 319 320 321 319 323 322 321 320 318
TS 26.97 27.66 28.92 27.17 27.06 26.77 27.13 28.72 27.45 27.05
IP 10 12 13 12 12 10 10 12 12 10

```

The first line gives the station, receiver, day number, time (hour, minute, second), azimuth, elevation, and frequency in MHz. The second line gives the results from the pattern. The data contains the source name, polarization, azimuth pointing offset, azimuth width, elevation peak, elevation width, and amplitude. The TP lines are total power, SP is switched power, and Ts is system temperature. The whole set repeats for the other polarization. The last set of data taken at the same time (there could be several sets if there are several frequency channels) ends with an IP line, which gives the integration time in seconds for the data in that column. If this were a spectral line source using the frequency difference method, there would be a line starting with LP after the TS line which contains the difference power. If there are other frequencies, they will be in new sets like the above. If the set is for the off line data for the difference method, the PTVLBA lines will say "Off-line data" rather than giving fit results. If the fit fails, the PTVLBA line will give an error message rather than a result.

One can see a lot in these power data. Columns 1, 5, 6, and 10 are the off-source positions and all should have rather similar values. Columns 3 and 8 are on-source and should be the highest values. The two should also be fairly similar. The other 4 are half power measurements and should roughly split the difference between the on and off source powers. An exception is for the lower frequency of the dual frequency sets S/X and 50/90, where the "half power" points are actually much closer to the center than the half power point because the position appropriate for the higher frequency is used.

Here is an example for a receiver that is having problems:

```

LA 7mm 14 324 03 05 58 AZ=-64.65 EL= 51.65 FQ= 43124
PTVLBA: DR21 R Power data incomplete
TP 1828 1828 1829 1828 1828 1828 1829 1829 1829 1828
SP 0 -0 -0 -0 -0 -0 -0 0 -0 0
TS 57.57 58.44 60.08 58.74 58.11 57.88 58.86 60.03 58.77 57.85
LA 7mm 14 324 03 05 58 AZ=-64.65 EL= 51.65 FQ= 43124
PTVLBA: DR21 L AZpk= -0.13 wid= 1.02 ELpk= 0.02 wid= 1.09 amp=215.23
TP 16196 16347 16403 16274 16201 16208 16320 16418 16335 16216
SP 526 526 525 526 524 524 525 525 525 526
TS 97.98 98.73 99.28 98.29 98.35 98.43 98.79 99.37 98.86 98.08

```

With the legacy system, the data are autoleveled so the power should have been set to around 16200 on a half power position before the pattern starts. With the RDBE, there will be more fluctuation in the power levels because the leveling is not done in the individual bands before the powers are measured. This example is legacy data, so the values near 1828 seen for RCP are vary bad. Something is wrong. Also, there is no sign of a source – all values are about the same while in LCP, the source is visible. On top of that, the switched power is zero. All this points to the RCP channel not working for the 7mm receiver at LA. Curiously, when the frequency was visited again about 12 minutes later, the data look

good. So I'm not sure what happened. Close examination of data often reveals curiosities – too many to track them all down with the available manpower.

An example of an incomplete pattern is:

```

LA 20cm 14 318 07 56 46 AZ=155.01 EL= 83.33 FQ= 1658
PTVLBA: 3C123 R Power data incomplete
TP 15024 15903 17404 15799 15096 14973 0 17330 15808 0
SP 477 483 475 478 485 482 0 476 481 0
TS 224.52 250.49 292.00 225.36 426.36 427.67 260.35 284.91 265.26 215.24
LA 20cm 14 318 07 56 46 AZ=155.01 EL= 83.33 FQ= 1658
PTVLBA: 3C123 L Power data incomplete
TP 15503 16155 17409 16492 15409 15543 0 17679 16536 0
SP 362 350 340 361 359 364 0 357 367 0
TS 321.04 308.62 267.19 319.42 437.27 494.76 270.08 294.00 272.61 314.01
IP 8 3 10 10 3 8 0 9 9 0

```

Here you see that the first half power position in elevation, and the final off source have zero integration time and the no power data. Looking a bit closer, you see that the same pattern is there in the azimuth scan, except that 3 seconds rather than zero were obtained. This suggests that the antenna was having trouble finishing the long azimuth slews to or from the off-source positions. A quick look at the elevation shows why. At 83 degrees, the change in azimuth (length of the drive) is actually rather large to go the required distance on the sky. Missing data at the start of a pattern can mean that the slew from the last source was not finished. Missing data in the last or last few positions means that the end time of the pointing loop was reached too soon – time to reexamine the scheduled scan durations.

The following is an example of 4cm data taken in rain.

```

NL 4cm 14 287 02 13 58 AZ=141.05 EL= 59.48 FQ= 8419
PTVLBA: 3C454.3 R High negative AZ value
TP 15001 14817 14708 14359 14259 14105 13898 13738 13445 13114
SP 157 154 155 155 156 157 154 155 156 158
TS 89.07 89.45 88.54 86.40 85.40 83.50 84.23 82.43 80.43 77.13
NL 4cm 14 287 02 13 58 AZ=141.05 EL= 59.48 FQ= 8419
PTVLBA: 3C454.3 L High negative AZ value
TP 14884 14689 14602 14292 14171 14027 13831 13660 13358 13055
SP 161 159 163 160 162 161 162 162 160 161
TS 88.23 88.06 85.63 85.06 83.59 83.21 81.64 80.44 79.64 77.48

```

As can be seen, even at this relatively low frequency, the power is unstable. The off source power levels fluctuate by more than the difference between the off and on source powers. If you imagine trying to subtract an interpolated off-source value from the on-source and half power positions, you can easily imagine that you are not left with a nice Gaussian and can see why FIT complained of a high negative value.

In summary, if the main FIT output file is showing an uncomfortable level of failed fits, or if the data seem weird for some reason, these station specific files, marked by an extension of the sort *.tsm.KP* or *.fit.KP*, are the place to go to try to understand what is going on.

Tcal files:

The measurements of T_{sys} for both pointing and for calibration of scientific observations depends on the

measurement of switched power. The noise cal signal, which is inserted just behind the feed of each receiver, is switched on and off at 80 Hz. The power levels are measured separately for the on state (P_{on}) and the off state (P_{off}). These measurements are made in the legacy BBCs with the help of voltage-to-frequency converters and in the RDBE using the digitized voltages. The system temperature is $T_{sys} = T_{cal} \times 2 \times (P_{on} - P_{off}) / (P_{on} + P_{off})$. A key element of this equation is the equivalent temperature of the noise cal signal (T_{cal}). For the VLBA, the T_{cal} values are maintained in files provided by the receiver engineers and kept in the directory /home/jansky3/vlbaops/TCAL. The values in the files are measured in the lab before the receivers are sent to the antennas. The files contain values for both polarizations at frequency intervals of 25 to 100 MHz, depending on band. The first few lines of a typical file (*4cm.kp*) are:

