An Adaptive ESPRIT Based on URV Decomposition*

K.J.R. Liu[1], D.P. O'Leary[2], G.W. Stewart[2], and Y.-J.J. Wu[3]

1. Electrical Engineering Department and ISR, University of Maryland, College Park, MD 20742
2. Computer Science Department, University of Maryland, College Park, MD 20742
3. Applied Mathematics, University of Maryland, College Park, MD 20742

Abstract

ESPRIT is an algorithm for determining the fixed directions of arrival of a set of narrowband signals at an array of sensors. Unfortunately, its computational burden makes it unsuitable for real time processing of signals with time-varying directions of arrival. In this work we develop a new implementation of ESPRIT that has potential for real time processing. It is based on a rank-revealing URV decomposition, rather than the eigendecomposition or singular value decomposition used in previous ESPRIT algorithms. We demonstrate its performance on simulated data representing both constant and time-varying signals. We find that the URV-based ESPRIT algorithm is effective for estimating time-varying directions-of-arrival at considerable computational savings over the SVD-based algorithm.

1 Introduction

The ESPRIT algorithm [7] is a method for determining directions-of-arrival (DOA) of a set of narrowband signals impinging on an array of m sensors with translational invariance. It handles array geometries almost as general as those of the MUSIC algorithm [8] at a significant computational savings.

A key limitation of both the MUSIC and ESPRIT algorithms is the work required to process a new sample. At the heart of the algorithms is the separation of the m dimensional sample space into a signal subspace and a noise subspace. Usually this separation is done by computing the eigendecomposition of the estimated covariance matrix or part of the singular value decomposition of the data matrix. Unfortunately, these decompositions require $O(m^3)$ operations to update, making the algorithms unsuitable for real-time computation. Some attempts have been made to reduce the updating complexity by maintaining an approximate singular value decomposition(e.g., [3, 6]), but we believe that better results can be obtained using an alternate decomposition.

Recently, Stewart [9] has introduced the rank-revealing URV decomposition, a new matrix decomposition that produces the signal and noise subspaces, but can be updated in $O(m^2)$ time sequentially and in $O(m)$ time on an array of m processors. This means that algorithms that previously depended on eigendecomposition or singular value decomposition may now be practical in real time applications, provided the URV decomposition can be successfully substituted. Boman, Griffin, and Stewart [1] have already exploited this fact to accelerate the MUSIC algorithm. The purpose of this paper is to investigate the use of the URV decomposition in time-varying signal processing using ESPRIT.

2 The ESPRIT Algorithm

Roy and Kailath [7] noted that the ESPRIT idea is independent of the choice of matrix decompositions used in its implementation. In order to highlight this basic idea, in this section we describe ESPRIT in terms free of reference to specific decompositions. We hope that this description will be of independent interest.

2.1 The background

We consider d narrow-band plane waves simultaneously incident on a planar array of m sensors (m even), arranged in $m/2$ doublet pairs. The displacement between sensors in a pair is constant in both direction and magnitude Δ, but the location of each pair is arbitrary. The wave sources are assumed to be located in the same plane, and the location of each source is specified by a single parameter $\theta_i \in [-\pi, \pi]$, the DOA of the ith source. Quantities related to the first and second sensors in each pair are subscripted by 1 and 2 respectively. All vectors are column vectors.

Given data from the array of sensors, the DOA estimation problem is to locate the directions of the sources. If the narrowband signals have the same known center frequency ω_0, then the DOA problem can be described by a simple model. The relationship between the unknown signal $s(t) \in C^d$ and the sensor output $x_1(t) \in C^{m/2}$ and $x_2(t) \in C^{m/2}$ is given by

$$x_1(t) = As(t) + e_1(t), \qquad (1)$$
$$x_2(t) = A\Phi s(t) + e_2(t), \qquad (2)$$

or

$$x(t) = \begin{pmatrix} A \\ A\Phi \end{pmatrix} s(t) + e(t),$$

where $e(t)$ is the measurement noise, and $A \in C^{m/2 \times d}$ is the unknown matrix of array responses or array steering vectors. The diagonal matrix Φ is also unknown, and is related to the phase delays between the sensors in each doublet pair:

$$\phi_i = e^{j\omega_0 \Delta \sin \theta_i / c}, \qquad i = 1, ..., d. \qquad (3)$$

Our task is to estimate the number of signals d and the directions-of-arrival, θ_i. For this it is sufficient to estimate the matrix Φ, which is the idea underlying ESPRIT.

