Supplementary Notes to Interferometry and Synthesis in Radio Astronomy

Dec. 16, 2010

Contents

*Asterisk indicates recent addition.

Trends Since 2001 (May 20, 2009)

(Relates mainly to Chs. 1 and 5)

Some general trends in the system design of interferometers and arrays have appeared during the years since the publication of the second edition of *Interferometry...*(2001), and such improvements continue to evolve. These include wider bandwidths, larger and faster correlators and greater use of large numbers of small antennas to optimize the cost per unit of collecting area. The continuing need for greater sensitivity in radio astronomy has led to the concept of an internationally funded instrument with very large collecting area, known as the Square Kilometer Array (SKA). The original Square Kilometer Array (SKA) specification was based on a system temperature of 50 K, and has since developed into what can be described as a substantial fraction of a square kilometer with system temperature ~ 20 K. The desirable frequency range goes from ~ 100 MHz to cover extreme HI redshifts to $20 - 30$ GHz to cover planets and some molecular lines. The astronomical requirement for such an instrument have been widely studied: see, for example, a series of papers in New Astronomy Reviews, 48, 979-1563, 2004. To help focus on the main performance requirements, five particular astronomical programs have been considered, as follows.

- Probing the "dark ages" and the epoch of reionization (EoR) through observations of highly redshifted HI and of molecular gas, dust, and star formation in the earliest galaxies (0.1-20 GHz).
- Galaxy evolution and cosmology, from measurements of neutral hydrogen (HI) and continuum emission from star forming regions, etc. (0.3-1.4 GHz)
- The origin and evolution of cosmic magnetism, through measurements of Faraday rotation and the Zeeman effect over a wide range of redshifts (0.3-10 GHz).
- Strong field tests of gravity using pulsars and black holes (0.5-15 GHz).
- "The Cradle of Life", i.e. the search for terrestrial type planets and astrobiology ($\geq$20 GHz).

The frequency ranges for the programs given above are from Carilli and Rawlings (2004).

A large fraction of the proposed programs for the SKA involve surveys over the observable sky for galaxies, pulsars, transient phenomena etc. and to map the redshifted HI distribution. Thus in addition to the requirement for sensitivity, speed of coverage of the sky for surveys has become a major consideration. Maximization of the efficiency of sky coverage improves statistical studies and increases the probability of success in detecting transient events such as supernovae. It also increases the liklihood of discovering new phenomena. These considerations are mainly applicable to the lower frequencies, i.e. less than ~ 2 GHz, at which antenna systems that provide instantaneous wide-angle coverage are most practicable.

The high end of the frequency requirement for the SKA is generally taken to be about 25 GHz, set by the desire to search for Earth-like planets and the study of astrochemistry. The low end of the frequency range is set by the detection of the decrease in the neutral hydrogen radiation at the epoch of reionization (EoR) at an estimated redshift $z \simeq 11$. The original specification of the collecting area A and system temperature T for the SKA was $A/T = 20,000\, m^2 K^{-1}$, e.g., $A = 1\,\mathrm{km}^2$ with $T = 50\,\mathrm{K}$. With lower system temperature it is practical to use, say, $A = 0.4\,\mathrm{km}^2$ with $T = 20\,\mathrm{K}$, and thus reduce the physical collecting area, and hence also the overall cost. A wide instantaneous field of view is an important factor in obtaining high survey speed, so design ideas tend toward use of large numbers of small antennas and multiple simultaneous beams. The frequency range envisaged for the SKA is too large to be covered efficiently by a single type of antenna.

SKA Antennas

1. Low-frequency range

At the low frequency end, up to about 300 MHz, arrays of broadband dipoles mounted over a ground-plane reflecting screen are the most practical scheme. Dipoles are robust and relatively cheap and crossed dipoles provide full polarization data. Low-noise transistor front ends can operate at ambient temperature at these

frequencies, where the system noise level is set largely by radiation from the sky. Signal from groups of dipoles are combined and phases adjusted to form beams which can be pointed as required without the need for moving parts. Phase adjustments can be performed using analog methods, or digitally if the signals have been digitized after amplification. If the spacing between the centers of the dipoles in such an array is greater than $\lambda/2$, the array is said to be sparse. The collecting area is maximized at $\lambda^2/4$ per element, but because of the spacing the grating sidelobes begin to be significant as $\lambda/2$ is exceeded. If the spacing is less than $\lambda/2$, the array is said to be compact. The effective area is then less than $\lambda^2/4$ per element, but grating lobes are avoided. These sparse or compact conditions apply to arrays of dipoles or any other elements that are spaced by about $\lambda/2$ and combined to form beams. The variation of the path length through the ionosphere is a serious problem in imaging at these low frequencies, but it is possible to calibrate the ionosphere over a wide angular range by forming beams in the directions of calibration sources for which the positions are accurately known. The Low Frequency Array (LOFAR) (Kassim et al. 1999, 2004, Stewart et al. 2004.) provides an example of this type of instrumentation.

2. Mid-frequency range

In the middle frequency range, approximately 0.3-2 GHz, two possible solutions are under development. For the frequencies up to about 1 GHz, systems known as aperture arrays are being developed. These can take the form of half-wave dipoles over a ground screen but, especially at the shorter wavelengths, arrays of Vivaldi antennas (Schaubert and Chio 1999) are being used. The Vivaldi elements are formed on strips of aluminum or of copper-clad insulating board. By using two sets of elements running in orthogonal directions, full polarization is obtained. Each Vivaldi element feeds a balanced transistor amplifier, and for convenience the system is produced in units of approximately one square meter size, referred to as Vivaldi tiles. The approximate spacing between adjacent Vivaldi elements is $\lambda/2$, and with this arrangement approximately four amplifiers are required for each square wavelength of collecting area, e.g. $\sim$44 amplifiers per square meter at 1 GHz. Aperture arrays provide multiple beams with rapid and flexible pointing. The number of beams available depends upon the beamformer and signal processing system. The sparse and compact conditions described above apply also to such arrays.

For the range from $\sim$ 700 MHz to $\sim$ 2 GHz, dish-type antennas with multiple beams become more practicable than aperture arrays since, for a given collecting area, they do not require such large numbers of low noise amplifiers and phasing components. Note that with feeds in the form of a *focal plane array*, i.e. an array of single pixel feeds in the focal plane of an antenna, it is usually not possible to get the feeds close enough together to avoid gaps between the individual beams. As a result it is necessary to interleave the antenna pointings to obtain complete sky coverage. Thus it is preferred to use *phased-array feeds* in which an array of closely spaced receiving elements are arranged in the focal plane. Any one beam is formed as a phased combination of the signals from a number of the feed elements, and such combinations can be chosen to provide optimum beam spacings for efficient sky coverage. The elements are individually terminated with matched amplifiers, but mutual coupling between the elements cannot be avoided, so the design and adjustment of phased-array feeds is generally more critical than for focal plane arrays. Two designs of phased-array feeds have been developed, one based on the Vivaldi system mentioned above, and one using a "checkerboard" conductor pattern (Hay et al. 2007). The checkerboard scheme can be envisaged as a series of conducting elements on a circuit board that are arranged like the black squares of checkerboard, the remainder of the board being blank. At each point where two corners of conducting squares meet, the corners do not touch but each feeds one input of a balanced amplifier. An interesting feature of the checkerboard scheme is that the patterns of conducting and nonconducting surfaces are identical and thus self-complimentary. A screen of this form in free space is well matched with load impedances of 377 ohms between the corner pairs of conducting squares where the amplifiers are connected, as follows from Booker's formula (which can be found in most antenna textbooks). For use as a feed array, the checkerboard screen is mounted over a ground plane which introduces some frequency variation of the impedance. Development of phased-array feeds is still continuing and those currently in use are mainly in "pathfinder" projects for larger arrays.

3. High-frequency range
At the highest frequencies, approximately ∼3-25 GHz, it is necessary to use cryogenically cooled amplifiers to obtain the lowest system temperatures, and hence the highest sensitivity. In practice this requires that the antennas use some form of concentrators that focus the radiation received onto small area, i.e. in practice parabolic reflectors. (Possible use of spherical dielectric-foam lenses, which allow simultaneous use of multiple feeds, was one of the early SKA considerations, but found to be impractical because of dielectric losses.) The need for survey speed is generally not considered essential in the high frequency range. To obtain a large collecting area, a large number of small dishes (LNSD) is more cost effective than a smaller number of large dishes. The optimum dish diameter in such systems generally in the range 6-15 m. The Allen Telescope Array (Deboer et al. 2004) provides an example of the LNSD concept. Feeds considered are generally single-pixel type (i.e. single beam), and cover a broad bandwidth (∼10:1 frequency range).

Bandwidth and Digitization of Signals
Receiver bandwidths continue to increase to provide greater continuum sensitivity and spectral line coverage. As an example, in the redesigned system for the VLA (which has now become the EVLA), the total instantaneous bandwidth is increased from 200 MHz to 16 GHz (Perley et al. 2009). A limiting factor is the cost and complexity of the correlators and signal processing required. Digitization of the signals at an earlier point within the receiver processing chain is another important trend. In large instruments like the EVLA and ALMA it is performed at the antennas and signal transmission to the correlator is in digital form. Effects of the frequency response of the transmission medium, and the possibility of crosstalk between signals, are both eliminated. Performing more of the signal processing digitally, for example using polyphase filters for channelization, is advantageous since replication of responses does not depend upon matching of analog circuit elements. With the wider frequency coverage it is generally expected that much of the radio frequency interference encountered can be mitigated by designing for high dynamic range so that the analog stages are not driven into nonlinearity, combined with the use of spectral correlators which allow contaminated channels to be rejected.

References

Carilli, C. R. and S. Rawlings, Motivation, key science projects, standards and assumptions, New Astronomy Reviews, 48, 979-984, 2004.

Deboer, D. and 16 coauthors, The Allen Telescope Array, Experimental Astronomy, 17, 19-34, 2004.

Hay, S. G., J. D. O'Sullivan, J. S. Kot, C. Granet, A. Grancea, A. R. Forsythe, and D. H. Hayman, Focal plane array development for ASKAP, Proc. European Conf. on Ant. and Prop.: EuCAP 2007.

Kassim, N. E., T. J. W. Lazio, P. S. Ray, P. C. Crane, B. C. Hicks, K. P. Stewart, A. S. Cohen, and W. M. Lane, The Low-Frequency Array (LOFAR): Opening up a New Window on the Universe, Planetary Space Sci., 52, 1343-1349, 2004.

Kassim, N. E., T. J. W. Lazio, K. W. Weiler, W. C. Erickson, and D. L. Jones, Opening a New Window on the Universe: High Resolution, Long Wavelength Radio Astronomy from the Ground (LOFAR) and from Space (ALFA), in Perspectives on Radio Astronomy: Science with Large Antenna Arrays, M. P. van Haarlem, ed., Astron, Netherlands, pp. 11-23, 1999.

Perley, R. and eleven coauthors, The Expanded Very Large Array, Proc. IEEE, in publication, 2009.

Schaubert, D. H. and T. H. Chio, Wideband Vivaldi Arrays for Large Aperture Arrays, in Perspectives on Radio Astronomy: Technologies for Large Arrays, A. B. Smolders and M. P. van Haarlem eds, Astron, Netherlands, pp. 49-57, 1999.

Stewart, K. P., B. C. Hicks, P. S. Ray, P. C. Crane, N. E. Kassim, R. F. Bradley, and W. C. Erickson, LOFAR antenna development and initial observations of solar bursts, Planet. Space Sci., 52, 1351-1355, 2004.

Further Information on the SKA

Science with the Square Kilometer Array, C. Carilli and S. Rawlings, eds., New Astronomy Reviews, 48 (nos iv-vii), pp. 979-1563, 2004. Also as a single volume, Elsevier, Amsterdam, 2004.

The Square Kilometer Array: An Engineering Perspective, P. J. Hall, ed., Experimental Astronomy, 17, (nos. 1-3), 2004. Also as a single volume, Springer, Dordrecht, 2005.

SKA Memo Series can be found at http://www.skatelescope.org/pages/documents_3htm

Reference on Source Confusion (May 11, 2006)

(Relates to Ch. 1, pp. 24-25)

Source confusion in early observations is briefly discussed on pp. 24-25 (p. 22 of first ed.). A classic reference on the subject that should be mentioned is:

Scheuer, P. A. G., A Statistical Method for Analyzing Observations of Faint Radio Sources *Proc. Cambridge Phil. Soc.*, **53**, 764-773, 1957.

For a random distribution of faint sources with a model intensity distribution, the statistics of the signal amplitudes observed with an interferometer are derived. This method was used in the analysis of early interferometer surveys to examine the intensity distribution of sources below the confusion limit.

Response of an Array in Matrix Format (May 11, 2006)

(Relates to Ch. 3)

The data gathered by a synthesis array can be expressed in the form of a covariance matrix, which is useful for certain analytical procedures. The discussion below largely follows that given by Leshem, van der Veen, and Boonstra (2000). We start from the expression for the interferometer response, using the approximation in which the w component is omitted. This is given in Eq. (3.9) [or, in the first edition, Eq. (4.5)]:

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

Here $\mathcal{V}$ is the complex visibility and u and v represent the projected baseline coordinates measured in wavelengths in a plane normal to the phase reference direction. We make four adjustments to the equation. First, we replace A_N by the product of the corresponding complex voltage gain factors $g_i(l,m)g_j^*(l.m)$, where i and j indicate the antennas: constants resulting from conversion of aperture to gain, etc. can be ignored since, in practice, the intensity scale is determined by calibration. Second, we allow the factor $(\sqrt{1-l^2-m^2})^{-1}$ to be subsumed within the intensity function $I(l,m)$. Third, we assume that the astronomical brightness function can be represented by a point-source model with a finite number of points p. To represent extended structure points may be spaced at intervals smaller than the angular resolution. For each point k the direction is specified by direction cosines (l_k, m_k). We replace the integrals with a summation over the points. Fourth, 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 $(u_i - u_j, v_i - v_j)$. The first and fourth modifications allow the parameters involved to be specified in terms of individual antennas rather than antenna pairs. Equation (1) can now be written as:

$$\mathcal{V}(u_i - u_j, v_i - v_j) = \sum_{k=1}^{p} I_k\, g_i(l_k, m_k) e^{-j2\pi(u_i l_k + v_i m_k)} g_j^*(l_k, m_k) e^{j2\pi(u_j l_k + v_j m_k)}, \tag{2}$$

where $I_k = I(l_k, m_k)$. Note that u and v do not vary with the point-source positions but are defined for the phase reference position (field center). Equations (1) and (2) 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 antenna gain and exponential of Eq. (2) (i.e. the terms associated with antenna i):

$$\mathbf{A} = \begin{bmatrix} g_1(l_1,m_1)e^{-j2\pi(u_1l_1+v_1m_1)} & g_1(l_2,m_2)e^{-j2\pi(u_1l_2+v_1m_2)} & \cdots & g_1(l_p,m_p)e^{-j2\pi(u_1l_p+v_1m_p)} \\ g_2(l_1,m_1)e^{-j2\pi(u_2l_1+v_2m_1)} & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ \cdot & \cdot & \cdots & \cdot \\ g_n(l_1,m_1)e^{-j2\pi(u_nl_1+v_nm_1)} & \cdot & \cdots & g_n(l_p,m_p)e^{-j2\pi(u_nl_p+v_nm_p)} \end{bmatrix}. \tag{3}$$

The antenna index increases downward across the n rows and the point-source 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:

$$\mathbf{B} = \begin{bmatrix} I_1 & & & & \\ & I_2 & & & \\ & & \cdot & & \\ & & & \cdot & \\ & & & & \cdot \\ & & & & I_p \end{bmatrix}. \tag{4}$$

Now we can write

$$\mathbf{R} = \mathbf{ABA}^H, \tag{5}$$

where the superscript $(^H)$ indicates the Hermitian transpose (transposition 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. (2), that is, the sum of responses to the p intensity points for a specific pair of antennas. For row i and column j the element is

$$r_{i,j} = \sum_{k=1}^{p} g_i(l_k,m_k)g_j^*(l_k,m_k)I_k\, e^{-j2\pi(u_il_k+v_im_k)}e^{j2\pi(u_jl_k+v_jm_k)}, \tag{6}$$

which represents the cross correlation of signals from antennas i and j. When the gain factors g are equal to unity the elements represent the source visibility $\mathcal{V}$. The diagonal elements are the n self-products ($i = j$), which represent the total power responses of the antennas. 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. In cases where the w component is important, a term of the form $w(\sqrt{1-l^2-m^2}-1)$ [as in Eq. (3.7), or Eq. (4.3) of the first edition], with appropriate subscripts, can be included within each exponent.

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.

The covariance matrix can also be formulated in terms of the complex signal voltages from the antennas of an array. Let the signal from antenna k be x_k, which is a function of time. For the array, the signals can be represented by a (column) vector $\mathbf{x}$ of dimensions $n \times 1$, each term of which corresponds to the sum of the terms in the corresponding row of the matrix in Eq. (3). The outer (or Kronecker) product $\mathbf{x} \otimes \mathbf{x}^H$ leads

to a covariance matrix:

$$\mathbf{R}' = \begin{bmatrix} x_1 \\ x_2 \\ \cdot \\ \cdot \\ \cdot \\ x_n \end{bmatrix} \otimes \begin{bmatrix} x_1^* & x_2^* & \cdot & \cdot & \cdot & x_n^* \end{bmatrix} = \begin{bmatrix} x_1x_1^* & x_1x_2^* & \cdot & \cdot & \cdot & x_1x_n^* \\ x_2x_1^* & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ x_nx_1^* & \cdot & \cdot & \cdot & \cdot & x_nx_n^* \end{bmatrix}. \quad (7)$$

The elements of the matrix $\mathbf{R}$, as in Eq. (6), represent the correlator outputs which involve a time average of the signal products. If the signal products in the elements of $\mathbf{R}'$ are similarly understood to represent time averaged products, then $\mathbf{R}'$ is equivalent to the covariance matrix $\mathbf{R}$.

The matrix formulation is commonly used in non-astronomical applications of array signal processing. An example of applications in radio astronomy is provided by the discussion of gain calibration by Boonstra and van der Veen (2003). Also, the eigenvectors of the matrix can be used to identify interfering signals that are strong enough to be distinguished in the presence of the noise. Such signals can then be removed from the data as discussed, for example, by Leshem, van der Veen, and Boonstra (2000).

References

Boonstra, A.-J. and A.-J. van der Veen, Gain Calibration Methods for Radio Telescope Arrays, *IEEE Trans. Signal Processing*, **51**, 25-38, 2003.

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.

Three Dimensional Imaging (April 10, 2007)

(Relates to Sections 3.1 and 11.9)

An observation by Sault et al. (1997) provides the first example of three dimensional imaging by Fourier synthesis. This is an image of the radiation belts surrounding the planet Jupiter, and the emission per unit volume of space is surveyed. The three dimensions in the image become possible because of the rotation of the planet which allows the distribution of emission as seen from aspect angles covering 360° to be combined. The location of a radiating element in the vicinity of the planet can be determined by viewing it from two different aspect angles, as in Fig. 1. Viewing from all aspect angles allows a mapping of the area surrounding the planet in the two dimensions of the diagram and the resolution of the synthesized beam also provides information in the third dimension normal to the page. In terms of the projection-slice theorem of Bracewell (see TMS2 pp. 65-66), the two-dimensional Fourier transform of an image provides a slice through the three dimensional Fourier transform of the planet and associated belts. Different aspect angles provide different slices from which the three-dimensional image can be obtained. It is assumed that the belts rotate with the planet, being 'frozen' into the magnetic field of the planet. The rotation period of the planet is known with respect to the magnetic field since details on the solid surface are not visible (Cox 2000).

For the three-dimensional observation the coordinate systems in which the astronomical object and the interferometer baselines are specified are different from those used in conventional two-dimensional imaging, which are formulated in Section 3.1. Instead of specifying the astronomical object in term of the ℓ and m coordinates of hour angle and declination, the Jovian belts are specified in three dimensions (x, y, z). The origin of these coordinates is the center of the planet with one axis aligned toward the planet's pole and the other two in the equatorial plane. The interferometer baseline is specified in a system that is related to the (x, y, z) coordinates in a manner similar to that used in conventional synthesis, which is shown in TMS Fig.

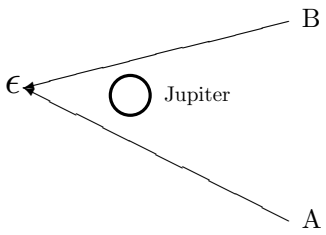

Figure 1: The relative position of a radiating element ϵ within the Jovian belts can be determined by observation from two different aspect angles A and B. The same principle is applied in tomography.

3.2. In the conventional case the baseline coordinate u is the projection of the baseline onto a plane containing the ℓ axis and the location of the observer[1] and the v coordinate is the baseline projection onto the plane containing the m axis and the observer. For Jupiter we represent the baseline coordinates by (U, V, W), using capitals to emphasize that they are not the same as the (u, v, w) used in conventional mapping [however, (U, V, W), like (u, v, w), are measured in wavelengths]. In the analysis of Sault et al., the interferometer baseline is first projected onto a plane normal to the direction of the planet. If $\mathbf{D}_\lambda$ represents the baseline vector and $\mathbf{r}$ is a unit vector in the direction of the planet as seen by the observer, the *projected* baseline is $\mathbf{D}_\lambda - (\mathbf{n} \cdot \mathbf{D}_\lambda)\mathbf{n}$. To define U the *projected* baseline is further projected onto the plane defined by the x axis and the array location, and V and W are similarly defined by projection onto the y and z axes. Note that in conventional synthesis the w component represents the component of the baseline in the direction of the field center, i.e., $(\mathbf{n} \cdot \mathbf{D}_\lambda)\mathbf{n}$, which is omitted in the Jupiter analysis. Thus it is implicitly assumed that the Jupiter case is one where the synthesized field is sufficiently limited that w can be ignored (see discussion in TMS Section 3.1). The use of three components in the Jupiter observation results entirely from the ability to examine the three dimensional nature of the object under study, in contrast to conventional synthesis in which the object is treated as a two-dimensional intensity distribution on the celestial sphere. The baselines and the positions within the source are expressed in conjugate systems (U, V, W) and (x,y,z), and the three dimensional visibility distribution can be written as

$$\mathcal{V}(U, V, W) = \left(\frac{R_0}{R}\right)^2 \int I_0(x, y, z)\, e^{-j2\pi(xU+yV+zW)/R} dx\, dy\, dz, \tag{1}$$

where R_0 is a distance of 4.04 AU conventionally used for standardization of Jupiter measurements, R is the true distance at the time of observation, and I_0 is the intensity of the radiation at a distance R_0. Imaging is performed by inverting the Fourier transform in Eq. (1). Note that it has been assumed that the radiating medium is optically thin, and that the primary antenna beam is wide enough that it can be ignored. Sault et al. (1997) observed Jupiter for 10 days using the Australia Telescope Compact Array at 13 and 22 cm wavelength. The thermal component from the planet was removed from the data and synthesized images of the radiation belts as seen from several angles can be found in Sault et al. (1997) and Leblanc et al. (1997).

The three dimensional approach requires that the direction of the rotation axis of the target object should not be close to the direction of observation, and in this respect Jupiter is nearly ideal as these two directions are within a few degrees of perpendicular. Also Jupiter's rotation period of 9.9 hours is very convenient. The requirement of the varying aspect angle limits the three dimensional technique to objects within the solar system, of which Jupiter is the best candidate. In the case of the sun, much of the interesting activity occurs on timescales shorter than the rotation period.

References

Cox, A. N., Editor, Allen's Astrophysical Quantities, Springer Verlag, N.Y., 2000

[1]The dimensions of the observer's array are assumed to be very small compared with the distance and dimensions of the object being observed.

Leblanc, Y.,G. A. Dulk, R. J. Sault, and R. W. Hunstead, The Radiation Belts of Jupiter at 13 and 22 cm, I Observations and 3-D Reconstruction, Astron. Astrophys. 319, 274-281,1997

Sault, R. J., T. Oosterloo, G. A. Dulk, and Y. Leblanc, The First Three-Dimensional Reconstruction of a Celestial Object at Radio Wavelengths: Jupiter's Radiation Belts, Astron. Astrophys. 324, 1190-1196, 1997.

Polarimetry (May 11, 2006)

(Relates to Section 4.8, p. 109)

A further paper on the Hamaker-Bregman-Sault algorithm:

Hamaker, J. Matrix self-alignment at work: Processing a simulated observation, ASTRON, March 17, 2003.

Low Frequency Antennas (April 10, 2007)

(Relates to Ch. 5)

A number of arrays are being brought into operation in the frequency range 10-100 MHz for study of the most distant universe including red-shifted neutral hydrogen just prior to the epoch of reionization, e.g. LOFAR (Kassim et al. 2004). Such arrays require large numbers of dipoles, often arranged in subarrays. The noise level that limits the sensitivity in this frequency range results mainly from the background radiation of the Galaxy, rather than from the receiver input stages as is generally the case at higher frequencies. Thus to achieve the maximum sensitivity it is necessary only to match the antennas to the receivers sufficiently well that the total noise is dominated by the component received by the antennas. This is an advantageous situation since it allows the dipoles to be used over a much wider frequency range than is possible when the impedance must be well matched. A review of the characteristics of dipole antennas under these conditions is given by Ellingson (2005).

Let γ be the power ratio of the background noise received from the sky to the noise contributed by the receiver. Then we have,

$$\gamma \simeq e_r \frac{T_{sky}}{T_{rec}}(1 - |\Gamma^2|), \tag{1}$$

where e_r (< 1) is an efficiency factor that results largely from the ohmic losses in the ground and in the dipole, T_{sky} is the noise brightness temperature of the sky, T_{rec} is the noise temperature of the receiver, and Γ is the voltage reflection coefficient at the antenna looking toward the receiver. Γ is given by

$$\Gamma = \frac{Z_{rec} - Z_{ant}}{Z_{rec} + Z_{ant}}, \tag{2}$$

where Z_{rec} and Z_{ant} are the impedances at the receiver and antenna terminals respectively. For dominance of the sky noise we take γ greater than ~ 10. T_{sky} is related to the intensity of the background radiation I_ν ($\mathrm{Wm^{-2}Hz^{-1}ster^{-1}}$), by $T_{sky} = c^2 I_\nu / 2k\nu^2$, where c is the speed of light and k is Boltzmann's constant.

A model for the intensity I_ν is given by Dulk et al. (2001) based on measurements by Cane (1979):

$$I_\nu = I_g \nu^{-0.52} \frac{1 - e^{-\tau(\nu)}}{\tau(\nu)} + I_{eg} \nu^{-0.80} e^{-\tau(\nu)}, \tag{3}$$

where $I_g = 2.48 \times 10^{-20}$ is the galactic component of the intensity, $I_{eg} = 1.06 \times 10^{-20}$ is the extragalactic component, and $\tau(\nu) = 5.0\,\nu^{-2.1}$. This model applies broadly over the sky except near the galactic plane where higher intensities are encountered. A graph of T_{sky} based on this model is shown in Fig. 1.

Ellingson (2005) considers performance of an inverted-V dipole design in which the conducting elements slope downward from the center at 45°. A wide frequency response is obtained with Z_{rec} in the range 200-800Ω. Computed responses indicate usable beamwidths in the range $120 - 140°$. Stewart et al. (2004) describe design of an inverted-V dipole in which the effective width of the conducting arms is increased in one dimension, which reduces the impedance variation with frequency compared to that of a dipole with single-wire elements.

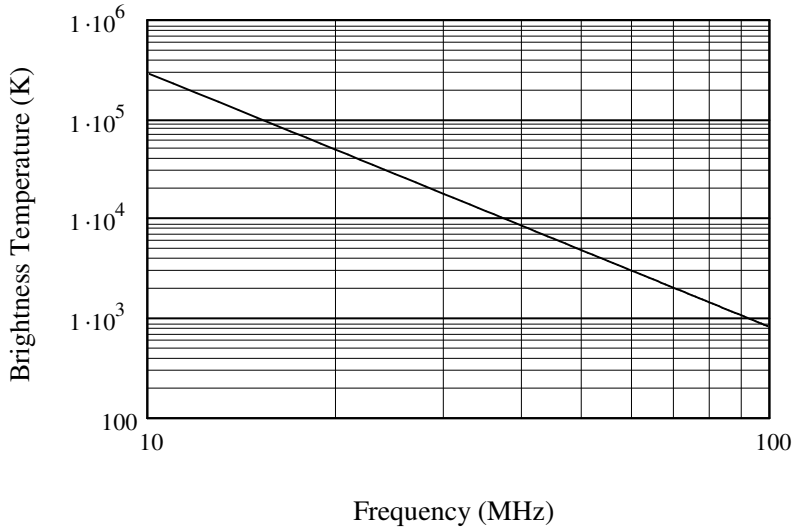

Figure 1: Background temperature T_{sky} derived from the intensity given in Eq. (3).

References

Cane, H. V., Spectra of the Nonthermal Radio Radiation from the Galactic Polar Regions, Mon. Not. Royal Astron. Soc., 149, 465-478, 1979.