```
! Tcals for the 4cm receiver at KP
! Data is: Freq in MHz, LCP Tcal and RCP Tcal.
RECEIVER 37 20000412.0000 20000413.1700 PP
 7400    1.85    1.54
 7500    1.88    1.63
 7600    2.01    1.95
 7700    1.99    2.15
 7800    2.09    2.04
```

This is for receiver 37 (the VLA and VLBA sequences are combined for the 4cm receivers) which was installed on April 12, 2000 at KP. The Tcal file was updated on April 13, 2000 by PP (presumably Peggy Perley). Note that the numbers are LCP followed by RCP – opposite from the order of polarizations everywhere else in the system. It is not uncommon to find this order is backwards which generally shows up as big RCP/LCP T_{sys} or T_a differences.

The calibration of VLBA data depends on being able to determine the correct SEFD (system equivalent flux density) for an observation by dividing the T_{sys} , measured at observe time and at the specific frequency used by the observer, by the gains, which are measured at certain standard frequencies in the pointing observations. The data are calibrated by multiplying the correlation coefficients by the geometric mean of the SEFDs for the stations in a baseline. This works, regardless of the quality of the T_{cal} values if the observing frequency is the same as the frequency at which the gain was determined because the T_{cal} divides out. Both the T_{sys} and the gain depend linearly on T_{cal} . Problems arise when the observing frequency and the gain measurement frequency differ. Both T_{sys} and gain can vary, in principle, with frequency. But also the T_{cal} varies with frequency. Those variations are a mixture of real variations and measurement errors. Cases are seen in some T_{cal} files where there seem to be unreasonable fluctuations or values. For example, at 6cm at KP, near 5 GHz, there are fluctuations of the LCP values that are almost certainly too large, and in fact that lead to calibrated flux densities that are more scattered than the correlation coefficients.

As we move to wide band receivers, the number of cases where observations are done at frequencies different from the ones at which the gains are measured are increasing, so it is becoming more important to provide appropriate gains and T_{cal} values. VLBA Scientific Memo 51 (Hronek & Walker, 1996) gives measurements over the full receiver bands of T_{sys} , gain, and SEFD. Examination of the plots shows that the SEFD values are considerably smoother than either the T_{sys} or gain. This suggests that the measured T_{cal} values have fluctuations that are not justified. Ignoring fluctuations of the scale that seem to be introduced by the T_{cal} variations, it is reasonably apparent that approximating the gain as a constant with frequency is not likely to be in serious error. More recent data on the new

6cm receiver (4-8 GHz) suggests that a first order slope with frequency might be needed in cases of new wide-band receivers. If we could use a gain that is constant with frequency, it would become relatively easy to distribute the gain values to users. One suggestion is to observe calibrators over the full frequency range of each receiver, determine an appropriate gain to apply at all frequencies (perhaps with a linear variation), and then determine adjusted Tcal values that would give that gain. That should then produce Tsys values that allow users to calibrate reasonably accurately at any frequency, plus will remove any undesired effects of errors in the Tcal determinations.

SARA:

Program SARA is a monitor data dumper for the legacy system that can occasionally be useful for helping understand mysterious problems. Long ago, the SARA output was the input to my old FORTRAN version of TSM (since rewritten in C), the program that generates the legacy calibration data. Now it is not much used except to examine monitor data by hand. But one program does use it. That is PN3ANAL which is described later in the section on PN3DB tests. To use SARA output, you need to know the monitor point meanings and addresses. These are given in VLBA Technical Report No. #5 by D. S. Bagri updated by W. Zamora in February 1996 (at least that is the latest I have).

PTANAL:

Once the pointing offsets and amplitudes are obtained by FIT for the individual pointing scans, the overall analysis for the observations is obtained using the program PTANAL. PTANAL is able to process data from multiple observing sessions, which is the preferred scheme to derive a new pointing equation that will actually be used on the antennas. It will try not to get confused if changes to the pointing parameters were made between the input observations. However it is not wise to analyze data that spans true hardware changes such as a receiver change. PTANAL performs the following tasks for each station:

1. Parse the user input, which is given in KEYIN format – the same as is used by SCHED. The details are given later.
2. Read the FIT output data. Determine the a priori pointing and focus/rotation values used for each scan. Get the pointing parameters and rail heights used during the observations. Get the weather data. The reference pointing offsets are added to the measured offsets so that the analysis is done with the total offset from the a priori pointing equation. If the nominal focus or rotation values are changed, only the data from after the change are used. These are the values from the FEARM records, not the actual set values. It will flag points with high wind, and points below a minimum elevation if one was given by the user. Separate flags are kept indicating whether the data point can be used for gain determinations and for pointing determinations.
3. It is possible that the pointing parameters were changed during the data under consideration, although this is unlikely unless multiple observations are being processed. Shortly after reading all the data, PTANAL adjusts the pointing offsets to be relative to the last pointing parameters seen by subtracting the difference between what those parameters would give as a priori offsets and what the actual parameters used for a particular scan would give.
4. Get the expected flux density for any sources for which one is available. For continuum sources, values from Perley & Butler (2013) will be used starting early in 2015. Previously the Baars flux densities were used, with some corrections from circa 1990. Subroutines have been

provided by Bryan Butler that give model flux densities for the planets, mainly Venus, Jupiter, and Saturn, which are the primary calibrators at high frequencies. DR21 is also used at high frequencies using flux densities determined in an analysis of many pointing observations that also use other calibrators. It is possible to give calibrator flux densities as user input, but that is not normally done.

5. Read any new rail height data.
6. Get the nominal gain, gain curve, and Tsys data from the gains file (generally *vlba_gains.key*). These will be compared with the new values to allow flagging of significant changes. Determine what frequency to report for the output (with line sources, there can be some variation).
7. Correct for any Tcal multiplication factors. In the early days, this allowed corrections for Tcal errors. It is rarely if ever used now.
8. Derive the squint (the difference in position between the RCP and LCP beams). Correct data for the squint so that the derived pointing will correspond to the point between the beams. Note that, because of this, single polarization data could give different results from dual polarization data. This does affect reference pointing which is done with one baseband channel.
9. Solve for the atmospheric contribution to Ts and the closely related parameter, opacity, by fitting the Tsys, corrected for spillover, as a function of $\sec(\text{zenith angle})$. While gains will be derived for both the opacity corrected and uncorrected cases, it is the opacity corrected values that will be distributed to users.
10. Use the amplitude data and the known flux densities of the calibrators, corrected for opacity, to solve for the gain, gain curve (if requested), and flux densities of the non-calibrators. Accumulate values for determination of averages over antennas.
11. For observations using many frequencies within a receiver, accumulate Tsys, gain, and SEFD as a function of frequency.
12. Derive the new pointing equation. The parameters of the equation are given in a later section. For many observations, only the collimation offsets will be determined (eg startup observations) because there is insufficient data distributed in azimuth and elevation to determine the full equation. The terms to fit can be controlled from the user input. The bands to use for a fit can be restricted. It is often useful to include data from all bands, but only fit terms other than collimation offsets to the high frequencies because of the poor accuracy of the low frequency pointing results. The ODRPACK routines are used for the non-linear least squares fit.
13. Get BBC specific data (was meant to help debug BBC issues).
14. Count all the points used for various functions.
15. Write the *.summary* and *.ptanal* files. More will be said about these later.
16. Plot the results. The plots will be described later.
17. After doing all stations, get experiment wide flux density averages and display offsets between antennas.

For many of the fits and averages done by PTANAL, robust methods are used. These generally involve doing a weighted fit or average, then down weighting any seriously deviant points, and repeating. This cycle can be repeated a few times, although usually no more than 1 or 2 repeats is needed. This allows bad data to be kept without corrupting the final results. Because of this, data editing is usually not done, saving a lot of work.

The user inputs for PTANAL are given in KEYIN free format, which is described in more detail in the SCHED manual. The input descriptions here are from the code, slightly updated. The type

indicators are “C*N” character with length N, “R or R*4” single precision real, and “L” logical (usually give the keyword with no argument). A number in parentheses indicates an array with the number of elements shown.

```

C   FILES:
C       Changes to file names will be ignored after first input group.
C   INFILE      C*80      Input FIT file name.
C   SUMFILE     C*80      Output summary file.
C   OUTFILE     C*80      Output data file name.
C   PLOTFILE    C*80      Plot file (PGPLOT conventions). A station code
C                       will be appended and one file made for each
C                       station (/ps and /vps files only. This keeps
C                       file sizes socially acceptable for printing.
C   GAINFILE    C*80      File with nominal gains, gain curves, and
C                       Tsys's. (usually vlba_gains.key with full path).
C   CHKFILE     C*80      File for messages about array problems. Sort
C                       of like on-line checker messages.
C   RAILFILE    C*80      Input file for rail height data. The file
C                       should have one line per azimuth with
C                       the station code, the astronomical azimuth,
C                       and the rail height in mm. Azimuths should
C                       be between 0 and 360 degrees, but can start
C                       anywhere and wrap. Lines starting with '!'
C                       are treated as comments. Several stations
C                       can be combined in one file which can
C                       simplify running the program on many stations.
C
C   PROGRAM CONTROL:
C   DEBUG       L          Give extra printout for debugging.
C   MODPLOT     L          Plot models along with data for debugging.
C   PSCALE      R   (20)   Plot scale for pointing offsets per band. Plots
C                       will go to plus and minus this value.
C   PLOTALL     L          Plot points regardless of flags for pointing
C                       offsets, beams etc, but not derived data.
C   PLDIFFS     L          Make a page of plots of point to point differences.
C                       This was meant to explore if the collimation offsets
C                       between bands could be treated separately from the
C                       pointing equation. Not normally used.
C   PLOTPG2     L          Considerably more plots are done for the bands used
C                       for a fit (FITBAND) than for ones for which only
C                       collimation offsets are being determined. PLOTPG2
C                       causes the larger group of plots to be made
C                       regardless.
C   TSYSMAX     R   (20)   Flag data with Tsys higher than this, per band.
C
C   EXTERNAL CALIBRATORS:
C   SOURCES     C*8 (20)   List of sources with provided fluxes to be
C                       used in addition to the Baars sources and planets
C                       for calibration.
C   CALBAND     C*6 (20)   Band name for the provided flux.
C   CALFLUX     R*4 (20)   Flux density for source and band above.
C
C   BANDS:
C   BANDS       C*6 (20)   Band names. (Max of 20 allowed)
C   TATM        R*4 (20)   Assumed Tatm for band. If not 0, this overrides
C                       the sec(z) fit.
C   TCALR       R*4 (20)   Revised RCP Tcal for band. If 0, use input
C                       file. Only for data in specified bands.
C   TCALL       R*4 (20)   Revised LCP Tcal for band. If 0, use input
C                       file. Only for data in specified bands.
C   TMULTR      R*4 (20)   A multiplier for each band for RCP system and
C                       antenna temperatures. Can be used as an
C                       alternative to TCALR to correct for cal. Will

```



```

C          be multiplied by the FT factor from the
C          GAINFILE. Rarely used.
C  TMULTL   R*4 (20) Like TMULTR for LCP.
C  TRXR     R*4 (20) Forced value for receiver temperature, rcp.
C          One per band.
C  TRXL     R*4 (20) Forced value for receiver temperature, lcp.
C  FREQLOW  R*4 (20) Lower limit for frequencies to be included in
C          this band. A data point will be accepted into
C          this band regardless of the band name in the
C          tsm file if its frequency falls between FREQLOW
C          and FREQHIG. Note that a data point will
C          also be accepted if the band names match.
C  FREQHIG  R*4 (20) Upper limit for frequencies in band.
C
C STATIONS:
C  STCODES  C*3 (20) A list of possible station codes.
C  TDIFF    R*4 (20) Use for the difference between the surface
C          temperature and the mean atmospheric temp.
C          One for each station in STCODES.
C
C FIT CONTROL:
C  FITBAND  C*6 (3) Bands for full pointing solution. Def: 4cm,
C          although usually 1cm and 7mm are used.
C  SOLVE    R*4 (50) Solve for specified parameters.
C          eg SOLVE=3,4,5,6 to leave tilt constant.
C          Parameter numbers given in POINTING FIT RESULTS
C          section of output. If SOLVE = 3,6 , only the
C          collimation offsets are wanted and they are
C          the same as the average offsets. This triggers
C          a special STARTUP mode that has slightly
C          different and more compact printout.
C  SETCOEFF R*4 (50) Forced changes to coefficients. All in
C          arcminutes. These are applied to the data
C          just before the fit and added to the
C          results of the fit. They should
C          have no effect on parameters that are fitted,
C          but force a result for parameters that are
C          held fixed (not listed in SOLVE). These are
C          changes to the old equation, not new totals.
C  PTTYPER  C*8      Specify pointing equation type. The options
C          are NTHETA, RAIL, BCA, and WIND. Normally
C          it is left blank. See PTEQNO for details.
C  TEMP     R*4      Assumed ambient temperature for opacity calc.
C          Will be used if weather data not present.
C  GAINCURV R*4      Order of gain curve to fit. 0=> don't fit.
C
C DATA EDITING:
C  TANT     R*4      Minimum Tant for editing data for pointing
C          solution. Default 0.3
C  MAXSECZ  R*4      Max Sec(z) for editing gain data. Default 3
C  RBEAM    R*4 (2) Lower and Upper limits on beam size for edit
C          as fractions of expected beam size.
C  ELRANGE  R*4 (2) Only process data in el range ELRANGE(1)
C          to ELRANGE(2). To get fluxes in presence of
C          gain curves.
C  MAXWIND  R*4      Throw out data when the wind is more than this.

```

To give a good flavor of how PTANAL is run, the following is the portion of the script *PtanalMulti* that runs PTANAL for the case of a data set with all stations and with all bands and enough 1cm and 7mm data for a pointing equation fit. The rest of the script only concerns finding the right execute module for the type of computer being used (setting \$MPBIN).

```

$MPBIN/ptanal <<eof
! Input file for ptanal for startup type pointing.

