Theoretical and Numerical Aspects of an SVD-Based Method for Band-Limiting Finite-Extent Sequences

Søren Hein and Avideh Zakhor

***Abstract*—We present an SVD-based method for band-limiting oversampled discrete-time finite-extent sequences. For this purpose, we show that finite-extent band limitation is best defined in terms of the discrete prolate spheroidal sequences rather than complex exponentials. Our method has maximum energy concentration as defined in the paper, its dimension agrees asymptotically with Slepian's dimension result, and the method specializes correctly to the discrete-time Fourier transform as the sample size tends to infinity. We propose an efficient computational method, based on the Lanczos algorithm, for computing only the necessary singular vectors. The SVD is signal-independent, only needs to be done once and can be precomputed. The SVD-based band limitation itself is not necessarily much slower than the fast Fourier transform for sample sizes on the order of 4096.**

I. Introduction

The problem of band-limiting a finite-extent discrete-time signal is seemingly well understood. The standard solution is to take the discrete Fourier transform (DFT), set the out-of-band coefficients to zero, and take the inverse DFT. However, the inherent periodicity assumption in the underlying DFT-based definition of band limitation is not always satisfactory. For instance, an infinite-extent pure sinusoid may have a band-limited discrete-time Fourier transform (DTFT), whereas a finite set of samples of the same signal may not have a band-limited DFT.

In this paper, we present an alternative approach to band-limiting finite-extent signals which is aimed at alleviating the sensitive frequency dependence of the DFT. Our approach is inspired by existing work on band-limited extrapolation [1]–[4], and we show that finite-extent band limitation is best *defined* in terms of the discrete prolate spheroidal sequences (DPSS's), which have been studied in detail by Slepian [1]. The DPSS's enjoy an optimal energy concentration property [1] which is preserved in our band limitation method.

Although our results are general, their original application was in nonlinear signal reconstruction for $\Sigma\Delta$ modulators [5]. In this problem, only a small or moderate number of samples are available, and the oversampling ratio (OSR) is large, where the OSR is defined as the ratio between the sampling frequency and the Nyquist frequency of the class of considered signals.

The paper is organized as follows. In Section II we propose our method for band limitation. In Section III we describe an efficient numerical method for calculating the required, signal-independent SVD, and we discuss computational complexity. In Section IV we present results to illustrate our method. Section V contains conclusions. A fuller account of the presented work is given in [6].

Manuscript received December 24, 1991; revised May 24, 1993. The associate editor coordinating the review of this paper and approving it for publication was Dr. Barry Sullivan. This work was supported by National Science Foundation PYI award MIP-9057466, ONR young investigator award N00014-92-J-1732, and Analog Devices.

S. Hein was with the Electrical Engineering and Computer Sciences Department, University of California at Berkeley, Berkeley, CA, USA 94720. He is now with Siemens AG, Corporate Research, Munich Germany.

A. Zakhor is with the Electrical Engineering and Computer Sciences Department, University of California at Berkeley, Berkeley, CA USA 04720.

IEEE Log Number 9216669.

II. SVD-Based Band Limitation

In Section II.A we present the mathematical background for our proposed method, drawing on previous work on band-limited extrapolation [1]–[4]. We arrive at an interpretation problem which is resolved in the following sections. In Section II.B we provide an alternative interpretation of the band-limited extrapolation method in [2], [3], and propose our band limitation method. In Section II.C we discuss the underlying problem of defining finite-extent band limitation.

A. Mathematical Background

In this section we describe some simple, but unsuccessful attempts to define the finite-extent band limitation problem. Consider an infinite-extent sequence $\mathbf{x} = \{x_n\}$ and its DTFT

$$X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} x_n e^{-j\omega n}, \quad -\pi < \omega \le \pi.$$

