Filterbank Implementation of Meyer's Wavelet

A Project in partial fulfillment of the requirements for EE392G
submitted by Jon Dattorro

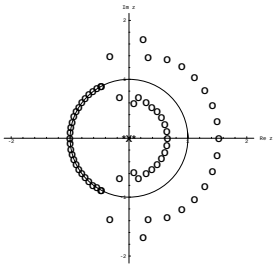

Wavelets
Prof. M. Vetterli
T.A. Tony Verma
Stanford University
June 10, 1998

Introduction

Meyer's wavelet construction is fundamentally a solvent method for solving the two-scale equation. Given a basis φ for the approximation space V_0, Meyer employed Fourier techniques to derive the DTFT of the two-scale equation coefficients, $g_o[n]$, from $\Phi(\omega)$;

$$G_o\left(e^{j\omega}\right) = \sqrt{2}\sum_k \Phi(2\omega + 4k\pi) \qquad \text{[Vetterli, (4.3.4)]}$$

Knowledge of $g_o[n]$ leads naturally to a filterbank interpretation. An implementation becomes attractive when $g_o[n]$ is associated with a rational system function of finite order. But because Meyer's construction has compact support in the frequency domain, there is no rational solution.

We propose approximating the rational transform of $g_o[n]$ for a Meyer-type construction specified to a total mean-square fidelity of greater than 100 dB. We will consider both FIR and IIR realizations. The FIR realizations will be shown to be more efficient than the IIR. The transforms are designed using a construction polynomial $\beta(x)$. The superiority of the FIR implementation is because of the consequent energy compaction in the impulse response which is controlled by the choice of $\beta(x)$. The IIR transforms are then determined from the FIR impulse response $g_o[n]$ via the technique known as Prony's method [Marple].

Generating the Continuous Filter $g_o(t)$

The first step in the process is to construct the continuous synthesis filter prototype which will be sampled to get $g_o[n]$. We do this with reference to [Mallat,pg.247] and [Vetterli,pg.226] from which we determine

$$g_o(t) = \sqrt{2}\left[\frac{\sin\left(\frac{\pi t}{3}\right)}{2t} + \frac{1}{\pi}\int_{\frac{\pi}{3}}^{\frac{2\pi}{3}} \cos\left(\frac{\pi}{2}\,\beta\left(\frac{3\omega}{\pi} - 1\right)\right)\cos(\omega t)\,d\omega\right] \qquad (1)$$

where

$$\beta(x) = x^5\left(126 - 420x + 540x^2 - 315x^3 + 70x^4\right) \qquad (2)$$

This choice of $\beta(x)$ has four derivatives equal to 0 at each endpoint of the range [0,1]. Equ.(2) provides more energy compaction of the impulse response than [Vetterli,(4.3.2)]. We observed empirically that going to polynomials of higher order than (2) did not offer significantly more compaction.

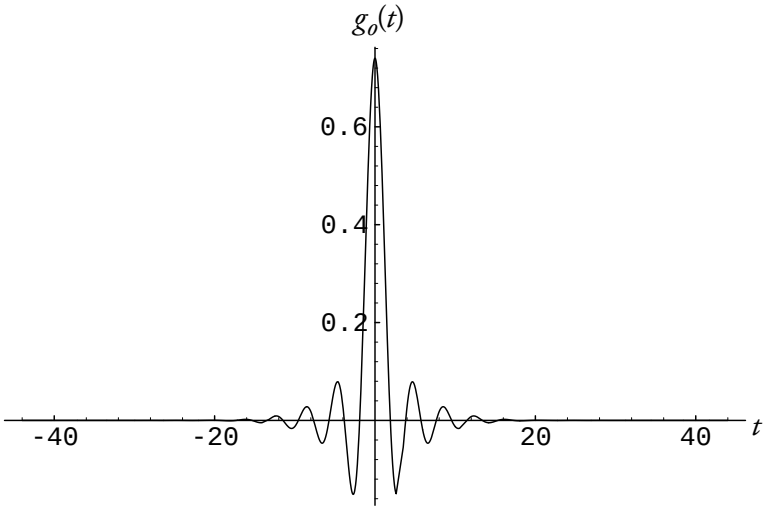

Figure 1. The continuous prototype synthesis filter.

Figure 1 shows only the significant portion of the infinite impulse response $g_o(t)$ of our selected prototype synthesis filter, while Figure 1A shows its Fourier transform.