INFILE   = $1.tsm           ! Input data file from FIT. (will be .fit for RDBE data)
SUMFILE  = $1.summary       ! Output summary file
OUTFILE  = $1.ptanal        ! Output data file
PLOTFILE = $1.plot/vcps     ! Output plot files
CHKFILE  = $1.check         ! Check output.
GAINFILE = /home/jansky/POINTING/vlba_gains.key ! Expected gains and Tsys.
! RAILFILE = "$RAILFILE"

STCODES  = SC,HN,NL,FD,LA,PT,KP,OV,BR,MK
BANDS    = '90cm', '50cm', '21cm', '18cm', '13cm', '13cmsx', '6cm', '7ghz',
           '4cm', '4cmsx', '2cm', '1cm', '24GHz', '7mm', '3mm'
FREQLOW  = 324, 609, 1400, 1640, 2240, 2240, 4960, 6400,
           8380, 8380, 15320, 22200, 23600, 43100, 80000
FREQHIGH = 328, 613, 1450, 1680, 2300, 2300, 5040, 6800,
           8460, 8460, 15400, 22350, 24000, 43300, 90000
GAINCURV = 1, 1, 1, 1, 1, 1, 1, 1,
           1, 1, 2, 2, 2, 3, 3
TSYSMAX  = 600, 500, 60, 60, 70, 70, 70, 70,
           70, 80, 100, 250, 250, 250, 300

! PSCALE  = 9.0 force scales
FITBAND  = '1cm', '7mm'

PLOTPG2 ! Make time plots for all bands.
! MODPLOT

! Set the rail height to what look to be good nominal values based on
! antennas that have significant rail fluctuations.
! May want to turn off 14 (el sin encoder offset) eventually.

SETCOEFF(19) = 0.32 SETCOEFF(20) = -0.40
SOLVE = 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 15, 16 ! , 19, 20

MAXWIND  = 20.0 ! Flag data during higher winds (mph).
TANT     = 0.2 ! Minimum antenna temperature.
MAXSECZ  = 5 ! Max secz for fits for receiver
           ! and atmosphere temperature.
TEMP     = 280 ! Default atm temp if met data missing.
RBEAM    = 0.6,1.5 ! Allowed deviation from expected beam size.
! The following allow ptanal to ignore most of the usual edits.
! PLOTALL ! Plot all points - ignoring flags.
! MAXOFF = 5.0 ! Maximum pointing offset as a fraction of expected
           ! beam size (def 0.7)

