Extension of the Pisarenko method to sparse linear arrays

Jean-Jacques FUCHS, Member IEEE
IRISA/Université de Rennes I
Campus de Beaulieu - 35042 Rennes Cedex - France
phone: +33 299 847 223, Fax: +33 299 847 171, e-mail: fuchs@irisa.fr

Abstract

When applied to array processing, the Pisarenko Harmonic Decomposition (PHD) method is limited to linear equispaced arrays. We present an approach that allows us to extend it to general arrays though for the ease of exposition, we consider only sparse linear arrays.

We exploit the fact that the PHD can be seen as a deconvolution or model-fitting approach that minimizes an ℓ_1 norm and can thus be implemented as a standard linear program. Looking at the PHD from this point of view has two advantages : it allows us to extend its applicability to arbitrary arrays and by diverging slightly from the basic philosophy it allows us to improve its performance which is often quite poor in its original version.

submitted for publication in the IEEE Transactions on Signal Processing
July 1996
revised version March 1997

EDICS SP 1.2.

1 Introduction

The covariance matrix of the outputs of a linear array with N equispaced sensors (denoted ULA for uniform linear array, in the sequel) is hermitian and Toeplitz. It has $2N-1$ real degrees of freedom and one can localize up to $N-1$ uncorrelated narrowband far-field sources in additive white noise. If the geometry of the array is arbitrary, the covariance matrix is hermitian with a constant diagonal, it can have up to $L = N(N-1)+1$ real degrees of freedom, and provided the array manifold is unambiguous, one can locate up to $N(N-1)/2$ uncorrelated sources. We restrict our attention to sparse linear arrays for the ease of exposition, though the localization technique we present applies to more general arrays. By sparse linear array, we mean a linear array for which the sensors are laid out on an underlying regular grid [1].

While for ULAs, there are parametric means (ME methods, PHD, min-norm, root-MUSIC, etc.) to achieve the localization, for non-equispaced arrays no such systematic approach exists and some kind of search and "peak-picking" has to be performed. MUSIC can be applied to such arrays if the number of sources is smaller than the number of sensors, but MUSIC does not take into account the array geometry in the estimation of the signal subspace. The conventional beamformer (CBF) can handle such arrays without limitation on the number of sources but its performance is poor for closely spaced sources.

We propose to apply a model-fitting algorithm to the output of a CBF that drastically improves this performance. In [2] such a model-fitting or deconvolution approach has been applied locally to spatial sectors; here the approach is global and is presented as an extension of the Pisarenko method to non-equispaced arrays. It is based on the same premises as the methods proposed in [3] and belongs to the correlation matching spectral estimation schemes introduced in [4]. Its implementation relies on a standard linear programming algorithm such as the simplex algorithm available in any scientific program library.

2 Modeling and problem statement

We consider a linear array having N omnidirectional sensors in position $[0, x_1, x_2, ..x_{N-1}]$. The sensor outputs are low-passed, sampled and Fourier-transformed, and we will only be concerned with the narrowband signal corresponding to a single frequency bin of this discrete Fourier transform (DFT) of the sensor data. We denote by X_k the complex order-N vector (snapshot) containing the corresponding components over the k–th time interval. An estimate of the covariance matrix $\hat{R}$ of X_k is then given by :

$$\hat{R} = \frac{1}{T} \sum_{k=1}^{T} X_k X_k^* \tag{1}$$

For sufficiently long time intervals, this is also an estimate of the spectral density matrix of the sensor outputs at the considered temporal frequency. Its exact value, denoted R, can be decomposed into the sum of R_s, the contribution of P sources and R_n, the contribution of the noise :

$$\begin{aligned} R &= R_s + R_n \\ R_s &= \sum_{p=1}^{P} a_p d\theta\left(\nu_p\right) d\theta(\nu_p)^* \\ R_n &= vI \end{aligned} \tag{2}$$

where v is the power of the ambient noise assumed spatially white, and $d\theta\left(\nu_p\right)$ is the steering vector associated with the p-th source whose power is denoted a_p. We assume the source signals to be uncorrelated. The signal-to-noise ratio (SNR) of the p-th source is then $\rho_p = a_p/v$. We take as unit of length $\lambda/2$, where λ is the wavelength at the considered temporal frequency. The spatial frequency ν is then related to the bearing ϕ taken with respect to broadside by $\nu = (\sin\phi)/2$. The steering-vector associated with spatial frequency ν admits the following model :

$$d\theta(\nu) = \left[1 \; e^{2i\pi\nu x_1} \; e^{2i\pi\nu x_2} \; \cdots \; e^{2i\pi\nu x_{(N-1)}}\right]^T \tag{3}$$

The exact covariance matrix R is hermitian with constant diagonal and element (k,l) $(k \neq l)$ equal to :

$$r_{k,l} = \sum_{p=1}^{P} a_p e^{2i\pi\nu_p \delta_{k,l}} \tag{4}$$

where $\delta_{k,l} = x_k - x_l$ denotes the inter-sensor distance. The elements of $\{\delta_{k,l}, k > l\}$ define the support of the so-called co-array or discrete co-array function. The co-array assumes values only at the discrete set of locations given by the pairwise differences of the sensor locations [1], [5] and its value at a given point of its domain (support) is an integer equal to the number of pairs k, l that give the same $\delta_{k,l}$. These integers thus vary between 1 and N and at the origin the value of the co-array is always N. For an arbitrary linear array, if all the elements $\delta_{k,l}$ are distinct the cardinality of the support is equal to $N(N-1)/2+1$ and the number of real degrees of freedom, we denote L, of the covariance matrix R is then $N(N-1)+1$. Note that L is always an odd integer. R can equivalently be represented by a complex vector $\underline{r}$ of dimension $(L+1)/2$ containing the ordered, distinct covariances. For ULAs this vector is of dimension N and is also the first column of the covariance matrix :

$$\underline{r} = [\, r_0 \;\; r_1 \;\; r_2 \;\; . \;\; . \;\; r_{N-1}]^T \tag{5}$$

