

Spatial Nulling for Attenuation of Interfering Signals

A. R. Thompson*

National Radio Astronomy Observatory, USA

7/8/03

Abstract

The attenuation of unwanted signals by forming nulls (i.e. minima in the gain) in the spatial response pattern of a radio astronomy array is described in a number of recent papers. This memorandum reviews the basic principles and attempts to examine factors that limit the performance of such techniques. In adaptive nulling, the removal of interference is limited by the ability to identify interfering components of the received power within short averaging times. After full integration over a long observing period, the level of the interference is comparable to, or perhaps a few times greater than, that of the noise. In deterministic nulling, the effectiveness of the suppression does not depend on the strength of the interference, but the responses of the antennas in the direction of the interferer must be closely matched from one antenna to another, or else must be calibrated. Since, in most cases, interference is received through the far sidelobes of the antennas, deterministic nulling introduces requirements on the sidelobe responses that are not usually encountered. The parameters of some satellite systems that cause interference to radio astronomy are considered, and numerical analysis is used to estimate the widths of nulls in some simple arrays.

1 Introduction

A number of papers have been published which examine the possibility of eliminating or mitigating interfering signals in radio astronomy arrays by generating nulls in the spatial response pattern in the directions of incidence of such signals. This technique will be important in the operation of the SKA. There are two basic ways in which a null can be applied. In *adaptive* nulling the interference components within the received signals are identified, and the signals from the antennas (or the cross products from pairs of antennas) are combined in a way that causes the interference vectors to cancel one another. The direction of the interfering sources on the sky need not be known. In *deterministic* nulling the direction of the interferer is known and a null is formed in that direction. It is not necessary to be able to identify the interfering components within the received signal, but if the angular responses in the direction of the null differ from one antenna to another, it is necessary to calibrate them in both phase and amplitude. There are two ways in which these two nulling procedures can be applied to a synthesis array. First, the nulls can be formed by adjusting the weights with which the cross-products of the outputs of pairs of antennas (i.e. the visibility values) are combined. In this case the nulls are formed in the synthesized beam, with which the resulting image is convolved. Second, in the case of synthesis arrays in which the elements between which cross-products are formed are themselves phased subarrays of antennas, the nulls can be formed individually for each subarray. In this case the nulls are formed in the subarray beam with which the image is multiplied.

Adaptive nulling in the synthesized beam is described by Leshem and van der Veen (1999, 2000), Leshem et al. (2000), and Fridman and Baan, (2001). Deterministic nulling in the subarray beam is discussed by

*athomps@nrao.edu

Smolders and Hampson (2002), Ellingson and Hampson (2002), and Ellingson and Cazemier (2003). The application of deterministic nulling to the Allen Telescope Array (ATA), a large phased array, is discussed by Welch (2000a,b), Bower (2001), and Harp (2002).

2 Elementary considerations

Consider an array of n nominally-identical antennas, each of which is connected through a phase shifter to an n -to-1 power combiner. Each picks up a power level p from an RF signal. In the power combiner, the power is divided n ways between the other antennas and the output. Thus each antenna contributes a power level p/n to the combiner output. The voltage contributions of the antennas can be represented by vectors of amplitude $\sqrt{p/n}$. If the phase shifters are adjusted so that the contributions combine in phase, the vectors are aligned and the output voltage is \sqrt{np} . The output power is np , as expected, since the total collecting area is n times that of a single antenna. Now suppose that the phase shifters are set so that the signal vectors combine with random phase angles. The combined voltage has an expectation of \sqrt{n} times that of a single antenna (this is equivalent to the classic random walk problem in which distance achieved is proportional to the square root of the number of steps taken). Thus the expectation of the combined power received is equal to that from a single antenna, p . Finally, consider the case where the phases are adjusted so that the vectors form a closed loop with zero resultant, thus producing a null in the direction of incidence of the signal. If each signal vector has a random error in amplitude and phase of relative rms amplitude ϵ , then the vector sum will fail to close by a amount equal to the sum of the errors, that is, $\epsilon\sqrt{p}$, resulting in a power level of $\epsilon^2 p$. Thus a null of depth x dB below the response of a single antenna requires $\epsilon = 10^{-x/20}$, e.g., $\epsilon = 0.03$ for a null depth of 30 dB. These requirements on the accuracy of the voltage responses apply to the interference components identified in adaptive nulling, and to the accuracy of the antenna responses in deterministic nulling. In closing the vector loop for a null, the shape of the loop is not constrained, so free parameters remain for forming beams or nulls in other directions.

In deterministic nulling, suppose that one starts by determining the phase adjustment required to close the vector loop on the assumption that the antennas are all ideal isotropic radiators. Then, to take account of the actual gain of the antennas, each signal vector has to be multiplied by a complex gain factor. If the gain factor is the same for each antenna, the size and orientation of the vector loop will be changed, but it will remain closed. Thus the response factor in the direction of the interferer need not be known so long as it is identical for all antennas. If the gain factor differs from one antenna to another, the loop will not remain closed unless the individual gain factors are known and taken into account.

