FAST ADAPTIVE BAYESIAN BEAMFORMING USING THE FFT

Chun-wei J. Lam, Andrew C. Singer

University of Illinois at Urbana Champaign
1308 West Main St., Urbana, IL 61801
{clam, acsinger}@uiuc.edu

ABSTRACT

A fast algorithm is developed to implement a Bayesian beamformer that can estimate signals of unknown direction of arrival (DOA). In the Bayesian approach, the underlying DOA is assumed random and its a posteriori probability density function (PDF) is approximated by a discrete probability mass function. A Bayesian beamformer then balances a set of beamformers according to the associated weights. To obtain a close approximation of the a posteriori PDF, the number of samples must be sufficiently large, incurring a significant computational burden. In this paper, we exploit the structure of a uniform linear array (ULA) to show that samples of the a posteriori PDF can be computed efficiently using the fast Fourier transform (FFT). This leads to a fast algorithm for the Bayesian beamformer, which operates in $O(MlogM + N^2)$ operations where M is the number of samples and N is the number of sensors.

1. INTRODUCTION

The problem of estimating a signal that arrives from an uncertain spatial direction has been investigated in a wide variety of areas such as radar, sonar, and wireless communications. In [1], a robust adaptive beamformer for this problem is introduced using a Bayesian approach. The algorithm performs conventional MVDR beamforming on a number of hypothesized, discrete values of DOA candidates, and then balances multiple outputs according to the discretized samples of the a posteriori PDF that describes the likelihood of the true DOA. When the number of observed data samples increases, the likelihood function becomes more acute and eventually converges to an impulse on the DOA that has the largest likelihood. The resulting beamformer is able to adapt to the underlying true DOA, given that the true DOA lies on one of these hypothesized candidate points. The authors also show that the a posteriori PDF can be approximated from a spatial power spectrum estimate.

The number of discrete samples along the a posteriori PDF plays an important role in the accuracy of the estimation process. When the number of samples is too small, the prior DOA set is unlikely to contain the true underlying DOA, and this results in performance degradation with respect to the true DOA. When a large number of samples are chosen, the computational cost becomes high because each sample point induces the cost of a single separate MVDR algorithm.

In this paper, we consider the Bayesian approach of [1] applied to a uniform linear array and develop a fast algorithm to compute discrete samples along the a posteriori PDF. We show that the a posteriori PDF can be completely represented by only a fixed number of critical samples, and that extra samples can be computed efficiently by interpolation. The new algorithm allows us to consider both a large number of sensors and a large number of PDF samples without inducing significant computational cost.

This work was supported in part by the Department of the Navy, Office of Naval Research, under grant number N00014-01-1-0117.

2. BACKGROUND

We consider a number of narrowband signals arriving at a uniform linear array of N sensors in the presence of additive noise. The array outputs are sampled at times $t = 0, T, 2T, \dots$, such that the received data vector $\mathbf{x}_k = \mathbf{x}(kT)$ has the form

$$\mathbf{x}_k = \mathbf{a}(\phi)s_k + \mathbf{n}_k, \tag{1}$$

where $\mathbf{a}(\phi)$ is the array manifold vector associated with spatial frequency ϕ of the desired signal s_k, and $\mathbf{n}_k$ is the independent interference-plus-noise component. The equivalent discrete-time baseband model is assumed of the form of Eq. (1) in which s_k is assumed to be a zero-mean white Gaussian processes and the noise is assumed to be independent of s_k and Gaussian with covariance $R_n = \sigma_w^2\mathbf{I} + \sum_{j=1}^{L}\mathbf{a}(\phi_j)\mathbf{a}(\phi_j)^H\sigma_j^2$, where ϕ_j corresponds to the DOA of the j-th interferer. Assuming plane wave propagation, the array vector has the form

$$\mathbf{a}(\phi) = [1\ e^{j\phi}\ \dots\ e^{j(N-1)\phi}]^T, \tag{2}$$

where ϕ is defined over $[-\pi, \pi]$. Note that the spatial frequency is related monotonically to DOA θ through $\phi \propto \sin\theta$.

