Sensitivity, Calibration, and Self-Calibration

Fourth INPE Advanced Course on Astrophysics: Radio Astronomy in the 21st Century

Rick Perley, NRAO/Socorro

Topics

- Noise, Power, and Temperature
- Antenna Sensitivity – System Equivalent Flux Density
- Radio Telescope Sensitivity – Total Power and Interferometric
- Synthesis Image Sensitivity
- Aspects of Calibration
- From Calibration to Self-Calibration
- Isoplanicity and Full-Sky Calibration – the cutting edge …

NRAO

Introduction

- All lectures so far have presumed the systems are perfect – no noise, and no errors in the electronics or propagation path.
- This lecture deals with the 'real' world:
  - The electronics – although meeting our requirements in so many ways – add noise to the signal.
  - The antennas – although marvelous – add a direction-dependent dependency in their own phase and amplitude response.
  - The atmosphere is turbulent, which 'corrugates' the incoming phase front.
- All these issues can be addressed – some better than others.

Noise, Power, and Temperature

- A simple view of a radio telescope is shown below:

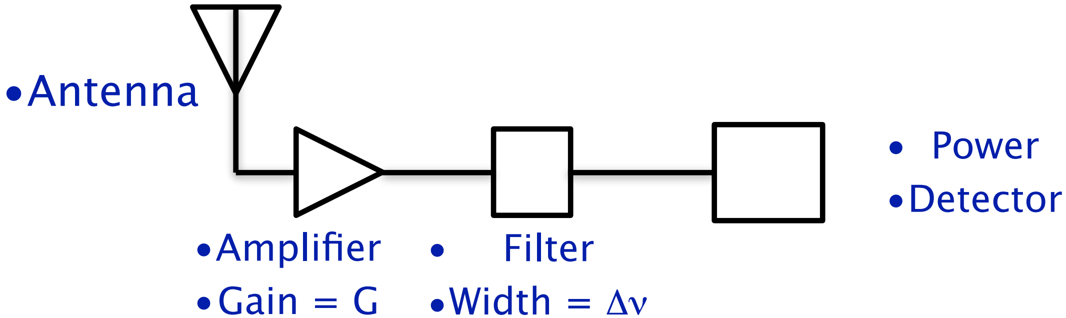

- Powers add, so that within the spectral width $\Delta\nu$, we can write for the power measured: $P_{sys} = G(P_{\nu,ant} + P_{\nu,rcv})\Delta\nu$
- Radio engineers (and radio astronomers) like to use temperature as a surrogate for spectral power. They can do this by:
  - A) Modelling the system with a matched resistor of temperature T_{ant} in place of the antenna, and an amplifier as a resistor followed by gain.

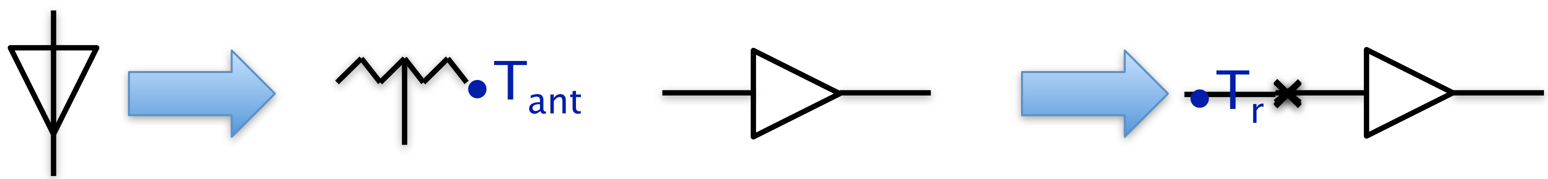

  - B) Using the relation: $P = kT\Delta\nu$, which describes the available power from a matched resistor of temperature T within bandwidth $\Delta\nu$.

System and Receiver Temperatures

- We can then talk about system temperature, rather than system spectral power, as a measure of system performance.
- By convention, all such 'temperatures' are referenced to the output of the antennas (input to the first amplifier).
- Then, we can write:

$$T_{sys} = T_{ant} + T_R$$

where the various temperatures are those of matched resistors, which, when located at the input, would produce the same noise power as that contributed by the device.

- Typically, the total receiver temperature is ~15 to 30K.
- And for most centimeter wavelengths, $T_R >> T_{ant}$

Antenna Temperature, and the 'System Equivalent Flux Density'.

- System Temperatures for modern telescopes are typically 25K.
- The Antenna Temperature is a measure of the power of the external source we are interested in measuring:

$$T_{ANT} = \frac{\eta_A A}{2k} S = KS$$

- Here, S = source flux density (watt/m/Hz), A is the antenna aperture area, and η = antenna aperture efficiency. k = Boltzmann's constant.
- For a 1 Jy source, and a 60% eff. 25-meter antenna, T_{ant}= 0.1K (!)
- A very useful measure of antenna sensitivity is the 'System Equivalent Flux Density':

$$S_E = \frac{T_{SYS}}{K} = \frac{2kT_{SYS}}{\eta_A A}$$

- In words: S_E is the flux density of a source which doubles the total system power (or system temperature).
- For an EVLA antenna, S_E~200 Jy (low freq.), and ~600 Jy (high freq.)

System Sensitivity – Total Power

- We can now move to determining system sensitivity.
- A simple argument provides the basic relationship:
  - In a band-limited signal of width $\Delta\nu$, there is an independent measure of the power every $\Delta\nu^{-1}$ seconds.
  - Hence, in T seconds we have $N=\Delta\nu T$ independent samples of the power.
  - Hence, the accuracy is: $\sigma_P = \dfrac{P}{\sqrt{\Delta\upsilon\, T}}$
  - From which, we deduce: $\sigma_S = \dfrac{S_E}{\sqrt{\Delta\upsilon\, T}}$ (units in Jy).
  - The square root factor can be large:
    - 10^6 for $\Delta\nu \sim 1$ GHz, T = 1000 sec
    - We can reach micro-Jy levels with GHz bandwidths and hours of integration.

Interferometer Sensitivity

- For a (digital) correlation interferometer, a somewhat more sophisticated argument provides the following relation.
- Let S_{E1} and S_{E2} be the SEFDs for the two antennas, and η_S be an 'efficiency' factor which accounts for losses due to the quantization and correlation.
- Then, analysis shows that for a single (COS or SIN) correlator,

$$\sigma_S = \frac{\sqrt{S_C^2 + S_T^2 + S_T(S_{Eq} + S_{E2}) + S_{Eq} S_{E2}}}{\eta_S \sqrt{2\Delta\upsilon\, T}}$$

where S_c is the fringe flux, and S_T is the total flux of the source.

- In almost all cases, $S_E >> S_T > S_C$ And, if all antennas are the same, then:

$$\sigma_S = \frac{S_E}{\eta_S \sqrt{2\Delta\upsilon\, T}}$$

Which shows that a 2-element interferometer has a noise ~41% worse than a single total power telescope of the same total collecting area.

These are simple gaussian statistics.

Complex Correlator, and the Noise Ball

- The two real correlator comprising the complex correlator are independent – the same noise analysis applies to both.
- In the complex Visibility plane, the visibility distribution is (ideally) a Gaussian 'ball' centered on the true source visibility, V.

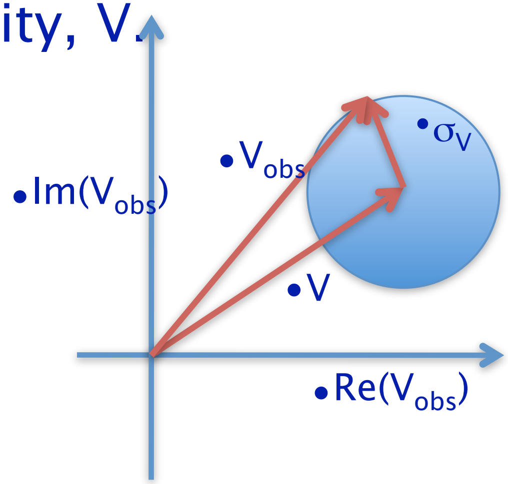

- One can utilize standard statistical techniques to generate PD functions of the amplitude and phase of the visibility, as a function of SNR.

And it Really Is This Way!

- I show here the Real and Imaginary parts of the visibility from a 10 Jy point source.
- The true visibility is 10 Jy (all real).
- The scatter is due to thermal noise.
- Each component should be gaussian, with equal width.