3 Adaptive nulls in the synthesized beam

Adaptive nulling in the synthesized beam is described by Leshem and van der Veen (1999, 2000), Leshem et al. (2000), and Fridman and Baan, (2001). The method described requires the correlator outputs for all antenna pairs to be presented in the form of the covariance matrix (also referred to as the coherency matrix), \mathbf{R} , which for n antennas has dimensions $n \times n$. The element r_{ij} represents the correlator output for antennas i and j , and since $r_{ij} = r_{ji}^*$, the matrix is Hermitian. The diagonal values represent the total power outputs, which must be included. Section A of the Appendix shows how the elements of \mathbf{R} can be derived from parameters of the antennas and incident signals. For a long observing period, \mathbf{R} must be determined repeatedly for averaging times short enough that the interference received does not vary significantly. As discussed in §4 and §5, in practice this time scale typically has a value in the range milliseconds to seconds. The next step is to perform an eigendecomposition of \mathbf{R} :

$$\mathbf{R} = \mathbf{U}\mathbf{A}\mathbf{U}^H, \tag{1}$$

where \mathbf{U} is an $n \times n$ unitary matrix whose rows are the eigenvectors, $\mathbf{\Lambda}$ is an $n \times n$ diagonal matrix that contains the eigenvalues, and superscript (H) indicated the Hermitian transpose (the complex conjugate of the transpose). It is assumed that over such short averaging periods the astronomical signals at the correlator output are small compared with the rms noise. Thus, eigenvectors for which the eigenvalues are greater than the rms noise level represent interference. It may also be possible to distinguish between interference and cosmic signals from the directions of incidence¹. The number of such interference sources must be less than the number of antennas. Since \mathbf{R} is Hermitian, the eigenvectors are orthogonal, and can be projected out, thus forming nulls in the directions of the corresponding signal sources. A matrix \mathbf{U}_N can be formed from \mathbf{U} by removing rows that correspond to eigenvalues significantly greater than the noise. The columns of $\mathbf{U}_N \mathbf{U}_N^H$ span the subspace occupied by the noise and the cosmic signals. A filtered form of the covariance matrix can then be formed:

$$\mathbf{R}_F = \mathbf{U}_N \mathbf{U}_N^H \mathbf{R} \mathbf{U}_N \mathbf{U}_N^H, \quad (2)$$

from which the identified interference has been projected out. The evaluations of \mathbf{R}_F for the series of short averaging times comprise the data from which an image, free from the identified interference, can be formed. Since each element of \mathbf{R}_F is composed of a weighted sum of the elements of \mathbf{R} , there is no longer a simple Fourier-transform relationship between \mathbf{R}_F and the required image. However, the weighting factors used in deriving the elements of \mathbf{R}_F are known, and possible approaches to image formation are discussed in the references given above.

4 Initial averaging times in adaptive nulling

In the process just described, the measured values in the covariance matrix are redetermined at intervals that we shall refer to as the *initial* averaging period. The data from all such periods during the particular observation are subsequently combined in the imaging process. Several effects limit the initial averaging period. In the absence of interference, the limit depends upon the angular dimensions of the field to be synthesized, measured in units of the synthesized beamwidth. A corresponding number of samples is required for the measurement of visibility across the (u, v) plane, so the averaging time must be small enough that independent values of visibility are measured at sufficiently fine intervals in u and v , as the baseline vectors follow the rotation of the Earth. An analysis of this effect is given in Thompson et al. (2001), see Eq. (6.81) [or (6.71) in the first ed.]. For a map with 10^3 beamwidths on a side the averaging time should not exceed 18 s if the visibility is not to be reduced by more than 10% for the longer (u, v) spacings. However, in practice, since sporadic interference is often a possibility, initial averaging times are usually no more than 1-10 s, so that such interference can be efficiently edited out.

If adaptive nulling is being used, a further concern is the stationarity of the interference, as discussed by Leshem et al. (2000). The initial averaging period should be short enough that during this period the interfering signal does not change significantly at the correlator output. The response to the interference will be modulated by the motion of the interfering transmitter through the fringe patterns of the array. The angular width of the narrowest fringe of an array is roughly equal to the width of the synthesized array beam. Some examples of the angular velocities of satellites are given in Table 1. Thus, for example, for an array with a 10 arcsec synthesized beam, the time interval for relative motion of a satellite through the narrowest fringes would be about 700 ms for a geostationary-orbit (GEO) satellite², 250 ms for a medium-earth orbit (MEO) satellite, and 5 ms for low earth orbit (LEO) satellite. To perform adaptive nulling we

¹Each eigenvector represents a component from a distinct spatial direction. Suppose that one derives the directions from the eigenvectors on the assumption that the antenna response in any direction is the same as that in the main beam. The wanted cosmic signals, if they are not much smaller than the noise, will have directions within the small spatial field under study. The interfering signals will be more widely scattered, and if the sidelobe patterns vary from one antenna to another the positions will be incorrect, but the probability of their falling within the study field will be small.

²The array fringe patterns track the center of the field under observation, resulting in fringe oscillations for a signal transmitted from a fixed position relative to the Earth.

need to preserve the amplitude of the interference responses so as to be able to determine the corresponding eigenvectors accurately. Thus for such an array the initial averaging time must be in the range of milliseconds to a few tenths of a second.