Since, in order to work with real data, we will mainly work with the beamformer outputs, let us introduce their definitions and expressions. The beamformer output at spatial frequency f is :

$$y(f) = \frac{1}{N^2}\, d\theta^*(f) R d\theta(f) \;\; . \tag{6}$$

For R as in (2), it can be decomposed as :

$$y(f) = \sum_{p=1}^{P} a_p F_0(f-\nu_p) + \frac{1}{N}\, v \tag{7}$$

where :

$$F_0(f-\nu) = \frac{1}{N^2}|d\theta(f)^* d\theta(\nu)|^2 = \frac{1}{N^2}|d\theta(0)^* d\theta(f-\nu)|^2 \tag{8}$$

is the output of the beamformer at spatial frequency f for a source with unit power located at spatial frequency ν. This is also known as the array pattern when the sensor weights all equal one. For an ULA, $F_\nu(f)$ is a shifted version of the discrete Fejer kernel, the discrete Fourier transform of the sampled *triangle*. Note that the shift property holds even for non-equispaced linear arrays.

If the exact covariance matrix R has L real degrees of freedom, we propose to take as observations L equispaced samples of the beamformer output :

$$y_k \;\; = \;\; y(f_k) \;\; \text{for} \;\; f_k = \frac{k}{L} \;, \;\; k = 0, \pm 1, \pm 2, \; .. \pm \frac{L-1}{2} \tag{9}$$

In practice, only an estimate $\hat{R}$ of R is available and the corresponding CBF-outputs (6, 9) are denoted $\hat{y}_k$. There might be a slight loss of information when performing this transformation since for some array geometries (e.g. ULA) it is not one to one.

3 The original Pisarenko method

The PHD method relies on a theorem of Caratheodory [6]. Let us recall it for completeness.

Theorem: Let $c_1, c_2, .., c_{N-1}$ be complex numbers (not all zero) with $N > 1$. Then there exists some integer P, $(1 \le P < N)$ and some real numbers a_p and ν_p $(p = 1, 2, ., P)$ such that $a_p > 0$, $\;-1/2 \le \nu_p < 1/2$ and $\nu_p \ne \nu_j$ for $p \ne j$ and such that the following representation of the sequence of numbers c_k is true :

$$c_k = \sum_{p=1}^{P} \; a_p \; exp \; (2i\pi k \nu_p) \qquad\qquad k = 1, .., N-1. \tag{10}$$

The constants a_p and ν_p are determined uniquely. $\diamond$

3.1 The basic formulation

To see the link with the PHD method, one must think of the c_k's as being the elements of a covariance sequence r_k of a scalar stationary complex time-series. Given the partial sequence $\{ r_1, r_2, .., r_{N-1}\}$, one can always adjoin to it a (real) positive variance r_0 such that the N-dimensional hermitian Toeplitz matrix, say R, associated with this sequence is (just) positive semi-definite. In terms of this Toeplitz matrix, the theorem then says that it admits a unique decomposition as the sum of P dyads $d\theta(\nu_p)d\theta(\nu_p)^*$:

$$R = \sum_{p=1}^{P} a_p d\theta(\nu_p) d\theta(\nu_p)^*$$

with $d\theta(\nu)$ as in (3) but for a ULA, i.e., with sensors at positions $x_k = k$ half-wavelengths :

$$d\theta(\nu) = \begin{bmatrix} 1 & e^{i2\pi\nu} & e^{i4\pi\nu} & \cdots & e^{2i\pi(N-1)\nu} \end{bmatrix}^T \tag{11}$$

For a variance r_0 larger than this minimal value, R is strictly positive definite and the above decomposition of R includes a white noise contribution (2) with variance $v = r_0 - \sum_p a_p$ equal to the minimal eigenvalue of R.

Using this interpretation of the Caratheodory's theorem, the algorithm proposed by Pisarenko to decompose a partial covariance sequence is then as follows :

◇ build R, the positive definite hermitian Toeplitz matrix of order N associated with the covariance sequence.

◇ compute v, the smallest eigenvalue, let q denote its multiplicity and form $R - vI$.

◇ if $q = 1$ keep this matrix, otherwise extract a principal minor of order $N + 1 - q$ from $R - vI$.

◇ in both cases the resulting matrix has a single zero eigenvalue, let $U = [\ u_0\ u_1\ ..\ u_P\ ]^T$ be an associated eigenvector, where $P = N - q$.

◇ the P roots of $u_o + u_1 z + ... + u_P z^P = 0$ are then distinct and of modulus equal to one, they can be rewritten as : $exp(2i\pi\nu_p)$ with $-1/2 < \nu_p \leq 1/2$

◇ the associated amplitudes a_p are obtained by solving a set of linear equations.

The basic PHD method has been developed for scalar complex time series and uses an estimated covariance sequence. Since the components of $d\theta(\nu_p)$ represent the covariance sequence associated with a cisoid (a complex sinusoid), the PHD method decomposes a given partial covariance sequence into the contribution of a minimal number of P cisoids and white noise with maximal variance.

