EVLA Memo #61
Removing RFI Through Astronomical Image Processing

Rick Perley and Tim Cornwell

July 21, 2003

Abstract

The development of broad-band radio astronomy systems will inevitably increase the vulnerability of radio astronomy to non-cosmic signals. Despite considerable efforts to design systems which can block strong RFI, it will not be possible to block all such signals, so that ultimately we must deal with the consequences in the imaging process. In this memo, we point out the similarity between the unwanted effects of RFI and that of a distant cosmic source. In the latter case, a well-developed methodology exists to characterize and remove such emissions. We show that the same methodology will work for RFI, provided it is not so strong as to saturate the system response, and has been sampled quickly enough to prevent phase smearing. We estimate the effectiveness of 'imaging' and removing RFI, and compare this approach to others. This approach to removing RFI has great potential, but will also be very expensive in temporary storage and processing requirements.

1 Introduction

The new generation of radio interferometers,– the EVLA, the ATA, LOFAR and the SKA, – are all being designed to utilize very wide frequency bandwidths. This is driven by the need for both maximum sensitivity and for access to any frequency where useful astronomy information may be found, as for example in red-shifted molecular or atomic transitions, or in obtaining more accurate information on spectral shapes or polarization characteristics.

But while utilizing maximum bandwidths is clearly beneficial from an astronomical viewpoint, there are very strong and negative consequences in operating instruments which will be open to all frequencies – the full radio spectrum is a very hostile environment for astronomy due to man-made signals. These signals are extremely strong in comparison to the radio astronomical signals we are seeking. For example, the spectral power flux density (SPFD) of the GPS satellite network (perhaps the weakest of all space-borne systems) corresponds to about 10 MJy, about 10 orders of magnitude above the noise levels we seek to reach with the EVLA in a 1 km/sec velocity width. Considerably worse are the DME ('Distance Measuring Equipment') signals used in air navigation, and which occupy the 1030 to 1150 MHz band. Within a 100 kHz filter width, the peak SPFD is 10^{12} Jy! Such signals will contribute a power – through the isotropic sidelobes of a 25-meter antenna – 100 to 1000 times the entire noise power of a cooled 1 - 2 GHz receiver. Even when factoring in the 50 dB forward gain of a 25-meter antenna at L-band, it is clear that these powerful signals will make astronomical imaging very difficult.

Ground-based imaging radio interferometers have two natural lines of defense against man-made signals – the antennas' forward gain, and the effects of differential fringe-winding. For a 25-meter antenna operating in the 1–2 GHz band, external signals entering through the isotropic sidelobes will be attenuated by approximately 50 dB in comparison to emission from the target source. Differential fringe winding can add from 15 to 40 dB more attenuation, depending on the baseline length, target source declination, and duration of the interfering signal (Perley, EVLA Memo #49). Typically, we might expect 70 dB or more of 'natural' attenuation of unwanted signals.[1]

This much attenuation is very good, as it reduces the typical MJy artificial sources down to a level more-or-less comparable to the stronger astronomical sources. However it is clear from the known signal strengths that attenuations of over 100 dB will be needed to permit astronomical observations at or near the frequencies of the stronger emissions. How can this be achieved? If the 'natural' isolation is only ∼70 dB, can we gain a further ∼30 dB? Are there means of subtracting such signals from the interferometer data which leave the astronomical information intact? This is an important problem for radio astronomy, which has attracted the attention of many people. Methods proposed include signal cancellation at the antenna, special weighting coefficients for large-N phased arrays which create spatial/frequency 'nulls' in the direction of interfering sources, or post-correlation removal through employment of a special RFI-tracking antenna which is employed to chracterise the amplitude and phase of the unwanted emission.

In this memo, we develop another very general approach which treats the external emission in the same manner as an unwanted astronomical source. The well-established methods of self-calibration can then allow calibration of the antenna gain and phase, and subsequent subtraction of the unwanted 'source' from the correlated astronomical data.

2 RFI as an Astronomical Source

Consider a source of external interference, mutually visible to two antennas of an array. A voltage due to this source arrives at the correlator with an amplitude and phase dependent upon the transmitter strength, transmission path, gain and phase of the antenna sidelobe through which the signal enters, and the characteristics of the telescope signal amplification and transmission system.