The DTFT is said to be band-limited to the frequency range $\Omega = (-\alpha, +\alpha)$, $0 < \alpha < \pi$, if and only if $X(e^{j\omega}) = 0$ for all $\omega \notin \Omega$. We define the oversampling ratio (OSR) to be π/α. We also define a noisy version $\mathbf{s} = \{s_n\}$ of $\mathbf{x}$. The noise on $\mathbf{s}$ outside of Ω can be rejected with an ideal low-pass filter, or equivalently by multiplying $\mathbf{s}$ by an infinite-dimensional matrix[1] $\mathbf{L} = \{\ell_{nm}\}$ with elements [2], [4]

$$\ell_{mn} = \begin{cases} \frac{\sin \alpha(m-n)}{\pi(m-n)} & \text{for} \quad m \neq n \\ \alpha/\pi & \text{for} \quad m = n \end{cases}, \quad m, n \in \mathcal{Z} = \{\cdots, -1, 0, 1, \cdots\}.$$

We use the subscript Λ to denote the time limitation of a vector to an index set $\Lambda = \{1, 2, \cdots, N\}$. Defining an $N \times \infty$ time limitation matrix $\mathbf{T} = \{t_{mn}\}$ as $t_{mn} = \delta_{mn}$ $(m \in \Lambda, n \in \mathcal{Z})$, where δ_{mn} is the Kronecker delta symbol, we thus have $\mathbf{s}_\Lambda = \mathbf{T}\mathbf{s}$ and $\mathbf{x}_\Lambda = \mathbf{T}\mathbf{x}$.

The simplest version of our band limitation problem is to estimate $\mathbf{x}_\Lambda$ from a set of observed samples $\mathbf{s}_\Lambda$. This can be related to the band-limited extrapolation problem [2], [4] of finding an infinite-extent band-limited sequence $\tilde{\mathbf{x}}$ which passes as closely as possible through $\mathbf{s}_\Lambda$

$$\tilde{\mathbf{x}} = \mathbf{L}\tilde{\mathbf{x}} \text{ and } \|\mathbf{s}_\Lambda - \mathbf{T}\tilde{\mathbf{x}}\|_\Lambda^2 = \|\mathbf{T}(\mathbf{s} - \tilde{\mathbf{x}})\|_\Lambda^2 \text{ is minimized.} \quad (1)$$

In (1), $\|\cdot\|_D$ denotes the 2-norm of a vector over an arbitrary index set D. Having solved the related problem (1), we can consider the time limitation $\tilde{\mathbf{x}}_\Lambda = \mathbf{T}\tilde{\mathbf{x}}$ to be the band limitation of $\mathbf{s}_\Lambda$. However, as might be expected from the sampling theorem,there are in fact infinitely many infinite-extent sequences $\tilde{\mathbf{x}}$ which have band-limited DTFT's and which also pass through any finite set of samples $\mathbf{s}_\Lambda$ [2]. Therefore the solution to (1) is not unique, and the optimal estimate of $\mathbf{x}_\Lambda$ is the trivial solution $\tilde{\mathbf{x}}_\Lambda = \mathbf{s}_\Lambda$.

In the context of band-limited extrapolation, the standard way to obtain a unique solution to (1) is to choose the minimum-energy one [2], [4]. This approach leads to the following least-squares problem

$$\min \|\mathbf{s}_\Lambda - \mathbf{T}\mathbf{L}\tilde{\mathbf{x}}\|_\Lambda^2 \text{ such that } \|\tilde{\mathbf{x}}\|_\mathcal{Z}^2 \text{ is minimized.} \quad (2)$$

It is easily shown that due to the minimization of the energy of $\tilde{\mathbf{x}}$, the solution to (2) is implicitly band-limited[2]. Rephrasing our band limitation problem in terms of (2), we can consider the time limitation $\mathbf{b} \triangleq \mathbf{T}\tilde{\mathbf{x}}$ to be the band-limited estimate of $\mathbf{x}_\Lambda$.

[1] The use of infinite-dimensional matrices is purely a notational convenience.

[2] We thank an anonymous referee for pointing this out.