In an array processing context, the PHD method only applies to ULAs in the presence of white noise and decorrelated sources. If the covariance matrix is estimated (1) from the snapshots, it is not Toeplitz. One then starts by transforming it into a Toeplitz matrix (for instance by averaging the diagonals, a non-optimal procedure). The resulting matrix is in general positive definite and its smallest eigenvalue is simple. If P, the number of sources, is known one should then extract a principal minor of order $P + 1$ and apply the steps described above. This procedure is clearly suboptimal since it uses only part of the array. If P is unknown, one applies the above described method to the whole matrix and localizes generically $N - 1$ sources. If P denotes the true number of sources, the $N - P$ smallest eigenvalues are almost equal and there are $N - P$ spurious sources whose amplitude is in general low, i.e., of the order of the estimation error standard deviation. One can also check, and somehow establish, that the spatial frequencies of these spurious sources are uniformly distributed and if one of these sources falls close to a true one, it catches part of its power and may thus be taken for a true one. This explains that the performance of the Pisarenko method can be quite poor. The extension we propose will not have these shortcomings. It will allways use all the information available and very few spurious sources will appear since (see section 6.2) we allow for discrepancies between the reconstructed model and the estimated covariances.

3.2 Another formulation

The basic PHD method described above can also be seen as an ingenious manner to solve the following nonlinear and non-trivial optimization problem for which the theorem of Caratheodory establishes the existence and uniqueness of a solution. Given $\underline{r} = [\ r_0, r_1, r_2, .., r_{N-1}\ ]^T$ a N-dimensional complex vector associated with a positive definite covariance matrix, solve :

$$\begin{aligned} &\text{Max } v \\ &\text{subject to} \quad \underline{r} - v\ e_1 = \sum_{p=1}^{P} a_p d\theta(\nu_p) \\ &\text{with} \quad a_p\ >\ 0 \ \text{ and } \ 0 \leq P < N \end{aligned} \tag{12}$$

where $e_1 =\ [\ 1, 0, 0, .., 0\ ]^T$ corresponds to the white noise model and $d\theta\,(\nu_p)$ is as in (11) the source contribution. The $2P + 1$ unknowns (with P itself unknown) are v and $\{a_p, \nu_p\}$ for $p = 1, 2, ..P$.

Let us look at this approach from a geometrical point of view [4] that will allow us to extend it to more general arrays. We consider $\underline{r}$ as a N-dimensional complex vector in $\mathbb{C}^N$. The set of all such vectors associated with positive Toeplitz matrices is then a closed positive cone with vertex at the origin. The generating elements of this cone are the steering-vectors (11) to which one associates the dyads $d\theta(\nu_p)\ d\theta(\nu_p)^*$, the rank-one Toeplitz matrices. A generic vector $\underline{r}$ associated with a positive definite hermitian Toeplitz matrix belongs to the interior of this cone while a vector $\underline{r}$ associated with a positive semidefinite matrix lies on the boundary of the cone. The Pisarenko method applied to an interior vector projects it onto the boundary of the cone where it admits a unique decomposition into generating elements (the steering vectors). The projection is performed by subtracting the white-noise contribution- a specific vector e_1 in the *middle* of the cone- with maximal amplitude. This is precisely what the nonlinear optimization problem (12) is about.

In order to solve this nonlinear problem (12) in a manner different from Pisarenko, let us discretize the spatial frequencies (the sources bearings) and consider that these frequencies can only take M preassigned values ν_m, equispaced in the interval $]-1/2\ ,1/2]$. This amounts to replacing the above-mentioned cone by a polyhedral cone that is the convex conical hull of the M equispaced generating elements. By maximizing then the noise variance, one obtains a vector $\underline{r}-v\ e_1$ that lies on the boundary of the polyhedral cone. A sufficiently small discretization step (large M) allows us to approximate the original cone with an arbitrary small error and thus to reduce the bias introduced by this approximation at will. The optimization problem that approximates (12) then becomes :

$$\begin{aligned} &\text{Max } v \\ &\text{subject to} \quad \sum_{m=1}^{M} a_m d\theta(\nu_m)\ +\ v\ e_1\ =\ \underline{r} \\ &\text{with} \quad a_m\ >\ 0. \end{aligned} \tag{13}$$

This is a linear program [7] in the $M + 1$ unknowns $(v, \{a_m, m = 1, .., M\})$ and its properties are well established. Note that in this formulation the variables ν_p disappear as well as P. The localization of the sources is deduced from the indices of the amplitudes a_m that are greater than zero. There are N linear complex equality constraints that are easily transformed into $L = 2N - 1$ real constraints. The standard simplex algorithm available in any scientific program library will yield a global solution of this convex optimization problem. Such a solution, will generically have L non-zero unknowns and thus, since v will be positive, $2(N-1)$ positive weights a_m's. This number is to be compared to P in (12) which is generically equal to $N-1$. The ratio two between the two numbers corresponds to the fact that the spatial frequency of a true source will generically fall in between two discretization points and both will be needed in the reconstruction of its contribution. This means that among the M potential spatial frequencies, in general, at most $N-1$ pairs of them will emerge from this procedure. One must keep in mind that we are approximating the standard PHD method which indeed seeks to reconstruct the observations as a sum of rays (cisoids). Among these $N-1$ locations only those having $significant$ energy will be kept.

It is this second manner (13) to approximatively solve (12) that lends itself to potential extensions. Note however that even in the case of a ULA, solving (13) may be computationnally cheaper than using the standard PHD method which requires the calculation of the minimal eigenvector, the rooting of the associated polynomial and the estimation of the powers. Moreover, as we will see later, it is this second approach that allows us to take into account the statistics of the estimation errors and thus to improve drastically the performance.

4 The extension to arbitrary linear arrays

When one considers arbitrary linear arrays, a number of problems arise that are either absent or trivial for ULA's. Ambiguity or uniqueness of the decomposition is one of them. Invertibility or identifiability is another one. We will define these issues and comment on them below, once we have presented the extension of the PHD method we propose.