Upon arrival at the correlator, this voltage is multiplied with those from the other antennas to produce correlation products, each with an amplitude and phase which are uniquely identified with the originating pair, and which can be factored into a product of antenna-based voltages and phases:

$$V_{ij} = g_i g_j e^{\phi_i - \phi_j} \tag{1}$$

where the gains g_i describe the amplitude of the signals (which can further be factored in a product of the RFI power, distance, antenna sidelobe and antenna electronic gain), and the phases ϕ_i describe the path lengths, sidelobe phase, and phase offsets due to the signal transmission system.

As the RFI source is spatially unresolved, there can be no loss of coherence in the multiplication of the signals, so that the $N(N-1)/2$ baseline products can be factored into products of N complex antenna gains, as indicated in the above equation. This is true even if the signal arrives at an antenna via multiple paths – the resulting amplitude and phase at the correlator is the appropriate vector sum over all signal paths. And it is true regardless of the polarization state of the radiation or the antenna.

This situation is the same as the astronomical data set resulting from observations of an unresolved object, where the correlation products can be factored into products of antenna-based complex gains. For such an observation, well-established methodologies allow an efficient solution for the antenna-based amplitude and gain, and subsequent imaging – or removal – of the emission from the dataset. It is posited here that the same procedures will work as well with a source of nearby unwanted emission, provided certain conditions are met.

Sources of man-made emission can be classified into 'stationary' and 'moving' classes. We present an elementary analysis of the practicalities of post-correlation removal of 'stationary' interferors. There

[1] We ignore the situation when an emission source passes through the antenna beam – this situation is clearly hopeless, and the only recourse will be to delete the affected data. Fortunately, durations of such situations are very short and fairly rare.

are, however, no essential differences between this class and that of the moving sources, and a minor extension to the method will make it applicable to all moving sources.

3 Imaging Stationary Interference

To an imaging interferometer, a stationary source of emission – regardless of where it is physically located – has the same phase rate – zero – as a source of astronomical emission located at the north pole. The difference between a 'real' source at $\delta = 90$ and a stationary source of RFI lies in the phase offsets, which are due to propagation path differences, and the amplitude and phase of the sidelobes through which the signal arrives. However, these differences are immaterial to the antenna-based calibration system used in interferometry, as RFI, like an astronomical source 'closes' as long as its time variability is appropriate sampled[2]. Thus, we can consider calibrating the RFI signal as if it were an astronomical source by instructing the software to find the antenna-based phase and amplitudes corresponding to an object of given strength located at the north pole. The resulting amplitudes and phases will, when applied to the data, permit a 'clean' image of the RFI to be made, located at the north pole. The visibility data corresponding to the source can be removed through UV subtraction, after which the calibration coefficients used can be removed, returning the data to its natural state, minus the RFI.

In fact, the phase and amplitude of a stationary interfering source, as measured at the antenna, will vary in response to variations in the propagation path and the gain of the antenna sidelobe through which the emission arrives. The results of the 'RFI-calibration' procedure will provide the amplitude and phase variations due to these effects. For the proposed procedure to be effective, the solution time-scale must be short enough to resolve these variations.

In a practical phase-tracking interferometer, the phase of the signal from each antenna is adjusted to reduce the phase rate of the correlator output to zero for emission originating from the moving astronomical object. This is done to permit integration of the correlator output, and results in a significant reduction in data volume. Signals from stationary sources will then have imposed on them a phase rotation corresponding to that of the 'moving' object. The natural (sidereal) fringe rate for any earth-based interferometer is

$$f_n = \omega_e u \cos\delta \qquad \text{Hz} \tag{2}$$

where ω_e is the angular frequency of the earth (7.27×10^{-5} rad/sec), u is the projected E-W component of the interferometer baseline, measured in wavelengths, and δ is the declination of the target source. For the VLA, this frequency can range from less than 1 Hz to a few kHz. This is generally a much shorter timescale than that imposed by variations in the propagation path or sidelobe gain. It is the difference in the phase rate timescale which allows the proposed method to separate 'real' from unwanted signals.

4 Closure Errors for Stationary Interferors

For the proposed procedure to work, the correlator products must 'close', meaning that the $N(N-1)/2$ correlations can be perfectly factored into N gain factors, associated with the constituent antennas, according to Eq. 1. The limitations in the suggested methodology for removing external interference will be set by the degree of departure of the real system from this assumption.

There are many ways for 'closure' to be destroyed. Two of the most important are non-linearities in the correlator, and the effects of finite time averaging. As no correlator is perfect, there will always be a point at which an input of sufficient power will result in a non-linear and non-closing result. The

[2]Note that this argument is unaffected by whether the source is in the near field.

level at which this occurs depends on the correlator and electronics characteristics, and will not be discussed here. The second effect, that imposed by time averaging, will affect all correlators, and will limit the applicability of the proposed method of RFI removal. A straightforwards analysis of this limitation is now developed.