/
eof

```

The pointing equation has been described in previous memos so won't be discussed in great detail here. For reference, the parameters of the pointing equation and their functional forms are listed in Table 1. For more information about the basic pointing equation, please see VLBA Test Memo 24 (Walker, 1991). The use of the rail heights to improve pointing is described in VLBA Test Memo 61 (Walker, 1999) on "Pointing Improvements Using Rail Height Information". The effect of the rail heights is a bit complicated because there are 4 wheels in a square with the optical axis half way between two of them. Lifting a wheel partly rotates the support structure on one side of the antenna, rotating the elevation encoder if it is on the encoder side and shifting the elevation bearing up and to the side causing a pointing shift in both azimuth and elevation.

No.	Parameter description	Azimuth function	Elevation function
1	West tilt	$\text{Cos}(Az) \text{ Sin}(El)$	$-\text{Sin}(Az)$
2	North tilt	$\text{Sin}(Az) \text{ Sin}(El)$	$\text{Cos}(Az)$
3	Azimuth Encoder Offset (typically near 180 degrees because of ACU zero point.)	1.0	---
4	Azimuth Collimation Offset FITBAND 1 (Collimation offsets are band dependent)	$\text{Cos}(El)$	---
5	Sag	---	$\text{Cos}(El)$
6	Elevation Collimation Offset FITBAND 1 (for first FITBAND)	---	1.0
7	Azimuth two theta term for $\text{Cos}(2Az)$. ("potato chip terms")	$\text{Cos}(2Az) \text{ Cos}(El)$	---
8	Azimuth two theta term for $\text{Sin}(2Az)$	$\text{Sin}(2Az) \text{ Cos}(El)$	---
9	Elevation two theta term for $\text{Cos}(2az)$	---	$\text{Cos}(2Az)$
10	Elevation two theta term for $\text{Sin}(2Az)$	---	$\text{Sin}(2Az)$
11	Axis non-perpendicularity (so far not used)	$\text{Sin}(El)$	---
12	Azimuth encoder axis offset, cosine term.	$\text{Cos}(Az) \text{ Cos}(El)$	---
13	Azimuth encoder axis offset, sine term.	$\text{Sin}(Az) \text{ Cos}(El)$	---
14	Elevation encoder axis offset sine term. (cosine term is the same as sag).	---	$\text{Sin}(El)$
15	Azimuth collimation offset, FITBAND 2	$\text{Cos}(El)$	---
16	Elevation collimation offset, FITBAND 2	---	1.0
17	Azimuth collimation offset, FITBAND 3	$\text{Cos}(El)$	---
18	Elevation collimation offset, FITBAND 3	---	1.0
19	Rail height model BCCR term.	See Test memo 61	
20	Rail height model BCB term	ditto	
21	Optional rail height model BCA (when PTTYPER='BCA')	ditto	
22	Optional rail height model BCD.	ditto	

Table 1: The pointing parameters.

In the fits, some of the pointing terms are highly correlated. This is especially true of the terms that are constant, sine, or cosine in elevation. This is because the range of elevations used is typically only between around 10 degrees and 80 degrees. That is not that large a fraction of a full 360 degrees of phase needed to keep the terms orthogonal. PTANAL gives some listings of combinations of these terms to help assess the significance of changes. It also shows the correlation coefficients.

One of the terms of the pointing equation, the axis non-perpendicularity, has not actually been used in practice. It is highly correlated with the azimuth collimation offset and the azimuth encoder axis offset so those terms can absorb most of any real offset. Early in commissioning of the VLBA, fits were made using the non-perpendicularity term, but it did not prove to be helpful, and, because of the high correlations, just made the other terms less stable between different fits.

The tilt has proved to be interesting at Pie Town, and to a lesser extent, at some other stations. Pie Town has seen a maximum rate of change of tilt of about an arc minute per year, although the rate has generally been smaller and the direction has changed so the accumulated change is smaller, a bit over 2 arcminutes. The changes are big enough to see in geodetic VLBI data (Petrov et al. 2009, Journal of Geodesy, 83, 859).

PTANAL writes several **output files**. Most of the useful results are in the *.summary* file. A more machine friendly table of point-by-point results is in the *.ptanal* file (which is not used much), and, for each station, there is a plot file. There is also a *.odr* file produced by the low level fitting routines. I suspect that file is never used. Finally there is a *.check* file that is meant to have a summary of problems identified by PTANAL, mainly by comparing the results with previous results found in *vlba_gains.key*, the standard VLBA gains file, or with generally expected results. All of these files, at least for processing done by operations, can be found in subdirectories of */home/vlbiobs/pointing*, as mentioned earlier and it would be useful to be looking at some while reading this section.

A run of PTANAL with a significant amount of data can produce a lot of output. As an example, the *.summary* file alone from an analysis of about 2 months of data is approximately 5000 lines long and full of useful data. The plot file for each station has 41 pages of plots with 8 plots on most pages. This is 2 pages of plots for each of the 15 standard bands (some receivers participate in 2 standard bands) plus several extra pages for the fit bands and a page of weather plots. Even simple startup runs involve all standard bands so produce about as many plots and summary lines. Perusal of all these data is both important for maintaining awareness of issues on the array and frustrating in the sense that generally one has to triage the problems. There just aren't the resources to address everything that might look suspicious.

The *.check* file summarizes the obvious problems, and tries to prioritize them. Startup tests are done immediately after maintenance days and are processed through FIT and PTANAL immediately by the operators. The *.check* file is their main resource to determine if there are any problems that need to be addressed immediately.

The *.ptanal* file is potentially useful for anyone trying to extract detailed pointing results. But I have not located any currently active programs that use it, so the details of its format will not be discussed here. The column headings are reasonably self-explanatory. One warning to someone trying to use it is that barely enough digits have been allowed for some parameters so adjacent numbers sometimes do not have intervening gaps. It must be read fixed format, or PTANAL needs to be fixed to make spaces. This may be mainly a problem for the lowest frequencies where the beams, offsets etc are all large.

The *.summary* file is the main repository for the numerical results of the pointing observations. It starts with a "RESULTS SUMMARY" which is a table of results by band, where the bands were defined in the user input. The table starts with an accounting of the number of data points attempted,

the number failed, the number used for pointing, the number on calibrators of known flux density, and the number used to determine gains and gain curves. The number attempted and failed give a good idea whether there were problems sufficiently severe to prevent obtaining pointing results at all. The number used for pointing and for gains are different because of somewhat different flagging criteria. The next items are the observing frequency and the expected beam size.

The next block of results in the table are the sensitivity items, starting with the Tcal value used and a scale factor (see user inputs above) which is nearly always 1.00. The key sensitivity results are next – the zenith system temperature, zenith atmosphere temperature, uncorrected gain, opacity corrected gain, and opacity corrected efficiency. SEFD values derived from the system temperature and gain values are given for the zenith, for 45 degrees elevation with atmosphere and for 45 degrees elevation with no atmosphere. For the higher frequencies with non-flat gain curves, the 45 deg elevation values attempt to give values more realistic for typical observations, although the values usually reported are the “zenith” values. The quotes around “zenith” recognize the fact that this is not really the value at the zenith. It is the zenith system temperature divided by the peak gain, which, at the highest frequencies, is not normally at the zenith. The 45 degree values do take into account the gain curves.

The next lines are the expected Tsys, opacity corrected gain, and zenith SEFD as obtained from the gains file (usually *vlba_gains.key*). These can be compared to the results obtained from the data being analyzed in order to look for abnormal values. Much of what is in the *.check* file is based on these comparisons. Columns containing discrepant values are marked with circumflex accent symbols (“^^^” - name found on Wikipedia).

The pointing items follow the sensitivity items. First are the squint and beam sizes derived from the data. Then come the collimation offsets. The first is the a priori value from the equipment records, the second is the average value actually found. Both of these are relative to the a priori pointing equation. The third set of values is relative to the newly fitted pointing equation, if there is one. These are followed by RMS deviations of the data from the a priori values, from the average values assuming the a priori pointing equation, and from the newly fitted pointing equation. Comparison of these values gives a good idea of whether changes are needed to the average collimation offsets or even to the pointing equation. To simplify such a comparison, the pointing error that produces a 5% gain loss is printed below the RMS values. Finally, the table ends with differences between the average and a priori collimation offsets, first assuming the a priori pointing equation for both and then assuming the new pointing equation for the new values. Again these help assess whether upgrades are needed to the collimation offsets or to the whole pointing equation.

The squint is an interesting diagnostic because it can detect cross polarization without having to find fringes. The squint is the offset between the left and right circular beams that is seen with antennas with the sort of offset optics (feed circles) that the VLBA and VLA have. If the polarizations are swapped by any of the very many possible ways to do that (like crossing cables), the sense of the squint is reversed. This shows up clearly in the pointing analysis. It was long thought that one could not test for a polarization swap by means other than performing interferometry with other antennas. But the squint provides a handy and quick measurement that can be done on a single antenna that will detect a polarization swap. PTANAL knows the expected squint direction for each band and flags, in the *.check* file, any bands that seem to be reversed. Note that the 50 and 90 cm systems are not so asymmetric with their dipoles at the center of the subreflector, so the squint at those bands is much smaller than frequency scaling from other bands would suggest and does not make a good diagnostic of

crossed polarization.

After the table, on the next page if one were to print the file (rarely done), the main user inputs are reflected. Then the pointing fit results are given if one or more FITBANDS were specified. Otherwise just the a priori pointing parameters are shown. If there is a new fit, the old values, a flag indicating whether the parameter was included in the fit, the change, the new values, and the formal error are given. Any large changes seen here can be a concern that should be understood, unless the changed parameters are the highly correlated ones. The table of fit results does not give the collimation offsets for all bands, just for the FITBANDS. For the others, see the main table of the first page for the postfit average collimation offsets.

For the correlated terms, the table after the pointing fit table shows the appropriate sums as a function of elevation for the old and new values and the change. If the change in these sums is small, then updating the values is probably not critical. The table of the correlated values is followed by the table of the correlation coefficients between the parameters. There will be 100% values on the diagonal, of course. Also, there will be values of 0 for any terms not fitted, such as usually the axis non-perpendicularity. The correlated values are clear, with correlation coefficients sometimes as high as 98 and 99% for some pairs, generally terms that are sine, cosine, and constant terms in elevation as discussed above.

Next comes a list of the rail height data used in the fit. Usually these will be the values from the monitor data. But when new rail height data are introduced, there will be tables of both the old and new data. The points are in mm deviation from a plane at 3 degree azimuth intervals starting at zero degrees. The data provided by the engineering division has had an overall tilt removed so as not to get tangled with the tilt terms in the pointing equation.

The next tables in the *.summary* file give the input and, if a fit was done, the fitted gain curves for each band for which something other than a flat curve is assumed – usually 2cm and higher frequencies. These curves are normalized to have a peak of 1.0, so the gain, as listed in the first page table, corresponds to the gain at the elevation of the peak of the curve. It has been found that, for the 2cm and 1cm receivers (including the extra 24 GHz band in the 1cm receiver), the gain can be approximated by a linear slope. This is a bit surprising to me, but seems to match the data. For the 7mm and 3mm receivers, more complex curves are needed. They generally peak at intermediate elevations. While it is not recognized in these tables, the gain curves are intimately tied to the management of focus and rotation. For both, there is a significant variation in the optimum position with elevation. The feeds for those two frequencies have been placed near (on in the case of 3mm) the elevation axis so that rotating the subreflector can compensate for the shifts due to rotation. We shift focus with elevation for 2cm and higher bands. But we have not yet implemented the rotation shift with elevation. This is a project that should be done soon. We were not able to do it in the past because the FRM would not successfully deal with commands to shift rotation by small amounts. But the FRM has been upgraded, and will be further upgraded in the future. So this should be revisited. For now, the shape of the gain curve, in particular where it peaks, will depend on the choice of the rotation value used.

Moving on down the page, the next segment is the new data in the proper format for installing in the legacy on-line systems at the VLBA stations. You'll note that it looks a lot like what came out in the FIT output file. Items marked with 2 asterisks (“**”) are not meant to change. See the description of the FIT file for the details of what each number is. This is here to simplify the job of operations. The

listing can be cut and pasted into the appropriate place to update the pointing. Note that some new scheme is likely to be needed once the antenna control is shifted to the new control computers.

Now we get to the detailed information about the sources. The table has one line per source per frequency per flux density. That last item in that list may sound weird, but it is the result of the fact that the distances to planets change, causing changes in apparent flux densities. So when PTANAL is given data from several different days, the planets get a line per day. Each line gives the source name, the station name (the table is for only one station, but including the station simplifies understanding the information found with search utilities like grep), the frequency, polarization, the number of observations, the derived flux density and scatter using data and gains not corrected for the atmospheric opacity (based on calibrator flux densities outside the atmosphere with no opacity corrections), the flux density and scatter using opacity corrected data and gains (using the fluxes of calibrators that have been reduced by the atmospheric opacity – the is the performance relative to what actually reaches the antenna), the assumed flux density for the calibrators corrected for beam size, the average pointing offsets in RA and Dec (helps look for position problems with sources, especially the ones that are fairly large in size) for both the a priori pointing equation (“Prefit”) and relative to the newly fitted equation, and finally the program internal index numbers for the band, the source, and the polarization (presumably to help some other program, but I’m not sure which).

The opacity corrections are based on fitting the Tsys measurements, with a nominal spillover subtracted, as a function of elevation do separate the contribution from the troposphere from constant terms such as the receiver. The nominal spillover was derived for each band by a careful examination of high SNR data early in the VLBA history. The tropospheric contribution is assumed to be simply the physical temperature times the opacity. For a nominal physical temperature, the surface temperature measured by the weather station is reduced by a nominal 10 degrees C, although that number can be controlled by user parameter TDIFF. The opacity is derived by first fitting all of the Tsys data for the constant (Trec) term. The fit is done with a robust method that is biased toward the lower values to try to survive weather based fluctuations. Then for each data point, the tropospheric contribution is the Tsys for that point with the constant term and the nominal spillover subtracted. That tropospheric contribution is divided by the derived physical temperature to get an estimate of the opacity.

For the gains, analysis is done for opacity corrected and uncorrected values. For the opacity corrected values, the calibrator flux densities are reduced by the expected absorption. Data calibrated using the opacity corrected gains, which are the ones we give to users, produced flux densities as received at the antenna. Those flux densities must then be adjusted upwards by an amount based on the tropospheric opacity at the time of the observations. APCAL in AIPS can do this. Unfortunately, many users do not appreciate that these opacity corrections are needed at all bands for correct calibration. The uncorrected values are derived by ignoring the troposphere. A typical uncorrected gain curve would drop sharply at low elevations. Data calibrated using the uncorrected gains would produce flux densities correct outside the troposphere if the opacity is the same on the days the gains were determined and the astronomy observation were made. At low frequencies, that might be a reasonable assumption. At high frequencies, it is not. Amplitude calibration is discussed in depth in VLBA Scientific Memo 37 by Walker.

It is useful to check this the source flux density table to be sure that all the sources are adequately strong and don't have a peculiar position in single dish measurements. It is also useful to check that the

derived flux densities match the assumed ones reasonably closely for the calibrators. If more than one calibrator is used, a good match is a consistency check between the calibrators. The scatter numbers can give an indication of the data quality.

Note that there are a few items of concern with sources. The 50/90 cm flux densities may be affected by the galactic background radiation in the sense that the “off-source” positions include significant flux density. It is especially likely that TauA is afflicted with this issue, so I have removed it from the list of calibrators. Another issue is that one of the favorite calibrators, 3C274 (commonly known as VirgoA or M87) is rather large so resolution corrections are needed at most bands. Those corrections were derived by calibrating 3C274 data using other calibrators, mainly 3C123. But the corrections are likely subject to some error. The planets also need resolution corrections, and Saturn needs a ring correction which may not have even been done in older data (I was recently warned of this by Bryan Butler and it does appear to be true in the code where function TBRING returns 0.0 always).