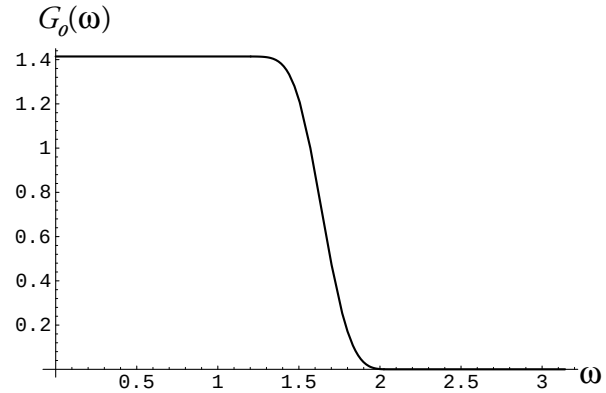

Figure 1A. Ideal response corresponding to Equ.(1).

We wish to sample $g_o(t)$ and use it in a discrete wavelet transform (DWT) implemented via filter bank. The key to the success of this proposed technique is to sample $g_o(t)$ using a shah function advanced by 1/2 sample.

$$\sum_{n=-K+1}^{K} g_o[n]\,\delta\left(t - nT + \frac{T}{2}\right) = g_o(t) \sum_{n=-K+1}^{K} \delta\left(t - nT + \frac{T}{2}\right) \qquad ;\ T = 1 \qquad \text{(2A)}$$

We call this *asynchronous sampling*. Then choosing an even number of even symmetric samples forces a zero at π in the frequency response $G_o(e^{j\omega})$. [O&S,pg.265] That zero is demanded by Rioul [Vetterli,pg.251] as a necessary condition for convergence of the iteration [Vetterli,(4.4.9)].

The impact of sampling in the manner of Equ.(2A) is to place a delay into the two-scale equation for the scaling function. That is, instead of [Vetterli,(4.2.9)], we now have

$$\Phi(2\omega)\, e^{j\omega/2} \;=\; \frac{1}{\sqrt{2}} G_o\left(e^{j\omega}\right)\Phi(\omega) \qquad (3)$$

as our two-scale equation.

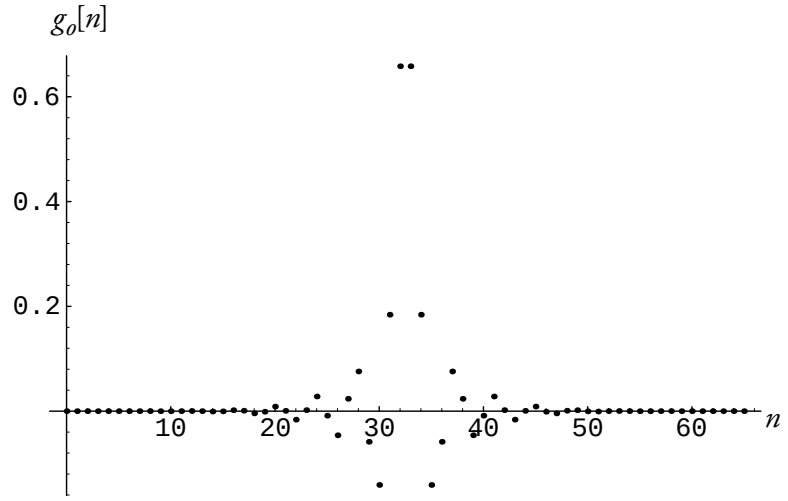

Figure 2. The sampled synthesis filter having even indexed samples discarded.

Figure 2 shows the result of sampling $g_o(t)$ using the asynchronous sampling scheme described above. The many samples of $g_o[n]$ which are close to zero attest to the energy compaction. This sampled signal contains only 66 significant samples and represents 99.99999999051% of the total energy in $g_o(t)$. Stated in deciBels, the sampled signal in Figure 2 represents $g_o(t)$ [Luthra] to within a total mean-square error of -100.227 dB. That particular error is completely due to support truncation of the infinite impulse response.

