

Applications of Truncated QR Methods to Sinusoidal Frequency Estimation

S.F. Hsieh, K.J.R. Liu and K. Yao

Electrical Engineering Department
University of California, Los Angeles
Los Angeles, CA 90024-1594

ABSTRACT

We propose three truncated QR methods for sinusoidal frequency estimation: (1) truncated QR without column pivoting (TQR); (2) truncated QR with pre-ordered columns (TQRR); and (3) truncated QR with column pivoting (TQRP). It is demonstrated that the benefit of truncated SVD (TSVD) for high frequency resolution is achievable under the truncated QR approach with much lower computational cost. Other attractive features of the proposed methods include the ease of updating, which is difficult for the SVD method, and numerical stability. Thus, the TQR methods offers efficient ways for identifying sinusoids closely clustered in frequencies under stationary and nonstationary conditions. Based on the FBLP model, computer simulations and comparisons are provided for different truncation methods under various SNR's.

1 INTRODUCTION

In recent years, there is much interest in seeking efficient methods for spectral estimation. In general, the criteria for an efficient method may include: frequency resolution capability, computational efficiency, updating and downdating capability, and implementable parallel processing structure so that real-time applications are possible. While the SVD-based method is well known for its robustness in resolving closely clustered sinusoids, it is not attractive from the other desirable features of an efficient method as considered above. In this paper, we consider several other promising approaches based on the truncated QR techniques.

In resolving closely spaced frequencies from limited amount of data samples, Tufts and Kumaresan [1] proposed a SVD-based method to solve the forward-backward linear prediction (FBLP) least squares (LS) problem. By providing an excess order in the FBLP model and then truncating small singular values to zero, this truncated SVD method yields a very low SNR threshold and greatly suppresses spurious frequencies. However, the massive computations required by SVD makes it unsuitable for *real time* super-resolution applications. We propose to use truncated QR methods which are more amenable to VLSI implementations, such as on systolic arrays, with insignificantly degraded performances as compared to the TSVD method. Three different truncated QR methods will be considered, depending on the ordering of the columns of the data matrix.

¹This work is partially supported by UC MICRO grant and NASA/Ames Grant NCC 2-374.

A FBLP model for estimating sinusoidal frequencies is formulated first, followed by an introduction of different truncation methods and the minimum-norm solutions. Finally, comparisons of these three QR methods to the TSVD method are given based on computer simulations.

2 FBLP MODEL

Consider a complex-valued data sequence of length n ,

$$\tilde{x}_i = \sum_{k=1}^p c_k e^{j2\pi f_k i} + w_i \equiv x_i + w_i, \quad i = 1, 2, \dots, n, \quad (1)$$

where p is the number of sinusoids, complex-valued c_k comprises the amplitudes and phases of each sinusoid, and w_i is an additive white Gaussian noise. We define the signal-to-noise ratio (SNR) as

$$\text{SNR (dB)} = 20 \log(\|x\|_2 / \|w\|_2). \quad (2)$$

It can be shown [1] that under noise-free conditions, the frequency locations can be obtained by finding the roots of

$$S(z) = 1 - \sum_{k=1}^{\ell} g_k z^{-k} = 0, \quad (3)$$

on the unit circle, where the complex-valued coefficients g'_k s, $k = 1, 2, \dots, \ell$, satisfy the following system of FBLP equations

$$\begin{bmatrix} x_{\ell} & x_{\ell-1} & \cdots & x_1 \\ x_{\ell+1} & x_{\ell} & \cdots & x_2 \\ \vdots & \vdots & \ddots & \vdots \\ x_{n-1} & x_{n-2} & \cdots & x_{n-\ell} \\ x_2^* & x_3^* & \cdots & x_{\ell+1}^* \\ x_3^* & x_4^* & \cdots & x_{\ell+2}^* \\ \vdots & \vdots & \ddots & \vdots \\ x_{n-\ell+1}^* & x_{n-\ell+2}^* & \cdots & x_n^* \end{bmatrix} \begin{bmatrix} g_1 \\ g_2 \\ \vdots \\ g_{\ell} \end{bmatrix} = \begin{bmatrix} x_{\ell+1} \\ x_{\ell+2} \\ \vdots \\ x_n \\ x_1^* \\ x_2^* \\ \vdots \\ x_{n-\ell}^* \end{bmatrix} \quad (4)$$

