UNLIMITED

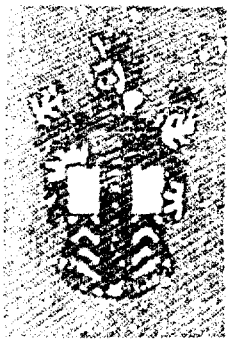

RSRE
MEMORANDUM No. 3864

ROYAL SIGNALS & RADAR ESTABLISHMENT

LEAST SQUARES SOLUTIONS IN SIGNAL PROCESSING USING THE SINGULAR VALUE DECOMPOSITION

Author: John L Mather

PROCUREMENT EXECUTIVE,
MINISTRY OF DEFENCE,
RSRE MALVERN,
WORCS.

UNLIMITED

ROYAL SIGNALS AND RADAR ESTABLISHMENT

Memorandum 3864

Title: LEAST SQUARES SOLUTIONS IN SIGNAL PROCESSING USING THE SINGULAR VALUE DECOMPOSITION

Author: John L Mather

Date: January 86

SUMMARY

We compare a number of advanced antenna array and time series signal processing techniques which either utilise singular value decomposition, or may be profitably interpreted in terms of it. Computer simulations of these techniques acting on spatio-temporal data from linear, conformal circular and focal plane arrays are presented and discussed. Means of introducing prior knowledge are described, and lead to the suggestion of a computationally efficient processing approach. We have made extensive use of a sequential processing strategy for multiple domain data and show by example that this potentially allows both detection and classification of additional emitters.

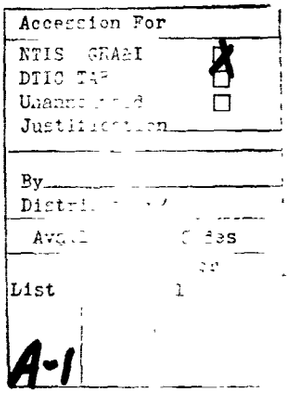

LEAST SQUARES SOLUTIONS IN SIGNAL PROCESSING USING THE SINGULAR VALUE DECOMPOSITION

John L Mather,

System and Chip Architecture - Techniques, SP2,

R.S.R.E.,

St Andrews Road,

Malvern.

CONTENTS

1. INTRODUCTION

2. SINGULAR VALUE DECOMPOSITION (SVD)

3. SPECTRAL ESTIMATION ALGORITHMS

3.1. SIGNAL PROCESSING AND PRIOR INFORMATION

3.2. MATCHED, INVERSE AND ORTHOGONAL FILTERING

3.3. THE ARRAY MANIFOLD

3.3.1. CALIBRATION

3.3.2. AMBIGUITIES

3.3.3. THE MANIFOLD AS PRIOR KNOWLEDGE

3.4. THE SAMPLE COVARIANCE MATRIX

3.4.1. UNCORRELATED SOURCES

3.4.2. THE SIGNAL AND NOISE SUBSPACES

3.4.3. CORRELATED SOURCES

2.4.4. THE MINIMUM DESCRIPTION LENGTH

3.5. RESOLUTION AND DISCRIMINATION

3.6. ALGORITHMS

3.6.1. THE BEAM SCAN ALGORITHM (BSA)

3.6.2. THE PSEUDO-INVERSE METHOD (PIM)

3.6.3. THE PRINCIPAL COMPONENTS BARTLETT METHOD (PCB)

3.6.4. THE CAPON MAXIMUM LIKELIHOOD METHOD (MLM)

3.6.5. THE THERMAL NOISE ALGORITHM (TNA)

3.6.6. THE MAXIMUM ENTROPY METHOD (MEM)

3.6.7. THE MULTIPLE SIGNAL CLASSIFICATION ALGORITHM (MUSIC)

3.6.8. THE PISARENKO MINIMUM EIGENVECTOR METHOD (PME)

3.6.9. THE KUMARESAN AND TUFTS SIGNAL CANCELLATION ALGORITHMS (KTSC)

3.7. THE SUB-APERTURE TECHNIQUE (SAT)

3.7.1. THE WINDOWED DATA MATRIX

3.7.2. SINGLE SNAPSHOTS

3.7.3. MULTIPLE SNAPSHOTS : COHERENT SOURCES

3.7.4. TIME SERIES PROCESSING

3.7.5. SUMMARY

3.8. REDUCING THE COMPUTATIONAL LOAD

3.8.1. COMPARISON OF THE ALGORITHMS

3.8.2. THE COEFFICIENTS METHOD

3.9. ITERATIVE PROCESSING

3.9.1. ITERATING THE PSEUDO-INVERSE

3.9.2. USING SUBSPACE ALGORITHMS AS PRIOR KNOWLEDGE

4. COMPUTER SIMULATIONS

4.1. THE SIMULATION FACILITY

4.2. DISCUSSION OF RESULTS

4.2.1. NOTATION

4.2.2. SENSITIVITY OF PIM AND BSA TO PRIOR ASSUMPTIONS

4.2.3. UTILISATION OF THE TIME DOMAIN INFORMATION

4.2.4. COMPARISON OF ALL ALGORITHMS ON SAME DATA

4.2.5. DEALING WITH COHERENT EMITTERS

4.2.6. MULTIPLE TARGET DISCRIMINATION

4.2.7. FOCAL PLANE PROCESSING

4.2.8. IMPROVING ANGULAR ESTIMATES THROUGH DOPPLER PROCESSING

4.2.9. ITERATIVE PIM PROCESSING

5. CONCLUSIONS

REFERENCES

1. INTRODUCTION

The numerical technique known as singular value decomposition (SVD) is presented here as a valuable tool for many signal processing applications. A description is given of a number of spectral analysis techniques which either utilise this fundamental tool, or may be usefully interpreted in terms of it. By comparison of these methods acting on simulated data we demonstrate the usefulness of SVD for performing data conditioning and, as a result, for obtaining improved signal estimation.

Section 2 provides a brief outline of SVD as a numerical tool for deriving least mean squares solutions to ill-conditioned or error-sensitive problems, and shows how the decomposition gives a measure of the contribution of orthonormal components to the matrix rank.

In Section 3, following an explanation of the means by which prior information may be introduced into the problem, a number of methods are discussed in some detail. The methods described are :

1. the Beam Scan algorithm, a simple matched filter method,
2. the Pseudo-Inverse method,
3. the Principal Components Bartlett algorithm,
4. the Capon Maximum Likelihood method,
5. the Thermal Noise algorithm,
6. the Maximum Entropy method,
7. Multiple Signal Classification,
8. the Pisarenko Minimum Eigenvector method,
9. the Kumaresan and Tufts Signal Cancellation algorithms,

and 10. the Sub-Aperture technique.

After a discussion of the computational load presented by the algorithms, we suggest possible means by which a number of them may be accelerated. This is followed by the description of a pre-processing step which makes use of knowledge concerning the redundancy in an array of sensors in order to compact the data before analysis. Consideration is also given to the possibility of utilising the output of a number of the algorithms as prior knowledge in some form of iterative processing, or to enable more detailed investigation of other data domains. This leads us to propose a computationally efficient sequential processing strategy for dealing with multiple domain data, which we have used extensively in our experiments. However, we show by examples presented in Section 4 that adoption of this approach implies that the order in which domains of the data are investigated can have an impact on the limits of discrimination.

A wide variety of typical results of computer simulations are presented and discussed in Section 4. These demonstrate the algorithms acting upon spatio-temporal data generated by simplified models of linear, conformal circular, and focal plane arrays of spatial signal detectors, and are used to illustrate many of the points developed earlier.

2. SINGULAR VALUE DECOMPOSITION (SVD)

Determination of the singular value decomposition of a matrix is a well established numerical technique for evaluating least-squares solutions to error-sensitive or ill-conditioned problems (Lawson and Hanson (1974), Golub and Reinsch (1970)).

The general matrix, $\underline{A}$, of dimensions $(m \times n)$, with $m \leq n$, may be decomposed as follows :

$$\underline{A} = \underline{U}\,\underline{S}\,\underline{V}^H \qquad (2.1)$$

where H denotes the Hermitian (complex conjugate) transpose, and

$$\underline{V}^H\,\underline{V} = \underline{U}^H\,\underline{U} = \underline{U}\,\underline{U}^H = \underline{I}_m \quad . \qquad (2.2)$$

$\underline{S}$ is a diagonal matrix containing the non-negative square-roots of the eigenvalues of $(\underline{A}\,\underline{A}^H)$, conventionally arranged in non-increasing order. The columns of the matrix $\underline{V}$ are the m orthonormalised eigenvectors corresponding to the m largest eigenvalues of $(\underline{A}^H\,\underline{A})$, and the columns of $\underline{U}$ are the orthonormalised eigenvectors of $(\underline{A}\,\underline{A}^H)$. The columns of $\underline{U}$ and of $\underline{V}$ are respectively known as the left and right hand singular vectors of $\underline{A}$, and the elements $\sigma_1, \ldots \sigma_m$ of the matrix $\underline{S}$ are called the singular values of $\underline{A}$. $\underline{I}_m$ is the $(m \times m)$ identity matrix.

If the matrix, $\underline{A}$ is square $(m = n)$ it is said to be full rank if the singular values all have values greater than zero. However, if the vectors comprising the rows or columns of $\underline{A}$ are not fully independent, then a number of the singular values will tend towards, or become, zero and the matrix may no longer be full rank. The rank of a matrix is defined as being equal to the number of non-zero singular values, and so for a rectangular matrix $(m \leq n)$ will always be limited by the smaller dimension, m. By setting any very small singular values to zero in equation 2.1, we can easily form a reduced rank approximation to $\underline{A}$. Similarly, by inverting all non-zero singular values, we can create a full rank or reduced rank pseudo-inverse (Penrose (1955)) of the original matrix as follows:

$$\underline{A}^{-} = \underline{V}\,\underline{S}^{-1}\,\underline{U}^H \quad . \qquad (2.3)$$

It can be shown that this procedure enables us to find a least squares solution to a general linear problem (Lawson and Hanson (1974)), and may therefore be considered as a potentially useful signal processing tool with which to average noisy data or generally to solve numerically sensitive problems (see, for example, Varah (1973), Cullum (1980)). The theory may also be extended to deal with inverse problems involving both continuous and discrete functions, and thus in the solution of certain types of integral equations (Hanson (1971), Pike et al (1984)).

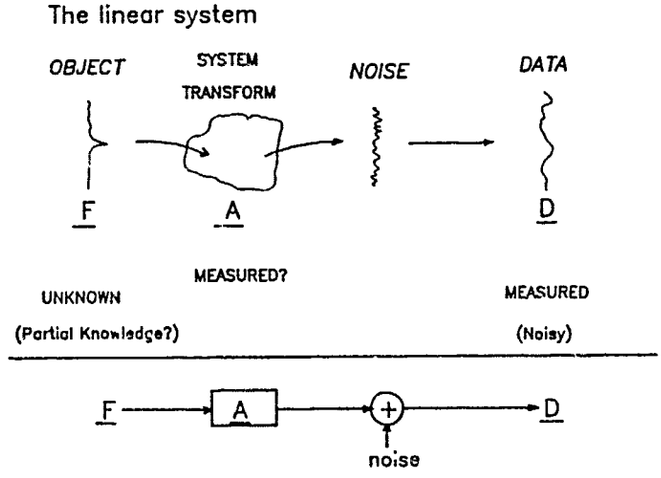

Fig 3 1 A general representation of a linear system, involving a mapping of object (F) to data (D) through the transformation A Often this will be a singular mapping of a continuous F to a discrete D

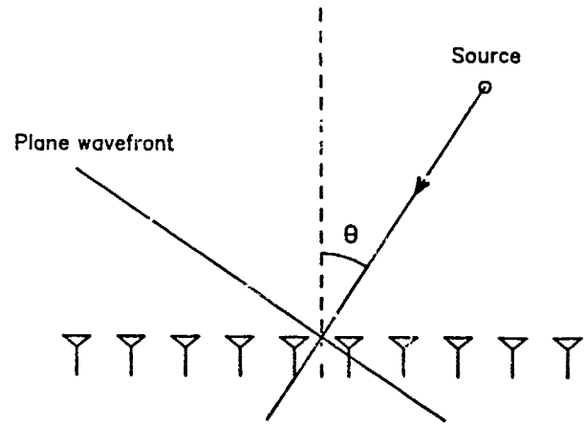

Fig 3 2 A line array of sensors is traversed by a plane wavefront at angle θ, emitted from a point source in the far field The array samples a portion of a spatial sinusoid which has its frequency related to the angle θ

3. SPECTRAL ESTIMATION ALGORITHMS

3.1. SIGNAL PROCESSING AND PRIOR INFORMATION

All signal processing relies on prior assumptions or information about the origin of the data. Presented simply with data in the absence of any other contextual information, we can say little, if anything, about its causation. In order to recover a transformed signal from received data, we need to know at least some of the characteristics of the transformation involved, such as the filter bandwidth, and the transform kernel (spatial Fourier transform, for example). We may also make use of information about the general nature of the solution required - the temporal or spatial extent of a signal, or its bandwidth for example - and the statistical form of the noise In general, signal processing should be regarded as a multi-dimensional problem, and it may be possible to utilise information from one domain or source of data to assist with the comprehension of other domains Such diverse information, which is often collectively referred to as prior knowledge, not only enables an attempt to be made at source reconstruction if used sensibly, but can also give an indication of the limiting factors involved if it is introduced in an explicit manner

The algorithms to be presented make use of varying levels of detail in the possible range of prior knowledge, from postulation only of the general nature of the transform, to the inclusion of knowledge about the statistical behaviour of the noise, and details of the type of solution required

3.2. MATCHED, INVERSE AND ORTHOGONAL FILTERING

Any general linear system may be represented as in Fig 3 1. The input vector, **f**, representing the source characteristics, is transformed by the system impulse response, $\underline{\mathbf{A}}$, to give a data 'snapshot' vector, **d**

$$\underline{\mathbf{A}}\,\mathbf{f} = \mathbf{d} \qquad (3.1)$$

A conventional method of signal processing attempts to transform the data into an image of the source by performing the following operation .

$$\underline{\mathbf{A}}^{H}\,\mathbf{d} = \mathbf{f}' \qquad (3\ 2)$$

This is loosely known as 'matched' or 'correlation' filtering : $\mathbf{f}'$ may be seen as representing the degree of correlation or match between those vectors which define the system transform and the data Generally defining the columns of $\underline{\mathbf{A}}$ as representing the response of the system to all possible impulses, each row of $\underline{\mathbf{A}}^{H}$ is a matched filter for at least one possible input

If we take as a specific example a multi-element linear array of spatial signal detecting elements, then our prior knowledge may be that $\underline{\mathbf{A}}$ represents a spatial Fourier transform, and $\underline{\mathbf{A}}^H$ is the inverse Fourier transformation This type of processing is only strictly valid if the data vector is infinitely long, and in the case of a Fourier transform, the effect of truncating the data is to create the familiar sin(x)/x response to a single frequency input signal. In general, therefore, we might believe that an exact inverse of $\underline{\mathbf{A}}$ would enable us to recover f more accurately, thus:

$$\underline{\mathbf{A}}^{-1} \mathbf{d} = \mathbf{f} \qquad (3.3)$$

This is known as inverse filtering, and incorporates more precise prior knowledge of the filtering characteristics of the mapping $\underline{\mathbf{A}}$. This is achieved through the implied constraint on the extent of f which is required for the formation of the inverse For a sample of data from an infinitely long line array, this becomes equivalent once again to the inverse Fourier transform, since $\underline{\mathbf{A}}\ \underline{\mathbf{A}}^H = \underline{\mathbf{I}}_m$ in this case.

Orthogonal filtering involves comparing the system response matrix with a vector which is orthogonal to the data. In the case of the spatial sensor array, the result is a vector with nulls in the directions of decorrelated sources. This vector is conventionally plotted inverted on a logarithmic scale, to give peaks at the appropriate angles. An example of such a filter would be an adaptive beamformer network (Applebaum (1976)) which is designed to filter out the spatial frequencies in the data which correspond to sources of interference

A number of the algorithms to be described make use of all the three principles outlined above, utilising the results from the orthogonal filter to provide support for the design of an inverse filter which is accurately constructed from vectors matched to the signals detected. However, before describing specific examples of algorithms, the system response matrix, $\underline{\mathbf{A}}$, will be discussed in greater detail. For a spatial array of sensors, this is commonly known as the array manifold as a result of its multi-dimensional vector space interpretation (Schmidt (1981)).

3.3. THE ARRAY MANIFOLD

3.3.1. CALIBRATION

The array manifold is a multi-dimensional characterisation of the impulse response of the particular configuration of sensors under consideration. We show first how this characterisation may be carried out for the example of a linear phased array antenna. The method is easily applied to alternative conformations.

As depicted by Fig 3.2, a signal source in the far field will result in a plane wave being received at a particular angle, $(\pi/2 \pm \theta)$, to the face of the array. For example, a source along a direction normal to the plane of the array ($\theta = 0$) will result in a wavefront parallel to the array. The simplifying assumption of uniform amplitude and equal phase weighting applied to the sensing elements of the array will then produce a data vector sampled at an instant in time whose terms are all equal. Similarly, a source located at a positive (negative) angle to the normal will result in positive (negative) frequency spatial pattern in the data vector. In the simplified case this pattern is a discrete complex exponential sampled at the element spacings and truncated by the spatial limits of the array.

Assuming that sources are in the far field and approximate to point emitters, complex data vectors may be recorded for all such possible source locations within the field of view of the antenna. These vectors can then be collected together as the columns, $\mathbf{a}(\theta)$, of a matrix, $\underline{\mathbf{A}}$, which forms a 'reference library' of signals expected for each incident direction. This library is a discrete space equivalent of the array manifold described by Schmidt (1981). Since each of the vectors which define the manifold may also be used to control the phases of each antenna element in a beamforming mode, they are often referred to as 'steering vectors'.

3.3.2. AMBIGUITIES

Clearly, ambiguities may exist in the interpretation of the steering vectors (Schmidt (1979)). Even assuming that spatial sampling rate criteria have been met, in the instance of a linear array of receiving elements as described, the array manifold will still contain one ambiguity - every column in the library may be interpreted either as the response to a source ahead of the array, or to one behind the array. This ambiguity may be resolvable by the addition of further information about the array, such as that signals are only to be expected from sources ahead of the array, or that the field of view is still further restricted by the directional properties of each element for example. The field of view may be specified in a general way by stating its extent in terms of a number of Rayleigh beamwidths (Cox (1973)) defined primarily by the overall aperture of the antenna. This is analogous to the information transmission capacity of the system (Shannon (1948)), and is described by its Shannon number.

Having demonstrated the way in which we may construct the array manifold, an example of the real (cosine) component of such a manifold is given in Fig 3.3. The array is a co-linear phased array of twenty elements

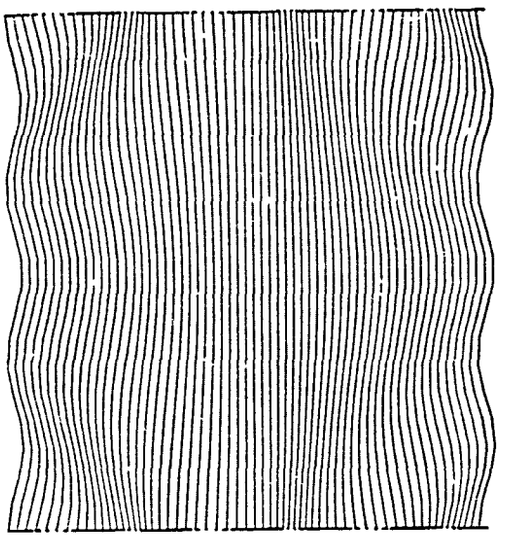

Fig 3 3 The real (cosine) component of a typical array manifold Each column represents the array response to a source at one of 60 angles in space This particular 20 element array has a Shannon number of 3 so that the limits to the angle of view are represented by waveforms of 1 5 cycles The ambiguity between left and right is resolved by the imaginary component

with uniform antenna weights and spacing, and having a Shannon number of six. The field of view, or object domain, is quantised into sixty positions across the six beamwidths. A measure of the angle of arrival of signals is the number of spatial cycles created across the face of the array by the source For this case the angular limits of the field of view thus correspond to ± 3 cycles.

3.3.3. THE MANIFOLD AS PRIOR KNOWLEDGE

Conventionally, interpretation of received data from a line array corresponds directly to performing a discrete Fourier transform. For a situation in which the array response does not approximate closely to the idealised linear example given above. Fourier techniques are not directly applicable The methods of analysis to be described are designed to take account of any such perturbations if they can be modelled or measured.