Dulk, G. A., W. C. Erickson, R. Manning, and J.-L. Bougeret, Calibration of Low-Frequency Radio Telescopes Using Galactic Background Radiation, Astron. Astrophys., 365, 294-300, 2001.

Ellingson, S. W., Antennas for the Next Generation of Low-Frequency Radio Telescopes, IEEE Trans. Antennas, Propagat., 53, 2480-2489, 2005.

Kassim, N. E., T. J. W. Lazio, P. S. Ray, P. C. Crane, B. C. Hicks, K. P. Stewart, A. S. Cohen, and W. M. Lane, The Low-Frequency Array (LOFAR): Opening up a New Window on the Universe, Planetary Space Sci., 52, 1343-1349, 2004.

Stewart, K. P., B. C. Hicks, P. S. Ray, P. C. Crane, N. E. Kassim, R. F. Bradley, and W. C. Erickson, LOFAR Antenna Development and Initial Observations of Solar Bursts, Planetary Space Sci. 52, 1351-1355, 2004.

Paraboloidal Antennas (May 20, 2009)

(Relates to Chapter 5)

Baars, J. W. M., The Paraboloidal Reflector Antenna in Radio Astronomy and Communication, Springer, 2007

The FFT (Fast Fourier Transform) Telescope (Dec. 14, 2010)

(Relates to Ch. 5)

Tegmark and Zaldarriaga (2009a) have pointed out that it is possible to make an array with a large number of elements in which the usual correlator function is implemented by use of the fast Fourier transform (FFT) algorithm, which could greatly reduce the overall computation required. This idea requires the antennas to be arranged in a planar rectangular grid, or hierarchy of such grids. The purpose of this note is to describe the basic principle in terms of functions defined by Bracewell (2000, ch. 3), which are convenient for consideration of analysis involving Fourier transforms, and also to compare the results of the method with those of the usual procedure.

As a start, it is useful to introduce two functions. The two-dimensional *autocorrelation* function of $f(x, y)$ is defined as

$$f \star f = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) f^*(x - u, y - v) dx dy, \tag{1}$$

where the star represents correlation and the superscript asterisk indicates the complex conjugate. The *self-convolution* of $f(x, y)$, represented by the in-line asterisk, is

$$f * f = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) f^*(u - x, v - y) dx dy. \tag{2}$$

The formation of the visibility values measured in a multi-element array can be expressed as an autocorrelation. To illustrate this concept, consider a rectangular array of antennas in the horizontal (x, y) plane, spaced at intervals Δx and Δy. There are n_x antennas in each row in the x-direction, and n_y in each row in the y-direction. The antennas receive signals from a radio source, and the receiver outputs are sampled at the Nyquist rate. At any instant the outputs indicate the complex signal amplitudes[1] at the antenna locations. These signals can be represented by a rectangular array of two-dimensional delta functions in the (x, y) plane, each of magnitude equal to the complex amplitude at the corresponding point. The autocorrelation function of this delta-function array is non-zero only for offsets (u, v) equal to integral numbers of increments Δx and Δy. For example, in the case where the offset is one spacing in each direction, i.e. Δx in the x direction and Δy in the y direction, the autocorrelation is the sum of the products of one signal amplitude with the complex-conjugate of the other, for every antenna-pair for which $u = \Delta x$ and $v = \Delta y$, where u and v are the usual baseline coordinates[2]. Thus the autocorrelation gives a series of values of the visibility $\mathcal{V}(u, v)$, weighted in proportion to the number of antenna pairs for each (u, v) baseline[3]. It produces

[1] Since visibility is a complex function and is formed by combining signal samples from pairs of antennas, the relative phases of the signals must be measured. In principle, the phases can be obtained from the fringe oscillations at the correlator output, but in most large instruments the phase oscillations are removed by adjusting the local oscillator phase. Two signal samplers are then used for each antenna, one of which is fed through a $\pi/2$ phase shifter, thus providing the signal amplitude as a complex quantity. Measurement of both components of the signal amplitude also provides a signal to noise ratio that is greater than that from the amplitude alone by a factor of $\sqrt{2}$.

[2] Here, for simplicity, u and v have the same units as x and y, whereas they are often measured in wavelengths

[3] Note that autocorrelation produces a function that is hermitian in u and v with a maximum at the origin, as required for the visibility function. Self-convolution does not, in general, have these properties.

an array of $(2n_x - 1)(2n_y - 1)$ numbers which represent the measured values of $\mathcal{V}(u, v)$ at intervals of Δx and Δy, including both positive and negative ranges of u and v. Thus the visibility array has dimensions approximately twice those of the antenna array.

Self-convolution differs from autocorrelation in that the sign of the values in the second function f in (2) is reversed in both the x and y coordinates. This sign reversal effectively reverses the order of the signal values in the $f(x, y)$ data array, and it is intrinsic to the convolution process. Self-convolution has the desirable property that it can be evaluated by taking the Fourier transform of $f(x, y)$, multiplying each resulting value by its complex conjugate and transforming back to the (u, v) domain: this follows from the convolution theorem of Fourier transforms. For N data points in the form of a rectangular array, as considered above, the fast Fourier transform (FFT) algorithm can be used with which the number of calculations is proportional to $N\log_2 N$, rather than to N^2 as when the products of the signal amplitudes for all possible pairings of antennas are calculated individually. Calculation by FFT can similarly be adapted to evaluation of the autocorrelation function by using the process for self-convolution with the addition of a reversal of the signs of the variables to cancel the intrinsic sign reversal in the convolution process. The basic relationship can be found in Bracewell (2000, p. 46). For any two functions g and h, $g * h = g(-) \star h$, where $(-)$ indicates that the sign of the variable in function g has been reversed (i.e. the order of the data values in the array has been reversed). If the sign of the variable of g is reversed once more the expression becomes $g(-) * h = g \star h$, which indicates that cross-correlation can be obtained by convolution with an appropriate sign reversal. In the case of autocorrelation, as required above, this becomes $f(-) * f^* = f \star f^*$. Thus the procedure is, for each sample time, to form an array of the complex amplitude values measured at the antennas, $f(x, y)$, and then derive both the reversed order version, $f(-)$, and the complex conjugate version, f^*. Next, to obtain the convolution of $f(-)$ and f^*, take their FFTs[4] and form a third array in which the terms are the pair products of corresponding terms from the two FFTs. Finally, take the inverse FFT of this third array. The result will be an array of visibility values for each baseline vector in the antenna array, weighted in proportion to the number of antenna pairs for which the baseline vector occurs. The data in the pair-product array, prior to the inverse FFT step, have the dimensions of brightness on the sky and may provide the required sky mapping. However, the corresponding visibility values have weights that are proportional to the number of antenna pairs that contribute to each of them. This weighting maximizes the sensitivity to a point source, but the decreasing weight with increasing baseline length degrades the angular resolution. The final inverse FFT step allows for adjustment of the visibility weighting, as necessary. Note that the autocorrelation of the component of system noise in the signal amplitude data results only in random noise variations except in the case of the zero baseline point ($u = v = 0$) where the noise is fully correlated. This visibility point should not be included in the derivation of an image unless correction for the correlated noise component can be made.

Reduction of the number of computing steps is clearly advantageous with regard to reduction of the computing load. However, as a system for use with arrays in general the FFT scheme would have significant limitations. For any baseline in the (u, v) plane, the responses of all antenna pairs that measure the same (u, v) component are combined as a sum in the FFT process. Thus the responses of individual antenna pairs are not available, as they are when using a conventional correlator system. So with the FFT system it is not practicable to calibrate individual antennas[5]. Analysis involving phase closure techniques, such as self calibration, which are important in maximizing dynamic range, would not be possible. With the FFT system compensation for a phase or gain factor in the response of an individual antenna could only be implemented by applying it to each Nyquist sample at the corresponding receiver output before taking

[4]It will be necessary to add zero padding to approximately double the array dimensions to accommodate the visibility data. For discrete input sequences the FT repeats in a circular manner, see e.g. Bracewell (2000, p. 283), so it is necessary to provide room so that data are not overwritten.

[5]The responses to the correlated signal components of individual antenna pairs could, in principle, be separated out by phase switching using different Walsh functions for each antenna. However, the uncorrelated system noise would not respond to this process and the noise level would be that of the combined antenna pairs, which would generally be too high for this idea to be useful.

the FFT. In contrast, with conventional signal processing, the products of individual antenna pairs are not immediately combined with those of other pairs, and adjustments for individual antennas can be applied after averaging over many ($\sim 10^5$ or more) Nyquist samples. However, Tegmark and Zaldarriaga propose the FFT system for use with arrays of dipoles or similar small elements with wide fields of view, for observations of background radiation and the red-shifted HI line. At these longer wavelengths it may be possible to make the individual antennas with sufficiently uniform characteristics and accurate positioning that individual adjustments would not be necessary.

In a second paper Tegmark and Zaldarriaga (2009b) show how the principle of forming cross correlations of the received signals by using the FFT algorithm can be extended to arrays in the form of blocks of antennas within larger configurations. Cross correlations for such a hierarchy of, for example, three levels, can be formed by using a six-dimensional FFT. This approach would allow the situation in which the cross-correlation computing load is proportional to $N\log_2 N$ to be extended to less compact configurations, thus improving the angular resolution. The overall linear dimensions of such an array are likely to be restricted by the difficulty of providing a large enough plane area. Note, however, that if a low frequency array is extended to baselines comparable to the dimensions of the isoplanatic structure of the ionosphere, compensation for the induced phase errors requires phase adjustment of the signals from some of the antennas. For observations at ~ 74 MHz and baselines of a few km, such techniques become necessary: see, e.g., Kassim, et al. (1993).

References

Bracewell, R. N. (2000), *The Fourier Transform and its Applications*, 3rd ed., McGraw-Hill, N.Y.

Kassim, N. E., R. A. Perley, W. C. Erickson, and K. S. Dwarakanath, Subarcminute Resolution Imaging of Radio Sources at 74 MHz with the Very Large Array, Astron. J., 106, 2218-2228, 1993.

Tegmark, M. and M. Zaldarriaga, The Fast Fourier Transform Telescope, arXiv:0805.4414v2 [astro-ph] 24 Apr 2009a. Also published in Phys. Rev. D, 79, 08530, 2009.

Tegmark, M. and M. Zaldarriaga, Omniscopes: Large Area Telescope Arrays with only N log N Computational Cost, arXiv:0909.0001v1 [astro-ph.CO] 31 Aug 2009b.

Speed of Surveying. (Feb. 26, 2008, May 20, 2009)

(Relates to Chapters 5 and 6)

The requirement for survey speed for the SKA and the use of multibeam antennas is discussed under Trends Since 2001 on p. 2. In the frequency range below ~ 2 GHz four of the SKA's key science applications that require imaging of a significant fraction of the sky are as follows.

- Searching for pulsars in binary combinations with neutron stars or black holes, for gravitational studies.
- Measurement of Faraday rotation in very large numbers of radio galaxies to determine the structure of galactic and intergalactic magnetic fields.
- Mapping of very large numbers of HI galaxies out to red shifts of $z \simeq 1.5$ to study galactic evolution and provide further constraints on the nature of dark energy.
- Detection of transient events such as afterglows of gamma-ray bursts.

Choice of parameters for optimization of speed in survey observations is not the same as for optimization of sensitivity in targeted studies (Bregman 2005). Consider first the case of targeted observations of individual

continuum sources which have angular dimensions small compared with a station[1] beam. We can adapt the expression for the rms noise [Eq. (6.62) in TMS2] as a measure of the minimum detectable flux density S_{min} observed with two stations in a time τ:

$$S_{min} = \frac{2kT_s}{A(\Delta\nu\tau)^{1/2}}, \tag{1}$$

where A is the collecting area of a station. For an array with n_s stations there are $n_s(n_s-1)/2 \approx n_s^2/2$ correlated pairs of signals, so the right-hand side of Eq. (1) is multiplied by a factor $\sqrt{2}/n_s$. The observing speed, i.e. the number of observations per unit time is

$$1/\tau = \frac{A^2\Delta\nu S_{min}^2 n_s^2}{8k^2T_s^2}. \tag{2}$$

Next consider the case of a survey in which we are concerned with the of speed of coverage of a specified solid angle of sky down to a sensitivity level S_{min}. Since A is the is area of a station, the solid angle of a station beam is λ^2/A ster, where λ is the wavelength. If each station forms n_{sb} simultaneous beams, the instantaneous field of view is

$$F_v = \lambda^2 n_{sb}/A. \tag{3}$$

The reciprocal of the time τ required to cover solid angle F_v down to flux density level S_{min} is given by Eq. (2), and the corresponding survey speed is

$$F_v/\tau = \frac{\lambda^2 A\Delta\nu S_{min}^2 n_s^2 n_{sb}}{8k^2T_s^2} \tag{4}$$

steradians per unit time. For surveys to detect spectral line features, $\Delta\nu$ in Eqs. (2) and (4) represents the bandwidth of the line. Then, if it is necessary to search in frequency, the bandwidth of the receiving system can be included as an additional factor in the expression for the speed in speed in Eq. (4).

Comparison of Eqs. (2) and (4) shows the effect of the field-of-view dependence in the survey case. The survey speed is proportional to the number of simultaneous beams and is less strongly dependent on the station aperture A. The wavelength-squared factor results from the increased beamwidth with decreasing frequency, but the benefit of lower frequency (increasing λ) on the survey speed applies only so long as the effect of the galactic background radiation on the system temperature is small. From the galactic background model of Dulk et al. (2004), the brightness temperature in the range 10-1000 MHz is approximately proportional to $\nu^{-2.5}$, so for frequencies at which it is the dominant contributor to T_s, the frequency dependence of the survey speed is approximately proportional to ν^3. For directions which are not close to the galactic plane, the background temperature is about 20 K at 500 MHz and 2 K at 1 GHz, so if the receiver contribution to T_s is ~ 20 K there is a broad maximum in survey speed between these two frequencies.

Note that the discussion above involves the assumption that the sensitivity is limited only by the system noise. If dynamic range is the limiting factor then the density of the (u, v) coverage, which improves with increasing n_{st} and τ, may become the most important consideration. In either case performance improves with increasing number of Stations.

Survey speed can be increased by increasing the number of stations as well as the number of station beams. However, the size of the correlator system for the full array is proportional to n_s^2 and to n_{sb}, so increasing the number of stations, or the number of station beams, requires an increase in the size of the correlator. Increasing the station aperture A is likely to require adding more antennas to the station subarray

[1]For an array with a large total collecting area it is generally more practical to use a large number of small antennas rather than a smaller number of large antennas. To limit the number of cross correlated signal pairs, the small antennas are located in groups. These groups are commonly referred to as *stations* and the signals from the antennas at a station are combined to provide a number of beams within the main beam of the individual antennas. For pairs of stations, the signals from corresponding beams are cross correlated to provide the visibility data.

and thus increases the station beamforming hardware. The only way of increasing the observing speed that does not increase the signal processing requirements is reducing the system temperature T_s. However, the complexity of phased-array feeds for forming multiple beams from a single paraboloid antenna can degrade the system temperature. If cryogenic cooling in necessary, the required cooling capacity is considerably greater for multiple beam feeds than for single ones. Thus, optimization of the array performance for a given overall cost requires a broad consideration of the performance of many parts of the receiving system.

References

Bregman, J. D., System Optimization of Multi-Beam Aperture Synthesis Arrays For Survey Performance, Exp. Astron., 17, 365-380, 2005.

Dulk, G. A., W. C. Erickson, R. Manning, and J.-L. Bougeret, Calibration of Low-Frequency Radio Telescopes Using Galactic Background Radiation, Astron. Astrophys., 365, 294-300, 2001.

Combination of Single-sideband and Double-sideband Receiving Systems

(May 11, 2006)

(Relates to Ch. 6, pp. 195-199)

Table 6.1 does not include the rather unusual case of cross correlation of the signals from two antennas, one of which is equipped with a single sideband (SSB) receiving system and the other with a double sideband (DSB) system. However, the discussion of such a hybrid instrument has arisen in connection with systems of both types which are in use in neighboring observatories on Mauna Kea. The local oscillator frequencies and IF responses would be the same for both systems. In the DSB system, only the sideband that matches the SSB system contributes to the measured correlation, and the effect of other one is to double the noise power in the DSB receiver. The hybrid system is thus equivalent to a single sideband system in which one antenna has its antenna temperature multiplied by a factor 2α (for a definition of α see Eq. (6.65) of TMS2). The relative sensitivity (signal-to-noise ratio) as considered in Table 6.1 is therefore $1/\sqrt{2\alpha}$.

The value of α for the case where the characteristics of the mixer elements of the single- and double-sideband systems are identical can be considered as follows. Consider the simple double-sideband mixer in Fig. 1a. The temperature of the input load is T_1 and the mixer output power is $2\beta T_1$, where the 2 results from the two input sidebands and β is a conversion loss factor. The mixer also adds a noise power level T_M. (This can be considered as including noise introduced by an IF amplifier which is sometimes mounted in a commom structure with the mixer). If we measure the Y-factor by sequentially replacing T_1 by hot (T_H) and cold (T_C) loads,

$$Y = \frac{T_H + T_{R(DSB)}}{T_C + T_{R(DSB)}} = \frac{2\beta T_H + T_M}{2\beta T_C + T_M}, \tag{1}$$

where $T_{R(DSB)}$ is the double-sideband noise temperature of the mixer. From Eq. (1), $T_M = 2\beta T_R$. Figure (1b) represents the signal path for one sideband of a single-sideband (or sideband-separating) mixer of the type shown in Fig. A7.1 of TMS2. We assume that the two mixers used are identical to the one in Fig. 1a. The input passes through two power dividers and one power adder before the single sideband output is obtained. The levels of signal and mixer noise at the various stages are indicated in the figure. Just one sideband of the input T_1 is considered. Note that in the adder the two components of mixer noise are incoherent and add in power, whereas the signal components T_1 are coherent and their voltages add. The output for one sideband is $\beta T_1 + T_M = \beta(T_1 + 2T_{R(DSB)})$. Again we replace the input T_1 with hot and cold loads and measure the Y-factor at the output for one sideband. Thus,

$$Y = \frac{T_H + T_{R(SSB)}}{T_C + T_{R(SSB)}} = \frac{\beta(T_H + 2T_{R(DSB)})}{\beta(T_C + 2T_{R(DSB)})}, \tag{2}$$

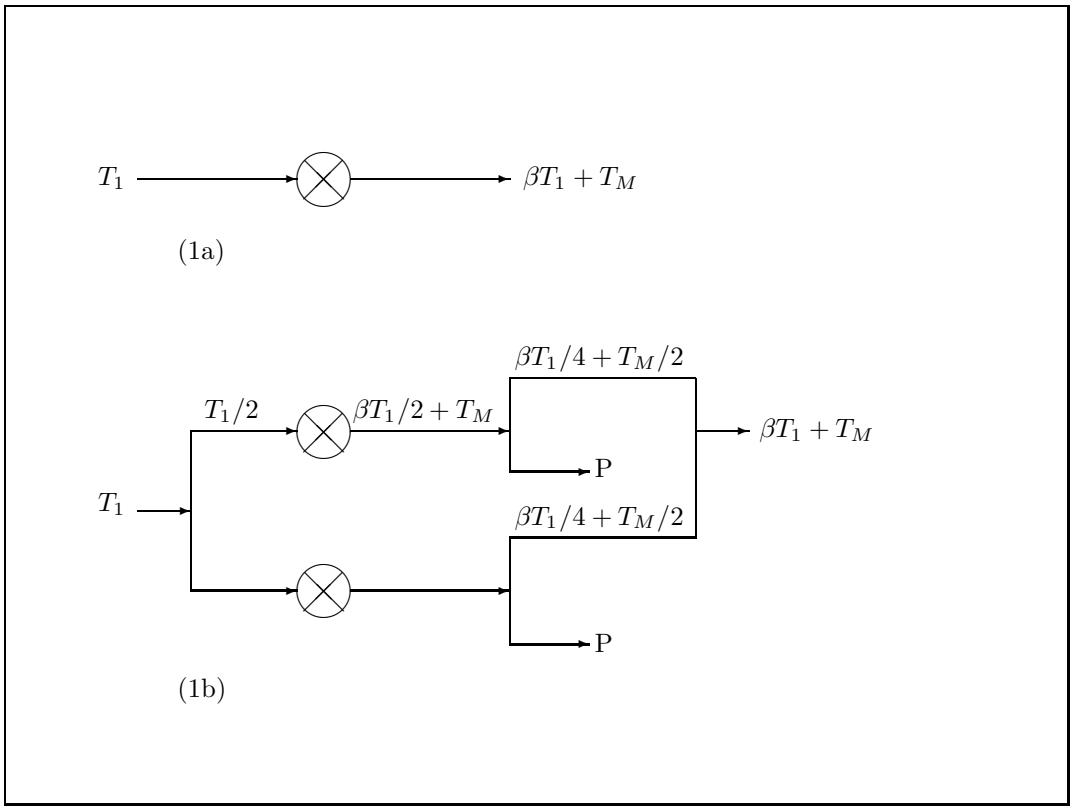

Figure 1: (a) Double-sideband mixer. (b) Single-sideband mixer. The expressions are proportional to the sum of the powers of the signal and the mixer noise. Outputs at points P provide the other sideband, with the use of a voltage-differencing network (for further details see Fig. A7.1 of TMS2).

from which $T_{R(SSB)} = 2T_{R(DSB)}$, i.e. $\alpha = 1/2$. This result may be useful in considering the comparative sensitivity of interferometers that use a mixer as the input stage, so long as the noise performance of the mixer elements (including any component of local-oscillator noise contributing to T_M) can be assumed to be similar in the different receivers, and losses in power dividers, etc. are small.

Tolerances on the Bandpass Frequency Response (Sept. 22, 2006)

(Relates to Ch. 7, pp. 235-238)

Table 7.2 gives tolerances on the deviation of the bandpass response of the receiving channels from an ideal rectangular form. L. R. D'Addario has calculated the tolerances for 0.2% maximum gain error which is more stringent than the 1% considered in Table 7.2. A detailed discussion including recommendations for the ALMA array are given in:

D'Addario, L. R., Passband Shape Deviation Limits, *ALMA Memo.* **452**, 2003.

ALMA Memos can be accessed at http://www.alma.nrao.edu/memos/index.html

See also the note that follows on sinusoidal ripple.

Sinusoidal Ripple in the Bandpass Response (June 26, 2007)

(Relates to Section 7.3)

Sinusoidal ripples in the amplitude response of a signal channel can result from two different effects. First we can consider variations in the gain of an amplifier that can be modeled as a periodic ripple in the amplitude frequency response, with no effect on the phase. Second, sinusoidal ripples in frequency can result from a reflection in a transmission line, in which case the amplitude variations are accompanied by periodic variations in the phase.

Consider first the case of pure amplitude variations. The frequency response of a signal channel with such variations can be represented by

$$H(\nu) = 1 + \gamma \cos[2\pi N(\nu - \nu_0)/\Delta\nu + \theta], \tag{1}$$

where γ is the (real) ripple coefficient, ν_0 is the center frequency and N is the number of ripples across the bandwidth $\Delta\nu$. The integral in the denominator of the expression for the relative sensitivity, Eq. (7.36) of TMS2 (Eq. (7.25) of TMS1), is[1]

$$\begin{aligned}\int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} H_m H_n d\nu &= \int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} \big[1 + \gamma_m \cos[2\pi N(\nu - \nu_0)/\Delta\nu + \theta]\big]\big[1 + \gamma_n \cos[2\pi N(\nu - \nu_0)/\Delta\nu + \theta]\big] d\nu \\ &= \Delta\nu(1 + \gamma_m\gamma_n/2), \end{aligned} \tag{2}$$

where for simplicity we have assumed that the phase, θ, of the ripples is the same in each channel and that N is an integer. In the denominator of Eq. (7.36) we have

$$\begin{aligned}\int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} |H_m|^2|H_n|^2 d\nu &= \int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} \big[1 + \gamma_m \cos[2\pi N(\nu - \nu_0)/\Delta\nu + \theta]\big]^2\big[1 + \gamma_n \cos[2\pi N(\nu - \nu_0)/\Delta\nu + \theta]\big]^2 d\nu \\ &= \Delta\nu\big[1 + \tfrac{1}{2}\gamma_m^2 + \tfrac{1}{2}\gamma_n^2 + 2\gamma_m\gamma_n + \tfrac{3}{8}\gamma_m^2\gamma_n^2\big]. \end{aligned} \tag{3}$$

From Eq. (7.37), the degradation factor is

$$\mathcal{D} = \frac{(1 + \gamma_m\gamma_n/2)}{\sqrt{1 + \tfrac{1}{2}\gamma_m^2 + \tfrac{1}{2}\gamma_n^2 + 2\gamma_m\gamma_n + \tfrac{3}{8}\gamma_m^2\gamma_n^2}}. \tag{4}$$

As an example of this result, for values of $\gamma_m = \gamma_n = 1/3$ (ripples with 6 dB power ratio between maxima and minima) the degradation factor is 0.913, but if the sign of one of the gammas is reversed, so that the ripples in the two channels are in antiphase, the degradation factor is 0.997. In the latter case the ripples largely cancel in the products of the signal samples at the correlator output, so the degradation is very small.

Next consider the effect on the frequency response of a reflection in a transmission line resulting in standing waves. The reflection causes a sinusoidal ripple in amplitude and a similar periodic effect in the phase. To represent the signal we use $\exp(-j2\pi\nu t)$ and the reflection adds a component with relative amplitude[2] γ and a time delay τ. Thus for the frequency response of one signal we have

$$H(\nu) = e^{-j2\pi\nu t} + \gamma e^{-j(2\pi\nu t + \tau)}, \tag{5}$$

and

$$|H|^2 = H(\nu)H^*(\nu) = 1 + 2\gamma\cos(2\pi\nu\tau) + \gamma^2. \tag{6}$$

[1] Symbolic evaluation of the integrals in this note was obtained using Mathcad.

[2] To be fully general γ should be complex with argument representing the phase change on reflection (which can vary depending upon the load and line impedances). Here, to obtain an approximate estimate of the magnitude of the ripple effect, γ is taken to be real.

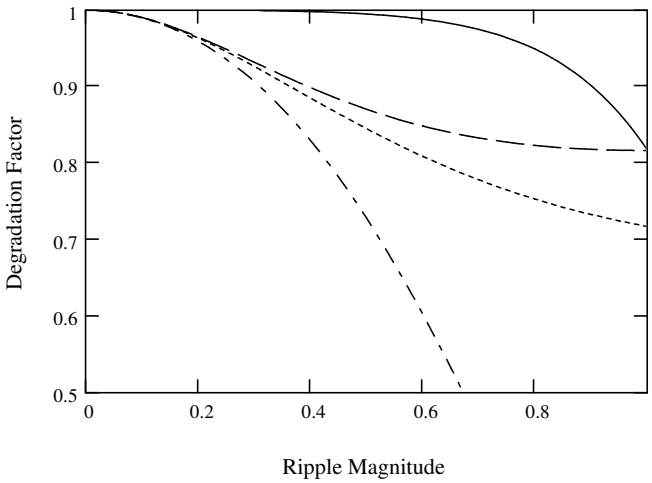

Figure 1: The ordinate is the degradation factor and the abscissa is the magnitudes of the ripple coefficients γ_m and γ_n, which are equal for each curve ($\gamma_m = \pm\gamma_n$). The curves for the pure amplitude ripples are the third down (short dashes) for in-phase ripples and the top curve (solid) for antiphase. The curves for the ripples resulting from reflection are the second down (long dashes) for in-phase and the bottom curve (long and short dashes) for antiphase.

In evaluating the integrals in Eq. (7.36) we use the same value of τ for each channel, and take $\Delta\nu\,\tau$ to be an integer which greatly simplifies the expressions:

$$\begin{aligned}\int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} H_m H_n^* d\nu &= \int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} [1+\gamma_m e^{-j2\pi\nu/\Delta\nu} + \gamma_n e^{j2\pi\nu/\Delta\nu} + \gamma_m\gamma_n] d\nu \\ &= \Delta\nu[1+\gamma_m\gamma_n]. \end{aligned} \tag{7}$$

For the integral in the denominator of Eq. (7.36) we have,

$$\begin{aligned}\int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} |H_m|^2|H_n|^2 d\nu &= \int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2} [1+2\gamma_m\cos(2\pi\nu/\Delta\nu)+\gamma_m^2][1+2\gamma_n\cos(2\pi\nu/\Delta\nu)+\gamma_n^2] d\nu \\ &= \Delta\nu[1+\gamma_m^2+\gamma_n^2+2\gamma_m\gamma_n+\gamma_m^2\gamma_n^2]. \end{aligned} \tag{8}$$

Thus from Eq. (7.36),

$$\mathcal{D} = \frac{1+\gamma_m\gamma_n}{\sqrt{1+\gamma_m^2+\gamma_n^2+2\gamma_m\gamma_n+\gamma_m^2\gamma_n^2}}, \tag{9}$$

which is the same as the expression for sinusoidal ripple in Table 7.1 of TMS (except that here we have considered only real values of γ_m and γ_n). Thus the results for ripple in Table 7.1 (and in Thompson and D'Addario, 1982) correspond to the case of ripple resulting from reflection in a transmission line. In the reflection case ripples of 6 dB magnitude result in degradation factors of 0.921 and 0.883 for ripples in phase and antiphase respectively. The decorrelating effect of the phase variation introduced by the ripples is greatest in the antiphase case, and overrides the reduction of the amplitude effect seen in the pure amplitude case.