with $\ell \geq p$ representing the order of the prediction model, and $*$ the complex conjugate. We will assume that $2(n - \ell) > \ell$. For simplicity, denote (4) as

$$Ag = b. \quad (5)$$

The rank of A is p if $2(n - \ell) \geq p$ and $\ell \geq p$. When the noise is present, we use an $\tilde{\cdot}$ on A and b , i.e., $\tilde{A} = A + E$ and $\tilde{b} = b + e$, to denote the noise-corrupted FBLP model and (5) now becomes the FBLP LS problem of

$$\tilde{A}g \approx \tilde{b}, \quad (6)$$

where \tilde{A} usually has full rank due to the perturbation of the noise. One standard approach [1] is to use the TSVD method on (5) to obtain a rank- p approximation of the FBLP matrix \tilde{A} , denoted by $\tilde{A}_{SVD}^{(p)}$, followed by solving for a minimum norm LS solution of g given by

$$\tilde{A}_{SVD}^{(p)} g \cong \tilde{b}. \quad (7)$$

Then the frequencies can be computed by finding the phases of the roots of (3) close to the unit circle or searching for the peaks on the pseudo-spectrum $1/|S(\exp(j2\pi f))|$, $0 \leq f \leq .5$. Notice that the proper choice of the prediction order ℓ depends on p , the number of sinusoids, which may or may not be known in advance. Fig. (1) depicts a flowchart diagram summarizing the estimation of harmonics frequencies based on the FBLP model. Next we consider the rank- p approximation of the FBLP matrix \tilde{A} and subsequently solve for the minimum-norm solution $\tilde{g}^{(p)}$.

3 TRUNCATION METHODS

For many LS problems, ill-conditioning can be troublesome, and truncation methods are known to be useful in stabilizing the solutions at the cost of slightly increased residual errors.

Let

$$\tilde{A} = \tilde{U} \tilde{\Sigma} \tilde{V}^H = [\tilde{U}_1 \ \tilde{U}_2] \begin{bmatrix} \tilde{\Sigma}_1 & 0 \\ 0 & \tilde{\Sigma}_2 \end{bmatrix} \begin{bmatrix} \tilde{V}_1^H \\ \tilde{V}_2^H \end{bmatrix} \quad (8)$$

and

$$\tilde{A}\Pi = \tilde{Q} \tilde{R} = [\tilde{Q}_1 \ \tilde{Q}_2] \begin{bmatrix} \tilde{R}_{11} & \tilde{R}_{12} \\ 0 & \tilde{R}_{22} \end{bmatrix} \quad (9)$$

be the SVD and QR decomposition (QRD) of the $2(n-\ell) \times \ell$ complex-valued matrix \tilde{A} respectively, where H denotes the Hermitian of a matrix or a vector. $\tilde{\Sigma}_1 = \text{diag}(\tilde{\sigma}_1, \dots, \tilde{\sigma}_p)$ and $\tilde{\Sigma}_2 = \text{diag}(\tilde{\sigma}_{p+1}, \dots, \tilde{\sigma}_\ell)$ represent nonincreasing singular values. $\tilde{R}_{11} \in \mathcal{C}^{p \times p}$, $\tilde{R}_{12} \in \mathcal{C}^{p \times (\ell-p)}$, and $\tilde{R}_{22} \in \mathcal{C}^{2(n-\ell)-p \times (\ell-p)}$, while \tilde{R} is an upper-triangular matrix.