The individual station entries in the *.summary* file end with a table of weather data. The column meanings are clear from the column headings. Any values of -99.0 means that the data were not collected. PTANAL actually uses the temperature (digital value) as part of the opacity calculation. The rest of the numbers are mainly useful in trying to understand the data quality. One thing to be aware of is that the “Rain” value is the accumulated precipitation. So it will step up during a shower. At midnight it gets reset to zero. An item to watch is the dew point relative to the temperature. The dew point is measured with a chilled mirror instrument which detects condensation on a mirror whose temperature is carefully controlled. That measurement can be corrupted if the mirror gets dirty. There is a cleaning cycle once per day (last I heard), so a step, probably at midnight, indicates that the cleaning had an effect. Occasionally the dew point starts accurately tracking the temperature with some fixed offset. That is not normal atmospheric behavior and most likely means that the instrument needs to be serviced. The site techs should be informed – they know what to do.

Note that the geodesy community usually carefully calibrates their pressure instruments at their stations by shipping around a standard. At one time we were doing something similar. As far as I know, this is no longer being done. This probably needs to be revisited. The pressure measurement is needed for accurate atmospheric delay calculations. Unfortunately, mostly due to confusions at Mauna Kea where the high altitude may be causing deviations from simple models, our effort long ago to gain confidence in such corrections and start using the data failed to reach a conclusion. We could likely improve our a priori atmospheric models significantly by starting to use the pressure. Note that the delay involved is from the dry atmosphere which has a far larger effect than the wet atmosphere, but is less variable. Ignoring the pressure changes cannot be good.

The *.summary* file repeats all of the above for each station.

Once all of the stations are finished, there is one final table that looks at the source flux densities globally. The table is broken into segments for each band. Then for each source, frequency, polarization, and expected flux density, the average opacity corrected flux density is given followed by, for each station, the ratio of the value found for the station divided by the average. This is a good way to look for consistency across the array. Any station that has values deviating significantly from 1.0 is probably having problems. One can get a feel for what is meant by “significantly” by looking over all the numbers in the table. For most cases, the stations tend to match to within 1 or 2 percent.

When run from some scripts, the *.summary* file ends with a reflection of the input control file for reference. This is added by the script, not PTANAL itself. If PTANAL is run from one of the standard scripts, there may be no specific control file in the analysis area, so one can look here to see what was requested.

The station plot files, given names like 14sep.plotLA.ps (for the multi-month analysis done in September 2014), provide graphics that make understanding the results far easier than looking at tables of numbers. Each plot file starts with a page of plots of the weather seen during the observations. The items plotted are clear from the labels. Note that the snow depth is not working. As noted earlier, it is a good idea to check if the dew point is tracking the temperature abnormally closely which indicates the sensor needs attention. Examination of the data can provide expectations of the data quality. If there is rain or high humidity, or high winds, the data may have problems.

The bulk of the plots in the file are for the individual bands. There are three basic options. If the observation is something like a startup test where only collimation offsets are determined, but there is no pointing equation fit, there is one page with 8 plots for each band. The first 4 plots are the azimuth and elevation collimation offsets, each plotted as a function of azimuth and elevation. In some cases, colors are used to distinguish between some of the calibrator sources. If reference pointing got used, which is the case if, and only if, there were 3mm observations, then both the offsets measured relative to the reference pointing position and the offsets relative to the a priori pointing equation are plotted. The former indicates how well the reference pointing works while the latter shows the quality of the a priori parameters. The other 4 plots relate to the amplitudes. The first is the system temperature plotted against $\text{Sec}(Z_a)$ which helps show how stable and how opaque the atmosphere is. The atmospheric model fitted to the data is plotted, with a solid line showing the total T_{sys} model including atmosphere, spillover, and “receiver” temperature (where “receiver” temperature includes all constant terms, not just the receiver) and a dashed line showing just receiver plus atmosphere. The plot below is the system temperature plotted against time. That helps determine if there are bad periods. If there are, and a lot depends on the results, one might go to the FIT file and edit out the offending time range. This is rarely done. The next plot shows the opacity corrected gain on the calibrators. These are the data used to determine the gain and gain curves, so this plot helps give an idea of the quality of the gain determination. The final plot is the antenna temperature over the measured flux density for all sources without opacity corrections. This will help show the consistency of the flux density measurements of the non-calibrator sources.

If a pointing equation fit is done, then there are more plots. For the bands that are not the ones being used to fit the pointing equation, there are 2 pages. The first page is the same as for startup tests. The second page for each band gives plots to help assess the quality of the pointing fit. The first plot shows the distribution of the points on the sky. Ideally, but rarely, those are spread all over the sky. To the right of that is a plot of the collimation offset against source flux density. This is meant to give an indication of the degree to which the quality of the individual measurements depends on the amount of signal available. Typically the scatter at high flux densities is lower. This could also be used to establish the minimum flux density for sources used in pointing (typically we stay above about 5 Jy).

Below those two plots, are the prefit azimuth and elevation offsets (prefit means relative to the a priori pointing equation) plotted against time and against time of day. The latter, for data sets that span many observations, helps look for such effects as day/night differences. The bottom plots are the postfit collimation offsets plotted against time of day (postfit means relative to the newly fitted pointing

equation). Generally these data will be somewhat tighter than the prefit values because of the supposedly improved pointing equation. Note that the pointing observations are usually made on sources scattered around the sky. The main structural members of the antenna mounts are fairly well insulated, so temperature changes of the steel will be moderately slow – slower than the time between scans. We rarely see a day/night difference in the pointing results, probably because the illumination of the structure changes on time scales less than the thermal time constant. Larger effects are seen if a single source is tracked, which is more like many normal observations.

If the band is one of those being used for the pointing parameter fit (FITBAND – usually 1cm and 7mm), the selection of plots is more extensive. The first page has the same 4 plots of collimation offsets that are shown for the other bands. But there are 4 additional plots that show the same azimuth and elevation collimation offsets but for the postfit data. Comparison of the prefit and postfit plots shows how much the new pointing equation has improved the pointing. Examination of these plots, especially the prefit plots, can give a feel for what terms in the pointing equation might be needing correction. Such examination can also reveal oddities in the pointing. For example, it is possible for the ACU (antenna control unit) to have an error that only affects a small portion of the full range. This might show up as collimation offsets that are not near zero over the affected range. We have seen cases where the pointing was discrepant over ranges like 10 or 20 degrees.

If reference pointing was used, the FITBAND bands will have a page with just 4 plots. These show the pointing offsets from the reference pointing positions. They indicate how well the reference pointing corrected errors in the prefit pointing (a priori equation) or errors from distortions of the antenna structure. They can be compared with the prefit plots on the preceding page. Recall that, once a reference pointing value has been determined, it is used for all pointing at all bands until it is replaced with a new value. There is no way to set it back to zero. Hence the use of reference pointing for all bands when it is used for any.

The third page for the FITBAND bands is the same as the second page for other sources. It will differ mainly in the density of points because there are typically many more observations of the FITBAND bands than for the others. The sky coverage will also likely be much better because more sources are used.

The fourth page for FITBAND bands gives the same 4 amplitude related plots as were given for other bands plus system temperature vs elevation, antenna temperature vs time, antenna temperature over measured fluxes with opacity corrections vs elevation, and antenna temperature over measured fluxes with gain curve and opacity corrections vs time. These are various different ways to look at, and help understand, the gains. Their usefulness increases when things don't seem to be behaving quite as expected.

The final page for the FITBAND bands has 8 plots that have proven useful for one reason or another when trying to understand data that has issues. Going left to right, then down, the first plot is the antenna temperature (not opacity corrected) over the opacity corrected flux density against the atmospheric temperature. The later is the system temperature minus the receiver temperature and spillover. Note that this is the excess temperature for the scan, not the atmospheric temperature derived from the Sec(Za) fit. Typically, the value plotted decreases as the atmospheric temperature increases, showing directly the effect of opacity. To the right is essentially the same plot, but for which the antenna temperatures have been corrected for opacity. Now the data should be flat, if the corrections

are all working properly.

The next row of plots starts with the zenith opacity (opacity of the scan over $\text{Sec}(Z_a)$) plotted against time. Weather changes will show up here rather clearly. To the right, is the elevation plotted against time. This can be useful to help make sure that some temporal effect is not actually an elevation effect.

The third row starts with the zenith opacity plotted against elevation. If all is well, this should be a scatter plot mainly determined by time effects that would be seen in the plot above. There should not be a trend or the elevation dependence has not been handled properly. To the right is the zenith atmospheric temperature vs time. That should look a lot like the plot of zenith opacity vs time in the second row.

The fourth row first shows the elevation beam size and the elevation squint (times ten) plotted against elevation. The right plot is the same for the azimuth beam and squint (times ten). These can be useful to look for odd effects, and to get an idea of the quality of the measurements.

After all the individual band plots, there are some summary plots. The first set is of the azimuth and elevation collimation offsets for both prefit and postfit data for all of the FITBAND bands together. This helps identify systematic vs band specific effects. The following pages are plots of the T_{sys} , gain, and SEFD in one plot and the beam width and squint, both as a function of frequency for each receiver for which more than one distinct frequency was measured. These plots are mainly useful for wideband tests which will be discussed later when we intentionally sample frequencies over the full RF range of each receiver. That's the end of the plots. For a full pointing analysis, of multiple epochs, there are 41 pages. For a startup test, there are 20 pages of plots.

Pointing equation maintenance:

Most pointing observations are a few hours in duration and involve roughly 100-200 actual pointing patterns. For startup tests, all 15 standard bands are sampled, so most see fewer than 10 observations while the key high frequency bands might see 20-50 observations. A dedicated pointing equation observation at 1cm and 7mm (KQ) might have 50-200 observations per band. These numbers of points are marginal for fitting for a 21 parameter pointing equation, especially considering that the sky distribution is rarely optimal. On top of that, an observation on any one day may be afflicted with poor weather or other atmospheric abnormalities. Therefore we do not like to derive pointing parameters that will actually be used on the antennas from just one set of observations.

PTANAL is happy to read an input file that is several FIT output files appended together, often with a simple Unix "cat" command with output directed to a file. This allows all the pointing results gathered over an extended period of time to be analyzed together. Since PTANAL uses robust fitting and averaging schemes, a few deviant points in such a data set have little impact on the results. Therefore, when new pointing parameters are desired, the usable data are gathered and processed through PTANAL. Such data sets can cover several months with 3 or 4 months being a reasonable duration. The resulting file will often have around a hundred points at each of the non-fit bands and around a thousand in the fit bands. Such a data set will produce pointing parameters that, assuming the plots don't show anything weird, can be trusted.

One issue when making such a solution is that data should not span real hardware changes. For example, if a receiver is changed, only data from one side (before or after) the change should be used and the change should be reflected in a break in entries in *vlba_gains.key*.

The process of gathering the desired data is handled by script *getfit*. *Getfit* should be edited to list all of the pointing observations to use and which stations to use from each observation. This is where a station that has had a change can be restricted to one side of the change, although it is a blunt scheme that cannot distinguish bands (this is an area of possible enhancement). The *getfit* script both gathers all of the FIT results into a single large file with the help of *getfit1* and runs PTANAL on that file by invoking the *PtanalMulti* script that was described earlier. *Getfit*, *getfit1*, and *PtanalMulti* are maintained in svn at <https://svn.aoc.nrao.edu/repos/VLBA/Pointing/trunk/Scripts/Pointing>. In their current form, *getfit1* and *PtanalMulti* are available in /users/cwalker/bin. Eventually all my programs that are actively used should be put somewhere that does not depend on the presence of my directories.

It is important to spend time looking through the *.summary* file and the plot files when a big run of PTANAL is done. Much can be learned about the condition of the array. Once a solution is obtained that appears to be correct and robust, the pointing parameters can be transferred to the on-line systems using a cut-and-paste of the appropriate segment of the *.summary* file as described earlier.

The processing of the large data sets is generally done in dated subdirectories of /home/vlbiobs/pointing/POINTING-EQUATION/. In an ideal world when the people responsible have enough time, this analysis should be done about every 3 months. The interval lately has tended to be much longer.

It is possible to obtain plots of the history of pointing parameters by using files, such as the *summary.pointing* file discussed in the following section, and the program PLOTPTG. The inputs for PLOTPTG are answers to questions it asks. This program is not used much as far as I am aware so it won't be discussed in detail, but it is available and is maintained in the same SVN area as the other programs discussed here. There is also a similar program called PLOTTILT that focuses on the tilt, which is of special interest because of what is happening at Pie Town. It is also in the SVN area, but would need to be compiled if someone wants to use it.

Gain history:

The gain history is maintained in a somewhat different manner from the pointing parameters. In principle, the gains could be obtained from the large pointing runs and there would be some advantages to doing it that way. However, partly for historical reasons, the mechanism that has been established is based on extracting Tsys, gain, and SEFD from the individual session pointing analysis (PTANAL) files and generating averages and plots. The process is carried out using the script *gethistory*. The master of *gethistory* is maintained at <https://svn.aoc.nrao.edu/repos/VLBA/Pointing/trunk/gethistory>. The actual analysis is done by operations in the directory /home/jansky/POINTING. The reader could look there for the current examples of the various files discussed in this section.

Gethistory gathers a list of all the *.summary* files in the pointing areas that were made recently. The files are listed in a file called *best_files*. Program GETSUMMARY is then used to extract the required amplitude information into *summary.amplitudes.new* and the pointing information into

summary.pointing.new. Every couple of years, those files are saved into dated versions in the *HISTORY* subdirectory to save the effort of reextracting the data from all the old *.summary* files. *gethistory* then combines the saved files and the new ones into *summary.amplitudes* and *summary.pointing*. The data in the amplitudes files are the station, band, polarization, year, day, decimal time in fractional years, number of data points, Tcal, scale factor, Tsys, atmosphere temperature, receiver temperature, gain, efficiency, and SEFD. There is one line per station/band/session.

The next action for *gethistory* is to create the new *vlba_gains.key* file. This is *gethistory*'s most important task. *vlba_gains.key* is the file from which PTANAL and the programs involved in calibration transfer get their information about the gains and expected Tsys, SEFD etc. The program that *gethistory* uses to do this job is NEWGAIN, which is maintained in the same svn area mentioned for other programs. The input for NEWGAIN is the *summary.amplitudes* file and a file called *vlba_gains.master*. The latter file is in the same format as *vlba_gains.key*, except that it doesn't contain accurate values for the measured parameters. What it does have is an entry for each block of time during which a band at a station was not expected to experience gain changes. So the time range is important. Also the gain curves are put into this file by hand based on the results from the big pointing fits. Whenever there is an event at a station that could cause gain changes, someone needs to edit *vlba_gains.master* to add a new entry with the appropriate band, station, comment, frequency, axis type, time range and gain curve. NEWGAIN uses the input data to make weighted, robust averages of the key parameters, especially the gain (DPFU for RCP and LCP)

The format of the *vlba_gains.key* file is appropriate for processing using KEYIN, the same parser used by SCHED. The parameters are described in the header for *vlba_gains.master* and *vlba_gains.key*. A typical pair of entries, for 7mm at LA, is:

```

BAND = '7mm'           TIMERANG = 2014,08,18,00, 2014,09,04,00
FREQ =43124.          COMMENT = 'receiver pulled for repairs dummy values'
TS = 87., 105.       TR = 74., 93. TCAL = 7.22, 6.27 FT = 1.000, 1.000
LA GAIN ALTAZ DPFU = 0.100, 0.096 POLY= 0.680, 1.579E-02,-1.950E-04 /

BAND = '7mm'           TIMERANG = 2014,09,04,00, 2100,01,01,00
FREQ =43124.          COMMENT = 'repaired receiver sn08 installed'
TS = 92., 98.       TR = 79., 86. TCAL = 7.39, 6.36 FT = 1.000, 1.000
LA GAIN ALTAZ DPFU = 0.097, 0.093 POLY= 0.680, 1.579E-02,-1.950E-04 /

```

The final entry for each station/band has an end date (last 4 numbers of *TIMERANG*) that is far in the future so that no data will not fall in a range. If all goes well with *gethistory*, the new *vlba_gains.key* can be made the standard. When that happens, a dated version (creation date) should be put in the *VLBA_GAINS* subdirectory and a copy should be put in */home/jansky3/vlbaops/TCAL*.

Next *gethistory* produces a table of the sensitivity related parameters for a specified time using program GAINTABLE. It also produces plots of the history of the Tsys, gain, and SEFD for the full history of the available data (essentially the history of the VLBA) and for a requested shorter period. The user will need to edit the *gethistory* script to designate the time for which to make the table and the time ranges for the plots.

The table file made by *gethistory* with GAINTABLE is called *gains.table*. It contains tables by band and station of the cal temperatures, the gains, the system temperatures, the receiver temperatures, the

SEFD (System Equivalent Flux Density)																
St	P	90cm	50cm	21cm	18cm	13cm	13cmsx	6cm	7ghz	4cm	4cmsx	2cm	1cm	24ghz	7mm	3mm
SC	RCP:	3757	4000	286	309	529	571	235	321	336	559	598	1066	883	1527	—
	LCP:	4780	5333	765	306	539	575	228	313	322	526	694	1086	894	1646	—
HN	RCP:	4848	4167	309	318	348	359	250	333	393	600	452	851	784	2060	—
	LCP:	3813	4184	327	330	349	349	254	316	380	584	426	859	768	1849	—
NL	RCP:	2167	2173	293	345	333	367	223	290	269	357	621	643	595	982	4842
	LCP:	2154	2103	315	357	349	361	202	282	265	380	583	652	623	910	5345
FD	RCP:	3615	4489	264	282	455	446	193	274	345	466	725	817	743	1352	4240
	LCP:	4818	7449	268	282	321	309	191	267	360	490	757	827	800	1232	5222
LA	RCP:	2133	2305	269	270	406	409	214	264	287	373	491	427	432	948	2087
	LCP:	2159	2235	275	260	365	375	205	257	250	336	526	525	485	1054	3186
PT	RCP:	2875	2638	255	287	290	301	204	271	281	388	532	453	393	989	2244
	LCP:	2775	2774	263	275	296	309	198	248	303	383	486	457	388	966	2395
KP	RCP:	2037	2167	263	261	276	289	237	340	246	362	482	460	452	1056	3931
	LCP:	2110	2326	263	272	275	301	234	306	253	351	533	404	385	933	4037
OV	RCP:	2299	2258	333	365	267	289	212	305	295	371	688	594	563	883	7571
	LCP:	2035	2061	305	333	290	303	229	330	372	453	550	657	535	953	5457
BR	RCP:	2391	2125	259	287	294	297	209	266	322	411	535	617	504	1155	4677
	LCP:	2222	2121	263	278	299	302	203	260	336	431	504	616	522	1167	4571
MK	RCP:	3015	2347	317	519	391	411	228	289	345	458	439	379	387	885	5118
	LCP:	2846	2469	306	427	407	416	239	311	348	480	557	357	410	1000	4895
AVG:		2942	3086	310	318	354	367	220	292	315	438	559	637	577	1177	4364
Sensitivity data for the OSS																
BW: Bandwidth total (twice BW per polarization) (MHz).																
FF: Fringe fit time in seconds.																
LT: Long track integration time in hours.																
FF Sen: - baseline sensitivity in fringe fit time (mJy).																
LT Sen: - array sensitivity in long track time (microJy).																
St	P	90cm	50cm	21cm	18cm	13cm	13cmsx	6cm	7ghz	4cm	4cmsx	2cm	1cm	24ghz	7mm	3mm
BW MHz:		64.	8.	512.	512.	512.	512.	512.	512.	512.	512.	512.	512.	512.	512.	512.
FF sec:		60.	60.	120.	120.	120.	120.	120.	120.	120.	120.	120.	120.	120.	60.	30.
LT hr:		8.	8.	8.	8.	8.	8.	8.	8.	8.	8.	8.	8.	8.	8.	4.
FF Sen:		42.0	124.5	1.1	1.1	1.3	1.3	0.8	1.0	1.1	1.6	2.0	2.3	2.1	5.9	31.1
LT Sen:		285.5	847.1	10.6	10.9	12.1	12.6	7.5	10.0	10.8	15.0	19.2	21.9	19.8	40.4	268.5

Figure 2: The SEFD table from a version of *gains.table*.

efficiencies, and the SEFDs. As an example, the table of SEFDs followed by averages extracted specifically to simplify filling in the table in the Observational Status Summary is shown in Figure 2.

The plots have one page for each of T_{sys} , gain, and SEFD for each band. Each page has all stations for easy comparison. The plots are called *tsplot_all.ps* and *tsplot_recent.ps*. The full time range plot goes back to 1992 so it is an interesting place to look for long term trends. For example, it is clear that one of our neglected systems, 50/90cm, degraded by around a factor of 2 at HN and SC in 2007-8 as did FD last year. The reason is unknown but is probably related to the dipoles. It is sometimes apparent when something probably went wrong with the T_{cals} . An example is LA at 18cm for which, in about 2010, there was a sudden onset of a large difference between RCP and LCP T_{sys} and efficiency, but not in SEFD. That sounds like the T_{cals} got swapped in the *tcals* file. Unfortunately there are many such oddities and we have by no means chased down and fixed all of them. At the higher frequencies, especially 1cm, there is a very strong annual effect that is very obvious in the plots of T_{sys} . It might be

useful, for when trying to estimate sensitivity, to use not a single predicted Tsys value, but one that is a function of time of year.

It is a very good idea to examine the tables and plots produced by *gethistory* when it is run. Often oddities like those described above are easier to spot in the averages and/or long term plots than when just looking at a snapshot in time.

Subreflector focus and rotation measurements:

Occasionally on general principles, and any time a receiver is changed or there is work done on the FRM, the optimum focus and rotation values should be measured. This is currently done, rather inefficiently, by doing full pointing patterns at the a priori position and at half power points in both focus and rotation. Thus a minimum of 5 full pointing patterns are required for a single measurement so it takes time, and is subject to data loss if one pattern fails. The pointing patterns are done at each focus and rotation position to deal with the possibility that the optimum pointing position is a function of one or both. In practice, the pointing does not change significantly with focus, but does change, in a reasonably predictable manner, with rotation. So an alternate scheme for the future might be to start with a reference pointing, then do on-off pointing at each FRM offset position. In principle SCHED knows how to create on-off schedules and FIT knows how to analyze them, but those options have not been exercised in very many years so they would need to be tested. Plus the analysis program would need modification.

The focus rotation measurements following the current scheme are scheduled using SCHED in a manner very similar to making a regular pointing schedule. The parameters that differ are ROTPAT, ROTOFF, and FOCOFF. ROTPAT tells SCHED how many different focus/rotation positions to observe while ROTOFF and FOCOFF give the offset in rotation and focus respectively for each position. The units are relative to a band dependent shift supplied by the program. A typical pattern is to start with a negative focus offset with zero rotation offset, then a negative rotation offset and no focus offset, then no offsets, then a positive rotation offset with no focus offset, and finally a negative focus offset with the rotation centered. One needs to be sure that there is adequate time specified for the “scan” to allow all the pointing patterns. The SCHED inputs for such a pattern (from *do7mmrot.com*) are