2.2 ESPRIT

To derive the ESPRIT algorithm, it is convenient to assume noise-free signals, so that certain subspaces are unambiguously defined. The ESPRIT algorithm exploits the array geometry in the following way. Let the n columns of

$$X = \begin{pmatrix} X_1 \\ X_2 \end{pmatrix}$$

form an ensemble of n snapshots, and let S be the corresponding matrix of signals. If the d columns of

$$\left(\begin{array}{c} V_1 \\ V_2 \end{array} \right)$$

form a basis for the column space of X, it follows from (1) and (2) that

$$\left(\begin{array}{c} V_1 \\ V_2 \end{array} \right) = \left(\begin{array}{c} ASB \\ A\Phi SB \end{array} \right) \tag{4}$$

where SB is a nonsingular matrix of order d.

Next let the rows of

$$(W_1 \ W_2)$$

form a basis for the row space of $(V_1 \ V_2)$. It follows from (4) that W can be written in the form

$$(W_1 \ W_2) = (CASB \ CA\Phi SB),$$

where CA is a nonsingular matrix of order d. Since CA and SB are nonsingular and Φ is diagonal, the diagonal elements of Φ are the eigenvalues of the pencil

$$W_2 - \lambda W_1 = (CA)\Phi(SB) - \lambda(CA)(SB).$$

Given Φ, the directions of arrival can be found from (3).

The above description of ESPRIT leads to many algorithms, even in the absence of noise, since there are many choices of bases. For example, if the bases $(V_1^H \ V_2^H)^H$ and $(W_1 \ W_2)$ are chosen to be orthonormal, then the resulting pencil is orthogonally equivalent to the pencil resulting from using the singular value decomposition to implement *total least squares* ESPRIT (the variant studied in this paper). In the presence of noise, one is faced with the additional problem of estimating the subspaces and their dimensions. The role of specific decompositions in ESPRIT algorithms is to make these estimations possible.

It is worth noting that the above derivation provides an alternate interpretation of ESPRIT: namely, the problem reduces to choosing the matrices B and C that define the required bases. In particular, in the absence of noise, virtually any full rank matrices B and C will do (technically, any full rank matrices except those in a set of measure zero). Although we do not pursue this approach here, we feel it has the potential to yield fast algorithms that will work when there is a good signal-to-noise ratio.

3 SVD ESPRIT

We now turn to the estimation of bases for the subspaces required by ESPRIT. In this section we consider the use of the singular value decomposition. We call the resulting algorithm SVD ESPRIT .

The SVD algorithm finds unitary matrices U and V such that

$$X^H = UDV^H, \tag{5}$$

where the diagonal elements of $D = \text{diag}(\delta_1, \delta_2, \ldots, \delta_m)$ are nonnegative and in descending order of magnitude. If δ_d is judged to be sufficiently large and δ_{d+1} is judged to be sufficiently small, then we conclude that there are d signals, and the first d columns of V furnish the required basis for the signal subspace.

The sum $\delta_{d+1}^2 + \cdots + \delta_m^2$ approximates the sum of squares of the projection of the error onto the orthogonal complement of the signal subspace (i.e., onto the noise subspace). If the individual components of the noise have variance σ^2, then the expected value of this sum is $n(m - d)\sigma^2$ for rectangular windowing and approximately $(m - d)\sigma^2/(1 - \mu^2)$ for exponential windowing, where μ is the forgetting factor. Consequently, it is reasonable to choose d so that

$$\sqrt{\delta_{d+1}^2 + \cdots \delta_m^2} \leq \psi_d \sigma \sqrt{n(m - d)}$$

for rectangular windowing and

$$\sqrt{\delta_{d+1}^2 + \cdots \delta_m^2} \leq \psi_d \sigma \sqrt{\frac{m - d}{1 - \mu^2}}$$

for exponential windowing. Here $\psi_d > 1$ is a factor chosen to make it unlikely that the dimension of the signal subspace is overestimated. We vary it in our experiments.

A basis for the row space of $(V_1 \ V_2)$ can also be determined by calculating a SVD. Specifically, let

$$(V_1 \ V_2) = TD_V W^H$$

be the SVD of $(V_1 \ V_2)$. Then the first d rows of W^H form the required basis $(W_1 \ W_2)$.