$$\tilde{U} = [\tilde{U}_1 \ \tilde{U}_2] = [\tilde{u}_1, \dots, \tilde{u}_p, \tilde{u}_{p+1}, \dots, \tilde{u}_\ell] \in \mathcal{C}^{2(n-\ell) \times \ell}, \quad (10)$$

$$\tilde{V} = [\tilde{V}_1 \ \tilde{V}_2] = [\tilde{v}_1, \dots, \tilde{v}_p, \tilde{v}_{p+1}, \dots, \tilde{v}_\ell] \in \mathcal{C}^{\ell \times \ell}, \quad (11)$$

and

$$\tilde{Q} = [\tilde{Q}_1 \ \tilde{Q}_2] = [\tilde{q}_1, \dots, \tilde{q}_p, \tilde{q}_{p+1}, \dots, \tilde{q}_\ell] \in \mathcal{C}^{2(n-\ell) \times \ell} \quad (12)$$

all have orthonormal columns, i.e., $\tilde{u}_i^H \tilde{u}_j = \tilde{v}_i^H \tilde{v}_j = \tilde{q}_i^H \tilde{q}_j = \delta_{ij}$.

In the absence of noise, $\tilde{\Sigma}_2 = \tilde{R}_{22} = 0$. Here the permutation matrix $\Pi = [\pi_1, \dots, \pi_\ell]$ is used to represent different methods of performing QRD with column interchanges. Now we want to preserve as much of the energy as possible (with respect to the Frobenius norm defined below) in the trapezoidal matrix $[\tilde{R}_{11} \ \tilde{R}_{12}]$ of (9). Equivalently, we want to leave as little as possible the energy residing in the lower right submatrix \tilde{R}_{22} , which will be truncated. This approach amounts to selecting the columns of \tilde{A} in an order such that those columns with largest linear dependency will be selected first.

For QRD with no pivoting, Π is simply an identity matrix. QRD with pre-ordered columns [2] determines Π according to a column index maximum-difference bisection rule. Here we select the first and the ℓ^{th} columns, followed by the column $\lceil \frac{1+\ell}{2} \rceil$ halfway between 1 and ℓ . Then we pick the columns that lie in the midway of those ones which are already selected, i.e.,

$\lceil (1 + \lceil \frac{1+\ell}{2} \rceil)/2 \rceil$, $\lceil (\lceil \frac{1+\ell}{2} \rceil)/2 + \ell/2 \rceil$, and so on. This selection rule does not depend on the real-time data in \tilde{A} . The underlying reason for this *ad hoc* fixed-ordering scheme is to provide the selected columns with a possibly maximum differences or minimum linear dependency among these columns. This scheme was motivated due to the nature of the matrix \tilde{A} arranged in the form of (4) consisting of perturbed sums of harmonic sinusoids. As an example, suppose there are 5 columns, then the pre-ordering strategy leads to [1, 5, 3, 2, 4]. Thus we have $\Pi = [e_1, e_5, e_3, e_2, e_4]$, where e_i is a dimension ℓ column vector with all zero components except for an one at the i^{th} position.

As for QRD with column pivoting [3, p. 233], Π is determined during the QRD process, where $\pi_1 = e_{d_1}$ and $d_1 \in [1, \ell]$ is the index such that \tilde{a}_{d_1} has the largest norm. Continuing with this column-pivoting process on the lower right submatrix yet to be triangularized, we can determine the permutation matrix Π which yields an optimum QRD column ordering strategy in the sense of preserving most energy in the upper trapezoidal submatrix. However, this Π is data-dependent and the extra cost for this pivoting makes it less desirable for some applications.

After forcing those *rank-weakly* quantities to be zero and preserving the most significant p -rank, we can obtain a rank- p approximate of \tilde{A} . For TSVD, $\tilde{\Sigma}_2$ is discarded and

$$\tilde{A}_{TSVD}^{(p)} = \tilde{U}_1 \tilde{\Sigma}_1 \tilde{V}_1^H. \quad (13)$$

Similarly, for TQR, the lower-right submatrix \tilde{R}_{22} is discarded and

$$\tilde{A}_{TQR}^{(p)} \Pi = \tilde{Q}_1 [\tilde{R}_{11} \ \tilde{R}_{12}]. \quad (14)$$

