Proceedings of The Institute of Acoustics

COMPUTER SIMULATION OF ADVANCED SIGNAL PROCESSING ALGORITHMS

Ira J. CLARKE, John L. MATHER

Royal Signals and Radar Establishment, Malvern, England.

INTRODUCTION

Under-determined problems commonly arise in many areas of signal processing due to restrictions on the physical and temporal apertures of practical devices and also on the number and on the bandwidth of sensors available. Within the tolerance or margin of uncertainty set by the noise floor of practical systems and also by inevitable inaccuracies in knowledge of the sensor response characteristics there are generally many valid solutions which 'fit' a given batch of received data. Usually just one of these solutions is required for further processing - most frequently to infer from it a model of the input scenario. Since the data itself does not indicate uniquely the correct model to choose, at least one additional criterion or preference must be applied. Various algorithms are now offered to the system designer and each apply a different constraint, model or assumption (maximum liklihood and maximum entropy are terms which are commonly used). However the underlying reasoning leading to the choice of one particular solution is not always obvious to the user of an algorithm, and consequently may not be related in any way to the real situation. Conceivably collateral knowledge, which could assist the choice of model, could arise in many different ways: for example research findings or characterisations of the scenario may be available off-line, other equipment may be gathering related data or there may be circumstantial evidence to be tested. The merging of data and application of useful predetermined assumptions are of interest in many areas of current research. For this reason we have set up a computer simulation which directly compares the performance of a selection of advanced algorithms and allows us to test basic ideas in this area quickly and easily.

A POINT MODEL USED AS COLLATERAL KNOWLEDGE

There is a class of algorithms, typified by MUSIC [1], which explicitly depend on a 'point target model' to describe a batch of data in terms of a few parameters. These algorithms generally first require the diagonalisation of a covariance matrix, estimated from incoming data. The orthonormal property of eigenvectors is used to separate 'unwanted' noise components from a signal subspace. If the point description is a valid assumption in the prevailing circumstances, then the one solution which the algorithm selects might be expected to be nearer the 'truth' than a relatively diffuse version possibly based on an arbitrary model. We demonstrate later that the MUSIC class of algorithms generate results which appear to defy the generally accepted limits of 'resolution' as postulated by Rayleigh, yet clearly cannnot violate the bounds set by information theory. It is clear that the real limitation must depend both on the signal content relative to the noise power in the data but supported by the knowledge that targets are points. We suggest that the application of a point model illustrates that simple forms of collateral knowledge can be applied effectively. It has already been shown by several workers that the addition of collateral knowledge in the form of a mathematical constraint or support function does modify resolution limits [8].

ADVANCED ALGORITHMS

The interpretation of incoming data into a point model by advanced algorithms can be divided into three distinct stages. The first operation involves filtering the data in some optimum way in order to emphasise the most useful components in the data – the information content. In the second stage the filter output is interpreted as being equally well represented by a few points – each of which is associated with previously characterised impulse response waveforms or vectors which are then regarded as basis components of the data. For example, in spatial processing each waveform/vector – often referred to as a steering vector – would correspond to the response of the sensor system to a 'point source' at a certain 'direction' of arrival. In the third stage, if required, we go back to the original data and use the model resulting from the previous stage. The simplest concept of this stage, for a model comprising n point sources, is that n–1 nulls are steered onto all the 'directions' except one in order to extract the signal from that direction. This is analogous to more conventional adaptive cancellation algorithms but the advantage is that the steering vectors for all the required signal direction are already available. In practice we directly decompose the data using the steering vectors as basis vectors, a straightforward mathematical operation requiring a pseudo inverse [1,6]. The two latter stages are common to all algorithms in our simulation, while for the first stage we have developed a generalised filter routine. This filter can easily be adapted to implement any one of a number of algorithms.

Assuming a previous diagonalisation of a covariance estimate and using a column notation for all vectors, the basic equation for the routine is:

$$Z\{ A^H . U . \alpha^n . P . U^H . z \}$$

where:

$Z\{\ldots\}$ performs an operation on a matrix or vector which we define as:

a) form the magnitude squared of each element and

b) for the case of a matrix average along the rows.