Figure 1 shows curves of the degradation factor for both cases with ripples in phase and antiphase. For identical ripples in the two responses the phase effects cancel, and the curves for the pure amplitude case and

for reflection are very similar, as shown by the two center curves in Fig. 1. However the amplitude responses are slightly different in the two cases so these curves diverge as γ increases. Under certain conditions it is possible to induce sinusoidal amplitude variations in the frequency spectrum of a signal without affecting the phase but, especially in an analog system, it is likely that some phase variation will also occur. The curves in Fig. 1 provide an indication of the range of the sensitivity degradation.

Thanks are due to D. T. Emerson and L. R. D'Addario for much helpful discussion on this subject.

Reference

Thompson, A. R. and L. R. D'Addario, Frequency Response of a Synthesis Array: Performance Limitations and Design Tolerances, Radio Science, 17, 357-369, 1982.

Quantization Noise and Bandpass Tolerances (May 11, 2006)

(Relates to Sections 7.3 and 8.3)

In receiving systems in which the analog signal voltage is sampled and quantized to allow further processing in digital form, the difference between the analog samples and their digital representation gives rise to a component of random quantization noise. The power spectrum of the quantization noise is close to being uniformly level across the receiver passband, even for large variation in the shape of the input spectrum. Thus in cases where the gain of the analog system varies across the passband, the addition of the quantization noise causes a variation in signal-to-noise ratio (SNR). This effect is most serious where the number of quantization levels is small and the quantization noise correspondingly large. A detailed discussion with examples of numerical calculations can be found in:

Thompson, A. R. and D. T. Emerson, Quantization Noise, EVLA Memo. 88, NRAO, Jan. 21, 2005.

EVLA Memos can be accessed at http://www.aoc.nrao.edu/evla/memolist.shtml

Delay and Phase Errors, and Fringe Rotation, in Single- and Double-Sideband Interferometer Systems (Sept. 22, 2006; Apr. 10, 2007; Aug. 17, 2007: Nov. 20, 2007)

(Relates to Ch. 7)

Delay setting tolerances are discussed on pp. 238-9 of TMS2, but it is useful to review the subject again to understand better the effect of double-sideband systems and also take account of approaches being used in new instruments. These include EVLA (the Extended Very Large Array) for which the receiving and data processing systems of the VLA are being completely replaced, and ALMA (the Atacama Large Millimeter Array (Wooten 2003)). Digital delay systems have proved to be much more accurate and economical than analog ones and the principal consideration is the mitigation of phase errors resulting from the discretely adjustable nature of the delays.

For an incoming wavefront from a source the path lengths to different antennas of an array are generally unequal. The relative time differences in the wavefront arrival at the antennas are referred to as the geometric delays, τ_g. To compensate for the different geometric delays, the signal received at each antenna is subjected to an instrumental delay τ_i that is continuously adjusted so that $\tau_g + \tau_i$ is the same for all antennas. Thus the signals at the correlator inputs are aligned in time with respect to a common wavefront incident from the phase reference position. The fringes at the correlator output result from the fact that the signals traverse the geometric and the instrumental delays at different frequencies, and that the phase shifts resulting from the delays vary continuously as the delays vary. In an ideal situation the instrumental delays would be continuously adjusted, and if there were no frequency changes within the receivers, no fringe oscillations would occur. In practice the situation is rather more complicated. The instrumental delays are inserted after the signals have been digitized and the sample interval τ_s provides a convenient coarse adjustment. For Nyquist sampling, $\tau_s = 1/2\Delta\nu$ where $\Delta\nu$ is the signal bandwidth.

Delay Errors

In an array with a number of antennas, the delays are adjusted relative to the delay for a designated reference antenna. We shall consider the reference antenna to be the one that is the last one encountered by the approaching wavefront and its instrumental delay remains fixed. The delay error for any antenna is the difference between the sum of the geometric and instrumental delays for that antenna and for the reference antenna. When the delay error becomes as large as $\pm\tau_s/2$, the delay is adjusted by an increment $\pm\tau_s$. Thus the delay error for a single antenna is uniformly distributed over $\pm\tau_s/2$. Coarse delay adjustments in units of the digital sampling interval are implemented in a FIFO (first-in-first-out) memory, possibly using RAM (random access memory). These provide the major part of the instrumental delay but the residual delay errors are large enough that they can cause serious loss of sensitivity if not mitigated. In the original VLA system, for example, finer steps are provided by an adjustment in the timing of the sampler action in steps of $\tau_0 = \tau_s/16$. The spacing of adjacent samples remains τ_s except when a delay adjustment is made and the sample occurs earlier or later by τ_0. When the delay error becomes equal to $\tau_0/2$, the instrumental delay is adjusted by τ_0 as represented by the staircase function in Fig. 1. One can see from Fig. 1 that the probability distribution of the delay error is uniform within a range $\pm\tau_0/2$. For a pair of antennas it can usually be assumed that the times of delay adjustment are unrelated (in general the rates of change of the geometric delay will be different for each antenna), so the probability distribution of their combined delay errors is a triangular function with extreme values of $\pm\tau_0$, as in Fig. 2. The rms value of this delay error is:

$$\left[\frac{\int_0^{\tau_0} p(\Delta\tau)\Delta\tau^2 d\Delta\tau}{\int_0^{\tau_0} p(\Delta\tau) d\Delta\tau}\right]^{\frac{1}{2}} = \frac{\tau_0}{\sqrt{6}}\,, \tag{1}$$

where $p(\Delta\tau)$ is the expression for the probability distribution of $\Delta\tau$ in Fig. 1.

Phase Errors and Degradation of Sensitivity

A delay error $\Delta\tau$ results in a phase error in a signal equal to $2\pi\Delta\tau\nu$ where, for systems using analog delays,

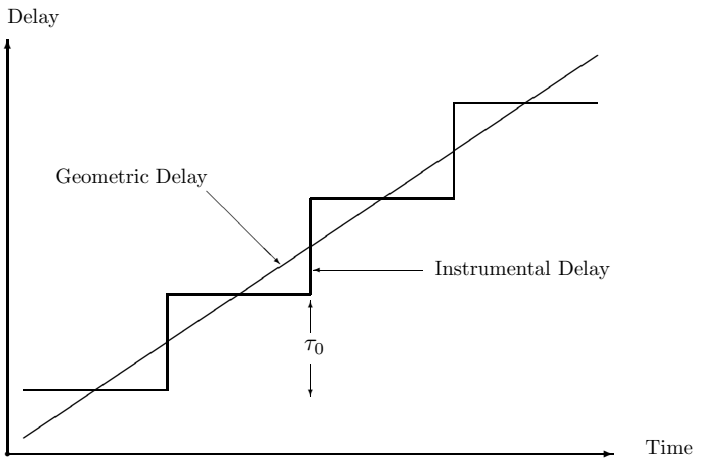

Figure 1: Adjustment of instrumental delay in steps τ_0 to compensate for geometric delay. The vertical sections of the staircase function indicate change of instrumental delay and the horizontal sections are the time intervals during which the signal is sampled. Over small time intervals the geometrical delay can be represented as a linear function of time. Both axes have the dimensions of time but for, e.g. a baseline of 1 km E-W, the maximum slope of the line representing space delay is 0.24 ns per second and the time scales of the two axes differ by a factor of order 10^{10}.

ν represents frequency in the IF band in which the delays are inserted. For systems in which the signal bandpass is Nyquist sampled, which is the most usual case, ν is a baseband frequency[1] in the range 0 to $\Delta\nu$, as explained in Appendix 1. With a spectral correlator, the highest frequency channel within such a band has a center frequency that is approximately equal to the high frequency edge, $\Delta\nu$. For frequencies in this top channel, the maximum delay error τ_0 for an antenna pair results in a phase error of $2\pi\Delta\nu\tau_0 = (\tau_0/\tau_s)\pi$. Thus the probability distribution of this phase error is a triangular function as in Fig. 2 with extreme values $\pm(\tau_0/\tau_s)\pi$. As shown for the delay error in Eq. (1), the rms phase error is the maximum value divided by $\sqrt{6}$.

To determine the effect of delay errors on sensitivity note that for frequency ν a delay error $\Delta\tau$ results in a phase error $2\pi\nu\Delta\tau$. Let α be the size of the fine step as a fraction of the coarse step τ_s. The sensitivity (i.e. the relative response) is determined by averaging the cosine of the phase error, weighted by the triangular distribution of delay error in Fig. 2,

$$\frac{2}{\alpha\tau_s}\int_0^{\alpha\tau_s}\left(1-\frac{\Delta\tau_s}{\alpha\tau_s}\right)\cos(2\pi\nu\Delta\tau)d\Delta\tau = \left[\frac{\sin(\pi\nu\alpha\tau_s)}{\pi\nu\alpha\tau_s}\right]^2. \tag{2}$$

This is the sensitivity for a very small bandwidth centered on frequency ν. The sensitivity for a band extending from $N\Delta\nu$ to $(N+1)\Delta\nu$, where N is any integer (including zero), is obtained by averaging over the baseband response,

$$\frac{1}{\Delta\nu}\int_0^{\Delta\nu\alpha\tau_s}\left[\frac{\sin(\pi\nu\alpha\tau_s)}{\pi\nu\alpha\tau_s}\right]^2 d\nu = \frac{2}{\alpha}\int_0^{\frac{\alpha}{2}}\left[\frac{\sin(\pi x)}{\pi x}\right]^2 dx. \tag{3}$$

Here we use $\Delta\nu\tau_s = \frac{1}{2}$, and have put $\nu\alpha\tau_s = x$ for convenience in numerical evaluation of the integral[2].

[1] This point was not properly taken into account in versions of this note earlier than Aug. 17 '07.

[2] In the case where the phase errors are small, it may be convenient to use $\langle\cos(\phi)\rangle \approx (1-\langle\phi^2\rangle/2)$, where $\phi = 2\pi\nu\alpha\Delta\tau$ and $\langle\,\rangle$ indicates the mean value. Then noting that $\Delta\tau$ and ν vary independently, $\langle\phi^2\rangle = (2\pi)^2\langle\Delta\tau^2\rangle\langle\nu^2\rangle$. From Eq. (1) $\langle\Delta\tau^2\rangle = \tau_0^2/6$ and for a baseband response $\langle\nu^2\rangle = \Delta\nu^2/3$.

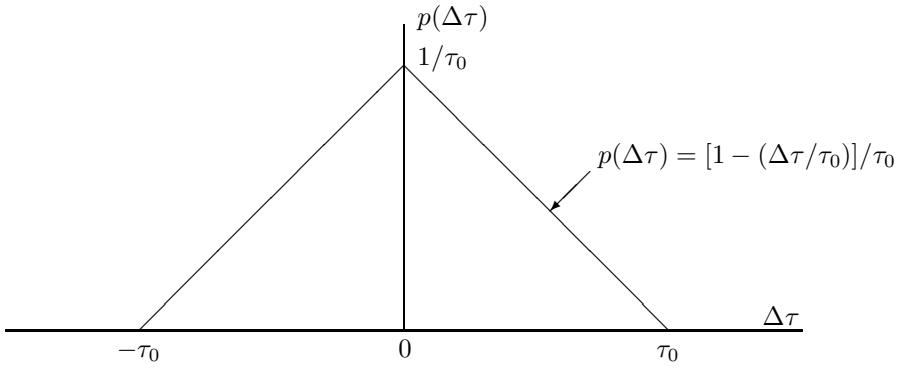

Figure 2: Probability distribution $p(\Delta\tau)$ of the delay error $\Delta\tau$ for a pair of antennas. τ_0 is the minimum increment of the instrumental compensating delay. The expression shown for $p(\Delta\tau)$ applies to the part of the probability function for which $\Delta\tau \geq 0$.

For the case where we use only the coarse delay steps ($\alpha = 1$) expression (3) is equal to 0.774 so, as noted earlier, the performance with coarse delay steps without further mitigation is not acceptable. Some values of rms phase error and sensitivity loss averaged across the bandwidth are given in Table 1.

$\alpha = \tau_0/\tau_s$	ϕ_{rms}	SNR loss
1/4	10.6°	1.7%
1/8	5.30°	0.43%
1/16	2.65°	0.11%
1/32	1.33°	0.027%

Table 1. Values of the loss in signal-to-noise ratio (sensitivity) for the full baseband response from 0 to $\Delta\nu$, as used in continuum observations.

In Table 1, sensitivity loss for $\alpha = 1/4$ is approaching an acceptable level. However, the maximum phase errors are $\sqrt{6} = 2.45$ times the rms value, i.e. $2.45 \times 10.6° = 26°$ in this case. Depending upon how fast the delay error is changing with time, the maximum error will be decreased somewhat in the data averaging after cross correlation. In the (u, v) plane, the rate of change of delay error goes through zero as the u-component of the baseline crosses the v axis. Thus averaged data in which the phase error is close to the maximum are to be expected, especially for short-baseline configurations. Hence in considering the acceptable delay errors, phase errors should be considered as well as sensitivity loss. The original VLA system used $\tau_0/\tau_s = 1/16$ and a baseband IF response.

Other Methods of Mitigation of Delay Errors in Continuum Systems

Conceptually the most straightforward way of keeping the loss in sensitivity resulting from delay errors within a tolerable limit[3] (say, $\sim 1\%$) is to use a small enough value for the minimum delay increment. This may not always be easy in systems with wide bandwidths which require correspondingly high sample rates in the digitization. A possible scheme to reduce phase errors (D'Addario 2003) is one in which whenever a delay

[3]Various effects in an interferometer system limit the sensitivity. There are some large effects like aperture efficiency and quantization efficiency, and more numerous smaller ones such as phase irregularities in frequency responses, local oscillator noise, timing errors, delay errors, etc. The combined effect of the smaller losses can become serious, so for each one it is reasonable to aim at a fairly stringent limit like the 1% figure suggested here.

increment is inserted or removed, a phase jump of magnitude $2\pi\nu_0\tau_0$, and opposite sign to the delay-induced phase jump, is inserted in the corresponding signal through a local oscillator[4]. Here ν_0 is the IF center frequency, for which the phase error is exactly canceled. The overall effect for the full bandwidth can be found by determining the value of $\langle(\nu-\nu_0)^2\rangle$, that is, the mean squared value of frequency measured with respect to the band center:

$$\langle(\nu-\nu_0)^2\rangle = \frac{1}{\Delta\nu}\int_{\nu_0-\Delta\nu/2}^{\nu_0+\Delta\nu/2}(\nu-\nu_0)^2 d\nu = \frac{\Delta\nu^2}{12}. \quad (4)$$

This result applies to any IF band of width $\Delta\nu$. Since the phase changes resulting from the changes in the instrumental delay provide a component of the frequency offset used to stop the interferometer fringes, it is also necessary to replace this by inserting a smooth component in the form of a frequency offset, $2\pi\nu_0\, d\tau_g/dt$, where τ_g is the geometric delay. This could be combined with the fringe-rotation offset in a local oscillator[5]. The combination of the inserted phase jumps and the frequency offset provide a sawtooth phase component that, at the band center, exactly cancels the phase sawtooth induced by the delay error. If this method were used with no fine delay steps, i.e. $\tau_0 = \tau_s$, the loss in sensitivity would be $\sim 13\%$ which is still unacceptably large. However, if used in combination with fine delay steps the steps need not be as fine as would otherwise be necessary.

Multichannel (Spectral Line) Correlator Systems

In multichannel correlators the input band is divided into many channels, and the signals for corresponding channels are cross correlated. The numbers of channels is commonly 1024 or more, and within any channel the relative variation of frequency is very small. Thus at any instant the effect of a delay error $\Delta\tau$ is to introduce a phase error $2\pi\Delta\tau\nu_c$, where ν_c is the center frequency of the channel. Since the frequency is almost constant across the channel, the loss in signal amplitude that occurs in a wide (continuum) band, resulting from the frequency variation of the phase error, is avoided. The time variation of the delay error results in a varying phase error that can be corrected by inserting a phase correction at the correlator. Thus with a multichannel correlator it is possible to avoid the need for delay increments finer than τ_s, so long as the extra processing steps to correct the phase can be incorporated. For a single antenna, the maximum delay error is $\tau_s/2 = 1/(2\Delta\nu)$, and for the highest channel, centered very close to frequency $\Delta\nu$, the maximum phase error is π. The time for the delay error variation to compete one cycle is $\frac{d\tau}{dt}/\tau_s = \frac{d\tau}{dt}\Delta\nu/2$. (Note that $\frac{d\tau}{dt}$ is different for each antenna.) This is greater than the time for one fringe cycle by a factor of $2\times$(signal frequency at the antenna)/(signal frequency in the baseband $0-\Delta\nu$). At any instant the phase error for the signal from an antenna is equal to $2\pi \times$ (delay error) $\times$ (channel frequency in the baseband $0-\Delta\nu$). If the correction is applied at the correlator output, the corrections for both antennas of each pair must be included.

Carlson and Dewdney (2002), describe a multichannel correlator designed to handle wide bandwidth signals (the WIDAR system). The signals from the antennas are Nyquist sampled and then the band is divided by digital FIR filters into a number of channels, N_c. The the Nyquist sample rate appropriate for each channel is equal to the original sample rate divided by N_c, and the sample rates are decimated to this value at the filter outputs. The outputs of the filters then go to separate cross correlators. In this way the total bandwidth that can be processes is not limited by the capacity of a single correlator. A value of $N_c = 32$ would be sufficient to reduce the loss in sensitivity resulting from the delay errors to an acceptably small value. Adjusting the phases of the signals at the correlator inputs is used to remove the phase errors resulting from delay errors and also to provide fringe stopping. Phase adjustment at this point is possible because the samples are in complex form, having been through the filtering process. Since multichannel

[4]This method of mitigation of the delay errors was considered but not implemented during the early development of the VLA. The original idea is attributed to B. G. Clark.

[5]In the case of a DSB system, the fringe rotation must be applied to the first LO, but the frequency offset required by the phase-error reduction scheme must be applied to the second or a later LO so that, like the delay-induced phase errors, the offsets are applied with the same sign to each sideband component within the IF band.

correlators provide a means of removing channels that are contaminated by interference, they are widely used for continuum as well as spectral-line observations.

Double-Sideband Systems

The considerations up to this point have applied to single-sideband (SSB) systems. For double-sideband (DSB) systems some differences must be considered (Thompson and D'Addario 2000). For an SSB system the main effect of a phase error is to cause a rotation of the correlation vector as indicated in Fig. 6.5a (TMS2), resulting in an error in the correlator output phase, as considered above[6]. For a DSB system, the delay error causes the components of the correlation vector resulting from the two sidebands to rotate in opposite directions in the complex plane, as shown in Fig. 6.5b (TMS2), where the line AB represents the phase angle when the delay error is zero[7]. The amplitude of the vector sum of the two components is proportional to $\cos(2\pi\nu_0\Delta\tau)$ where ν_0 is the IF center frequency, but the phase of the correlation is not changed by a variation in the instrumental delay.

Consider a case in which the geometric delay is varying rapidly enough that the delay error changes sign several times during the minimum averaging time at the correlator output. For an SSB system (TMS2 Fig. 6.5a), the phase of the correlation vector swings back and forth, following the difference of the error patterns for the two antennas (the small arrows indicating variation of the vector phase in Fig. 6.5 reverse direction when the sign of the phase error changes). For a DSB system (TMS2 Fig. 6.5b) the phase angles of the vectors representing the two sideband responses move in opposite senses. In both the SSB and DSB cases components of the correlation that are normal to the vector time-average (in Fig. 6.5b the line AB) cancel, and the magnitude of the correlation is proportional to the time-average of the cosine of the phase measured with respect to the mean phase. Over an averaging period in which the SSB phase error changes sign, the loss in sensitivity is effectively the same for the SSB and DSB systems. Note, however, that in the SSB case the loss in sensitivity occurs in the averaging, whereas in the DSB case the loss occurs immediately in the correlation process. Thus in the SSB case there is a chance to correct for phase errors after cross correlation, but in the DSB case only if the sideband responses can be separated (see pp. 154-155 of TMS1, or pp. 181-3 of TMS2).

If we are considering delay errors that are quasi-constant, or only vary slowly with time, the tolerance on the errors is more stringent in the DSB case. Such errors were more important in early interferometers with analog delay systems using coaxial cable or ultrasonic elements (see, e.g. Coe 1973), which could be temperature sensitive and difficult to calibrate accurately. The statement (TMS1 p. 154, TMS2 p. 181) that for DSB systems increased accuracy of delay setting is required applies to such cases. In digital systems the delays are controlled by a highly accurate master clock and the only significant errors result from the incremental nature of the adjustment. With digital delays the effects are in most respects the same for SSB and DSB systems, except that for DSB systems, again, post-correlation corrections are possible only if the sidebands can be separated.

Fringe Rotation

The fringe oscillations in the data from the correlator must be removed before a radio image can be formed. This process is sometimes referred to as fringe stopping (i.e. stopping the motion of the fringes with respect to the astronomical sky), and sometimes as fringe rotation (i.e. causing the fringed to rotate with the sky). As described in Ch. 6 (p. 149 of TMS1 and pp. 173-4 of TMS2), this can be achieved by inserting a fringe-frequency offset on a local oscillator. For a multi-antenna array, the offset for each antenna is chosen to stop the fringes for that antenna when combined with a common reference antenna. It is also possible to

[6]There is also a relatively small decrease in the amplitude, which results from the variation of the phase error with frequency across the IF band, and is proportional to $\mathrm{sinc}(\Delta\nu\Delta\tau)$ where $\Delta\nu$ is the IF bandwidth. This results from the averaging of the varying phase over time. The same sinc function appears as part of Eq. (6.9) (TMS1 and TMS2) and is shown by the broken line in Fig. 6.4.

[7]To measure the phase of the cross correlation of both sidebands in combination it is necessary to periodically insert a $\pi/2$ phase shift into the IF signal of one antenna, or not to stop the fringes and fit a sine wave to the fringe function (TMS2, p. 178).

stop the fringes by inserting corrections in the phase of the signals at the correlator, and this procedure may be preferred as it reduces the complexity of the LO system. If the corrections are inserted before cross multiplication they can be applied to each of the N_a antennas of the array (see, e.g., Carlson and Dewdney 2002), whereas after cross multiplication the corrections must be applied to all of the $N_a^2/2$ antenna pairs. However, corrections inserted before cross multiplication must be applied to each signal sample that goes to the cross correlator, whereas corrections applied to the cross products can be performed after some time averaging of the products. The averaging must not be so long that the fringe oscillations are attenuated. The effect of time averaging on the fringes is to convolve the sinusoidal fringe function of frequency ν_f with a rectangular function of width equal to the averaging time τ_{av}. A 1% loss in sensitivity occurs for $\nu_f\tau_{av} = 0.078$ and 2% loss for $\nu_f\tau_{av} = 0.111$. As an approximate criterion, the averaging time should be no more than $\sim 1/10$ of a fringe period. For a double-sideband system the LO offset must be applied to the first LO (see pp. 180-1 of TMS2), and if the fringe stopping is applied at the correlator the sidebands must first be separated (see pp. 154-155 of TMS1, or pp. 181-3 of TMS2).

Appendix 1, Signal Frequency After Sampling

After Nyquist sampling of a signal band $n\Delta\nu$ to $(n+1)\Delta\nu$, where n is an integer, the frequency band of the sampled data is 0 to $\Delta\nu$, and does not depend on the frequency of the band at the sampler input[8]. To illustrate this situation, consider a Fourier component $A\sin(2\pi\nu t + \phi)$, with arbitrary amplitude and phase, within a band 0 to $\Delta\nu$. The band is sampled at the Nyquist rate, the sample times being $t = m/(2\Delta\nu)$ where $m = 0, 1, 2, ...$ The sampled values of the component are $A\sin(\phi)$, $A\sin(\frac{\pi\nu}{\Delta\nu} + \phi)$, $A\sin(\frac{2\pi\nu}{\Delta\nu} + \phi), ...,$ Now consider the case where the same input band has been converted to the range $\Delta\nu$ to $2\Delta\nu$. The frequency is higher by $\Delta\nu$, so the original component becomes:

$$A\sin[2\pi(\nu + \Delta\nu)t + \phi] = A\sin(2\pi\nu t + \phi)\cos(2\pi\Delta\nu t) + A\cos(2\pi\nu t + \phi)\sin(2\pi\Delta\nu t) \quad (5)$$

Again, sampling at times $m/(2\Delta\nu)$ we obtain for the component: $A\sin(\phi)$, $-A\sin(\frac{\pi\nu}{\Delta\nu} + \phi)$, $A\sin(\frac{2\pi\nu}{\Delta\nu} + \phi), ...$ The result is the same as before except that the sign is reversed for odd values of m. Further investigation shows that this sign reversal occurs when n has an odd value. Since the sign reversal occurs for both signals of a cross correlated pair, it has no effect on the product. Thus for any value of n the result at the correlator output is the same as for a baseband input to the sampler. Thus sampling of the band $n\Delta\nu$ to $(n+1)\Delta\nu$ has the effect of converting the band down by $n\Delta\nu$ and in this respect acts like a local oscillator.

References

D'Addario, L. R., Notes on Delay-Tracking for ALMA: Resolution and Tolerance. Unpublished memorandum, Feb. 2, 2003.

Carlson, B. R. and P. E. Dewdney, Efficient Wideband Digital Correlation, Electronics Letters, 36, 987-988, 2000.

Coe, J. R., NRAO Interferometer Electronics, Proc. IEEE, 61,1335-1339, 1973.

Thompson, A. R. and D'Addario, L. R., Relative Sensitivity of Double- and Single-Sideband Systems for Both Total-Power and Interferometry, ALMA Memo. 304, April 22, 2000.

Wooten, A., The Atacama Large Millimeter Array (ALMA), Proc. SPIE, 4837, 110-118, 2003.

[8]This is the case, for example, in both the EVLA and ALMA where a 1:2 frequency ratio between the lower and upper edges of the final analog IF response is used because it is easier to maintain uniform gain than with a baseband response.

Timing accuracy in Phase Switching (Feb. 27, 2008)

(Relates to Section 7.5)

As described in Section 7.5 (pp. 240-248 of TMS2) phase switching is used to mitigate spurious correlation that results from unwanted components that infiltrate the receiving channels of different antennas. At low levels these can be introduced through power supplies, local oscillators, monitoring circuitry, etc. Phase switching also removes the effect of any DC offset in the zero level of the digitizer. The synchronous detector action in Fig. 1.8 (p. 20 of TMS2) is usually effected by a second application of the first π-shift phase switching waveform at a later point in the signal path, which thereby cancels the first switching of the wanted signals but modulates any unwanted components that have infiltrated the system after the first phase switching. After digitization the signals are not vulnerable to such infiltration and the second switching can be applied by reversal of the sign bit of the samples. Thus, at the correlator input, the unwanted components are modulated in phase but the wanted ones are not. For each antenna a different switching function is used, any pair of which are orthogonal, resulting in elimination of the unwanted components in the averaging at the correlator output.

In designing a phase switching system timing tolerances should be considered. In general the accuracy should be much smaller than the minimum interval between transitions in any function. For example, in ALMA (the Atacama Large Millimeter Array) two phase switching actions are used, one nested within the other (Emerson 2007). For sideband separation $\frac{\pi}{2}$-shift phase switching is used, with Walsh functions from a 128-element set with time base 2.048 s and minimum interval between transitions of 16 ms. A second 128-element set with time base 16 ms and minimum interval 125 μs is used for π-shift switching. Thus timing errors must be very small compared with 125 μs, that is $\sim 1\mu$s.

In general the orthogonality of Walsh functions requires that there should be no relative time shifts between the functions. The first switching occurs at the antenna location, that is, as early in the signal path as possible. Digitization may occur at the antenna or after transmission of the signal in analog form to a central processing location. The major system delays are shown in Fig. 1, in which τ_g is the geometric delay, τ_{tr} is the transmission delay (antenna to processing location), and τ_i is the instrumental compensating delay. Delays in the analog or digital circuitry are generally small enough to be neglected with regard to the timing of phase switching. There are three main timing requirements in the receiving system.

1. The total delay from the incident wavefront to the correlator input, $\tau_g + \tau_{tr} + \tau_i$, must be the same for all antennas to preserve the correlation of the wanted signals. This is implemented through adjustment of τ_i.

2. The corresponding transitions in the first and second phase switchings should be aligned in time so that the phase switching of the wanted signals is precisely canceled. For example, if the second phase switching is done at the central location, the timing of the second switching should be delayed relative to the first one by τ_{tr}.

3. Both the first and second switchings in any signal path should be delayed by the geometric delay of the corresponding antenna τ_g so that, at the correlator input, the switching transitions in the unwanted components are aligned in time from one antenna to another. The delay τ_g varies with time as the antennas track a source.