Interference in the received signal can only be identified and adaptively nulled if it is distinguishable in the presence of the noise after averaging for the initial period. Also, the accuracy of the nulling will be limited by the noise. Thus we can assume that interference that produces components at the correlator output that are equal to or less than the rms noise (after averaging for the initial period) will not be removed by the nulling. As the data are further averaged over the full observing time, the noise will be reduced in proportion to the square root of the total averaging time. Since the interference responses are sinusoidally modulated by the fringe patterns, they will also average down, but to an extent that depends on the angular motion of the interfering source, the configuration of the antennas, the constancy of the radiated power, and similar details.

An analysis for a case of a geostationary interferer can be based on a result by Thompson (1982): see §B of the Appendix to this memorandum. Consider, for example, the VLA in configuration C, in which the spacings are comparable to those being considered for the central core of the SKA. The longest spacing between antennas is 3.37 km. Let the operating wavelength be 20 cm. For a low declination observation, the highest fringe frequency for interference from a stationary transmitter is 1.13 Hz. If we take 1/8 of the fringe period (a range of 45° in fringe phase) as a measure of the stationarity, the initial averaging period must not exceed 0.11 sec. Then, from Eq. (11) in the Appendix, an interfering signal that produces a component in the correlator output equal to the rms noise after averaging for 0.11 sec, will result in an interference-to-noise ratio (INR) of ~ 3 after averaging over 12 h of observation³. There are many variables in a practical situation, so the precise value of this result should not be over-interpreted. However, it can be concluded that over a long observing time the interference can average down in a way that approaches that of the noise. This can be understood physically as follows. Since the initial averaging time was chosen so that the resulting frequency response is just starting to roll-off at the highest interference fringe frequency, the initial averaging will reduce the frequency range of the noise components to something very similar to the range of the fringe frequencies. Thus, in the long period averaging, it can be expected that the interference and noise will be reduced by factors of a similar order of magnitude, but since the interference cannot be assumed to have a normal distribution in amplitude, or to have as flat a spectrum as the noise, it may not be reduced as efficiently.

Table 1

System	Frequency	Height (km)	Angular vel ^a . (Deg/sec)	SPFD ($\text{Wm}^{-2}\text{Hz}^{-1}$)	Reqd. atten. ^b (dB)
TV	10.7-12.5 GHz	35,860	geostationary ^c	-166	74
GPS	1565-1585 MHz	20,000	1.1×10^{-2}	-195	56
Glonass	1597-1617 MHz	19,000	1.2×10^{-2}	-192	59
Iridium	1621.35-1626.5 MHz	780	0.55^d	-170	81

^aThe angular velocity shown is the maximum value for the satellite, as seen by an observer on the surface of the Earth. This maximum occurs when the satellite is near the zenith.

^bThe required attenuation is the amount by which the SPFD exceeds the detrimental thresholds for radio astronomy given in Table 1 of Recommendation ITU-R RA.769. These threshold levels correspond to interference received through sidelobes of gain 0 dBi and averaging times of 2000 s.

^cFor a geostationary satellite the rotation of the earth introduces an angular velocity relative to a source under observation at declination δ of $4.17 \times 10^{-3} \cos \delta$ deg/sec.

^dThe spectral power flux density (SPFD) for Iridium is an approximate estimate for high traffic conditions.

³Note that this result holds for observing declinations between about $\pm 70^\circ$. Averaging becomes less efficient in reducing interference at higher latitudes because the fringe frequencies that result from Earth rotation are proportional to the cosine of the observing declination.

Practical situations may be more complicated. For example, a small value of the initial averaging time may be necessary to accommodate a fast-moving satellite when other interference is present that would be better handled with a larger value. Also, the analysis above does not take account of the reduction in the correlation of interference that can occur as a result of the time delays from signals received in two antennas not being matched when the direction of incidence is not close to the field center under observation⁴. This decorrelation would not apply to interference that is intrinsically narrow-band or processed through a narrow spectral-line channel.

5 Effect of residual interference

The correlator system for the SKA will provide measurements of the coherency matrix for n_c individual frequency channels, say, for example, $n_c = 4096$. It is clearly most effective to apply adaptive nulling to bandwidths similar to the expected widths of the interfering signals, since this maximizes the probability of identifying them in the presence of noise. Also the depth of a spatial null is limited in frequency. If adaptive nulling is performed for each channel, the example above suggests that, for any channels that contain interference, the residual INR after full time-averaging will be reduced to some small number. For spectral line observations, the effect of adaptive nulling will probably not be sufficient to prevent some deterioration in sensitivity if the frequency of the interference is close to that of the line of interest. For continuum observations the data are averaged in frequency over the spectral channels to form a single image. Suppose that there is interference in n_i channels. The INR for the full bandwidth, after averaging over frequency, is equal to that for a single channel with interference multiplied by $\sqrt{n_i}/\sqrt{n_c}$, that is, multiplied by the square root of the fraction of channels that contain interference. For example, interference in 10% of channels would result in INR of 0.3 of that for a single channel. Decorrelation of interference due to timing mismatches can also be helpful for continuum observations.