To account for the effect due to truncation, we define the *fractional truncated F-norm* as

$$\mathcal{F}^{(p)} = 1 - \|\tilde{A}^{(p)}\|_F / \|\tilde{A}\|_F \quad (15)$$

where $\|\cdot\|_F$ is the Frobenius norm given by

$$\|A\|_F^2 = \sum_i \sum_j |a_{i,j}|^2. \quad (16)$$

Thus we have

$$\mathcal{F}_{TSVD}^{(p)} = 1 - \sqrt{\Sigma_{j=p+1}^\ell \tilde{\sigma}_j^2 / \Sigma_{j=1}^\ell \tilde{\sigma}_j^2} \quad (17)$$

and

$$\mathcal{F}_{TQR}^{(p)} = 1 - \left(\frac{\|\tilde{R}_{11}\|_F^2 + \|\tilde{R}_{12}\|_F^2}{\|\tilde{R}_{11}\|_F^2 + \|\tilde{R}_{12}\|_F^2 + \|\tilde{R}_{22}\|_F^2} \right)^{1/2}. \quad (18)$$

The relationship among different truncation methods can be shown as follows:

$$0 \leq \mathcal{F}_{TSVD}^{(p)} \leq \mathcal{F}_{TQR}^{(p)} \leq \mathcal{F}_{TQRR}^{(p)}, \mathcal{F}_{TQR}^{(p)} \leq 1. \quad (19)$$

Therefore from the point of view of preserving the Frobenius norm(energy) of a matrix, SVD provides the optimum truncation, with TQR being next, while TQR and TQRR truncate even more.

4 MINIMUM-NORM SOLUTIONS

After truncation, the FBLP LS problem becomes rank-deficient, hence the minimum-norm LS solution is desired. It can be shown that a minimum-norm solution vector $g^{(p)}$ must lie in the row space of the rank-reduced matrix $\tilde{A}^{(p)}$, namely, the row space of

\tilde{V}_1^H or $[\tilde{R}_{11} \ \tilde{R}_{12}]$. For TSVD it is given by

$$g_{TSVD}^{(p)} = \tilde{V}_1 \tilde{\Sigma}_1^{-1} \tilde{U}_1^H \tilde{b}, \quad (20)$$

and for TQR it is

$$g_{TSVD}^{(p)} = \sum_{j=1}^p \frac{\tilde{u}_j^H \tilde{b}}{\tilde{\sigma}_j} \tilde{v}_j. \quad (21)$$

To obtain $g_{TQR}^{(p)}$, we can perform QRD on the right of the trapezoidal upper-triangular matrix in (14) to zero out \tilde{R}_{12} and also obtain the orthonormal row space, \tilde{Z}_1^H , of $[\tilde{R}_{11} \ \tilde{R}_{12}]$. That is

$$\tilde{A}_{TQR}^{(p)} \Pi = \tilde{Q}_1 [\tilde{R}_{11} \ \tilde{R}_{12}] = \tilde{Q}_1 \tilde{T}_{11} \tilde{Z}_1^H, \quad (22)$$

where $\tilde{Z}_1 = [\tilde{z}_1, \dots, \tilde{z}_p] \in \mathcal{C}^{\ell \times p}$ has orthonormal columns and $\tilde{T}_{11} \in \mathcal{C}^{p \times p}$ is an upper-triangular matrix. This is sometimes called a *complete orthogonal factorization* [3, p. 236], and we can consider it as a two-sided direct unitary transformations on a rank-deficient matrix to compress all the energy of a matrix into a *square upper-triangular* matrix. This resembles the SVD method where two-sided iterative unitary transformations are applied to reduce a matrix into a *diagonal* matrix. Then the minimum-norm solution follows by

$$g_{TQR}^{(p)} = \Pi \tilde{Z}_1 \tilde{T}_{11}^{-1} \tilde{Q}_1^H \tilde{b}. \quad (23)$$