Requirement 2 above is concerned only with the relative accuracy of switchings within the same signal path from one antenna to the correlator. This is the simplest case because it is concerned only with offsets in two switchings of the same Walsh function. Consider the effect of a small time offset δ in the relative timing of the first and second switchings. For each transition, the timing difference causes the correlator output voltage to be reversed for a period δ and thereby cancels an equivalent interval of the unreversed output. Hence, for each transition, there is an effective loss of signal for a period 2δ. The average fractional loss of

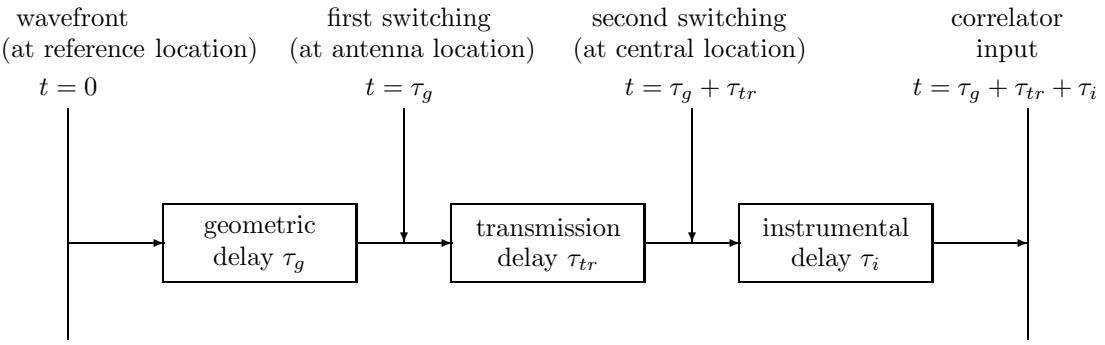

Figure 1: Delays in an array that are large enough to affect the timing of the Walsh functions used in phase switching. Here the second switching is shown after the transmission delay, as applies when the signals are transmitted in analog form to the central signal processing location. If digitization occurs at the antenna location both the first and second switchings can be applied before the transmission delay. t is time relative to a signal wavefront at the point where it intercepts the delay reference antenna.

sensitivity is $2n_t\delta/\tau_{tb}$ where n_t is the number of transitions within the time base τ_{tb} (i.e. twice the Walsh sequency). Thus for a tolerable limit of, e.g., 1% correlation loss, the tolerable value of δ can be determined for any given time base and maximum sequency used. Since the correlation loss is proportional to n_t, use of the lowest sequencies within the Walsh set helps to minimize loss in sensitivity. For arrays in which the numbers of antennas and the baseline lengths are not too long, the delaying of the switchings by τ_g (as noted in the third requirement above) can often be neglected. This introduces a timing error τ_g which is greatest for the longest baselines. The effect this error can be minimized by using the lowest values of n_t for the antennas for which the geometric delay is greatest.

Requirements 1 and 3 listed above are concerned with the relative timing of transitions at different antennas, i.e. between different Walsh functions. The effect of a timing offset on the rejection of the unwanted components depends on the loss in orthogonality of the Walsh functions used for different antennas. This is more complicated than the effect of an offset on two identical Walsh functions discussed above. The loss in orthogonality depends upon the sequencies of the two functions involved and is greatest for sequencies in the middle range of the Walsh set, as shown by Emerson (2005). Pairs consisting of a function with an even sequency and one with an odd sequency remain orthogonal in the presence of time shifts, but such combinations are possible for no more than half the pairs in a complete Walsh set. Of the other pairs, some remain orthogonal with time offsets, as can be shown by numerical trials, and some do not [see Fig. 3 in Emerson (2005)]. It is clearly beneficial to use equal numbers of odd and even sequencies in an array, so that for approximately half of the antenna pairs the orthogonality is independent of time offsets.

References

Emerson, D.T., Walsh Function Demodulation in the Presence of Timing Errors, Leading to Signal Loss and Crosstalk, ALMA Memo. 537, NRAO, 2005. (ALMA Memos can be found at www.alma.nrao.edu/memos/index.html)

Emerson, D.T., Walsh Function Definition for ALMA, ALMA Memo. 565, NRAO, 2007. (ALMA Memos can be found at www.alma.nrao.edu/memos/index.html)

Quantization Efficiency (May 11, 2006, Apr. 10, 2007, May 14, 2009)

(Relates to Section 8.3)

In a system that includes digital signal processing, quantization of the digitized (i.e. sampled) values of the received signal results in the addition of a component of random noise. This degrades the signal-to-noise ratio relative to that for a fully analog system. The quantization efficiency η_Q is defined as the ratio of the signal to noise amplitudes (SNR) at the output of the correlator of a digital system, divided by the same quantity for the corresponding analog system. The quantization efficiency is also equal to the variance of the noise at one input of the correlator for an ideal analog system divided by the equivalent quantity for a digital system.

Let x represent the voltage of the unquantized signal samples, which have a Gaussian probability distribution with variance σ^2. Let y represent the quantized values of x. The difference $x - y$ represents an inequality introduced by the quantization. The inequality contains a component that is correlated with x, and an uncorrelated component that behaves as random noise. Consider the correlation coefficient between x and $\Delta = x - \alpha y$, where α is a scaling factor. The correlation coefficient is

$$\frac{\langle x\Delta\rangle}{x_{rms}\Delta_{rms}} = \frac{\langle x^2\rangle - \alpha\langle xy\rangle}{x_{rms}\Delta_{rms}}. \tag{1}$$

Here the angle brackets $\langle\,\rangle$ indicate the mean value. If $\alpha = \langle x^2\rangle/\langle xy\rangle$, then the correlation coefficient is zero, and Δ represents purely random noise. We refer to this random component as the quantization noise, equal to $x - \alpha_1 y$ where $\alpha_1 = \langle x^2\rangle/\langle xy\rangle$. Without loss of generality, we take $\sigma^2 = \langle x^2\rangle = 1$ in this note, and use $\alpha_1 = 1/\langle xy\rangle$. The variance of the quantization noise is

$$\langle q^2\rangle = \langle (x - \alpha_1 y)^2\rangle = \langle x^2\rangle - 2\alpha_1\langle xy\rangle + \alpha_1^2\langle y^2\rangle = \alpha_1^2\langle y^2\rangle - 1. \tag{2}$$

The total variance of the digitized signal is $1 + \langle q^2\rangle$ and the quantization efficiency η_Q is equal to the variance of the unquantized signal expressed as a fraction of the total variance. Thus,

$$\eta_Q = \frac{1}{(1 + \langle q^2\rangle)} = \frac{1}{\alpha_1^2\langle y^2\rangle} = \frac{\langle xy\rangle^2}{\langle y^2\rangle}. \tag{3}$$

A formula for η_Q for eight and higher numbers of levels is given in Ch. 8 of the second edition (pp. 273-276). However, the quantization inequality was used as an approximation for the quantization noise as defined above, i.e. in effect α was taken to be 1. This is a fairly good approximation if the number of quantization levels is not too small. Here we provide exact formulas using Eq. (3).

Consider the case for an even number of equally spaced levels, as in the 8-level case in Fig. 1. It is convenient to define $\mathcal{N}$ as half the number of levels. We first determine $\langle xy\rangle$. Note that for each sample value, x and y have the same sign, so xy is always positive. Let ϵ represent the difference between adjacent quantization levels. The values of x that fall within the quantization level between $m\epsilon$ and $(m+1)\epsilon$ are assigned values $y = (m + 1/2)\epsilon$, and their contribution to $\langle xy\rangle$ is

$$\frac{1}{\sqrt{2\pi}}\int_{m\epsilon}^{(m+1)\epsilon} (m + \tfrac{1}{2})\epsilon\, x\, e^{-x^2/2}\, dx. \tag{4}$$

The contribution from the level between $-m\epsilon$ and $-(m+1)\epsilon$ is the same as the expression above, so to obtain $\langle xy\rangle$ we sum the integrals for the positive levels and include a factor of two:

$$\langle xy\rangle = \sqrt{\frac{2}{\pi}}\left[\left(\sum_{m=0}^{\mathcal{N}-2}\int_{m\epsilon}^{(m+1)\epsilon} \big(m + \tfrac{1}{2}\big)\epsilon\, x\, e^{-x^2/2}\, dx\right) + \int_{(\mathcal{N}-1)\epsilon}^{\infty} \big(\mathcal{N} - \tfrac{1}{2}\big)\epsilon\, x\, e^{-x^2/2}\, dx\right]. \tag{5}$$

The summation term contains one integral for each positive quantization level except the highest one. The integral on the right covers the highest level and the range of x above it, for both of which the assigned value is $y = (\mathcal{N} - 1/2)\epsilon$. Then performing the integration, Eq. (5) reduces to

$$\langle xy \rangle = \sqrt{\frac{2}{\pi}}\,\epsilon\Big(\frac{1}{2} + \sum_{m=1}^{\mathcal{N}-1} e^{-m^2\epsilon^2/2}\Big). \tag{6}$$

To evaluate the variance of y, again consider first the contribution from values of x that fall between $m\epsilon$ and $(m+1)\epsilon$. For this level the quantized data y all have the value $(m+1/2)\epsilon$. The variance of y for all values of x within this level is

$$(m+\tfrac{1}{2})^2\,\epsilon^2 \frac{1}{\sqrt{2\pi}} \int_{m\epsilon}^{(m+1)\epsilon} e^{-x^2/2}\,dx. \tag{7}$$

For negative x we again include a factor of 2, sum over all positive quantization levels except the highest, and add a term for the highest level and the range of x above it. Thus the total variance of y is:

$$\langle y^2 \rangle \;=\; \sqrt{\frac{2}{\pi}}\Bigg[\Big(\sum_{m=0}^{(\mathcal{N}-2)} (m+\tfrac{1}{2})^2\,\epsilon^2 \int_{m\epsilon}^{(m+1)\epsilon} e^{-x^2/2}\,dx\Big) + (\mathcal{N}-\tfrac{1}{2})^2\,\epsilon^2 \int_{(\mathcal{N}-1)\,\epsilon}^{\infty} e^{-x^2/2}\,dx\Bigg]. \tag{8}$$

The right-hand side of Eq. (8) can be simplified by expressing the integrals in terms of the error function erf(): $\mathrm{erf}(\xi/\sqrt{2}) = \sqrt{2/\pi}\int_0^\xi \exp(-t^2/2)\,dt$. Then using Eqs. (3) (6) and (8) we obtain,

$$\eta_{Q(2\mathcal{N})} = \frac{\frac{2}{\pi}\Big(\frac{1}{2} + \sum_{m=1}^{\mathcal{N}-1} e^{-m^2\epsilon^2/2}\Big)^2}{\big(\mathcal{N}-\frac{1}{2}\big)^2 - 2\sum_{m=1}^{\mathcal{N}-1} m\,\mathrm{erf}\big(\frac{m\epsilon}{\sqrt{2}}\big)}. \tag{9}$$

For the case where the number of levels is odd the thresholds of the levels occur at values that are an integer plus $\frac{1}{2}$, as in the 9-level case in Fig. 1. We represent the odd level number by $2\mathcal{N}+1$. Consider the values of x that fall within the quantization level between $(m-\frac{1}{2})\epsilon$ and $(m+\frac{1}{2})\epsilon$. These are assigned the value $m\epsilon$. For this level the contribution to $\langle xy \rangle$ is

$$\frac{1}{\sqrt{2\pi}} \int_{(m-\frac{1}{2})\epsilon}^{(m+\frac{1}{2})\epsilon} m\epsilon\, x\, e^{-x^2/2}\,dx. \tag{10}$$

Summing over all levels, as in Eq. (5), we obtain

$$\langle xy \rangle = \sqrt{\frac{2}{\pi}}\Bigg[\Big(\sum_{m=1}^{\mathcal{N}-1} \int_{(m-\frac{1}{2})\epsilon}^{(m+\frac{1}{2})\epsilon} m\epsilon\, x\, e^{-x^2/2}\,dx\Big) + \int_{(\mathcal{N}-\frac{1}{2})\epsilon}^{\infty} \mathcal{N}\epsilon\, x\, e^{-x^2/2}\,dx\Bigg]. \tag{11}$$

Then, as in Eq. (8) we determine $\langle y^2 \rangle$:

$$\langle y^2 \rangle = \sqrt{\frac{2}{\pi}}\Bigg[\sum_{m=1}^{(\mathcal{N}-1)} \Big(\int_{(m-\frac{1}{2})\epsilon}^{(m+\frac{1}{2})\epsilon} (m\,\epsilon)^2\, e^{-x^2/2}\,dx\Big) + \int_{(\mathcal{N}-\frac{1}{2})\epsilon}^{\infty} (\mathcal{N}\,\epsilon)^2\, e^{-x^2/2}\,dx\Bigg] \tag{12}$$

Performing the integration in Eqs. (11) and (12) we obtain

$$\eta_{Q(2\mathcal{N}+1)} = \frac{\frac{2}{\pi}\Big(\sum_{m=1}^{\mathcal{N}} e^{-\left(m-\frac{1}{2}\right)^2\epsilon^2/2}\Big)^2}{\mathcal{N}^2 - 2\sum_{m=1}^{\mathcal{N}}(m-\frac{1}{2})\,\mathrm{erf}\Big(\frac{(m-\frac{1}{2})\epsilon}{\sqrt{2}}\Big)}. \tag{13}$$

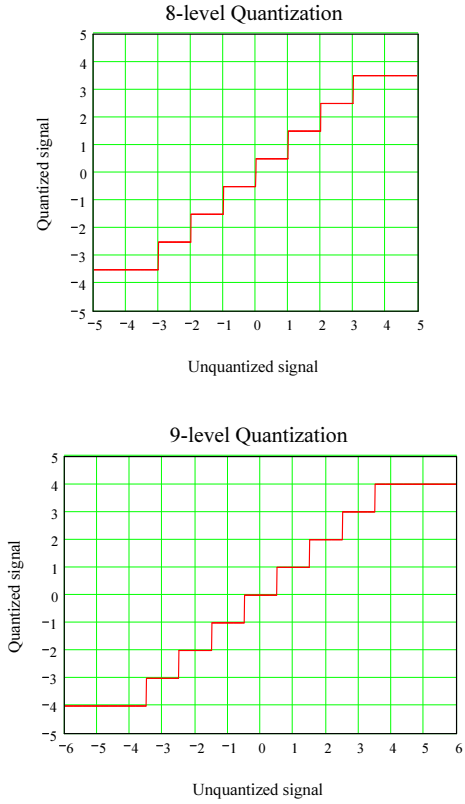

Figure 1: Examples of quantization characteristics with, top, an even number of levels (eight), and bottom, an odd number of levels (nine). In each case $\mathcal{N} = 4$. Units on both axes are equal to ϵ. Note that for even numbers of levels the the level thresholds occur at integral values on the abscissa, whereas for odd numbers of levels the thresholds occur at values that are an integer plus $\frac{1}{2}$.

Equations (9) and (13) can be easily evaluated by computer and provide precise values of quantization efficiency for any number of levels.

Since no significant approximations are made, the same method can be used for cases where the number of quantization levels is small and consequently the quantization noise is relatively large. Values of η_Q for two, three and four levels are derived on pp. 261-272 (pp. 217-229 in the first edition) by considering the cross-correlated products. The same values can be obtained by considering the effect of the quantization noise at a correlator input, following the method used above. In cases such as that for four levels, for which the assigned values for the levels are chosen to optimize η_Q, or for which the spacing between the level thresholds is not uniform, the formulas derived above cannot be applied directly. However, the same general approach can be used and, for example, for 3-level quantization the expressions are

$$\langle xy \rangle^{-1} = \sqrt{\frac{2}{\pi}} \left[\int_{0.612}^{\infty} x\, e^{-x^2/2}\, dx \right] \tag{14}$$

$$\langle y^2 \rangle = \sqrt{\frac{2}{\pi}} \left[\int_{0.612}^{\infty} e^{-x^2/2} dx \right]. \tag{15}$$

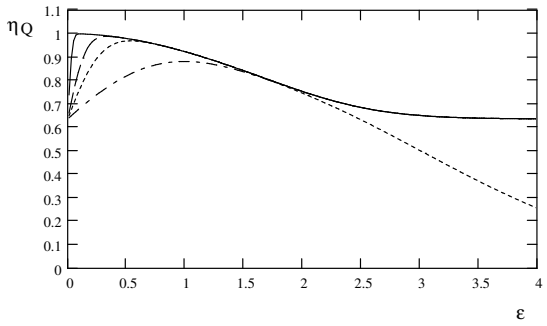

Figure 2: Quantization efficiency as a function of threshold spacing in units of σ. The curves are for 64-level (solid line), 16-level (long-dashed line), 9-level (short-dashed line), and 4-level (long-and-short-dashed line).

and

$$\eta_Q = \frac{\sqrt{\frac{2}{\pi}}\, e^{-0.612^2/2}}{erf(\infty) - erf(\frac{0.612}{\sqrt{2}})} \tag{16}$$

The threshold level of 0.612 can be varied to confirm that this value provides a maximum for η_Q.

Examples of results derived using Eqs. (9), (13) and (16) are shown in Table 1. In each case the values of ϵ are chosen empirically to maximize η_Q. Over the range of accuracy shown, the values of η_Q are in exact agreement with the equivalent values derived by Schwab (2005).

Table 1

No. of Levels	$\mathcal{N}$	ϵ	η_Q
3		0.612	0.809826
8	4	0.586	0.962560
9	4	0.534	0.969304
16	8	0.335	0.988457
32	16	0.188	0.996505
256	128	0.0308	0.999912

Figure 2 shows curves of η_Q as a function of ϵ for four different levels. As ϵ becomes very small the output of the quantizer depends mainly on the sign of the input, so the curves meet the ordinate axis at the two-level value of $\eta_Q = 2/\pi$. As ϵ increases, more of the higher (positive and negative) levels contain only values in the extended tails of the Gaussian distribution, so the number of levels that make a significant contribution to the output decreases, and the curves merge together. The curves for even level numbers move asymptotically to the two-level value, and curves for odd level numbers move toward zero.

If the constant voltage spacing between adjacent thresholds for both input and output values is not maintained, the individual levels can be adjusted to obtain an improvement in η_Q of a few tenths of a percent, decreasing with increasing number of levels. Level values optimized in this way are given by Jenet and Anderson (1998) for several numbers of levels. Threshold optimization is also included in a detailed analysis of quantization effects (Schwab 2007).

In recent designs of radio telescopes the level increment ϵ is frequently chosen so that signals at levels much higher than the rms system noise can be accommodated within the range of levels of the quantizer. This preserves an essentially linear response to interfering signals so that they can be eliminated or mitigated by

further processing. For example, with 256 levels (8-bit representation) and $\epsilon = 0.5$, we find that $\eta_Q = 0.9796$. The range of ± 128 levels then corresponds to $\pm 64\sigma$, i.e. ± 36 dB above the system noise. Thus a 2% sacrifice in signal-to-noise ratio allows for 36 dB of "headroom" above the normal rms level.

The principal results on quantization efficiency given above are also published in Thompson, Emerson, and Schwab, 2007.

References

Jenet F. A. and S. B. Anderson , The Effects of Digitization on Nonstationary Stochastic Signals with Applications to Pulsar Signal Baseband Recording, *Pub. Astron. Soc. Pacific*, 110, 1467–1478, 1998.

Schwab, F. R., Optimal Quantization Functions for Multi-Level Digital Correlators, *in preparation*, 2007.

Thompson, A. R., D. T. Emerson, and F. R. Schwab, Convenient Formulas for Quantization Efficiency, Radio Sci., 42, RS3022, doi:10.1029/2006RS003585, 2007[1]

Linearity Correction for Quantized Data and the Effect of Oversampling on Quantization Efficiency (Dec. 4, 2010)

(Relates to Ch. 8)

1. Introduction.

The cross-correlation of two signals that is calculated from quantized samples should not be assumed to be a precise estimate of the true correlation of the unquantized signals. Chapter 8 of TMS2 contains a discussion of quantization of the received signals, largely in terms of sampling with no more than four levels. Curves indicating nonlinearity in the cross correlation of quantized signals are shown in Fig. 8.6. Table 8.1 gives estimates of the effect of two-times oversampling on the quantization efficiency for 3 and 4 level sampling, which depend upon assumptions of linearity in the quantized correlation curves. More recently, advances in digital technology are facilitating a gradual increase in the number of levels used in large instruments, so the purpose of this note is to derive accurate general-case expressions for determining the nonlinearity in the cross correlation and also for calculation of the effect of oversampling over a range of oversampling factors and quantization levels.

Quantization is necessary to allow the use of digital techniques for the application of compensating delays and for computing cross correlation, which are most efficiently performed digitally. The digital values of cross correlation deviate from the true correlation of the unquantized signals to an extent which is most serious for two-level sampling, and decreases as the number of levels is increased. Correction for this effect requires determination of how the quantized cross-correlation, here designated R, is related to the true cross-correlation, ρ. Interferometry is concerned with the correlated component of the signals received in two antennas. To examine the effect of quantization, we consider the effect of a time offset τ on two Gaussian waveforms that are otherwise identical. In the case of two-level sampling the required relationship is given by the Van Vleck equation, (Eq. (8.25) in TMS2),

$$R_2(\tau) = \frac{2}{\pi} \sin^{-1} \rho(\tau). \tag{1}$$

For more than two quantization levels the relationship is more complicated, and although the nonlinearity of the quantized correlation becomes less serious with increasing number of levels, correction may still be necessary. As very large new instruments come into operation, for example the SKA (square-kilometer array)

[1] In Eq.(11) in the Radio Science printing, there is an error in the function erf() in the denominator. Compare with Eq. (13) given above.

project, etc., it will be necessary to remove the responses to strong radio sources in order to study the fainter emission from the most distant regions of the universe. This will require very accurate calibration of the received signal strengths. Similar considerations occur in the case of separating the responses of strong radio lines from those of weaker ones. A general-case derivation of the correlation coefficient for quantized data, from which numerical results can easily be computed, is given in the following section. In section 3 these results are also used to determine the effect of oversampling on quantization efficiency.

2. Correlation Coefficient for Quantized Data.

Let x and y represent two Gaussianly-distributed data streams that differ only by a time offset, τ. The correlation coefficient, $\rho(\tau)$ is equal to $< xy > / < x^2 >$ where the brackets $<>$ indicate mean values. The quantized values of x and y are identified by primes, i.e. x' and y'. The correlation coefficient of the quantized variables is

$$R(\tau) =< x'y' > / < x'^2 > . \tag{2}$$

To determine $R(\tau)$ as a function of the correlation coefficient ρ, we need to consider the probabilities of occurence of the unquantized variables x and y within each quantization interval. First consider the case where the number of quantization intervals is even, and equal to $2N_q$. Thus there are N_q positive intervals plus N_q negative ones. For simplicity, take the width of these intervals to be unity within the range $(N_q - 1)$ to $-(N_q - 1)$. Values that fall outside this range are within the extreme intervals $[(N_q - 1)$ to $\infty]$ or $[-(N_q - 1)$ to $-\infty]$. The mean value of the products of pairs of the quantized values, $< x'y' >$, is obtained by considering each of the ${4N_q}^2$ possible pairings of the levels of x' and y'. Only half of these need be calculated, since if the x and y values are interchanged the probability remains the same. The probability of the unquantized variables x and y falling within any pair of intervals is given by integration of the Gaussian bivariate probability distribution, $p(x, y,)$ over the corresponding ranges of x and y, where,

$$p(x, y) = \frac{1}{2\pi\sigma^2\sqrt{1-\rho^2}} \exp\left[\frac{-(x^2 + y^2 - 2\rho xy)}{2\sigma^2(1-\rho^2)}\right]. \tag{3}$$

Here σ is the variance of the noise waveforms, ρ is the correlation coefficient for the unquantized data, which is a function of the time difference between samples, $n\tau_s$, where τ_s is the time interval between adjacent samples and n is the number of time intervals between the samples considered. For a rectangular passband of width $\Delta\nu$, the correlation coefficient is given by

$$\rho(n\tau_s) = \frac{\sin(\pi n\tau_s)}{\pi n\tau_s}. \tag{4}$$

To calculate $< x'y' >$ for each combination of two quantization intervals, the joint probability of the required unquantized variables falling within these intervals is multiplied by the product of the corresponding values assigned in the quantization process. These results are then summed for all the pairs of intervals. Since the probability distributions of x and y are both symmetrical about zero, first consider the case where both these variables are positive and run from zero to ∞. As noted above we take the step size to be unity. Let $L(i)$ be the series of $N_q + 1$ values that define the positive quantization steps, i.e., $0, 1, 2...(N_q - 1), \infty$. Thus, for $i = 1$ to N_q, $L(i) = i - 1$, and $L(N_q + 1) = \infty$. For y there are an identical series of levels represented as $L(j)$. Then the component of $< x'y' >$ that results from the positive ranges of x and y is

$$\sum_{i=1}^{N_q}(i - 1/2)\left[\sum_{j=1}^{N_q}(j - 1/2)\int_{L(j)}^{L(j+1)}\int_{L(i)}^{L(i+1)} p(x, y)\, dx\, dy\right], \tag{5}$$

where $(i - 1/2)$ and $(j - 1/2)$ are the values assigned to the corresponding quantization intervals. The case where both x and y are negative provides an equal component of $< x'y' >$. Thus the component of $< x'y' >$

for cases where x and y have the same sign is

$$\frac{1}{\pi\sigma^2\sqrt{1-\rho^2}}\sum_{i=1}^{N_q}(i-1/2)\left[\sum_{j=1}^{N_q}(j-1/2)\int_{L(j)}^{L(j+1)}\int_{L(i)}^{L(i+1)}\left[\exp\left(\frac{-(x^2+y^2-2\rho xy)}{2\sigma^2(1-\rho^2)}\right)\right]dx\,dy\right]. \tag{6}$$

For the cases where x and y have opposite signs, one of either $(i-1/2)$ or $(j-1/2)$ is negative, and the sign of either x or y within the exponential function in (6) is negative. When the corresponding expression is added to (6) we obtain

$$\begin{aligned}< x'y' >= \frac{1}{\pi\sigma^2\sqrt{1-\rho^2}}\sum_{i=1}^{N_q}(i-1/2)\Bigg[\sum_{j=1}^{N_q}(j-1/2)\int_{L(j)}^{L(j+1)}\int_{L(i)}^{L(i+1)}\Bigg[\exp\left(\frac{-(x^2+y^2-2\rho xy)}{2\sigma^2(1-\rho^2)}\right)\\ -\exp\left(\frac{-(x^2+y^2+2\rho xy)}{2\sigma^2(1-\rho^2)}\right)\Bigg]dx\,dy\Bigg].\end{aligned} \tag{7}$$

Equation (7) shows how $< x'y' >$ is derived from the usual form of the bivariate distribution in Eq. (3). An equivalent form of the function in Eq. (3) is given by Abramowitz and Stegun (1964, see Eq. 26.3.2), which avoids the use of double integrals, and using this Eq. (7) can be written as follows:

$$\begin{aligned}< x'y' >= \frac{1}{\sqrt{2\pi}\sigma}\Bigg[&\sum_{i=1}^{N}\Bigg[(i-1/2)\sum_{j=1}^{N}\Bigg[(j-1/2)\int_{L_i}^{L_{i+1}}\operatorname{erfc}\left(\frac{(L_j-\rho(q)x)}{\sigma\sqrt{2(1-\rho(q)^2)}}\right)\exp\left(\frac{-x^2}{2\sigma^2}\right)dx\Bigg]\Bigg]\\ &-\sum_{i=1}^{N}\Bigg[(i-1/2)\sum_{j=1}^{N}\Bigg[(j-1/2)\int_{L_i}^{L_{i+1}}\operatorname{erfc}\left(\frac{(L_{j+1}-\rho(q)x)}{\sigma\sqrt{2(1-\rho(q)^2)}}\right)\exp\left(\frac{-x^2}{2\sigma^2}\right)dx\Bigg]\Bigg]\\ &-\sum_{i=1}^{N}\Bigg[(i-1/2)\sum_{j=1}^{N}\Bigg[(j-1/2)\int_{L_i}^{L_{i+1}}\operatorname{erfc}\left(\frac{(L_j+\rho(q)x)}{\sigma\sqrt{2(1-\rho(q)^2)}}\right)\exp\left(\frac{-x^2}{2\sigma^2}\right)dx\Bigg]\Bigg]\\ &+\sum_{i=1}^{N}\Bigg[(i-1/2)\sum_{j=1}^{N}\Bigg[(j-1/2)\int_{L_i}^{L_{i+1}}\operatorname{erfc}\left(\frac{(L_{j+1}+\rho(q)x)}{\sigma\sqrt{2(1-\rho(q)^2)}}\right)\exp\left(\frac{x^2}{2\sigma^2}\right)dx\Bigg]\Bigg]\Bigg],\end{aligned} \tag{8}$$

where erfc is the complimentary error function. Although this form is a little less concise than Eq. (7), the avoidance of double integrals allows it to run significantly faster, and also the result was found to be less dependent on the infinity limit used[1].