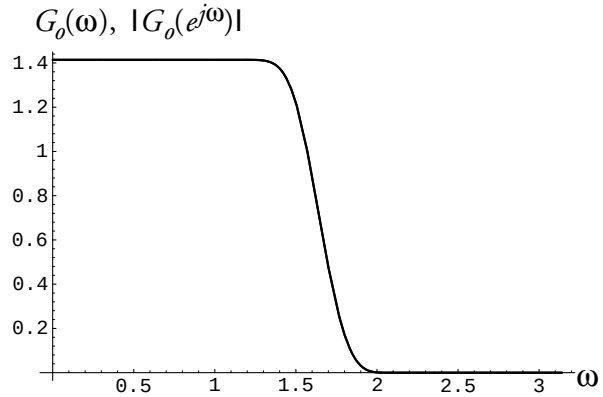

Figure 2A. Overlaid frequency response of the ideal continuous-time prototype and its discrete-time approximation.

We demonstrate in Figure 2A that on an absolute scale, the discrete-time approximation to the synthesis filter prototype is indistinguishable from its continuous-time counterpart. Of course on a dB scale, the discrete-time approximation is seen to be non-ideal but it does have an absolute zero at π, shown in Figure 2B.

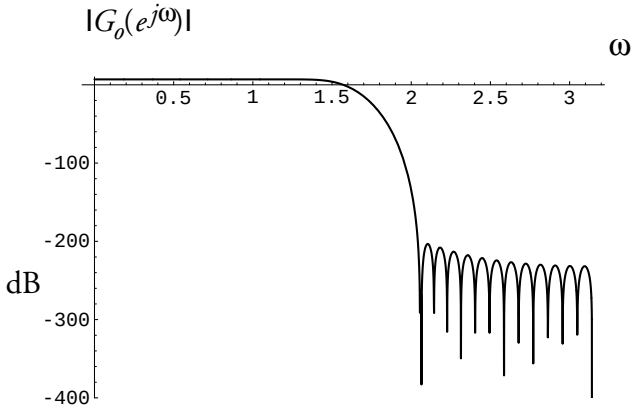

Figure 2B. Actual frequency response of the discrete-time synthesis filter.

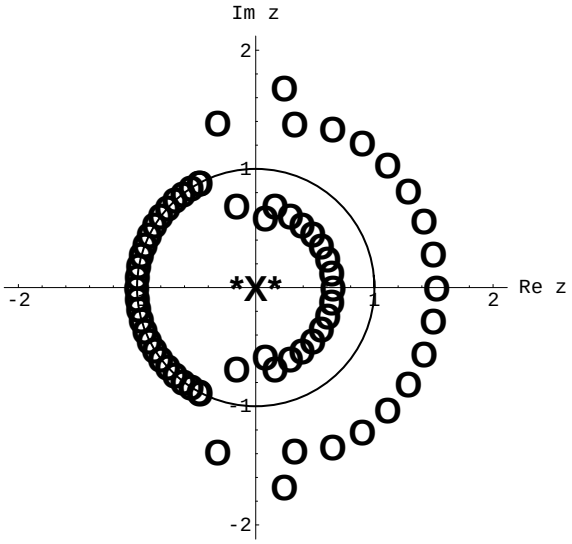

Figure 2C. Pole/zero constellation of the discrete-time synthesis filter.

Figure 2C shows the zero locations for the discrete-time synthesis filter. The zero at π is clearly present.

Expected Shape of the Wavelet

When the filter bank is iterated in an octave-band tree as in [Vetterli,Fig.4.14,pg.238], the two most remote branches have impulse responses that respectively converge (under the proper conditions) to the scaling function and the wavelet. Referring to [Vetterli,Fig.4.10,pg.228], due to the asynchronous sampling scheme we see that the wavelet spectrum no longer has even symmetry. Since the wavelet spectrum now shows odd symmetry[1] as in Figure 3, we therefore expect the wavelet to have odd symmetry as well.

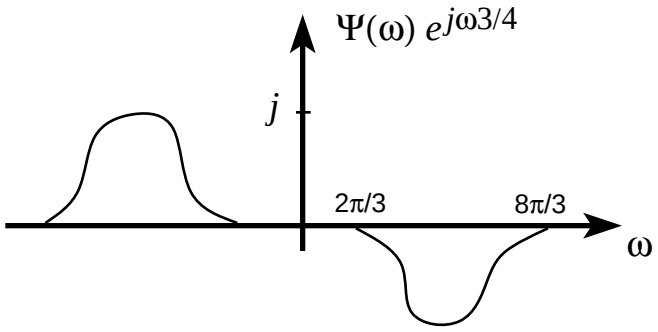

Figure 3. The symmetry of the wavelet spectrum is now odd.