It is noted that if no truncation is performed at all and \tilde{A} has full column rank ℓ , then the LS solution from the FBLP model is either obtained from SVD as

$$\begin{aligned} \tilde{g} &= \tilde{V}_1 \tilde{\Sigma}_1^{-1} \tilde{U}_1^H \tilde{b} + \tilde{V}_2 \tilde{\Sigma}_2^{-1} \tilde{U}_2^H \tilde{b} \\ &= \sum_{j=1}^p \frac{\tilde{u}_j^H \tilde{b}}{\tilde{\sigma}_j} \tilde{v}_j + \sum_{j=p+1}^{\ell} \frac{\tilde{u}_j^H \tilde{b}}{\tilde{\sigma}_j} \tilde{v}_j \\ &= g_{TSVD}^{(p)} + \sum_{j=p+1}^{\ell} \frac{\tilde{u}_j^H \tilde{b}}{\tilde{\sigma}_j} \tilde{v}_j, \end{aligned} \quad (24)$$

or from QRD as

$$\tilde{g} = \Pi \begin{bmatrix} \tilde{R}_{11} & \tilde{R}_{12} \\ 0 & \tilde{R}_{22} \end{bmatrix}^{-1} \begin{bmatrix} \tilde{Q}_1^H \tilde{b} \\ \tilde{Q}_2^H \tilde{b} \end{bmatrix}. \quad (25)$$

Because $\tilde{\Sigma}_2$ and \tilde{R}_{22} are both nearly zero under a high SNR condition, slight variations on them will cause significant perturbations in the solution vector \tilde{g} and hence leads to many spurious frequencies in the pseudo-spectrum. Now it becomes clear why one truncates these rank-weakly quantities to remedy these ill-conditions from the view point of numerical stability and also prefilters some stray noise in an attempt to guard against possible contaminations in the pseudo-spectrum.

For many problems, the conservative approach of over-modeling (i.e., $\ell \gg p$) is preferred [1, 4] to taking $\ell \gtrsim p$, since we can later *truncate* some noises that reside in the null space which is orthogonal to the signal space. The advantage of over-modeling is to provide some extra dimensions to *trap* the stray noises and then remove them by truncation. This is an effective way of enhancing the SNR. However, there is always the danger that some signal has been mistakably truncated in low SNR cases where ambiguous changes in the truncated F-norm is possible. On the other hand, spurious frequencies are still very likely to occur when there is insufficient truncation of the rank of the data matrix.

5 SIMULATION RESULTS

Finally, we present various computer simulations based on the following model. Let $\tilde{x}_i = \cos(2\pi f_1 i) + \cos(2\pi f_2 i) + w_i$, $i = 1, 2, \dots, 32$, with $f_1 = .15$, $f_2 = .2$, $\ell = 11$ and w is a white Gaussian random sequence. 100 independent simulations are used for each SNR varying from 0 to 40dB in an increment of 5 dB. The frequencies are determined by the phases of complex roots closest to the unit circle. For TQRR, we pre-permute the columns of the FBLP matrix in the order of: {1 11 6 3 9 4 8 2 5 7 10} as specified by [2]. Fig. (2) gives the average fractional truncated Frobenius norms in (16) versus SNR when we preserve only the four most significant ranks of the FBLP matrix for the four different methods. This confirms their relationship in (19) and also shows that truncated energy decreases monotonically as SNR increases. Fig.(3) and (4) show the averages of the frequency biases for the two harmonic frequencies. We define the average frequency bias as $E(\hat{f}_k) - f_k$, $k = 1, 2$, where $E(\hat{f}_k)$ is the ensemble average of \hat{f}_k , which is an the estimated frequency for f_k . Fig. (5) and (6) show the standard deviations of \hat{f}_1 and \hat{f}_2 . We can see that TQRP competes quite well with TSVD, while TQRR performs slightly worse than TQRP but better than TQR without pivoting.

6 CONCLUSIONS

While a myriad of researches have been focused on SVD and eigen-decomposition analysis of narrowly spaced harmonic frequency estimations [1, 4, 5], very few have been directed towards the QRD approaches. Owing to the iterative massive computations and the difficulty encountered in updating the decompositions [6] when new data are acquired under time-varying conditions, these SVD and eigen-based approaches are not suited for real time applications. It is well known [3] that QRD is numerically as stable as SVD, requires much less computational cost, easy to update(and/or downdate), and amenable to VLSI implementations. The slightly degraded performance for these truncated QR methods is greatly compensated for all the benefits mentioned above.

Table 1 summarizes the comparisons among different truncation methods. We conclude that TQR is the simplest and can be performed easily in a real time updating, but may suffer significant degradation. TQRP provides almost the same performance as SVD, but is not easy to implement in real time processing in that the difficult column reshuffling is required whiling performing QRD with pivoting. TQRR provides a good compromise between these two and can also be implemented for systolic array processing.