To calculate $< x'^2 >$ the Gaussian probability function for a single variable is used, taking double the expression for the positive range of x:

$$\begin{aligned}< x'^2 >&= \frac{\sqrt{2}}{\sqrt{\pi}\sigma}\sum_{i=1}^{N_q}(i-1/2)^2\int_{L(i)}^{L(i+1)}\exp\left(\frac{-x^2}{2\sigma^2}\right)dx\\ &=\sum_{i=1}^{N_q}(i-1/2)^2\left[\operatorname{erfc}\left(\frac{L_i}{\sqrt{2}\sigma}\right)-\operatorname{erfc}\left(\frac{L_{i+1}}{\sqrt{2}\sigma}\right)\right].\end{aligned} \tag{9}$$

[1]In numerical calculation of $< x'y' >$ using Eq. (7) the choice of the number to represent infinity in the integration over the top quantization interval may be important, depending upon the behavior of the algorithms used in evaluating the integrals. In the case of two-level sampling, values of the correlation coefficient R_2 are given by Eq. (1), and can be used to check the same results derived from Eqs. (7) and (9). Experience with Mathcad showed that the best agreement with Eq. 1 was obtained with values in the range 10-20 for the for the infinity limit of the integrals, and within this range the two-level values generally agreed with Eq. 1 within eight decimal places.

Thus for a given value of the time interval between samples, the correlation coefficient for the quantized data is as given in Eqs. (2), (7) or (8), and (9). Note that the dependence on the time interval $\tau = n\tau_s$ is introduced through ρ in Eq. (4).

In many cases the number of quantization levels is even, and often a power of two, but an odd number can also be used. For an odd number of levels, $2N_q + 1$, the quantization intervals for positive values are $\frac{1}{2}$, $1\frac{1}{2}$, $2\frac{1}{2}$ etc., i.e. for $i = 1$ to N_q, $L(i) = i - \frac{1}{2}$, and $L(N_q + 1) = \infty$. For the interval $-\frac{1}{2}$ to $\frac{1}{2}$, the assigned (quantized) value is zero, so this interval makes no contribution to the expressions for $< x'y' >$ and $< x'^2 >$. Thus the assigned values of the quantized data, $(i - \frac{1}{2})$ and $(j - \frac{1}{2})$ in Eqs. (7) and (8), become i and j and the values of $L(i)$ and $L(j)$ are redefined, but otherwise these equations remain the same for either $2N_q$ or $(2N_q + 1)$ levels.

Figures 1 and 2 show examples of the relationship between the correlation of the quantized signals and the true signal correlation, using Eqs. (2), (8), and (9). For a few cases, similar curves are shown in Fig. 8.6 of TMS2, but a general expression for the for the correlation coefficient of quantized signals was not derived. To facilitate correction for quantization as part of the signal processing, it may be helpful to generate expressions that give the true correlation as a rational function of the measured correlator output, as in the examples for the four-level case in Appendix 8.3 of TMS2. The equations derived above can also be used to examine effects such as the variation in the rms level σ with respect to the step size, which can occur as a result of gain variation in the analog amplification that precedes the quantization.

3. Oversampling.

As discussed in Ch. 8 of TMS2, sampling of signals at the Nyquist rate results in no loss of information, but subsequent quantization causes the loss in sensitivity represented by the quantization efficiency. General formulas for quantization efficiency are derived in the supplementary notes under Quantization Efficiency and in Thompson, Emerson, and Schwab (2007). Some of the loss due to quantization can be recovered by oversampling, that is, sampling faster than the Nyquist rate. A general analysis of the effect of oversampling involves use of the correlation coefficient for quantized data derived in section 2 above.

For sampling of random noise with an ideal rectangular spectrum of width $\Delta\nu$, the time interval between adjacent Nyquist samples is $1/2\Delta\nu$. With Nyquist sampling the noise within each sample is uncorrelated with respect to the noise in any other sample, and when such data are combined the noise combines additively in power. Consider the case of oversampling, where the number of samples per second is β times the Nyquist rate. When the sample rate exceeds the Nyquist rate, the samples are no longer independent and for any particular sample there are components of the noise within other samples that are correlated with the noise in the sample considered. [Note, however, that for any two samples spaced by β times the sample interval (i.e. spaced at the Nyquist interval), or by an integral multiple of the Nyquist interval, the noise is uncorrelated[2]]. The correlated components of the noise in different samples combine additively in voltage, rather than additively in power as is the case for uncorrelated noise.

To illustrate how the components of noise combine, consider one pair of antennas and, for example, just four consecutive samples at the correlator output. Let a_1, a_2, a_3 and a_4 be these voltages, which are proportional to the product of the voltages at the correlator inputs. Then we have for the squared sum of these correlated noise voltages, i.e. the total noise power,

$$[a_1 + a_2 + a_3 + a_4]^2 = a_1^2 + a_2^2 + a_3^2 + a_4^2 + 2(a_1a_2 + a_1a_3 + a_1a_4 + a_2a_3 + a_2a_4 + a_3a_4). \qquad (10)$$

The autocorrelation coefficient of the quantized signals at the correlator input is $R(n\tau_s)$, where n is an integer and τ_s is the spacing in time between adjacent samples. The output of the correlator consists of values that are the product of two input samples, so the autocorrelation coefficient of the samples at the correlator output is $R^2(n\tau_s)$. The mean noise power is given by the mean of the terms in the right-hand side

[2]It can be assumed that the noise components of the signals from any two antennas are uncorrelated, because noise from the sky background that is received in separate antennas is resolved, and generally the antennas are sufficiently far apart that noise cross-talk can be ignored.

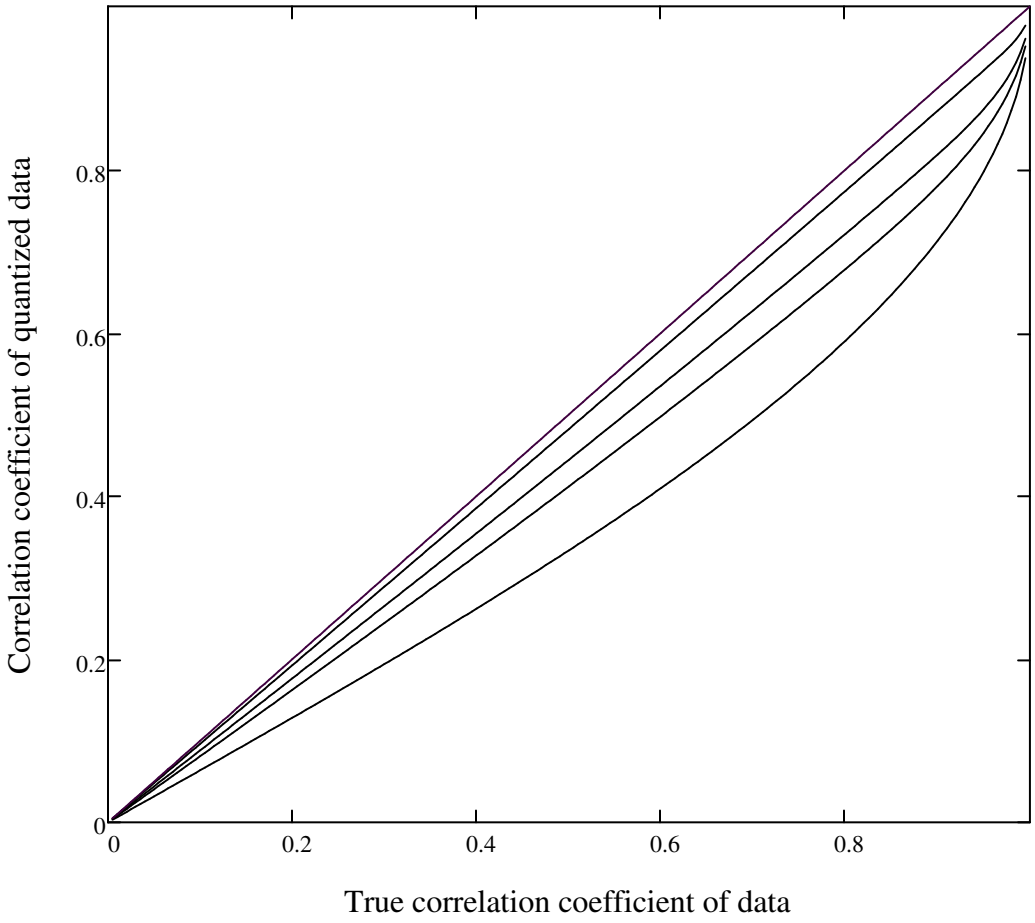

Figure 1: Curves of the correlation coefficient of quantized data as a function of the true correlation (i.e. the correlation of the unquantized data). The lowest curve is for two-level quantization, and moving upward the curves are for three, four, and eight levels, and the top line indicates the ideal response.

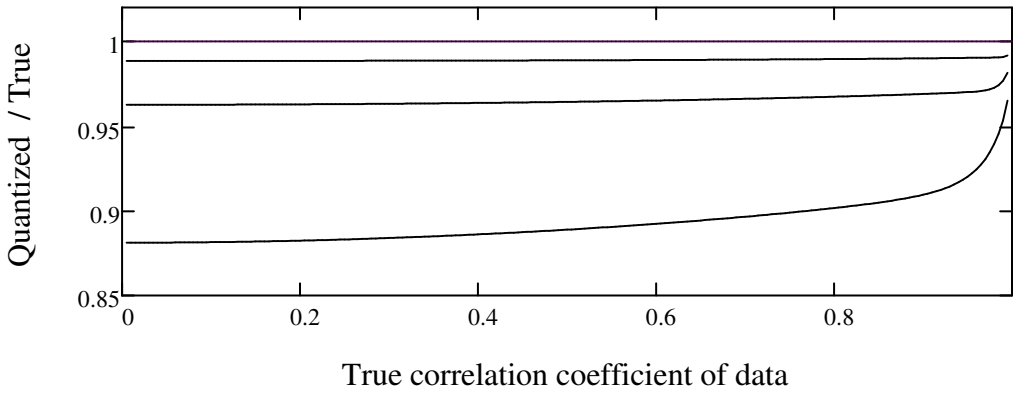

Figure 2: Curves of the correlation coefficient of quantized data expressed as a fraction of the true correlation, plotted as a function of the true correlation. The lowest curve is for four-level quantization, and moving upward the curves are for eight and 16 levels, and the top line indicates the ideal response.

of (10), in which each of the a_n^2 terms can be replaced by the mean squared noise amplitude $< a^2 >$, and each of the $a_m a_n$ terms by $< a^2 > R^2(|n-m|\tau_s)$. Thus the squared sum of the four noise voltages becomes

$$4 < a^2 > +2 < a^2 > [3R^2(\tau_s) + 2R^2(2\tau_s) + R^2(3\tau_s)]. \tag{11}$$

If the four noise terms were uncorrelated, i.e. if the R^2 terms were zero, the noise power would be the sum of the individual noise powers, $4 < a^2 >$. The effect of the correlation of the noise is to increase the averaged noise power by a factor equal to (11) divided by $< 4a^2 >$:

$$1 + 2[(3/4)R^2(\tau_s) + (1/2)R^2(2\tau_s) + (1/4)R^2(3\tau_s)]. \tag{12}$$

In the general case, averaging a total of N samples at the correlator output, this factor becomes

$$1 + 2\left[\left(\frac{N-1}{N}\right)R^2(\tau_s) + \left(\frac{N-2}{N}\right)R^2(2\tau_s) + \left(\frac{N-3}{N}\right)R^2(3\tau_s) + ... \left(\frac{1}{N}\right)R^2[(N-1)\tau_s]\right]. \tag{13}$$

In practice, in radio astronomy, the rate at which the data are sampled is in the range of MHz to GHz. The averaging times are in the range $\sim$10 milliseconds to seconds, so N is likely to be within the range 10^4 to 10^9. The autocorrelation coefficient decreases as the time interval between samples increases, and in practice $R^2(n\tau_s)$ becomes very small for $n\tau_s >\sim 200$ times the Nyquist sample interval. Thus, for the terms within the square brackets in (12), those after about the first 200β can be neglected. Since, in most cases, $N \gg 200\beta$, the squared sum of the noise voltages simplifies to

$$1 + 2[R^2(\tau_s) + R^2(2\tau_s) + R^2(3\tau_s) + ...] = 1 + 2\Sigma_{n=1}^{\sim 200\beta} R^2(n\tau_s). \tag{14}$$

Equation (14) is the fractional increase in the squared noise voltage (i.e. the noise power) that results from the fact that the noise in the samples is no longer independent when the data are oversampled. The quantization efficiency η_Q is equal to the quantization efficiency for Nyquist sampling, η_{QN}, multiplied by $\sqrt{\beta}$ to take

account of the increase in the number of samples, but divided by the square root of (14) because the noise in different samples is no longer independent. Thus, noting that $\tau_s = 1/(2\beta\Delta\nu)$, we obtain

$$\eta_Q = \frac{\eta_{QN}\sqrt{\beta}}{\sqrt{1 + 2\sum_{n=1}^{\sim 200\beta} R^2\left(\frac{n}{2\beta\Delta\nu}\right)}}. \tag{15}$$

For the case of two-level sampling the equivalent result is given in TMS2 (see Eq. 8.32), but Eq. (15) is not limited to any particular number of levels.

Table 8.1 of TMS2 gives values of the quantization efficiency for two-times oversampling ($\beta = 2$). For two level sampling the result is based on the exact expression for R_2 in Eq. 1, above, but for 3 and 4 levels the values involve the assumption that for adjacent samples the correlation is no higher than ~ 0.6 and varies in a closely linear manner with the correlation coefficient for unquantized data. However, as the oversampling factor β increases, the interval between adjacent samples decreases and the correlation of the noise within adjacent samples increases. Thus, in general, to determine the effect of oversampling on the quantization efficiency it is preferable to use precise expressions for the correlation coefficient of quantized data, as derived in Section 2 above. Examples of the quantization efficiency, η_Q, derived from Eqs. (2), (8), (9), and (15) are shown in Table 1 for 2-, 3-, 4-, 8- and 16-level sampling, and values of $\beta = 2, 4, 8, 16$, and 32. In each case the value of σ used is the one that maximizes η_Q for Nyquist sampling as given in Thompson, Emerson, and Schwab (2007)[3], and the values for $\beta = 1$ (η_{QN}) are also taken from this reference. Note that as β is increased the improvement gained by each further increase declines, because the correlation between adjacent samples increases and thus the new information provided by finer sampling becomes progressively smaller.

Table 1, Values of Quantization Efficiency.

No. of Levels	σ	$\beta = 1$	$\beta = 2$	$\beta = 4$	$\beta = 8$	$\beta = 16$	$\beta = 32$
2	1.0	0.6366	0.744	0.784	0.795	0.798	0.799
3	0.817	0.8098	0.882	0.912	0.920	0.922	0.923
4	1.005	0.8812	0.930	0.951	0.958	0.960	0.960
8	1.706	0.9626	0.980	0.987	0.991	0.991	0.992
16	2.985	0.9885	0.994	0.996	0.998	0.998	0.998

We thank F. R. Schwab for drawing attention to the alternate form of the bivariate probability function.

References

Abramowitz, M. and I. A. Stegun, *Handbook of Mathematical Functions*, National Bureau of Standards, Washington, DC, 1964, repr. by Dover, New York, 1965.

Thompson, A. R., D. T. Emerson, and F. R. Schwab, Convenient formulas for quantization efficiency, Radio Science, 42, RS3022, doi:10.1029/2006RS003585, 2007.

[3]In this reference, σ is taken to be unity and the size of the quantization steps is chosen to maximize the quantization efficiency. To use the same ratio of step size to σ in the present discussion (in which the step size is unity) the value of σ is the reciprocal of the step size (ϵ) that maximizes the quantization efficiency as given in Thompson, et al. (2007). In the two-level case any value can be used for σ.

Polyphase Filters (Sept. 22, 2006)

(Relates to Section 8.7)

Polyphase filtering is a digital signal processing technique that has been developed for applications such as separation of signals in multichannel communication systems, and has applications in radio astronomy in correlators for spectral-line observations. Tutorial papers on polyphase filters can be found in Harris (1999) and Harris, Dick, and Rice (2003), of which the latter provides a detailed analysis and contains a discussion of how the polyphase system relates to an equivalent analog process. This note is a brief review of the basic polyphase principle based largely on these references.

Consider a signal band of width $\Delta\nu$ that contains M equally-spaced channels, each of bandwidth $\Delta\nu/M$. In a communication system this might represent M modulated carriers, or in radio astronomy a signal band to be processed as M spectral channels. We suppose that the received signal has been converted to an IF band and digitally sampled at a Nyquist rate of $f_s = \Delta\nu$, each sample containing an I (in-phase) and a Q (quadrature) component. One way of separating the channels would be to perform an FFT (Fast Fourier transform) on each sequence of M samples to convert from time to frequency. However, the resulting channels then have a sinc function profile in frequency, which in communications systems results in low-level cross talk between channels. In radio astronomy it can cause difficulty in detecting weak lines that are close to strong lines or interfering signals. In the polyphase system further filtering is introduced to control the response of the channels. Figure 1 shows a basic polyphase filter system. The input digital samples are distributed by a commutator to the inputs of a set of M partitions[1] of a digital FIR filter (the partitions may thought of as a set of M individual filters). The sample frequency at the input to each partition is thus down-sampled to f_s/M, which is the appropriate Nyquist value for a single channel. Each partition shown implements two filtering actions, one for each of the I and Q components. A specific requirement of the polyphase scheme is that the center frequency of each channel must be an integral multiple of the output frequency f_s/M from each the filter partition. In practice this places a constraint on the choice of the frequencies of the IF band that goes to the digitizing sampler.

Let $h(n)$ be the impulse response of a digital lowpass filter for which the frequency response is the Fourier transform $H(\theta)$. Here n is the input sample index and thus is a measure of time in units of $1/f_s$. θ is the frequency, i.e. the phase change per time interval f_s^{-1}. For the equivalent bandpass filter centered at frequency θ_0 the frequency response is $H(\theta + \theta_0)$ and the impulse response is[2] $h(n)e^{-j\theta_0 n}$. Let r_p be the index of the filter partitions ($0 \le r_p \le M-1$) and r_t be the index of the filter tap weights[3] ($0 \le r_t \le L-1$, where L is the number of taps). We can represent the bandpass characteristic of the filter partitions in Fig. 1 by the lowpass impulse response $h(n)$ multiplied by a factor $exp(-j\theta_0 n)$. If we take $n = 0$ to represent the time at which a data point enters the first partition (for which $r_p = 0$), then the time for a data point to reach partition r_p is represented by $n = r_p$. As a data point shifts through its partition, it encounter tap points at values of n equal to $r_p,\ r_p + M,\ r_p + 2M,\ \ldots r_p + (L-1)M$, i.e. in general at $n = r_p + Mr_t$. Thus for a bandpass characteristic the impulse response of a partition, $h(n)e^{-j\theta_0 n}$, is a series of terms in r_t (one for each tap), each of which can be represented by

$$h(r_p + Mr_t)e^{-j\theta_0(r_p + Mr_t)} = h(r_p + Mr_t)e^{-j\theta_0 r_p}e^{-j\theta_0 Mr_t}. \tag{1}$$

As mentioned above, the center frequencies of the channels are chosen to be integral multiples of f_i/M. For each channel the center frequency is $\theta_0 f_s$ radians/s ($= \theta_0 f_s/2\pi$ Hz). Thus $\theta_0 f_s/2\pi$ is chosen to be equal to kf_s/M where k is an integer. (Note that θ_0 and k have different values for each channel.) As a result, $M\theta_0 = k \times 2\pi$. Hence the factor $e^{-j\theta_0 Mr_t}$ in the right-hand side of Eq. (1) is equal to unity and can be ignored. The remaining exponent on the right-hand side can be written as $exp(-j\frac{2\pi k}{M}r_p)$ and represents r_p units of

[1] Also sometimes referred to as segments.

[2] This follows from the shift theorem of Fourier transforms. Note also that with respect to the sign of the exponent we follow the convention used in TMS in which $f(\nu) = \int F(t)exp(-j2\pi ft)dt$.

[3] See TMS2, p.298.

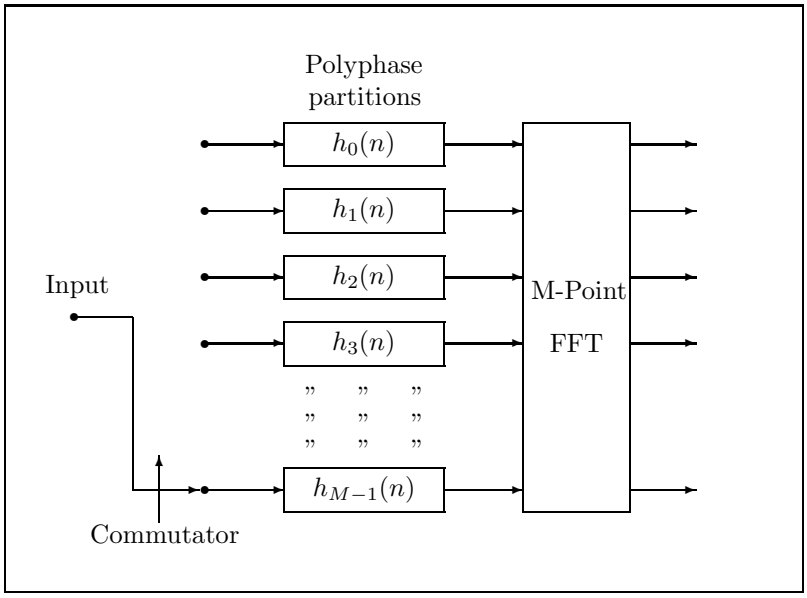

Figure 1: Polyphase filter system. The input stream with sample frequency f_s is distributed between the M filter partitions by the commutator, which starts at the r_p index $M-1$ and decrements to index 0. Thus the sample rate within the filter is down-sampled by the factor M. In each partition h represents the impulse response and it's subscript indicates the value of r_t for the particular partition. The M FFT outputs at the right-hand side are the output sample streams corresponding to the frequency channels.

phase, of which M units is an integral number of turns. Since the phase is proportional to the partition index r_p, it can be implemented by performing an M-point discrete Fourier transform, which forms appropriately phase-shifted sums of the of the outputs of the partitions. In practice this is implemented as an FFT. Note that the boxes that represent the partitions in Fig. 1 actually have identical low-pass characteristics, and the bandpass responses are effected when the FFT is performed. Although the input digital data represent a wide band of frequencies, $\Delta\nu$, they can pass through the lowpass (baseband) responses of the partitions since the down-sampling to digital frequency f_s/M results in aliasing of the channel center frequencies to zero. The FFT outputs correspond to the required M frequency channels. The term *polyphase filtering* refers to the series of phases $\theta_0 r_t$ applied to the partitions in the FFT.

To optimize the efficiency of the FFT, the number of partitions should be an integral power of two. In addition, the down-sampling implemented by the commutator and the choice of channel center frequencies as multiples of f_s/M result in a computationally efficient system. In the discussion above, to avoid complication, we have considered only the case where the length of the FFT is equal to the down-sampling ratio. Other cases are possible but require more complicated manipulation of the data, see, e.g., Harris (1999).

References

Harris, F., 1999, Tutorial on Polyphase Transforms, http://www.techonline.com/pdf/pavillions/icspat/1999/w211/pdf.

Harris, F., Dick, C., and Rice, M., 2003, Digital Receivers and Transmitters Using Polyphase Filter banks for Wireless Communications, *IEEE Trans. Microwave Theor. Tech.,* **51**, 1395-1412.

Advances in VLBI (May 20, 2009)

(Relates to Ch. 9)

As of 2009, the use of magnetic tape as the storage medium in very long baseline interferometry has been almost entirely replaced by disks of the type developed for computers. These provide higher capacity and are more convenient to handle. Errors and loss of data that can occur when getting the tapes synchronized at playback are avoided. Direct connection from the antenna stations to the correlator by using optical fiber is also being used in some cases. The European VLBI system has been a leader in using this real-time correlation technique. A possible disadvantage is that it is not possible to re-correlate the data with changes in the relative timing, as may be necessary to adjust the location of the field center. The relative cost of disks and of renting bandwidth in fiber systems will be a major factor in the choice between these schemes as development continues.

Since about 2008, increases in the sensitivity have extended the astrometric use of VLBI, allowing trigonometric parallax and proper motion measurements of maser emissions from star forming regions. Accuracies of ~ 10 μas are being achieved, allowing for measurements of distances across major regions of the galaxy by trigonometric parallax, thus providing increased accuracy in mapping of galactic structure (see, e.g. Reid et al., 2009). In the US there are plans for receivers with frequency response 4-8 GHz for the Very Long Baseline Array (VLBA), to further increase the sensitivity and include the 6.7 GHz maser transition of methanol.

References

Reid, M. J. et al., Astrophys. J., in publication, 2009

Fringe Rotation in Digital Correlators as used in VLBI (April 10, 2007)

(Relates to Section 9.7)

A detailed analysis of the fringe rotating scheme commonly used in VLBI is given by Roberts (1997). This is the system shown in Fig. 9.17(c) of TMS involving two cross correlators and two fringe rotation functions. Calculations are given for the efficiency (signal-to-noise) factors for two-, three-, and four-level signal quantization schemes with a three-level fringe rotation function. In comparing the results with the similar figures in Table 9.6 (Table 9.5 of TMS1) it should be noted that the values listed under *Fringe Rotation Loss* in the TMS table include a factor of $\sqrt{2}$ to take account of the increased signal-to-noise ratio that results from combining the outputs of the real and imaginary correlators. In the system that Roberts considers the outputs of each of the two correlators is processed by a fast Fourier transform which converts the data from functions of time to functions of frequency. A $\pi/2$ phase shift is applied at the output of the FFT processor on the imaginary correlator and the two FFT signals are then combined, providing a single sideband response. Roberts notes that a practical requirement of his analysis is that many fringe periods are contained within the integration time applied to the data. This is an important point since, as a result, the method is not adaptable to arrays with short baselines: the possibility of performing the fringe rotation in a computer rather than in the local oscillator system (see Ch. 6) might otherwise offer simplification of the hardware. In VLBI the fringe rates are usually some hundreds of Hertz or more and the number of fringe periods required presents no problem.

The requirement on the number of fringe periods does not result from the quantization of the waveforms and can be investigated simply using analog functions, as follows. Consider the system in Fig. 9.17(c) and assume that the signals from the two antennas that are inserted at the inputs at the top of the figure have been converted by identical LO signals so that the phase relationships resulting from the fringes are the same as at the antennas. We can represent the signal input on the left side of the figure by $\sin[2\pi(\nu_0 + \nu_f)t + \theta]$,

where ν_0 is the IF signal frequency, ν_f is the fringe frequency[1], and θ is the visibility phase which also includes a possible instrumental phase contribution that can later be removed by calibration. The signal input on the right-hand side is $\sin(2\pi\nu_0 t)$. For the fringe functions $\mathcal{F}_R$ and $\mathcal{F}_I$ we use $\cos(2\pi\nu_f)t$ and $\sin(2\pi\nu_f t)$, respectively. In addition to reducing the fringe frequency to zero the mixing action of the fringe functions produces an image response at a frequency of $\nu_0 - \nu_f$ at the left-hand input. However, this signal would not be correlated with the frequency ν_0 at the right hand input and would introduce only noise. Thus, in what follows the input at the unwanted image frequency is not included. The product of the three functions that are combined in the *real* correlator is

$$\sin[2\pi(\nu_0+\nu_f)t+\theta]\sin(2\pi\nu_0 t)\cos(2\pi\nu_f t) = \tfrac{1}{2}\{\cos(2\pi\nu_f t+\theta) - \cos[2\pi(2\nu_0+\nu_f)t+\theta]\}\cos(2\pi\nu_f t). \quad (1)$$

Similarly the product of the functions combined in the *imaginary* correlator is

$$\sin[2\pi(\nu_0+\nu_f)t+\theta]\sin(2\pi\nu_0 t)\sin(2\pi\nu_f t) = \tfrac{1}{2}\{\cos(2\pi\nu_f t+\theta) - \cos[2\pi(2\nu_0+\nu_f)t+\theta]\}\sin(2\pi\nu_f t). \quad (2)$$

On the right-hand sides of Eqs. (1) and (2) the terms $\cos[2\pi(2\nu_0+\nu_j)t+\theta]$ represent IF components that are filtered out in the correlation process. Thus the output of the real correlator is

$$\tfrac{1}{2}\cos(2\pi\nu_f t+\theta)\cos(2\pi\nu_f t) = \tfrac{1}{4}[\cos(4\pi\nu_f t+\theta)+\cos(\theta)], \quad (3)$$

and the output of the imaginary correlator is

$$\tfrac{1}{2}\cos(2\pi\nu_f t+\theta)\sin(2\pi\nu_f t) = \tfrac{1}{4}[\sin(4\pi\nu_f t+\theta)-\sin(\theta)]. \quad (4)$$

The $(4\pi\nu_f t+\theta)$ terms in the right hand sides of Eqs. (3) and (4) tend toward zero as the correlator output data are averaged over time. From the remaining terms the combined output of the two correlators is seen to be $\frac{1}{4}[\cos(\theta) - j\sin(\theta)]$. The fringe phase is θ. The amplitude factors $\frac{1}{2}$ and $\frac{1}{4}$ result from our arbirarily taking unit amplitude for the input functions in Eqs. (1) and (2), and hereafter are omitted.