6 Adaptive nulls in phased-subarray beam patterns

In several concepts for the SKA, the elements of the synthesis array are themselves phased subarrays of small antennas. Several advantages of producing the nulls in the subarray beams, compared with producing them in the synthesized beam, are apparent. First, distortion in the synthesized beam is greatly reduced. (The accuracy with which the synthesized beam is known limits the dynamic range that can be achieved in the deconvolution of the image.) Second, since the spacings of the antennas in a subarray are much shorter than those in the full synthesis array, the fringe frequencies corresponding to antenna pairs within the subarray are lower, so the initial averaging period can be longer. Third, the required computation in the application of nulling to a ‘large-n’ array is simplified if small groups of antennas can be treated separately.

For the various antenna pairs within the phased subarray, the cross correlations that correspond to the elements of the covariance matrix are not usually measured. However, Ellingson and Hampson (2002) present a method known as *subspace tracking* which allows development of estimates of the eigenvectors and eigenvalues of the covariance matrix \mathbf{R} without explicitly estimating \mathbf{R} . They illustrate the method using an iterative method by Yang (1995) called “projection approximation subspace tracking with deflation” (PASTd). The steering vectors (i.e. complex weighting factors required in the combination of the antenna outputs), have to be redetermined on a time scale set by the effective rate of motion of the source of interference, as discussed in §4. Ellingson and Cazemier (2003) consider the choice between several different algorithms for the application of steering vectors, to minimize distortion of the main beam responses. Use of more sophisticated algorithms can improve upon the basic method of projecting out the interference.

⁴More precisely, individual antenna pairs should be considered. For each antenna pair, time delays are matched for directions of incidence represented by a small circle on the celestial sphere that lies in a plane normal to the baseline and passes through the field center.

7 Deterministic nulls in phased-array beam patterns

In deterministic nulling, the effectiveness with which a signal can be suppressed does not depend on the strength of the signal, so the problem of identification of weak interference that occurs with adaptive nulling is avoided. The direction of incidence of the interference must be accurately known, and this requires maintaining a catalog of potentially interfering satellites. This should not be difficult with the SKA, since a small number of antennas could be devoted to checking satellite orbital data. Note that for most of the time, interference from satellites will be incident on the far sidelobes of the antennas, tens of degrees or more from the main beam axis. Deterministic nulling is discussed by Smolders and Hampson (2002) with respect to the European SKA concept (Aperture Array), and by Welch (2000a,b), Bower (2001), and Harp(2002) with respect to the Allen Telescope Array (ATA).

There are several well-known procedures for generating nulls in antenna patterns, (e.g. Welch 2000b). In deterministic nulling, the sidelobe responses (in amplitude and phase) of the individual antennas in the direction of the source of interference must either be equal (but need not be known), or else the sidelobe patterns of the individual antennas must be calibrated. The required accuracy of the equality or the calibration is as considered in §2, for example 3% in amplitude or 1.7° in phase for a null of -30 dB (with respect to the sidelobes). In the Aperture Array concept of the SKA the individual elements are dipoles or Vivaldi radiators, which have very wide beams, and Smolders and Hampson (2002) describe a method of calibrating the response of phased groups of elements within a 1 m^2 tile. They also show examples of measured patterns in which nulls of about 20 dB with respect to an isotropic radiator have been achieved.

8 Time scale for adjustment of parameters in deterministic nulling

Widths of nulls for some simple arrays are examined in §C of the Appendix. It is found that the width of a single null at a depth of, say, 30 dB is roughly 0.06 of the beamwidth of the same array. Then, for example, for an array with a 10-arcsec beam observing a low-declination source, the null would be traversed by a geostationary-orbit (GEO) satellite in 40 ms. For a medium-earth-orbit (MEO) satellite, like GPS, the null would be traversed in about 15 ms, and for a low-Earth-orbit (LEO) satellite, like Iridium, the null would be traversed in about 0.3 ms. (These figures are based on the angular velocities in Table 1.) For large zenith angles the angular velocities are reduced, the reduction being greatest for the satellites in the lowest orbits. Also, the nulls on the sky may be elongated and the satellite may not be moving normal to the elongation. Thus the crossing times above are worst-case values. To track a satellite with a deterministic null, the pointing vectors for the null must be adjusted on corresponding time scales. However, this should be possible if the satellite position data and the antenna positions and responses are determined in advance.

9 Sidelobe requirements for deterministic nulling

Deterministic nulling provides a way to make use of the predetermined direction of incidence of an interfering signal. However, the requirement that the sidelobe patterns of the individual antennas be closely matched, or else calibrated, may present a substantial problem. Consider, for example, paraboloid antennas widely used in radio astronomy, such as the 25-m VLA antennas. For such antennas the sidelobe level at angles from the main beam of about 20° or more results mainly from scattering from feed supports, surface errors, feed spillover, and similar unwanted effects which are larger than the $J_1(x)/x$ oscillations resulting from the diffraction at the circular aperture. The individual surface panels are adjusted after the main structure is assembled, and, since the surface adjustment is concerned entirely with optimizing the characteristics close to the main-beam axis, it is likely that the detailed structure of the far sidelobes differs from one nominally-identical antenna to another. If, for example, the positions of the points where the quadrupod legs intersect the main reflector differ by a centimeter or two from one antenna to another, this would have virtually no

effect on the main beam, but could have a significant effect on the precise response in any direction in the far sidelobes. Also, for example, the phase of spillover sidelobes will vary with focusing adjustments in a different way from the phase of the main beam. Thus although the sidelobes of such antennas are *statistically* identical, the precise responses as a function of direction may be quite different. It would not be surprising if deterministic nulling in uncalibrated sidelobes of such antennas failed to work. Antennas suggested in the US concept for the SKA are about 12 m diameter and hydroformed from a single sheet of aluminum, so the situation with respect to matching of the sidelobes can perhaps be expected to be better. The performance of the 6-m ATA antennas will provide information in this area.