	Freq. est.	Comput. cost	VLSI	updating
TSVD	excellent	very high	complex	difficult
TQRP	very good	medium	medium	medium
TQRR	good	fair	fair	easy
TQR	fair	fair	fair	easy

Table 1: Table of comparisons

References

- [1] D. W. Tufts and R. Kumaresan. "Estimation of frequencies of multiple sinusoids: Making linear prediction perform like maximum likelihood". *Proc. IEEE*, 70:975-989, Sept 1982.
- [2] J. P. Reilly, W. G. Chen, and K. M. Wong. "A fast QR-based array-processing algorithm". In *SPIE Vol. 975 Advanced Algo. and Arch. for Signal Proc III*, pages 36-47, 1988.
- [3] G. H. Golub and C. F. Van Loan. "*Matrix computations*". Johns Hopkins University Press, Baltimore, MD, 2 edition, 1989.
- [4] B. D. Rao. "Perturbation analysis of an SVD-based linear prediction method for estimating the frequencies of multiple sinusoids". *IEEE Trans. Acoust., Speech, Signal Processing*, 36(7):1026-1035, July 1988.
- [5] S. L. Marple Jr. "A tutorial overview of modern spectral estimation". In *IEEE ICASSP*, pages 2152-2157, 1989.
- [6] J. R. Bunch and C. P. Nielsen. "Updating the singular value decomposition". *Numerische Mathematik*, 31:111-129, 1978.




Fig. (1) Diagram of the FBLP model.




Fig. (3) Frequency biases for estimating $f_1 = .15$.

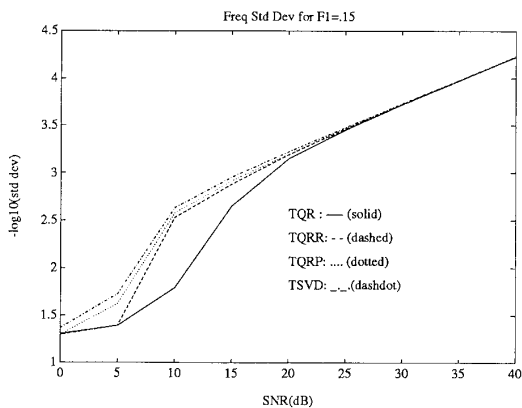


Fig. (5) Standard deviations for estimating $f_1 = .15$.

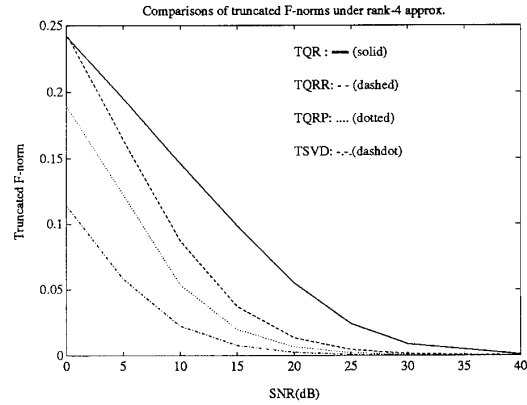


Fig. (2) Average fractional truncated Frobenius norms.

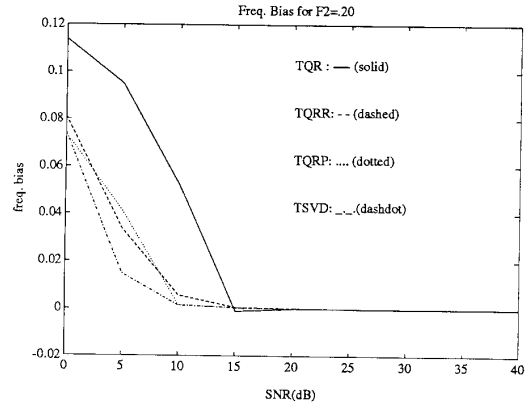


Fig. (4) Frequency biases for estimating $f_2 = .20$.




Fig. (6) Standard deviations for estimating $f_2 = .20$.