It can clearly be seen that the array manifold contains much prior knowledge useful for interpreting the received data. Firstly, it holds the system response characteristics for a large number of directions of arrival, and secondly it may also implicitly contain knowledge about the parameters associated with particular signals, or the angular bounds between which signals are to be expected.

3.4. THE SAMPLE COVARIANCE MATRIX

3.4.1. UNCORRELATED SOURCES

If we have a sequence of data snapshots of a source scenario consisting of uncorrelated point emitters it should be possible to make use of the additional information that the temporal modulation in the data provides It can be shown that this extra information may be optimally recovered by examining the eigenvalue spectrum of the data covariance matrix (van Trees (1968))

Assuming that signals and noise are uncorrelated, and that the noise is zero mean Gaussian at each antenna element, we represent the process of detection as :

$$\underline{D} = \underline{A}\,\underline{F} + \underline{N} \qquad (3\ 4)$$

where $\underline{D}$ is the matrix of snapshots, the columns of $\underline{F}$ contain the unknown target parameters as they vary with time, and $\underline{N}$ is the record of noise samples.

For uncorrelated signals with Gaussian properties, knowledge of the correlation function of the process is sufficient to define the required spectrum. However, in general we do not have sufficient samples to form an

unbiased estimate of this function, and so we compute the data covariance matrix, which, in the large sample limit tends towards the form of a Toeplitz correlation matrix (Makhoul (1975)).

The (m x m) spatial covariance matrix of the data may be estimated in practice by a variety of techniques, such as simple block averaging, sliding window averaging (see section 3.7), time decay averaging, forward-backward averaging, etc (Gabriel (1984)). Since these methods may all be interpreted as differences in the construction of the data matrix, they can be generally defined by the sample covariance estimate, $\underline{R}$, given by the equation,

$$\underline{R} = \underline{D}\,\underline{D}^H = \underline{A}\,\underline{H}\,\underline{A}^H + \underline{Q} = \underline{W} + \underline{Q} \qquad (3.5)$$

where $\underline{H}$ is $\underline{F}\,\underline{F}^H$, the 'signal-in-space' covariance matrix, $\underline{W}$ is the signal covariance matrix, and $\underline{Q}$ is the noise covariance matrix. Over a long period of observation the uncorrelated nature of the noise will result in $\underline{Q}$ becoming a diagonal matrix, which will simply add to the eigenvalues of the matrix $\underline{W}$. If the noise is white, equation 3.5 may be written,

$$\underline{R} = \underline{W} + \sigma_o^2 \underline{I} \quad , \qquad (3.6)$$

where $\underline{I}$ is the identity matrix, and σ_o^2 is the noise power which will equal the minimum eigenvalue, λ_{min}, of the generalised eigen equation,

$$\underline{R}\,\Phi_i = \lambda_i\,\underline{Q}\,\Phi_i \quad , \qquad 1 \leq i \leq m \quad , \qquad (3.7)$$

3.4.2. THE SIGNAL AND NOISE SUBSPACES

The signal covariance matrix may be considered as a sum of outer (dyadic) products of the columns of the $\underline{A}$ matrix, weighted according to the power of the corresponding signal associated with each direction. Thus, since the rank of each outer product is one, the rank of the total signal correlation matrix is, in the absence of noise, and given a sufficient number of independent data snapshots, equal to the number of spatially and temporally uncorrelated sources, N_S, for $N_S \leq m$. In the presence of noise, an evaluation of the pseudo rank of $\underline{R}$ must be made in order to form an estimate of the minimum number of sources which are needed for the model.

In general, if we have a good measurement of the second order noise statistics in the form of the matrix $\underline{Q}$, we may solve for the eigenvalues of $\underline{R}$ in the metric of the noise (equation 3.7). Cox (1973) shows that this is equivalent to ensuring that emphasis is given to those components containing least noise so that the choice of pseudo rank for $\underline{P}$ becomes less subjective. The eigenvectors of $\underline{R}$ associated with the largest eigenvalues are known as the "signal subspace" basis vectors, and those associated with the smaller eigenvalues are referred to as the "noise subspace" basis vectors. Because of the mutual orthogonality of singular vectors (the eigenvectors of the covariance

matrix are the singular vectors of the data) the two subspaces are also mutually orthogonal. We therefore appear to have a means by which the signals and noise may be largely separated if required. It is important to remember, however, that the signal and noise as such are not necessarily completely orthogonalised, as the subspace terminology might be taken to imply. Given a limited sample of data, it may be the case that the signal subspace eigenvectors still contain a significant component of noise, and vice versa

3.4.3. CORRELATED SOURCES

Implicit in the above discussion is the assumption that the autocorrelation function is stationary, which is the same as requiring that the covariance matrix tends to the Toeplitz form in the large sample limit. This is only the case if it is true both that the sources are stationary in space for the duration of the observation and that they appear to be decorrelated within that interval.

Correlated sources produce fields in space which are non-stationary (White (1979)) and the data matrix which is formed does not have a proportional relationship between rank and number of emitters. Thus in general, processing intended for the estimation of stationary signals does not produce useful results, unless the data may be pre-processed in such a way as to derive an approximation to a stationary correlation function. A possible means of achieving this is described in section 3.7.

3.4.4. THE MINIMUM DESCRIPTION LENGTH

Numerous methods have been devised for automatically evaluating the dimension of the signal subspace by examination of the eigenvalues of $\underline{\underline{R}}$ These variously make use of curve fitting techniques (Gabriel (1984)), maximum likelihood estimates of the model order (Akaike (1974)), and information theoretic criteria (Rissanen (1978)). We have chosen to implement the minimum description length (MDL) criterion based on the work of Rissanen (1978) and described by Wax and Kailath (1983). This criterion forms the estimate by calculating the function

$$MDL(k) = - N_p \log\left(\prod_{i=k+1}^{m} \sigma_i^2 / [1/(m-k) \sum_{i=k+1}^{m} \sigma_i^2]^{m-k} \right) + \tfrac{1}{2}k(2m-k)\log(N_p) \qquad (3.8)$$

where N_p is the number of observations or snapshots, for k = 0 to m-1, and finding the value of k for which it is minimised. The first term of this equation is the log of the maximum likelihood estimates of the model parameters (van Trees (1968)) and the second is a bias correction term. This criterion has been used in our simulations and has consistently given estimates of subspace

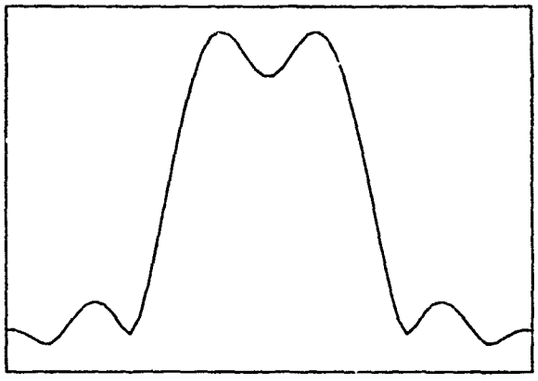

Fig 3 4a A typical amplitude function consisting of a sum of two sinc functions demonstrating 'classical' resolution.

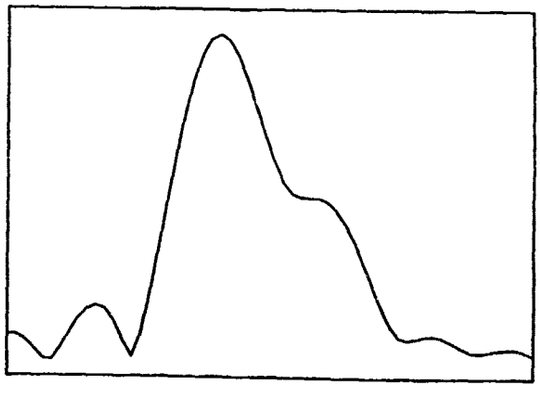

Fig 3.4b A typical function consisting of two sinc functions of unequal amplitude In a classical sense, resolution has not been achieved.

dimension which agree with those which would have been selected by a manual operator.

3.5. RESOLUTION AND DISCRIMINATION

Cox (1973) defines 'resolution' in the classical sense simply as recognition that an observed effect is due to two separate sources instead of one. In practice, perhaps the best we can say is that the observation is the result of a minimum of two sources, since the effect of additional sources may be lost in the noise. However, the definition elaborated by Cox is based on the assumption that resolution may be measured by the ability of a processor to produce distinct peaks in its response corresponding to each independent source of signal (Fig 3.4a). This immediately implies a direct 'operator' interpretation of the spectrum in terms of point targets corresponding to such peaks.

In the example of an ideal regular linear array, we know that this is an incorrect interpretation of the end product of Fourier processing. Prior knowledge of the transform kernel in this case tells us that the correct interpretation is as a sum of sinc functions of unknown amplitudes and positions. In this sense we are able to say that the Fourier beamformer has resolved when it is apparent that the response differs from the single sinc function in some way such as illustrated in Fig 3.4b. Clearly this does not lead to a very useful definition of resolution in practice since it gives us no direct measure of relative target position, other than that the sources are placed within the conventional Rayleigh limit. We therefore prefer to use the term 'discrimination' to describe the ability of an algorithm to produce distinct peaks or nulls for each source in its output.

3.6. ALGORITHMS

3.6.1. THE BEAM SCAN ALGORITHM (BSA)

The beam scan algorithm is otherwise loosely known as the correlation or matched filter, since it is composed of a number of parallel filters (the rows of $\underline{A}^H$), each matched to a specific possible input. We refer to it as a beam scan algorithm because the signal processing is analogous to scanning the beam of a radar array through the data. The data is compared with each of the beam steering vectors contained in the array manifold in turn, and the resulting product forms the spatial power estimate. If we have more than one

data vector (multiple snapshots), the sample covariance estimate, $\underline{R}$, may be calculated,

$$\underline{R} = \underline{D}\,\underline{D}^H \quad . \tag{3.9}$$

and the power estimate may be written as,

$$P_{BSA}(\theta) = \mathbf{a}^H(\theta)\,\underline{R}\,\mathbf{a}(\theta) \quad . \tag{3.10}$$

where θ is the angle of arrival of the signal, and $\mathbf{a}(\theta)$ is the corresponding steering vector.

The ability of the BSA to discriminate multiple spatially distributed point emitters by producing independent peaks in $P_{BSA}(\theta)$ is limited for short data records because of the resulting broad 'beam' shape and high sidelobes, but nevertheless provides the best possible (maximum likelihood) linear estimator of the direction of a single emitting source from data containing additive white Gaussian noise (van Trees (1968)).

3.6.2. THE PSEUDO-INVERSE METHOD (PIM)

As stated in section 3.2, an improved processor for a given system should be an exact inverse of the impulse response matrix, $\underline{A}$ in order to recover f thus:

$$\underline{A}^{-1}\mathbf{d} = \mathbf{f} \tag{3.11}$$

In the case of the linear detector array, because the vectors contained in the library matrix described in section 3.3 consist of *truncated* complex exponentials, they are not necessarily fully independent. This results in the inversion of $\underline{A}$ being an ill-conditioned problem since a number of the singular values become very small. The operation described by equation 3.11 implies the multiplication of components within the data by the inverted singular values. For the transform kernels of interest to us here this would possibly lead to the amplification of any noise contained in the data and of uncertainties in the characterisation of the array (Cullum (1980)). The array manifold must therefore be approximated in some suitable way The computational problem is further exacerbated by the generalisation to a rectangular $\underline{A}$ matrix, which represents a limited number of sample points in relation to an infinite number of possible source locations, and may be thought of as resulting in an underdetermined system of equations

One means of overcoming both of these problems is to recombine only a small number of singular values and vectors from the decomposition of $\underline{A}$ in order to provide a least-squares solution to the problem (Lawson and Hanson (1974)). A suitable *upper* limit on the number of singular vectors used could occur when the calculated singular values fall to the level of errors introduced by their computation. In the example of the twenty element array used earlier,

with its Shannon number of six, truncating the singular value series when σ_i/σ_1 falls below say 10^{-5} leaves us with only eleven 'significant' singular vectors from the initial set of twenty. The array may be said to have only eleven significant degrees of freedom (Toraldo di Francia (1969)). Truncation of the singular value series now allows us to find a pseudo-inverse in the general case of a rectangular $\underline{\mathbf{A}}$ matrix

Using the singular value decomposition of $\underline{\mathbf{A}}$,

$$\underline{\mathbf{A}} = \underline{\mathbf{U}}\,\underline{\mathbf{S}}\,\underline{\mathbf{V}}^H \quad , \tag{3.12}$$

we may re-write equation

$$\underline{\mathbf{A}}\,\mathbf{f} = \mathbf{d} \tag{3.13}$$

in the form

$$\underline{\mathbf{S}}\,\underline{\mathbf{V}}^H\,\mathbf{f} = \underline{\mathbf{U}}^H\,\mathbf{d} = \mathbf{y} \quad , \tag{3.14}$$

by decomposing the data onto the orthonormal basis provided by the columns of $\underline{\mathbf{U}}$. Finally the result is obtained from

$$\mathbf{f} = \underline{\mathbf{V}} \sum_{i=1}^{m} (y_i/\sigma_i)\,\mathbf{e}_i = \underline{\mathbf{V}}\,\underline{\mathbf{S}}^{-1}\mathbf{y} \tag{3.15}$$

where the y_i are components of $\mathbf{y}$, and $(\mathbf{e}_i)$ are the columns of an identity matrix. Thus,

$$\mathbf{f} = \underline{\mathbf{V}}\,\underline{\mathbf{S}}^{-1}\,\underline{\mathbf{U}}^H\,\mathbf{d} \tag{3.16}$$

Clearly, if we are finding a rank-reduced pseudo-inverse of $\underline{\mathbf{A}}$ by setting some of the σ_i to zero, we must assume that the y_i decrease faster than do the σ_i with increasing i. If this is the case we may then truncate the σ_i when they fall to the level of the noise in the system, since this implies that eigenvector components of the data have been attenuated by the transformation to such a degree that they may no longer be recovered.

For multiple snapshots of data, the power estimate as a function of angle of arrival may be written in the same form as for the BSA,

$$P_{PIM}(\theta) = \mathbf{a}^{-1}(\theta)\,\underline{\mathbf{R}}\,\mathbf{a}'(\theta) \quad , \tag{3.17}$$

where $\mathbf{a}^{-1}(\theta)$ is a row of the matrix $\underline{\mathbf{A}}^{-1}$, and $\mathbf{a}'(\theta)$ is a column of the rank-reduced system transform.

The 'beamwidth' of this processor varies according to the number of singular values included in the pseudo-inverse calculation, and under conditions of low signal to noise ratio, with appropriate truncation of the σ_i, broadens until it approaches that of the BSA. Very high S/N ratios have to be sustained in order to achieve beamwidth reductions of up to a factor of 3 or 4, and the algorithm is much more sensitive to the validity of the assumed spatial extent of the object domain than the BSA, as will be shown by examples presented in Section 4.

3.6.3. THE PRINCIPAL COMPONENTS BARTLETT METHOD (PCB)

This signal matching method estimates the spatial power spectrum by a similar approach to the BSA It differs from BSA in using a rank-reduced data matrix as the object to be scanned by the beam

In section 3 4 it was shown that the sample covariance matrix could be diagonalised in order to define a signal subspace, with basis vectors corresponding to those eigenvectors with large eigenvalues, and an orthogonal noise subspace defined by those eigenvectors associated with small eigenvalues The PCB algorithm involves the formation of an estimate of $\underline{R}$ which includes only the signal space, or principal eigenvectors, and sets the corresponding eigenvalues to unity,

$$\underline{R}_S = \sum_{i=1}^{k_{max}} \Phi_i \Phi_i^H \qquad (3\ 18)$$

Inherent in this description is a strong interpretation of the eigenvalue spectrum We are asserting that the large eigenvalues *do* correspond to signal, and that the rest *do* correspond to noise as far as subsequent processing is concerned Particularly in the case of under-estimation of the signal subspace dimension incorrect partitioning of the spaces can have significant influence on the form of the spectrum obtained

In the formation of $\underline{R}_S$ the PCB algorithm has thrown away the estimate it had about the relative signal powers in the form of the discarded eigenvalues This may even be considered a valid step, particularly in the case of short data records, when the estimate may have been a poor one, and is in fact the means by which the algorithm ensures discrimination of low amplitude emitters in proximity to sources of higher power By fitting functions of equal magnitude to the data, the algorithm circumvents the difficulty in the definition of resolution described by Cox (1973) which arises under such circumstances

We may now simply form an evaluation of the directions of arrival of signals,

$$P_{PCB}(\theta) = \mathbf{a}^H(\theta)\ \underline{R}_S\ \mathbf{a}(\theta) \qquad (3.19)$$

If we now wish to interpret the resulting spectrum in terms of point targets, it is necessary to locate the N_{pk} peaks of the generally very broad estimate. Having achieved this, the steering vectors corresponding to such peaks are taken as indicating the directions of the point emitters, and may be stored in a full rank matrix, $\underline{B}$ Forming a pseudo-inverse of $\underline{B}$ and multiplying by the data matrix,

$$\underline{T} = (\underline{B}^H\ \underline{B})^{-1}\underline{B}^H\ \underline{D} = \underline{B}^{-1}\ \underline{D} \qquad (3\ 20)$$

produces a matrix containing N_{pk} rows, each of which is a time series of data associated with a particular direction in space, information from the other directions being cancelled by the alternative steering vectors in $\underline{B}$ As

described in section 3.7, further processing of $\underline{I}$ can provide more detailed information about the target scenario, such as the number of emitters at each direction. Target cross and auto-powers can be derived by examining the elements of the reduced target data covariance matrix, $(\underline{I}\ \underline{I}^H)$. The powers determined in this fashion are the total powers corresponding to each direction of arrival and may be of use in discriminating against spurious peaks in the response which are the result of noise.

3.6.4. THE CAPON MAXIMUM LIKELIHOOD METHOD (MLM)

The maximum likelihood method (Capon et al (1967)) may be interpreted as an orthogonal filtering technique, since it utilises the inverse of the sample covariance matrix to produce a vector with nulls in the direction of sources, and with null depths corresponding to estimates of signal strengths. The spatial spectral estimate is then derived from the inverse of the component magnitudes of the vector containing the nulls.

The approach to this algorithm is common to all the orthogonal methods to be discussed. An attempt is made to minimise the mean square array output power under certain constraints. It is the constraints which differentiate the various algorithms. In the case of MLM the constraint is that the array should have unit gain in a particular direction, whilst simultaneously performing the power minimisation. This has the effect of optimally rejecting power from directions other than that which defines the constraint (Cox (1973)).

For the case of Gaussian noise, the average power output from the array for a plane wave arriving from an angle θ is given by

$$P_{MLM}(\theta) = \mathbf{w}(\theta)^H \underline{R}\, \mathbf{w}(\theta) \quad , \tag{3.21}$$

where $\mathbf{w}(\theta)$ is the vector of array weights. The unity gain constraint requires that $\mathbf{w}(\theta)$ should satisfy

$$\mathbf{w}(\theta)^H \mathbf{a}(\theta) = 1 \quad . \tag{3.22}$$

The optimum $\mathbf{w}$ vector is the Wiener weight vector, which maximises the output signal to interference ratio,

$$\mathbf{w}(\theta) = \lambda\, \underline{R}^{-1}\, \mathbf{a}(\theta) \quad , \tag{3.23}$$

where λ is a complex number. Solving for λ in equation 3.23, we find that λ and P_{MLM} are the same. Substituting for λ and solving for $\mathbf{w}(\theta)$ produces the MLM power estimate,

$$P_{MLM}(\theta) = \left(\mathbf{a}(\theta)^H\, \underline{R}^{-1}\, \mathbf{a}(\theta) \right)^{-1} \quad . \tag{3.24}$$

Interpreting this result in terms of an eigenvector decomposition of $\underline{R}$, the matrix inverse is calculated using all of the eigenvector and eigenvalue information. There is no attempt to impose a signal subspace interpretation onto the decomposition. The result may be understood as an attempt to create

a filter which is maximally orthogonal to the signal, since noise eigenvectors will receive greater weight in the calculation of $\underline{R}^{-1}$, without actually having to decide where to partition the signal and noise spaces.