A is a matrix containing a complete set of impulse responses or steering vectors for the sensor system, and should span the whole signal space.

U is a complete orthonormal set of eigenvectors of the covariance estimate derived from the batch of data.

α is a diagonal matrix containing the corresponding singular values of the data, (square root of the eigenvalues of the covariance estimate).

P is a rank deficient identity operator selecting the 'noise' (P_n) or 'signal' (P_s) vectors respectively from U using relative singular value as a partitioning criterion.

z is an operator which selects one column (usually the first) from the preceeding matrix.

X^H denotes the hermitian transpose of an arbitrary matrix X.

n takes a value 1, 0, –1 or –2

In our routine, P, z and α can be replaced independently by identity matrices in order to carry out any or all of the following algorithms:

MUSIC [1]	$Z\{ A^H . U . P_n . U^H \}$
TUFTS & KUMARESAN [2]	$Z\{ A^H . U . P_n . U^H . z \}$
MAXIMUM LIKELIHOOD(Capon) [3]	$Z\{ A^H . U . \alpha^{-1} . U^H \}$
MAXIMUM ENTROPY [4]	$Z\{ A^H . U . \alpha^{-2} . U^H . z \}$

In this standardised form the latter two filters generate minima in signal directions by virtue of the α^{-1} or α^{-2} weighting against the larger components. For large signals, the depths of minima depend closely on the inverse of signal amplitude but it is the 'directions' of these minima which are of interest to us. Note that weak signals may have shallow minima which are subject to noise and masking by the 'sidelobes' of deeper minima.

The first two algorithms also generate minima in signal directions by partitioning $U.U^H$ (an identity matrix) into two complementary subspaces – the signal partition and the noise partition [5]. It can be assumed that noise, given a sufficiently long sample of data, will appear to be uncorrelated (white) so that in the absense of signals, the filter response to the complete noise subspace should be reasonably flat. We see that the partitioning process not only extracts components corresponding to the 'wanted signals' but also removes signal-like waveforms from the noise. Projecting the incomplete noise subspace onto matrix A containing 'all possible signals' then shows minima, each corresponding to a 'hole in the noise'. Ideally these minima indicate the best matches between vectors in the matrix A and each of the signal waveforms.

Note that both the MUSIC algorithm and the KUMARESAN and TUFTS variant do not include a weighting depending on singular value or signal amplitude. The 'sidelobes' of large components should not therefore mask small signals. We regard this as a very significant pointer to the reason for the potentially superior performance of these algorithms in a realistic scenario where targets may be of unequal amplitude.

Signals, in general, appear to be partially correlated due to both physical and temporal aperture limitations and cannot be associated with specific (orthogonal) eigenvectors. Consequently signal power may leak into the noise subspace resulting in reduced null depths. This suggests the mechanism which limits the ability of this type of algorithm to identify waveforms arising from closely spaced sources or weak sources.

The generalised equations we give for the four first stage filters show up other similarities. In particular we see that both the MAXIMUM ENTROPY technique and the TUFTS and KUMARESAN version each select the first column of a 'modified covariance matrix', $U.\alpha^n.P.U^H$. Disregarding potentially useful information by using only one column appears wasteful but does improve performance in specific applications notably the case of a complex exponential kernel [2]. This can be explained by realising that using an end column is equivalent to applying a linear prediction filter of maximum window length. This is known to give the best model in frequency domain analysis. We see that both the other algorithms utilise all columns – forming an average with equal weighting. These are appropriate in the general case where signals are localised in the aperture.

For reference we also include in the results section the performance of a bank of conventional matched filters, defined simply by A^H, and implemented via our standard routine, as:

SIMPLE FILTER $\quad z(A^H.U.\alpha.U^H)$

It is of interest to realise that, in order to eliminate the under-determined aspect of the problem, the collateral knowledge implied by using this filter is that the amplitude of data extrapolated outside the collecting aperture should be minimised, clearly a basic assumption which is not valid in most applications. The filter is easy to implement and is also robust but cannot give a good representation of a 'point source' input scenario. The well known sidelobe masking of weak signals by larger signals means that a subsequent simple application of a point model (stage two) is not able to extract detail.