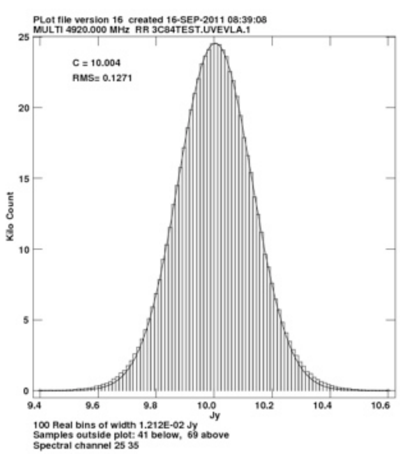

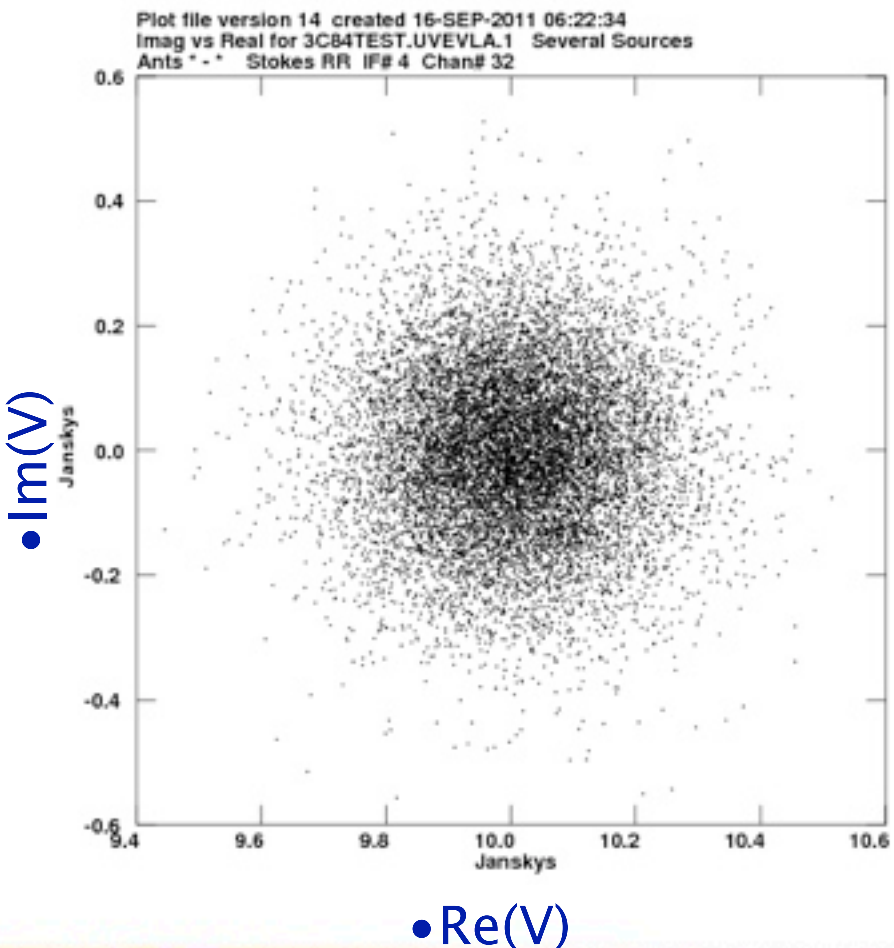

How's this for Gaussians!

- Shown below are the histograms for the real and imaginary parts …

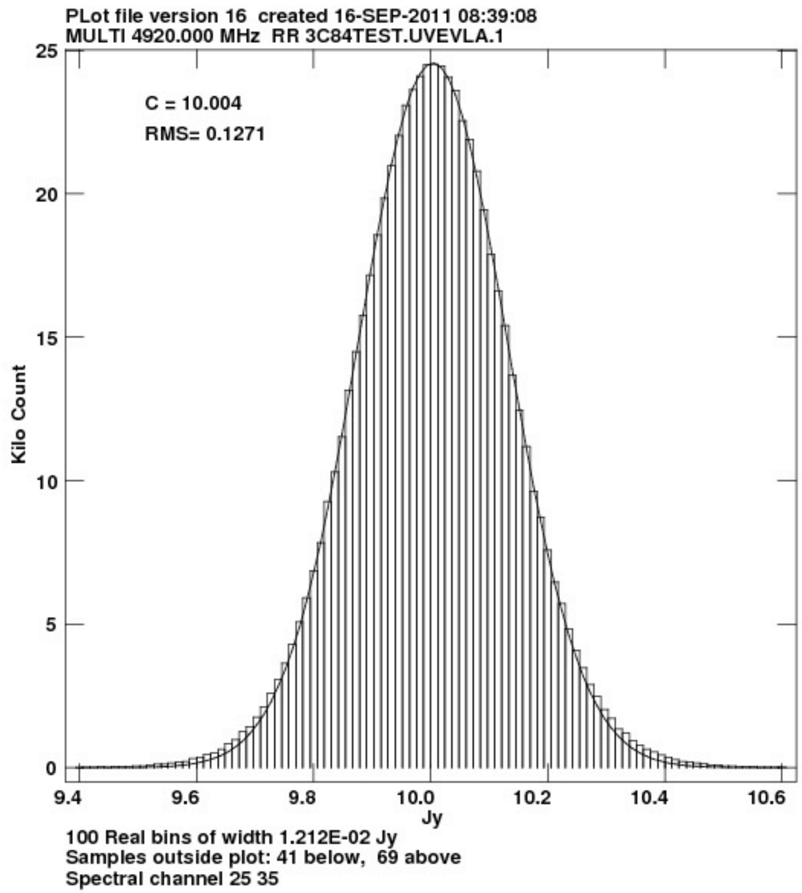

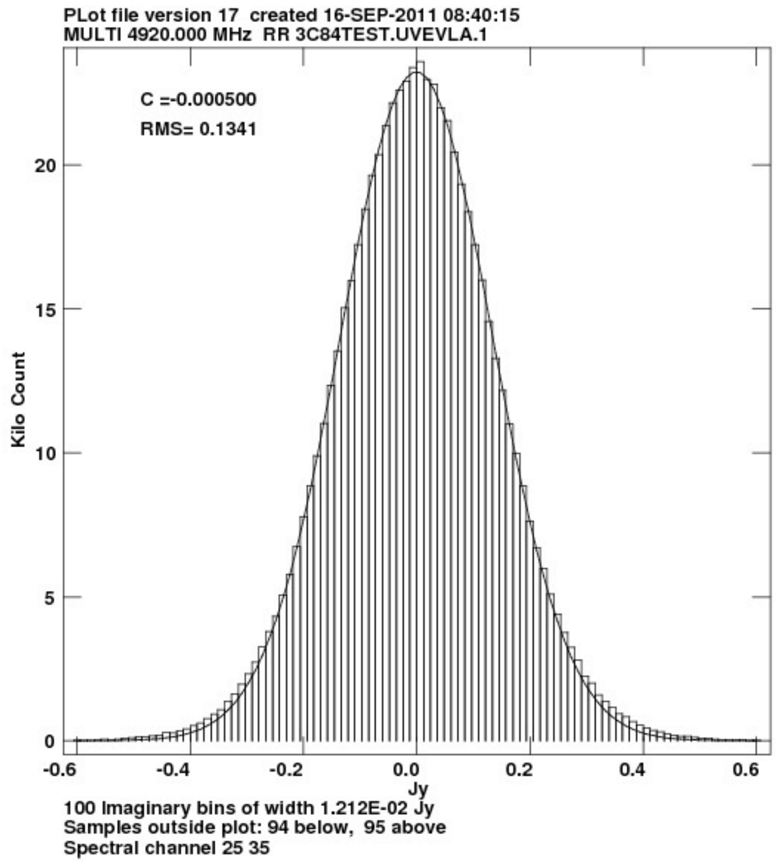

- Real Part Imaginary Part

Sensitivity of a Synthesized Image

- It is straightforwards to generate the expected sensitivity of a synthesized image, since the statistics of the visibility data are simple (gaussian).
- The easy case is 'natural' weighting with no taper, where every visibility is counted equally.
- Simple approach:
  - Noise is noise – every cell must give the same noise statistics.
  - Pick the central cell (l=m=0), for which $\exp[i2\pi(ul+vm)]=1$.
  - Then, for each baseline, the noise contribution at this cell is:

$$\sigma_S(0,0) = \frac{V(u,v)+V^*(-u,-v)}{2} = \frac{\sigma_C + i\sigma_S + \sigma_C - i\sigma_S}{2} = \sigma_C$$

  - Each of the $N_B=N(N-1)/2$ baselines has independent noise, so the noise in the image plane is:

$$\sigma_S = \frac{S_E}{\eta_S\sqrt{N(N-1)\Delta\upsilon\,T}} \approx \frac{S_E}{\eta_S N\sqrt{\Delta\upsilon\,T}}$$

Image Noise for Stokes I, Q, U, and V

- The preceding analysis holds for a single complex correlator (presuming natural weighting, and no taper).
- For Stokes I, Q, U, or V, we use two independent measurements: e.g.

  I = (RR + LL)/2*
- Each of the four correlations (RR, LL, RL, LR) should have independent and equal noise, so the noise in a real Stokes image is:

$$\sigma_I = \frac{S_E}{\eta_S \sqrt{2N(N-1)\Delta\upsilon\, T}}$$

- Numerous tests show excellent gaussian envelopes to the noise in blank field images, at the level expected from known values of the antenna SEFD, bandwidth, and time averaging.
- When tapered or non-natural weighting is used, the image noise rises, typically by 10 – 20% for the VLA.
- * This is the 'AIPS' definition of I, which is not the same as in standard physics texts. T The noise result is correct, nonetheless.