It should be noted that the estimate in equation 3.24 is the true maximum likelihood estimate only for the assumption of a single plane wave. The solution for multiple plane waves becomes progressively more complicated, as discussed by Evans et al (1982). Statistical estimation theory shows that the true maximum likelihood estimate is the optimal estimate for the Gaussian signal detection problem, and enables discrimination which approaches the theoretical limit for such unbiased estimators defined by limits such as the Cramér-Rao bound on the variance of the estimate (van Trees (1968))

The MLM estimate has very low 'sidelobe' characteristics for the case of incoherent signals, since $w(\theta)$ places nulls at all directions, unless θ corresponds to the direction of arrival of a plane wave, when the unity gain constraint prevents a null being produced. The magnitude of the response then gives an estimate of the power of the emitter.

3.6.5. THE THERMAL NOISE ALGORITHM (TNA)

In attempt to devise an algorithm which had the desirable characteristics of low 'sidelobe' response coupled with sharp 'resolution' peaks, Gabriel (1980) derived his thermal noise algorithm.

Taking the dot product of the optimum weight vector derived above with itself, and defining this as the adapted thermal noise power N_O.

$$N_O = \lambda \, w(\theta)^H . w(\theta) \qquad (3.25)$$

where λ is a quiescent noise power level constant, we derive the new spectral estimator

$$P_{TNA} = (\, a(\theta)^H \, \underline{R}^{-2} \, a(\theta) \,)^{-1} \quad , \qquad (3.26)$$

from the reciprocal of N_O. The peaks produced by this algorithm are proportional to the square of the source power levels.

If we examine this method as we did for MLM, by looking at its use of the eigenfunctions of $\underline{R}$, we see that the spatial estimate may be written

$$P_{TNA}(\theta) = (\, a(\theta)^H \, \underline{\Phi} \, \underline{\Lambda}^{-2} \, \underline{\Phi}^H \, a(\theta) \,) \quad . \qquad (3.27)$$

where $\underline{\Phi}$ is the matrix of eigenvectors of $\underline{R}$, and $\underline{\Lambda}$ is the diagonal matrix of its eigenvalues. The TNA may therefore be interpreted as placing a stronger weighting on the eigenfunctions which are most likely to belong to a noise subspace. As with MLM, the decision as to the partitioning of the subspaces has been avoided, but in this case the implication that the small eigenvalues correspond to the noise is reinforced by the squaring of the inverted eigenvalues.

Clearly the idea of manipulating the effective eigenvalue spectrum in this way may be extended to enable the processor to impose even stronger assumptions about the nature of the data on the result. The MUSIC and PCB algorithms are examples of a step function weighting being applied to the eigenvalues.

3.6.6. THE MAXIMUM ENTROPY METHOD (MEM)

The maximum entropy method was developed as a means of extrapolating a limited estimate of an autocorrelation function in such a way that the appropriate probability density function has its entropy maximised at each stage. Its intention is to form a 'maximally noncommittal' estimate (Jaynes (1957)). In common with MLM, TNA, MUSIC, KTSC, and many other similar algorithms, however, MEM is equivalent to a constrained least squares fitting of an all-pole model to the available data (van den Bos (1971)) and may be outlined without recourse to information theoretic concepts for the simple case of the regular linear array.

For brief data records, the usual unbiased estimator is the data covariance matrix, and we obtain the required filter coefficients by solving the equation

$$\underline{R}\,\mathbf{w} = [\,e\ 0\ 0\ \ldots\ 0\,]^T = \mathbf{e} \qquad (3\ 28)$$

where e is the so-called prediction error (Makhoul (1975)), with the weight vector,

$$\mathbf{w} = [\,w_1\ w_2\ w_3\ \ldots\ w_m\,]^T \qquad (3\ 29)$$

We obtain the solution for w of the form

$$\mathbf{w} = \underline{R}^{-1}\,\mathbf{e} \qquad (3\ 30)$$

and scaling w and e appropriately we may choose e to be the first column of the identity matrix. Thus the filter w is equal to the first column of the inverted covariance matrix (Nuttall (1976)).

Interpreting this filter in terms of eigenfunctions, as we did with MLM, it can be seen that we again have derived a function which gives greatest weight to the noise components of the data without the requirement for any additional information or assumptions about the subspace dimensions. If we compare this filter against the library of 'expected' waveforms in the $\underline{A}$ matrix, nulls should be produced in the result at directions corresponding to the signal directions. Thus our spectral estimate may be written,

$$P_{MEM}(\theta) = (\,\mathbf{a}(\theta)^H\,\mathbf{w}\,\mathbf{w}^H\,\mathbf{a}(\theta)\,)^{-1} \qquad (3.31)$$

It is worth noting the similarity in form between this estimator and P_{TNA}, in which $\underline{R}^{-1}\mathbf{e}$ is replaced by $\underline{R}^{-1}$. In fact the TNA result may be interpreted as an average of MEM estimates derived from alternative columns of the inverse

covariance estimate. Burg (1972) has also demonstrated a link between MLM and an average of MEM estimates resulting from different filter lengths, under certain circumstances. Note also that the form of equation 3.31 is only strictly the maximum entropy solution for a regularly sampled process and that the magnitude of $P_{MEM}(\theta)$ does not necessarily relate to the received powers

3.6.7. THE MULTIPLE SIGNAL CLASSIFICATION ALGORITHM (MUSIC)

Schmidt (1979) has described a signal subspace algorithm which he calls MUSIC (MUltiple Signal Classification). This is fully described by him for the generalised case where the data covariance matrix is analysed in the metric of the previously estimated noise (equation 3.7). Clearly the noise covariance matrix must be measured in the absence of other signals and then assumed to be stationary for a further period whilst the generalised eigenvalue problem is being solved for each block of data. Alternatively, in order to reduce the required computation, the noise may be assumed Gaussian, and an estimate of rank made from the decomposition of $\underline{R}$ in an identity metric.

Having obtained the singular vectors of the data, and formed an estimate of the number of identifiable independent sources present from the matrix rank estimate, the mutually orthogonal signal and noise subspaces can be defined. Defining $\underline{\Phi}_N$ to be the matrix of noise subspace basis vectors, the directions of arrival are estimated by computing the Euclidean distance, d, between the vectors of the array manifold and the noise subspace and then plotting $1/d^2$ against angle of arrival, θ :

$$P_{MU}(\theta) = (a(\theta)^H \underline{\Phi}_N \underline{\Phi}_N^H a(\theta))^{-1} \quad . \qquad (3.32)$$

This algorithm clearly has similarity to MLM if $(\underline{\Phi}_N^H \underline{\Phi}_N)$ is considered as an approximation to $\underline{R}^{-1}$ which simply uses the normalised noise eigenvectors (Gabriel (1984)). Alternatively, interpreting $(\underline{\Phi}_N^H \underline{\Phi}_N)$ as an approximation to $\underline{R}$, the similarity to PCB becomes apparent. Kay and Demeure (1984) show that the MUSIC and PCB estimators are linked by the simple relationship,

$$P_{MU}(\theta) = 1 \; / \; (1 - P_{PCB}(\theta)) \quad , \qquad (3.33)$$

for $P_{PCB}(\theta)$ normalised to a maximum value of 1. Therefore, although MUSIC is capable of producing an estimate which contains extremely narrow peaks at signal directions, whereas PCB results in a very flat response, the 'resolution' capability of the two algorithms is identical.

It should be noted that, in the case of covariance estimates calculated from very short data records, the noise subspace eigenvalues are not equal, even when the sample is from a white Gaussian noise process. MUSIC may be thought of as assuming equality of these values, perhaps on grounds of irrelevance or on the basis that computer rounding errors render them

inaccurate. We have discovered that the estimate,

$$P_{EMU}(\theta) = (\mathbf{a}(\theta)^H \underline{\Phi}_N \underline{\Lambda}_N^{-1} \underline{\Phi}_N^H \mathbf{a}(\theta))^{-1} \quad . \qquad (3.34)$$

often results in slight enhancement of the ability of MUSIC to discriminate sources. This estimator has the effect of increasing the weighting applied to those eigenvectors likely to contain the least amount of signal. The EMU estimate has additional advantages over MUSIC in that it more accurately reflects non isotropic noise backgrounds, and that it also degrades gracefully to give the estimate $P_{MLM}(\theta)$ if no estimate of signal subspace dimension can be formed (Johnson and de Graaf (1982)).

As demonstrated by Alsup (1984) both MUSIC and, by implication, the PCB filter are capable of extremely good discrimination for the case of uncorrelated sources when the noise is perfectly characterised This situation is hardly likely to arise in practise. The results presented here adopt a more realistic simulation in which the noise is *assumed* to be uncorrelated Gaussian by the processor, and the short data record is analysed in an identity metric.

3.6.8. THE PISARENKO MINIMUM EIGENVECTOR METHOD (PME)

The Pisarenko minimum eigenvector method may be considered as a special case of the MUSIC algorithm, in which the noise subspace dimension equals one. The eigenvector chosen, Φ_m, is that associated with the minimum eigenvalue of the covariance matrix In the case of white noise corrupting an infinitely long data record (so that the noise subspace eigenvalues are all equal), this can become an arbitrary choice. The form of the estimate is

$$P_{PME}(\theta) = (\mathbf{a}(\theta)^H \Phi_m \Phi_m^H \mathbf{a}(\theta)^H)^{-1} \quad . \qquad (3.35)$$

Although potentially easier to compute than the MUSIC response, it can be seen that, because the minimum eigenvector is orthogonal both to the signal and also to any other noise related eigenvalues, the result may contain spurious noise-generated peaks in positions unrelated to the signal directions

3.6.9. THE KUMARESAN AND TUFTS SIGNAL CANCELLATION ALGORITHMS (KTSC)

Many variants of the basic signal subspace technique are discussed in the literature. Two methods by Kumaresan and Tufts (1983) for example, produce a single filter vector which is effectively a weighted combination of noise subspace eigenvectors. This vector is then used to formulate the directional estimates by comparison with the array manifold as before, but in some cases improved performance over the MUSIC algorithm has been demonstrated

The two methods have been derived by the authors from linear predictive algorithms used for speech processing (Makhoul (1975)). As with MEM and many other algorithms, the differences are simply a matter of notation and

terminology and the fundamental approach of autoregressive modelling employed is equally suitable for processing data of any origin which may be satisfactorily described by second moment statistics (van Trees (1968)), given constraints which are suitable for the individual problem. We refer to the algorithms as signal cancellation algorithms because of the form of the constraint and the resulting similarity to a null-forming adaptive canceller equation.

The basic form of the equation for both methods is

$$\underline{D}^H \mathbf{w} = \mathbf{0} \quad , \tag{3 36}$$

where the weight vector is constrained to be

$$\mathbf{w} = [\, 1 \; w_2 \; w_3 \; \ldots \; w_m \,] \quad . \tag{3 37}$$

and **0** is the null vector. In radar terms, we wish to derive an antenna weight vector with unity gain on the end element, **w**, which cancels the data. Since $w_1 = 1$, equation 3.36 may be solved by taking the first column of $\underline{D}^H$, which we denote **c**, across to the right-hand side, and solving for elements w_2 to w_m of **w** by forming a pseudo inverse of the remaining columns of $\underline{D}^H$, which we denote $\underline{C}$.

$$\mathbf{w}' = [\, w_2 \; w_3 \; \ldots \; w_m \,] = -1 \; \underline{C}^{-1} \mathbf{c} \tag{3 38}$$

Setting the smaller singular values of $\underline{C}$ to zero ensures that the solution for **w** is of minimum length. This is equivalent to discarding the noise subspace vectors of the reduced data matrix. An excellent and detailed description of the operation of this algorithm appears in the paper by Sibul (1984).

Using the singular value decomposition of $\underline{D}^H$, equation 3.36 may be written in the form

$$(\underline{Y} \; \underline{S} \; \underline{\Phi}^H) \; \mathbf{w} = \mathbf{0} \quad , \tag{3 39}$$

where $\underline{Y}$ and $\underline{\Phi}$ are the appropriate basis vectors. Clearly this may be again re-written as

$$\underline{\Phi}^H \mathbf{w} = \mathbf{0} \; . \tag{3 40}$$

Utilising only those eigenvectors which correspond to the noise subspace, a stable result for the prediction filter **w**, may be derived, which can be shown to be of the form

$$\mathbf{w} = \underline{\Phi}_N \, \Phi_n \, / \, (\Phi_n \, \Phi_n^H) \quad , \tag{3 41}$$

where, in this case, Φ_n is the first row of the noise eigenvector matrix, $\underline{\Phi}_N$. Clearly, this produces a filter vector which is a complex weighted linear combination of the noise subspace basis vectors, and compares with MUSIC, which effectively employs unity weighting to combine the eigenvectors.

Finally, the spatial estimate is obtained by correlating the orthogonal filter vector with the array manifold,

$$P_{KTSC}(\theta) = (\, \mathbf{a}(\theta)^H \, \mathbf{w} \, \mathbf{w}^H \, \mathbf{a}(\theta) \,)^{-1} \tag{3 42}$$

Subsequently we refer to the latter form of the algorithm presented here as KTSC1, and the former as KTSC2. In the majority of cases the results obtained from these two forms are indistinguishable. As with other methods, the peak magnitudes in the result do not coincide with the received power and so power levels corresponding to each of the directions required are derived by interpreting the result in terms of point sources located at the positions of the peaks and forming the appropriate reduced $\underline{A}$ matrix, as described in section 3.6.3.

The results produced by this method are similar to those given by MUSIC, but with apparently slightly better noise smoothing and consequently improved discrimination. However, because of the type of constraint employed they are only ideally suited as presented here for regularly spaced and illuminated arrays of detectors. Care must be taken in general to specify a suitable form of constraint for the particular problem. An alternative approach, which we have used with some success on non-linear arrays is described in section 3.8.

3.7. THE SUB-APERTURE TECHNIQUE (SAT)

3.7.1. THE WINDOWED DATA MATRIX

The sub-aperture technique is not a self-contained method of analysis in itself, but rather a useful appendage to the algorithms discussed above. Its use is derived from work on time series analysis, and it is essentially a method of forming an improved covariance matrix estimate from limited data. It has important application to spatial processing for its ability to assist many of the spectral estimators described earlier to resolve coherent sources and multipath, and also to convert single snapshot data into a suitable format for processing by multi-snapshot algorithms.

The method consists of taking the data vector, and observing it through a small 'window' vector so that the autocovariance structure of the series may be examined. For every new observation, we slide the window vector further through the data, and create a new row in a data matrix. Thus, if the data vector is

$$\mathbf{d} = (d_1\ d_2\ d_3\ \ldots,\ d_n) \tag{3.43}$$

and this is examined with an r-element window, we can build up the following matrix:

$$\underline{D} = \begin{bmatrix} d_1 & d_2 & d_3 & \ldots & d_r \\ d_2 & d_3 & d_4 & \ldots & d_{r+1} \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ d_{n+1-r} & \ldots & \ldots & & d_n \end{bmatrix} \tag{3.44}$$

Measuring the pseudo-rank of $\underline{D}$ by singular value decomposition gives the estimate of the number of significant independent components contributing to the data series if this is less than the minimum dimension of $\underline{D}$.

The forward-backward technique of analysis which seems popular in the field of speech processing (Makhoul (1975)) and has recently appeared more frequently in the general signal processing literature (for example Evans et al (1982) and Kumaresan and Tufts (1983)) creates a data matrix which, in addition to the 'forward' moving window described above, utilises a sampling of the original series of data in the reversed direction. The reversed data is complex conjugated and included in the same matrix as the forward samples. It is clear that subsequent formation of the sample covariance matrix involves extra averaging of the data and better rejection of noise, particularly in the case of data which contains sinusoidal signal components.

3.7.2. SINGLE SNAPSHOTS

If **d** is a data vector received at a single instant in time from a uniformly weighted and spaced linear array of receiving elements, or perhaps if it may be suitably transformed to resemble this, we may proceed by continuing the processing using a subspace approach. The window vector is now the sub-aperture described by Gabriel (1984), and the algorithm must be provided with an array manifold for the sub-array rather than for the total number of elements. This is simply computed as a sub-section of the original translationally invariant manifold. Having done this, a startling improvement in the ability to discriminate closely spaced point emitters is possible.

3.7.3. MULTIPLE SNAPSHOTS : COHERENT SOURCES

Sub-aperture processing may also be beneficially applied to multiple snapshot data, particularly for its effect on data from coherent sources. In this case it is perhaps easier to think of using an $(r \times m)$ rectangular sampling matrix $\underline{K}$ of the form,

$$\underline{K}_i = [\ \underline{0}\ |\ \underline{I}\ |\ \underline{0}\ ] \quad , \tag{3.45}$$

where the subscript i denotes the column in which a 1 first appears, to extract a principal submatrix, $\underline{R}_i$, of the original sample covariance matrix,

$$\underline{R}_i = \underline{K}_i\ \underline{R}\ \underline{K}_i^H \quad . \tag{3.46}$$

Averaging the submatrices produced in this manner then has the desired result (Evans et al (1982)).

For the case of data originating from coherent sources, the improvement attainable with this approach may be understood as deriving from the 'internal averaging' of the normally computed correlation function (Evans et al (1982))

which causes the non-stationary terms to tend toward zero. Alternatively, Gabriel (1982) suggests that the improvement may be interpreted as arising from the movement of the sub-array phase center with respect to the targets, which introduces an effective doppler phase shift between the sources.

3.7.4. TIME SERIES PROCESSING

Clearly, such processing can have wide applicability, and may perhaps be used on data from domains other than the spatial domain. It was suggested in section 3.6.3, for example, that more information about the nature of the received signals could be deduced by investigating the time series component of the data in isolation. This is a potentially useful post-processing stage after reducing the number of time series to be examined by means of the spatial spectral estimators described previously Taking an individual time series from the matrix $\underline{T}$ (equation 3.20), corresponding to a particular direction identified as worthy of further investigation, it is possible to estimate the number of decorrelated signals present at that direction by forming a covariance estimate using the sub-aperture technique and processing the result by the same methods used for spatial data (Ziegenbein (1979)). In fact, if a library of expected time domain signals is available, the individual waveforms which are present may be estimated. In the absence of such a manifold, the eigenvalues of the covariance matrix, $(\underline{D}\ \underline{D}^H)$, may simply be used to give an estimate of the relative powers of the different component modes, given sufficient samples (van Trees (1968)) The corresponding problem of estimating source locations, given the emission frequencies may also be tackled in this way

The benefits of using such a sequential approach to the processing are illustrated by many of the examples presented in Section 4. If data has been gathered in more than two domains (for example along two spatial axes and in time) it is clear that this sequential strategy may be extended to process it in a simple manner.

3.7.5. SUMMARY

It can be seen that the sub-aperture technique enables the detection of a larger number of independent sources than the rank of the usual data matrix, limited as it is by the minimum dimension of the matrix, might initially suggest. We have shown, for example, that it is possible both to locate coherent emitters which are sufficiently spatially uncorrelated and also the frequencies of spatially correlated targets if they are decorrelated in the time domain. This suggests that the introduction of data from additional measurement domains in which targets are decorrelated, will not only allow

more detailed characterisations of sources already detected, but will also increase the total number of sources which the processor is able to detect. If a sequential processing strategy is adopted, tackling the problem which has the smallest dimension first enables the remaining load to be accordingly reduced. However, if targets are not easily discriminated in that domain, some information may be lost or inaccurately recovered, as will be discussed in section 3 9 2, and demonstrated in section 4.2.8.

3.8. REDUCING THE COMPUTATIONAL LOAD

3.8.1. COMPARISON OF THE ALGORITHMS

If computational simplicity is the major criterion in choosing a suitable algorithm for signal processing, then clearly the optimum choice is either the BSA or the PIM algorithm. Both involve pre-computation of the filter matrix, and so data processing can take place on line, acting on each snapshot as a simple matrix-vector multiplication, or on blocks of data in the form of the covariance matrix estimate. If the enhanced discrimination of the other algorithms presented is desired, however, the computational load is significantly increased and calculations may have to be carried out off line

The true maximum likelihood method, although optimal from the point of view of estimation theory (van Trees (1968)), rapidly becomes computationally intractable as the number of assumed wavefronts is increased, since it involves both matrix inversions and a simultaneous multi-parameter search for each new sample covariance matrix. It is more common to find it used, as we have done for the current work, in the form given in section 3.6 4 which has been derived on the assumption of a single plane wave arriving at the antenna.