The solution to (2) can be expressed in terms of the N orthogonal infinite-extent discrete prolate spheroidal sequences (DPSS's) $\{\mathbf{U}_n\}$ [1]. For any finite N, all the N DPSS's depend only on N and the OSR, and have DTFT's that are band-limited to Ω. With suitable normalization, the truncated DPSS's $\{\mathbf{u}_n = \mathbf{T}\mathbf{U}_n\}$ are also orthonormal over Λ [1]. The truncated DPSS's are the left singular vectors in a singular value decomposition (SVD) of the matrix $\mathbf{A} = \mathbf{T}\mathbf{L}$, or equivalently, the eigenvectors of the $N \times N$ matrix $\mathbf{L}_\Lambda = \mathbf{A}\mathbf{A}^T = \mathbf{T}\mathbf{L}\mathbf{T}^T$ [7]. The ordered singular values of $\mathbf{A}$ are denoted by $\{\sigma_1, \cdots, \sigma_N\}$ and satisfy $1 > \sigma_1 > \cdots > \sigma_N > 0$ [1]. $\mathbf{L}_\Lambda$ has full rank. Returning to (2), the band-limited extrapolation solution is known to be [1]

$$\tilde{\mathbf{x}} = \sum_{n=1}^{N} (\mathbf{u}_n^T \mathbf{s}_\Lambda) \cdot \mathbf{U}_n \quad (3)$$

from which we again obtain the trivial band limitation solution $\mathbf{b} = \mathbf{T}\mathbf{L}\tilde{\mathbf{x}} = \mathbf{s}_\Lambda$. The solution is also clear from (2) itself, because the least squares constraint only singles out the extrapolate with minimum energy that passes through the given samples $\mathbf{s}_\Lambda$. Thus, even with the energy minimization, any finite-extent sequence appears to be band-limited. We resolve this central conflict with intuition in the next sections.

B. Dimension Considerations

It has been proposed to replace (3) by the truncated summation [2], [3]

$$\tilde{\mathbf{x}} = \sum_{n=1}^{r} (\mathbf{u}_n^T \mathbf{s}_\Lambda) \cdot \mathbf{U}_n, r \approx N/\text{OSR}. \quad (4)$$

The change is often suggested as a way to improve numerical stability [2], [3]. Specifically, it is shown in [1] that approximately N/OSR of the singular values $\{\sigma_n\}$ are close to 1, and the rest are close to 0. Thus, (4) amounts to discarding the part of $\mathbf{L}_\Lambda$ without full numerical rank.

In this section we make the alternative interpretation that a dimension consideration suggests the use of (4). We base this claim on the fact that the right-hand side of (3) tends to $\mathbf{s}_\Lambda$ as $N \to \infty$, whereas the expected result is $\tilde{\mathbf{x}} = \mathbf{L}\mathbf{s}_\Lambda$. Due to Slepian's dimension result, the discrepancy can be removed by replacing (3) with (4). Slepian's result is summarized as follows in [1]: "For large N the set of sequences of bandwidth W that are confined to an index set of length about N has dimension approximately $2WN$."

In our case, the product $2WN$ equals $N/\text{OSR} = N\alpha/\pi$, so $r = N/\text{OSR}$ orthonormal vectors are asymptotically necessary to span the space of sequences on Λ that are band-limited to Ω. By choosing $r \approx N/\text{OSR}$, we obtain the correct asymptotic specialization to the DTFT, since for $N = \infty$, the matrix $\mathbf{L}_\Lambda$ becomes the ideal low-pass matrix $\mathbf{L}$ which has the complex exponentials $\exp(j\omega n)$ as its eigenvectors. The fraction $1/\text{OSR}$ of the infinite-extent eigenvectors are band-limited to the frequency range Ω, and correspond to the eigenvalue 1 of the infinite-extent matrix $\mathbf{L}$. The remaining eigenvectors contain strictly high-frequency energy, and correspond to the eigenvalue 0. The cost of replacing (3) with (4) is that (3) does not exactly solve the extrapolation problem (2) for finite N.

Our re-interpretation of (4) suggests the use of (4) as the basis of a band limitation method. Equation (4) can be viewed as connecting the well-known problem of band-limited extrapolation to the present problem of band limitation. We propose the band limitation method

$$\mathbf{b} = \mathbf{T}\tilde{\mathbf{x}} = \sum_{n=1}^{r} (\mathbf{u}_n^T \mathbf{s}_\Lambda) \cdot \mathbf{u}_n, r \approx N/\text{OSR} \quad (5)$$

that is, a projection onto the r truncated DPSS's with the largest singular values. Like (4), the band limitation in (5) is not a solution to a dimensionally unconstrained minimization problem such as (1) or (2).