Array Brightness Sensitivity

- I close this section with a short but important discussion of sensitivity to extended sources.
- Consider an extended source of surface brightness I being observed with an interferometer of baseline B at wavelength λ.
- Introduce the Brightness Temperature: $T_B = \frac{\lambda^2}{2k} I$
- The synthesized beam solid angle is: $\Omega \sim (\lambda/B)^2$.
- The map amplitude is: $S \sim I\Omega = I(\lambda/B)^2 = 2kT_B/B^2$
- The map noise is: $$\sigma_I \approx \frac{S_E}{\eta_S N\sqrt{2\Delta\upsilon\, T}} = \frac{2kT_{SYS}}{\eta_S \eta_A N\sqrt{2\Delta\upsilon\, T}}$$
- The condition: $S \sim \sigma_S$ leads to a relation for the Brightness Temperature Sensitivity, σ_T: $$\sigma_T \approx \frac{B^2}{NA}\frac{1}{\eta_A \eta_C}\frac{T_{SYS}}{\sqrt{2\Delta\upsilon\, T}}$$

- Array Filling Factor
- Efficiency Factor
- Radiometer Sensitivity

Calibration

- Although the designers of radio telescope arrays strive to have the output data be in a form ready for imaging, this is rather difficult to achieve to the accuracy needed for today's demanding requirements.
- Most calibration must still be done 'off-line'.
- Instrumental/Atmospheric characteristics we must correct include:
  1. Delay Error
  2. Bandpass Function (instrumental frequency structure)
  3. Atmospheric phase irregularities
  4. Atmospheric absorption (especially at high frequencies)
  5. Antenna gain variability (especially at high frequencies)
  6. Antenna polarization impurities
- In this section, I briefly discuss each, with examples.

Calibration Formalism

- For a given visibility measurement, we can write, very generally:

$$V_{ij}^{Obs} = G_{ij} V_{ij} + \varepsilon_{ij} + \eta_{ij}$$

- Observed Visibility
- Baseline Gain
- True Visibility
- Correlator Offset
- Random Noise

- For a well-designed system, we can write:

$$G_{ij} = G_i G_j^*$$

$$\eta_{ij} = 0$$

- Giving us: $$V_{ij}^{Obs} = G_i G_j^* V_{ij} + \eta_{ij}$$

- Clearly, in order to determine the G terms, we must observe calibrator objects for which $|V| >> |\eta|$
- Calibration is most easily done with 'point' sources – for which every baseline sees the same visibility.

Least Squares Solution for Gains

- Suppose we observe an unresolved calibrator source, of flux density S.
- The calibrated visibilities should then give, for every baseline, every frequency channel, and at every time:

$$V = Ae^{i\phi} = S$$

- We then seek, for every antenna, every frequency channel, for every time, complex numbers G_i and G_j which minimize the following function:

$$\chi^2 = \sum_{i,j} \left| V_{ij} - G_i G_j^* S \right|^2$$

- These gain values are then applied to the target source, using some sort of interpolation.
- It is thus assumed that these solutions – valid for the calibrator – are equally valid for the target.
- More on this subject, later …

Major Calibration Parameters

- Although instrument builders strive to provide fully calibrated data, some errors (usually small) will slip through, requiring calibration.
- The important ones are (in the order which I use to calibrate them):

1. **Delay Error.** If the inserted delay (to equalize propagation times for the two paths) is in error by δt seconds, there results a phase slope over frequency given by:

$$\frac{d\phi}{d\upsilon} = 2\pi\delta t$$

- radians/Hz

- In words: A 1 nsec delay error causes a slope of 360 degrees over 1 GHz

2. **Bandpass** This is the change in amplitude/phase as a function of frequency. In general, we determine this for those frequency channels which are digitally generated for a single 'chunk' of analog spectrum. Some people consider bandpass calibration identical to amplitude/phase calibration. Bandpass functions are usually very stable.
3. **Phase Error.** There are two origins: Electronic (generally constant in time), and atmospheric (can be large and very fast). For both, the LSQ solutions from sources of known position provides a phase which can be applied to the data. This is the most difficult parameter to establish.

More Calibration Parameters

1. **Amplitude Error.** Same comments as for phase, except that the calibrators must have a known flux density. In general, well designed systems have a very stable amplitude gain.
2. **Polarization Leakage.** The polarizers used to separate the two orthogonal components of the EM wave (either orthogonal linear, or opposite circular) are not perfect. Some signal from one 'hand' leaks into the other. If stable, this contamination can be determined and removed. I'll discuss this in the polarization lecture.

- I show in the succeeding slides examples of each of these.
- The data are from an EVLA test observation (last week!) of the point source 3C84 at C-band (6cm wavelength).
- There are 3 minutes' data, in two short 1.5 minute scans.
- EVLA was in A-configuration – longest baseline 35 km.
- Two separate tunings with 1.024 GHz bandwidth each.

Delay

- Shown is the amplitude and phase of the bandpass of antenna #6.

- Frequency span is 1.024 GHz.
- Channel resolution is 2 MHz.
- Sloping line is phase.
- The smooth slope tells us there is a delay error.
- Easily measured and removed.

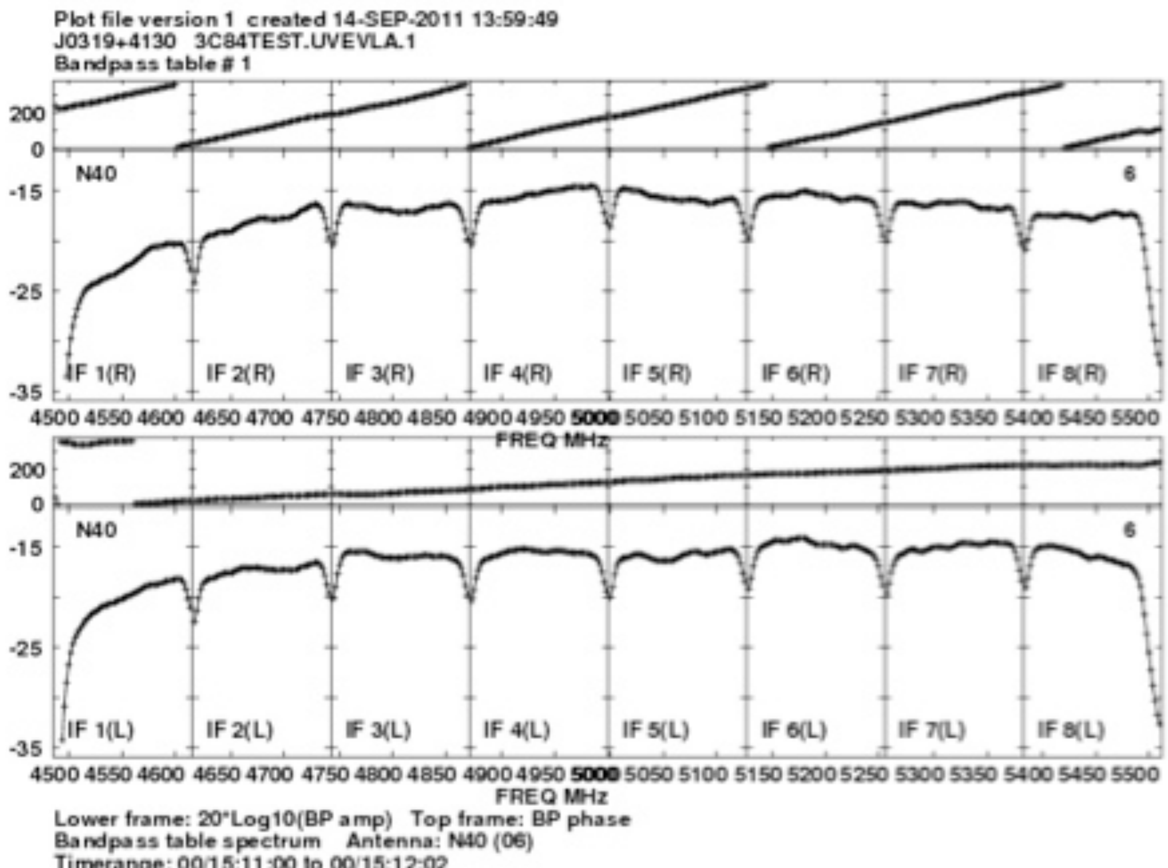

- RCP
- LCP
- 1 GHz

Bandpass

- Same antenna, but with the delay removed.

- Phase slope is now zero.
- Amplitude function shows 8 sub-bands.
- Notches are due to digital filters.
- The shape is very stable.
- One solution usually works for full day.