Does the Iterated Filter Bank Converge?

We tested convergence to the scaling function out to 8 iterations as in [Vetterli,(4.4.9),pg.243]. Figure 4 shows the result of the iteration.

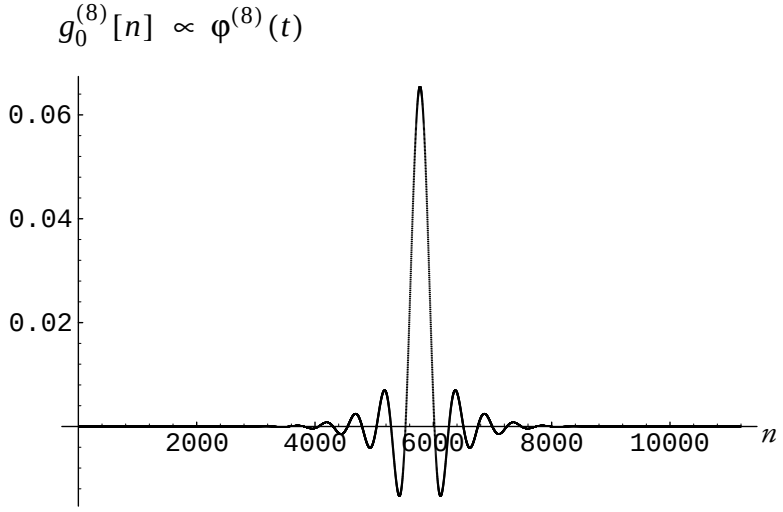

Figure 4. The result of 8 iterations on $G_o(e^{j\omega})$. Note that this is a rendering of a discrete plot; i.e., there are no lines drawn.

[1]We do not derive this result here, but we believe that the sign sense in correct. See 'Glory' n.b. pg.66.

The result in Figure 4 bears a remarkable resemblance to Figure 1 to within a constant factor which we ignored. [Vetterli,(4.4.10),ibid.] Using fewer iterations, we constructed the wavelet branch of the filter bank via the following substitution, as in [Vetterli,(3.2.52)],

$$G_1(z) = z^{-1}\, G_o^*\left(-z^{-1^*}\right) \qquad (4)$$

The leading sign reversal, with respect to [Vetterli,ibid.], is dictated by the sense of Figure 3.

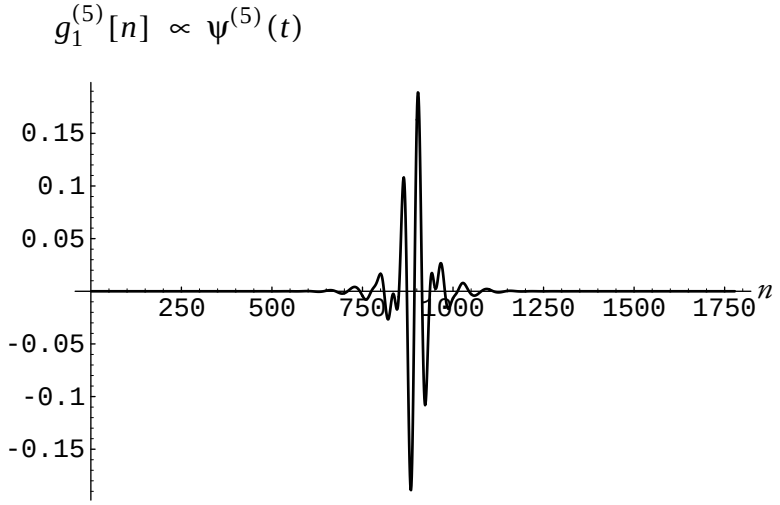

Figure 5. The result of $G_o^{(4)}(e^{j\omega}) \cdot G_1(e^{j\omega\, 2^4})$.

We presume that in some sense $g_1^{(\infty)}[n]$ converges to the wavelet; e.g., as in the sense of [Vetterli,Sec.4.4.2] or [Vetterli,Sec.4.5.3]. The odd symmetry of the wavelet is in accordance with its predicted Fourier transform as sketched in Figure 3.

Figure 6 shows the DTFT of the wavelet in Figure 5. Figure 6 conforms to the approximate response expected for the fifth iteration [Vetterli,pg.240].