C. Definition of Band Limitation

Equation (5) is a special case of the general linear band limitation method with dimension r

$$\mathbf{b}_{\text{general}} = \sum_{n=1}^{r} (\mathbf{a}_n^T \mathbf{s}_\Lambda) \cdot \mathbf{a}_n \quad (6)$$

where $\{\mathbf{a}_1, \cdots, \mathbf{a}_r\}$ is an arbitrary set of orthonormal N-dimensional "baseband" vectors. DFT-based band limitation is also a special case of (6) with complex exponentials as the baseband vectors. We view (6) as *defining* a class of band-limited finite-extent sequences for each choice of $\{\mathbf{a}_n\}$, namely the span of $\{\mathbf{a}_n\}$.

We now show that the optimal choice of baseband vectors $\{\mathbf{a}_n\}$ in an energy concentration sense is the truncated DPSS's. This result explains the fundamental difference between the DFT and the method (5). The result is based on the fact that among all band-limited infinite-extent sequences, the first DPSS, $\mathbf{U}_0$, is the one with the largest possible fraction, namely σ_0^2, of its energy on the set Λ. The second DPSS, $\mathbf{U}_1$, is the band-limited sequence which is orthogonal to $\mathbf{U}_0$ and has the largest fraction of its energy, namely σ_1^2, on Λ, and so on [1]. We define the energy concentration of the general band-limitation method (6) as a function of $\mathbf{s}_\Lambda$ to be

$$C_{\text{general}}(\mathbf{s}_\Lambda) \triangleq \frac{\|\mathbf{b}_{\text{general}}\|_\Lambda^2}{\|\mathbf{b}_{\text{BLX, general}}\|_\mathcal{Z}^2} \quad (7)$$

where $\mathbf{b}_{\text{BLX, general}}$ is the minimum-energy band-limited extrapolate of $\mathbf{b}_{\text{general}}$ given by (4). If $\mathbf{b}_{\text{BLX, general}}$ is the zero vector, (7) is undefined. Maximizing $C_{\text{general}}(\mathbf{s}_\Lambda)$ is reasonable, since $\|\mathbf{b}_{\text{general}}\|_\Lambda^2$ should contain as much as possible of the energy of the true band limitation solution $\mathbf{s}_\Lambda$ in relation to the energy of the infinite-extent band-limited extrapolate. We show in the appendix that for the general method (6)

$$\min_{\mathbf{s}_\Lambda} C_{\text{general}}(\mathbf{s}_\Lambda) \leq \sigma_r^2. \quad (8)$$

We also show in the appendix that the SVD-based method (5) has optimal energy concentration in that it achieves the upper bound in (8). Thus, choosing the vectors $\{\mathbf{a}_n\}$ to be the truncated DPSS's $\{\mathbf{u}_n\}$ indeed optimizes energy concentration in a maximin sense.

III. COMPUTATIONAL ASPECTS

The SVD of the matrix $\mathbf{L}_\Lambda$ only needs to be done once for each sample size and OSR. The key observations are that: (1) We only need about $r = N/\text{OSR}$ of the eigenvectors of $\mathbf{L}_\Lambda$, and (2) $\mathbf{L}_\Lambda$ is a Toeplitz matrix, so we can multiply a vector by $\mathbf{L}_\Lambda$ in $O(N \log N)$ operations using a $2N$-point fast Fourier transform (FFT).

Observation (1) implies that the Lanczos algorithm is ideally suited for the eigen-decomposition [7], [8]. This algorithm generates a series of tridiagonal matrices, starting with a scalar and ending with an $N \times N$ matrix. The smaller matrices tend to have eigenvalues and eigenvectors that are good approximations to the largest and smallest eigenvalues and corresponding eigenvectors of $\mathbf{L}_\Lambda$. In addition, the algorithm can provide error bounds on its estimates [9]. The Lanczos algorithm requires modest storage, because no storage for intermediate matrices is necessary. In addition, explicit storage of $\mathbf{L}_\Lambda$ is not necessary, as the only requirement is a subroutine that can multiply an arbitrary vector by $\mathbf{L}_\Lambda$. The $N \times N$ matrix $\mathbf{L}_\Lambda$ is completely specified by N numbers, and the time required for multiplications by $\mathbf{L}_\Lambda$ is reduced through observation (2) above.