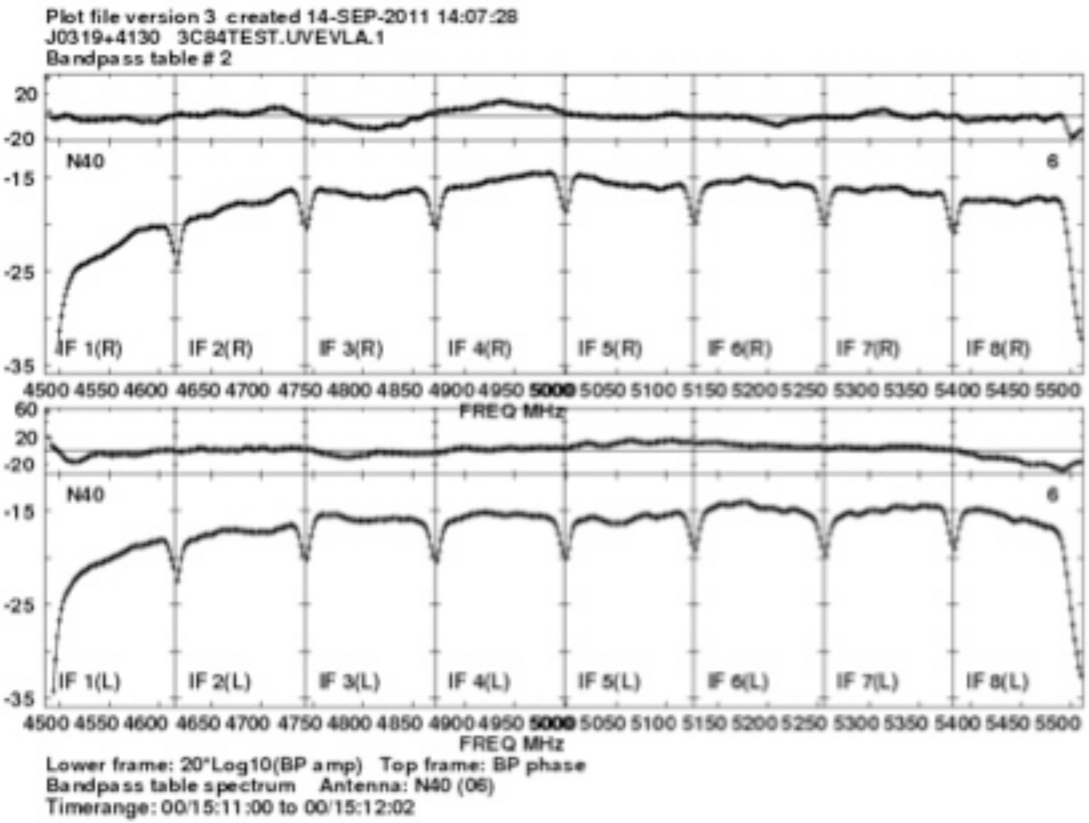

Phase and Amplitude

- After the delay and bandpass variations are removed, we must measure the variations in phase and amplitude.
- 3C84 is unresolved, and has 10 Jy flux density.
- All baselines must show same amplitude and zero phase, after calibration.
- Shown here are the amplitude solutions.
- Time interval is 1 second!
- Amplitude stability is very good!

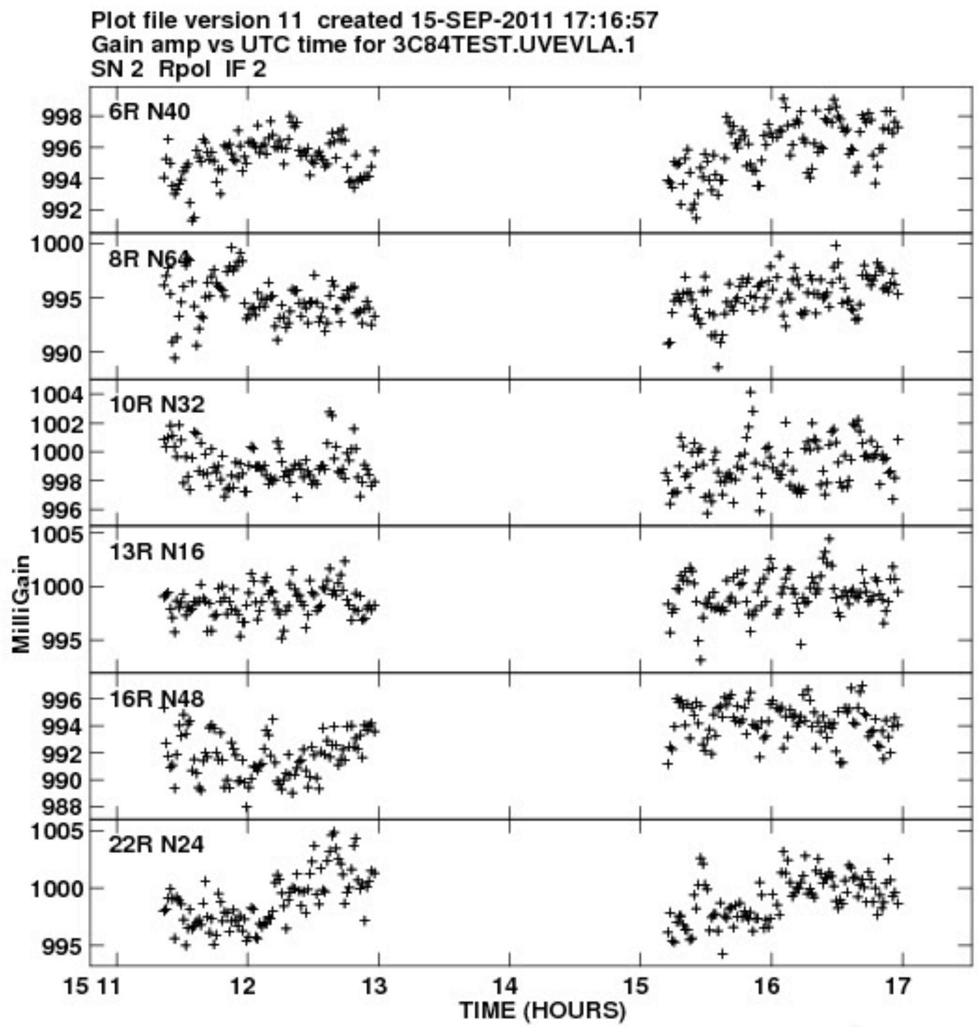

•6 minutes

Phase

- More critical is the phase – since it is more variable.

- Shown here are the 1-second phase solutions (reference is the array center).
- These 6 antennas are on the north arm, maximum is ~15 km.
- The ~10 degree variations are typical for a good day.
- Variations are almost certainly atmospheric in origin.
- Calibration of target source probably accurate to ~5 degrees – good enough for most purposes.

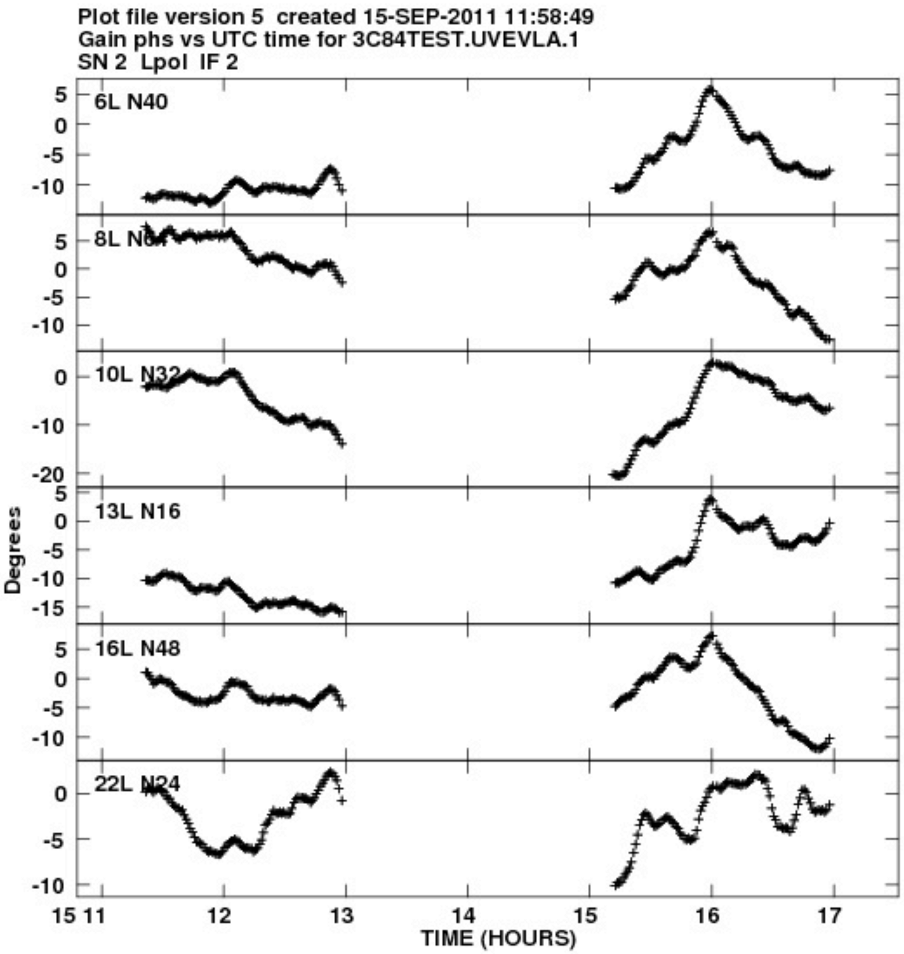

Polarization

- 3C84 is known to be unpolarized, so we can use it to determine the 'D' terms
- Here they are for antenna 1.