4.1 The extension

The extension of the PHD method to arbitrary linear arrays is now straightforward: it consists simply in adapting the linear program (13) to these arrays. For an arbitrary linear array with N sensors and L real degrees of freedom, the vector $\underline{r}$ is now of dimension $(L+1)/2$ (see section 2) and contains all the distinct covariances present in R. The steering vectors present in (13) are similarly replaced by corresponding $(L+1)/2$-dimensional complex vectors that are extended versions of (3). Of course, as in (13), to solve the linear program all these vectors have to be transformed into L-dimensional real vectors, but we do not further detail this formulation since we will later (see section 5) apply a linear one-to-one transformation to it, to obtain a real formulation that is easier to handle.

For the moment, let us consider the structure of the covariance matrices associated with these arrays. One can think of them as matrices or as L-dimensional vectors in $\mathbb{R}^L$.

The covariance matrix R of the output of an arbitrary linear arrays with N sensors and L real degrees of freedom can be decomposed (2) into the contribution R_s of the sources and the contribution vI of the white noise. Since R_s is a positive linear combination of dyads built upon the steering vector (3) associated with the array, R_s belongs to the closed convex conical hull of the set of the dyads. This is again a closed convex cone with vertex at the origin whose generating elements are the dyads. Since the contribution of the additive white noise can be constructed as a convex combination of L equispaced dyads, it clearly belongs to the interior of the same cone. The global matrix R thus always lies in the interior of the cone.

A major difference with the cone associated with linear and equispaced arrays is that, except for very specific sparse linear arrays [5], the interior and the boundary of this cone are difficult to characterize in terms of the covariance matrix. The elements on the boundary of the cone include, but are not limited to, the rank deficient matrices. Similarly it is difficult to characterize the set of matrices that are potential covariance matrices of a given array, unlike in the ULA case where these are simply the set of positive hermitian Toeplitz matrices. However, since we are only concerned by the decomposition of covariance matrices associated with true arrays, it is possible to be more specific. We will present in section 4.2 and 4.3 below, the conditions under which the proposed extension of the PHD method will work. For the moment, let us say that for P small enough, one is in the degenerate situation where the global source-contribution lies on the boundary of the cone. To identify it, one again maximizes the contribution of the noise and this forces the remainder to lie on the boundary.

4.2 The ambiguity issue

For arbitrary linear arrays, nothing guarantees that an element on the boundary, which represents the global contribution of the P sources, has a unique decomposition into P dyads. While in the equispaced sensor case the uniqueness has been proven by Caratheodory, for arbitrary linear arrays no general result exists and the uniqueness of the decomposition is a difficult question that concerns the ambiguity of the array manifold. Indeed, even for linear arrays, the study of ambiguity is a difficult task and only a few sufficient conditions for non-ambiguity or necessary conditions for ambiguity are available. Ambiguity appears trivially for ULAs with inter-sensor distance greater than $\lambda/2$. This corresponds to the aliasing phenomenon induced by undersampling. Some specific unambiguous sparse linear arrays are known as for instance the minimum redundancy linear arrays [5]. A non-trivially ambiguous non-equispaced linear arrays is analysed in [8]. In what follows we will mainly consider un-ambiguous arrays as for instance minimum redundancy linear arrays (MRLAs). It is worthwhile to note that some care has to be taken when applying MUSIC-like algorithm to such *fully augmentable* arrays since these algorithms have to be implemented on the augmented array in order to be indeed un-ambiguous [9].

4.3 The invertibility issue

The invertibility issue concerns, assuming the array is unambiguous, the conditions under which the knowledge of R allows us to reconstruct the underlying scenario. The investigation of (2,4) leads to the following conclusions :

• If $P < N$, the covariance matrix R (2) is rank deficient in the absence of noise and one can recover the scenario in a number of different ways, MUSIC for instance. It is easy to realize that the noise power is simply the value of the possibly multiple, minimal eigenvalue of R.

• If $N < P \leq (L-1)/2$, the covariance matrix (2) is full rank even in the absence of noise and the task is less trivial. However since the number $2P+1$ of unknowns $(\{a_p, \nu_p\}_{p=1,..,P}, v)$ is smaller than the number L of real degrees of freedom, it is possible to identify the scenario. In the absence of noise, one has to find the solution, assumed to be unique, of a system of L nonlinear equations in $2P+1$ unknowns. In the presence of noise one has precisely to solve the non-linear optimization problem (12) we propose to replace by (13).

• If $P > (L-1)/2$, the task is impossible. The covariance matrix lies in the interior of the cone even in the absence of additive noise. The unique decomposition into the sum of an element on the boundary and the white noise element still exists but this decomposition has no relation with the true scenario.

The maximal number of identifiable (uncorrelated) sources is thus $(L-1)/2$.

5 A model-fitting approach

Before describing a further extension intended to improve the statistical performance of the method, we propose to make some modifications in the implementation of (13). A possible interpretation of the linear constraints in (12,13) is that one seeks a reconstruction of the observed covariance sequence $\underline{r}$ as a positive combination of a (small) number of elementary covariance sequences, each associated with a source, and a white noise covariance sequence with maximal variance. The PHD method is thus also a model-fitting approach using as a criterion the maximization of the noise variance. But looking at the first linear constraint in (13) which reads :

$$\sum_m a_m + v = r_0, \qquad\qquad a_m \geq 0$$

one notes that maximizing v is indeed equivalent to minimizing $\sum_m a_m$, the sum of the sources powers. One can thus replace the criterion in (13) and rewrite the program as :

$$\begin{aligned} &\text{Min} \ \sum_{m=1}^{M} \ a_m \\ &\text{subject to} \ \ \sum_{m=1}^{M} \ a_m d\theta(\nu_m) \ + \ v \ e_1 \ = \ \underline{r} \\ &\text{with} \ \ a_m \ \geq \ 0. \end{aligned} \tag{14}$$

This is an interesting re-interpretation of the PHD method as a model-fitting approach that minimizes the sum of the source powers. Since these a_m's are greater than or equal to zero, Min $\sum a_m$ is further equivalent to minimizing their ℓ_1-norm.