SIMULATION COMPARISONS

We show examples of simulation results for a relatively complex target scenarios. The system response matrix, A, for the first case has a (Sin(x)/x) impulse response covering 8 sample points uniformly spaced at intervals equal to the width of the sinc (between nulls). This example could represent a focal plane antenna array used for direction finding or perhaps a sampled aperture in the time domain used for time delay estimation - any filter with a sinc response. The second example has a complex exponential (Exp(ix)) impulse response and might be used to model a spatial direction finding problem based on a linear array of 16 antenna elements. Alternatively it could equally well represent any system requiring frequency domain analysis of signals obtained in the time domain. As indicated for example by Kumaresan and Tufts [2], and due to the symmetry of Exp(ix), we can discriminate against noise by reversing data in the x dimension and adding the resulting new covariance matrix to the original estimate. The information content in both covariance matrices is identical but some integration against noise improves the emphasis needed in stage one. Whenever appropriate we have utilised this technique, which can be regarded as implementing both backward and forward linear prediction filters. Note that, in the examples we give, the matrix A defining the system response is known exactly by the signal processing algorithm.

Our simulation program is able to process 3D data [6] and we have therefore been able to use the same targets for both the above examples; the two dimensional positions are given in the table below. The position units are dimensionless and correspond to the distance between nulls for the sinc case and correspond to the number of cycles within the aperture in the Exp(ix) example. One unit or cycle is equivalent the conventional resolution limit. Amplitudes are given in decibels relative to the noise power of a single sensor without integration. The sources are uncorrelated in a third domain (5 sample points).

Source	1	2	3	4	5	6
Amplitude(dB)	9	5	15	20	14	5
Sin(x)/x	3.5	2.6	.3	.3	-.2	-1.9
Exp(ix)	1	1	.1	-.3	-1.2	-1.2

Figure 1 shows the eigenvalues generated for the sinc domain and demonstrates the partitioning process. The curve in figure 2a, for the sinc domain, shows a typical 1D filter response for each of the five algorithms mentioned. The positions of the minima, from which the 'directions' of the point sources must be estimated, indicates a clear preference for the MUSIC algorithm followed perhaps by MAXIMUM LIKELIHOOD. Note that the MUSIC algorithm can be implemented by estimating the positions of peaks using the complementary signal subspace eigenvectors $(A^H.U.P_S.U^H)$. These peaks are visually comparatively broad and intervening dips appear to be very shallow but the resulting model is identical [5]. Care should therefore be taken in assessing performance on the basis of a natural preference for visual sharpness. The scale for these plots is 20dB per division with an arbitrary offset. The crosses mark the directions of the simulated sources used to generate the data. The dB scale refers to the power level of sources relative to unintegrated gaussian noise added to both in phase and quadrature components at each sample point. Figure 2b, for the complex exponential kernel, illustrates the superior performance in this case, of the KUMARESAN and TUFTS filter in comparison to MAXIMUM ENTROPY. Application of a point target model (stages two and three) to both the main dimensions of the input data leads to Figure 3. This 3 dimensional representation of the model and is obtained by applying the relevant psuedo inverses to the respective dimensions of the input data. The vertical lines represent the estimated power levels and are the average of the 5 samples in the third dimension of the data. The positions in the horizontal plane indicate estimated direction parameters.

DISCUSSION

The impact of the point target model is clear. The ouput model can be compared with the simulated input scenario and although there are several spurious signals indicated the model generated appears to be potentially useful. It could in principle be refined by further processing. The data reduction/compression achieved in the modelling process also has potential applications.

In both examples, as expected, an algorithm utilising the subspace/partitioning approach gives the best performance for the idealised point target scenario used in the simulation - particularly in respect of weak signals. Although results for only sample of data are shown, it can be seen that 'directions' are established for multiple sources, to within a fraction of the width of the relevant impulse responses without recourse to any form of conventional monopulse or 'split gate' technique. The level of discrimination achievable not only allows the power levels of sources of unequal amplitude and close spacing to be estimated but also enables the individual waveforms of sources in other domains to be analysed.