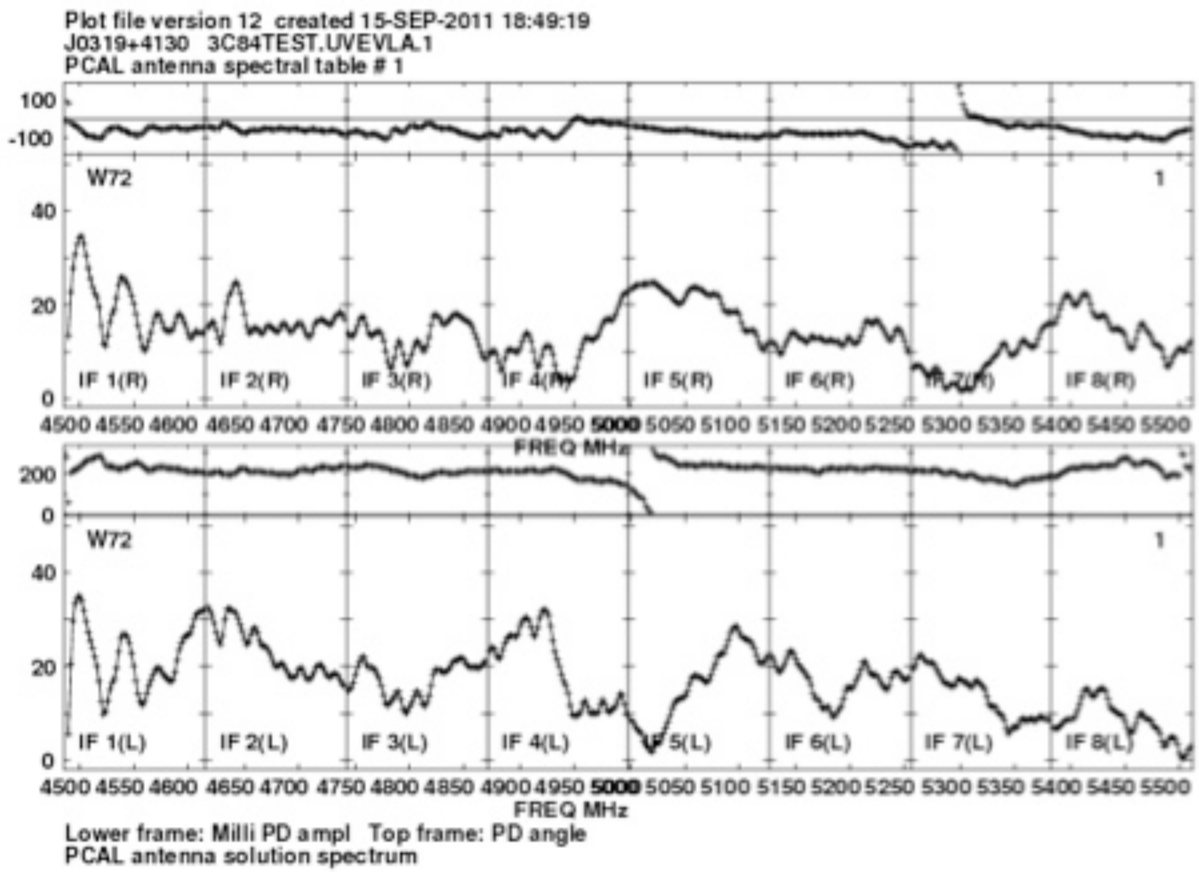

- Phase and amplitude of leakage of LCP into RCP.
- Typically 2%.

- Phase and amplitude of leakage of RCP into LCP.
- Typically 2%.

Result: Calibrated Visibilities!

- After applying these various phase and amplitude corrections, we can plot the visibility amplitudes (left) and phases (right).

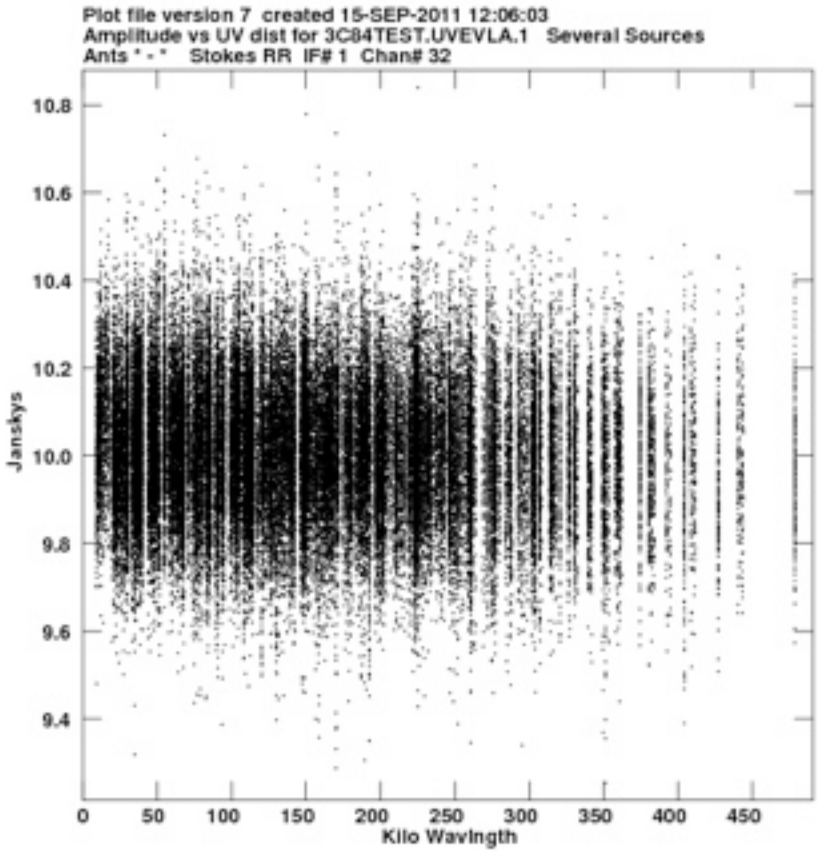

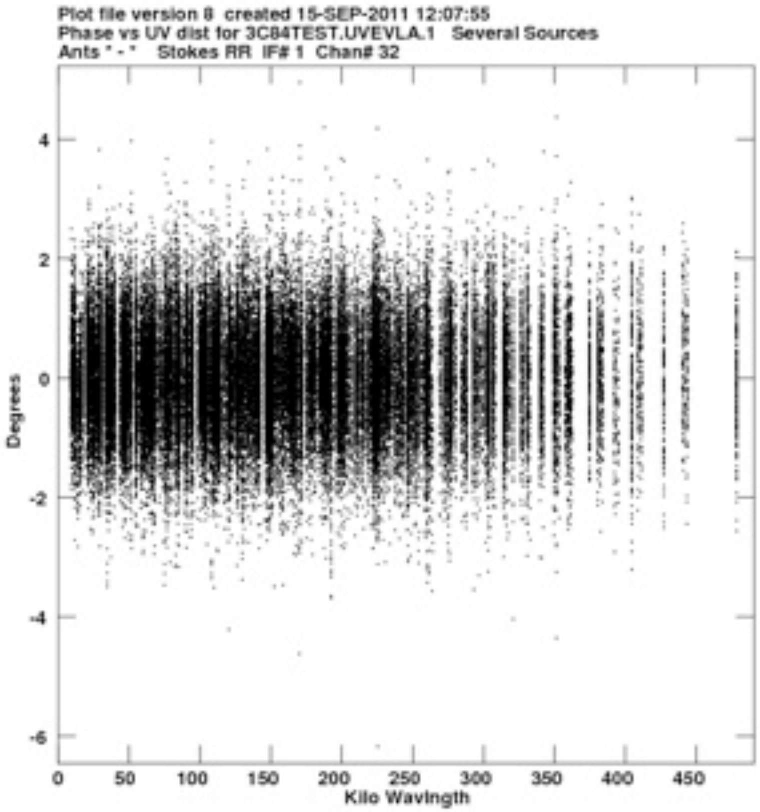

Self-Calibration

- Phase calibration is the most difficult aspect of array calibration.
- The phases determined on a calibrator at a particular time and place are not, in general, the same as those for the nearby target source.
- The differences can vary from a few degrees (not serious for much work) to many tens of degrees (disastrous for all).
- The cause is the unsteady atmosphere – this causes phase corrugations with angular scales of ~degrees, and time scales of seconds.
- Fast calibration cycles help – but are generally not sufficient for high accuracy imaging, for which phase calibration accuracy of less than 1 degree are needed.
- What to do?
- High dynamic range imaging means strong and bright sources – which implies good SNR – which makes us wonder: Can the source itself serve as its own calibrator?
- YES! --- but with some limitations.