If it were necessary to calibrate the sidelobe pattern of a 12-m antenna, a fairly large amount of data would be involved. At 20 cm wavelength, the beamwidth is about one degree. The angular scale of the structure in the sidelobes is similar to that of the main beam, so one would need measurements at least every half degree, or 8×10^4 measurements to cover one hemisphere (approx 2×10^4 sq. deg.). The results would have to be interpolated to get precise gains in intermediate directions, and calibration at smaller angular intervals would be advantageous. Pointing corrections with elevation might also be needed. For satellites just above the horizon, blockage by other antennas (in the case of a close-spaced array) will complicate the situation. Also, the sidelobe structure will vary with frequency: for an antenna of diameter n_λ wavelengths, a change in relative frequency by $1/n_\lambda$ can be expected to make a significant change in the far sidelobes. Overall, it appears that calibration of the far sidelobes is hardly practical. Thus, we are left to consider the extent to which the sidelobe patterns are matched from one antenna to another. Computer programs exist for detailed calculation of the effect of scattering of radio waves from complicated structures such as aircraft. It should be possible to use such techniques to examine how sidelobes vary as specific dimensions of the antenna are varied within the range of engineering tolerances. Thus it should be possible to identify structural dimensions that are most critical to the matching of sidelobe responses.

Some different aspects of sidelobe behavior can be expected with the Luneburg-lens type of antenna. Widely-spread sidelobes may result from low-level reflections at the surface of the lens since there is a small abrupt change in permittivity from air to dielectric material. It is important to examine whether the sidelobe level varies with the condition of the surface, for example if it is wet or becomes abraded. Also, in the design of the Luneburg-lens antennas suggested in the Australian concept for the SKA, the feed moves around the lens to implement pointing. Thus any departures from spherical symmetry in the surface of the lens could cause variations in the sidelobe pattern as a function of main-beam pointing. If more than one feed arm is used on each antenna, the sidelobe pattern of one feed is likely to be affected by scattering or blockage by other feeds.

10 Conclusions

Conclusions drawn from these considerations are necessarily rather tentative, since practical interference situations are complex and modeling requires simplifying assumptions. For spectral line observations, adaptive nulling may not provide sufficient mitigation when interference occurs close to the frequency of the line under study. For continuum observations the situation is better, provided that the fraction of the observing band that contains interference is not large. Deterministic nulling is likely to be limited by the variation of the sidelobe responses from one antenna to another.

The incorporation of wideband receiving systems into several major radio astronomy instruments under development introduces the risk that sensitivity will not always be limited by system noise, as in narrower-band instruments, but by a combination of noise and residual interference. To avoid degradation of sensitivity, the rms level of interference in an image should be no more than a fraction of the rms noise. (In calculations of detrimental thresholds of interference to radio astronomy in the ITU literature, an interference level of 1/10 the rms noise is used as a criterion for tolerance.) A number of mitigation techniques depend upon the identification of interference in the radio-telescope output data. These include the common practice of editing time sequences or spectra, as well as adaptive nulling. Elimination of interference in these cases is necessarily

limited by the system noise that is present. Techniques that are not directly limited by the noise include deterministic nulling, which makes use of independently-determined information on the interferer’s position, and adaptive cancelation, in which an additional antenna is used to obtain a version of the interference that is relatively free from noise (Barnbaum and Bradley 1998), or information about the expected signal content is used (Ellingson et al. 2001). These latter techniques are likely to be particularly useful in special cases, such as when interference occurs at the frequency of a spectral line under study. To facilitate deterministic nulling, the question of whether there are practical steps that can be taken to improve the matching of the sidelobe patterns of the elements of large- n arrays should be considered.

Acknowledgments Thanks are due to L. R. D’Addario and F. R. Schwab for helpful comments and discussions.

Appendix

A Interferometer response in matrix format

This derivation the covariance matrix from parameters of the source and antennas essentially follows that of Leshem et al. (2000). We start from the expression for the interferometer response, using the approximation in which the w component is omitted. This is given in Thompson, Moran and Swenson (2001), see Eq. (3.9) [or, in the first edition, Eq. (4.5)]:

$$\mathcal{V}(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{A_N(l, m) I(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-j2\pi(ul + vm)} dl dm. \quad (3)$$