```
rotpat = 5
focoff = -1.0, 0, 0, 0, 1.0
rotoff = 0, -1.7, 0, 1.7, 0
dwell = (5*5:30) + 15
ptdur = 15
ptslew = 15
```

Any of the standard scripts that have “*foc*” or “*rot*” in their name can be used as an example of how to use these parameters. They are also described in the SCHED manual. SCHED will not schedule focus and rotation patterns for 50, 90, or 20 cm bands. The results would likely be worse than using the a priori positions and the radius of the feed circle is small compared to the movements needed so the pattern would be a bit problematic.

ROTANAL is the program used to reduce the focus/rotation patterns. First FIT is run in the same manner as for normal pointing patterns. The one difference is that there are a lot of FRM lines in the output as the focus and rotation are adjusted. ROTANAL is one of the relatively simple programs that

simply asks questions of the user and gets answers from the terminal, or from a script. The questions are:

1. Input File Name (From FIT):
2. Output Text File Name: (usually will be the experiment code with extension *.txt*)
3. Plot File Name (use *nnn.ps/vps* for postscript):
4. Band to process:
5. Plot options (choose among 1: Summary, 2: Every focus pattern (focus only data), 3: Every rotation pattern (rotation and focus data). Typically give 1 3).

The two output files are the plot file and the text file. The plot file starts with the same page of weather plots that PTANAL provides. Then there is a summary of the focus and rotation results. The top four plots are the azimuth (left) and elevation (right) collimation offsets as a function of rotation (upper) and focus (lower). The typical result is that the collimation offset does not depend much on the focus, but is sensitive to the rotation. On the plots vs rotation, the a priori model (nominal rotation adjusted by the offset in rotation times the expected variation of 0.6666 arcminutes per degree of subreflector rotation) is shown as a dashed line while the fit to collimation offsets vs. rotation are shown as a solid line. The fit is a simple linear slope. Examination of some of the plots for lower frequencies shows a very non-linear effect in the model and the actual data showing that the linear approximation is not so good when the curvature of the feed ring kicks in strongly. Unfortunately, while the model and the actual data have similar characteristics, they don't reproduce all that accurately. If we do shift to a on-off scheme with blind pointing after a reference point, the model will need to be improved. Note that the model pointing offsets are used by SCHED to try to center the pointing patterns.

The bottom two plots of the six on the second page are the derived optimum focus vs elevation and the optimum rotation vs elevation. There is a clear elevation dependence of both optimum focus and rotation. The plots show the a priori, average, and fitted curves. They also list the "Tolerance" which is the amount by which the focus or rotation can deviate from optimum and still have the amplitude stay within 1% of the optimum value. If the deviations are larger than that, improved performance is available from better setting of focus or rotation.

For actual operations, we adjust the focus with elevation for the optimum setting at 2cm and shorter wavelengths. The necessary inputs are there in the control system and equipment records to do the same for the rotation. When it was tried early in the VLBA history, it ran into an issue that the FRM would not successfully make small position adjustments. It was fine when switching bands. But the small tweaks required for using the rotation vs elevation adjustments often failed, and raised out-of-position flags that caused other problems. So it is not being done. In recent years, improvements have been made to the FRM so there is a good chance that the rotation vs elevation could be implemented now. Someone just needs to try. ROTANAL provides the values and all the tools are in place to use them. But it just has not risen sufficiently high on anyone's priority list yet to get done.

The rest of the plot pages show the amplitudes (Tant) vs focus or rotation for the patterns. These are the raw data used for the fits. If all is well, the central cluster of points should be the highest, with the side ones down. If the pattern is symmetric, the a priori value was good. It is not uncommon for one of the edge settings to be high, which means the pattern is not centered, or there are other issues with the data. If there are a lot of those, it is advisable to change the a priori setting in that direction, then repeat the observations.

An example file, chosen at random, is */home/vlbiobs/pointing/nov14/FOCROT/rnov07q.ps*. That, or any other .ps files in that, or similar, directories can be examined to get a feel for the data. Note that the focus/rotation results are usually in *FOCROT* subdirectories of the monthly pointing directories.

The other file from ROTANAL is the text file which is traditionally given the extension *.txt*. That file has a segment for each station even when there are no data. For a station with data, first the fitted linear equations relating the collimation offsets to the rotation are given. These could be fed back to scheduling, especially if going to an on-off scheme. Next comes a table of results from focus patterns (which might be embedded in the focus/rotation patterns). There is a line for each source/polarization starting with bookkeeping data such as time, elevation, rotation and pattern numbers. That is followed by the peak amplitude, derived focus, and focus width, all with fit errors. Flags indicate if the pattern is good for focus and also whether it is measured close enough to the rotation peak. Then the a priori focus position based on the equipment records is given, along with the difference between the derived and a priori values (called "ERROR"). That is followed by the focus value constant term when used in an equation with a slope variation of 0.75 times $\text{Cos}(Za)$. At the end of the line, the station and band are given to aid with using grep on files.

After the focus table is a summary of fit results which are self explanatory, but basically are results for various values of the focus variation with zenith angle. A fairly reasonable value for most antennas is 0.75 so that value is forced for some fits. But the last line gives the results with a fitted dependence on $\text{Cos}(Za)$. In practice, each station does use different $\text{Cos}(Za)$ parameters so likely those showed up at some time on that last line.

Next on the page is the table of rotation results which are similar to the focus results. Typically the scans are the same. Results are not shown for rotation vs zenith angle. That may need to be added some day if the use is implemented. The program clearly does the fit because the results show on the second plot page, but they have not made it to the printout.

The station data ends with the nominal (input) and average derived rotation values plus the azimuth and elevation collimation offsets that will need to be set if the average rotation is accepted. Note that, without the collimation adjustments, any change in the rotation setting will introduce a pointing error. Also note that a change in the rotation will likely have a significant effect on the gain curve so that will need to be redetermined.

Focus is the subject of VLBA Test Memo 67 by R. C. Walker. (2001)

Wideband tests:

There are two kinds of wideband test. The first is the subject of this section. These are single dish pointing observations designed to measure the system temperature, gain, and SEFD at a moderately dense sample of frequencies across the full receiver bandwidths. The other type are actually RFI tests that try to obtain autocorrelation data over the full range of frequencies available to the receivers. Those are interferometer tests and are evolving quickly at the moment so they will not be covered in this memo, which is about single dish observing.

The processing of wideband tests happens in monthly directories under */home/vlbiobs/wideband*. Wideband tests are much like normal pointing observations, although usually are confined to a small number of sources, perhaps only one in the ideal case. They do have more and rather different setup files that are used by SCHED compared to normal pointing. The VLBA system has IFs that are 512 MHz wide. Therefore, the maximum frequency range that can be covered at any one time is 512 MHz.

The scheme for doing wideband tests is in something of a state of flux as of this writing. The standard scripts for making the observations are meant for use with 8 legacy BBCs. The sites only still have 4. The RDBE based FIT program uses the RDBE data (either DDC or PFB) but only does as many channels as exist for the legacy system. Eventually, the RDBE based FIT should be able to do all of the specified PFB channels, or better yet, all 32 channels actually derived by the PFB filters. That is a future project. For now, the scripts should be rebuilt for 4 channels, which will mean fairly coarse sampling, or a lot of scans. Regardless, the channels specified can only spread across the 512 MHz VLBI IFs. Multiple setups are needed to cover the full frequency range of most receivers. For the cases, such as 6cm, where 4 IFs are available, they can be used and with fewer frequencies covered within each IF.

The command files used to generate the wideband schedules are kept in the same place as other standard SCHED command files (*/users/cwalker/bin* for now) and maintained in svn at <https://svn.aoc.nrao.edu/repos/VLBA/Pointing/trunk/Scripts/Pointing>. The setup files are contained within the command files. As noted above, the command files need updating for the current reduced capabilities.

For these tests, PTANAL is run in close to the standard mode but with inputs that either allow the full frequency range of each receiver to be used, or don't even specify frequency ranges. Two scripts of interest are *ptanalw* and *wptanal* although the latter is somewhat out of date and probably should not be used. *Wptanal* is not in the svn area and is kept in */home/vlbiobs/bin*.

The graphical version of the results is shown on the last pages of the PTANAL plot output files where the Tsys, gain, and SEFD are plotted against frequency. There is also a second plot of the squint terms and the beam width as a function of frequency, more for diagnostic purposes.

The wideband tests could, with a small amount of software work, be used to help derive new Tcal values as described in the Tcal section earlier.

PN3DB tests:

The PN3DB tests, for Peter Napier 3db (half power) tests (originally used on the VLA), are a way of testing the stability of the telescope tracking. The basic idea is to point at the half power position in either elevation or azimuth for an extended period and watch the power level. This is done on a very strong source. At the half power position, very small fluctuations in pointing position cause large power changes – much larger than would be seen if the source were at the center of the beam because you are on a steep part of the power vs pointing position curve. To first order, it is only pointing fluctuations in the direction away from the source position that count because those fluctuations are along the beam gradient. Pointing fluctuations across the gradient don't cause power changes until

second order effects kick in. For this reason, tracking at the elevation half power point tests the elevation tracking. Tracking at the azimuth half power point tests the azimuth tracking. The encoders have various natural wavelengths related to the mechanical structure of the device which can cause resonances. The analysis is set up to find those resonances.

There is a script, kept with the other scripts discussed earlier, called *pn3db.com* which can be used to schedule pn3db tests. The script invokes the PN3DB mode in SCHED. From the description in the SCHED manual: PN3DB tells SCHED to write a VLBA tracking test sequence in the VLBA control file for this scan. NOPN3DB tells SCHED to stop writing such sequences. This parameter should only be of interest to VLBA operations staff. The sequence, which is not a loop, starts with a scan of length PTSLEW to set levels at the half power point, then does a 15 second scan at the same place to fix those levels, then does one scan each of length PTDUR at off source and on peak positions, then tracks the source at the azimuth half power point. That is followed by an off source, on source, and half power in elevation tracking sequence. PTINCR in the setup file is used to set the distance to the half power point. The periods of tracking are the full scan duration minus all the level fixing, off, and on scan durations, divided by two. Scan durations of about an hour are appropriate.

The source used for a pn3db test should be very strong and, for maximum sensitivity to tracking issues, the frequency should be high. Probably the best source to use is the water maser in W49 at 1cm, which when observed with a 16 MHz bandwidth, triples or more the system temperature and has a flux density integrated over 16 MHz (the lines are much more narrow) of about 2400 Jy.