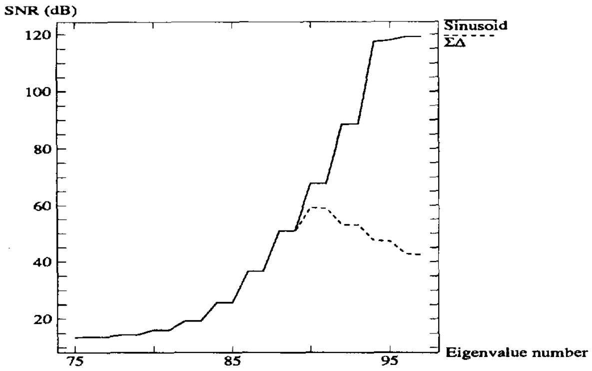

Fig. 1. SNR for the SVD method as a function of eigenvalue number r of the matrix $\mathbf{L}_\Lambda$ for $N = 4096$, OSR = 48. The input signals are a sinusoid and its $\Sigma\Delta$ encoded version. The input frequency is $\omega = \frac{9}{32}\alpha$, and the input amplitude is at −6 dB relative to the quantizer step size of the $\Sigma\Delta$ modulator.

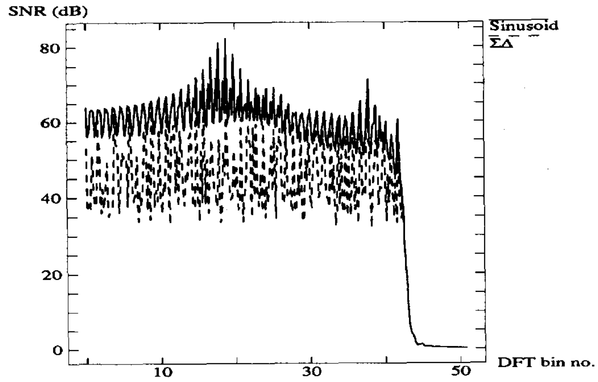

Fig. 2. SNR for the SVD method as a function of input frequency for a fixed number $r = 91$ of eigenvectors, $N = 4096$ and OSR = 48. The input signals are a sinusoid and its $\Sigma\Delta$ encoded binary version. Frequencies are not limited to integer bins. The input amplitude is at 6 dB relative to the quantizer step size of the $\Sigma\Delta$ modulator.

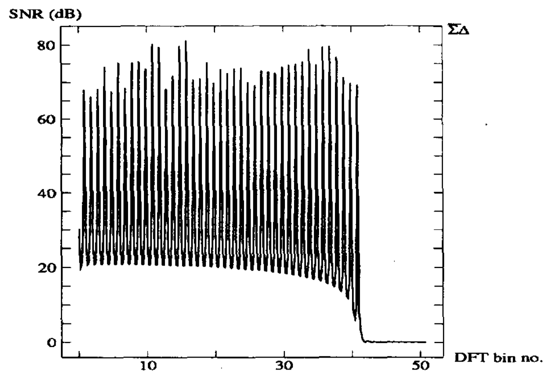

Fig. 3. SNR for the DFT method as a function of input frequency. Parameters are as in Fig. 2, but the curve for sinusoidal input is omitted. The input amplitude is at −6 dB relative to the quantizer step size of the $\Sigma\Delta$ modulator.

Further comments on the Lanczos algorithm can be found in [6]. On a SUN Sparc IPC, a Lanczos-based partial SVD of the $\mathbf{L}_\Lambda$-matrix takes 27 min when $N = 4096$ and OSR = 48. The disk storage requirement is $O(N^2/\mathrm{OSR})$ or 3 Mbytes. For comparison, a full decomposition even in the former case takes days and requires 128 Mbytes of virtual memory. Performing an SVD-based band limitation requires $O(N^2/\mathrm{OSR})$ time, which is only somewhat slower than the FFT-based method for small to moderate N. For $N = 4096$ and OSR = 48, FFT and SVD band limitations take 0.4 s and 1.2 s, respectively, on a SUN Sparc IPC.