If a full singular value analysis of the data matrix is needed for the simpler KTSC algorithm, for example, this might be accomplished by diagonalisation of the covariance matrix using standard routines. Householder tridiagonalisation of the matrix, followed by calculation of all the eigenvalues and eigenvectors using a QR routine (Wilkinson (1965), Wilkinson and Reinsch (1971)) would involve $2m^3/3$ multiplications and m square roots for the first stage, and a further $4m^2$ multiplications with m-1 square roots for each iteration of the QR algorithm (generally the number of iterations required is of the order of two per eigenvalue for the tridiagonalised matrix). In addition, the formation of rank-reduced pseudo-inverses and matrix approximations requires yet more processing for each new block of data, if it is to be carried out on line.

Because of the usual m^3 dependency of the processing load, the requirement escalates rapidly with the size of the problem, and so from this

point of view algorithms which reduce the size of the matrix to be dealt with are potentially very valuable. An example of such an algorithm is the sub-aperture technique, which reduces the covariance matrix to the dimension of the window used, at the expense however of a possibly arbitrary reduction in the number of available degrees of freedom for modelling the targets.

It is worth noting, however, that for problems involving only a small number of sensors, the eigen analysis may actually require less computation than is then needed for the location of the sources, if this involves searching a large number of points Instead of exhaustive searching, more efficient algorithms must be employed For searches involving more than one data domain, our sequential processing approach, which uses the results of a search in each domain as support for a reduced search in the next domain, has proved to be effective in reducing the complexity of the computation (Clarke et al (1985)).

3.8.2. THE COEFFICIENTS METHOD

We know from section 3.6.2 that the system response matrix, $\underline{A}$, may be a less than full rank transform and that our data, $\underline{D}$, will therefore generally contain redundant or noisy components. Analysis of the data in terms of the appropriate basis vectors of $\underline{A}$, and rejection of components at or below the noise floor allows us to represent the useful data in a more compact form

As a result of the singular transformation,

$$\underline{A}\,\underline{F} + \underline{N} = \underline{D} \quad , \tag{3 47}$$

$\underline{D}$ is a matrix of dimensions $(m \times N_p)$, where m is the number of elements in the array, and N_p is the number of snapshots in the block of data to be processed.

Premultiplying both sides of this equation by the transposed set of left-hand singular vectors of $\underline{A}$ gives

$$\underline{U}^H (\underline{A}\,\underline{F} + \underline{N}) = \underline{U}^H \underline{D} = \underline{C} \quad . \tag{3 48}$$

In general, the size of the matrix $\underline{U}$ may be reduced by examining the singular values of $\underline{A}$ in relation to the estimated noise floor This number will be less than m, and may be denoted k Thus, the data is now represented by its projection onto the "significant" basis vectors of the array manifold, and is stored in a matrix, $\underline{C}$, of dimensions $(k \times N_p)$. This has allowed a reduction in the size of the effective data array to a minimum size consistent with retaining all potentially useful degrees of freedom.

It is now possible to use the matrix, $\underline{C}$, as the input to a signal-subspace processor. However, the assumed array manifold must also be projected onto the same basis vectors, and becomes $(\underline{U}^H \underline{A})$. The vector or set of vectors

orthogonal to the signal subspace which the algorithms produce can be compared against the columns of this transformed manifold in order to identify source directions. In this way, it is possible to significantly reduce the computational load in poor signal to noise environments, when the rank of **A** is effectively very much lower than the number of detector elements.

We refer to this procedure as the coefficients method since the new data matrix may be thought of as consisting of coefficients of the known array manifold basis vectors.

The projection of data onto an alternative basis set of vectors in this way, which may be interpreted as an additional linear transformation of the data, has proved of some use in processing the data from nonlinear arrays of receivers. In the case of a circular array of detectors, operating only in the plane of the array and with the rear semi-circle assumed to be shielded from a source at a particular angle, the data matrix contains a large number of zero entries. If the KTSC as described is understood as producing a sidelobe canceller type of behaviour, then it can be seen that this format of data matrix leads to an insoluble equation if we cancel against a master element which is hidden from the main emitters. Transforming the data by the mapping just described removes such difficulties in addition to decreasing the time spent in computation.

3.9. ITERATIVE PROCESSING

3.9.1. ITERATING THE PSEUDO-INVERSE

The pseudo-inverse method described in section 3.6.2 depends partly on the definition of limits to the angular extent of sources for its ability to enhance the discrimination of close emitters. The proportionate reduction in the width of the system impulse response over that of the BSA matched filter increses slightly as the defined angle of view is restricted for a given signal to noise level. Thus it is possible to use the result of a reconstruction which assumed a wide angle of view to provide a form of prior information for an iterative form of processing of the data in the hope that this will lead to improved discrimination.

Such processing might be carried out by equating values in the output of the initial estimate which fall below a threshold level, ϵ, to that same value, and all remaining values to unity. The resulting function would then be used to weight the corresponding columns of the **A** matrix before forming a new inverse. This will require intensive computation for the processing of each block of data, thus removing one of the major advantages of PIM. It will be

shown by results presented in section 4 that, although this method ensures that sources resulting in measured amplitudes which fall above ϵ are not further attenuated, any low amplitude sources, which might be easily detected by alte native algorithms, could easily fall below a simple threshold based on the sidelobe levels of PIM and thus would be strongly suppressed The ensuing reconstructions would be seriously in error. This approach also results in spectra which contain sharp discontinuities at the switching boundaries.

An alternative approach could perhaps involve setting the below threshold regions of the reconstruction to ϵ, and all other values to a level reflecting the magnitude of the latest reconstruction at that point. It is not clear from our results that this gives any significant gain in discrimination although it may sometimes be used to beautify the result. As we have found in practice, this method also tends to suppress low amplitude sources still further with each successive iteration A similar result would obtain from applying such a weighting to an iterated BSA and would be considerably easier to evaluate

If we interpret the original $\underline{\mathbf{A}}$ matrix as implying a 'top hat' weighting on the space of all possible solutions (ie unity weighting to directions within the 'top hat' and zero weighting to all other directions), the methods described in this section may be interpreted as placing a further emphasis on certain directions within the space of solutions so that the weighting function reflects the shape of the reconstruction at each iteration.

These approaches seem to have much in common with the yet more easily computed practice of squaring the magnitude of results in order to sharpen already noticeable peaks and to suppress sidelobe effects This too gives additional weighting to particular directions in the space of possible solutions, as is borne out by results which we present in Section 4 2 9

3.9.2. USING SUBSPACE ALGORITHMS AS PRIOR KNOWLEDGE

Since a computer cannot store the infinite number of vectors required for a complete description of a continuous space, the array manifold, $\underline{\mathbf{A}}$ is generally stored in discrete form. For this reason the steering vectors used in the initial estimates may be slightly mismatched to the actual source directions It is, however, a simple matter to improve the estimates by iterative interpolation of steering vectors to search for maximum peak amplitude or null depth in the regions identified by the initial coarse processing

Having obtained estimates of the number of signals and their directions of arrival by means of the signal subspace algorithms, these can be interpreted as prior knowledge for further processing designed to extract greater knowledge of the target characteristics, as described in sections 3 6 3 and 3 7 The

precise directional estimates obtained from signal subspace algorithms may also be taken as prior knowledge for the choice of additional steering vectors to be used in the array manifold, $\underline{A}$, as used for the PIM processing described in section 3.6.2

It is also worth noting here that the order in which the data domains are investigated by a sequential implementation of the signal subspace algorithms is of particular importance in certain cases. Singular value analysis of a data matrix may, for example, suggest the presence of two incoherent sources which are not subsequently discriminated by orthogonal filtering in one of the domains. For example, two sources which are highly correlated in the spatial frequency domain, but which are decorrelated in time may not be discriminated spatially, despite the data matrix clearly having a pseudo rank of two. In this case, the sources may be discriminated first of all by examining the time domain singular vectors. The resulting frequency information can then be used to extract corresponding independent data series from the data matrix, enabling the emitter directions to be derived with potentially greater accuracy and ease.

Finally, it should be noted that, whatever the validity of interpretation of the data in terms of point targets (this is unlikely to be a suitable model for all possible radar targets for example), it is undeniably a very much more compact representation than the data matrix itself. The exercise may therefore be viewed in its alternative role as simply resulting in data compaction.

4. COMPUTER SIMULATIONS

4.1. THE SIMULATION FACILITY

A computer simulation facility has been developed in order to test and compare the performance of the techniques described. It enables the characterisation of co-linear and polygonal aperture plane and co-linear focal plane arrays of detectors for the cases of both ideal and perturbed weight vectors. Array manifolds may be analysed in terms of their singular value spectra, and the relevant data stored for future use on a library disc. A menu driven structure allows the creation of target scenarios either for immediate use or for storage on the library disc and experiments have been carried out using simulations of both regularly and randomly sampled sine wave emitters, and also of narrow band f.m. sources. Data generated by running such simulations with the addition of spatially uncorrelated Gaussian noise may be processed by any of the algorithms described in Sections 3.6 and 3.7 of this report and the results may be directly compared by running each technique with identical data samples. A representative selection of the results obtained is presented.

4.2. DISCUSSION OF RESULTS

4.2.1. NOTATION

Our results are generally presented in the form of plots of power against number of cycles. The number of cycles relates to the spatial or temporal frequency of interest, as discussed in Section 3.3.2. Vertical lines surmounted by diamond shapes indicate the signal frequencies and powers used to create the data simulations. Vertical lines surmounted by horizontal bars indicate the signal frequencies and total powers as perceived by the signal subspace algorithms, taking peak centres as the measure of frequency.

4.2.2. SENSITIVITY OF PIM AND BSA TO PRIOR ASSUMPTIONS

The capability of the PIM beamformer to discriminate the presence of even a single source of power depends on the formation of a suitably supported matrix inverse. The angular limits to the field of view must be defined before singular value decomposition can be carried out, and knowledge of the signal to noise ratio is required for the formation of a stable matrix inverse, as will be shown. Strictly, the signal to noise ratio at each component singular frequency must be known in order to find the best possible inverse, but our

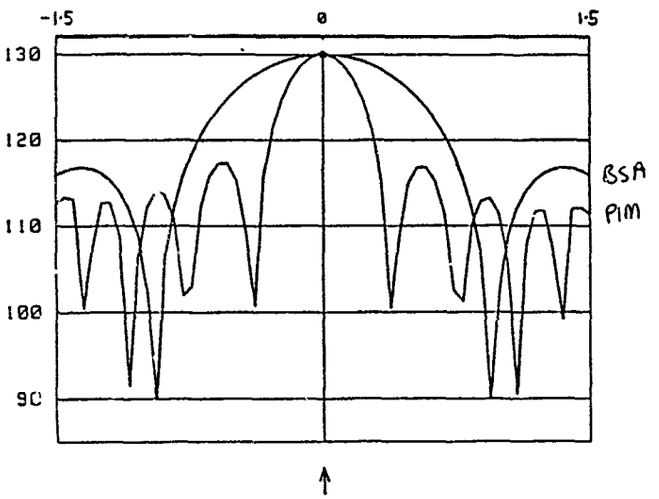

Fig 4.1a. Response of PIM and BSA to a single emitter. Signal to noise ratio is 130 dB. The array is a 20 element linear array, and 20 snapshots of data have been processed.

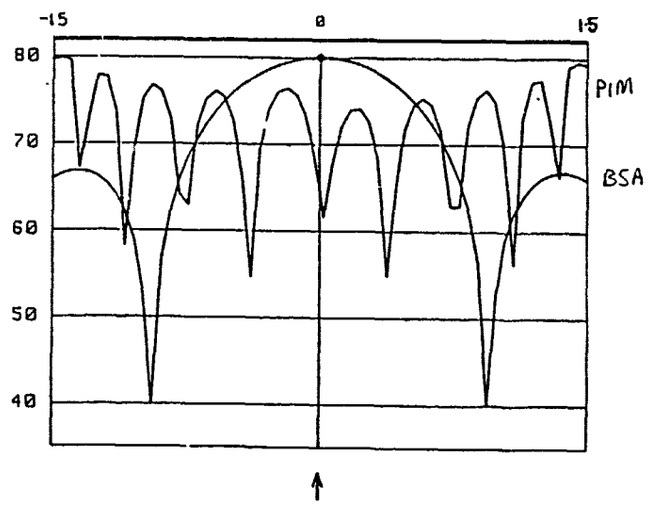

Fig 4.1b. Response of PIM and BSA to a single emitter. Signal to noise is 80 dB, but the PIM reconstruction has assumed a noise level of only -130 dB.

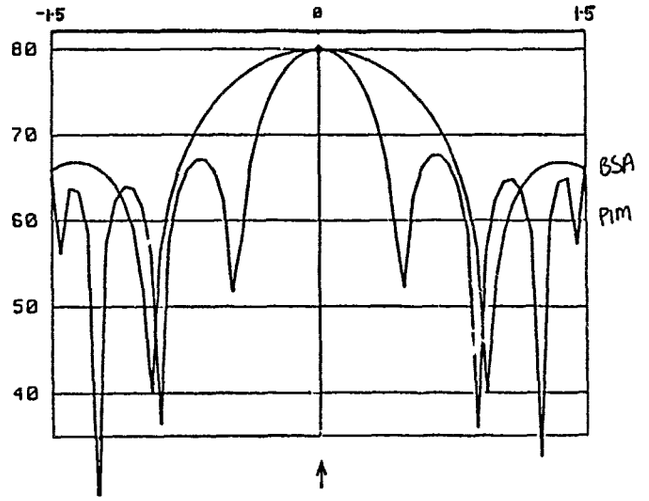

Fig 4.1c. Response of PIM and BSA to a single emitter. Signal to noise is 80 dB. and the PIM beamformer has been calculated appropriately.

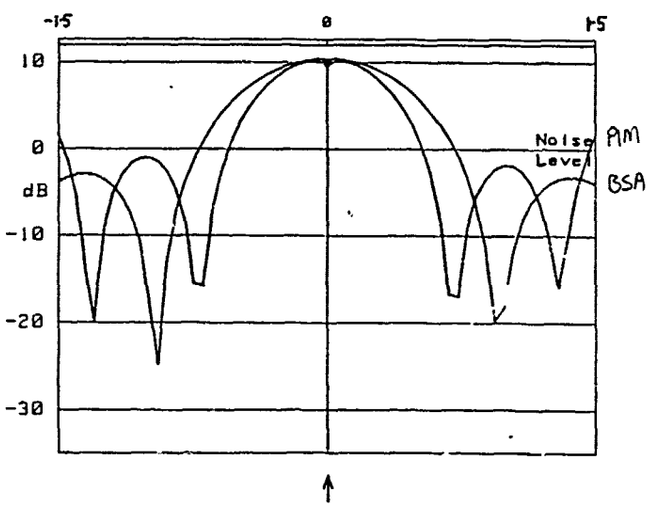

Fig 4 1d Response of PIM and BSA to a single emitter Signal to noise is 10 dB. with appropriate PIM truncation.

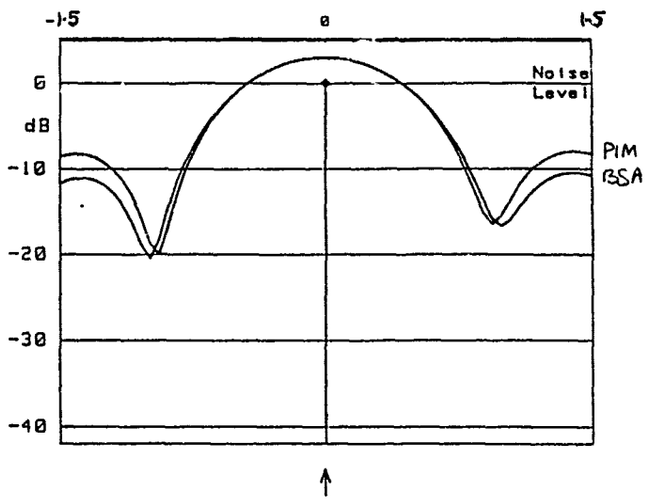

Fig 4 1a. Response of PIM and BSA to a single emitter. Signal to noise is 0 dB. The beamwidths of the two methods are virtually identical

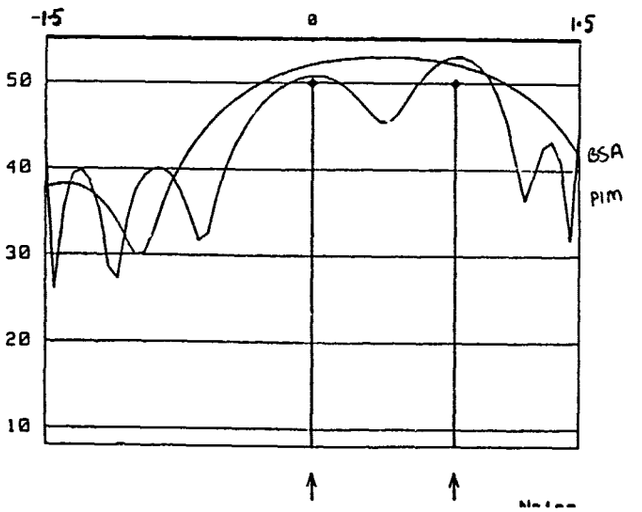

Fig 4 2a Response of PIM and BSA to two equal power sources, separated by 0 8 conventional beamwidths PIM has achieved 'classical' resolution.

simulations have assumed this information not to be available.

Fig 4.1 shows the response of the PIM and BSA beamformers to a single emitter, located broadside to a co-linear array of Shannon number three, for a wide range of signal to noise ratios. We assume initially (Fig 4.1a) that the singular value spectrum of the array manifold has been truncated at a level corresponding to the presence of a noise floor at the unrealistically low level of -130 dB. The response in this virtually noise free environment displays the narrower beamshape of PIM over that of BSA. Clearly, even at this S/N the difference in width is not great.

Fig 4.1b shows the result of an incorrect assumption about the level of additive noise in the formation of the pseudo inverse. The plot of $P_{PIM}(\theta)$ has been computed assuming a noise level of -130 dB, but with a simulation incorporating noise at -80 dB. The result is a meaningless amplification of the noise components of the data. Fig 4.1c shows the result obtained by using a more suitable truncation of the singular values. It is perhaps worth noting that, even at this still very high signal to noise ratio, we have only needed ten of the original twenty array manifold singular values for the reconstruction. At the very much more realistic S/N of 10 dB (Fig 4.1d) and using only the six most significant singular values, the beamwidth of the PIM estimator is appreciably broader than before, and for a S/N of 0 dB, using three singular values, becomes indistinguishable from that of the BSA (Fig 4.1e).

Since the enhancements obtainable from PIM are only clearly observed in conditions of high signal to noise, ensuing simulations in this section will be demonstrated in relatively noise free environments.

First of all, Fig 4.2a demonstrates the classical "resolution" problem, with two emitters of 50 dB, separated by 0.8 cycles across the aperture. In the conventional sense, taking only the major peaks as corresponding to resolved emitter positions, PIM demonstrates superior resolution capability when compared to the BSA. However, Fig 4.2b demonstrates very clearly the difficulty with this conventional view of resolution, when the second emitter is of lower power than the first. The only indication of the presence of the second target, as can also be seen in the BSA plot of Fig 4.2a, is in the asymmetry of the result about the main lobe of the response. We know this because we have prior knowledge about the expected shape of the response to a single emitter ; that we expect a sin(x)/x response from the BSA, for example. Given at least two independent snapshots of the data, a signal subspace algorithm such as MUSIC would hand this example with ease, producing sharp clearly defined peaks corresponding to the source directions.

Moving the lower power interfering source outside the limiting angle of view

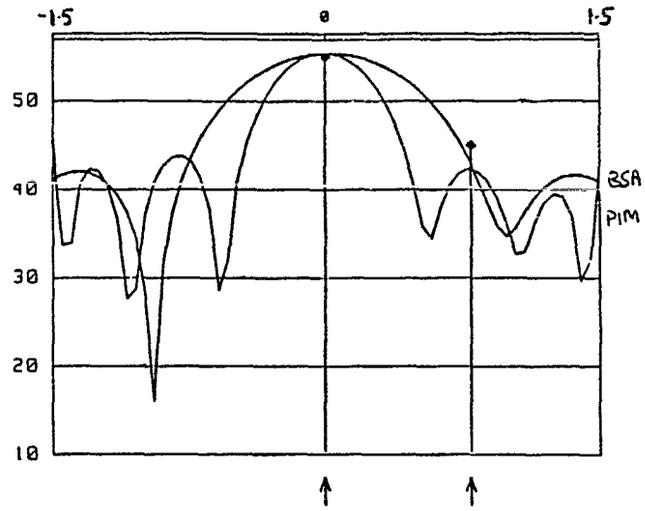

Fig 4 2b Response of PIM and BSA to two unequal power sources, separated by 0 8 conventional beamwidths. The presence of the lower power source is masked by sidelobes of the main response.