CONCLUSIONS

We have used a general equation to represent a number of advanced signal processing methods in order to compare and contrast algorithms and to aid our understanding of the principles which underpin advanced techniques. We have been able, by simulation, to demonstrate the superior discrimination potential of the MUSIC class of algorithms on multi-dimensional data and to illustrate the importance of utilising appropriate collateral knowledge in applications where a high level of discrimination between input vectors is required. The technique obtains much of its enhanced performance by utilising efficiently decorrelation between signal sources in more than one domain. The KUMARESAN and TUFTS version, although appearing to discard certain information, performs better than MUSIC in some applications, notably analysis in terms of complex exponentials, where the signal power of each component is essentially uniformly distributed over all the sample points.

REFERENCES

[1] Schmidt, R. O., A Signal Subspace Approach to Multiple Emitter Location and Spectral Estimation Ph.D. Thesis, Stanford University, November 1981.
[2] Kumaresan, R., Tufts, D. W., "Estimating the angles of arrival of multiple plane waves", IEEE Trans AES. Vol. AES-19, No. 1 January 1983 pp134-139
[3] Capon, J., Greenfield, R. J., Kolker, R. J., "Multidimensional maximum liklihood processing of a large aperture seismic array", Proc IEEE Vol. 55(2), 192-211, 1967.
[4] Nuttall, A. H., "Spectral analysis of a univariate process with bad data points, via maximum entropy and linear predictive techniques", NUSC Technical Report 5305, March 1976.
[5] Kay, S., Demeure, C., "The high resolution estimator - a subjective entity", Proc IEEE Vol 72(12), pp 1815-1816, 1984.
[6] Clarke, I. J., de Villiers, G. D., Mather J. L., "Resolution limits of a two dimensional antenna array", SPIE Proceedings Vol 564 1985. (to be published)
[7] Bowdler, H., Martin, R. S., Reinch, C., Wilkinson, J. H., "The *QR* and *QL* Algorithms for Symmetric Matrices", Numer. Math Vol 11 pp293-306, 1968.
[8] Pike, R. E., McWhirter, J. G., Bertero, M., de Mol, C. "Generalised

Information theory for inverse problems in signal processing", Proc IEE, Vol 131(F), p 660, Oct 1984

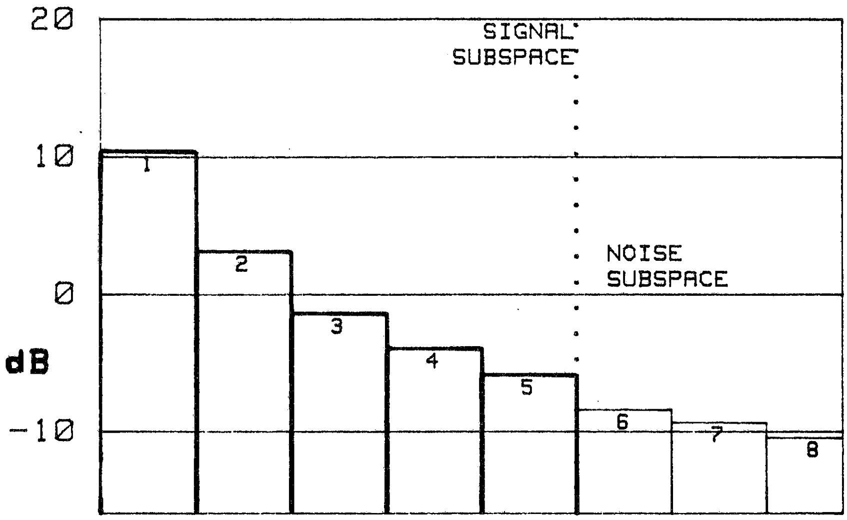

Figure 1. Eigenvalues - decreasing.