IV. RESULTS

This section presents simulation results for the SVD-based and DFT-based methods for band limitation. The sample size is 4096 throughout. We consider two classes of input signals. The first is the class of pure baseband sinusoids that are not necessarily at bin frequencies for the DFT. The second class consists of binary encoded versions of signals in the first class. Our encoder is the fourth order interpolative $\Sigma\Delta$ modulator described in [10] which has an OSR of 48. The noise introduced by the $\Sigma\Delta$ modulator is nonwhite and strongly high-pass. The SNR in approximating a signal $\mathbf{x}$ with the estimate $\hat{\mathbf{x}}$ is defined by $10\log_{10}(\|\mathbf{x}\|_\Lambda^2/\|\mathbf{x}-\hat{\mathbf{x}}\|_\Lambda^2)$.

As shown in [1], about $N/\mathrm{OSR} = 85$ of the eigenvalues of $\mathbf{L}_\Lambda$ are close to 1, and the rest are close to 0. The problem of choosing the number r of eigenvectors in (5) remains, although asymptotically we must have $r \approx N/\mathrm{OSR}$. In the context of band-limited extrapolation, it is well known that r controls a trade-off between the accuracy with which noise-free signals can be represented, and the noise sensitivity [3], [11]. The same observation holds for (5): Larger values of r increase the dimension of the band-limitation subspace, which implies that more of a signal, but also more noise can be represented. A good choice of r thus depends on the expected amount of noise. There exist formulas to aid the choice for band-limited extrapolation under various assumptions [3]. In this paper we will choose the exact value of r empirically.

Fig. 1 shows SNR curves for the SVD method as a function of eigenvector number. Two signals are considered, namely a sinusoid at the DFT bin frequency $\omega = \frac{9}{32}\alpha$ and its $\Sigma\Delta$ encoded binary version. For the very noisy $\Sigma\Delta$ encoded signal, the optimum value of r is 90 or 91. The optimum value of r in general depends on the input frequency.

Fig. 2 shows SNR curves as a function of input frequency for the SVD-based method with a fixed number $r = 91$ of eigenvectors. Again, the input sequences are baseband sinusoids and their $\Sigma\Delta$ encoded versions. Both bin and non-bin frequencies for a corresponding 4096-point DFT are considered. The figure shows that the SNR oscillates between about 58 dB and 65 dB for sinusoidal inputs, and between about 38 dB and 58 dB for $\Sigma\Delta$ encoded inputs.

Fig. 3 is analogous to Fig. 2 and shows SNR curves as a function of input frequency when the band limitation is done with the DFT. The SNR oscillates between 21 dB and infinity for sinusoidal inputs, and between 21 dB and 75 dB for the $\Sigma\Delta$ encoded inputs. The SNR depends strongly on whether or not the frequency is a DFT bin frequency. For both input types, the SVD-based method exhibits less frequency dependence than the DFT, but the largest SNR is smaller for the SVD-based method than for the DFT-based one.

The minimum SNR can be as low as $\mathrm{SNR}_{\min} \approx 10\log_{10}\{N\pi^2/(4\mathrm{OSR})\}$ for the DFT-based method [6]. For $N = 4096$ and OSR = 48, $\mathrm{SNR}_{\min} = 23$ dB which agrees well with 21 dB as observed above. This minimum SNR increases by only 3 dB/octave with the sample size. Due to distortion at the signal frequency, windowing only exacerbates the problem [6].

V. CONCLUSION

In this paper we related the problem of band-limiting finite-extent sequences to that of band-limited extrapolation, and we derived an

SVD-based method for band limitation. In doing so, we redefined finite-extent band limitation. Our method was chosen to have the correct dimension specified by Slepian's result [1], to specialize correctly to the DTFT as the sample size tends to infinity, and to have maximum energy concentration. We presented an efficient method, based on the Lanczos algorithm, for performing the proposed band limitation of sequences with up to 4096 or 8192 samples on present work stations. The SVD-based method is less dependent on input frequency than the DFT, but requires storage for the low-frequency singular vectors.