In determining the visibility amplitude and phase from the averaged outputs of the two correlators it is convenient to consider averaging time in units of the fringe period, that is in units of $n = \nu_f t$. Values of the time-averaged outputs of the real and imaginary correlators, R_{av} and I_{av} respectively, are given by

$$R_{av} = \frac{1}{N}\int_0^N [\cos(4\pi n+\theta)+\cos(\theta)]dn = \frac{1}{4\pi N}[\sin(4\pi N+\theta)-\sin(\theta)] + \cos(\theta), \quad (5)$$

and

$$I_{av} = \frac{1}{N}\int_0^N [\sin(4\pi n+\theta)-\sin(\theta)]dn = \frac{-1}{4\pi N}[\cos(4\pi N+\theta)-\cos(\theta)] - \sin(\theta), \quad (6)$$

for an averaging time of N fringe periods. R_{av} and I_{av} have been computed in increments of $N/100$ for four values of θ. From these the fringe amplitude $\sqrt{R_{av}^2+I_{av}^2}$ and the phase $\tan^{-1}(I_{av}/R_{av})$ have been derived. These are plotted in Fig. 1 for averaging up to five fringe periods. It can be seen that the curves contain oscillations at twice the fringe frequency. For small values of N these cause serious errors in the estimated values of the amplitude and phase, but decrease with increasing N. The oscillations in the amplitude curves fall to 1% peak-to-peak at an averaging time of approximately 18 fringe periods, and in the phase they fall to 1° peak-to-peak at approximately 9 fringe periods. Depending on the value of θ, the oscillations may be centered on the true value of amplitude or phase, or the extremes of the oscillations may be aligned with the true value. Thus for the averaging times of 18 and 9 fringe periods, errors due to the oscillations can be as high as ±1% or ±1° respectively. The right-hand sides of Eqs. (3) and (4) represent the real and imaginary terms of a Fourier analysis of the fringe frequency at the correlator outputs, so the results derived also apply to the case where an FFT is applied at that point.

[1]The fringe frequency corrresponding to the field center is accurately known from the observing frequency and the calibrated antenna locations, but the fringe amplitude and phase for the source under observation are to be determined.

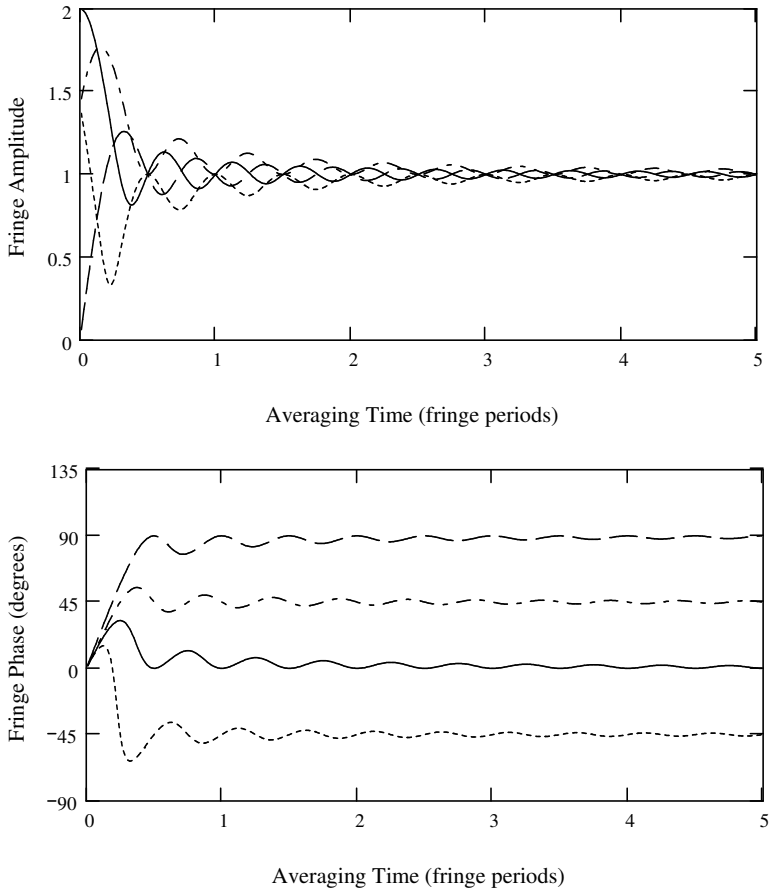

Figure 1: Fringe amplitude (upper diagram) and phase (lower diagram) as a function of averaging time shown in units equal to the fringe period. Curves are for values of θ of 0 (solid line), $\pi/4$ (dots), $\pi/2$ (dashes), and $3\pi/4$ (dash-dots).

For multielement VLBI arrays such as the VLBA it is simpler to apply the fringe rotation to the signals from the 10 antennas than to the 45 cross correlated signal pairs. For any correlated pair, both signals have undergone fringe rotation, so in systems using only a few quantization levels in the fringe rotation functions the signal-to-noise loss is incurred twice, as indicated under fringe rotation loss in Table 9.6. In the VLBA, however, eight-level fringe rotation functions are used and the fringe rotation loss is therefore negligible. The VLBA correlator is an FX type which can accommodate the increase in the number of levels in the signals resulting from the rotation functions: see Romney (1999), which also contains a review of correlator types and characteristics.

References

Roberts, P. P., Calculating Quantization Correction Formulae for Digital Correlators with Digital Fringe Rotation, *Astron. Astrophys. Suppl.*, 126, 379-383, 1997.

Romney, J. D., Cross Correlators, *Synthesis Imaging in Radio Astronomy II*, G. B. Taylor, C. L. Carilli, and R. A. Perley, Eds., ASP Conference Series, 180, 57-78, 1999.

Ghost Images (Jan. 12, 2007)

(Relates to Ch. 10, pp. 405-6)

Figure 10.6 on p. 405 of TMS2 illustrates how bandpass ripples are introduced into the visibility as a function of frequency, as a result of the sharp edges in the cross power spectrum. An effect of this discussed by Bos (1984) is the introduction of "ghost" images into the map derived from the observations. The ghost structure appears at a position that, relative to the true structure, is diametrically opposite with respect to the field center. For each spectral channel the amplitude of the ghost structure is proportional to the amplitude of the ripple component. Thus it is most serious for the channels at the edges of the receiver passband, as can be seen from Fig. 10.6.

The ghost phenomenon is most easily explained by considering a simple example. Suppose that we have a point source of unit amplitude at position $(\ell, m) = (\ell_1, 0)$, where $(0, 0)$ is the field center, and observe it over a range of baselines u. The fringe visibility is the FT (Fourier Transform)[1]with respect to ℓ of a delta function at ℓ_1, which is

$$\mathcal{V}_1(u) = e^{-j2\pi u\ell_1} = \cos(2\pi u\ell_1) - j\sin(2\pi u\ell_1) \tag{1}$$

We suppose that a multichannel spectral correlator is used so there is a visibility data set for each spectral channel. The the ripples across the spectrum in Fig. 10.6 have the effect that the relative amplitudes of the sine and cosine components are no longer equal, as they are in Eq. (1), so we rewrite Eq. (1) as

$$\mathcal{V}_1(u) = \cos(2\pi u\ell_1) - j(1+\Delta)\sin(2\pi u\ell_1) \tag{2}$$

where a component of relative amplitude Δ has been added to the imaginary component, which has the most severe ripples. Δ is positive for a channel in which there is a peak in the imaginary-component ripple. To determine the effect of the term $-j\Delta\sin(2\pi u\ell_1)$ in the image, we take its FT of with respect to u, which is $\Delta[\delta(u+\ell_1) - \delta(u-\ell_1)]/2$. Thus the ripple adds to the image a delta function of amplitude $\Delta/2$ at $-\ell_1$, which is the ghost, and subtracts a delta function of the same amplitude from the true image[2] at ℓ_1. For a source at the field center the ghost and the true image combine, providing a correct measure of the source intensity.

Since the visibility data are usually not calibrated prior to the spectral filtering, the relative amplitudes of the real and imaginary components in Eq. (1) result from the instrumental phases introduced by the receiving system, as well as from the structure of the source. If these instrumental phase data are lost after calibration of the visibility, which is commonly the case, precise removal of the ghost is not possible. However, the effect of the ripples can be reduced by use of smoothing functions on the spectral data before creating the image, as discussed on p. 406 (TMS2). If the spectral data are averaged to provide a continuum result before assigning (u, v) values, the effect of the frequency difference of the channels with high amplitude ripples at the two edges of the passband may be sufficient to separate the ghost into two components, as shown by Bos (1985). This separation will not occur if the (u, v) values are individually assigned for each spectral channel.

Bos (1984) points out that the ghost can be removed, or substantially attenuated, by $\pi/2$ switching of the relative phase between each signal pair before cross correlating, and restoring the phase before transformation of the visibility data to form a map. For the source considered in Eq. (1), the introduction of $\pi/2$ into the differential phase for an antenna pair results in the visibility

$$\mathcal{V}_2(u) = je^{-j2\pi u\ell_1} = j\cos(2\pi u\ell_1) + \sin(2\pi u\ell_1). \tag{3}$$

[1]In the Fourier transformations used here we follow Bracewell (1965) who, for the delta (impulse) function, defines a "transform in the limit" by considering two Gaussian function, $|a|e^{-\pi a^2\ell^2}$ and $e^{-\pi(u/a)^2}$ that are a Fourier pair. As $a \to \infty$ the first Gaussian tends toward a delta function at the ℓ-origin and the second toward unity. For a delta function at ℓ_1 we use the shift theorem and multiply by $e^{-2\pi u\ell_1}$.

[2]Bos (1984, 1985) refers to the ripple-induced component at the true image position as the "hidden component".

The imaginary part consists of the cosine components, which are the real part in Eq. (1). Adding the visibility term resulting from the ripples in the imaginary part of the spectrum, as in Eq. (2), we have

$$\mathcal{V}_2(u) = je^{-j2\pi u\ell_1} = j(1+\Delta)\cos(2\pi u\ell_1) + \sin(2\pi u\ell_1). \tag{4}$$

To remove the effect of the quadrature phase switch we multiply Eq. (3) by j. The visibility term introduced by the ripple then becomes $-\Delta\cos(2\pi u\ell_1)$, and taking the FT with respect to u we find that the contribution of the ripple to the image is $-\Delta[\delta(u+\ell_1)+\delta(u-\ell_1)]/2$. Again there are delta functions at $\pm\ell_1$, but in this case they both have the same sign. Thus the result of averaging the images with the two positions of the phase switch is to cancel the ghost but double the amplitude loss of the true image. Note that we have assumed that the quadrature phase shift introduced by the switch can be represented by the factor j in Eq. (3): if the sign of the phase shift is such that the factor is $-j$, then the sign of the right-hand side of Eq. (3) must be reversed. If the sign is wrong, the effect is to double the amplitude of the ghost but restore the amplitude of the image. Bos (1984) notes that in using $\pi/2$ phase switching on the Westerbork Array, the suppression decreases as the distance from the field center increases. This behavior is related to the use of rather slow switching, with intervals equal to the data-averaging time of 10 s, since for any antenna pair the rate of change of the measured visibility with time increases with distance of a source from the field center.

Most large synthesis arrays do not incorporate phase switching to mitigate ghosts. As larger and more sensitive instruments are being planned it is necessary to consider whether such phase switching should be included. Deep surveys will require removal of strong sources and any ghosts to allow accurate statistics of the more distant, faint sources. Spectral correlators will be used to allow excision of interference, so the ghost mechanism is likely to be present. Switching of the relative phase for each antenna pair in an array with a large number of antennas requires the use of Walsh functions or some similar scheme, as described in Section 7.5. Switching by π rad is usually required to reduce spurious responses, etc. and switching by $\pi/2$ may be required for sideband separation (TMS2 pp. 181-2). Each application of phase switching usually requires an independent set of switching functions. The longest switching interval is determined by the necessity for each set of Walsh functions to complete one cycle within the data averaging time, and the shortest interval by the requirement that it be long enough compared with the transition time of the function being switched that there is no significant loss of observing time. Thus with large numbers of antennas it can be difficult to accommodate all the switching functions desired.

References

Bracewell, R. N., *The Fourier Transform and its Applications*, McGraw-Hill, New York, 2000 (and earlier editions).

Bos, A., On Ghost Source Mechanisms in Spectral Line Synthesis Observations with Digital Spectrometers, in *Indirect Imaging*, J. A. Roberts, Ed., Cambridge Univ. Press, 1984, pp. 239-243

Bos, A., *On Instrumental Effects in Spectral Line Synthesis Observations*, Ph. D. Thesis, Univ. of Groningen,1985, see Section 10.

Interferometer Observation of the CMB Fine Structure and Polarization

(May 11, 2006, Jan. 12, 2007)
(Relates to Ch. 10)

1. Introduction

Observation of the structure of the Cosmic Microwave Background (CMB) radiation[1] is briefly discussed in Chapter 10, section 10.4 of TMS2 (2^{nd} ed.). In view of the importance of the results subsequently obtained, this subject requires further examination. The observations with the COBE and WMAP (Bennet et al. 2003a) satellites have shown that the fine structure of the CMB, resulting from the structure at the surface of last scattering, is at a level of order 10^{-5} of the mean temperature (2.725 K). These data, using total-power beam-switching techniques revealed the major peak in the angular spectrum of the background fluctuations at at $\sim 1.6^\circ$. Interferometry offers advantages for the study of the higher resolution peaks which, like the major peak, are attributed to acoustic waves in the early photon-baryon plasma. Since interferometers do not respond to uncorrelated signals, such as those generated within the Earth's atmosphere, and gain variations are more easily calibrated than in the case of total-power measurements, it is possible to use ground-based interferometers for the finer angular structure. Early attempts to observe the structure using general-purpose synthesis arrays provided only upper limits (see, e.g., Subrahmanyan et al. 1993), because the structure is very faint on the arcminute angular scale provided by the shortest antenna spacings of such large instruments. As a result, a number of special instruments have been developed specifically to cover structure of angular range $\sim 3^\circ$ to $\sim -0.1^\circ$. To illustrate these we consider two examples, the Degree Angular Scale Interferometer (DASI) (Leitch et al. 2002a; Halverson et al. 2002; Pryke et al. 2002), and Cosmic Background Imager (CBI) (Pearson et al. 2000; Padin et al. 2002, 2005). Both of these systems are located at sites with low atmospheric absorption: DASI at the South Pole and CBI at the Llano de Chajnantor in Chile.

In the study of the fluctuations in the CMB, it is the statistics of the temperature variations rather than images of specific fields on the sky that are of greatest interest for comparison with theoretical models. Model power spectra are usually given in terms of spherical harmonics, that is, the amplitudes of multipole moments of the temperature variation. Measurements of the angular spectrum of the CMB in this form can be derived directly from the Fourier components measured by interferometry without forming images of the structure on the sky. It is assumed that the CMB spectrum can be expressed as a function with circularly symmetry (rotational inveriance), that is, there is no preferred direction in the structure on the sky. Thus, characteristics of the CMB lead to some design considerations that differ from those for general-purpose synthesis arrays. The individual antennas need to be large enough to allow accurate phase and amplitude calibration with observing times of a few minutes, using strong discrete sources[2]. With regard to the antenna configuration, the main requirement is to obtain sampling in a radial coordinate in the (u, v) plane, rather than uniform sampling in two dimensions as required for imaging. To obtain sufficiently fine sampling in $\vec{u}$, the antennas in DASI and CBI are configured so that, considered pair-wise, the spacing between centers from the closest to the most widely spaced increases in increments that are smaller than the diameter of an antenna. This can be achieved, for example, by the curved arm configuration shown for CBI in Fig. 5.24 of TMS2. For the angular structure of scale $\sim 1^\circ$ and finer, it is sufficient to observe fields no more than a few degrees wide, and the data can then be analysed using a "flat sky" (small field) approximation as discussed for imaging applications in Section 3.1 of TMS.

In CMB measurements it is also essential to be able to separate out the effects of all foreground sources, i.e. those that lie in front of the surface of last scattering. Strong sources within the CMB fields can be located using existing catalogs or special surveys and subtracted from the observations. Details of foreground sources removed in the WMAP mission are given by Bennett et al. (2003b). Foreground signals can be identified

[1] For a general discussion of CMB theory see, e.g. Lachieze-Rey and Gunzig (1999).

[2] For CBI Jupiter and Taurus A were used, and for DASI, with it's south-polar location, the compact galactic H_{II} region PKS J0859-4731, the Carina Nebula, and Centaurus A.

by their spectral characteristics, which for synchrotron or optically-thin thermal emissions, differ from the black-body spectrum of the CMB. Thus another requirement for CMB interferometry is sufficient frequency coverage to allow the spectral characteristics of signals to be determined. Both DASI and CBI use 10 GHz-wide receiving bandwidths of 26-36 GHz, subdivided into channels 1-GHz wide. These frequencies are chosen to be high enough to take advantage of the increasing of CMB flux density with frequency and also to avoid H_2O and O_2 atmospheric absorption lines.

DASI is designed to provide measurements over a range of multipole moments $\ell = 100 - 900$ and uses 13 antenna of diameter 20 cm with baselines 0.25-1.21 m. For CBI the range of ℓ is $400 - 4250$, and 13 antennas of diameter 90 cm with a range of baselines 1-5.51 m are used. Each array is small enough to allow the antennas to be mounted on a mechanically rigid faceplate, which can be pointed in azimuth and altitude so that the normal to the faceplate the tracks the center of the field under observation. The faceplate can also be rotated about its axis, which allows control of the parallactic angle of the interferometer fringe patterns on the sky. The faceplate tracking results in significant simplification of the instruments since no delay system or fringe rotation are needed, although both instruments include phase switching to remove instrumental offsets. The antennas are arranged in patterns with three-fold symmetry and thus a rotation of the faceplate through 120° causes the configuration of the antennas to repeat relative to the sky (see Fig. 5.24 of TMS2). This property is very useful since the response to the sky remains unchanged after such a rotation, and variations in the signals resulting from unwanted effects such as residual cross-talk between antennas can be identified and removed. When imaging is required, rotation of the plane through angles smaller than 120° provides increased density of the visibility coverage in the (u, v) plane. For DASI the polar location causes the parallactic angle to remain constant during tracking of a point on the sky, without rotation of the faceplate.

A further problem at the high levels of sensitivity required to observe the CMB structure results from thermal radiation from the ground and nearby objects, incident through the antenna sidelobes This can introduce a serious unwanted contribution in the responses of the more closely spaced antenna pairs, but the effect decreases with increasing antenna spacing and is much less serious for the longest spacings of DASI and CBI. For CBI the ground signal, as a function of azimuth and elevation, was found to be sufficiently stable to be removable by differencing the visibility measurements. Visibilities were measured for two fields at the same declination, but separated by 8 min of right ascension, by alternating the pointing on an 8-minute timescale (Padin et al. 2005). For DASI a similar principle was used, facilitated by the constant-elevation tracking at the polar location (Leitch et al. 2002a). The the effect was further mitigated by providing a conducting ground shield in the vicinity of the instrument so that sidelobe responses were reflected up toward the relatively cold sky (Leitch et al. 2002b).

2. Analysis

The analysis that follows is based largely on Hobson et al. (1995) and White et al. (1999a). Following the convention used by most authors in discussion of CMB observations, some vector notation is used. Angular frequency, [(u, v) in TMS, see Fig. 3.2] is usually referred to as wavenumber[3]or wave-vector in the CMB literature, and is represented here by a 2-vector ($\vec{u}$) in the (u, v) plane[4]. For position on the sky, $\hat{x}$ is a unit 3-vector[5]. The CMB structure is quantified in term of $\Delta T/T_{CMB}$ where T_{CMB} is the mean brightness temperature of the CMB and ΔT is the amount by which the temperature at any point differs from T_{CMB}. It has become usual to express theoretical and observational CMB data in terms of spherical harmonics. Thus in any direction $\hat{x}$ the structure can be denoted by:

$$\frac{\Delta T}{T_{CMB}}(\hat{x}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{x}), \tag{1}$$

[3]Wavenumber is defined as the reciprocal of wavelength and as applied to CMB structure has the dimensions of angular frequency.

[4]In general one would use a 3-vector to represent components in (u, v, w), but here we are considering antennas mounted on a tracking faceplate for which $w = 0$.

[5]Here a circumflex (ˆ) indicates a unit vector (other vectors are indicated by an arrow (¯).

where $a_{\ell m}$ expresses the amplitude of the spherical harmonic $Y_{\ell m}$ in direction $\hat{x}$. ℓ and m are related to the harmonics in the azimuthal and polar directions represented by $\hat{x}$. Note that multipole $\ell = 0$ represents the mean level of the CMB temperature over the sky, and $\ell = 1$ represents the dipole moment which results largely from the motion of the Earth relative to the distant background. Thus the summation starts with $\ell = 2$ since the quadrupole is the lowest moment entirely attributable to the CMB structure. ℓ and m are integers, and for any value of ℓ there are $2\ell + 1$ harmonics corresponding to $-\ell \leq m \leq \ell$. The coefficients $a_{\ell m}$ are:

$$a_{\ell m} = \int_{4\pi} \frac{\Delta T}{T_{CMB}}(\hat{x}) Y_{\ell m}(\hat{x}) d\Omega. \tag{2}$$

For a given value of ℓ, the sum over m of products of the harmonics for two directions is

$$\sum_m Y_{\ell m}(\hat{x}_1) Y_{\ell m}(\hat{x}_2) = \frac{2\ell + 1}{4\pi} P_\ell(\cos\theta), \tag{3}$$

where P_ℓ is the Legendre polynomial of order ℓ and $\cos\theta = \hat{x}_1 \cdot \hat{x}_2$. With Gaussian statistics, which we believe represent the CMB structure, the coefficients of the spherical harmonics are uncorrelated, and thus the power represented by harmonics of order ℓ can be written as:

$$C_\ell = a_\ell^2 = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2 = \langle |a_{\ell m}|^2 \rangle. \tag{4}$$

For the CMB temperature structure we need to consider correlation products as follows:

$$C(\hat{x}_1 \cdot \hat{x}_2) = \left\langle \frac{\Delta T}{T_{CMB}}(\hat{x}_1) \frac{\Delta T}{T_{CMB}}(\hat{x}_2). \right\rangle \tag{5}$$

Since there is no preferred direction in the background structure, the function in Eq. (5) is assumed to be circularly symmetrical and can be represented by a a single radial variable θ. $(\hat{x}_1 \cdot \hat{x}_2)$ is equal to the cosine of θ, the angle between directions $\hat{x}_1$ and $\hat{x}_2$. The Fourier transform of the function in Eq. (5) is described as representing the power spectrum of the CMB structure on the sky [here "power" refers to the *mean squared* temperature variations, and should not be confused with the physical dimension of the correlations which is (voltage)^4 = (power)^2]. In the form of a covariance matrix, these data are the cross correlation of the visibility measured in direction $\hat{x}_i$ with baseline $\vec{u}_i$ and that in direction $\hat{x}_j$ with baseline $\vec{u}_j$. The matrix has n_a^4 elements, where n_a is the number of antennas (n_a antennas provide n_a^2 visibility measurements, including zero-baseline and complex conjugate values). As a result of the rotational invariance of the CMB structure, the non-diagonal elements are zero: the diagonal elements represent cross correlation of $(\Delta T/T_{CMB})$ values made in different directions but with the same baseline $\vec{u}$. This diagonal feature applies for the true structure of the CMB, but in measurements of the structure the finite size of the antenna apertures can introduce cross correlation between visibilities measured with different antenna pairs, as discussed later with respect to the "window function".

We denote the CMB power spectrum [the Fourier transform of $C(\theta)$] by $S(|\vec{u}|)$. Both $C(\theta)$ and $S(|\vec{u}|)$ are circularly symmetrical. The two-dimensional Fourier transforms of functions that can be represented by a single radial coordinate are related by the zero-order Hankel transform [see Appendix 1] and

$$S(|\vec{u}|) = 2\pi \int_0^\infty C(\theta) J_o(2\pi |\vec{u}| \theta) \theta d\theta. \tag{6}$$

$C(\theta)$ can be expanded as a series of values C_ℓ of the squared amplitudes of spherical harmonics corresponding to multipole moments ℓ,

$$C(\theta) = \frac{1}{4\pi} \sum_{\ell=2}^{\infty} (2\ell + 1) C_\ell P_\ell(\cos\theta). \tag{7}$$

From Eqs. (6) and (7),

$$S(|\vec{u}|) = \tfrac{1}{2}\sum_{\ell=2}^{\infty}(2\ell+1)C_\ell \int_0^\infty P_\ell(\cos\theta)J_0(2\pi|\vec{u}|\theta)\theta d\theta, \tag{8}$$

and the integration can be done using the following result from Gradshteyn and Ryzhic [1980, Eq. 7.251(3)],

$$\int_0^1 xP_n(1-2x^2)J_0(xy)dx = y^{-1}J_{2n+1}(y) \quad (y>0). \tag{9}$$

In Eq. (9) we put $x = \theta/2$ and with the small-field approximation $(1-2x^2) \simeq \cos(\theta)$. Also the small-field range of θ allows the upper limit of the integral to be extended to ∞. Then Eq. (8) becomes,

$$S(|\vec{u}|) = \frac{1}{4\pi|\vec{u}|}\sum_{\ell=2}^{\infty}(2\ell+1)C_\ell J_{2\ell+1}(4\pi|\vec{u}|). \tag{10}$$

As a measure of the power (mean squared amplitude) of the structure as a function of ℓ, it is usual to use $|\vec{u}|^2 S(|\vec{u}|)$. This is because a factor $|\vec{u}|$ is required to take account of the area $\propto |\vec{u}|\, d|\vec{u}|$ of a circular element in the (u,v)-plane. Also the angular frequencies of the components of the structure are inversely proportional to ℓ, so a second factor $|\vec{u}|$ (which is proportional to ℓ) is used to compensate for the bandwidth variation between successive (integer) values of ℓ.

A simplifying approximation can be made because, for Bessel functions of large integer order, the value of the function remains close to zero as the argument increases from zero to a value close to the order of the function, at which point a sharp peak occurs followed by rapid oscillations about zero that gradually decrease in amplitude as the argument tends to infinity. Thus the values of the terms in the summation in Eq. (10) that make a significant contribution are limited to a small range of ℓ near the Bessel function peak and, noting that $|\vec{u}| = \ell/2\pi$, we can write

$$|\vec{u}|^2 S(|\vec{u}|) \simeq \frac{1}{(2\pi)^2}\ell(2\ell+1)C_\ell \int_0^\infty J_{2\ell+1}(2\ell)d\ell. \tag{11}$$

Then using using $\int_0^\infty J_n(bx)dx = 1/b \;\; (n>-1, b>0)$ [Gradshteyn and Ryzhic, 1980, Eq. 6.511(1)], we obtain

$$|\vec{u}|^2 S(|\vec{u}|) \simeq \frac{\ell(\ell+1)}{(2\pi)^2} C_\ell\Big|_{\ell=2\pi|\vec{u}|}. \tag{12}$$

This result provides a relation between the power spectrum expressed as $S(|\vec{u}|)$ and expressed as multipole moments C_ℓ. It is accurate to a few percent for standard cold dark matter models when $|\vec{u}| \geq$ approx. 10 or $\ell \geq$ approx. 60 (White et al. 1999a).