Here it is necessary to say a word about the computation of the SVD. There exist algorithms that compute the SVD directly from the original matrix, and in applications with very small singular values where high accuracy is required they are the algorithms of choice. However, they are expensive, and for the DOA problem there are less costly alternatives. Specifically, from (5), it follows that

$$XX^H = VD^2V^H.$$

Thus V is the matrix of eigenvectors of the cross-product matrix XX^H and can be computed using any of many programs to solve the Hermitian eigenvalue problem. This approach is particularly attractive, because the cross-product matrix can be easily updated as signals arrive.

Even with these economies, the SVD ESPRIT algorithm is expensive, requiring the $O(m^3)$ solution of an eigenvalue problem with each snapshot. Unfortunately, updating the eigendecomposition results in another $O(m^3)$ algorithm, though the order constant is smaller [2]. Recently techniques for approximately updating an eigendecomposition have been proposed [3, 6] , and they show some promise. However, in this paper we consider an alternative decomposition that can be updated in $O(m^2)$ time.

4 URV ESPRIT

The rank-revealing URV decomposition expresses X^H in the form

$$X^H = U \left(\begin{array}{cc} R & F \\ 0 & G \end{array} \right) V^H \tag{6}$$

where the columns of U and V are orthonormal, R and G are upper triangular of orders d and $m - d$, and F and G are small in norm. This decomposition reveals that X is within $\sqrt{\|F\|^2 + \|G\|^2}$ of the matrix $\hat{X}$ of rank d obtained by setting F and G in (6) to zero. The column space of $\hat{X}$ is the space spanned by the first d columns of V, and those columns are therefore a natural candidate for the basis required by the ESPRIT algorithm. The matrix U is unnecessary and is not stored or updated.

In the same way, the URV decomposition of the matrix $(V_1 \ V_2)$ can be used to determine the matrices W_1 and W_2 of the ESPRIT algorithm. We call the resulting algorithm the URV ESPRIT algorithm.

The URV decomposition can be updated in $O(m^2)$ time (and in $O(m)$ time on a linear array of m processors). The updating procedure consists of two parts: an incorporation step and a deflation step. The incorporation is analogous to the standard update of a QR decomposition [4]; however, special care is taken that only the first column of F and G increases in norm. This corresponds to the fact that the addition of a row to a matrix can increase its rank by at most one. After the update, a condition estimator [5] is used to test R for rank degeneracy, and a deflation step reduces the norm of the last column of R. If a degeneracy is detected, a refinement step is performed to bring the decomposition closer to diagonal form. All transformations are accomplished by plane rotations,

and the algorithm is stable. We state the URV ESPRIT algorithm in the appendix.

So far as rank determination is concerned, the quantity $\|G\|$ is the analogue of $\sqrt{\delta_{d+1}^2 + \cdots \delta_m^2}$ for the SVD. Consequently we attempt to choose d so that

$$\|G\| \leq \psi_d \sigma \sqrt{n(m-d)}$$

for rectangular windowing and

$$\|G\| \leq \psi_d \sigma \sqrt{\frac{m-d}{1-\mu^2}}$$

for exponential windowing. These criteria are applied at the deflation step.

However, during the incorporation step a decision must be made as to whether G has grown in norm due to an increase in rank. Here we use the same criterion, but with a different factor ψ_u replacing ψ_d. As ψ_u grows, the signal subspace changes less frequently. Thus, ψ_u can be seen as a factor that controls the accuracy of the approximate signal subspace. In applications where only low accuracy is required, it may be taken large with a resulting savings in work. An alternative of this kind is not available for the SVD.

5   Experimental Results

In this section, we present some simulation results that compare the performance of the two algorithms: SVD ESPRIT (using the estimated covariance matrix) and URV ESPRIT.

We use a five-pair ($m = 10$) linear array with pair spacing $\lambda/4$. The pairs are equally spaced on a line with relative locations $[0, 1, 2, 3, 4]$ λ. The two signals are narrow-band with signal-to-noise ratio (SNR) 23dB and 20dB, respectively. The noise is Gaussian, and the algorithms were tested with duplicate data samples in order to make a fair comparison.

We say that an algorithm *failed* at a particular time if it estimated more or fewer than two signals.

The first example concerns two fixed signal sources located at 24° and 29° and with 50% initial correlation. For each trial, we estimated the DOAs based on 100 data samples, and we ran 2000 trials. Figure 1 shows a histogram and tabular summary of the results. Both algorithms were quite successful. This shows that we are not sacrificing much accuracy by substituting the more economical URV for the SVD.