APPENDIX
ENERGY CONCENTRATION

We first consider the energy concentration of the SVD-based method given by (5), and its corresponding minimum-energy band-limited extrapolation method (4). The energy concentration for the SVD-based method is

$$C_{\rm SVD}(\mathbf{s}_\Lambda) \triangleq \frac{\|\mathbf{b}_{\rm SVD}\|_\Lambda^2}{\|\mathbf{b}_{\rm BLX}\|_\mathcal{Z}^2} = \frac{\sum_{n=1}^{r}(\mathbf{u}_n^T\mathbf{s}_\Lambda)^2}{\sum_{n=1}^{r}(\mathbf{u}_n^T\mathbf{s}_\Lambda)^2/\sigma_n^2}$$

where we have used the fact that $\|\mathbf{U}_n\|_\mathcal{Z}^2 = 1/\sigma_n^2$ and the orthogonality of the the DPSS's $\{\mathbf{U}_n\}$ over $\mathcal{Z}$. As $1 > \sigma_1^2 > \cdots > \sigma_N^2 > 0$, a variational argument shows that $C_{\rm SVD}(\mathbf{s}_\Lambda)$ is minimized by making $\mathbf{s}_\Lambda$ proportional to $\mathbf{u}_r$, in which case (8) is satisfied with equality.

Consider now the general method (6), which can be rewritten as

$$\mathbf{b}_{\rm general} = \sum_{m=1}^{N}\left(\sum_{n=1}^{r}(\mathbf{a}_n^T\mathbf{s}_\Lambda)\cdot(\mathbf{u}_m^T\mathbf{a}_n)\right)\cdot\mathbf{u}_m.$$

By (4), the minimum-energy band-limited extrapolate of $\mathbf{b}_{\rm general}$ is

$$\mathbf{b}_{\rm BLX,\ general} = \sum_{m=1}^{N}\left(\sum_{n=1}^{r}(\mathbf{a}_n^T\mathbf{s}_\Lambda)\cdot(\mathbf{u}_m^T\mathbf{a}_n)\right)\cdot\mathbf{U}_m.$$

The energy concentration (6) as a function of $\mathbf{s}_\Lambda$ is defined in (7). We will show that (8) holds, that is, no band-limitation method of the form (6) has a larger minimum energy concentration than the SVD method. To show this, we choose $\mathbf{s}_\Lambda$ to be an arbitrary nonzero vector $\mathbf{s}_0$ in the intersection $I = {\rm span}(\mathbf{u}_r, \mathbf{u}_{r+1}, \cdots, \mathbf{u}_N) \bigcap {\rm span}(\mathbf{a}_1, \mathbf{a}_2, \cdots, \mathbf{a}_r)$. Since I is defined as the intersection of an $(N-r+1)$-dimensional and an r-dimensional linear subspace of an N-dimensional space, I must have at least dimension 1 and thus must contain nonzero elements. For an arbitrary nonzero $\mathbf{s}_0 \in I$, we then have $\mathbf{b}_{\rm general} = \mathbf{s}_0$ by construction. It also follows by the definition of I that $\mathbf{u}_n^T\mathbf{s}_0 = 0$ for $1 \le n < r$. Therefore

$$\min_{\mathbf{s}_\Lambda} C_{\rm general}(\mathbf{s}_\Lambda) \le C_{\rm general}(\mathbf{s}_0) = \frac{\sum_{n=r}^{N}(\mathbf{u}_n^T\mathbf{s}_0)^2}{\sum_{n=r}^{N}(\mathbf{u}_n^T\mathbf{s}_0)^2/\sigma_n^2} \le \frac{\sum_{n=r}^{N}(\mathbf{u}_n^T\mathbf{s}_0)^2}{\frac{1}{\sigma_r^2}\sum_{n=r}^{N}(\mathbf{u}_n^T\mathbf{s}_0)^2} = \sigma_r^2.$$

ACKNOWLEDGMENT

The authors thank Prof. J. Demmel for suggesting the Lanczos algorithm for eigen-decomposition and Prof. B. N. Parlett for volunteering Lanczos code as well as valuable implementation advice.