It is well known that MVDR beamformer at direction ϕ has the form [2]

$$\mathbf{w}_{\mathrm{MV}}(\phi) = \frac{R_x^{-1}\mathbf{a}(\phi)}{\mathbf{a}(\phi)^H R_x^{-1}\mathbf{a}(\phi)}. \tag{3}$$

The MVDR estimate is closely related to the minimum mean square error (MMSE) estimate of s_k is given by

$$\hat{s}_k = E[s_k|\mathbf{x}_k;\phi] = \left(\sigma_s^2 R_x^{-1}\mathbf{a}(\phi)\right)^H \mathbf{x}_k \tag{4}$$

$$= \mathbf{w}_{\mathrm{MS}}(\phi)^H\mathbf{x}_k. \tag{5}$$

It can be shown that $\mathbf{w}_{\mathrm{MV}} = \mathbf{w}_{\mathrm{MS}}(1 + \mathrm{SNR}(\phi)^{-1})$, such that at high signal-to-noise ratio (SNR), they become equivalent (where $\mathrm{SNR}(\phi) = \sigma_s^2\mathbf{a}(\phi)^H R_n^{-1}\mathbf{a}(\phi)$.) The minimum variance (MV) spatial power spectrum estimate is a measurement of signal power in the spatial domain. In [1], the authors show that the weighting coefficients of the Bayesian mixture, i.e., the discretized a posteriori PDF, can be approximately computed from the MV power

0-7803-7997-7/03/$17.00 © 2003 IEEE

spectrum estimate. The MV power spectrum estimate can be derived directly from the output power of the MVDR beamformer by

$$P(\phi) = \mathbf{w}_{MV}(\phi)^H R_x \mathbf{w}_{MV}(\phi) \tag{6}$$

$$= \frac{1}{\mathbf{a}(\phi)^H R_x^{-1} \mathbf{a}(\phi)} \tag{7}$$

$$= \sigma_s^2 (1 + \text{SNR}(\phi)^{-1}). \tag{8}$$

The underlying structure of the MV power spectrum estimate is investigated as follows. Expanding the denominator of the above expression, we have

$$\mathbf{a}(\phi)^H R_x^{-1} \mathbf{a}(\phi) = \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} A_{m,n} e^{j(m-n)\phi} \tag{9}$$

where $A_{m,n}$ is the mth row and nth column element in the matrix R_x^{-1}. Since R_x^{-1} is Hermitian, $A_{m,n}$ is the complex conjugate of $A_{n,m}$. Collecting the common exponential terms together gives

$$\mathbf{a}(\phi)^H R_x^{-1} \mathbf{a}(\phi) = q_0 + \sum_{n=1}^{N-1} [q_n^* e^{jn\phi} + q_n e^{-jn\phi}] \tag{10}$$

$$= q_0 + 2 \sum_{n=1}^{N-1} |q_n| \cos(n\phi - \angle q_n) \tag{11}$$

where q_n is the sum of the diagonal elements of R_x^{-1}:

$$\begin{aligned} q_0 &= A_{0,0} + A_{1,1} + \ldots + A_{N-1,N-1} \\ q_1 &= A_{0,1} + \ldots + A_{N-2,N-1} \\ &\vdots \\ q_{N-2} &= A_{0,N-2} + A_{1,N-1} \\ q_{N-1} &= A_{0,N-1}. \end{aligned} \tag{12}$$

The nth harmonic component in (11) can be written as

$$|q_n| \cos(n\phi - \angle q_n) = \Re\{q_n^* e^{jn\phi}\} \tag{13}$$

The reciprocal of the power spectrum (11) becomes

$$\mathbf{a}(\phi)^H R_x^{-1} \mathbf{a}(\phi) = 2\Re\{\mathbf{a}(\phi)^H \mathbf{q}\} - q_0. \tag{14}$$

The term $\mathbf{a}(\phi)^H \mathbf{q}$ is identical to the discrete time Fourier transform (DTFT) of $\mathbf{q}$. Therefore, the MV power spectrum estimate can be computed directly from the sequence $\mathbf{q}$.