Other experiments concerned time-varying DOAs. We used exponential windowing, with a forgetting factor $\mu = .9$. One data set was used for close sources located at

$$\begin{matrix} 10° + 5° \sin(2\pi n/360) \\ 20° + 5° \sin(2\pi n/240) \end{matrix} \quad n = 1, 2, \ldots, 719.$$

This corresponds to a sampling rate of 1 data point per .11°, .08°, or .05° change in angle. Typical radar applications produce 1 point per $10^{(-5)}$ ° change, so our experimental setup is much more demanding.

Experiments varying the rank determination tolerances $\psi_d \in \{3, 6\}$ and $\psi_u \in \{0, 1, 2\}$ for signals of different signal-to-noise ratio, separation angle, and rate of change of DOA showed that results improved as ψ_u was decreased, but that the value of ψ_d is more problem dependent. The error model we use accounts for random errors in the measurements, but not for movement of the sources. As the rate of change in the DOA increases, the value of ψ_d must be increased in order to account for this extra source of error. The experiments presented here use the values $\psi_d = 3$ and $\psi_u = 1$.

Figure 2 gives the results for the close sources. The DOAs were updated every two data samples. Both algorithms perform quite

well, and give similar numbers of failures and average error. There seems to be no reason to prefer the more expensive SVD-based algorithm over the URV , and the results indicate that in practice, low sampling rates can be tolerated well. Extra data could be used for noise reduction through averaging, feeding the time-averaged data to ESPRIT.

As a final example, to demonstrate tracking of instantaneously changing signals, we assumed that there were two signal sources located at 24° and 29°, each with SNR 23 dB, and that the signals alternatively appear and disappear. We took a rectangular window of size 10. Figure 3 shows the similar good performance of the two algorithms.

These experimental results lead us to believe that the URV-based ESPRIT algorithm can be successfully used for real-time tracking of time-varying signals.

6   Summary

We have presented a new variant of the ESPRIT algorithm that has potential for real-time tracking of moving signals. It has the following features:

- The storage requirement is $O(m^2)$ (plus mn for rectangular windowing).

- The work per update is $O(m^2 + d^3)$.

- Its performance is similar to SVD-based algorithms, at greatly reduced computational cost, and it admits an efficient parallel realization.

Appendix: The Time-Varying URV ESPRIT Algorithm

Suppose that we already have a rank-revealing URV decomposition of the data matrix from the previous time. For rectangular windowing, we also save the most recent n data samples.

1) Obtain the new data sample x.

2) Update the previous rank-revealing URV decomposition of the matrix of data samples by downdating and updating the previous factors if rectangular windowing is used, or updating the previous factors if exponential windowing is used.

3) Estimate the number of sources d (i.e., the rank of X) using a tolerance of ψ_d times the expected value of the noise. (The parameter ψ_d is chosen by the user.)

4) The basis for the signal subspace (the range of X^H) is E_X, equal to the first d columns of the V factor in the URV decomposition.

5) Partition E_X into $m/2 \times d$ blocks corresponding to the two sets of sensors:

$$E_X = \left(\begin{matrix} E_1 \\ E_2 \end{matrix} \right)$$

and compute a rank-revealing URV decomposition of (E_1, E_2):

$$(E_1, E_2) = W \Psi V^H.$$

6) Partition V into $d \times d$ blocks as

$$V = \left(\begin{matrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{matrix} \right)$$

and calculate the eigenvalues ϕ_i of $-V_{12} V_{22}^{-1}$ (or, equivalently, the eigenvalues of $V_{11}^{-H} V_{21}^{H}$).

7) Estimate the DOAs from ϕ_i using (3).

References

[1] E. C. Boman, M. F. Griffen, and G. W. Stewart. Direction of arrival and the rank-revealing URV decomposition. In *Proceedings of ACASSP-91*, Washington, DC, 1991. IEEE.

[2] J. R. Bunch and C. P. Nielsen. Updating the singular value decomposition. *Numerische Mathematik*, 31:111–129, 1978.

[3] W. Ferzali and J. Proakis. Adaptive SVD algorithm with application to narrowband signal tracking. In R. J. Vaccaro, editor, *SVD and Signal Processing, II*, pages 149–160, Amsterdam, 1990. Elsevier Science Publishers.

[4] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 1989.

[5] N. J. Higham. A survey of condition number estimation for triangular matrices. *SIAM Review*, 29:575–596, 1987.

[6] M. Moonen, P. Van Dooren, and J. Vandewalle. Combined Jacobi-type algorithms in signal processing. In R. J. Vaccaro, editor, *SVD and Signal Processing, II*, pages 177–188, Amsterdam, 1990. Elsevier Science Publishers.