Correlation is done at the data sampling rates – typically hundreds of MHz to GHz. There then follows an averaging of the result, as the relevant timescales are now those set by differential phase rotation, (i.e. the phase rate differential between the target source and that of nearby objects of astronomical interest) or variations in phase path set by the electronics or signal propagation – a typical averaging time is about 1 second. This averaging will result in a non-closing error.

Consider a complex correlation product with amplitude A and phase $\phi(t) = ft$, where f is the fringe frequency. The complex representation is

$$V = Ae^{i2\pi ft}. \tag{3}$$

The correlator integrates this over some length of time, Δt. The result of this operation is

$$V_{avg} = \frac{A}{\Delta t} \int\limits_{-\Delta t/2}^{\Delta t/2} = A\,\mathrm{sinc}\,(f\Delta t) \tag{4}$$

where

$$\mathrm{sinc}\,x = \frac{\sin \pi x}{\pi x}. \tag{5}$$

This result is to be compared to the 'true' signal at $t = 0$: $V_t = A$. The error resulting by the averaging

$$\delta V_{avg} = A(1 - \mathrm{sinc}\,(f\Delta t)) \tag{6}$$

is solely to reduce the measured amplitude – the phase is unaffected. (The phase would also be incorrect if a 2nd order time dependency were involved – as there will be for sufficiently long averaging times. We ignore this fine point). This amplitude reduction is *not* factorable into a product of two antenna gains, and thus will result in a 'closure' error that will limit any antenna-based calibration effort to remove the external interference.

It is evident that in order to accurately calibrate and remove strong interference, the integration time must be shorter than the shortest timescale for phase or amplitude variations on that baseline. In nearly all cases, this will be set by the 'natural' fringe rate $f_n = \omega_e B \cos\delta/\lambda$. In this case, we can expand Eqn. 6, and express the amplitude loss as

$$\delta V_{avg} = A[1 - \mathrm{sinc}\,(f_n \Delta t)] \sim A\frac{(\pi f_n \Delta t)^2}{6} = \frac{A}{6}\left(\frac{\pi\omega_e \Delta t B \cos\delta}{\lambda}\right)^2 \tag{7}$$

The tolerable size for this loss will be the level of the thermal noise σ on that particular baseline, within the averaging time and frequency width of the data: $\delta V_{avg} < \sigma$. This condition leads to the requirement

$$\begin{aligned} \Delta t \quad &< \quad \frac{\lambda}{B\cos\delta}\frac{1}{\pi\omega_e}\sqrt{\frac{6\sigma}{A}} && (8) \\ &< \quad \frac{1.1\times 10^4}{\sqrt{SNR}}\frac{\lambda}{B\cos\delta} \qquad \text{seconds} && (9) \end{aligned}$$

where $SNR = A/\sigma$ is the ratio of the strength of the interference in a baseline in units of the noise level, within the time and frequency width of the data.

Table 1: The integration time, in msec, required to permit removal of a strong signal from a stationary source from EVLA data, when the signal strength is 100 times the noise. The time averaging is inversely proportional to $\sqrt{SNR}$.

EVLA Integration Times, in msec					
	Wavelength Band				
Config.	90cm	20cm	6cm	2cm	7mm
E	3861	858	257	86	30
D	965	214	64	21	7.5
C	305	68	20	6.8	2.4
B	96	21	6.5	2.2	.75
A	30	6.8	2.0	.68	.24
NMA	3.0	.68	.20	.068	.024

Table 1 shows the required averaging time for $SNR = 100$, for various EVLA configurations and wavelengths. The table is based on the longest baseline contained within the configuration listed, and thus represents a worst case. In practice, an averaging time more representative of the median baselines will likely be sufficient, in which case the required integration times might be doubled or tripled.

A simple practical example will illustrate the potential. Presume a short-spacing, low frequency experiment (say, 20 cm in D configuration) in which RFI is expected to be found within the band. The spectral resolution employed is 1 kHz (0.3 km/sec), and RFI at a level 30 dB above the noise is anticipated. For this level, we see from Table 1 that an integration time of $\sim$ 100 msec will be needed. Within this time and bandwidth, the rms noise will be 20 Jy for the EVLA, from which we can estimate that a signal with SPFD of 20000 Jy can be removed. When referenced to the forward gain of the antenna, the actual strength is 2×10^9 Jy – well above the level of GPS satelites, and probably representative of satellite-borne emissions in general, but still well short of nearby DME signals.