3. BAYESIAN BEAMFORMER

In this section we consider the case when the source direction is unknown. The conditional mean of s_k given a collection of k data points $\mathbf{X}_k = \{\mathbf{x}_1, \ldots, \mathbf{x}_k\}$ yields the MMSE estimate, that is,

$$\hat{s}_k = E[s_k | \mathbf{X}_k]. \tag{15}$$

For notational simplicity, the time index k is dropped, and we assume sufficient SNR such that the MVDR and MMSE beamformers are equivalent. Expanding the conditional mean gives

$$\hat{s} = E_\phi[E[s|\mathbf{X}, \phi]] \tag{16}$$

$$= \int_{-\pi}^{\pi} p(\phi|\mathbf{X}) E[s|\mathbf{X}, \phi] \, d\phi \tag{17}$$

$$= \int_{-\pi}^{\pi} p(\phi|\mathbf{X}) \mathbf{w}_{MV}(\phi)^H \mathbf{x} \, d\phi \tag{18}$$

which is a Bayesian mixture of MVDR beamformers over all ϕ from $-\pi$ to π. The a posteriori PDF $p(\phi|\mathbf{X})$ can be written as

$$p(\phi|\mathbf{X}) = \frac{p(\phi) p(\mathbf{X}|\phi)}{\int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) \, d\phi}. \tag{19}$$

The a posteriori PDF depends on both ϕ and $\mathbf{X}$, which vary according to the a priori PDF $p(\phi)$ and the likelihood function $p(\mathbf{X}|\phi)$. The former carries prior information about the true DOA. In practice, the source signal is usually known to arrive only within a particular range $[\phi_a, \phi_b]$. The a priori PDF can be set to zero outside this range of interest.

Recall that at any given ϕ, the MVDR beamformer is given by

$$\mathbf{w}_{MV}(\phi)^H \mathbf{x} = \frac{\mathbf{a}(\phi)^H R_x^{-1} \mathbf{x}}{\mathbf{a}(\phi)^H R_x^{-1} \mathbf{a}(\phi)} \tag{20}$$

$$= P(\phi) \cdot \sum_{n=0}^{N-1} \tilde{r}_n e^{-jn\phi}, \tag{21}$$

where $\tilde{r}_n$ is the nth element in the Nx1 vector $R_x^{-1}\mathbf{x}$ and $P(\phi)$ is the MV power spectrum estimate. Together with (18), the Bayesian signal estimate $\hat{s}$ has the following form:

$$\hat{s} = \sum_{n=1}^{N-1} \tilde{r}_n \frac{\int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) P(\phi) e^{-jn\phi} \, d\phi}{\int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) \, d\phi}. \tag{22}$$

The numerator has the structure of the conjugate of an inverse DTFT. Define variables $\mathbf{g}$ and κ as

$$g_n \triangleq \int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) P(\phi) e^{jn\phi} d\phi, \quad n = 0, \ldots, N-1 \tag{23}$$

$$\kappa \triangleq \int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) d\phi, \tag{24}$$

where g_n is the nth element of the Nx1 vector $\mathbf{g}$ and κ is a scalar. In terms of these newly defined variables, (22) has the form

$$\hat{s} = \frac{1}{\kappa} \mathbf{g}^H R_x^{-1} \mathbf{x}. \tag{25}$$

If the source signal power is unknown and to be estimated, (for example for the MMSE form of the beamformer), then similar techniques can be used. Using the MV power spectral estimate, the Bayesian power estimate is

$$\begin{aligned} \sigma_B^2 &= E_\phi[P(\phi)|\mathbf{X}] \\ &= \frac{\int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) P(\phi) \, d\phi}{\int_{-\pi}^{\pi} p(\phi) p(\mathbf{X}|\phi) \, d\phi} \\ &= \frac{g_0}{\kappa}, \end{aligned} \tag{26}$$

where g_0 is the first element of the vector $\mathbf{g}$. With this substitution, the Bayesian signal estimate becomes

$$\hat{s} = \sigma_B^2 \, \mathbf{a}_B^H R_x^{-1} \mathbf{x}, \tag{27}$$

where $\mathbf{a}_B$ is defined to be $\mathbf{g}/g_0$.