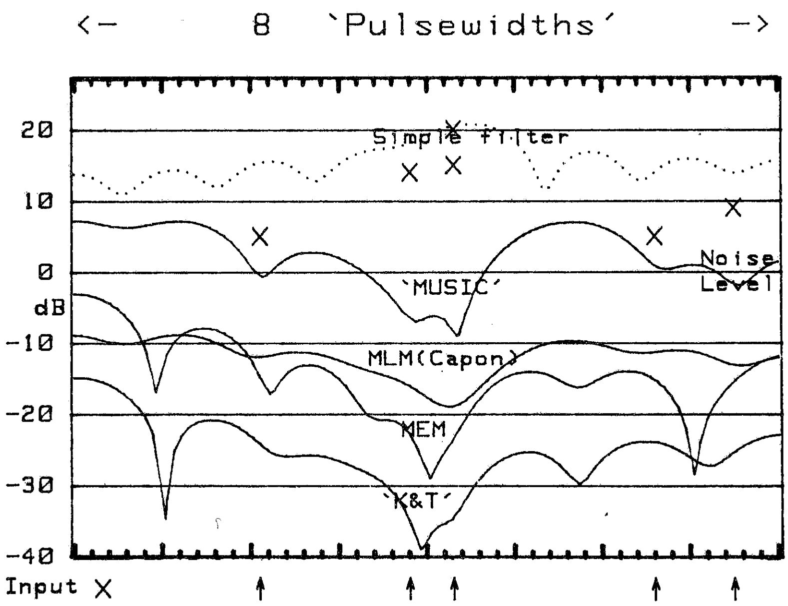

Figure 2a. Sinc analysis.

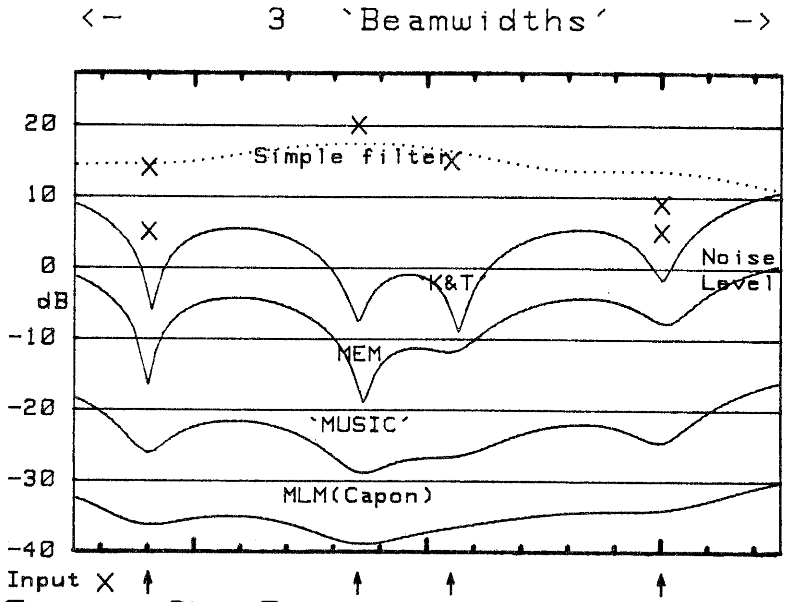

Figure 2b. Frequency analysis.

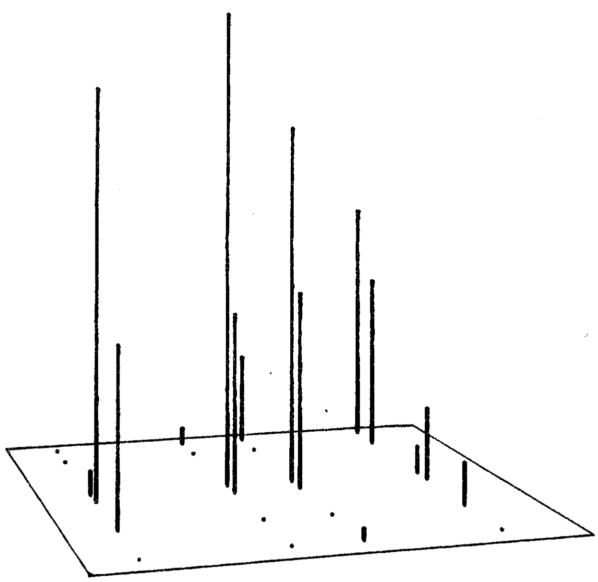

Figure 3. `3D' representation.