5 The Costs of Signal Subtraction

Although this method of signal excision has the potential to effectively remove corrupting data without SNR degradation, it will not come without considerable cost. For the 27-antenna VLA, and assuming the 16384 channels produced by the WIDAR correlator, the resulting data production rate for 1 msec output averaging is typically 100 GB/sec! Even for 'easy' cases such as D-configuration at 20cm, the required data output rate will be a relatively prodigious 500 MB/sec.

These are high data rates, requiring large storage capacity. But the problem is compounded by extensive processing requirements. Fortunately, the operations themselves are relatively straightforward. For each affected frequency channel, and for each affected integration time, we must:

1. Phase rotate the data so as to undo the phase rotation imposed by the fringe-tracking. This step will reduce the phase rate for the RFI to zero, and impose the 'natural' rate on the astronomical source. For a stationary source of RFI, the appropriate phase is the negative of that of the target source. For a moving source, the appropriate phase to be added must correspond to the difference between the moving source and the phase of the target source. This is a straightforwards complex multiply.

2. Solve, using a least-squares algorithm, for the RFI amplitude and phase for each antenna, for each polarization product. The model is a point source of unit amplitude located at the field center (i.e., the north pole). The solution interval must be sufficiently short to resolve any residual variations in the RFI phase and amplitude with an accuracy as good as that used to impose the

sample time, but long enough to permit the natural fringe rate to 'erase' the visibilities of the target source.

3. Apply the solution to the data through a complex multiplication of each data point by the derived antenna solution gains.
4. Subtract the interferor from each polarization. As the data are now calibrated and phase-centered, this is simply a unit subtraction from the Real component.
5. Remove the amplitude and phase calibration from the data by division by the same gains used to 'calibrate' the data.
6. Re-rotate the data back to the original phase reference frame by complex multiplication by the appropriate phase rotation.
7. Average the data down to a manageable size, appropriate for the target source and field of view.

Ideally, these operations would be done in close to real-time, so that the data need not be archived at the fast sampling rates required to permit the RFI subtraction.

A calculation of the required processing for practical cases remains to be done.

6 Suggested Applications to the EVLA

The EVLA's WIDAR correlator has the capability to output data at a rate which should allow much of this type of processing to be done. The current plan will permit all 16384 channels to be output at a maximum rate of 100 msec. Even this relatively slow rate will permit accurate removal of RFI for 90 and 20 cm data in D-configuration, and for 90 cm data in C-configuration. These are the configuration/frequency combinations with the most serious RFI problems.

The maximum data output archiving rate with the installed baseline boards will be about 11 msec, permitting removal (at SNR = 100) for all wavelength/baselines in the upper left half of Table 1. Thus, 90cm emissions out to A-configuration, and 23 GHz emission in the D-configuration, should be removable at the 100:1 level, with the planned correlator.

The theoretical maximum output rate of 0.7 msec will permit operations on most EVLA array/frequency combinations. Those remaining (lower right quadrant) will likely see sufficient natural fringe rotation and RFI suppression, so we may expect there not to be a need for post-correlation subtraction. Clearly, the high-frequency, long-baseline combinations will benefit considerably (typically 40 dB or greater) from natural suppression. The greatest need for post-correlation subtraction is for the short-spacing configuations at the lower frequencies, where the natural supression is lower and the RFI spectrum more severe. Fortunately, these are precisely the cases where the required data averaging is long enough to permit post-correlation subtraction techniques to be developed and applied.

7 A More Advanced Algorithm

In the proposed method, the characteristics of the desired source have been completely ignored. We have relied on the natural fringe rotation to effective 'rub out' the visibilities of the target field. However, this process is imperfect – and in any event is of little help when the target source is close to the north pole – and will result in a residual 'closure error' which will degrade the accuracy of the proposed subtraction process.

A better method must account for the visibility of the target source as well as that of the unknown interfering source. Were the brightness distribution of the target source and the forward gains of the antennas known, the corresponding visibilities could be computed and removed from the database prior to solution for the interfering signal's characteristics. This would allow a more accurate determination of the RFI characteristics.