Thus, the Bayesian beamformer is given by

$$\mathbf{w}_B = \sigma_B^2 R_x^{-1} \mathbf{a}_B \tag{28}$$

$$= \frac{1}{\kappa} R_x^{-1} \mathbf{g}, \tag{29}$$

which yields the identical estimate to that given in Eq. (25), such that when the signal power is estimated in this manner, the MMSE and MVDR beamformers are indeed identical.

4. APPROXIMATION

Several assumptions are made here so that the Bayesian beamformer can implemented efficiently. We assume that there are no interference signals within the range of interest $[\phi_a, \phi_b]$. That is, the source signal must be the only signal in the region of support of the a priori PDF. With this assumption, the likelihood function $p(\mathbf{X}|\phi)$ in (19) can be approximated by [1]

$$p(\mathbf{X}|\phi) \approx \hat{p}(\mathbf{X}|\phi) = c \cdot \exp(k\gamma P(\phi)), \qquad (30)$$

where,

$$\gamma \triangleq \frac{N}{\sigma_w^2} \frac{N\sigma_s^2/\sigma_w^2}{(1 + N\sigma_s^2/\sigma_w^2)}. \qquad (31)$$

The constant c is a normalization factor and k is the number of data samples used to approximate R_x; both variables are independent of ϕ. The variable γ is a function of N, σ_s^2, and σ_w^2, and controls the amplification of the spectral variation along the a posteriori PDF and the power spectrum. For simplicity, γ can be fixed as a constant. This avoids the need to estimate the unknown variables σ_s^2 or σ_w^2. The value chosen should be such that the spectral variation is not overly amplified or damped.

5. IMPLEMENTATION USING FFT

To approximate the continuous integrals, we discretize ϕ and evaluate the DTFT using the discrete Fourier transform (DFT). The effectiveness of this approach is illustrated by the use of the FFT when the data length is chosen to be a power of 2.

According to (14), the MV power spectrum estimate can be written as

$$P(\phi) = \frac{1}{2\Re\{\mathbf{a}(\phi)^H \mathbf{q}\} - q_0} \qquad (32)$$

where $\mathbf{q}$ is described in (12). The term $\mathbf{a}(\phi)^H\mathbf{q}$ is in the form of

$$\mathbf{a}(\phi)^H \mathbf{q} = \sum_{n=0}^{N-1} q_n e^{-jn\phi} = Q(\phi) \qquad (33)$$

which is equivalent to the DTFT of $\mathbf{q}$. Since $\mathbf{q}$ has finite length N, taking an N-point DFT on $\mathbf{q}$ yields $Q(\phi_m)$ that are samples along $Q(\phi)$ evaluated at uniformly spaced points $\phi_m = \frac{2\pi m}{N}, m = 0, ..., N-1$ [3]. The spectral samples $P(\phi_m)$ can then be computed directly from $Q(\phi_m)$ by

$$P(\phi_m) = \frac{1}{2\Re\{Q(\phi_m)\} - q_0}, \quad m = 0, \ldots, N-1. \qquad (34)$$

To obtain more sample points along the spectrum, extra points can be obtained by interpolation between the N critical samples $Q(\phi_m)$, and this can be done easily by performing an M-point DFT on a zero-padded Mx1 vector $\tilde{\mathbf{q}}$ where $M \geq N$:

$$\tilde{\mathbf{q}} = \begin{bmatrix} \mathbf{q} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \qquad (35)$$

This in turn generates M spatial samples $P(\phi_m)$ for $m = 1, ..., M$ along the MV power spectrum estimate $P(\phi)$. Figure 1 illustrates the relation between the original continuous MV power spectrum estimate and its spectral samples obtained from interpolation. The solid line represents the continuous power spectrum, the circles represent the N critical spectral samples, and the crosses represent the $M \geq N$ interpolated samples.