The pn3db data are analyzed using PN3ANAL. That program reads the power data from the output of program SARA which was described earlier. PN3ANAL asks for user input from the terminal. The inputs are:

1. The input SARA file.
2. The output plot file name. Used by pgplot so should be of the form nnn.ps/ps for a landscape postscript file.

PN3ANAL accumulates the data and creates 2 pages with a total of 5 plots for each full scan length (one hour in the scenarios described above). The first 3 plots are the ratio of total over switched power (closely related to Tsys) plotted against time, azimuth, and elevation. The last two plots, which are on the second page, do a Fourier analysis in wavelengths (in degrees of pointing direction) of the power data plotted in the first 3 plots. Any systematic oscillations should show up. Certain specific frequencies are marked, and are often the locations of peaks in the spectrum. In the first page plots, there is a tick that marks the range of amplitudes that would be seen if there is a pointing change of 10 arc seconds.

The special frequencies marked relate to the fine vs coarse encoder. The encoder is a 512 pole inductosyn which has a characteristic wavelength of $360/512 = 0.703$ deg. Outputs are signals for the sine and cosine for 1/256 turn. If there is an imbalance in the gains of those channels, there can be systematic errors that show up as ripples vs pointing direction. Wavelengths are marked that are at 1/4, 1/2, 1, and 2 times that value the 0.703 degree wavelength and peaks are often seen there. Also 1/816 of a turn (0.441 degrees) and half that are marked as there is sometimes a signal there too. I'm not sure of the reasons for those signals. Pn3db tests are discussed in VLBA Test Memo 64 by C. Walker (2001) and VLBA Test Memo 65 by W. Brisken (2001). When there are big oscillations, they can be fixed by balancing the encoder channels.

Summary list of observation types:

The following is a quick summary of the types of tests that are/should be done routinely. The tests are scheduled and analyzed by VLBA operations, mainly Jim Ogle in current practice. I have provided scientific oversight, established the nature of the tests, provided the scripts and provided most of the programs used. I try to look at the results occasionally, especially the major pointing equation, gain, and focus/rotation results. Operations can identify many of the issues revealed, but it is also important for someone with a science background to look through the data occasionally as the urgency and/or cause of some problems is not always obvious. So far, my successor in this area has not been identified, although Vivek Dhawan knows enough about it to deal with crises situations.

1. **Startup pointing:** All antennas, a few scans per band for all bands. Done weekly after maintenance day to check telescope operation. Data contribute to gain and pointing determinations made from averaging results from many days.
2. **Receiver checks:** Usually one antenna at one band to check receivers and maybe determine gain after a hardware change. Do as needed.
3. **“KQ” tests:** Check pointing equation at between one and all antennas. These projects are at 1cm and 7mm and observe many sources to get good sky coverage. There should be at least one “KQ” run at each antenna per month.
4. **“QW” tests:** Specific test targeting the 3mm system, and the 3mm/7mm offset needed when 7mm reference pointing is used for 3mm observations. Do before global 3mm runs and occasionally as needed if the all band pointing (startup mainly) is not proving to be adequate.
5. **Focus/Rotation:** Measure focus and rotation values for one or more receivers at one or more stations. Should be done if there are hardware changes. Should be checked at 6cm and shorter wavelengths about every 3 months. Note the values are elevation dependent so either use the ability to change with elevation (focus is done this way) or pick a value that would be good at an intermediate elevation – be aware that this affects the gain curve.
6. **Pointing equation determination:** This is a processing operation where all data collected for several months are used to determine the pointing equation parameters. It should be done about every 3 months and should use at least startup and “KQ” type observations to have good sky coverage and have some data for all bands so that the collimation offsets are appropriate for the rest of the equation. Focus/rotation data should not be used.
7. **Gain determination:** Make the new *vlba_gains.key* file, new gains tables and new long term gain history plots. This uses all or most of the data ever taken. Should be done about every 3 months. Focus/rotation data should not be used.
8. **PN3DB:** These encoder tests should be done about every 3 months at each station.
9. **Wideband/RFI measurements:** These should be done for any new receiver and for all receivers at all stations about twice per year.

Summary list of programs and scripts:

Note that SARA and the legacy version of FIT are part of the standard legacy software package and are written in C. The new FIT for RDBE data was written by Walter Brisken. All other programs were written by Craig Walker in FORTRAN and are maintained in svn. The scripts were also written by Craig Walker.

1. **SCHED:** VLBI scheduling program (used world-wide).
2. **doxxx.com:** Many scripts for running SCHED to create various types of pointing schedules.
3. **FIT:** Reads monitor power data and produces pointing offsets, antenna temperatures, system temperatures and much other data. There are separate versions for legacy and RDBE data.
4. **SARA:** Legacy system monitor data dumper.
5. **PTANAL:** Pointing analysis program.
6. **ptanalxx:** Scripts to analyze various types of pointing observations with PTANAL.
7. **getfit:** Gather multi-session pointing data and run and analysis.
8. **getfit1:** Gather a station's data from a pointing analysis. Used by *getfit*.
9. **gethistory:** Update *vlba_gains.key* and create various summary plots and files.
10. **NEWGAIN:** Called by *gethistory* to do the update to *vlba_gains.key*.
11. **GAINTABLE:** Called by *gethistory* to produce a compact table to summarize the sensitivity information for a given time.
12. **PLOTTS:** Called by *gethistory* to produce plots of the sensitivity related information over time scales of years.
13. **ROTANAL:** Analyze focus and rotation patterns.
14. **PN3ANAL:** Analyze pn3db tests.

Summary list of file locations:

1. */home/vlbiobs/pointing* – Scheduling and analysis of pointing observations in dated subdirectories.
2. */home/vlbiobs/pointing/POINTING-EQUATION* – Analysis of blocks of months worth of pointing data.
3. */home/vlbiobs/pointing/mmmyy/FOCROT* – Scheduling and analysis of focus and rotation observations.
4. */home/vlbiobs/wideband* – Area for reduction of wideband observations.
5. */home/vlbiobs/pn3db* – Area for reduction of PN3DB observations. Not much there yet. Analysis sometimes done in personal areas.
6. */home/jansky/POINTING* – Where *gethistory* is run and *vlba_gains.key* is maintained. Also used for some other pointing related activities.
7. */home/jansky3/vlbaops/TCAL* – Where Tcal files, rail files, and *vlba_gains.key* are kept and accessed by various tools from the on-line system to FIT and PTANAL.
8. */users/cwalker/bin* – Accessible location of many scripts and some program executables. This should be moved some day.
9. <https://svn.aoc.nrao.edu/repos/VLBA/Pointing/trunk> – SVN repository location for most of the pointing related programs and scripts provided by Craig Walker.
10. <https://svn.aoc.nrao.edu/repos/VLBA/rcwlib/trunk> – SVN repository for slalib, odrpack, and sublib (my routines, some from SCHED, some specific to the pointing programs) routines that are used by many of the pointing routines.

Note that PGPLOT is used by most of the routines, so it needs to be available to the linker.

Summary list of files:

This is a list of the main files associated with pointing and gains along with their locations and formats.

For the formats, mostly the reader is referred to the detailed sections earlier in this memo. These are mainly the files that are used by multiple programs or users or that might need attention as part of a switch from the legacy control system to the Executor based system at the VLBA antennas. Some comments related to that switch are included. The many SCHED output files and the various bookkeeping files associated with GETHISTORY are described earlier in the memo, but not noted here. Other than the *.vex* file from SCHED, they are not expected to change.

1. **SCHED scripts:** These are the keyin input files for SCHED or, more commonly for pointing, the scripts that create and run SCHED input files from simplified input. The details for the SCHED inputs are given in the SCHED user manual. The master scripts are mostly in the SVN area listed above and are simply C shell unix scripts. There is usually a copy of the one actually used in area where the observations were scheduled as described in the locations section above. There they are owned by vlbaops but can be read by anyone. A very big issue for the switch to the new control system is how to specify pointing scans. With the legacy system, capabilities available to SCHED through the *.crd* files, such as loops and the ability to explicitly set incremental collimation offsets and focus/rotation values, are used. Those are not available in VEX so something rather different will be needed. That has not yet been designed.
2. **FIT output:** The main FIT output file (often with an extension *.tsm* for older data) is described in detail in the FIT section above. The station dependent files giving power levels through the pattern are also described in the FIT section. These files are usually left in the dated subdirectories of */home/vlbiops/pointing*, although they can also be found in */home/vlbiops/pointing/POINTING-EQUATION*, */home/vlbiops/pointing/mmmmyy/FOCROT*, and *home/vlbiops/wideband*. There they are usually owned by vlbaops but can be read by anyone. The main fit output file is accessed by PTANAL and ROTANAL. A FIT for the new system already exists so these files should not be an issue in the control system update.
3. **SARA files:** These are used for the PN3ANAL program and for examining legacy monitor data by hand. See the SARA and PN3DB sections above for more information. These files can appear in a variety of places, depending on the use and generally have a name extension of *.sara*. As we convert from the legacy control system to the executor based system, these files will likely need to be replaced by a combination of the interactive monitor data displays that already exist and some to-be-determined file to pass the required information to PN3ANAL. I would not suggest trying to duplicate the SARA files.
4. **Tcal files:** These files are described in the TCAL section above and are maintained by operations (user vlbaops) in the */home/jansky3/vlbaops/TCAL* directory. They are needed by anything that converts powers to system temperatures, which includes the legacy and new operator interfaces, FIT, TSM, and the programs that create the calibration transfer data.
5. **Rail files:** These hold the rail height survey data for all stations for a date. For each file, one station is likely to have new data. They are also in the */home/jansky3/vlbaops/TCAL* directory. They are generally used once when they are first read during a PTANAL processing run. PTANAL reformats to a form appropriate to feed to the on-line system, mostly by interpolating to integer 3 degree points in azimuth. The rail file lines show the actual bolt azimuths and the height in mm. See the FIT section for more about this.
6. **Gain files:** *vlba_gains.key* and *vlba_gains.master* are described in the "Gain history" section above. They are keyin format (same as SCHED input) files that hold the derived gains, gain curves and other information for the antennas. They are built and maintained in */home/jansky/POINTING*. The current released version of *vlba_gains.key* available at */home/jansky3/vlbaops/TCAL* for use by the calibration transfer programs and others that might

need it. PTANAL can use it there but the scripts tend to look for it in */home/jansky/POINTING*. It is made available to users via ftp at */home/ftp/pub/vlba_gains.key* since it can be used by AIPS to rebuild gain tables.

A few To Do items:

This is by no means comprehensive, but is just some items that came up while writing the memo.

1. Modify PTANAL to complain more vigorously if no equipment records are found. It already complains, but it shouldn't finish.
2. Do some PN3DB tests.
3. Activate and test the rotation vs elevation variation.
4. Come up with a way to control pointing observations once the VME is gone. VEX does not allow the level of control used by the current system.
5. Redetermine the slew rates. I think SCHED's values are out of date.
6. Rationalize the locations of all operations software.
7. Update the wideband command files for the reduced BBCs available.
8. Set up to use all VEX channels, or better, all 32 PFB channels in wideband pointing.
9. Make it possible to use the frequency difference method for reference pointing.
10. Revisit the details of the pointing source catalog that comes with SCHED. Positions, especially of masers should be checked.