In order to work with real quantities having a physical meaning, rather than considering the real and imaginary part of $\underline{r}$, we propose to work with the outputs of the CBF (6); these are real and positive quantities with an obvious physical meaning. The two approaches are equivalent, i.e., will give the same solution, provided the two quantities are in a one-to-one relation. This is the case if the number of evaluated beams is equal to the number L of real degrees of freedom of the covariance matrix.

From a technical point of view, this amounts to applying a linear transformation to all the (column vectors that appear in the) constraints of the linear program (14). When applied to e_1, the model of the white noise contribution, it yields a constant vector the components of which all equal to $1/N$; we denote it $\underline{w}$. The number of unknowns, their meanings and the location of the optimum are not affected by such a transformation. Note that for ULAs or minimum redundancy linear arrays (MRLAs), the array pattern when the sensor weights all equal one is the DFT of the discrete co-array function. For arbitrary linear arrays the support of the co-array is irregular and this is no longer true.

Instead of fitting a model to the covariance matrix (or sequence) we now fit a model to the outputs of the CBF. Indeed the description given in section 3 and 4 about the structure of the set of the covariance matrices can be transposed on the CBF-side : all potential CBF-outputs belong to the positive cone whose generating functions are $F_0(f-\nu)$ (8). And any CBF-output that corresponds to an invertible scenario can be represented as the sum of P weighted and shifted generating functions and the contribution of the noise (7).

In this new setting and to meet the standard notation of LP's, we rewrite (14) as :

$$\begin{aligned} &\text{Min } \|X\|_1 \\ &\text{subject to : } \quad A\,X + v\ \underline{w} = \underline{b} \\ &X \geq 0, \quad v \geq 0 \end{aligned} \tag{15}$$

where$\|X\|_1$ denotes the ℓ_1-norm of X, the M-dimensional vector of the unknown source powers (a_m) and $\underline{w}$ is the model of the noise contribution, $\underline{b}$ is the L-dimensional vector containing the outputs of the CBF (9), A is an (L, M) matrix $(M \gg L)$ the columns of which contain the shifted and sampled generating functions (8). An inequality $X \geq 0$ means that all the components of X have to be greater than or equal to zero.

Using the discrete version of Parseval's theorem, one can show that for ALEs and MRLAs, all the columns of A have the same ℓ_1 (and ℓ_2) norm. Since moreover all the components of A, $\underline{w}$, $\underline{b}$ are positive, this can be seen as meaning that no spatial frequency should be privileged.

This is our first complete formulation as an implementable linear program (LP) of the extended PHD method. We will further improve upon this method in the next section by introducing the statistical properties of $\underline{b}$. Some remarks concerning properties and implementation issues of this LP are in order [7].

• If there are feasible points, the optimization problem has a solution, say (X^*, v^*) at a vertex of the domain. This means that there is a solution with at most L strictly positive components. Since v^* is always positive, there are at most $L-1$ components of the M-dimensional (weighting) vector X^* that are positive.

• Due to the discretization of the spatial frequencies (bearings) with a step-size $1/M$, a single source will generically be reconstructed by the two neighbouring nodes. $L-1$ positive components in X^* thus allow us to localize at most $(L-1)/2$ sources. This is also the maximal number of admissible sources for invertibility reasons (see section 4.)

• The spatial frequency to be attributed to a source associated with two such positive X^*-components is obtained by linear interpolation of the "indices" of the components. The amplitude is obtained by adding the values of the components.

• Strictly speaking, the two neighbouring elements never allow us to exactly reconstruct a missing generating function, other elements (distributed over the domain) contribute also. The amplitude of these further contributing elements can be monitored by choosing the discretization step $1/M$. This step should also be chosen in accordance with the accuracy one expects from the array.

• The solution (X^*, v^*) given by a simplex algorithm may have less than L positive components. Such a solution is termed degenerate and standard software can handle this case.

• In practice only an estimate $\hat{\underline{b}}$ of $\underline{b}$ is available. The maximal number $(L-1)/2$ of sources will then almost always be needed to reconstruct $\hat{\underline{b}}$. The difficulty is then to decide whether a source with low amplitude is indeed a source or is induced by the estimation error and/or the discretization step. Deciding how many sources there are, or which sources are true ones, is a detection problem. An *a priori* evaluation of the detection threshold of the array should thus be performed so that one can make a decision by looking at the amplitudes. This is a standard difficulty for localization algorithms if the number of sources is not assumed to be known.

6 Further extensions

In all the previous formulations (12, 14, 15) as well as in the basic PHD method, the model-fitting is performed exactly, i.e., without any discrepancy between the observations and the reconstructed model. There is no real justification or need for that since the observations, the CBF outputs $\hat{\underline{b}}$ in (15), are corrupted by estimation errors :

$$\Delta \underline{b} = \hat{\underline{b}} - \underline{b}$$

due to the fact that only T snapshots are available. There is indeed no reason to ask for a perfect fit under such circumstances and the maximum likelihood technique, for instance, is precisely one way to specify how the model should be made to fit the observations. In these parametric approaches the structure of the model (the number of sources) has to be known as well as a very precise initial point. We do not make this assumption here and handle simultaneously the detection and estimation problem.