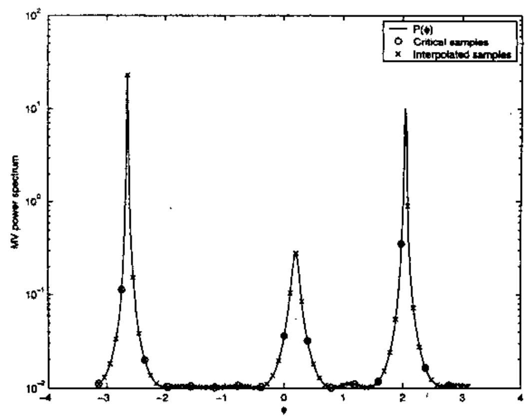

Fig. 1. Samples along the MV power spectrum estimate.

The computation of $\mathbf{g}$ involves the inverse DTFT, which can be implemented by inverse DFT. Using the approximate a posteriori PDF $\hat{p}(\mathbf{X}|\phi)$, the vector $\mathbf{g}$ can be approximated by evaluating the inverse DFT of the sequence $c\ p(\phi_m)P(\phi_m)\exp(k\gamma P(\phi_m))$; $m = 1, \ldots, M$, and the variable κ can be approximated by the sum of the sequence $c\ p(\phi_m)\exp(k\gamma P(\phi_m)); m = 1, \ldots, M$. Due to the zero-padding of $\mathbf{q}$, the inverse transform results in a length-M sequence. As the $M - N$ samples are the consequence of zero-padding and carry only redundant information, the sequence can be truncated, and the first length-N sequence is returned as $\hat{\mathbf{g}}$. The approximate Bayesian beamformer is thus

$$\mathbf{w}_B \approx \hat{\mathbf{w}}_B = \frac{1}{\hat{\kappa}} R_x^{-1} \hat{\mathbf{g}}. \qquad (36)$$

Note that the normalization constant c is canceled and need not be evaluated.

6. ALGORITHM

Let $\mathbf{F}$ be the M-point DFT operator, and let $\hat{\mathbf{P}}$ and $\mathbf{p}$ be Mx1 vectors that represent the discretized power spectrum estimate and the a priori PDF, respectively. The steps to perform Bayesian beamforming are listed as follows:

1. $\hat{R}_x = \frac{1}{k}\sum_{i=1}^{k} \mathbf{x}_i\mathbf{x}_i^H$
2. Compute $\hat{\mathbf{q}}$ from $\hat{R}_x^{-1}$ by (12)
3. Extend the length of $\hat{\mathbf{q}}$ to M by zero-padding
4. $\hat{\mathbf{P}} = (2\Re\{\mathbf{F}\hat{\mathbf{q}}\} - q_0)^{-1}$
5. $\hat{\mathbf{g}} = \mathbf{F}^H(\hat{\mathbf{p}} \odot \hat{\mathbf{P}} \odot \exp(k\gamma\hat{\mathbf{P}}))$
6. $\hat{\kappa} = \sum(\hat{\mathbf{p}} \odot \exp(k\gamma\hat{\mathbf{P}}))$
7. $\hat{s} = \frac{1}{\hat{\kappa}}\hat{\mathbf{g}}^H\hat{R}_x^{-1}\mathbf{x}$,

where, in 4, the inverse is taken element-wise and the symbol $\odot$ stands for Schur-Hadamard multiplication. When M is highly composite, the FFT algorithm can be used to evaluate the DFT.