3. Window Function

A conceptually simple way to estimate C_ℓ from interferometer observations is to average the the square of the measured visibility amplitudes in concentric incremental rings centered on the (u,v)-plane origin. This gives $C(|\vec{u}|)$ convolved with a "window function" and also with a bias resulting from the system noise that can be corrected. The window function results from the sensitivity of antenna pairs to values of u (i.e. the "window" is in the angular frequency domain). To visualize this, note that for antennas of diameter d_λ (measured in units of the wavelength λ), with distance $|\vec{u}|$ between the centers, spacings in the range $u - d_\lambda$ to $u + d_\lambda$ can be found. An interferometer with tracking antennas responds to sky brightness function multiplied by the antenna power beam. Thus the measured visibility is the true visibility function convolved with the Fourier Transform of the antenna power beam. The visibility (measured in units of flux density) is

$$\mathcal{V}(\vec{u}) = \frac{2k\,T_{CMB}\,\nu^2}{c^2}\int \frac{\Delta T(\vec{x})}{T_{CMB}} A(\vec{x}) e^{2\pi\vec{u}\cdot\vec{x}} d\vec{x}, \tag{13}$$

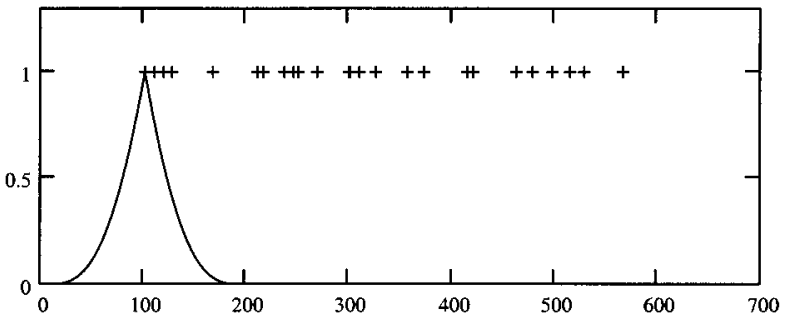

Figure 1: The abscissa values refer to $|\vec{u}|$: multiply by 2π to obtain the corresponding values of ℓ. The curve shows the model window function W_{ii} [Eq. (19)] for the shortest-spacing antenna pair of the CBI configuration in Fig 5.24 of TMS2. The ordinate is the normalized amplitude. The crosses indicate the spacings for the other antenna pairs. The sensitivity of the array to background structure expressed as multipole moments is represented by a set of similar window functions centered on the crosses for the antenna pairs of CBI.

where $A(\vec{x})$ represents the antenna collecting area (i.e. the power beam) and for simplicity we have used the Rayleigh-Jeans flux-density/temperature relation. The Fourier transform of $A(\vec{x})$, $\tilde{A}(\vec{u})$ is the autocorrelation (self-convolution assuming that $\tilde{A}$ is a symmetric function) of the voltage illumination of the antenna aperture which, for a circular aperture, is zero outside a circle of diameter $2d_\lambda$. At this point we introduce the visibility covariance matrix $C^{\mathcal{V}}$ ($\mathcal{V}$ being the visibility). The elements of the visibility covariance matrix are pair-products of the visibility (Hobson et al. 1995, see Appendix A),

$$\begin{aligned} C_{ij}^{\mathcal{V}} &= \langle \mathcal{V}(\vec{u}_i)\mathcal{V}^*(\vec{u}_j)\rangle \\ &= \Big(\frac{2k\,T_{CMB}\,\nu^2}{c^2}\Big)^2 \Big\langle \int \tilde{A}(\vec{y})\Delta T(\vec{u}_i-\vec{y})d^2\vec{y}\int \tilde{A}^*(\vec{z})\Delta T(\vec{u}_j-\vec{z})d^2\vec{z}\Big\rangle \\ &= \Big(\frac{2k\,T_{CMB}\,\nu^2}{c^2}\Big)^2 \int \tilde{A}(\vec{y})d^2\vec{y}\int \tilde{A}^*(\vec{z})\big\langle \Delta T(\vec{u}_i-\vec{y})\Delta T(\vec{u}_j-\vec{z})\big\rangle d^2\vec{z}. \end{aligned} \tag{14}$$

Then with the change of variables $\vec{v}=\vec{u}_i-\vec{y}$ and $\vec{w}=\vec{u}_j-\vec{z}$ [all of which are vectors in the (u,v)-plane],

$$C_{ij}^{\mathcal{V}} = \Big(\frac{2k\,T_{CMB}\,\nu^2}{c^2}\Big)^2 \int \tilde{A}(\vec{u}_i-\vec{v})d^2\vec{v}\int \tilde{A}^*(\vec{u}_j-\vec{w})\big\langle \Delta T(\vec{v})\Delta T^*(\vec{w})\big\rangle d^2\vec{w} \tag{15}$$

The term in angle brackets in Eq. (15) represents the generalized power spectrum of the background intensity fluctuations and can be written as $S(\vec{w})\delta(\vec{v}-\vec{w})$. The delta function allows us to perform the second integration in Eq. (15) and obtain

$$C_{ij}^{\mathcal{V}} = \Big(\frac{2k\,T_{CMB}\,\nu^2}{c^2}\Big)^2 \int \tilde{A}(\vec{u}_i-\vec{w})\tilde{A}^*(\vec{u}_j-\vec{w})S(\vec{w})d^2\vec{w}. \tag{16}$$

In the integral, $\tilde{A}(\vec{u}_i)\tilde{A}^*(\vec{u}_j)$ is convolved with $S(\vec{u})$ using $\vec{w}$ as a variable of integration. In terms of the asterisk notation for two-dimensional convolution used in TMS,

$$C_{ij}^{\mathcal{V}} = \Big(\frac{2k\,T_{CMB}\,\nu^2}{c^2}\Big)^2 \tilde{A}(\vec{u}_i)\tilde{A}^*(\vec{u}_j) ** S(|\vec{u}|). \tag{17}$$

Thus the pair-products of the measured visibilities, which are the elements of the visibility covariance matrix, represent the power spectrum of the fluctuations convolved with the squared modulus of the Fourier transform

of the power beam. The window function W is the part of Eq. (17) that is convolved with the power spectrum $S|\vec{u}|$:

$$W_{ij}(|\vec{u}|) = \tilde{A}(\vec{u}_i)\tilde{A}^*(\vec{u}_j). \tag{18}$$

As a simple model representing the function $\tilde{A}(\vec{u})$, consider the case of a circular antenna aperture that is uniformly illuminated. Then $\tilde{A}(\vec{u})$ is the autocorrelation of a circular pillbox function of diameter d_λ,

$$\tilde{A}(\vec{u}) = \frac{d_\lambda^2}{2}\left(\arccos\frac{|\vec{u}|}{d_\lambda} - \frac{|\vec{u}|\sqrt{d_\lambda^2 - |\vec{u}|^2}}{d_\lambda^2}\right), \qquad |\vec{u}| \le d_\lambda, \tag{19}$$

that is, the area of overlap of two circles of diameter d_λ and spacing between centers $|\vec{u}|$. Note that $\tilde{A}(0) = \pi d_\lambda^2/4$, the area of an antenna aperture. For both DASI and CBI the 13 antennas provide 78 baselines but, as a result of the three-fold symmetry, no more than 26 different baseline magnitudes. The curve in Fig. 1 represents the model window function W_{ii} for the closest spacing in the CBI array, with $d_\lambda = 93$, which coresponds to the diameter of the antennas at 31 GHz. The crosses in the figure indicate the spacings of the antenna centers for the other pairs. The combination of the window functions for all of the pairs, one centered on each of the crosses, represents the full sensitivity of the array in the $\vec{u}$ domain. The overlapping of the functions indicates the continuity of the sensitivity in $\vec{u}$ over the range defined by the antenna spacings. In reality there is some tapering of the illumination of the antennas, which has the effect of decreasing the sharpness of the peak of the function of Eq. (19). Similar diagrams for DASI are given by White et al. (1999a, Fig. 1b; 1999b, Fig. 2).

4. Estimation of the Power Spectrum from the Visibility Measurements

A simple method of obtaining an estimate of the structural power spectrum is to average the observed values of $|\mathcal{V}(u)|^2$ in circular increments of u in the (u, v) plane. This gives estimates of C_ℓ at $\ell = 2\pi u$ that contain a bias resulting from the noise and also contain the effect of convolution with the window function. The noise bias can be removed by subtraction, i.e. taking $|\mathcal{V}(u_j)|^2 - \sigma_{Nj}^2$, where σ_{Nj}^2 is the variance of the noise for the j^{th} visibility value. For testing results of a model against this observed spectrum, the model values can be convolved with the window function to allow a direct comparison. A more sophisticated approach used by White et al. (1999a) involves starting with a model of the power spectrum, essentially an initial guess, and calculating the expected visibility values with the effects of noise and the window function included. The observed visibilities are binned and averaged using, e.g., intervals of $\Delta u = d_\lambda$. A simple initial model can be based on a power spectrum $\ell(\ell+1)\mathcal{C}_\ell$ that is piecewise constant within each of the binning intervals used for the observed visibilities. The likelihood of the observed values, given the initial model, is then calculated and the model parameters varied to obtain a maximum likelihood fit. The likelihood function for the Gaussianly distributed structure ΔT and *complex*[6] values of $\mathcal{V}$ is, (Hobson et al. 1995, White et al. 1995a)

$$\mathcal{L}(\{C_\ell|\}) = \frac{1}{\pi^n |C_\ell|} e^{[-\mathcal{V}^*(\vec{u}_i) C^{-1}_{\ell\, ij} \mathcal{V}(\vec{u}_j)]}\ , \tag{20}$$

where n is the number of bins into which the visibility values are distributed and the matrix C_ℓ, of which the inverse and the determinant are included, represents the model. Both the observed and model visibilities have zero mean, the phases being random, and variance $\langle \mathcal{V}_i^* \mathcal{V}_i \rangle = \langle C_{ii} \rangle$. In matrix form these data are concentrated along the diagonal, but in neighboring elements the signal components of the two visibilities may also be correlated as a result of the convolution with the window function. The term $\mathcal{V}^*(\vec{u}_i) C^{-1}_{\ell\, ij} \mathcal{V}(\vec{u}_j)$ in the exponent is equal to the χ^2 statistic (see, e.g. Bevington and Robinson 1992), and can be used to estimate the quality of fit of the model to the data. The maximum likelihood technique can also be extended to estimation of cosmic parameters that are intrinsic to the assumed model (see, e.g., Pryke et al. 2002),

[6] White et al. (1999b) consider the real and imaginary parts of the visibility (which are uncorrelated) separately, and in effect have twice the number of simple quantities. The expression for the liklehood in that case is proportional to the square root of Eq. (20) (Hobson et al. 1995).

and much detailed analysis has been publishes as the accuracy of the observations has been increased by further measuerements. Observations using the mosaicking technique have been used to extend the width of the observed field and thereby enhance the resolution in ℓ, see, e.g., Pearson et al. (2003). Because the CMB structure is so faint, very large data sets are accumulated to obtain satisfactory signal-to-noise ratios, and various techiques have been considered to reduce such data sets efficiently. For example White et al. (1999a) discuss the use of optimal subspace filtering and Wiener filtering in imaging applications.

5. Polarization

Although the CMB radiation is randomly polarized as emitted, subsequent Thomson scattering from electrons generates a component of linear polarization. The predicted level of polarization is 5-10% of the CMB variations. As a qualitative explanation, following Hu and Dodelson (2002), consider an electron at the origin of the coordinate system in Fig. 2a. Unpolarized radiation incident on the electron from the x and y directions is represented by orthogonal linearly polarized components indicated by the arrows. Scattering of the components of the incident radiation for which the polarization vectors are parallel to the z axis produces no radiation in the z direction. Thus an observer located on the z axis sees scattered radiation from the horizontal component from y and the vertical component from x. . Thus if the intensities of the radiation from the x and y directions are not equal, the observer sees a net linear polarization. Radiation from the directions $-x$ and $-y$ behaves in the same way. Therefore the observed polarization is parallel to the y axis if the intensity from the directions $\pm x$ is greater than that from directions $\pm y$, and parallel to the y axis if the intensities are reversed. In either case the polarization results from a quadrupole moment in the intensity of the radiation incident upon the electron. Note that if the radiation from the x direction is stronger than that from the y direction, the radiation scattered in the direction normal to the z axis (which the observer does not see) is stronger in the direction of increasing y than of increasing x, so the scattered radiation tends to equalize the energy in the CMB.

Figure 2b represents a Fourier component of the CMB structure. The + signs indicate ridges of maximum brightness temperature and the − signs the corresponding minima. The wavevector $\vec{k}$ represents the direction and angular frequency of the component: it can also be envisaged as representing the direction and spacing (in wavelengths) of an antenna pair that responds to that particular Fourier component. In a direction parallel to the wavevector the average brightness temperature corresponds to the mean level of the sinusiodal oscillation. Thus for a scattering electron located at any point on the ridges of maxima, the incident radiation from directions normal to $\vec{k}$ is greater than that from directions parallel to $\vec{k}$. Then the scattered radiation contains a component with linear polarization in a direction parallel to $\vec{k}$. Similarly, for an electron located on a ridge of minima the direction of incidence of the stronger radiation is parallel to $\vec{k}$, resulting in polarization in a direction normal to $\vec{k}$. In CMB theory, polarization is represented by modes E and B. The E mode is a measure of the polarization aligned with the wavevector of the total radiation (positive E), or orthogonal to the wavevector (negative E)[7]. Thus Thomson scattering of the radiation results in an E component that is strongly correlated with the fluctuations in the background temperature. Mode B is a measure of the polarization at $\pm 45^\circ$ to the wavevector (Hu and Dodelson 2002, Fig. 7). It can be generated by scattering associated with gravitational-wave fluctuations.

The following equations give the relation between E and B and Stokes parameters Q and U, in the flat sky approximation (Seljac 1997, Kovac et al. 2005):

$$\begin{aligned} \tilde{Q}(\vec{u}) &= \cos(2\theta_u)\tilde{E}(\vec{u}) - \sin(2\theta_u)\tilde{B}(\vec{u}) \\ \tilde{U}(\vec{u}) &= \sin(2\theta_u)\tilde{E}(\vec{u}) + \cos(2\theta_u)\tilde{B}(\vec{u}), \end{aligned} \tag{21}$$

from which,

$$\tilde{E}(\vec{u}) = \cos(2\theta_u)\tilde{Q}(\vec{u}) + \sin(2\theta_u)\tilde{U}(\vec{u})$$

[7]The E and B vector representation of polarization is such that a 180° rotation in position angle leaves the polarization unchanged. A sign reversal indicates a rotation through 90°. E and B should not to be confused with the electric and magnetic field components, although there is an analogy in the terminology since E and B represent orthogonal Fourier components of the polarized structure, as can be seen from Eqs. (22).

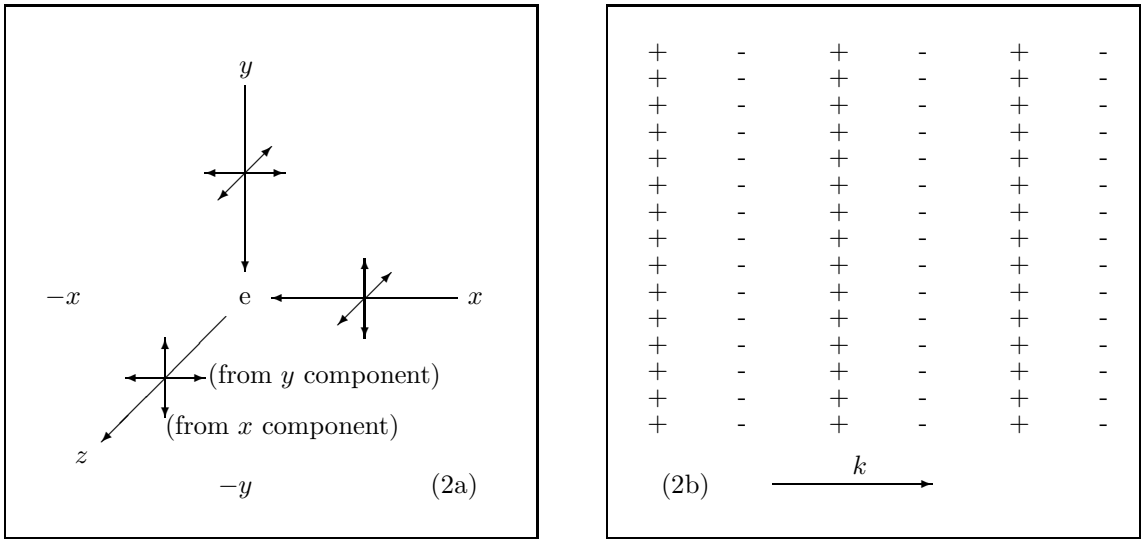

Figure 2: (a) Randomly polarized components of the CMB are incident on an electron (e) from directions x and y. Scattered radiation in direction z contains components of linear polarization from the x and y directions as indicated. If the intensity from the $\pm x$ directions differs from that from the $\pm y$ directions a net linear polarization is seen by an observer in the z direction. (b) Ridges in the background intensity in the (x, y) plane corresponding to a Fourier component with wavevector k. These result in a quadrupole moment in the background intensity and a component of linear polarization in the z direction that is strongly correlated with the intensity variations. See further explanation in the text under Section 5. Polarization.

$$\tilde{B}(\vec{u}) \quad = \quad -\sin(2\theta_u)\tilde{Q}(\vec{u}) + \cos(2\theta_u)\tilde{U}(\vec{u}), \tag{22}$$

where the tilde accent indicates the Fourier transform with respect to $\vec{u}$ and $\hat{x}$ (i.e. the tilde-accented quantities are functions in $\vec{u}$-space) and θ_u is $\arg(\vec{u})$. The antennas of DASI and CBI are circularly polarized, but each antenna receives only one polarization at a time. Cross correlations of signals with the same sense, RR and LL, are sensitive to Stokes parameters $I \pm V$, and combinations with opposite senses, RL and LR, are sensitive to $-jQ \pm U$ (see TMS1 Eqs 4.49 or TMS2 Eqs. 4.35). If T represents the amplitude of the structure (ΔT) non-zero power spectra are to be expected for the mode combinations TT, EE, BB, and TE, but not for TB or BE. Both DASI and CBI have detected E-mode polarization consistent with theoretical expectations (Leitch et el. 2002b; Readhead et al. 2004).

Thanks are due to T. K. Sridharan for correction of an error in an earlier version.

References

Bennett, C. L. and Fourteen Coauthors, The Microwave Anisotropy Probe Mission, Astrophys. J. 583, 1-23, 2003a.

Bennett, C. L. and Fifteen Coauthors, First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Foreground Emission, Astrophys. J. Suppl., 148, 97-117, 2003b.

Bevington, P. R. and D. K. Robinson, Data Reduction and Error Analysis for the Physical Sciences, 2nd ed., McGraw-Hill, New York, 1962.

Bracewell, R. N, The Fourier Transform and its Applications, McGraw-Hill, New York, 1965.

Gradshteyn, I. S. and I. M. Ryzhik, Table of Integrals Series and Products, Academic Press, New York, 1980.

Halverson, N. W. and Twelve Coauthors, Degree Angular Scale Interferometer First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum, Astrophys. J., 568, 38-45, 2002.

Hobson, M. P., A. N. Lazenby and M. Jones, A Batesian Method for Analyzing Interferometer Observations of Cosmic Microwave Background Fluctuations, Mon. Not. R. Astron. Soc., 275, 863-873, 1995.

Hu, W. and S. Dodelson, Cosmic Microwave Background Anisotropies, Ann. Rev. Astron. Astrophys. 40, 171-216, 2002.

Lachieze-Rey, M. and R. Gunzig, the Cosmological Background Radiation, Cambridge Univ. Press, Cambridge, UK, 1999.

Leitch, E. M. and Sixteen Coauthors, Experiment Design and First Season Observations with the Degree Angular Scale Interferometer, Astrophys. J., 568, 28-37, 2002a.

Leitch, E. M., J. M. Kovac, C. Pryke, J. E. Carlstrom, N. W. Halverson, W. L. Holzapfel, M. Dragovan, B. Reddall, E. S. Sandberg, Measurement of Polarization with the Degree Angular Scale Interferometer, Nature, 420, 763, 2002b. (astro-ph/0209476 v1)

Padin, S. and Fifteen Coauthors, The Cosmic Background Imager, Pub. Astron. Soc. Pacific, 114, 83-97, 2002.

Padin, S. and Fourteen Coauthors, First Intrinsic Observations with the Cosmic Background Imager, Astrophys. J., in publication, 2005. (astro-ph/0012211 v2)

Pearson, T. J., B. S. Mason, S. Padin, A. C. S. Readhead, M. C. Shepherd, J. Sievers, P. S. Udomprasert, and J. K. Cartwright, The Cosmic Background Imager, IAU Symp. 201, New Cosmological Data and the Values of the Fundamental Parameters, ed. A. Lazenby and A. Wilkinson, San Francisco: ASP, 2000. (astro-ph/0012212 v1, 2000)

Pearson, T. J. and Twenty-five Coauthors, The Anisotropy of the Microwave Background to $\ell = 3500$: Mosaic Observations with Cosmic Background Imager, Astrophys. J., 591, 556-574, 2003.

Pryke, C., N. W. Halverson, E. M. Leitch, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, and M. Dragoven, Cosmological Parameter Extraction from the First Season of Observations with the Degree Angular Scale Interferometer, Astrophys. J. 568, 46-51, 2002.

Readhead, A. C. S. and Twenty-five Coauthors, Polarization Observations with Cosmic Background Imager, Science, 306, 836-844, 2004. (astro-ph/0409569 v2, 2004)

Seljak, U., Measuring Polarization in the Cosmic Microwave Background, Astrophys. J., 482, 6-16, 1997.

Subrahmanyan, R., R. D. Ekers, M. Sinclair, and J. Silk, A Search for Arcmin-scale Anisotropy in the Cosmic Microwave Background, Mon. Not. R. Astron. Soc. 263, 416-424, 1993.

White, M., J. E. Carlstrom, M. Dragovan, and W. L. Holzapfel, Interferometric Observations of Cosmic Microwave Background Anisotropies, Astrophys. J. 514, 12-24, 1999a (astro-ph 9712195 v2, 1998).

White, M., J. E. Carlstrom, M. Dragovan, and W. L. Holzapfel, Analyzing Data from DASI, astro-ph/9912422v1, 1999b.

Appendix 1, Hankel Transform.

The zero-order Hankel transform can be explained as follows (Bracewell 1965). $C'(\ell, m)$ and $S'(u, v)$ are a Fourier transform pair and are circularly symmetric: $C'(\ell, m) = C(r)$ and $S'(u, v) = S(q)$, where

$\ell^2 + m^2 = r^2$ and $u^2 + v^2 = q^2$. Let $\ell = r\cos(\phi),\ m = r\sin(\phi),\ u = q\cos(\psi),$ and $v = q\sin(\psi)$. Then,

$$\begin{aligned} C(r) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} S'(u,v)e^{-i2\pi(u\ell+vm)}du\,dv \\ &= \int_{0}^{\infty}\int_{0}^{2\pi} S'(u,v)e^{-i2\pi qr\cos(\phi-\psi)r}dq\,d\phi \\ &= \int_{0}^{\infty} S(u)\Big[\int_{0}^{2\pi} e^{-i2\pi qr\cos(\phi)d\phi}\Big]rdr \\ &= 2\pi\int_{0}^{\infty} S(u)J_0(2\pi qr)rdr, \end{aligned} \tag{23}$$

since $J_0(z) = \frac{1}{2\pi}\int_0^{2\pi} e^{-iz\cos\alpha}d\alpha$. For the inverse Fourier transform the sign in the exponent is reversed, but since J_0 is an even function we can write

$$S(u) = 2\pi\int_{0}^{\infty} C(r)J_0(2\pi qr)qdq. \tag{24}$$

Size of the Isoplanatic Patch in the Ionosphere as a Function of Frequency

(May 11, 2006)

(Relates to Ch. 10, p. 401; Ch. 11, p. 460; and Section 13.3)

The isoplanatic patch for the ionosphere (pp. 401, 460, 555) can be defined as an area of sky over which the variation of the excess phase at a given frequency ν is small compared with 2π radians, that is, the excess path length is small compared with the wavelength. The excess phase on any path is proportional to ν^{-1}, which results from the variation of the refractive index of the ionosphere (proportional to ν^{-2}), and the path length. Thus, as frequency is decreased, smaller structures in the ionosphere become important and one can take the patch size as approximately proportional to frequency. This result is given in Table 13.6. To be more precise, the structure of the ionospheric irregularities also affects the frequency dependence, and should be taken into account. The atmospheric irregularities are characterized by the structure function, which gives the variance of the phase difference for two paths separated by a distance d. This is discussed (pp. 534-539) for the case of Kolmogorov turbulence in the neutral atmosphere and results are shown in Table 13.2. For distances up to a few tens of km (i.e. small compared to the thickness of the ionosphere) it is appropriate to consider 3D turbulence. For this case the structure function is proportional to $d^{5/3}$, so the rms phase difference is proportional to $d^{5/6}$. [Note that in Fig. 13.11 β is equal to $(\alpha-2)/2$, and for the 3D regime the measured value of the exponent of d, $\beta_{obs} = 0.85$, is very close to $5/6(= 0.83)$.] In the ionosphere, the frequency variation of the refractive index also introduces a further dependence on ν. As noted above, for a fixed path in the ionosphere the phase is proportional to ν^{-1}. Thus the rms phase difference for paths with separation d is proportional to $(d^{5/6} \times \nu^{-1})$, so for a constant rms phase difference d is proportional to $\nu^{6/5}$. Since the isoplanatic patch is defined in terms of an rms phase difference its dimensions in length or angle are proportional to $\nu^{6/5}$. The simpler frequency dependence in Table 13.6 is often satisfactory in practice, since estimates of the isoplanatic patch size are generally only approximate.

Wide Field Imaging by W-Projection (May 11, 2006, Jan. 12, 2007)

(Relates to Ch. 11, §11.8)

W-projection (Cornwell, Golap, and Bhatnagar, 2003) is a new and more efficient method of handling the problem of non-coplanar baselines. This problem occurs when the width of the synthesized field is sufficiently large that the w term in the exact visibility equation (4.3 in TMS1, 3.7 and 11.25 in TMS2) cannot be neglected. Section 11.8 (TMS2) describes the situation and earlier imaging algorithms that can be used in such cases. Cornwell et al. also give a review of alternate approaches.

In w-projection we start by rewriting the visibility equation as

$$\mathcal{V}(u,v,w) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{A_N(l,m)I(l,m)}{\sqrt{1-l^2-m^2}} G(l,m,w)\exp\left[-j2\pi(ul+vm)\right]dl\,dm, \tag{1}$$

where

$$G(l,m,w) = \exp[-j2\pi w(\sqrt{1-l^2-m^2}-1)]. \tag{2}$$

Note that the w dependence is contained within $G(l,m,w)$, and the other parts of Eq. (1) represent $\mathcal{V}(u,v,w=0)$. If $\mathcal{G}(u,v,w)$ is the Fourier transform of $G(l,m,w)$ with respect to (u,v) and (l,m), Eq. (1) can be written as a two-dimensional convolution in (u,v),

$$\mathcal{V}(u,v,w) = \mathcal{V}(u,v,w=0) * * \mathcal{G}(u,v,w). \tag{3}$$

One can visualize (u,v,w) space with the u and v axes in a horizontal plane with w increasing vertically upward. The measured visibility values are scattered throughout a block of (u,v,w) space of dimensions limited by the longest antenna spacings and the geometry of the observations. Generally the observations are designed to optimize the uniformity of sampling of the visibility in the u and v dimensions, but the sampling in w is usually relatively sparse. The procedure in w-projection is to project the three-dimensional visibility data onto the $(u,v,w=0)$ plane, from which a two-dimensional Fourier transform provides an image in l and m. The $(u,v,w=0)$ plane is parallel to the tangent plane on the celestial sphere at the field center, and thus represents data for which the ray paths from a source at the field center to the corresponding pair of antennas are of equal length. Data for which w is non-zero are those for which the ray paths differ in length by w wavelengths. To use such data to obtain visibility in the $(u,v,w=0)$ plane it is necessary to account for the additional path length to one antenna. In propagating the extra distance in space the radiation from a point is spread by diffraction, so a single (u,v,w) point is spread into a diffraction pattern at $w=0$. This spread of the pattern results from the width of the convolution function $\mathcal{G}(u,v,w)$ in Eq. (3), and is approximately proportional to $|w|$.

If we use the approximation $\sqrt{1-l^2-m^2} \approx 1-(l^2+m^2)/2$, Eq. (2) becomes

$$G(l,m,w) \approx \exp[j\pi w(l^2+m^2)]. \tag{4}$$

Fourier transformation then gives

$$\mathcal{G}(u,v,w) \approx \frac{j}{w}\exp[-j\pi(u^2+v^2)/w]. \tag{5}$$

The visibility $\mathcal{V}(u,v,w)$ is entirely determined by $\mathcal{V}(u,v,w=0)$ through convolution with $\mathcal{G}$. Thus $\mathcal{V}(u,v,w=0)$ contains all the data that are required to provide an accurate image, limited only by the synthesized (dirty) beam. Nothing essential to the image is lost in the transition from three dimensions to two. The same convolution function $\mathcal{G}$ applies to projection in both directions, i.e. from $(u,v,w=0)$ to (u,v,w) and vice versa. Note that the convolving function is different for each (u,v,w) data point.

The imaging procedure is as follows. First the visibility data are gridded in (u, v, w) and then projected onto the $(u, v, w = 0)$ plane. In the projection the data are spread in (u, v) by the convolution[1], and thus regridding in the $(u, v, w = 0)$ plane is required. A two-dimensional Fourier transform then provides the dirty image, from which the dirty beam must then be deconvolved by the CLEAN algorithm or some alternate procedure (see §11.1 in TMS1, §11.2 and 11.3 in TMS2). CLEAN requires numerous transpositions of data between the visibility and image domains. In going from the model image to visibility, a two dimensional transform provides $\mathcal{V}(u, v, w = 0)$, from which projection gives values at (u, v, w) points required for comparison with the observations. For the regridding steps, convolution with a spheroidal or other gridding function (see pp. 327-332 in TMS1 or pp. 392-398 in TMS2) is required. Since convolution is commutative and associative it can be computationally efficient to convolve the spheroidal function with the projection function $\mathcal{G}$ and thus store the combined convolution functions for use with each (u, v, w) grid point. Convolution of $\mathcal{G}$ with the spheroidal function has the additional benefit of damping the behavior of $\mathcal{G}$ as $w \to 0$.

Cornwell et al. (2003) provide details of a simulated example of wide field imaging using w projection. They compare the results with the method of image-plane facets (similar to polyhedron mapping, §11.8 in 2^{nd} ed.), and also uvw-space facets (similar to mosaicking, §11.6 in 2^{nd} ed.) which projects the (u, v) space, rather than image space, onto tangent plane facets. (A paper by Sault, Staveley-Smith, and Brouw (1996), in particular the appendix, is a useful reference on mosaicking and the non-coplanar baseline problem.) Hitherto, the facets methods have been perhaps the most widely used for wide field imaging, and Cornwell et al. conclude that, with regard to computing load, they are roughly competitive with w-projection for images of low dynamic range, but that w-projection is distinctly superior when high sensitivity and dynamic range are required.