Self-Calibration Rationale

- The array has N antennas – we are thus attempting to determine N-1 phases. (There is no absolute phase in interferometry --- all visibility phases are relative).
- The N antennas provide N(N-1)/2 measures. For even modest values of N, there are many more measures than unknowns.
- If we have even an approximate idea of the structure of the target source, we can use this information, to generate a estimate of the antenna phases, which can then be used to generate an improved model, which is used to provide improved calibration …
- In practice, for the VLA, this cycle converges astonishingly quickly.

Self-Calibration Formalism

- The idea is to use any approximate model for the source.
- This can come from another telescope, or it can be the result of making a map with regular calibration.
- Denote this model as I^M.
- From this model, we can generate (through Fourier Transform) model visibilities. Call these V^M.
- We can then, for each time interval and each frequency slot, solve for antenna-based gain parameters by minimization of:

$$\chi^2 = \sum_{i,j} w_{ij} \left| V_{ij}^{OBS} - G_i G_j^* V_{ij}^M \right|^2$$

- This returns values of G_i and G_j, which can be used to (re-)calibrate the data, which gives us a new image, which is used to generate new G_i, G_j.
- Does this converge! YES!!!

Self-Calibration, (cont.)

- In practice, we divide the observed (and calibrated) visibilities by the model visibilities, to form a ratio:

$$R_{ij} = \frac{V_{ij}^{OBS}}{V_{ij}^{M}}$$

- Then derive the new gains from minimization of:

$$\chi^2 = \sum_{i,j} w_{ij} \left|V_{ij}^{M}\right|^2 \left|R_{ij} - G_i G_j^*\right|^2$$

- This procedure has, in effect, converted the source into a point-source. When the model matches the (calibrated) data, the ratio R=1, and the (incremental) gains are also 1: G=1.
- This process works – it's like magic!
- Why does it work so well?

Why Self-Cal Works …

- There are two key reasons why self-cal (for the VLA especially) works so well:
  1. The system is over-determined: The number of baselines greatly exceeds the number of unknowns.
  2. The model and the data disagree.
- Both points are needed, but the latter is critical. The provided model cannot agree with the (FT of the) data for successful convergence.
- Suppose you made a snapshot image (single integration) of a source with a direct transform (no gridding). Self-calibrating with this cannot return any changes – the data and the model agree perfectly.
- Suppose you now make the image with an FFT, using the usual anti-aliasing gridding. For the shorter spacings, different baselines will end up in the same cell – these will have different errors, and the convolution/gridding will generate an image which disagrees with the data. Remarkably – this alone will generate improved calibration parameters.

Why Self-Cal Works (cont.)

- Now consider a long VLA integration. There are 351 baselines all moving through the (u,v) plane.
- A given (u,v) cell can have many different baselines passing through it, and they pass through at different times. Each baseline, for each time (while in that cell) have the same astronomical information, but each has different phase errors.
- The image is made from the average value – which is in disagreement with each of the individual visibilities from which it came.
- Convergence is possible – even using the 'Dirty' Image!!!
- But use of 'Dirty' images is not optimal – we can impose some real intelligence (our own!) through judicious use of deconvolved images.

Using Deconvolved Images for Self-Cal

- In practice, we attempt to deconvolve our 'dirty' images to provide a model for self-cal.
- Since deconvolution is constrained to match the observed visibilities, you might wonder why this helps.
- Answer: We don't use *ALL* the deconvolved image – only those parts we 'like'.
- What do we like?
  1. The all positive components. (Negative components are un-physical)
  2. Those components which lie in regions where we know the emission originates.
- We then use the acceptable CLEAN components for the model – not the full deconvolved image. This difference is important!
- The CC components will certainly disagree with the data, and convergence is greatly speeded up.

Finding and Correcting, or Removing Bad Data – a simple example.

- I show some ‘multiple snapshot’ data on 3C123, a fluxy compact radio source, observed in D-configuration in 2007, at 8.4 GHz.
- There are 7 observations, each of about 30 seconds duration.
- For reference, the ‘best image’, and UV-coverage are shown below.
- Resolution = 8.5 arcseconds. Maximum baseline ~ 25 kλ

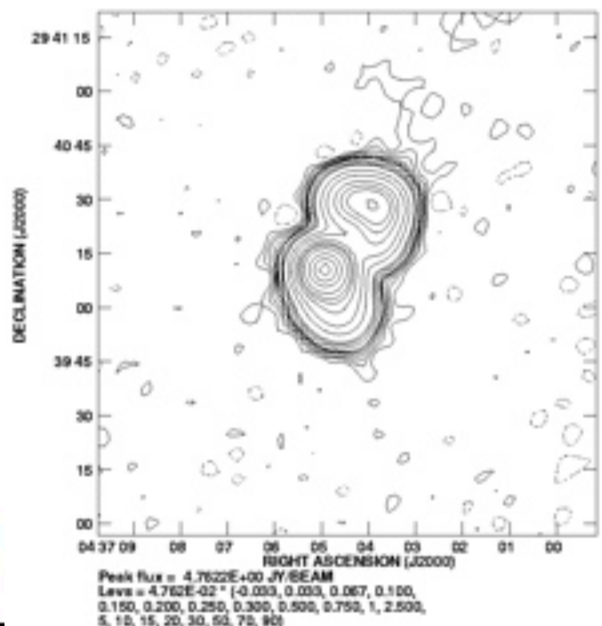

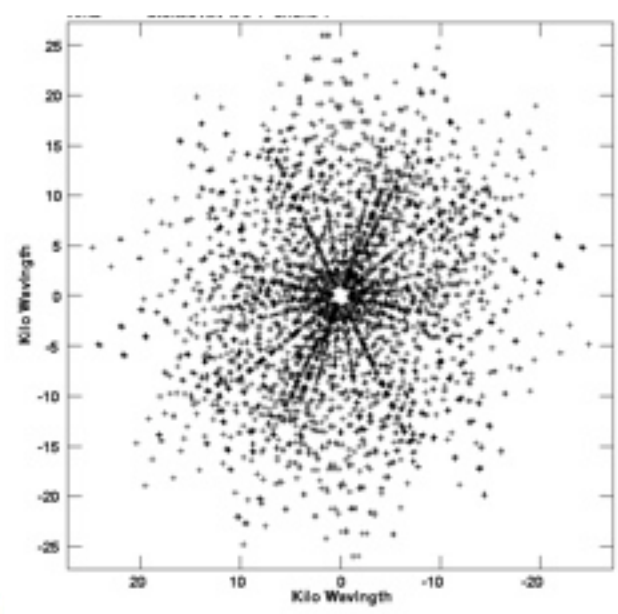

- Following standard calibration against unresolved point sources, and editing the really obviously bad data, the 1-d visibility plots look like this, in amplitude and phase:

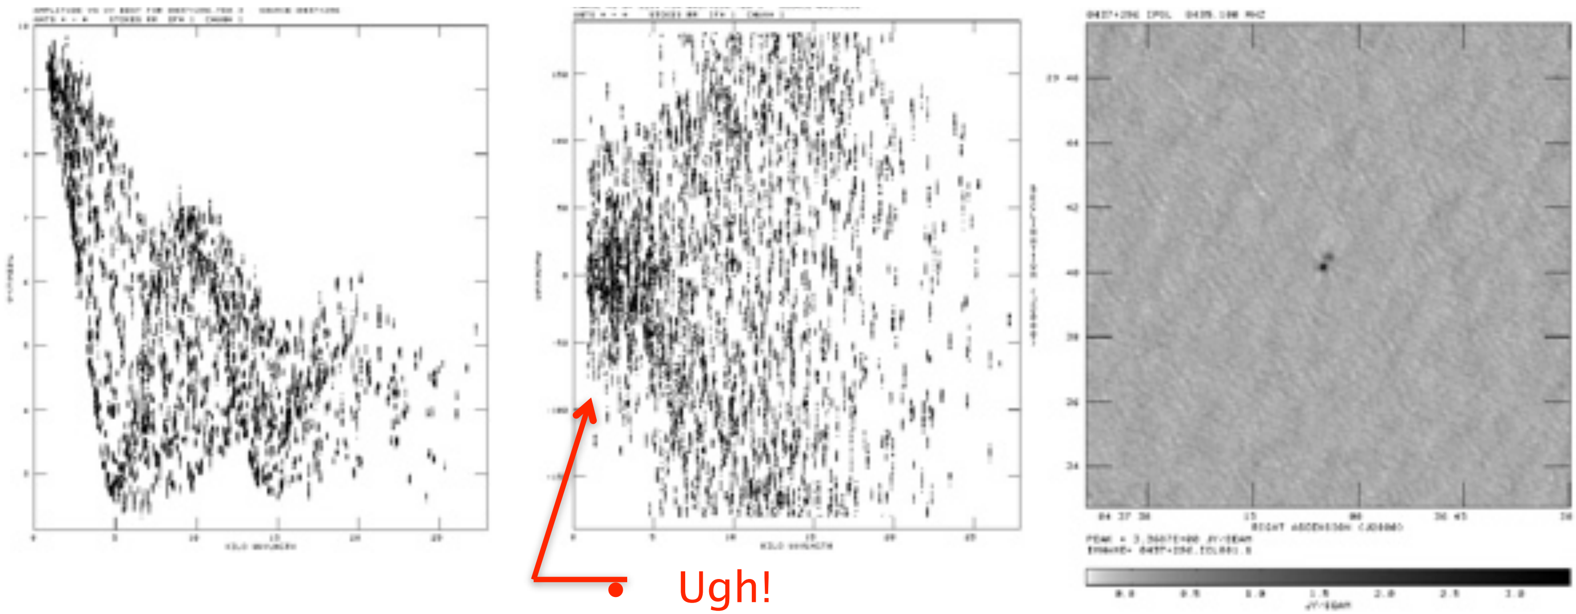

- Note that the amplitudes look quite good, but the phases do not.
- We don’t expect a great image.
- Image peak: 3.37 Jy/beam; Image rms = 63 mJy.
- DR = 59 – that’s not good!