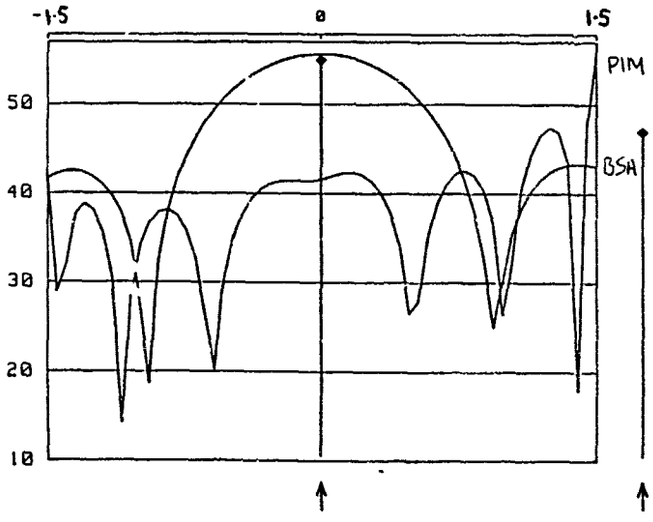

Fig 4 3a The response of PIM and BSA when a second source is present outside the defined angle of view, demonstrating the sensitivity of the PIM method to incorrect prior information.

defined for PIM, we obtain Fig 4.3a Whilst the BSA beamformer remains unperturbed by this, the PIM plot demonstrates once again the extreme sensitivity of this algorithm to the validity of the prior information required for the formation of the matrix pseudo inverse. The directional estimate produced by MUSIC for the central emitter (Fig 4.3b) is extremely precise, but the corresponding estimate of signal power, since it utilises a pseudo inverse supported with only a single delta function, is in error, as demonstrated by the high power attributed to a low level spurious peak.

We now demonstrate how KTSC, or a similar method such as MUSIC, may be used to gain prior knowledge for the support of the PIM pseudo inverse required by the above problem. Since the weight vector derived by the KTSC algorithm is derived from the data alone, the assumption of a narrow field of view required by PIM is not neccessary, and as we saw in Fig 4 3b the accuracy of the directional estimate is therefore not susceptible to such assumptions Simply by scanning a much wider angle of view we may obtain the spatial frequency of the interfering source. This frequency corresponds to a column of a much broader **A** matrix than was used for the PIM estimate. If we simply add this single column to the original **A** matrix, and re-compute the pseudo inverse, the normalised result shown in Fig 4.3c is obtained. The deleterious effect of the interfering source has been all but completely removed For as long as this source remains spatially stationary, the new pseudo inverse may be used to provide a rapid spectral estimate by simple matrix multiplication on each block of data received.

4.2.3. UTILISATION OF THE TIME DOMAIN INFORMATION

If the sources we are attempting to detect emit sinusoidal waves of slightly different frequencies, it is clear that the spatial power distribution simply consists of a three dimensional interference pattern which shifts as the sources appear, on a snapshot by snapshot basis, to change their relative phases This is perceived across the face of the linear array of antennae by a sinusoidal variation of signal amplitude sweeping across the array As this occurs, if we perform a PIM type reconstruction of each snapshot as it is received, we see that the result may be dependent on that relative phase between the two targets. When the pair appears in antiphase as observed at the antenna phase centre, this is manifested as a null in the power distribution across the elements, and a maximised power gradient across the array This enables the beamformer to more easily discriminate the presence of two targets. Conversely, when the sources appear to be in phase, the power at the array phase centre is a maximum, and the power gradient in the local

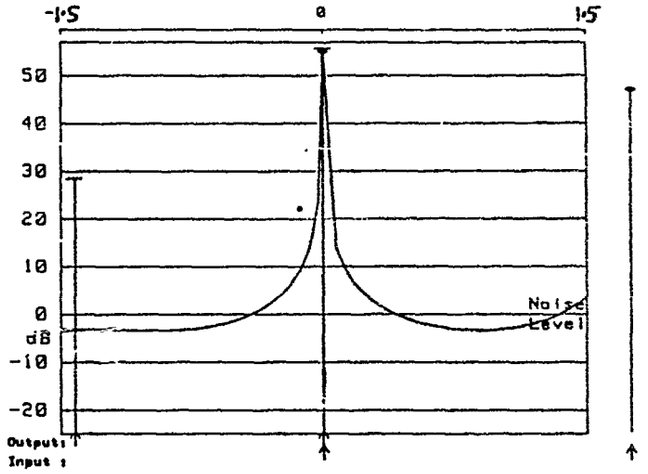

Fig 4.3b. The sensitivity of MUSIC to incorrect prior information. The direction finding ability of this algorithm is unaffected by the presence of the source outside the defined angle of view, but the power estimation, which utilises a pseudo-inverse, is in error.

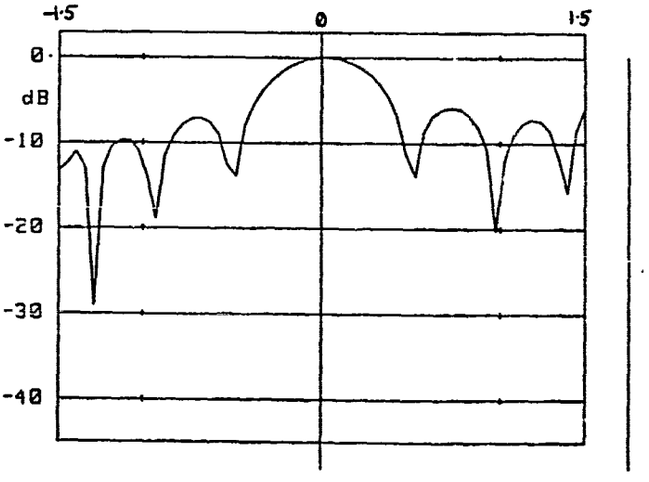

Fig 4 3c. The response of PIM to the sources of Fig 4 3a, but modified by knowledge of the position of the second source.

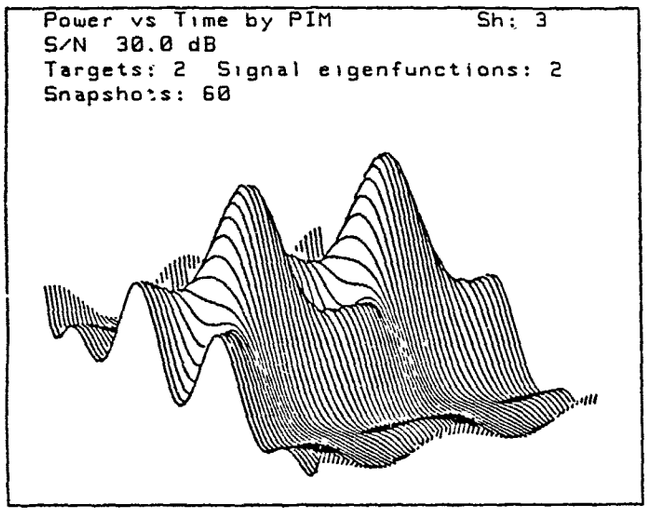

Fig 4.4a. A snapshot by snapshot reconstruction using PIM to demonstrate the effect of relative target phase for two slowly beating sources.

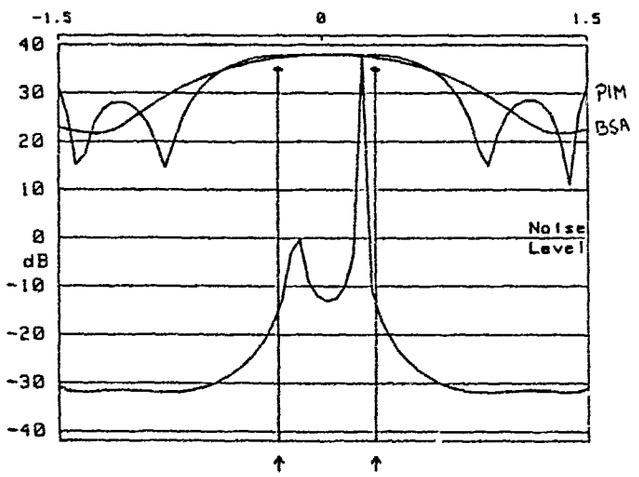

Fig 4.4b The result of PIM and BSA operating on the covariance matrix formed from the sixty snapshots of Fig 4.4a. Data is taken from a twenty element line array of sensors spaced by λ/2. The lower curve is the function $1/(1-P_{PIM}(\theta))$ to show the positions of the major peaks of $P_{PIM}(\theta)$

space is a minimum. Such a situation provides a difficult problem for the beamformer. Fig 4.4a shows a series of such single snapshot PIM reconstructions, with power plotted along a vertical linear axis. Such a plot tells us that there must be more than one point emitter in the object space, but not where the sources might be located, or the possible source powers. We will now attempt to derive this information from our multiple snapshot algorithms.

Combining the snapshots through the formation of a sample covariance matrix, we can plot the multiple snapshot results from the PIM and BSA processors for this example (Fig 4.4b). Also plotted in this figure is the function 1/(1-f(x)), where f(x) is the normalised PIM result. This shows the positions of the main PIM peaks which have just been resolved, and indicates that they do not correspond to the directions from which the signals originated. Fig 4.4c demonstrates that this case is handled with ease by MUSIC, which has determined both signal powers and directions with great accuracy. We conclude from this that examination of the data matrix rather than the array manifold enables superior resolution of the type of problem involving decorrelated point emitters, through better utilisation of the additional information available from the time domain decorrelation (a single point source would not result in the pattern shown in Fig 4.4a).

4.2.4. COMPARISON OF ALL ALGORITHMS ON SAME DATA

We have carried out numerous comparative simulations of the algorithms covered by this report, and present here, as an example, P(e) plots for each of the processors acting on data from a simple scenario involving three emitters (Fig 4.5). The simulation is of a twenty element half wavelength spaced collinear array, with a Shannon number of six. Data has been sampled at a very low rate to simulate the undersampling likely to occur in a real digital to analogue conversion system, and has thus allowed the sources ample time to decorrelate. The sources consist of two 12 dB sinusoidal emitters in close proximity to a 60 dB random phase jamming signal. In all but the BSA and PIM cases we have derived time series from the data which correspond to each of the peaks identified, and analysed them with a ten element moving sub-aperture, as described in section 3.7.4. The eigenvalues of these analyses have been taken as estimates of the power of independent sources, and plotted as arrow heads on the lines corresponding to the respective directions. This enables us to examine the amount of leakage of the jamming signal into the identified signal directions.

Referring to Fig 4.5, plot a shows the response of the MUSIC processor

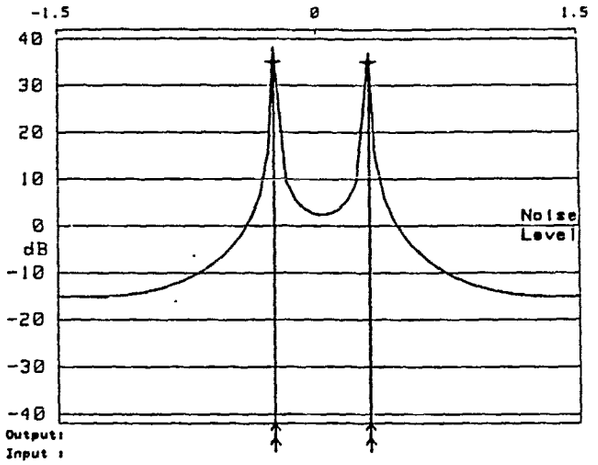

Fig 4 4c The result obtained by processing the data of Fig 4 4a with MUSIC

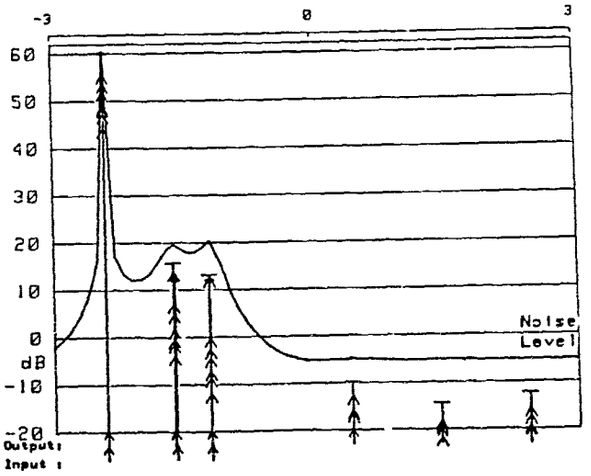

Fig 4 5a Emitters are located at -1 2 -1 6 and -2 4 spatial cycles with powers indicated in the figure by small diamond shapes Twenty snapshots of data from twenty sensors have been processed to determine both source location and power with the results of a time series analysis shown by the arrow-heads, to assess the degree of jammer leakage into source directions The data has been interpreted by MUSIC

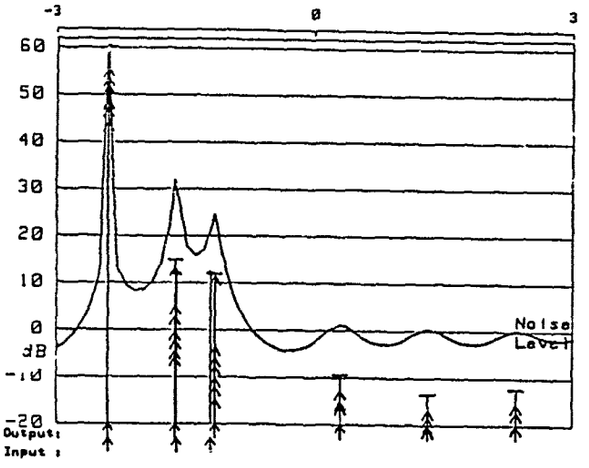

Fig 4.5b. The data generated for Fig 4.5a interpreted by the KTSC1 processor

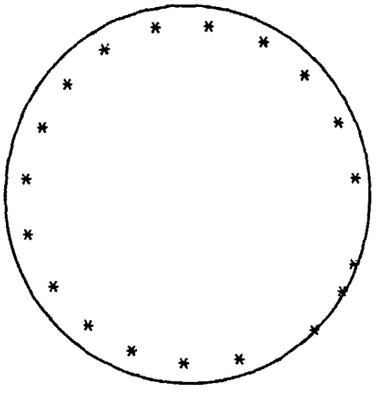

Fig 4.5c. The weight vector generated by the KTSC algorithm may be interpreted as a polynomial in the z-plane. The zeroes of the weight vector produced for Fig 4.5b are shown here. Signal zeroes lie on or close to the unit circle and all extra zeroes are within the circle.

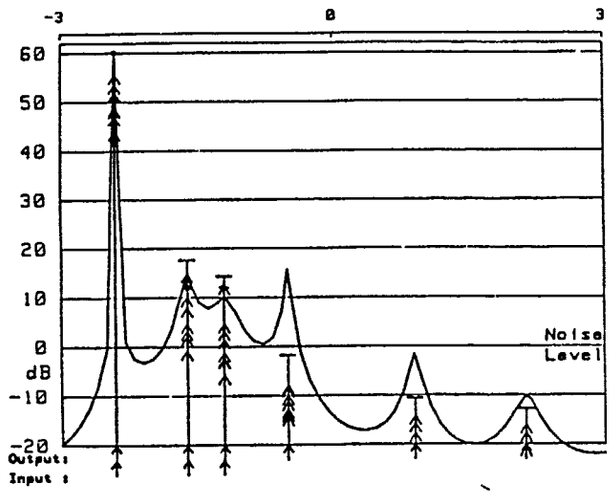

Fig 4.5d The result produced by MEM acting on the data of Fig 4.5a.

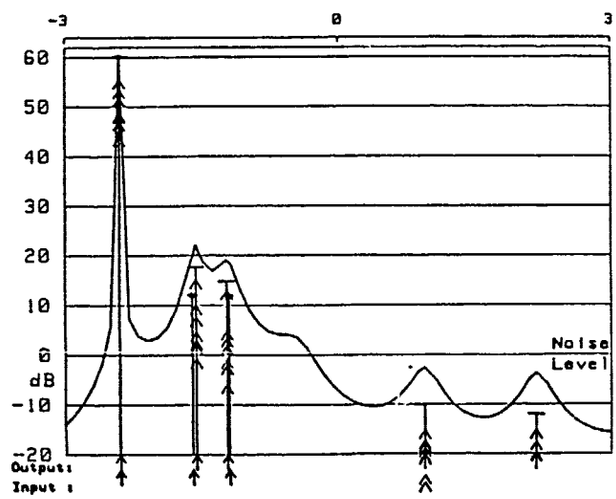

Fig 4.5e The data generated for Fig 4 5a processed by MLM

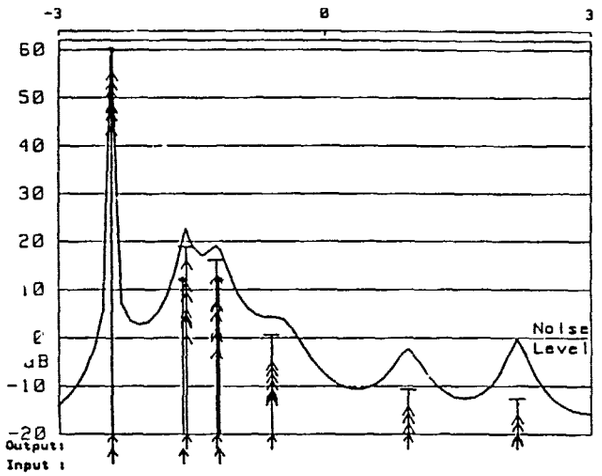

Fig 4 5f The data of Fig 4 5a processed by TNA

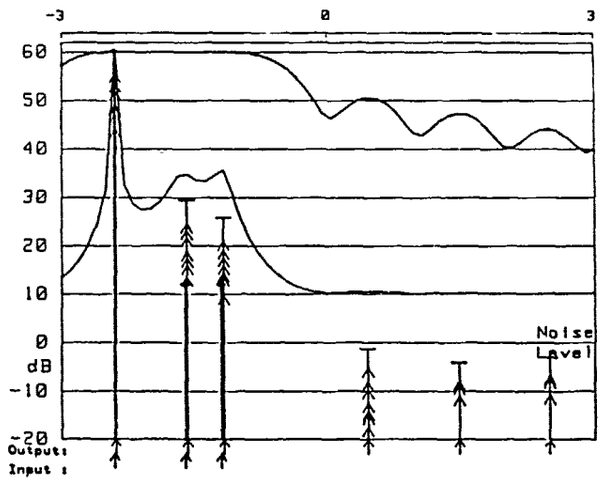

Fig 4 5g The results obtained from processing the same data as the previous figure, using the PCB algorithm (shown by the upper curve) The lower curve shows the function $1/(1-P_{PCB}(\theta))$ to demonstrate its similarity to the MUSIC curve of Fig 4.5a

to this simulation. The two signal directions are accurately determined, although only just, as indicated by the shallowness of the dip between the relevant peaks The area away from the signal sources is depicted as being very flat, as a result of the averaging of the seventeen equally weighted noise eigenvectors used by MUSIC in this example. The three very shallow peaks identified in this region are found to be associated with total power levels significantly below the noise level, and can safely be ignored. The cluster of high eigenvalues derived from the time series in the direction of the jammer indicates the presence of a locally very high noise floor. In each of the signal directions, a single eigenvalue stands out higher than the rest, indicating the presence of the sinusoidal emitters in a heavily rejected noise background from the jammer.

Plot b depicts the result of processing the same data by the KTSC1 algorithm. Utilising exactly the same eigenvectors as did MUSIC, we obtain a different result. The signal sources have resulted in two clearer peaks, although the particular sample of noise present in the data has resulted in slightly poorer estimates of direction in this case. The background ripple is more pronounced than in MUSIC, although such peaks may still be rejected on the basis of the low power levels detected there. Interpreting the KTSC1 weight vector as a polynomial in the z-plane (Kumaresan and Tufts (1983)), plot c shows the stable regular distribution of zeroes around the interior of the unit circle The signal zeroes lie almost on the circle.

Plot d shows the result obtained using MEM on this data sample. Although both signal and jammer directions have been accurately determined, note the typical presence of pronounced peaks at other directions. The proximity of one of these peaks to the signal frequencies, and the consequent relatively high power level associated with it has resulted in higher levels of noise leakage from the jammer into the signal directions as perceived through the incorrectly supported pseudo inverse. A simple iteration which removed the spurious peaks from the support for this operation would result in improved power estimation

MLM (plot e) has tended to average out the major spurious peak, at the expense in this case of less accurate determination of the signal directions. The result of TNA processing, shown in plot f has, not surprisingly, similarities to both MEM and MLM.