7. COMPUTATIONAL COST ANALYSIS

We compare the computational cost involved in the proposed beamformer with that in [1] by finding the number of complex multiplications involved. Here, the cost of computing $\hat{R}_x^{-1}$ and γ is not considered because the two variables are common in the implementation of the two beamformers. We ignore the cost to evaluate scalar inverses and scalar exponentials.

Let the number of DOA candidates be M where M is a power of 2, and let the number of sensors be N. The beamformer in [1] performs MVDR beamforming ($O(N^2)$) on each of the M candidate points. The total operation time is in $O(MN^2)$. In the proposed implementation, a direct computation of an M-point DFT using the FFT algorithm requires $M \log M$ complex multiplications. The power spectrum vector P and PDF $\mathbf{p}$ are real and thus only invoke real operations. Putting them together, the proposed algorithm runs in $O(M \log M + N^2)$. For any value of $N > O(\sqrt{\log M})$, (equivalently, for M as large as $O(2^{N^2})$), the cost of the FFT-based algorithm is lower than that in [1]. Note that the computational cost estimate of the proposed scheme is conservative because it includes the trivial multiplications (by zeros) involved in FFT.

8. SIMULATION

In this simulation we illustrate the beampattern and a posteriori PDF of the proposed beamformer as the number of observed data samples k increases. For the simulation results shown in figure 2, the number of sensors N is 16; the spatial direction range of interest is $[-0.5, 0.5]$, and the a priori PDF is chosen to be uniform over this interval; the variable γ is chosen to be 0.3; the SNR is 0dB; Two strong interference signals are present at -1 and 0.7 respectively. On the right column, the solid line indicates the a posteriori PDF and the dotted line indicated the a priori PDF.

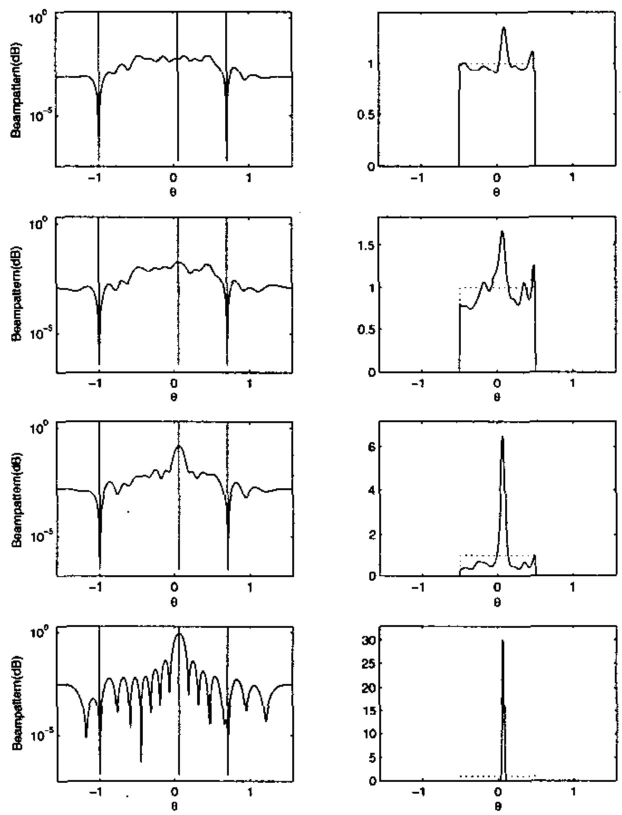

Fig. 2. Beampattern (left) and a priori and a posteriori PDF (right) of the proposed beamformer for $k = \{10, 30, 100, 500\}$ from top to bottom.

9. REFERENCES

[1] K. L. Bell, Y. Ephraim, and H. L. V. Trees, "A Bayesian approach to robust adaptive beamforming," *IEEE Trans. Signal Processing*, vol. 48, pp. 386–398, Feburary 2000.

[2] H. L. V. Trees, *Optimum Array Processing*. New York: Wiley Interscience, 2002.

[3] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*. Englewood Cliffs, New Jersey: Prentice-Hall, 1989.