- Using our good reference image, we do an 'amplitude and phase' self-cal.
- The resulting distributions and image are shown below.

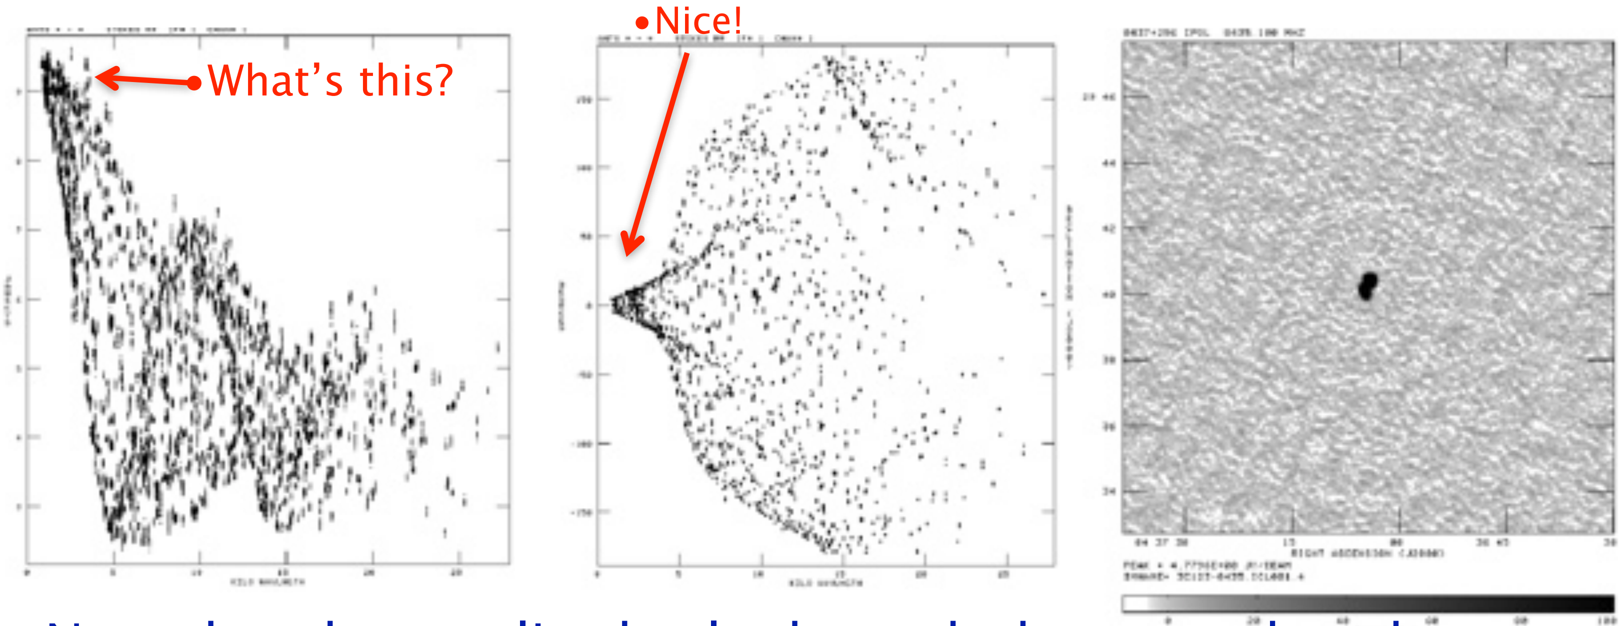

- Note that the amplitudes look much the same, but the phase are much better organized..
- Image peak: 4.77 Jy/beam; Image rms = 3.3 mJy.
- DR = 1450 – better, but far from what it should be…

- When self-calibration no longer improves the image, we must look for more exotic errors.
- The next level are ‘closure’, or baseline-based errors.
- The usual step is to subtract the (FT) of your model from the data.
- In AIPS, the program used is ‘UVSUB’.
- Plot the residuals, and decide what to do …

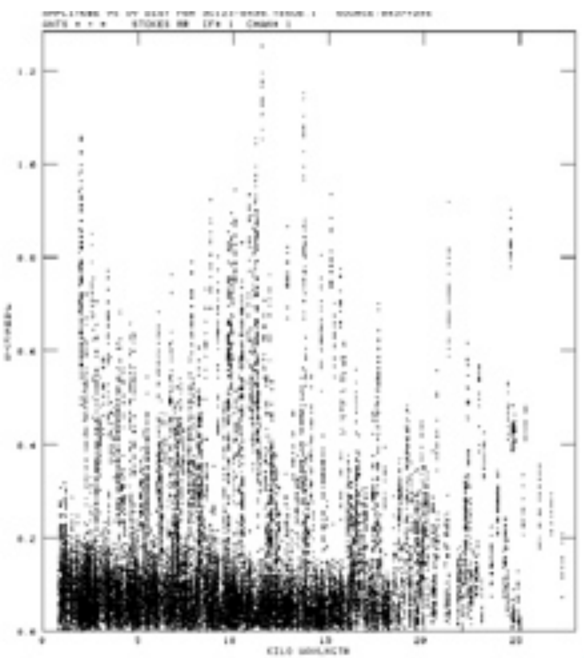

- If the model matches the data, the residuals should be in the noise – a known value.
- For these data, we expect ~50 mJy.
- Most are close to this, but many are not.

- These are far too large
- These are about right.

Removing or Correcting Baseline-based Errors

- Once it is determined there are baseline-based errors, the next questions is: What to do about them?
- Solution A: Flag all discrepant visibilities;
- Solution B: Repair them.
- **Solution A:**
  - For our example, I clipped ('CLIP') all residual visibilities above 200 mJy, then restored the model visibilities.
  - Be aware that by using such a crude tool, you will usually be losing some good visibilities, and you will let through some bad ones …
- **Solution B:**
  - Use the model to determine individual baseline corrections.
  - In AIPS, the program is 'BLCAL'. This produces a set of baseline gains that are applied to the data.
  - This is a powerful – but *dangerous* tool …
  - Since 'closure' errors are usually time invariant, use that condition.

- On Left – Image after clipping high residual visibilities. 20.9 kVis used.
- On Right – Image after correcting for baseline-based errors.