6.1 Statistical properties of the beams

We now present the statistical properties of the CBF-output and more precisely those of $\hat{\underline{b}}$. They will allow us to modify the optimization problem. Under the standard model for the snapshots presented in section 2, it is generally assumed that $\hat{R}$ in (1) is such that $T\hat{R}$ is a sample of a complex Wishart distribution with T degrees of freedom and parameter matrix R. Defining then $\Delta R = \hat{R} - R$, $\Delta r_{ij} = \hat{r}_{ij} - r_{ij}$ and r_{ij} the (i,j) – element of R one has [10] :

$$E\left(\Delta r_{ij} \Delta r_{mn}\right) = \frac{1}{T}\, r_{mj} r_{in}. \tag{16}$$

In a more compact form, this relation implies :

$$E\left(tr\left(\Delta R A\right) tr\left(\Delta R B\right)\right) = \frac{1}{T} tr\left(RARB\right) \tag{17}$$

with A and B arbitrary matrices and $tr(A)$ the trace of the square matrix A. Rewriting the CBF output $\hat{y}_k$ (6), (9) as :

$$\begin{aligned} \hat{y}_k &= \frac{1}{N^2}\, d\theta^*(f_k)\hat{R} d\theta(f_k) \\ &= \frac{1}{N}\, tr\left(\hat{R}\frac{d\theta(f_k)d\theta^*(f_k)}{N}\right) \\ &= \frac{1}{N}\, tr\ (\hat{R}P_k) \end{aligned}$$

with P_k a projection matrix and denoting $\Delta y_k = \hat{y}_{f_k} - y_{f_k}$, relation (17) yields :

$$\begin{aligned} E\left(\Delta y_k \Delta y_\ell\right) &= \frac{1}{TN^2} tr\left(RP_kRP_\ell\right) \\ &= \frac{1}{TN^4} |d\theta_k^* R d\theta_\ell|^2 \end{aligned} \tag{18}$$

and for $k = \ell$, one obtains :

$$E(\Delta y_k^2) = \frac{1}{T} y_k^2,$$

a standard result [11] in power spectrum estimation.

In what follows, we shall assume that the value of T, the number of snapshots, is large enough for the outputs $\hat{y}_k$ of the beams to be considered as real Gaussian random variables with mean y_k and covariance given by (18), though this only holds asymptotically in T. These relations then fully define the statistical properties of $\hat{\underline{b}}$. Let us denote Σ the covariance matrix of $\hat{\underline{b}}$ whose elements are defined in (18). A consistent estimate $\hat{\Sigma}$ of Σ is obtained by replacing R by $\hat{R}$ in the evaluation of the elements in (18).

6.2 Improving the performance

We now propose to use the estimate $\hat{\Sigma}$ to whiten the observation vector $\hat{\underline{b}}$. We define $\hat{\hat{\underline{b}}} = \hat{\Sigma}^{-1/2}\hat{\underline{b}}$ and apply the same transformation to the other elements of the LP (15), i.e., we similarly premultiply A and $\underline{w}$ by $\hat{\Sigma}^{-1/2}$ to obtain $\tilde{A}$ and $\tilde{\underline{w}}$. The components of the second member $\hat{\hat{\underline{b}}}$ of the linear constraints are now decorrelated and have unit variance.

This linear transformation does not change the location of the optimum if no further modification is introduced. It makes, in some sense, all the constraints equally valid or important. In order to introduce some flexibility in the equality constraints we propose to replace them by constraints having a lower and upper bound. The idea behind this modification is the following. Since the observations in $\hat{\hat{\underline{b}}}$ are corrupted by estimation errors there is no reason to ask the reconstructed model to fit exactly these observations; it is more adequate to ask the reconstructed model to pass within -say- δ standard deviations of the observations once they have been whitened. This amounts to put a *tube* around the observations and to ask the reconstructed model to stay within it. It is an easy way to remain within the linear programming framework while taking into account the estimation errors. The LP (15) thus becomes :

$$\begin{aligned} &\text{Min } \|X\|_1 \\ &\text{subject to : } \ \hat{\hat{\underline{b}}} - \delta I \leq \tilde{A}\, X + v\ \tilde{\underline{w}} \leq \hat{\hat{\underline{b}}} + \delta I \\ &X \geq 0, \quad v \geq 0 \end{aligned} \tag{19}$$

where I denotes an L-dimensional column-vector of ones and all the inequalities between vectors have to be considered component-by-component.

The parameter δ, the radius of the tube, is obviously an important one. Let us indicate how to choose its value. In order to be able to ignore for the moment the errors due to the discretization of the spatial frequencies, assume that the true scenario generating $\underline{b}$ can exactly be described by our model, i.e, there exists X^* such that the reconstructed second member denoted $\underline{b}(X^*)$ equals $\underline{b}$. We then want to choose δ such that, despite the additive estimation errors, this exact representation belongs to the solution set of (19). Taking into account the whitening step, we have to choose δ such that, with probability close to one :

$$|\ \tilde{\underline{b}}(X^*) - \hat{\hat{\underline{b}}}\ | \leq \delta.$$

Since $E(\hat{\hat{\underline{b}}}) = \tilde{\underline{b}}(X^*))$, the components of the difference $\tilde{\underline{b}}(X^*) - \hat{\hat{\underline{b}}}$ are, asymptotically in T, zero-mean independent normalized Gaussian random variables. We thus have to take δ so that the probability that the supremum of these L random variables be larger than δ in absolute value is close to zero. It is easy to show that δ has then to be of the order of $\sqrt{Log(L)}$, i.e., the maximum of L independent standard Gaussian variable *increases* as $\sqrt{Log(L)}$.

Incidentally, since the optimum, say $\hat{X}$, of (19) is a feasible point it satisfies :

$$\mid \tilde{\underline{b}}(\hat{X}) - \hat{\tilde{\underline{b}}} \mid \leq \delta$$

and combining the two last inequalities one has :

$$\mid \tilde{\underline{b}}(X^*) - \tilde{\underline{b}}(\hat{X}) \mid \leq 2\delta$$

which implies, taking into account the fact that Σ is of order T^{-1}, that

$$\mid \underline{b}(X^*) - \underline{b}(\hat{X}) \mid \leq 2\delta/\sqrt{T}.$$

The variance of $\underline{b}(\hat{X})$ and thus of $\hat{X}$ is of order δ^2/T. This is how the usual asymptotic order of convergence, with respect to the number T of snapshots, enters in this scheme.