REFERENCES

[1] D. Slepian, "Prolate spheroidal wave functions, Fourier analysis, and uncertainly—v: The discrete case," *Bell Syst. Tech. J.*, vol. 57, pp. 1371–1430, May/June 1978.
[2] A. K. Jain and S. Ranganath, "Extrapolation algorithms for discrete signals with application in spectral estimation," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 830–845, Aug. 1981.
[3] B. J. Sullivan and B. Liu, "On the use of singular value decomposition and decimation in discrete-time band-limited signal extrapolation," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-32, pp. 1201–1212, Dec. 1984.
[4] E. J. Diethorn and D. C. Munson, Jr., "A linear time-varying system framework for noniterative discrete-time band-limited signal extrapolation," *IEEE Trans. Signal Processing*, vol. 39, pp. 55–68, Jan. 1991.
[5] S. Hein and A. Zakhor, "Reconstruction of oversampled band-limited signals from Sigma Delta encoded binary sequences," *IEEE Trans. Signal Processing*, vol. 42, no. 4, pp. 799–811, Apr. 1994.
[6] S. Hein and A. Zakhor, *Sigma Delta Modulators: Nonlinear Decoding Algorithms and Stability Analysis.* Norwell, MA: Kluwer, 1993.
[7] G. H. Golub and C. F. V. Loan, *Matrix Computations.* Baltimore, MD: Johns Hopkins University Press, 1989.
[8] B. N. Parlett, *The Symmetric Eigenvalue Problem.* Englewood Cliffs, NJ: Prentice-Hall, 1980.
[9] B. N. Parlett (private communication), Aug. 1991.
[10] K. C. Chao, S. Nadeem, W. L. Lee, and C. G. Sodini, "A higher order topology for interpolative modulators for oversampling A/D converters," *IEEE Trans. Circuits, Syst.*, vol. 37, pp. 309–318, Mar. 1990.
[11] T. S. Huang and P. M. Narenda, "Image restoration by singular value decomposition," *Appl. Opt.*, vol. 14, pp. 2213–2216, Sept. 1975.

Two-fold Normalized Square-Root Schur RLS Adaptive Filter

Peter Strobach, *Senior Member, IEEE*

***Abstract*—A square-root Schur RLS adaptive filter with two-fold (input and residual) normalization is presented. The algorithm has several attractive features such as a fully systolic structure based on elementary hyperbolic plane rotations. All internal variables are bounded in the unit interval and fully utilize it in successive stages due to an inherent " autoscaling " property of the algorithm. This work is presented in a condensed form because it extends the previous work of this author in the area of Schur RLS adaptive filtering.**

I. ARBITRARILY WEIGHTED RLS USING SCHUR RECURSIONS

Recursive least squares (RLS) adaptive filters based on Schur's recursions [1] have been introduced in [2]–[6] as temporal adaptive filtering techniques which allow the incorporation of arbitrarily shaped windows or "forgetting functions" in a true RLS adaptation scheme. In its unnormalized form, a Schur RLS adaptive filter solves, at each time step, the following weighted forward/backward prediction error filtering problem in a true least squares sense:

$$\boldsymbol{e}_m(t) = \boldsymbol{W}^{1/2}\boldsymbol{x}(t) - \boldsymbol{W}^{1/2}\boldsymbol{X}_m(t-1)\boldsymbol{a}_m(t), \quad (1a)$$

$$\boldsymbol{r}_m(t) = \boldsymbol{W}^{1/2}\boldsymbol{x}(t-m) - \boldsymbol{W}^{1/2}\boldsymbol{X}_m(t)\boldsymbol{b}_m(t) \quad (1b)$$

Manuscript received March 25, 1992; revised August 2, 1993. The associate editor coordinating the review of this paper and approving it for publication was Dr. James Zeidler.
The author is with Fachhochschule Furtwangen, Furtwangen, Germany.
IEEE Log Number 9216700.