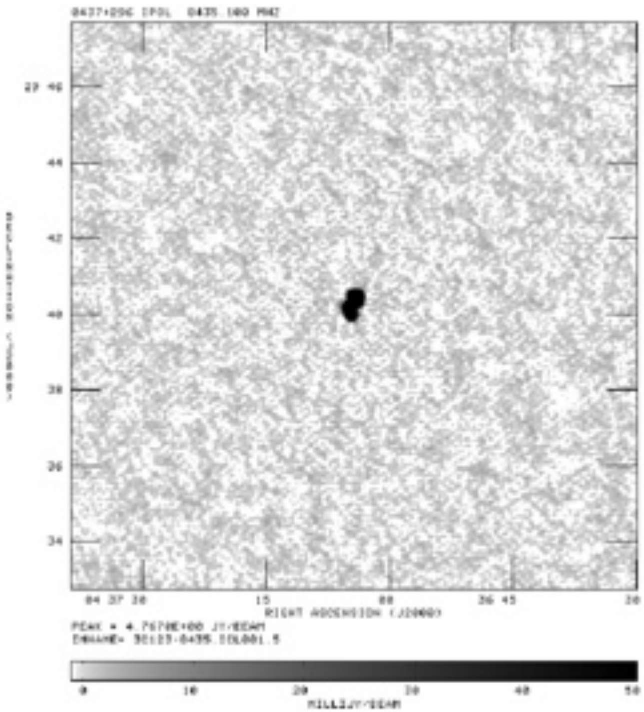

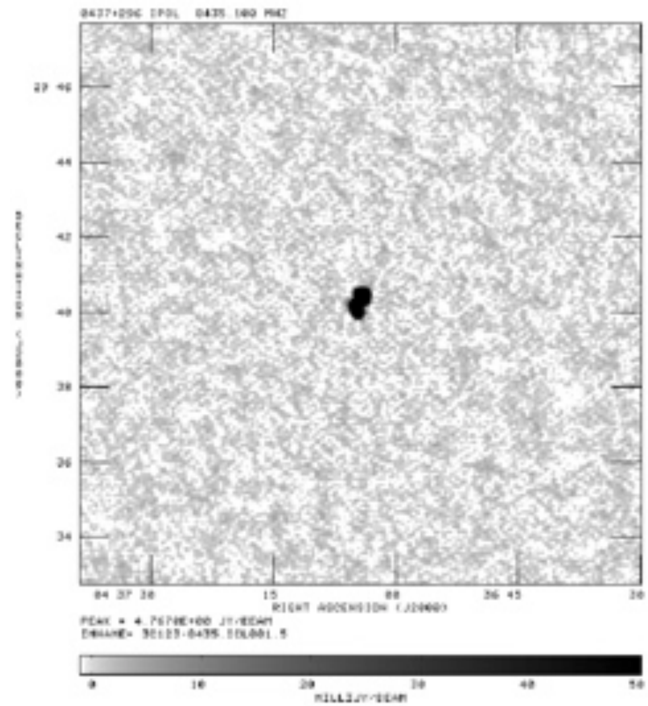

- Peak = 4.77 Jy s = 1.2 mJy Peak = 4.76 Jy s = 0.83 mJy
- DR = 3980 DR = 5740

Law of Diminishing Returns
or
Knowing When to Quit

- I did not proceed further for this example.
- One can (and many do) continue the process of:
  - Self-calibration (removing antenna-based gains)
  - Imaging/Deconvolution (making your latest model)
  - Visibility subtraction
  - Clipping residuals, or a better baseline calibration.
  - Imaging/Deconvolution
- The process always asymptotes, and you have to give it up, or find a better methodology.
- Note that not all sources of error can be removed by this process.

NRAO

Limitations to Self-Calibration

- Self-calibration is indeed an amazing algorithm, largely responsible for many (if not most) of the amazing image (and science) produced by the VLA (and other arrays).
- But it can't solve every problem:
    1. It cannot accurately locate the source. The accuracy of the position is only as good as that of the initial model – normally pretty poor, since it was made with grotty data.
    2. It cannot find the flux scale. Since the total flux contained in the model is nearly always less (and usually MUCH less) that the actual total, self-cal has the tendency to lower the flux density of the final image. (There are ways to minimize this, fortunately).
- Although self-calibration is well understood theoretically, applying it is an art, which is best learned through trial-and-error.
- Self-calibration is intimately tied with deconvolution – which means that errors in the deconvolution process will influence the outcome…

In-Beam Calibration – the best way?

- Calibration works only if the phase towards the calibrator is relevant for the target source.
- Atmospheric turbulence on many scales – cumulus clouds (for example) are a few degrees in angle – cannot expect accurate calibration to work if the source – calibrator separation is this large.
- VLA calibrator list ~ 1500 objects. Mean separation ~ 4.5 degrees – we need more sources.
- Best solution: 'In-Beam' calibration – using background sources within the antenna beam to calibrate the phases.
- The EVLA's extraordinary sensitivity (1 μJy, typical) will allow 'in beam' calibration at some bands: L (1 – 2 GHz) and S (2 – 4 GHz) for sure. Maybe at C (4 – 8 GHz).
- Not possible at higher frequencies.
- Must account for variation of phase and amplitude within each beam.

Fast-Switching and H_2O Radiometers

- Interferometer phases vary fairly rapidly with time.
- If in-beam calibration not possible, can try oscillating rapidly between source and calibrator to 'freeze' out the variations.
- Timescale of < 1 minute generally needed.
- Only partially successful.
- Most phase variations due to water vapor. If we can measure the column density of water vapor, can link this to phase.
- Some success in utilizing water vapor radiometers for this:
    - ALMA will use such devices at ~170 GHz
    - EVLA may (no funds at this time) using 22 GHz transition.
    - Need one for every antenna!

Low Frequency Calibration, and the isoplantic patch

- At low frequencies (say, below 1 GHz), we are *guaranteed* to have enough sources in the beam for in-beam calibration.
- But a new problem arises: The phase screen (primarily the ionosphere) is changing the phase on an angular scale smaller than the beam width.
- The angular scale of which the interferometer phase changes by ~1 rad is called the 'isoplanatic angle', θ_{iso}.
- If $\theta_{iso} << \lambda/D$, then a single phase calibration solution is not sufficient.
- Typically, $\theta_{iso} \sim 1$ degree – and strong dependent on frequency and baseline length.
- New low-frequency telescopes (especially LOFAR) need to develop algorithms to simultaneously solve for multiple gains, and interpolate these solutions to the target sources over the field of view.
- Much effort for this underway.

2π Steradian Calibration?

- The new low-frequency arrays will see ~1 steradian at a time.
- There will be thousands of isoplanatic patches within each…
- Instead of solving for each antenna, derive the full phase screen!

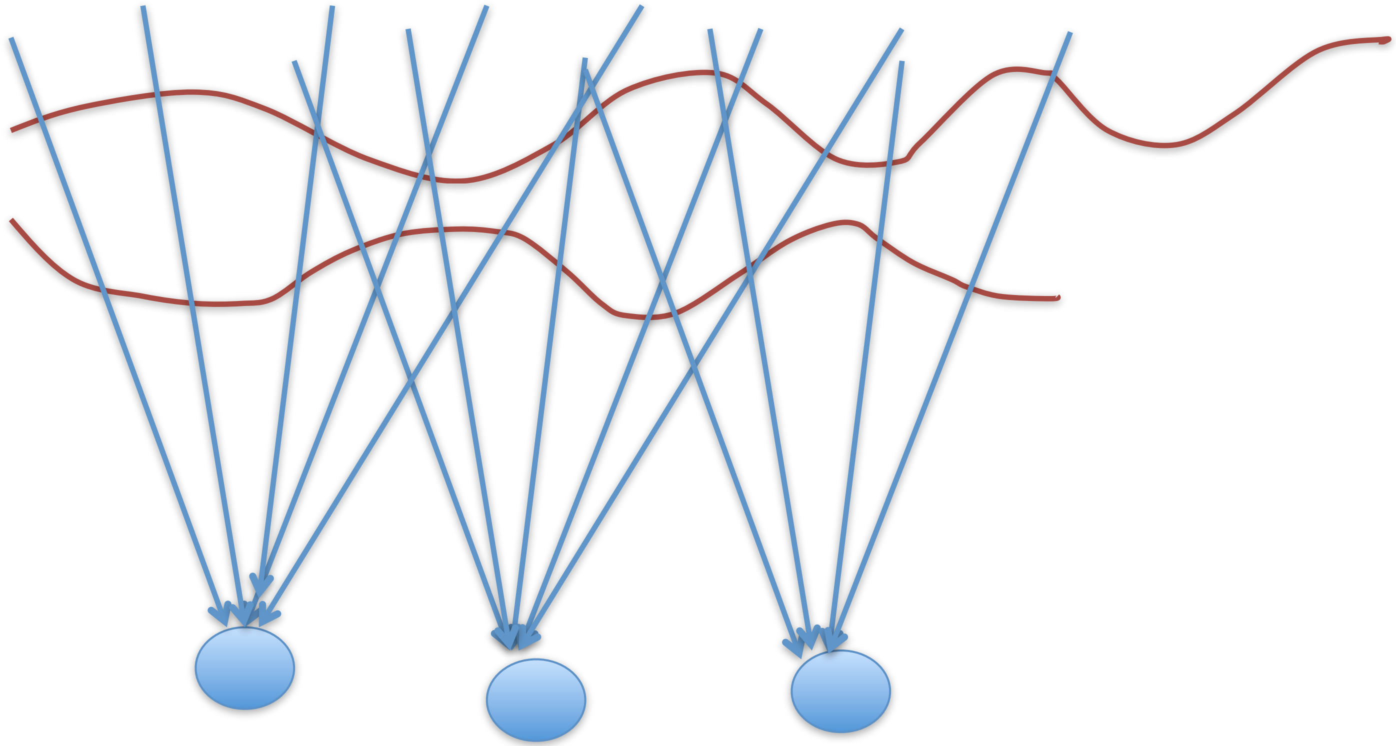

Summary

- Calibration is straightforward, in principle.
- In-beam calibration will work with modern arrays at lower frequencies.
- Fast-switching, and water vapor radiometers help – but not perfect.
- Self-calibration extremely effective – but is dependent on a good model.
- Low frequencies – and the ionosphere – present special, wide-field problems.
- Much progress on wide-angle ‘3rd generation’ calibration.
- In principle, we can describe, and utilize, the moving phase screen!

- ….

- I’m out of time … ☺