[7] R. Roy and T. Kailath. ESPRIT – estimation of signal parameters via rotational invariance techniques. In F. A. Grünbaum, J. W. Helton, and P. Khargonekar, editors, *Signal Processing Part II: Control Theory and Applications*, pages 369–411. Springer-Verlag, New York, 1990.

[8] R. O. Schmidt. *A Signal Subspace Approach to Multiple Emitter Location and Spectral Estimation*. PhD thesis, Stanford University, 1981.

[9] G. W. Stewart. An updating algorithm for subspace tracking. Technical Report CS-TR 2494, Department of Computer Science, University of Maryland, 1990. To appear in *IEEE Transactions on Signal Processing*.
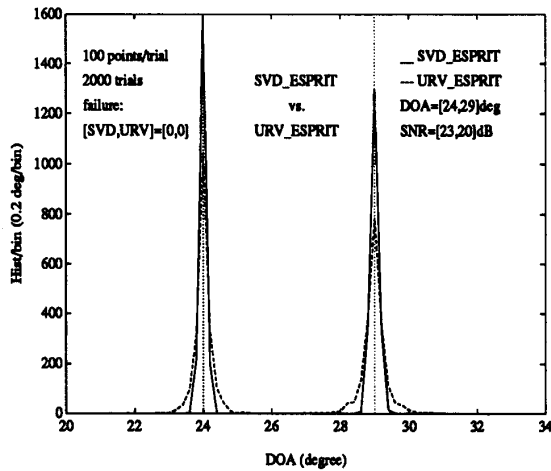
Algorithm	signal 1	signal 2
SVD _ESPRIT	23.9982 ± 0.0822	29.0020 ± 0.1145
URV _ESPRIT	24.0007 ± 0.2384	28.9957 ± 0.3454

Figure 1: Histogram and estimated means and variances from ESPRIT(SVD) and ESPRIT(URV) for fixed sources
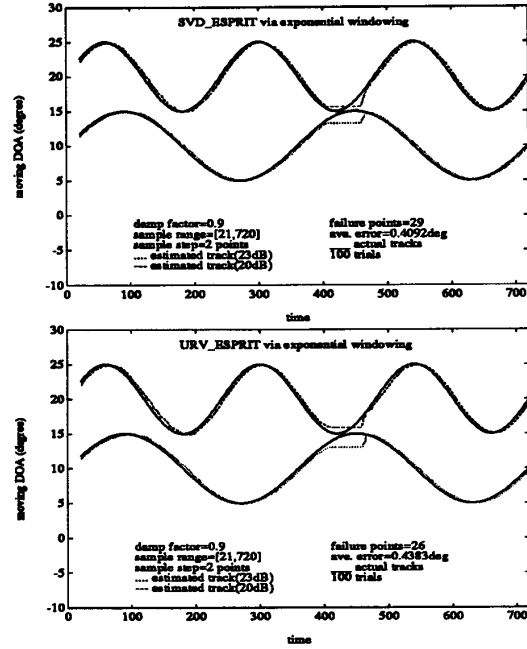

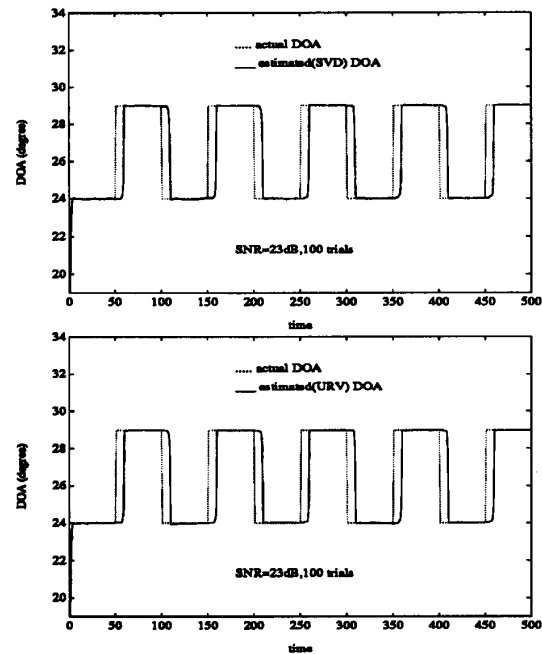Figure 2: Estimated time-varying DOAs for close sources using exponential windowing.

Figure 3: Estimated time-varying DOAs for instantaneously changing signals using rectangular windowing.