References

Cornwell, T. J., Golap, K., and Bhatnagar, S., *W-Projection: A New Algorithm for Non-Coplanar Baselines*, EVLA Memo 67, NRAO, 2003.
(EVLA Memos can be accessed at http://www.aoc.nrao.edu/evla/memolist.shtml)

Cornwell, T. J., Golap, K., and Bhatnagar, S., The Non-Coplanar Baselines Effect in Radio Astronomy. The W Projection Algorithm, *in publication*.

Sault, R. J., Staveley-Smith, L., and Brouw, W. N., An Approach to Interferometric Mosaicing, *Astron. Astrophys. Suppl.*,120, 375-384, 1996.

[1] Cornwell et al. (2003) point out that, interestingly, this spreading in (u, v) shows that in general a single antenna pair responds to a range of spatial frequencies, except when $w = 0$, i.e. when the baseline is normal to the direction of incidence of the wavefront.

Refraction Equations (13.6 and 13.73) (Jan. 12, 2007)

(Relates to Ch. 13, pp. 509 and 529)

Precise modelling of the propagation path length through the atmosphere depends on accurate measurements of the refractivity of the troposphere, i.e., moist air. Rüeger (2002) has studied the available laboratory measurements and the supporting theoretical framework, and has suggested that the following three term equation be adopted for the refractivity of moist air:

$$N = 77.6890\frac{p_d}{T} + 71.2952\frac{p_w}{T} + 375463\frac{p_w}{T^2}, \quad (1)$$

where T is the temperature in K and p_d and p_w are the partial pressures or dry and wet air in mbars, respectively. The dry air term includes the effect of 375 ppm of carbon dioxide. The three coefficients in this equation are slightly different from those quoted in Eqs. (13.6) and (13.73), which were based on the work of Thayer (1974). Although the changes in the two rightmost coefficients are 10 and -0.6 percent, respectively, the net change in the refractivity value is only -0.1 percent for a nominal temperature of 273 K. Rüeger claims that the accuracy of the wet part of the refractivity, as described in the equation above, is 0.2 percent. For this error, the uncertainty in the propagation path for a condition of 1 mm of precipitable water vapor is 12 microns (see TMS2, Eq. 13.19).

The two rightmost terms in the equation are particularly relevant for attempts to estimate propagation path fluctuations due to water vapor through the use of water vapor radiometry (see TMS2, pp. 541-543, and Stirling et al. 2005). However, until better laboratory measurements of both the index of refraction and absorption of water vapor can be made, the best strategy for propagation corrections made from WVR measurements may be to determine the required correlation coefficients empirically.

References

Rüeger, J. M., *Refractive Index Formulae for Radio Waves*, FIG XXII International Congress Washington D.C. April 19-26 2002. (An html version can be found on the web at http://geodesia.ufsc.br/Geodesia-online/index.htm, by searching for "Refractive index formulae".)

Stirling, A., Richer, J., Hills, R., and Lock, A., *Turbulence Simulations of Dry and Wet Phase Fluctuations at Chajnantor. Part 1: The Daytime Convection Boundary Layer.* ALMA Memo 517, NRAO, 2005. (ALMA Memos can be found at http://www.alma.nrao.edu/memos/index.html.)

Thayer, G. D., An Improved Equation for the Radio Refractive Index of Air, *Radio Science*, 9, 803-807, 1974.

New Model for Millimeter/Submillimeter Atmospheric Transmission

(May 11, 2006)

(Relates to Sections 13.1 and 13.2)

Pardo, J. R., J. Cernicharo, and E. Serabyn, Atmospheric Transmission at Microwaves (ATM): An Improved Model for Millimeter/Submillimeter Applications, *IEEE Trans. Antennas. Propagat.*, **49**, 1683-1694, 2001.

Near Field and Far Field Distances (May 11, 2006)

(Relates to Ch. 14, p. 601)

Derivation of formulas using Mathmatica can be found in:

Cornwell, T., Approximate formulas for the distance term in far and near field interferometry, EVLA Memo 75, NRAO, April 2004.

EVLA Memos can be accessed at http://www.aoc.nrao.edu/evla/memolist.shtml

Mitigation of Radio Frequency Interference (June 26, 2007)

(Relates to Chapter 15)

Radio frequency interference (RFI) has caused loss of radio astronomy data since the early days of the science, and over time has become increasingly serious. The problem is exacerbated by the necessity to observe outside the narrow bands reserved for passive scientific use, in order to observe spectral lines over wide redshift ranges and to use wider bandwidths to obtain greater sensitivity in continuum observations. Numerous approaches to interference mitigation have been suggested, reflecting the wide variety of man-made radio transmissions and radio astronomy techniques. Only a few of these are in routine use (as of 2007) but further development of mitigation techniques cannot be avoided as new arrays and improved receiving systems are brought into use in the coming decade. Descriptions of many of these approaches can be found in a special issue of Radio Science (2005). They include: (1) those that simply delete receiver output data that are believed to be contaminated by interference, (2) those that cancel or reduce the interference without removing the astronomical data that occurs at the same time and frequency, and (3) those that involve spatial filtering in which a null is generated in the antenna reception pattern in the direction of an interferer. This note is intended only to outline the main principles involved in RFI mitigation.

Detection and Data Deletion

A basic problem is the identification of the contaminated data. In the simplest case this is a matter of examining the output of a correlator or detector and deleting data in which the signal amplitude is larger than expected, or does not vary with time or antenna pointing in the manner expected for astronomical sources. In earlier times this usually meant deleting the full bandwidth received, but more recent radio astronomy instruments use spectral processing of the data, and only those frequency channels that show evidence of interference require deletion. Use of channel bandwidths comparable to the bandwidth of an interfering transmitter also has the advantage of maximizing the interference-to-noise ratio within the channel and thus improving the detectability. The greatest difficulty is the detection of weak interference. When inspecting the data for variations that indicate the presence of interference, data averaging times of seconds or minutes are often appropriate when the interference varies on similar time scales. However, the astronomical measurements may require averaging of data over periods of hours to obtain the required sensitivity, so interference at levels that can introduce errors in the data may be too weak relative to the noise to be easily detected. With the high data output rates produced by most synthesis arrays it is impractical to examine all of the data manually and algorithms by which contaminated data can be flagged by computer are important. Some examples of techniques to detect the presence of interference are as follows.

(1) Use of monitoring receivers with antennas pointed toward likely sources of interference, such as toward the horizon for terrestrial transmitters (e.g. Rogers et al. 2005).

(2) Comparison of data taken simultaneously at two observatories that are sufficiently widely separated that interference from any transmitter is unlikely to be received at both. This has been used in the search for pulses and other transient astronomical emissions (e.g. Bhat et al. 2005).

(3) Detection of transmissions with cyclostationary characteristics, that is, transmission in which some characteristic repeats at intervals τ_c in time. Examples are the frame cycles in TV signals and repeated data cycles in GPS transmissions. Values of τ_c can be determined for the expected signal environment and the occurrence of components of the data with cyclostationary characteristics can be investigated by performing an autocorrelation and looking features that repeat at intervals τ_c. Bretteil and Weber (2005) find that searching for a Fourier component at frequency $1/\tau_c$ in the data is an advantageous method.

(4) Use of the closure relationships [Eq. (11.20) in TMS1, Eq. (10.32) in TMS2] can provide an indication of interference. If observations are made of a point source in the absence of interference, the visibility ratios on the right hand sides of the equation are equal to unity and the values of r_{ij} on the left and side are proportional to the corresponding main-beam gains. A signal from an interfering source will add a component to the output, the magnitude of which will depend upon the sidelobe gains of the antennas which will probably not be matched and will vary with time as the antennas track. Thus variation of the closure gain values is an indication of possible interference. Note, however, that the correction for the fringe frequency of the target source will cause variation of the same frequency in the response to interference from a stationary transmitter, so the effect of the interference will decrease with increase of the baseline component u (see *Fringe-Frequency Averaging* in Ch. 14 of TMS1 and Ch. 15 of TMS2).

(5) In the case of interference from radar pulses, the individual pulses are often strong enough to be seen by looking at the output of a detector, especially if the transmitter is close enough that a direct signal path exists. It may be possible to determine the timing pattern of the pulses and generate blanking pulses for the radio astronomy receivers. The situation may be such that it is necessary to extend the blanking to include reflections from nearby aircraft, etc. See, e.g., the discussion by Dong et al. (2005). A buffer memory for the data allows blanking to begin just before the pulse is detected, to ensure effective blanking.

(6) Examination of the kurtosis of the receiver output data. Kurtosis is defined as μ_4/μ_2^2, where μ_2 and μ_4 are the mean values of the second and fourth powers of the data with respect to the mean value. For Gaussian noise kurtosis has a value of 3.0 and higher values are an indication of non-Gaussian data.

Interference Cancelation

Mitigation of interference by cancelation, leaving the astronomical data intact, is clearly preferable to simple deletion when possible. Cancelation requires not only detection but also an accurate estimation of the signal in order remove it. In adaptive cancelation (see, e.g. Barnbaum and Bradley 1998), a separate antenna (usually smaller than the astronomy antenna) is pointed toward the interferer. The received signal from this antenna is digitized, passed through an adaptive filter, and combined with the signal from the astronomy antenna. The combined outputs are processed by an algorithm that provides a control of the adaptive filter in such a way as to cause the interfering signal voltages from the two antennas to cancel one another. Of various algorithms that could be used to control the adaptive filter, Barnbaum and Bradley used a least-mean-squares algorithm which is computationally simple and thus easily adapted to follow the relative variation of the astronomical and interfering signals as the astronomy antenna tracks. All of this takes place before the signals reach the correlator or detector. Briggs and Kocz (2005) give an example of an interference cancelation scheme in which the outputs of the astronomy and interference antennas are cross-correlated to provide a control for the adaptive filter. For a detailed discussion of methods involving cross correlation of astronomy antenna outputs with those of axillary antennas directed toward the source of interference see Briggs, Bell, and Kesteven (2000). In some cases where the structure of the interfering signal is known in detail, as in the case of the GLONASS navigational satellite signals, it is possible to recreate the interfering signal from the interference received in the astronomy antenna with sufficient accuracy for cancelation. In an example discussed by Ellingson, Bunton, and Bell (2001) interference from GLONASS was reduced by 20 dB.

Spatial Nulling

Spatial nulling involves using a group of antennas in which a null in the combined spatial response is formed in the direction of the source of interference. In low frequency arrays in which the individual receiving

elements are dipoles with beams covering much of the sky, nulling may also result in loss of astronomical sky coverage. If high gain antennas are used and the nulls are formed to reject interference entering the far sidelobes, the astronomical signals in the main beam are essentially unaffected, so the result is similar to cancelation of interference. Nulling is discussed in more detail in the note that follows, see *Spatial Filtering for Attenuation of Interfering Signals*.

References

Barnbaum, C. and R. F. Bradley, A New Approach to Interference Excision in Radio Astronomy: Real-Time Adaptive Cancellation, Astron. J., 115, 2598-2614, 1998.

Bhat, N. D. R., J. M. Cordes, S. Chatterjee, and T. J. W. Lazio, Radio Frequency Interference Identification and Mitigation, Using Simultaneous Dual-Station Observations, Radio Science, 40, (DOI 10.1029/2004RS003172), 2005.

Bretteil, S. and R. Weber, Comparison of two Cyclostationary Detectors for Radio Frequency Interference Mitigation in Radio Astronomy, Radio Science, 40, (DOI 10.1029/2004RS003124), 2005.

Briggs, F. H., J. F. Bell, and M. J. Kesteven, Removing Radio Interference from Contaminated Astronomical Spectra Using an Independent Reference Signal and Closure Relations, Astron. J., 120, 3351-3361, 2000.

Briggs, F. H., and J. Kocz, Overview of Approaches to Radio Frequency Interference Mitigation, Radio Science, 40, (DOI 10.1029/2004RS003160), 2005.

Dong, W., B. D. Jeffs., and J. R. Fisher, Radar Interference Blanking in Radio Astronomy Using a Kalman Tracker, Radio Science, 40, (DOI 10.1029/2004RS003130), 2005.

Ellingson, S. W., J. D. Bunton, and J. F. Bell, Removal of the GLONASS C/A Signal from OH Spectral Line Observations Using a Parametric Modeling Technique, Astron. J. Suppl., 135, 87-93, 2001

Ellingson, S. W. and G. A. Hampson, A Subspace-Tracking Approach to Interference Nulling for Phased Array-Based Radio Telescopes, IEEE Trans. Antennas Propagat. 50, 25-30, 2002

Rogers, A. E. E., P. Pratap, J. C. Carter, and M. A. Diaz, Radio Frequency Interference Shielding and Mitigation Techniques for a Sensitive Search for the 327 MHz Line of Deuterium, Radio Science, 40, (DOI 10.1029/2004RS003157), 2005.

Spatial Filtering for Attenuation of Interfering Signals (May 11, 2006, June 26, 2007)

(Relates to Ch. 15.)

1. Introduction

One way in which interfering signals in radio astronomy can be eliminated or mitigated is by spatial filtering (nulling), i.e. generating nulls in the spatial response pattern in the directions of incidence of interference. This technique is generally well suited for use with synthesis arrays. There are two basic ways in which a null can be generated. In *deterministic* nulling the direction of the interferer is known and a null is formed in that direction by weighting the signals received. Weighting factors (in amplitude and phase) can be applied to the signals from individual antennas if they are being combined as in a phased array, or to the correlated products from antenna pairs before they are combined to form an image. It is not necessary to be able to identify the interference 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 the antenna responses, which may not be practicable. Deterministic nulling can be applied to a synthesis array in two ways. First, the nulls can be formed by adjusting the weights with which the cross-products of the outputs of pairs of antennas (the

visibility values) are combined. In this case the nulls are formed in the synthesized beam pattern, i.e. most likely in the sidelobes of the synthesized beam unless the direction of the interferer is within the synthesized field. Second, in the case of synthesis arrays in which the elements between which cross-correlations are formed are themselves phased subarrays of antennas, the nulls can be formed in the subarray beams. In this case the weighting is applied directly to the signals from the individual antennas. In *adaptive* nulling the interference components within the received signals are identified analytically and the cross products of signals from pairs of antennas are combined in a way that causes the interference vectors to be projected out. The direction of the interfering sources on the sky need not be known, but the interference must be strong enough to be detectable within the noise. Use of nulling requires the development of extensive software, and the intent here is only to explain the principles involved. The following discussion is largely based on Leshem, van der Veen, and Boonstra (2000).

2. Deterministic nulling

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. 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, i.e. $\epsilon\sqrt{p}$, 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 forming a null in a given direction, suppose that one starts by determining the complex gain factors 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 further complex gain factor. If this second gain factor is the same for each antenna, the size and orientation of the vector loop may 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, as is likely to be the case for signals received through far sidelobes, the loop will not close unless the individual gain factors are known and taken into account. Calibrating the far sidelobes of a high-gain reflector antenna over a large fraction of 4π steradians, and also as a function of frequency across receiver bandwidth, would be a task of problematic practicality, which is likely to limit the usefulness of deterministic nulling in such cases. Deterministic nulling is discussed by Smolders and Hampson (2002), Ellingson and Hampson (2002), and Ellingson and Cazemier (2003).

3. Adaptive nulls in the synthesized beam

Let $\mathbf{R}$ be the $n \times n$ covariance matrix of the cross correlations of the signals received by n antennas, as discussed in Supplementary Notes under *Response of an Array in Matrix Format*. In the absence of interference the covariance contains only a component $\mathbf{R}_v$ which represents the cross correlation of the

cosmic signals, i.e. the desired visibility data, and a component of Gaussian noise $\sigma^2\mathbf{I}$. The noise variance σ^2 is known from calibration, and it is generally assumed that $\mathbf{R}_v << \sigma^2\mathbf{I}$. If, in addition, an interferer with spatial signature $\mathbf{a}$ and power σ_s^2, is present. Then

$$\mathbf{R} = \mathbf{R}_v + \sigma^2\mathbf{I} + \sigma_s^2\mathbf{a}\mathbf{a}^H, \tag{1}$$

where the superscript $(^H)$ indicated the Hermitian transpose (the complex conjugate of the transpose). Both the interference and the visibility values vary with time. Typically the interference is stationary over periods of ∼10-100 ms, and when it is removed the resulting data can be averaged for ∼1-10 s, as limited by the stationarity of the visibility, which depends upon the variation of the projected baselines resulting from earth rotation and possible astronomical considerations. At this point we take the terms in Eq. (1) to represent data that are time averaged only for periods of interference stationarity. To remove the interference we can project it out by using a filtering function $\mathbf{P}$ that is orthogonal to the dimensions of the interference:

$$\mathbf{P} = \mathbf{I} - \mathbf{a}(\mathbf{a}^H\mathbf{a})^{-1}\mathbf{a}^H. \tag{2}$$

The spatial signature of the interference $\mathbf{a}$ depends upon the direction of incidence and the response of the array in the corresponding direction, which usually involves the response of the sidelobes of the individual antennas. Thus $\mathbf{a}$ is often not known. However, the spatial signatures of interfering signals can be obtained by eigenanalysis of $\mathbf{R}$, assuming that the astronomical signals are small compared with the noise and that the interference is relatively strong. Since $\mathbf{R}$ is Hermitian we can write:

$$\mathbf{R} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^H, \tag{3}$$

where $\mathbf{U}$ is an $n \times n$ (n = number of antennas) unitary matrix whose rows are the eigenvectors and $\mathbf{\Lambda}$ is an $n \times n$ diagonal matrix that contains the eigenvalues. To identify the interference within the data, it is generally assumed that eigenvectors for which the eigenvalues are greater than the rms noise level represent interference, although more detailed methods are possible[1]. Let q be the number of interferers, which we assume is less than n (i.e. the number of interference sources is less than the number of antennas, which is likely to be the case for large arrays). Then $\mathbf{\Lambda}$ will contain only q eigenvalues greater than σ^2. (In the absence of noise $\mathbf{\Lambda}$ would have rank q.) We introduce subscripts s to indicate interference and n to indicate noise which here includes both the system noise and the astronomical signals. The eigenvalues that are deemed to indicate interference can then be collected in a $q \times q$ diagonal matrix $\mathbf{\Lambda}_s$, and the corresponding eigenvectors in a $n \times q$ matrix $\mathbf{U}_s$. Also the $n - q$ interference-free eigenvectors can be collected in an $n \times (n - q)$ matrix $\mathbf{U}_n$, and then we can write,

$$\mathbf{R} = [\mathbf{U}_s\ \mathbf{U}_n]\begin{bmatrix} \mathbf{\Lambda}_s + \sigma^2\mathbf{I}_q & 0 \\ 0 & \sigma^2\mathbf{I}_{n-q} \end{bmatrix}\begin{bmatrix} \mathbf{U}_s \\ \mathbf{U}_n \end{bmatrix}. \tag{4}$$

Note that σ^2, the variance of the noise in the cross product terms, is assumed to be the same for all antenna pairs.

The interference can be removed by projection as proposed, for example, by Leshem et al. (2000). Since $\mathbf{R}$ is Hermitian, the eigenvectors are orthogonal, and those that represent interference can be projected out. The columns of $\mathbf{U}_n\mathbf{U}_n^H$ span the interference-free subspace occupied by the astronomical signals plus noise. A filtered form of the covariance matrix, from which the interference has been projected out, can then be formed:

$$\mathbf{R}_{F1} = \mathbf{U}_n\mathbf{U}_n^H\,\mathbf{R}\,\mathbf{U}_n\mathbf{U}_n^H = \mathbf{P}\mathbf{R}\mathbf{P}, \tag{5}$$

[1]For example, it may also be possible to distinguish between interference and cosmic signals from the directions of incidence. Each 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 of the main beam. 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 field of the cosmic signals under study will be small.

where $\mathbf{P} = \mathbf{U}_n \mathbf{U}_n^H$ is the projection (filtering) matrix. $\mathbf{R}_{F1}$ is an $n \times n$ matrix, but the subspace no longer includes that of the interferers and the noise does not remain white[2]. An alternate projection scheme uses $\mathbf{U}_n$ from which the rows of the eigenvector that represent the interference have been removed. The resulting filtered covariance matrix is

$$\mathbf{R}_{F2} = \mathbf{U}_n^H \, \mathbf{R} \, \mathbf{U}_n. \tag{6}$$

$\mathbf{R}_{F2}$ has dimensions $(n-q) \times (n-q)$ and the noise remains white. Also the data in $\mathbf{R}_{F2}$ represents mainly the interference, whereas $\mathbf{R}_{F1}$ also includes eigenvectors that represent noise (including the astronomical signals). In using either $\mathbf{R}_{F1}$ or $\mathbf{R}_{F2}$, components with the spatial signatures of the interference are removed and thus nulls are formed in the array response pattern in the corresponding directions of incidence. The beam patterns formed using $\mathbf{R}_{F1}$ and $\mathbf{R}_{F2}$ should both provide effective nulls, but the additional noise in $\mathbf{R}_{F1}$ can result in more unwanted variation of the beam pattern from one beamforming cycle to the next, as discussed by Ellingson and Hampson (2002). Note also that, since each element of the filtered covariance $\mathbf{R}_F$ (i.e. $\mathbf{R}_{F1}$ or $\mathbf{R}_{F2}$) 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. Two approaches to forming an image are outlined as follows.

We introduce k as the index of the time averages that are limited by interference stationarity, and the hat accent ($\hat{}$) to indicate quantities that have been averaged for the longer time allowed by the stationarity of the visibility data. Then from Eq. (5),

$$\hat{\mathbf{R}}_F = \frac{1}{N} \sum_{k=1}^{N} \mathbf{P}_k \mathbf{R}_k \mathbf{P}_k, \tag{7}$$

where N is the number of periods of interference stationarity within the period of visibility stationarity. To proceed further, use is made of the matrix identity $vec(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A}) \, vec(\mathbf{B})$, where $vec(\ )$ denotes stacking of the columns of a matrix and $\otimes$ indicates the outer (Krönecker) product. This allows the two-sided multiplication in Eq. (7) to be written in single-sided form:

$$vec(\tilde{\mathbf{R}}_F) = \frac{1}{N} \sum_{k=1}^{N} (\mathbf{P}_k^T \otimes \mathbf{P}_k) \, vec(\mathbf{R}_k). \tag{8}$$

The expectation of the right-hand side of Eq. (8) is

$$\left[\frac{1}{N} \sum_{k=1}^{N} (\mathbf{P}_k^T \otimes \mathbf{P}_k) \right] vec(\mathbf{R}_v + \sigma^2 \mathbf{I}) = \mathbf{C} \, vec(\mathbf{R}_v + \sigma^2 \mathbf{I}), \tag{9}$$

where $\mathbf{C}$ denotes the terms within the square brackets. [Note that if we use the projection scheme in Eq. (6), and put $\mathbf{P} = \mathbf{U}_n$, then in the expression for $\mathbf{C}$ (in the square brackets) $\mathbf{P}_k^T$ is replaced by $\mathbf{P}_k^*$, where the asterisk indicates the complex conjugate.] We now apply the inverse of $\mathbf{C}$ to $vec(\tilde{\mathbf{R}}_F)$,

$$\mathbf{R}_v + \sigma^2 \mathbf{I} = unvec[\mathbf{C}^{-1} vec(\mathbf{R}_{F1})]. \tag{10}$$

to obtain an unbiased estimate of $\mathbf{R}_v + \sigma^2 \mathbf{I}$. $\mathbf{C}$ is generally a large matrix, $n^2 \times n^2$ for the first projection scheme and $(n-q)^2 \times n^2$ for the second one, so for large n the inversion of $\mathbf{C}$ can be a significant computing task. Also, since the invertibility of a matrix requires that it be of full rank, $\mathbf{C}$ will only be invertible if the spatial signature of the interference $\mathbf{a}$ is sufficiently variable during the integration periods limited by visibility stationarity. For interference from an aircraft or non-geostationary satellite this is likely to be the case. For a geostationary satellite or terrestrial transmitter the motion of the interferometer fringes as they

[2]The term *white* is frequently used in discussions of this subject and essentially indicates that the variance is known.

track the field under observation plus the effect of atmospheric fluctuations for low elevation signals can be expected to provide some variation. This approach, based on the inversion of $\mathbf{C}$, is described by Raza et al. (2002) and analyzed in more detail by van der Tol and van der Veen (2005).

A second imaging approach follows from the fact that data in $\mathbf{R}_F$ are linearly related to those in $\mathbf{R}$, so if an image is made by Fourier inversion of the filtered covariance $\mathbf{R}_F$, the response to any point-source component of the field can be expressed in terms of a synthesized beam. That is, the synthesized image can be thought of as the true sky brightness scanned by a synthesized beam (i.e. the point-source response or, in CLEAN terminology, the dirty beam). Because the nulls must remain fixed in the directions of the interferers, it is apparent that the point-source response beam must vary as it is applied to different points in the sky, i.e. the dirty beam is spatially varying: see, e.g. Leshem and van der Veen (2000), who show this as follows. Since each element of the filtered covariance matrix $\mathbf{R}_F$ is a linear combination of other elements, each element can be represented by $\sum_j c_{ij}\mathcal{V}(u_j, v_j)$, where c_{ij} are the weights with which each element of the unfiltered covariance enters into the elements of the filtered matrix, and $\mathcal{V}$ is the visibility. Since $\mathbf{R}$ contains n^2 elements, each of which can contribute to each of the n^2 elements of $\mathbf{R_F}$, both i and j run from 1 to n^2. Let $I_D(l, m)$ represent the intensity distribution obtained by Fourier inversion of $\mathbf{R}_F$ (the "dirty" image). Then,

$$I_D(l,m) = \sum_i \sum_j c_{ij}\mathcal{V}(u_j, v_j)e^{2\pi j(u_i l + v_i m)}. \tag{11}$$

In Eq. (11) the visibility $\mathcal{V}(u_j, v_j)$ can be replaced by the Fourier transform of the true intensity distribution I,

$$\begin{aligned} I_D(l,m) &= \sum_i \sum_j c_{ij}\left[\int\int I(l',m')e^{-2\pi j(u_j l' + v_j m')}dl'dm'\right]e^{2\pi j(u_i l + v_i m)} \\ &= \int\int I(l',m')\left[\sum_i \sum_j c_{ij}e^{-2\pi j(u_j l' + v_j m')}e^{2\pi j(u_i l + v_i m)}\right]dl'dm' \qquad (12) \\ &= \int\int I(l',m')\mathcal{B}(l,m,l',m')dl'dm', \qquad (13) \end{aligned}$$

where $\mathcal{B}(l, m, l', m')$ is the dirty beam and is equal to the expression in square brackets in (12). Since $\mathcal{B}$ is a function of (l, m), it is space-varying. Thus the effect of the synthesized beam on the astronomical brightness distribution cannot strictly be described as convolution, but can be envisaged as a similar process with a beam whose shape varies as its center is moved from point to point within the field of view[3]. Since the steps in going from $\mathbf{R}$ to $\mathbf{R}_F$ are known, i.e. they are represented by the c_{ij} factors, the point-source response can be derived for any (l, m) point, and it provides a space-varying dirty beam that can be removed using an adaptation of the CLEAN algorithm. Note also that the maximum of this dirty beam does not necessarily occur at the origin of the coordinates of the beam function. A CLEAN-type process is clearly feasible, as discussed in detail by Leshem and van der Veen (2000). However, the effect of the variation of the (u, v) coordinates of each baseline as the Earth rotates will cause the values of u and v corresponding to each c_{ij} to vary. So for observations longer than the period of stationarity of the visibility, one cannot simply continue to accumulate mean values of c_{ij} to provide a dirty beam. However, the dirty beam functions can be averaged further in the form of gridded values of the point-source response on the sky (l, m) plane. Since the dirty beam can be different for each point in the image, the number of data values required to accumulate the time average of these beam functions is of the order of the square of the number of pixels in the image. Leshem and van der Veen (2000) also show how this process of conventional imaging (i.e. imaging by Fourier transformation) and the use of the CLEAN algorithm can be reformulated in terms of an iterative beamforming procedure, which provides further imaging algorithms.

[3]If the null is formed in the sidelobes of a high-gain antenna it should have less effect on the main beam than in the case of a low frequency dipole in which the null may be in the main lobe.

In conclusion, adaptive spatial filtering offers an important method of RFI mitigation for arrays of antennas, but at the expense of a considerable increase in the computation required for imaging. How this approach can best be adapted for different types of arrays and interfering signals will become more apparent as new instruments covering wider spectral bandwidths are brought into operation.

References

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.

Leshem, A. and A.-J. van der Veen, Radio-Astronomical Imaging in the Presence of Strong Radio Interference, *IEEE Trans. Inform. Theory,* **46**, 1730-1747, 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.

Raza, J., A-J Boonstra, and A-J van der Veen, Spatial Filtering of RF Interference in Radio Astronomy, *IEEE Signal Process. Lett.,* **9**, 64-67, 2002.

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.

van der Tol, S. and A-J. van der Veen, Performance Analysis of Spatial Filtering of RF Interference In Radio Astronomy, *IEEE Trans. Signal Process.,* **53**, 896-910, 2005.

A.R.T., J.M.M., G.W.S.