This parameter δ also plays an important role in the sparsity issue. Obviously, if it is taken much too large, $X = 0$ is the optimum (no positive weight), while if it is taken too small (or zero) the solutions will be non-degenerate, i.e., will have L positive weights (which is the situation for the PHD method). In general the true solution lies in between these two extremities. This indicates that the value of δ influences the number of sources (positive weights) in the solution. The simulations show that this is indeed the case. Fortunately for values around $\sqrt{Log(L)}$ the sensitivity is reasonably low and we do not need to adapt this parameter to the noise realisations. We further comment on these issues in the simulations.

7 Simulation results

We apply the approach described above to simulated data. Thus only an estimate $\hat{R}$ (1) of R, obtained from T observations (the snapshots), is available.

The implementation of our method is extremely easy and straightforward. The simplex algorithm from any scientific program library can be used. There are only two parameters that have to be tuned. One of them is M, the number of columns in A, that fixes the discretization step of the spatial frequencies. The results are quite insensitive to its value provided it is taken large enough. If it is too small, the reconstruction of an intermediate column by its two neighbours will be poor and other columns will be required creating spurious sources. We take M as a multiple of L the numbers of rows in A, the number of real degrees of freedom in R. Below , we take $M = 10L$. It is important to link the discretization step to the expected accuracy of the array which is somehow proportional to its size and this is an easy way to achieve it. The second parameter is δ the radius of the tube; we commented about its value at the end of section 6.2. and below we take it equal to .8 $\sqrt{Log(L)}$.

Let us also mention again that our method often produces spurious sources i.e. more sources are detected than are actually simulated. Their amplitudes are in general quite small and both their number and amplitudes decrease as M and/or δ increases. Remember that for ULAs, the PHD method - which corresponds to our method with M infinite and $\delta = 0$ - always has a maximum number $N - P - 1$ of spurious peaks. The results we present are obtained by keeping among the detected sources, the P sources with highest amplitude.

7.1 First example

We consider a minimum redundancy linear array with 4 sensors [5] [3]. The sensors are at positions $\{0, 1, 4, 6\}$ (half wavelengths). The exact covariance matrix R has $L = 13$ real degrees of freedom. An estimate $\hat{R}$ is obtained from $T = 100$ snapshots. It is used to evaluate the CBF at 13 equispaced bearings (9), the components of $\underline{\hat{b}}$. M is taken equal to $10L$, and the matrix A is thus of dimension (13,130). The discretization step in spatial frequency is thus $\frac{1}{M} \simeq 7.10^{-3}$.

The scenario [12] consists of 2 sources that have a common SNR of 0 dB with respect to the white noise $a_1 = a_2 = 1$ and $v = 1$ in (2). We simulated 3 different situations : -in the first case the 2 sources are at spatial frequencies 0 and .07 (0 and 8.0 degrees), -in the second case the spatial frequencies are 0 and .06 (0 and 6.9 degrees) and in the third case the spatial frequencies are 0 and .05 (0 and 5.7 degrees). The means and standard deviations of the estimated spatial frequencies, obtained from 400 independent trials, are given in **Table 1**. In row 5 and 6 of this table, the LP (19) with $\delta = 0$ is used (this is equivalent to using (15)), and in row 7 and 8 δ is taken equal to $\delta = .8\ \sqrt{Log(L)} = 1.3$.

Only the very last decimal is affected when the radius δ varies between 1 and 2. This parameter is thus not really sensitive for these scenarios. Note however that this radius has a direct influence on the estimates of the amplitudes of the sources which are systematically underestimated and whose bias increases with δ. In the last set of simulations, the spatial frequency separation was too small for the PHD method -(15) or (19) with $\delta = 0$- to work correctly, for some realizations the location of one of the two strongest sources had little to do with the true location and the global result is not given, since it had no meaning.

source number	1	2	1	2	1	2	δ
source power in dB	0	0	0	0	0	0	
spatial frequency ν	0	.07	0	.06	0	.05	
CR stdt dev. on ν	.0066	.0066	.0080	.0080	.0103	.103	
estimated spatial freq.	-.0043	.0727	-.0051	.0673	-	-	0
estimated stdt dev.	.0176	.0169	.0227	.0237	-	-	
estimated spatial freq.	.0008	.0710	.0003	.0622	.0002	.0552	1.3
estimated stdt dev.	.0087	.0087	.0099	.0103	.0114	.0125	

Table 1: Results over 400 trials for 3 different 2 two-sources scenario. The array has 4 sensors at position (0, 1, 4, 6). Rows 5 and 6 present the results for the basic Pisarenko extension. Rows 7 and 8 present the results for the improved extension.

It is also interesting to mention that, as opposed to the maximum likelihood method, the initialization has no influence on the results since the problem is convex and the optimum is unique.

For the array simulated in this example, the covariance matrix R contains indeed the covariances $\{r_k, k = 0, 1, .., 6\}$. This means, that using the components of the order 4 matrix $\hat{R}$, one can build an order 7 Toeplitz covariance matrix [13] [14] and apply standard localization algorithms. This is of course not a necessity for our approach (see the second example below) and has been chosen to ease comparisons and to be sure that the array presents no ambiguity. The reader can check with Table 1 that other algorithms have quite poorer performance. Even for the easiest case (the two first columns of Table 1), a MUSIC-type algorithm sometimes fails to solve the two sources. This is not in contradiction with the asymptotic optimality of these approaches but simply means that in the pre-asymptotic domain [14] MUSIC does not have the best possible performance.