Plot g shows both the PCB estimate and the function $1/(1-f(x))$ derived from equating $f(x)$ to the normalised matched filter estimate. This result clearly emphasises the fact that the visual appearance of an estimate is not a reliable indicator of its ability to resolve targets in the sense defined by Rayleigh. The inter-peak dips of the matched filter form of this estimate are of only a

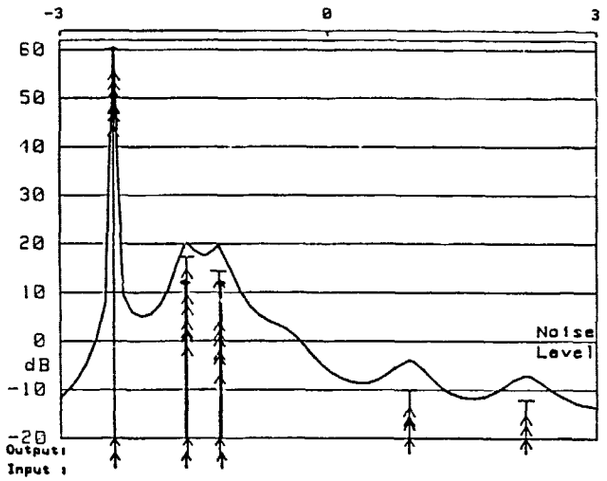

Fig 4 5h The results shown are those derived by the EMU processor, which weights the noise eigenvectors with the inverse eigenvalues The data is as used in Fig 4 5a

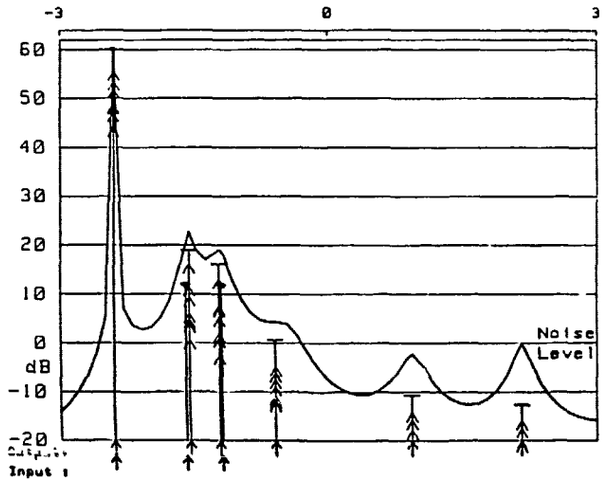

Fig 4 5i The data used for Fig 4 5a processed by the Pisarenko (PME) method

fraction of a dB, and yet they are stable because of the removal of noise in the rank reduced sample covariance estimate. The major error evident in this plot (the hopelessly inaccurate estimate of signal powers) has arisen not as a result of a deficiency of the PCB algorithm, but through the inability of the simple peak locating routine employed to accurately find the centre of the extremely broad response. This has caused a slight inaccuracy in the pseudo inverse support for the important jammer direction, and hence significant leakage of this relatively high power signal into the signal time series. This only serves to demonstrate, however, the high degree of cancellation which may possibly be achieved by this form of processing if properly executed.

$P_{EMU}(\theta)$, shown in plot h is very similar to the MUSIC result in this case, but shows by the presence of more pronounced peaks in the background floor that the flat floor of MUSIC is partially caused by the assignment of uniform weighting to each of the eigenvectors in the noise subspace. The plot obtained from the PME estimator (plot i) demonstrates that the flat floor of MUSIC is also the result of averaging many noise eigenvectors. All targets are clearly detected by all the algorithms so far discussed, and their spatial frequencies determined to within a fraction of a beamwidth. This behaviour is typical.

Returning to the BSA and PIM beamformers, we see that the results (plot j) are limited in application because of the high sidelobe levels. Although the BSA plot, for example, contains exactly the same information as the MLM estimate, it is not presented in a useful fashion for the location of multiple point target emitters, because the prior information that such sources result in a $\sin(x)/x$ response from the correlation action of the filter has not been used : the process of parameter estimation is only half completed. We will see that MUSIC and its related processors are capable of effecting this task of fitting multiple $\sin(x)/x$ functions optimally to the data when we examine the results obtained from a focal plane array which has an array manifold consisting of just such waveforms.

Finally in this section we present the result of processing the compacted data which results from the coefficients method discussed in section 3.7.2 by the MUSIC algorithm. The original data is exactly the same as used for all the other plots, but has been projected onto only the most significant ten eigenvectors of the array manifold (corresponding to a S/N ratio of 50 dB in the case of a Shannon number of six). The eigen analysis required for MUSIC is thus significantly faster for the resulting (10 × 10) covariance estimate than for the original (20 × 20) matrix. Apart from a slight discrepancy in the estimated signal powers, the result plotted in Fig 4.5k is virtually

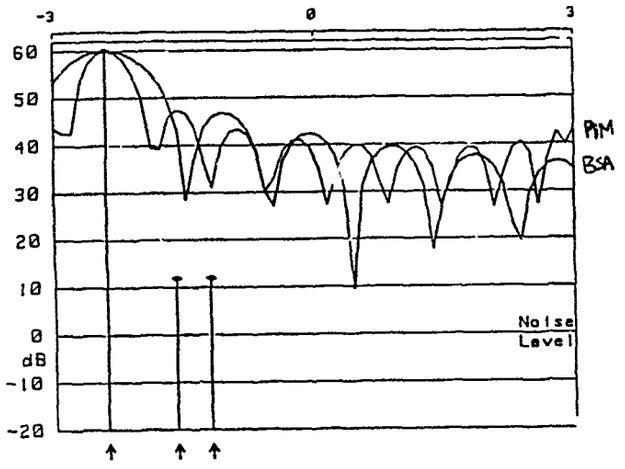

Fig 4 5j. The data used for Fig 4.5a processed by PIM and BSA. No time series analysis has been performed for this figure.

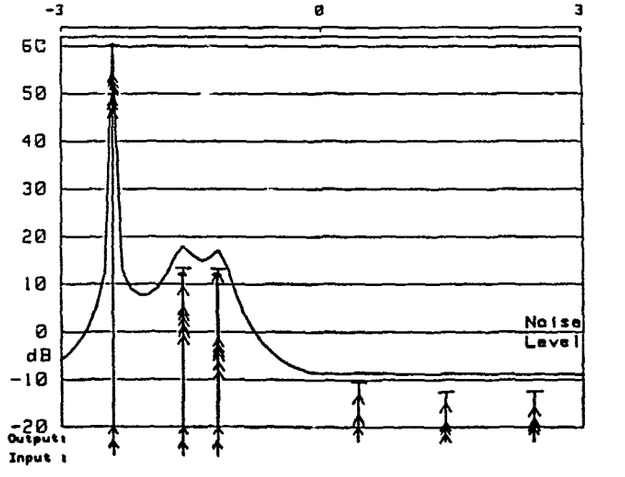

Fig 4 5k. The data generated for Fig 4.5a has been pre-processed by the coefficients method before being analysed by MUSIC in order to decrease the computation time. The original twenty elements of data per snapshot have been reduced to ten.

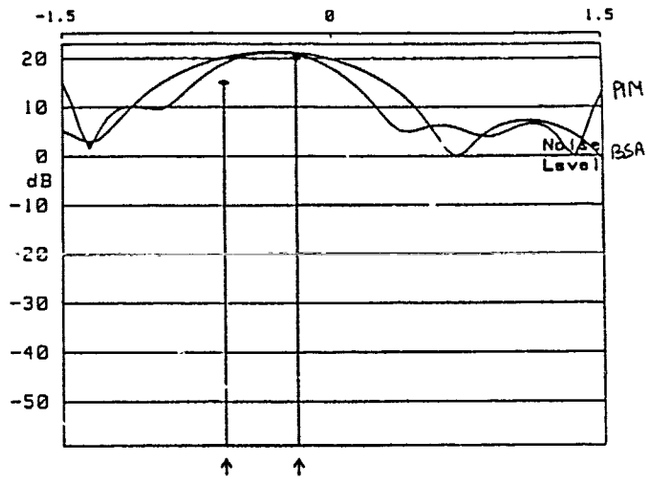

Fig 4 6a Simulation of a coherent multipath problem The thirty snapshots of data are generated by simulating two coherent emitters separated by 0.4 cycles measured across the twenty sensor aperture The results plotted are those given by BSA and PIM

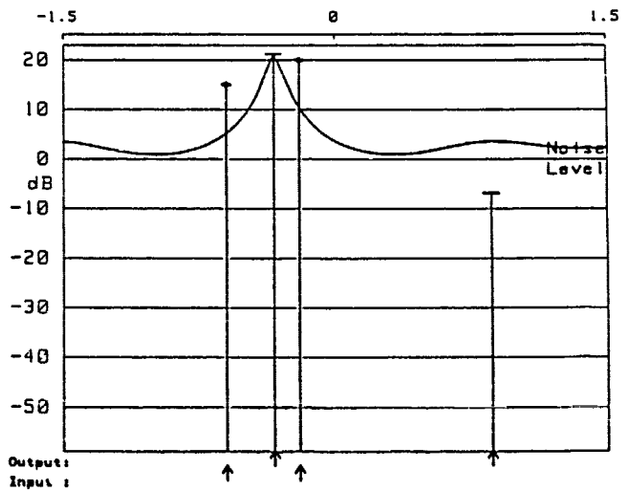

Fig 4 6b The KTSC1 algorithm acting directly on the data of the previous figure is unable to discriminate the two signals.

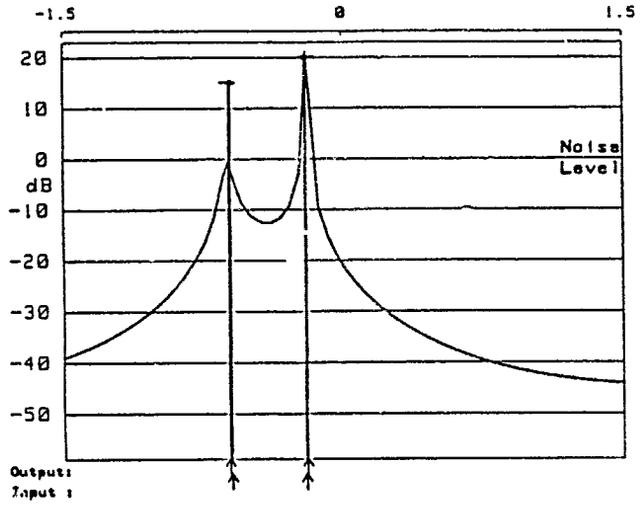

Fig 4 6c. Utilising a sub-aperture of ten elements to pre-process the data, the KTSC1 algorithm succeeds in discriminating the two coherent signals

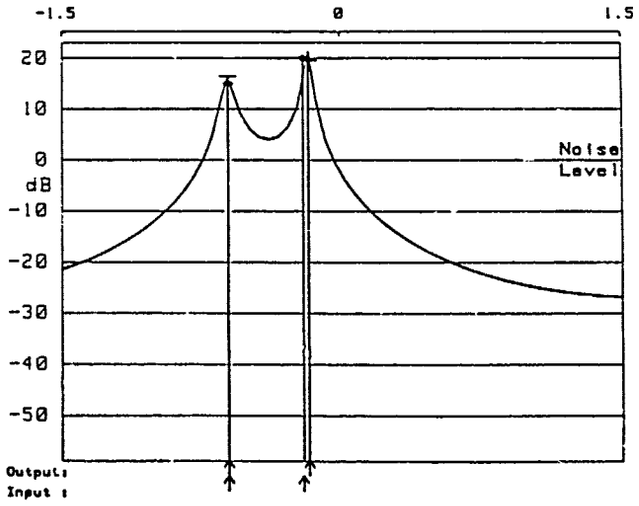

Fig 4 6d Taking only a single snapshot of data from the twenty element array, and pre-processing with a sub-aperture, the KTSC1 algorithm again succeeds in separating the two coherent sources of Fig 4 6a This result is more sensitive to noise than that based on thirty snapshots of data

indistinguishable from that displayed in plot a.

4.2.5. DEALING WITH COHERENT EMITTERS

Our next simulation deals with perhaps the most difficult discrimination problem of all, that of two coherent targets which appear to have a constant phase difference at the antenna phase centre. This could be the result of coherent multipath reflection from a fixed surface. We have simulated a source of 20 dB and a reflection of 15 dB separated by 0.4 conventional Rayleigh beamwidths, and with a relative phase difference of $\pi/3$ radians measured at the centre of a 20 element linear array of half wavelength spaced elements.

Fig 4.6a shows the result of processing twenty snapshots of information from this scenario. The effect of processing many snapshots in this way is simply to make the resulting estimate relatively stable against noise, and does not afford any improvement in discrimination performance. As can be seen, the BSA and PIM beamscans result in very similar responses at this low S/N ratio, and neither gives a clear indication of the presence of the pair of signals.

Fig 4.6b shows the result obtained from the KTSC1 processor acting on precisely the same data. The sample covariance estimate has a single high eigenvalue, and the reconstruction produces a single peak between the two sources, and slightly closer to the main signal.

If we now employ the sub-aperturing technique to pre-process the data covariance matrix, as discussed in section 3.7.3, the new covariance estimate has two significantly high eigenvalues. Forming the appropriate weighted sum of noise eigenvectors using the KTSC1 algorithm produces the minimum energy solution to the problem given in Fig 4.6c. Both signals are clearly and accurately located, and their associated powers determined. Employing forward-backward shifting of the sub aperture gives enhanced noise rejection and an even greater degree of precision in the frequency estimation. In fact, because the data in each of the snapshots is exactly the same but for the noise and an amplitude scaling, the sub aperture may be applied to any individual snapshot, although with resulting greater sensitivity to noise perturbations. Fig 4.6d shows the SAT operating successfully on single snapshot data from the even more difficult situation when the two signals appear to be in phase at the array phase centre.

Source	Power (dB)	Position (cycles)
1	9	−1.1
2	5	−1.1
3	12	−0.4
4	15	−0.4
5	17	−0.4
6	4	0.4
7	11	0.5
8	50	1.2
9	45	1.3

Table 4 1. Source powers and locations used in the simulations discussed in section 4 2.6 and illustrated by Fig 4.7(a-f).

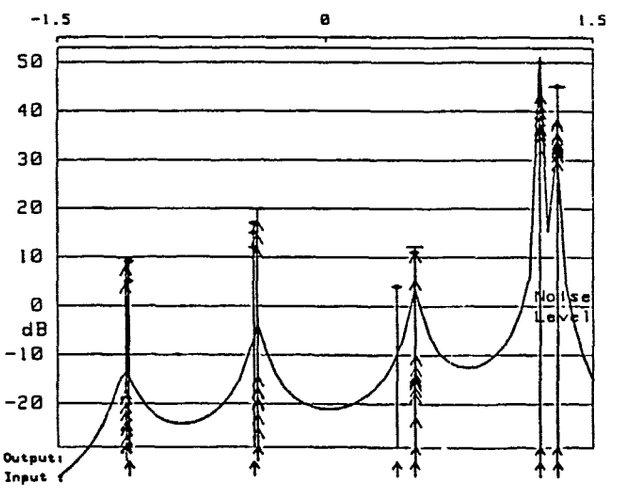

Fig 4 7a. See Table 4.1 for details of the sources modelled for this simulation. The array is again of twenty elements, separated by λ/2, and thirty snapshots of data have been processed. The algorithm used for this reconstruction was KTSC1, followed by a time series analysis.

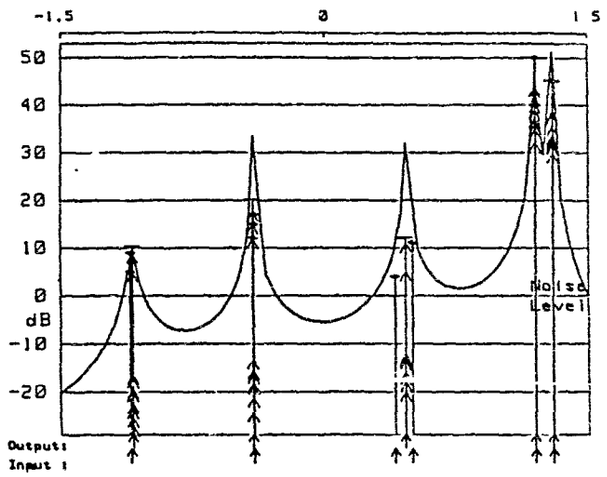

Fig 4.7b Sources as listed in Table 4.1. Data has been analysed by the forward-backward KTSC1 method.

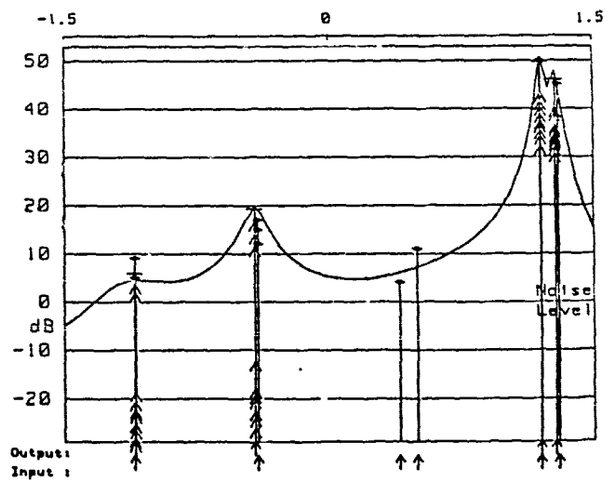

Fig 4.7c Sources as listed in Table 4.1 Data has been processed by MUSIC

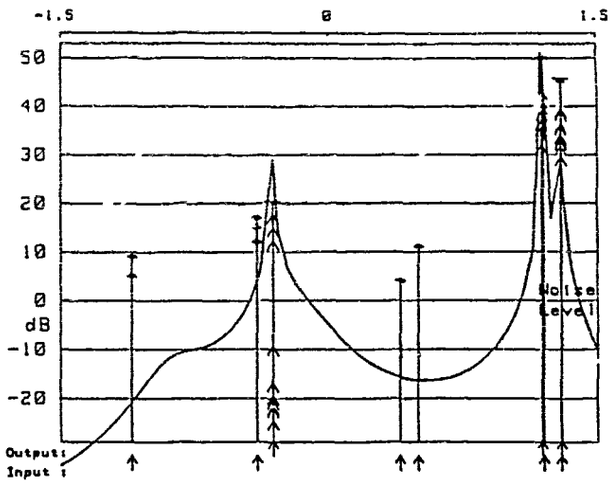

Fig 4 7d. Source details are given in Table 4.1. The result is that given by MEM.

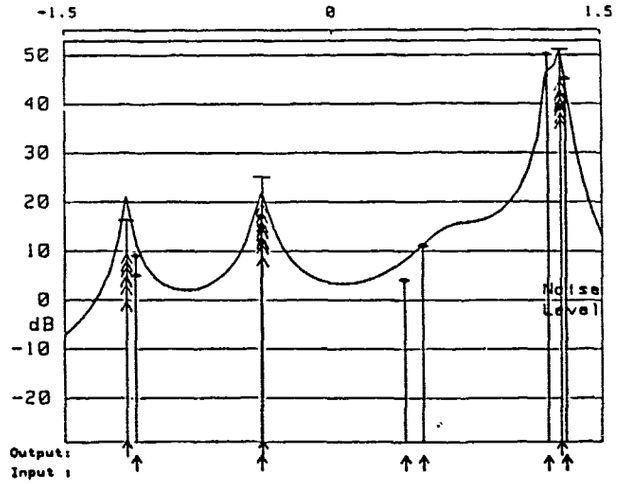

Fig 4 7e Details of emitters are given in Table 4.1 Data has been processed by KTSC2.

4.2.6. MULTIPLE TARGET DISCRIMINATION

Fig 4.7 presents the results of a number of simulations which have attempted to locate the powers and spatial frequencies of seven signal sources and two random phase jamming emissions located within a space of only three Rayleigh beamwidths, using a twenty element linear phased array. Processing has been carried out sequentially, to initially locate the sources in space, followed by sub aperture analysis of the resulting time domain information Details of the sources are given in Table 4 1.