Here \mathcal{V} is the complex visibility, u and v represent the projected baseline coordinates measured in wavelengths in a plane normal to the pointing direction, which is defined by direction cosines l and m . Three adjustments to the format are necessary. First, for simplification, we allow the factor $\frac{A_N(l, m)}{\sqrt{1 - l^2 - m^2}}$, which represents the antenna beam pattern, to be subsumed within the brightness function $I(l, m)$. Second, we assume that the astronomical brightness function can be represented by a finite number of values, p , at discrete pointings, spaced at intervals no smaller than the desired angular resolution. The pointing direction k is specified by direction cosines (l_k, m_k) . We replace the integrals with a summation over the pointings. Third, for each antenna, we specify the components in the (u, v) plane relative to a reference point which can be chosen, for example, to be the center of the array. The (u, v) values for a pair of antennas i and j then become $(u_i - u_j, v_i - v_j)$. This last modification allows the parameters of an antenna array to be specified in terms of individual antennas rather than the baselines. Equation (3) now becomes

$$\mathcal{V}(u_i - u_j, v_i - v_j) = \sum_{k=1}^p I_k e^{-j2\pi(u_i l_k + v_i m_k)} e^{j2\pi(u_j l_k + v_j m_k)}, \quad (4)$$

where $I_k = I(l_k, m_k)$. Note that u and v do not vary with the pointing k over the source but are defined for the phase reference position (field center). Equations (3) and (4) represent the visibility as measured by a single pair of antennas.

Next, for an array of n antennas, we define an $n \times p$ matrix containing terms corresponding to the first exponential of Eq. (3). To cover the general case in which the antennas have different response patterns, a

complex gain factor $g_i(l_k, m_k)$ is included for each antenna k and pointing direction (l_k, m_k) :

$$\mathbf{A} = \begin{bmatrix} g_1(l_1, m_1)e^{-j2\pi(u_1l_1+v_1m_1)} & g_1(l_2, m_2)e^{-j2\pi(u_1l_2+v_1m_2)} & \cdot & \cdot & \cdot & g_1(l_p, m_p)e^{-j2\pi(u_1l_p+v_1m_p)} \\ g_2(l_1, m_1)e^{-j2\pi(u_2l_1+v_2m_1)} & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ g_n(l_1, m_1)e^{-j2\pi(u_nl_1+v_nm_1)} & \cdot & \cdot & \cdot & \cdot & g_n(l_p, m_p)e^{-j2\pi(u_nl_p+v_nm_p)} \end{bmatrix}. \quad (5)$$

The antenna index increases downward across the n rows and the pointing index increases toward the right across the p columns.

To generate the covariance matrix we first define a $p \times p$ diagonal matrix containing the intensity values of the p source-model points:

$$\mathbf{B} = \begin{bmatrix} I_1 & & & & \\ & I_2 & & & \\ & & \cdot & & \\ & & & \cdot & \\ & & & & \cdot \\ & & & & & I_p \end{bmatrix}. \quad (6)$$

Now we can write

$$\mathbf{R} = \mathbf{A}\mathbf{B}\mathbf{A}^H, \quad (7)$$

where the superscript (H) indicates the Hermitian transpose (transpose plus complex conjugation). \mathbf{R} is the *covariance matrix*, which is Hermitian with dimensions $n \times n$. Each element of \mathbf{R} is of the form of the right-hand side of Eq. (4), that is, the sum of responses to the p brightness points for a specific pair of antennas. For row i and column j the element is

$$r_{i,j} = \sum_{k=1}^p g_i(l_k, m_k)g_j^*(l_k, m_k)I_k e^{-j2\pi(u_i l_k + v_i m_k)} e^{j2\pi(u_j l_k + v_j m_k)}, \quad (8)$$

which represents the cross correlation of signals from antennas i and j . (When the gain factors g are real and equal to unity the elements represent the source visibility \mathcal{V} .) At this point, the summation terms can be replaced by integrals over the source as in Eq. (3). The diagonal elements are the n self-products ($i = j$), which correspond to the total power response of each antenna. Note that \mathbf{R} is Hermitian: $r_{i,j} = r_{j,i}^*$. \mathbf{R} contains the full set of correlator output terms for an array of n antennas for a single averaging period and a single frequency channel. These data, when calibrated as visibility, can provide a snapshot image.

If the response patterns of the antennas are identical, i.e. $g_i = g_j$ for all (i, j) , then $g_i g_j^* = |g|^2$, and this (real) gain factor can be taken outside the matrix \mathbf{R} . Thus, to determine the angle of incidence (l, m) of a signal from the covariance measurements [the (u, v) values being known], the gain factors need not be known if they are identical from one antenna to another, but otherwise must be known.

B The effect of time averaging on the response to a geostationary interferer

An analysis of the interference-to-noise ratio (INR) for an interfering transmitter in a fixed location is given in Thompson (1982). The ratio of the rms levels of interference and noise in a synthesis map for which all

(u, v) tracks cross the u axis (i.e. an observing time of 12 h) is given by⁵

$$\frac{(|r_i|)_{\text{rms}}}{(|r_n|)_{\text{rms}}} = \frac{F_i c^2}{4\pi k T_s \nu^2 \sqrt{2\Delta\nu\omega_e \cos\delta(1 + |\sin\delta|)}} \frac{1}{\sqrt{\frac{1}{n_r} \sum_{n_r} q'}}, \quad (9)$$

where F_i is the flux density of the interfering signal, k is Boltzmann's constant, T_s is the system temperature, $\Delta\nu$ is the bandwidth of the receiver or spectral channel, ω_e is the angular velocity of the Earth (7.29×10^{-5}) rad s⁻¹, δ is the declination of the field center, n_r is the number of baselines, and q' is the length of a baseline projected onto the Earth's equatorial plane and measured in wavelengths. Here we are interested in the INR after integration for the full observing period, under the condition that the power of the interfering signal at the receiver input produces a component in the correlator output equal to the rms noise after averaging over the initial period, τ_i , discussed in §4. This condition is given by