In general, neither the brightness distribution of the target source nor the antennas' forward gains will be known, so a methodology for simultaneously solving for the antenna complex gains in two (or more) directions must be developed. It is worth noting that that this same need arises in calibration of low-frequency interferometric data, where the isoplanatic angle due to the ionosphere is much smaller than the antenna primary beam. In this situation, the antenna-based phase corrections required for coherent imaging vary as a function of both time and angle. These algorithms are being developed by various groups, and there appears to be no problem in principle with this approach.

8 Discussion, and Comparison to Other Proposed Methodologies

There are a number of advantages to this proposed method:

- No extra antennas, hardware, or expanded correlator are needed. The method uses the same data taken for astronomical applications.
- For stationary sources, the direction of the RFI does not need to be known. However, for moving sources, an approximate direction and angular velocity will be required.
- The basic proposed methodology uses well-tested software which has successfuly been employed to remove 'natural' interfering sources. However, it is likely this 'simple' method will not be sufficient, so the joint-solution approach needs to be developed.
- It is not necessary to know the polarization state of the signal nor of the antennas. All that is relevent is the observed correlated signals.

On the other hand, processing requirements will be large, so that the increased post-processing computing will be required. The method will not work if the signals are sufficiently strong to cause non-linearities anywhere in the signal chain.

The proposed means of RFI amelioration is one of many. Other suggested methodologies include:

- Cancellation of interfering signals at the antenna. This method, which is required for single-dish applications, uses an auxilliary antenna which is directed at an interfering source. The phase of this signal is inverted, and the signal added to the radio astronomy signal, after suitable adjustment of gain. The method requires an individual antenna for each interferor, and thus will be expensive for a large-N array with many sources of interference. Because the method adds information from an auxilliary antenna, there will be a degradation in the SNR of the astronomical signal.
- Post-correlation excision, using an auxiliary antenna of known gain. Here, an auxiliary antenna points to the source of interference, and this signal is correlated against all the other astronomical antennas' signals. The characteristics of the signal are extracted, using the auxilliary antenna as a reference, and interering data removed from the astronomy data by vector subtraction. This method is very similar to that proposed here, and differs mainly in that a gain solution solution is not required. Processing requirements should be lower, but instrumental costs higher, as an auxiliary antenna is required for each interferor, and the correlator must be able to accommodate the extra correlations. The direction to each interferor must be known.
- Phased-array nulling. If a large-N array is used in a phased-array mode (*i.e.* the antenna phases are adjusted to form a real-time beam, or beams, on the sky), small gain perturbations can

be introduced to place array nulls in known sky directions at particular frequencies, at a small cost in forward senstivity. Roughly, each null in both space and frequency costs 1/N in forward gain. This loss becomes significant if the number of nulled interferors becomes comparable to the number of antennas – guaranteed to be the case for an array of only 27 antennas like the VLA! It is also hard to see how this method can be applied for RFI coming in at a low angle, as the phase and amplitude of these data are modified both by the propagation path, and by the antenna sidelobes through which the signals arrive.

- Signal processing of baseline-based data. In this method (see Roshi and Perley, 2003), a signal whose frequency is significantly different than that of the desired astronomical signal is blocked by digital filtering. For stationary interference, for example, the undesired signal has a well-known frequency which is in general different than that of the desired signal. To be effective, the duration of the data stream must be much longer than the inverse of the frequency difference between the desired and undesired fringe rates. During this (potentially long) period, the phase characteristics of the unwanted signal must be very stable, or the filtering process will be ineffective. This method has the advantage of dealing separately with each data stream, and could probably be employed in real-time. The required processing is likely to be high.

9 Feasibility Studies

At the time of writing, this proposed method is simply a concept, and no demonstrations have been made. It is suggested that two be done:

1. The situation is easy to simulate. The response to two sources – one at the north pole and one at a desired declination– can be simulated with an integration time short enough to ensure minimal phase averaging. Reasonable phase and amplitude perturbations can be trivially added to the data. The available software can easily solve for the apparent properties of the stationary source, and the resulting strength of the target source be derived, for various values of the interferor and source strengths, compared to system noise. This proposed trial should tell us very quickly of the applicability of the proposed method, and of how cleanly it can separate interference for a real astronomical source located near the north pole.

2. There is no substitute for real data, so the ultimate test will be to remove some interference from real VLA data. The efficicy of this will be limited by the limited capabilities of the VLA correlator, as this cannot simultaneously give high time and frequency resolution. However, as the critical test of the method involves low frequency, short-spacing situations, the present correlator can be employed to produce databases appropriate for test.

An effort is now underway to do both tests. The results of these will be published in forthcoming memos.