Figs 4 7a and 4.7b show the results of processing by the KTSC1 algorithm on the forward data and on the forward backward data respectively. In both instances the algorithm has made a commendable estimate of the required parameters. Neither processor has discriminated the two sources lying closest to the jammers in space as individual peaks, although the subsequent analysis of the time series associated with the single peak has accurately determined the presence of two emitters and their powers.

In this example, MUSIC (Fig 4.7c), operating on an unweighted combination of the noise space basis vectors, has not performed as successfully as KTSC1. It has completely failed to register the presence of the two sources at 4 and 11 dB, and has only just recognised the sources located at -1 1 cycles MEM (Fig 4.7d) displays considerably worse discrimination of the sources in this spatial analysis.

Fig 4 7e is interesting, as it shows again the errors which arise from inaccurate estimates of the positions of high power sources. This result has surprisingly been derived by the KTSC2 processor, and has notably failed to discriminate the main jamming source. Such a response has resulted in incorrect matching of the pseudo inverse to the problem, and the poor cancellation of the random jamming signal in target directions

Finally, in Fig 4.7f, we show the unsuitable responses given by the BSA and PIM algorithms to this data sample No attempt to process doppler domain information has been made.

4.2.7. FOCAL PLANE PROCESSING

We now show the results of processing a 'simple' high signal to noise problem involving data received at seven equally spaced elements in the focal plane. The array samples spatially at the Nyquist rate The data has been generated from a model of two emitters separated by 0 3 beamwidths spatially, and by only 0 1 cycles across the time domain aperture in order to emphasise the sensitivity of the successful algorithms. The other intention here, however, is to reveal the unsuitability of the KTSC algorithms for handling such data, as

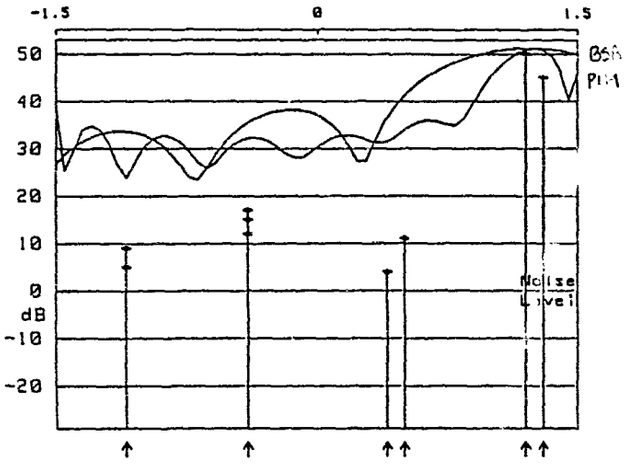

Fig 4.7f. Data generated by a simulation of the sources listed in Table 4.1 has been analysed by PIM and BSA. No subsequent time-series processing has been attempted.

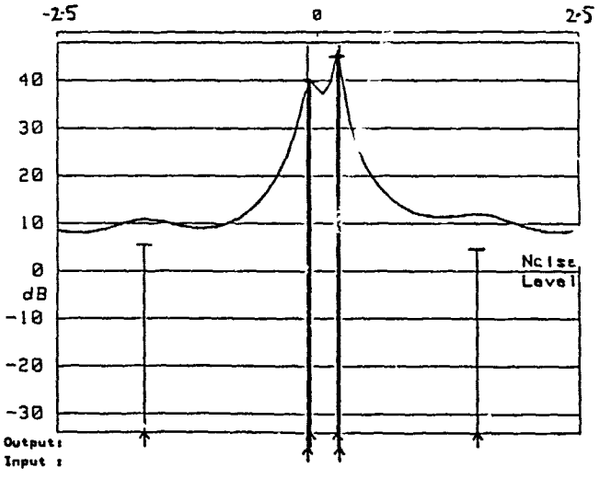

Fig 4.8a A simulation of data received from two point sources at 0.3 beamwidths separation by a seven element Nyquist spaced focal plane antenna has been processed by MUSIC. Sources are separated by only 0.1 cycles in the time domain aperture. Twenty snapshots of data have been used in the covariance estimate.

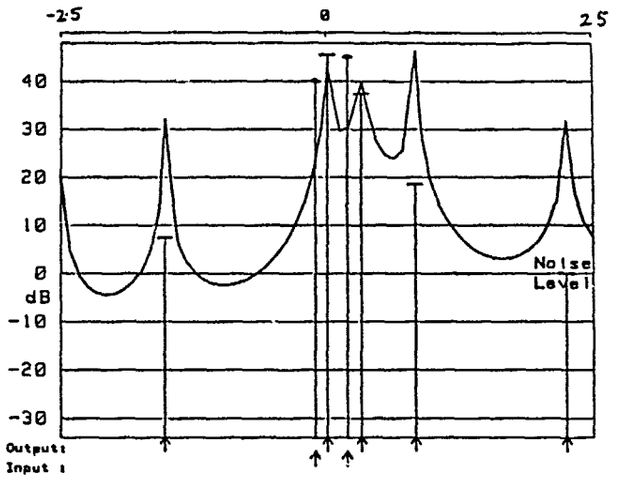

Fig 4 8b The data of Fig 4 8a processed by KTSC1, showing multiple spurious peaks.

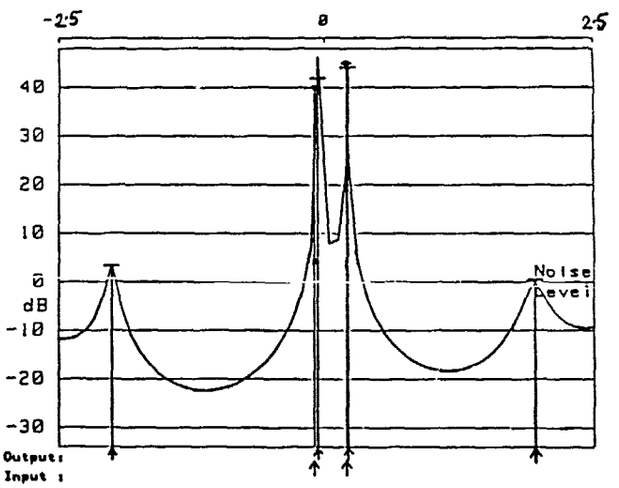

Fig 4 8c The data of Fig 4 8a processed by MEM

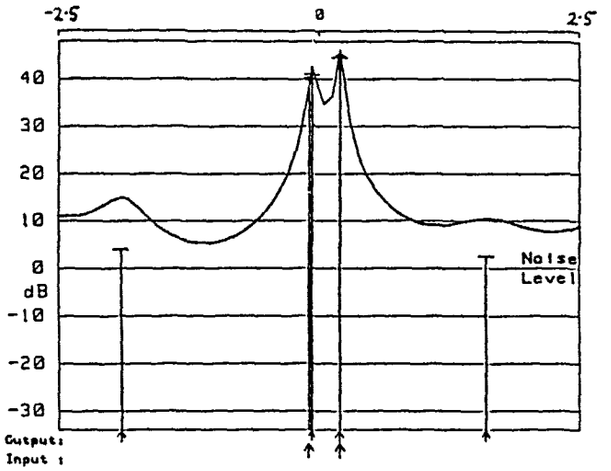

Fig 4.8d. The data of Fig 4.8a processed by MLM.

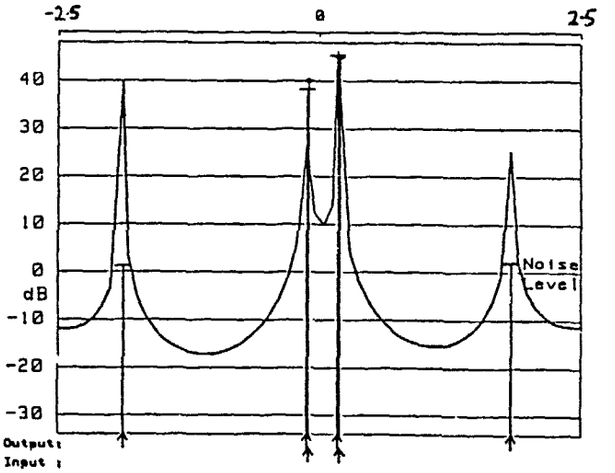

Fig 4.8e The data of Fig 4.8a pre-processed by the coefficients method before analysis by KTSC1 The number of spurious peaks is reduced.

a result of the form of constraint employed.

Fig 4 8a indicates that MUSIC interprets the data with ease, despite the close proximity of the sources in the two dimensional space-time aperture Spurious peaks are well suppressed It is also worth noting that MUSIC is effectively interpreting the data in terms of sin(x)/x functions across the focal plane The response of the KTSC1 processing is depicted in Fig 4 8b, and displays a large number of spurious peaks, in addition to incorrectly estimating the locations of the signals. This is as a result of the unsuitable form of constraint employed by this particular form of the signal subspace algorithm, which requires that all receiving elements in the sensor are equally weighted and irradiated Both MEM and MLM (Figs 4 8c and 4 8d) give superior results, when compared with KTSC1, although do not achieve the same accuracy as MUSIC in target location and power estimation Again, neither algorithm, as implemented here, is optimal for the problem of focal plane data

Fig 4.8e shows the result of analysis using the coefficients method of data compaction followed by the KTSC1 algorithm. Although the spurious peaks are pronounced, the estimated signal powers at these directions are relatively lower than determined by the previous analysis of the data matrix itself, and the locational estimates of the signal sources are considerably more accurate The coefficients representation of the data thus appears to have made the resulting total algorithm less sensitive to the original data format Although it is not presented here, we note that the coefficients/MUSIC method produces identical results to MUSIC alone

Finally in this section, we demonstrate how the same seven element focal plane array can locate at least seven emitters within the five beamwidth field of view The scenario consists of five partially decorrelated sources and two high power random phase jamming signals, as listed in Table 4.2 Fig 4 9a shows the eigenvalue spectrum of the sample covariance estimate MDL has determined a signal subspace dimension of four, and so MUSIC has operated on the remaining three noise subspace eigenvectors to produce the result shown in Fig 4.9b Four of the five directions have been determined to within a fraction of a beamwidth, and subsequent doppler processing of the data has evaluated the powers associated with each of the signals, and demonstrated the excellent rejection of the jamming signals in the source directions

Source	Power (dB)	Position (cycles)
1	9	−1.5
2	5	−1.5
3	15	0.25
4	12	0.25
5	25	0.2
6	55	1.2
7	50	0.9

Table 4.2. Source powers and locations used in the focal plane array simulation discussed in section 4.2.7 and illustrated by Fig 4.9(a-b).

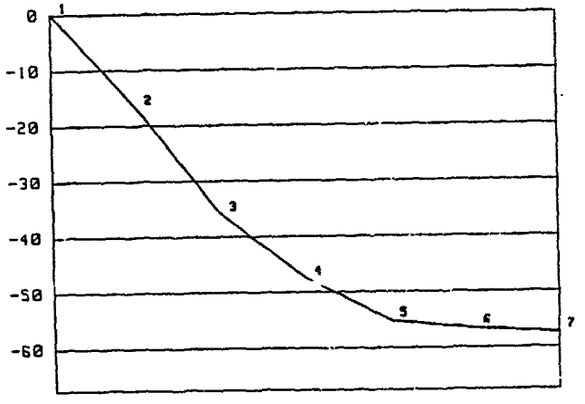

Fig 4.9a The sources listed in Table 4.2 have been simulated to give data from an array of the same specification as given in Fig 4.8a. The diagram shows the eigenvalue spectrum. MDL analysis indicates that the first four eigenvectors define the signal subspace.

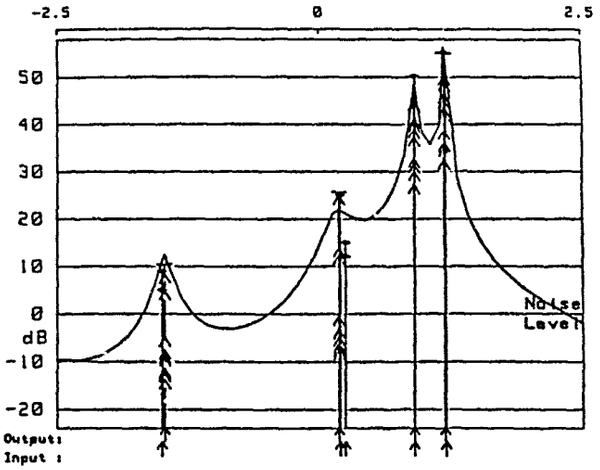

Fig 4.9b Reconstruction of the sources listed in Table 4.2 as attempted by MUSIC and subsequent time series analysis.

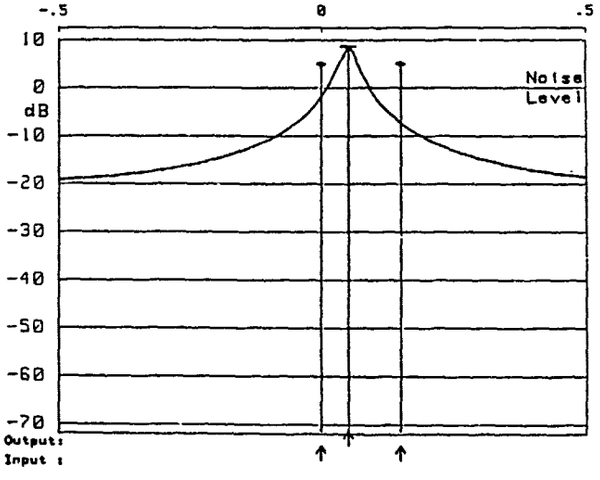

Fig 4 10a The diagram shows the result of processing a simulation of two 5dB sine wave sources, separated spatially by 0 15 cycles and by 0 9 cycles in the time domain aperture The twenty snapshots of data have been analysed by the KTSC1 method, solving the equation $\underline{\phi}_s^H w = 0$, where $\underline{\phi}_s$ is the matrix of spatial domain signal eigenvectors Although the signal subspace dimension was clearly two, the result contains only a single peak

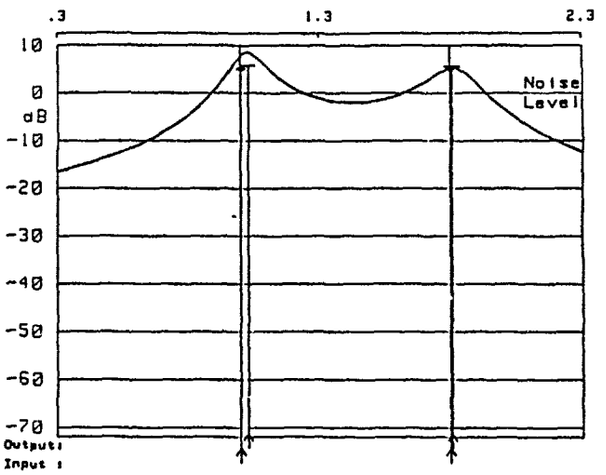

Fig 4 10b. Processing the time series associated with the single direction obtained from Fig 4.10a, using a moving window/KTSC1 algorithm, allows us to determine the two signal frequencies.

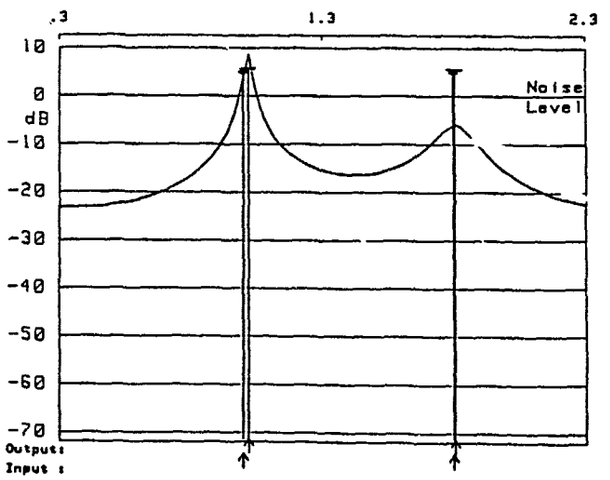

Fig 4 10c. The data matrix analysed in Fig 4.10a may alternatively be used to determine the time domain signal subspace eigenvectors, $\underline{\gamma}_s$. This matrix may then be used to derive the time domain estimate shown here by solving the KTSC1 equation, $\underline{\gamma}_s^H \underline{w} = 0$.

4.2.8. IMPPOVING ANGULAR ESTIMATES THROUGH DOPPLER PROCESSING

The results given in Fig 4.10 demonstrate that the apparent limits of angular discrimination for a particular algorithm may be exceeded if the sources may first be discriminated in another domain, such as the doppler domain in this example We have modelled two 5dB sources separated across the spatial aperture by only 0 15 cycles, but by 0.9 cycles in the time domain aperture.

Eigenanalysis of the data covariance estimate gives a signal subspace dimension of two Analysis using the two singular vectors corresponding to the dimension of the antenna aperture, using KTSC1, produces only a single peak, approximately centrally located between the two sources (Fig 4.10a) because of the correlation between the vectors of the array manifold A moving window analysis of the single time series derived from this directional estimate reveals the frequencies of the two sources with only a slight error (Fig 4.10b). If, however, our interest is to locate the sources spatially, this result is not suitable. We now show that, because of the decorrelation present in the time domain, it is possible to find the directions of the emissions with some degree of accuracy

Instead of using the spatial domain singular vectors, we first of all use the two signal space vectors in the time domain Processing, as before, with the KTSC1 algorithm, we determine the two frequencies to within a fraction of a cycle (Fig 4.10c). Subsequent analysis of each of the data series associated with these frequencies in turn enables the locations of both of the emitters to be found Such analysis may be carried out using the familiar moving window, followed by eigen analysis and determination of the appropriate signal subspace singular vectors (Fig 4 10d). Alternatively in this situation, where each independent data series corresponds to only a single target, the targets may be located by a simple BSA examination of each appropriate series. Fig 4 10e shows plots of the function $1/(1-f(x))$, where $f(x)$ is the normalised BSA result, to show the peak positions more clearly. This latter method is clearly much faster to compute than the moving window, but is obviously not capable of finding multiple sources at a particular frequency with the same accuracy

4.2.9. ITERATIVE PIM PROCESSING

Using an example of a twenty element circular array of detectors, looking only in the plane of the antenna, examples of the results of two different forms of iteration on the PIM and BSA beamformers are presented in Figs 4.12 and 4.13.

The target scenario consists of three low amplitude sources sufficiently

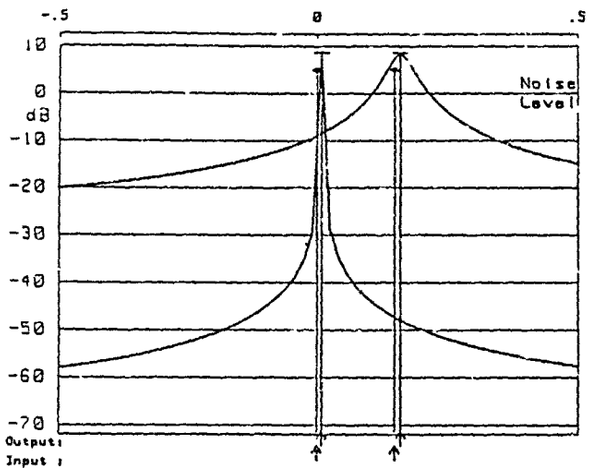

Fig 4 10d Utilising the knowledge of the signal frequencies derived in the previous figure the spatial locations of the two emitters may be determined independently using a moving window/KTSC1 algorithm on the data associated with each frequency in turn.