7.2 Second example

We take an example of a sparse linear non redundant array [1] for which there is a gap in the co-array. The array has 5 sensors at positions $\{0, 1, 4, 9, 11\}$ (half wavelengths). The exact covariance matrix R has $L = 21$ of real degrees of freedom. All the covariances going from r_1 to r_{11} are present once in R except for r_6 which is missing. We simulate 5 uncorrelated sources in white noise at spatial frequencies $0, .08, .16, .24$ and $.32$ all with SNR equal to $0dB$. The number of snapshots is equal to $T = 100$. We present the results obtained from (19). in **Table 2** for 400 independent trials. They are reasonably close to the Cramer-Rao bound which is also given in the table.

source number	1	2	3	4	5
source power in dB	0	0	0	0	0
spatial frequency ν	0	.08	.16	.24	.32
est. spatial freq. ν	-.0104	.0890	.1603	.2341	.3189
est. stdt dev. on ν	.0045	.0064	.0055	.0050	.0037
CR stdt dev. on ν	.0023	.0028	.0030	.0028	.0023

Table 2: Results over 400 trials for a five sensors-five sources scenario. The array has 5 sensors at positions (0, 1, 4, 9, 11). It is nonredundant with a gap in the coarray.

Note that since the number of sources is equal to the number of sensors the matrix is full rank, even in the absence of noise, so that an augmentation technique [13] [14] has to be used. But here since r_6 is missing, the extended array can only be of limited aperture and the difficulty remains the same. Since the sources are too close to be separated by the CBF, there does not seem to be an existing method, other than the ML approach, able to handle this case. Our approach allows us to separate the sources without any specific tuning of the two parameters M and δ. The same parameters as before are used : $M = 10L$ and $\delta = .8\ \sqrt{Log(L)})$.

8 Conclusions

We have presented an extension of the Pisarenko method that allows us to handle sparse linear arrays. Its computational cost is in general quite reasonable. The number of operations required to perform one iteration of a simplex algorithm in order to solve a linear program with L constraints and M unknowns is proportional to the square of $\min(L, M)$ and the number of iterations is, in general, considered to be linear [16] in $\min(L,M)$. The computational load is thus mainly cubic in $\min(L,M)$ and thus quite similar to these of high resolution methods. Our approach has however several advantages over these methods. There is no need to know *a priori* the number of sources. Though some sort of detection scheme has to be added to our procedure to detect and discard weak spurious peaks, knowing the number of sources is actually of no help in our approach. For arbitrary arrays and in difficult situations (such as closely spaced sources, see section 7.1.) it has in general better performance than other methods. In even more difficult situations (closely spaced sources whose number is equal to the number of sensors, see section 7.2.) it continues to work with reasonable performance while no other standard high resolution method actually applies. The maximum likelihood method can, of course, handle such situations. But even if one assumes the number of sources to be known, the difficulty lies in the necessary initialization procedure. Without an excellent initial point, the maximum likelihood method does not converge to the true global maximum making it useless.

Some simulation results have been presented to allow us to apprehend the performance of the method. The simulations have been performed on linear arrays but more general arrays can be handled as well. When the number of sources equals or exceeds the number of sensors, it does not seem that competing methods actually exist, except the conventional beamformer whose performance in resolution is poor. The statistical performance is quite good, and our method actually outperforms the basic Pisarenko method when it is applicable. Theoretical analysis of the performance of our method is under investigation.

Several modifications of the approach or extensions to other problems can be developed. As far as the localization problem considered here is concerned, one can generalize the approach to handle additive noise with arbitrary unknown spatial spectral density [15].

References

[1] D.H. Johnson and D.E. Dudgeon. *Array Signal Processing: Concept and Techniques.* Prentice Hall, 1993.

[2] J.J. Fuchs and H. Chuberre. "A deconvolution approach to source localization". *IEEE Trans. on Sig. Proc.*, 42(6):1462–1470, June 1994.

[3] D.R. Fuhrmann and F.C. Robey. Structured covariance estimation via maximal representation of convex sets. In *Proc. ICASSP*, pages 3249–3252, 1991.

[4] J.H. McClellan and S.W. Lang. Spectral estimation for sensor arrays. *IEEE-T-ASSP*, 31:349–358, April 1983.

[5] A.T. Moffet. Minimum redundancy linear arrays. *IEEE Trans. on Antenna and Propagation*, 16:172–175, March 1968.

[6] U. Grenander and G. Szego. *Toeplitz Forms and Their Applications.* University of California Press, 1958.

[7] D.G. Luenberger. *Linear and Non-Linear Programming.* Addison-Wesley, 1973.

[8] A. Manikas and C. Proukakis. Study of ambiguities of linear arrays. In *Proc. ICASSP*, volume IV, pages 549–55, 1994.

[9] Y.I. Abramovich et al. Ambiguities in DOA estimation for nonuniform linear antenna arrays. In *Proc. ISSPA*, pages 631–634, 1996.

[10] K.S. Miller. *Complex Stochastic Processes.* Addison-Wesley, 1974.

[11] R.W. Schafer and A.V. Oppenheim. *Digital Signal Processing.* Prentice-Hall, 1975.

[12] J.J. Fuchs. Extension of the Pisarenko method to sparse linear arrays. In *Proc. ICASSP*, volume III, pages 2100–2103, May 1995.

[13] F. Haber, S. Pillai and Y. Bar-Ness. A new approach to array geometry for improved spatial spectrum estimation. *Proceedings of the IEEE*, 73(10):1522–1524, 1985.

[14] Y.I. Abramovich et al. Positive definite Toeplitz completion for fully augmentable nonuniform linear arrays. In *Proc. ICASSP*, vol V, pages 2551–2554, May 1996.

[15] J.J. Fuchs. Linear programming in spectral estimation: application to array processing. In *Proc. ICASSP*, vol VI, pages 3161–3164, May 1996.

[16] A. Schrijver. *Theory of linear and integer programming.* John Wiley, 1986.