$$\frac{F_i \lambda^2}{4\pi} = \frac{k T_s \Delta\nu}{\sqrt{\Delta\nu \tau_i}}, \quad (10)$$

where $\frac{\lambda^2}{4\pi}$ is the collecting area of an isotropic radiator which is used for the interference reception in Eq. (9). Then from Eqs. (9) and (10),

$$\frac{(|r_i|)_{\text{rms}}}{(|r_n|)_{\text{rms}}} = \frac{1}{\sqrt{2\omega_e \tau_i [\frac{1}{n_r} \sum_{n_r} q']}}, \quad (11)$$

where, to remove the dependence on δ , the factor $\sqrt{\cos\delta(1 + |\sin\delta|)}$ has been replaced by unity, resulting in an error of less than 1 dB for $0 < \delta < 71^\circ$. For configuration C of the VLA, $\frac{1}{n_r} \sum_{n_r} q'$ is equal to 1.43×10^3 m divided by the wavelength.

C Numerical illustration of the filtering process and widths of the nulls

A simple program written in Mathcad was used to demonstrate the filtering of interference by adaptive nulling as described in §3. The number of antennas considered was five or nine, The u and v coordinate values were for an east-west, linear, configuration. The four model configurations were as follows.

1. Five antennas spaced at intervals of 100λ in an east-west line, with three interferers.
2. Five antennas, as in (1) but with locations proportional to powers of 2 from one end, i.e. values of u equal to 0, 25, 100, 225, 400.
3. Five antennas, as in (1) but with only two interferers, with \mathbf{A} as a 5×2 matrix.
4. Nine antennas, east-west, spaced at intervals 50λ , with three interferers.

To simulate the reception of three interfering signals through the far sidelobes, values for direction cosines l and m for three interferers were chosen to represent essentially random angles of a few tenths of a radian from the center of the field ($l = m = 0$) at the zenith. The sidelobe responses were assumed to be different for each antenna and each interference direction, and random numbers generated by Mathcad were used to provide values of (voltage) gain g and phase ϕ for the sidelobes. From the numerical values of (u, v, l, m, g, ϕ) for the antennas and interferers, elements $a_{i,j}$ of the 5×3 , or 9×3 , matrix \mathbf{A} were derived as indicated in

⁵This can also be found as Eq. (14.11) in Thompson et al. (1986) and Eq. (15.12) in the revised edition (2001).

Eq. (5). Values representing the interfering signals were inserted into a diagonal matrix \mathbf{B} as in Eq. (6). Then from Eq. (7) the covariance matrix \mathbf{R} was obtained. Noise can be included by adding values representing the self products of the system noise (real numbers) to the diagonal elements of \mathbf{R} . However, the noise was omitted since it limits the depth of the null that can be investigated.

The values of the elements of \mathbf{R} represent the power of the interfering signals in the cross correlations measured with the array. As a measure of the overall interference, the (Euclidian) norm of \mathbf{R} was used, which is equal to the square root of the sum of the squares of the elements. For a Hermitian matrix, this is a purely real number equal to the root-sum-square of the arguments of the elements. The matrix \mathbf{U} containing the eigenvalues of \mathbf{R} was then calculated and, as expected, indicate two or three interferers as inserted in \mathbf{A} . The left two or three columns of \mathbf{U} , representing the interference, were removed, to obtain \mathbf{U}_N . Following Eq. (2), the filtered matrix \mathbf{R}_F was obtained. The degree of suppression of the interference was measured by taking the norm of \mathbf{R}_F divided by the norm of \mathbf{R} which was of order 10^{-15} , showing that the interference was effectively eliminated. This verified that adaptive nulling works in cases where the interference is received in sidelobes for which the phase and gain are different for each antenna.

To investigate the angular width of a null as a function of the depth (attenuation), an offset in the l value for the first interferer was introduced, while maintaining the values of \mathbf{U} to be the same as without the offset, so as not to allow the null to track the movement of the interferer. Again the ratio of the norms of \mathbf{R}_F and \mathbf{R} was used to determine the fraction of the power that was received from the first interferer as a result of it no longer being centered on the null. (Here the norm of \mathbf{R} with only the first interferer present was used.) Note that the depth of the null as calculated here is with respect to the gain levels of the sidelobes through which the interference is received. Thus the null depth is roughly with respect to an isotropic radiator, i.e. dBi in decibel form. The width of the null clearly depends upon the spacings of the antennas, since l and m occur in product terms with u and v , so it is useful to express the width as a fraction of the width of the synthesized beam, which is taken to be equal to λ/D where λ is the wavelength and D is here the longest spacing. Curves of null depth in dBs as a function of width in units of λ/D are given for four different configurations of antennas in Fig. 1. From bottom-right to top-left, these correspond to for the array models in order listed above (the curve for model 2 is dotted).

For any particular width of null in Fig. 1 the depth varies over a range of approximately 10 dB and, for example, for a 30 dB null the mean of the range of width values is approximately 0.06 in beamwidth units. Only single nulls were considered, and nulls that are wider in both angle and frequency can be formed by placing two or more close together (e.g. Harp 2002). However, in some cases distortion of the main beam tends to be increased if nulls are clustered (Ellingson and Cazemier 2003).