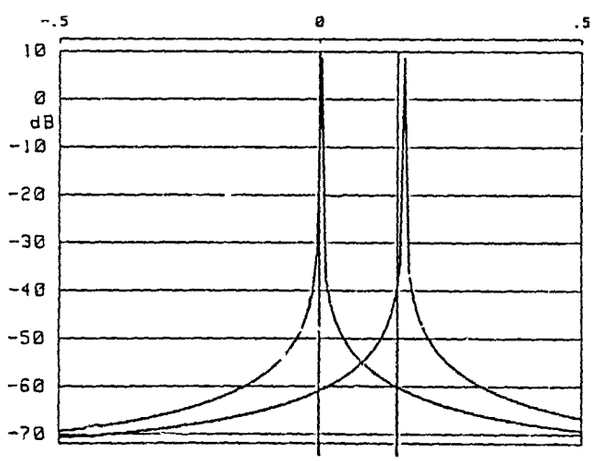

Fig 4 10e Given the estimate of signal frequencies derived in Fig 4.10c, the spatial locations may be determined by a simple BSA operation on the data derived from each in turn. The curves plotted here are of the function $1/(1-P_{BSA}(\theta))$. In order to show the peak positions more clearly

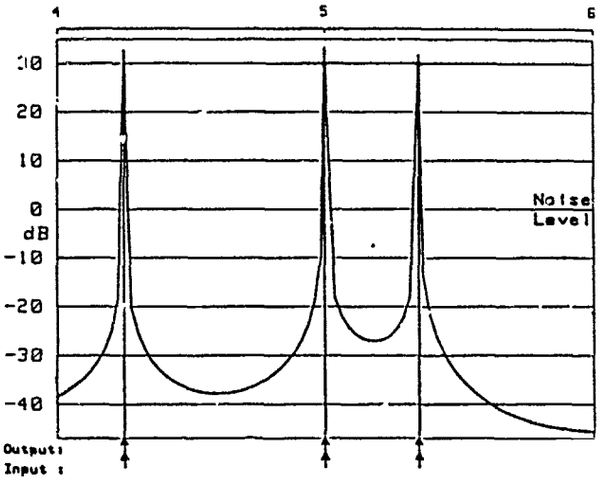

Fig 4.11a. The signals received at a 20 element five wavelength radius polygonal array over a period of 30 snapshots have been interpreted here by MUSIC. Signal locations and powers are accurately determined.

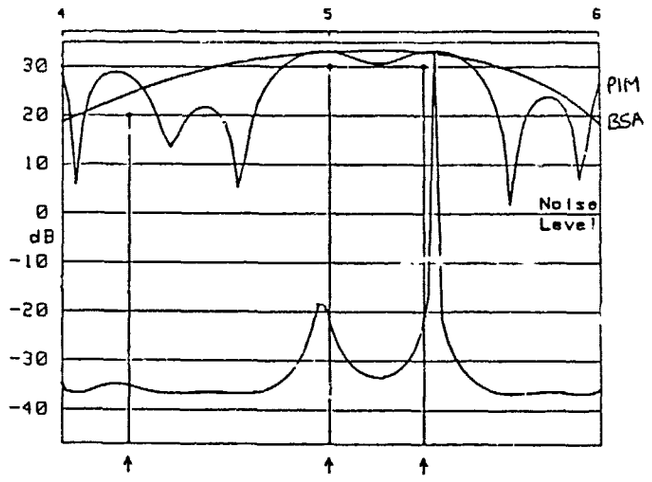

Fig 4 11b The same data as used for Fig 4 11a have been analysed by PIM and BSA to give the results shown here The lower curve shows the function $1/(1-P_{PIM}(\theta))$ and indicates the positions of the main peaks of the PIM response The PIM beamformer has achieved 'classical' resolution

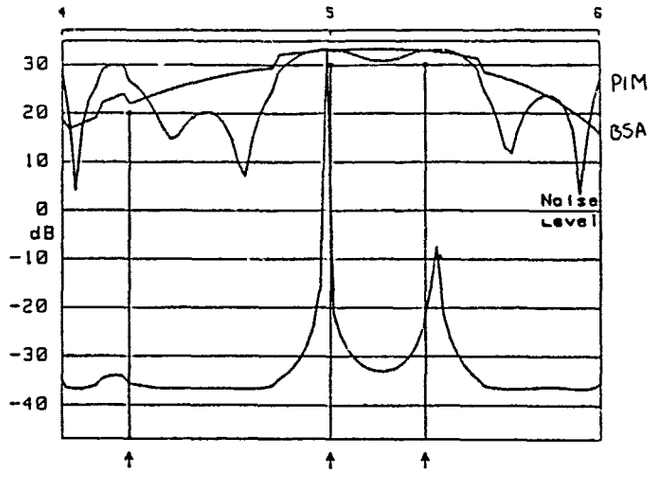

Fig 4.12a. This figure shows the first stage of an iterated reconstruction of the targets. A weighting function has been applied to the original array manifold, which assigns unity weight to those regions of the original picture which fall above a threshold, as described in the text.

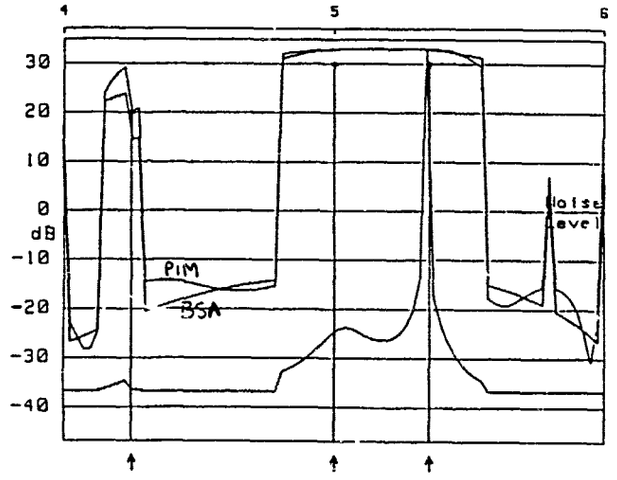

Fig 4.12b The eighth stage of the iteration started in Fig 4.12b. Showing the presence of a spurious peak caused by 'incorrect' thresholding at one point

separated in aperture that they are initially just discriminated by PIM, but not quite by the BSA The result of processing the data by MUSIC using the normal pre-iteration array manifold is given for reference in Fig 4 11a, and compares with the equivalent PIM and BSA responses in Fig 4 11b. The latter figure also shows the function 1/(1-f(x)) derived from the PIM estimator

Whereas MUSIC has easily determined both signal locations and powers, the BSA algorithm has failed to produce a clear indication of anything, having only a single broad response PIM has just managed to separate the two 30 dB signals, and the possible presence of the third signal is indicated by a peak only slightly above the level of the sidelobes, but at an incorrect direction.

Fig 4 12a shows the first iteration of this scene using a weighting function across the $\underline{A}$ matrix consisting of unity weight for regions of the plot which initially fell above a threshold value of -5 dB, and the same threshold value across the rest of the picture Estimates of the locations of the two central emitters are slightly more accurate than before, and the main sidelobe levels are slighlty suppressed However, the unity weighting in the region of the 20 dB signal has not been applied to the correct area because the corresponding lobe of the initial response did not accurately reflect the position of this source Fig 4 12b shows the result of the eighth iteration, based on intermediate results utilising successively lower threshold levels We see that, although the plot perhaps looks 'cleaner' than before because of the suppression of the sidelobes, no higher degree of discrimination has been achieved In fact, the response of the PIM beamformer has apparently been flattened in the region of the two central sources, and a spurious peak has appeared caused by an 'incorrect' thresholding decision at one stage of the iteration Although the thresholding error which occured here was the result of human 'error', it could equally have happened if the thresholding had been carried out on the basis of some other prior assumptions The only iteration guaranteed not to admit such a possibility of error is one supported by exact knowledge of the scene to be reconstructed

Fig 4 13 shows the eighth iteration of the same scene achieved using a weighting function which was derived from the square root of the amplitude of each succeeding result where it fell above the threshold level, and set to the threshold elsewhere This has apparently resulted in an enhancement of discrimination performance of both beamformers, in that both now have formed a clear estimate of the directions of the two central sources and the sidelobes have been suppressed to very low levels However, the improved discrimination in the central region has primarily occurred because of the dip

in the initial PIM estimate had that not been present, the result after any number of iterations would have been only a single central peak. We also note that the peak corresponding roughly to the lower power target, although now clear of the background, has been considerably suppressed itself, and is no more accurately placed than the lobe in the original response Such a clean response has only been obtained because the operator had clear knowledge of the locations of the simulated targets, and could easily have resolved spurious peaks as in Fig 4 12b without such knowledge.

It is clear that such iterative processing is extremely prone to error, and is unable to cope with the location of sources which fall below the sidelobe or interference level present in the initial estimate. MUSIC, on the other hand has dealt with the same data extremely efficiently and accurately, on the basis of the assumption that the targets may be represented as point emitters.

For comparison, Fig 4.14 shows the result of raising the values obtained by the initial PIM estimate to the power of eight This simple operation has clearly had almost as much impact on the result as did the iterations We conclude that iterative processing of the pseudo inverse in this manner does not appear to provide a useful means of improving discrimination of decorrelated point targets

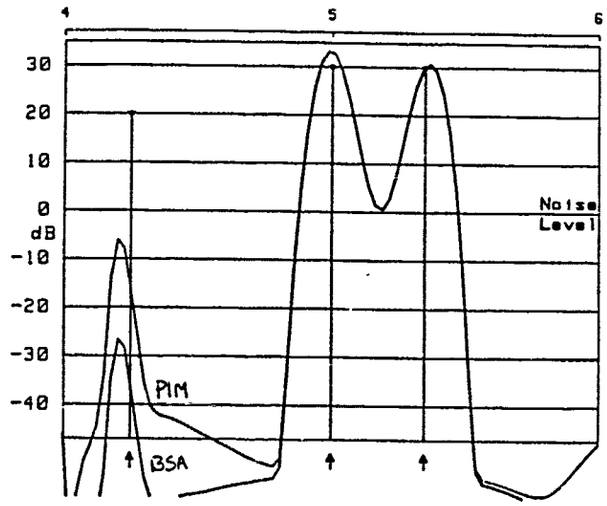

Fig 4.13 The eighth stage of an iteration of Fig 4.11b, in which the weighting function employed applies a weight equal to the square root of the amplitude of the previous iteration when this falls above the threshold level, as described in the text. The result is a smoother, 'cleaner' looking result. The BSA result has apparently improved discrimination, although this is because the original PIM was used as the starting point for the process

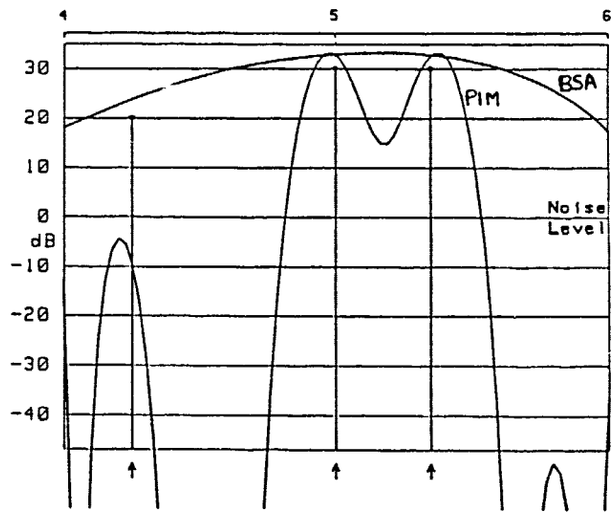

Fig 4.14 This diagram shows the original BSA result of Fig 4.11b along with the corresponding PIM reconstruction which has been plotted raised to the power eight

5. CONCLUSIONS

We have demonstrated the successful application of a number of modern signal processing algorithms acting on short samples of data from simulated far field point-like sources subject to observation noise. These achieve improved signal discrimination and noise reduction over conventional techniques through the application of suitable constraints to the solution. Such constraints are satisfied in a least squares sense, which we have shown may be achieved or understood through the application of singular value decomposition.

It is our belief that algorithms such as MUSIC offer worthwhile gains in both the quantity and quality of information which may be retrieved from radar phased array and time series data, through the clear and controlled introduction of suitable prior knowledge about the nature of the problem to be solved We have demonstrated the algorithms acting on data of a fairly idealistic origin Our signals have originated from point sources, and have been subject to additive Gaussian noise, and we have assumed in the results presented in this report that we have knowledge of any perturbations present in the array manifolds used However, the results obtained strongly suggest that further effort should be directed into investigation of their behaviour when handling real data Phased array radars have many inherent advantages in the form of robustness, beam agility, and so on, but it seems short-sighted not to seriously experiment with improved methods of handling the vast amounts of information which they are obviously capable of providing Although the processing load is heavy, we have demonstrated possible approaches which might possibly result in faster algorithms, and the development of solid state devices is advancing at such a rate that the realisation of on line signal processors performing eigenanalysis is a real possiblity.

We are currently evaluating the performance of these algorithms for a variety of applications such as resolving multipath duplicated images and counting jammers for raid strength analysis. Work is also continuing on possible methods of improving the discrimination of sources through the inclusion of additional prior knowledge It is of vital importance that the performance of such algorithms should be tested on digitised signals from real phased arrays, in order to assess their sensitivity to factors which are not present to date in our simple simulations

Of the algorithms tested here, MUSIC, EMU and KTSC1 appear to offer most promise As we have shown, they appear to be very closely related Although by suitable weighting of the noise subspace eigenvectors, KTSC1 is capable of better noise rejection than MUSIC when dealing with short samples

of data, we have demonstrated with the example of the focal plane data in Section 4 2 7 that, in its current form, it is only suitable for aperture plane or time series data analysis. MUSIC and the eigenvalue weighted version of the algorithm, EMU, appear to be most widely applicable methods, since they do not make use of assumptions about the conformation of the sensor array. True maximum likelihood parameter estimation remains a more distant goal because of the associated weight of processing. However, it seems from recently published results (Böhme (1985), Hudson (1985)) that methods of accelerated approximate maximum likelihood estimation may be possible, and it is certainly the case that such approaches warrant more detailed investigation and characterisation.

REFERENCES

Akaike,H. (1974) - A new look at the statistical model identification
I3E Trans. Autom. Contr. 19,716-723,(1974)

Alsup,J.M. (1984) - High resolution techniques for two-dimensional
estimation of angle of arrival for planar array
I3E ICASSP (1984) pp 39B.4.1-39B.4.4

Applebaum,S.P. (1976) - Adaptive arrays
I3E Trans AP-24(5),585-598,(1976)

Böhme,J.F. (1985) - Source parameter estimation by approximate maximum
likelihood and nonlinear regression
I3E J. Oceanic Engineering OE-10, July 1985

Burg,J.P. (1972) - The relationship between maximum entropy spectra
and maximum likelihood spectra
Geophysics 37(2),375-376,(1972)

Capon,J., Greenfield,R.J., Kolker,R.J. (1967) - Multidimensional
maximum likelihood processing of a large aperture seismic array
Proc I3E 55(2),192-211,(1967)

Clarke,I.J., de Villiers,G.D., Mather,J.L. (1985) - Resolution limits
of a two dimensional antenna array
to be published in Proc SPIE, 564, (1985)

Cox,H. (1973) - Resolving power and sensitivity to mismatch of
optimum array processors
Acoust. Soc. Amer. 54(3),771-785,(1973)

Cullum,J. (1980) - Ill-posed deconvolutions: regularisation and
singular value decompositions
I3E Proc. 19th Decision and Control Conf. pp 29-35,(1980)

Evans,J.E., Johnson,J.R., Sun,D.F. (1982) - Application of advanced signal processing techniques to angle of arrival estimation in ATC navigation and surveillance systems
MIT Technical Report 582, (FAA-RD-82-42), June 1982

Gabriel,W.F. (1980) - Spectral analysis and adaptive array superresolution techniques
Proc IEE 68(6),654-666,(1980)

—— (1982) - Tracking closely spaced multiple sources via spectral estimation techniques
NRL Report 8603, June 23, 1982

—— (1984) - A high resolution target tracking concept using spectral estimation techniques
NRL Report 879, May 31, 1984

Golub,G.H., Reinsch,C. (1970) - Singular value decomposition and least squares solutions
Numer. Math. 14,403-420,(1970)

Hanson,R.J. (1971) - A numerical method for solving Fredholm integral equations of the first kind using singular values
SIAM J. Numer. Anal. 8(3),616-622,(1971)

Hudson,J.E. (1985) - Near-optimum superresolution with a numerically efficient algorithm
Proc I.O.A. 7(4),105-109,(1985)

Jaynes,E.T. (1957) - Information theory and statistical mechanics
Physical Review 106(4),620-630,(1957)

Johnson,D.H., de Graaf,S.R. (1982) - Improving the resolution of bearing in passive sonar arrays by eigenvalue analysis
IEE Trans ASSP-30(4),638-647,(1982)

Kay,S., Demeure,C. (1984) - The high-resolution estimator - a subjective entity
Proc IEE 72(12),1815-1816,(1984)

Kumaresan,R., Tufts,D.W. (1983) - Estimating angles of arrival of multiple plane waves
I3E Trans AES-19(1),134-139,(1983)

Lawson,C.L., Hanson,R.J. (1974) - Solving least squares problems
Englewood Cliffs, N.J. : Prentice Hall

Makhoul,J. (1975) - Linear prediction : a tutorial review
Proc I3E 63(4),561-580,(1975)

Nuttall,A.H. (1976) - Spectral analysis of a univariate process with bad data points, via maximum entropy and linear predictive techniques
NUSC Technical Report 5305, 26 March 1976

Penrose,R. (1955) - A generalised inverse for matrices
Proc. Cambridge Philos. Soc. 51,406-413,(1955)

Pike,E.R., M^CWhirter,J.G., Bertero,M., de Mol.. (1984) - Generalised information theory for inverse problems in signal processing
Proc IEE (F) 131(6),660-667,(1980)

Rissanen,J. (1978) - Modelling by shortest data description
Automatica, 14,465-471,(1978)

Schmidt,R.O. (1979) - Multiple emitter location and signal parameter estimation
RADC Spectrum Estimation Workshop, Oct 3-5, 1979, pp 243-258

—— (1981) - A signal subspace approach to multiple emitter location and spectral estimation
PhD Thesis, Stanford University, November 1981

Shannon,C.E (1948) - A mathematical theory of communication
Bell Syst. Tech. J. 27,379-424,623-657,(1948)

Sibul,L.H. (1984) - Application of singular value decomposition to adaptive beamforming
I3E ICASSP (1984) pp 45.5.1-45.5.4

Toraldo di Francia,G. (1969) - Degrees of freedom of an image
J. Opt. Soc. Am. 59(7),799-804,(1969)

van den Bos,A. (1971) - Alternative interpretation of maximum entropy analysis
Trans IEE IT-17,493-494,(1971)

van Trees,H.L. (1968) - Detection, estimation, and modulation theory, I
Wiley, N.Y., 1968

Varah,J.M. (1973) - On the numerical solution of ill-conditioned linear systems with applications to ill-posed problems
SIAM J. Numer. Anal. 10(2),257-267,(1973)

Wax,M., Kailath,T. (1983) - Determining the number of signals by information theoretic criteria
Proc IEEE ASSP Spectrum Estimation Workshop II,192-196,(1983)

White,W.D. (1979) - Angular spectra in radar applications
IEE Trans AES-15(6),895-899,(1979)

Wilkinson,J.H. (1965) - The algebraic eigenvalue problem
Clarendon Press, Oxford, 1965

Wilkinson,J.H., Reinsch,C. (1971) - Linear algebra handbook for automatic computation, vol 2
Springer Verlag, 1971

Zeegenbein,J. (1979) - Spectral analysis using the Karhunen - Loeve transform
IEEE ICASSP (1979) pp 182-185

ADA 168249

DOCUMENT CONTROL SHEET

Overall security classif cation of sheet .....UNCLASSIFIED..... .....

(As far as possible this sheet should contain only unclassified information. If it is necessary to enter classified information, the box concerned must be marked to indicate the classification eg (R) (C) or (S))

1. DRIC Reference (if known)	2. Originator's Reference Memorandum 3864	3. Agency Reference	4. Report Security Classification U/C
5. Originator's Code (if known)	6. Originator (Corporate Author) Name and Location Royal Signals and Radar Establishment		
5a. Sponsoring Agency's Code (if known)	6a. Sponsoring Agency (Contract Authority) Name and Location		

7. Title
LEAST SQUARES SOLUTIONS IN SIGNAL PROCESSING USING THE SINGULAR VALUE DECOMPOSITION

7a. Title in Foreign Language (in the case of translations)

7b. Presented at (for conference papers) Title, place and date of conference

8. Author 1 Surname, initials	9(a) Author 2	9(b) Authors 3,4...	10. Date	pp. ref.
Mather J L				

11. Contract Number	12. Period	13. Project	14 Other Reference

15. Distribution statement
Unlimited

Descriptors (or keywords)

continue on separate piece of paper

Abstract
We compare a number of advanced antenna array and time series signal processing techniques which either utilise singular value decomposition, or may be profitably interpreted in terms of it. Computer simulations of these techniques acting on spatio-temporal data from linear, conformal circular and focal plane arrays are presented and discussed. Means of introducing prior knowledge are described, and lead to the suggestion of a computationally efficient processing approach. We have made extensive use of a sequential processing strategy for multiple domain data and show by example that this potentially allows both detection and classification of additional emitters.

S80/48