References

- Barnbaum, C. and R. F. Bradley, A New Approach to Interference Excision in Radio Astronomy: Real-Time Adaptive Cancellation, *Astron. J.*, **116**, 2598-2614, 1998.
- Bower, G. C., Simulations of narrow-band Phased Array Null Formation for the ATA, *ATA Memo 37*, SETI Institute, 2001.
- Ellingson, S. W., J. Bunton, and J. F. Bell, Removal of the Glonass C/A Signal from OH Spectral Line Observations using a Parametric Modeling Technique, *Astrophys. J. Suppl.*, **135**, 87-93, 2001.
- Ellingson, S. W. and W. Cazemier, Efficient Multibeam Synthesis with Interference Nulling for Large Arrays, *IEEE Trans. Antennas Propag.*, **51**, 503-511, 2003.
- Ellingson, S. W. and G. A. Hampson, A Subspace Tracking Approach to Interference Nulling for Phased Array Based Radio Telescopes, *IEEE Trans. Antennas Propag.*, **50**, 25-30, 2002.
- Fridman, P. A. and W. Baan, RFI Mitigation Methods in Radio Astronomy, *Astron. Astrophys.*, **378**, 327-344, 2001.

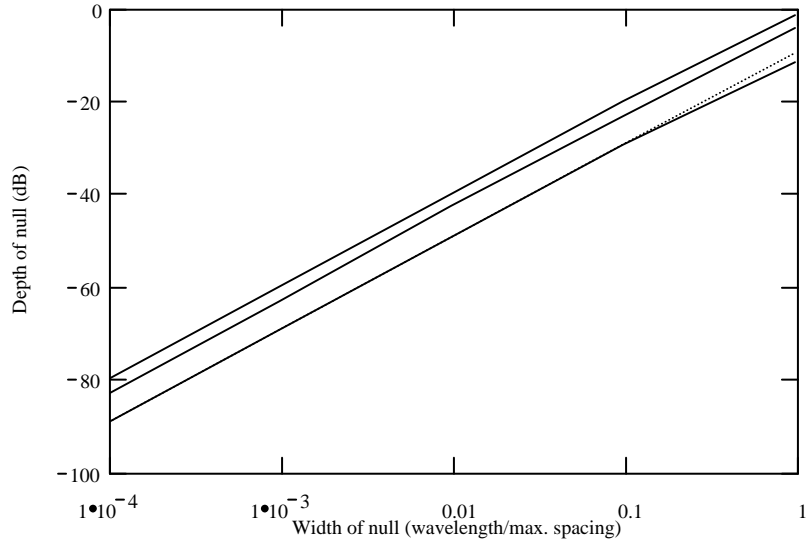


Figure 1: Ordinate is depth of null in decibels, that is, attenuation relative to the gain of a single antenna for the same direction. The abscissa is the angular width of a null in units of (wavelength) \div (maximum array spacing) which is approximately the half-power width of the array beam. Curves are for four model arrays described in the text.

Harp, G. R., Customized Beam Forming at the Allen telescope Array, *ATA Memo 51*, SETI Institute, 2002.

International Telecommunication Union (ITU), *Handbook on Radio Astronomy*, ITU Radiocommunication Bureau, Geneva, 1995.

International Telecommunication Union (ITU), *Recommendation ITU-R RA.769, Protection Criteria used for Radioastronomical Measurements*, ITU Radiocommunication Bureau, Geneva.

Leshem, A. and A.-J. van der Veen, Introduction to Interference Mitigation Techniques in Radio Astronomy, in *Perspectives in Radio Astronomy - Technologies for Large Antenna Arrays*, A. B. Smolders and M. P. van Haarlem, eds., Netherlands Foundation for Radio Astronomy 1999.

Leshem, A. and A.-J. van der Veen, the Effect of Adaptive Interference Suppression on Radio Astronomical Image Formation, in *Radio Telescopes*, H. R. Butcher, Ed., *Proc. S.P.I.E.*, **4015**, 341-352, 2000.

Leshem, A., A.-J. van der Veen, and A.-J. Boonstra, Multichannel Interference Mitigation Techniques in Radio Astronomy, *Astrophys. J. Suppl.*, **131**, 355-373, 2000.

Smolders, B. and Hampson, G., Deterministic Nulling in Phased arrays for the Next Generation of Radio Telescopes, *IEEE Antennas Propag. Mag.*, **44**, No. 4, 13-22, 2002.

Thompson, A. R., The Response of a Radio Astronomy Synthesis Array to Interfering Signals, *Proc. IEEE Trans. Antennas Propag.*, **AP-30**, 450-456, 1982.

Thompson, A. R., J. M. Moran, and G. W. Swenson Jr., *Interferometry and Synthesis in Radio Astronomy*, Wiley, New York, 1986, 2001.

Welch, W. J., Forming Nulls toward Satellites, *ATA Memo 7*, SETI Institute, 2000a.

Welch, W. J., Forming a Broadband Null with the ATA using the Tree Algorithm, *ATA Memo 13*, SETI Institute, 2000b.

Yang B. Projection Approximation Subspace Tracking, *IEEE Trans. Signal Processing*, **43**, 95-107, 1995.