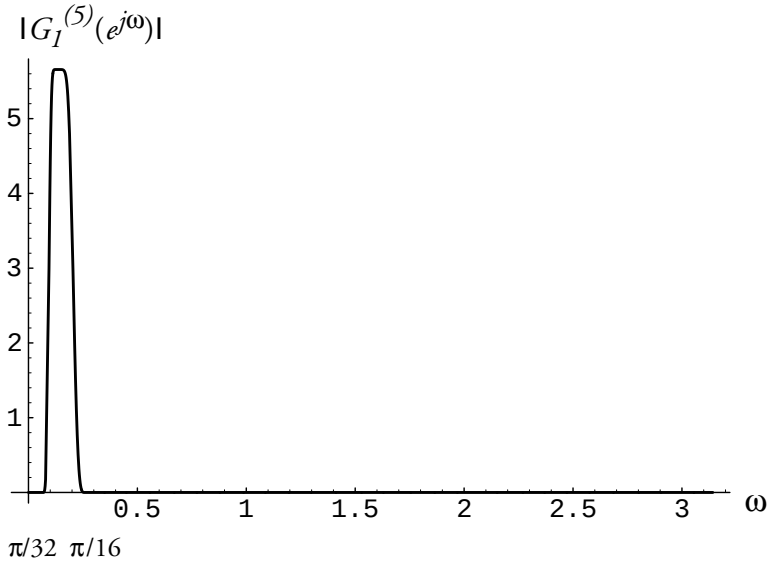

Figure 6. Frequency response of remote highpass branch of octave-band tree.

Implementation and the Virtual Pole

Given that each FIR filter has only 66 taps, the required computation is not excessive.[2] Even so, it might be advantageous in some circumstance to implement each FIR as a truncated IIR. [Wang] The noncausal impulse response may be divided into its causal and anticausal parts; each part being implemented separately by a causal circuit using an appropriate delay. Unlike the FIR implementation, finite word-length in the signal paths consequently becomes an issue because the "anticausal" IIR filters have exponentially growing, albeit truncated, impulse response.

To implement the truncated noncausal impulse response $g_o[n]$, we introduce the concept of the *virtual pole* applied separately to the anticausal and causal part, each of length K=33 . A virtual pole is nothing more than a transfer of the form

$$\sum_{n=0}^{K-1} a_i^n \, z^{-n} \;=\; \frac{1 - a_i^K \, z^{-K}}{1 - a_i \, z^{-1}} \qquad ;\; 0 \le |a_i| < 1 \tag{5}$$

Equ.(5) represents the non-delayed i^{th} exponential component of the causal part of the noncausal response. The delayed anticausal part would be represented by a conjugate-reciprocal transfer of the form

$$a_i^{*K-1} \sum_{n=0}^{K-1} a_i^{*-n} \, z^{-n} \;=\; a_i^{*K-1} \, \frac{1 - a_i^{*-K} \, z^{-K}}{1 - a_i^{*-1} \, z^{-1}} \qquad ;\; 1 \le \left|a_i^{-1}\right| < \infty \tag{6}$$

[2]Indeed, the author has put a 4096-tap filter into commercial production. [Dattorro] [Andreas] On the other hand, very small scale analysis can be achieved only at high sample rates when using 66 filter taps.

which is theoretically stable but presents some practical considerations in a finite-precision machine.

The role of the numerator in both (5) and (6) is to truncate the impulse response after K samples have passed. It does so by cancelling the state of the recursive memory element. Implementation in this manner forces the length of the complete noncausal filter to be 2K which insures that there always exists a zero at π, thus encouraging convergence to a Meyer-type scaling function and wavelet in an iterated filter bank.

We employ Prony's method [Marple], implemented using singular value decomposition (SVD) [Numerical] [Laplace] (see Appendices for C programs), to determine the pole locations (the a_i) of the sampled prototype IIR ($g_o(t)$,Figure 1) from a support-truncated version ($g_o[n]$,Figure 2) which is compact in energy spread. Prony's method is essentially a Laplace analysis over the entire z-plane. The poles found are those of the infinite impulse response. As previously discussed, support truncation to 2K=66 samples introduces an error in fidelity at about 100 dB below the total energy of the ideal prototype. SVD is insensitive to support-truncation error.

Table 1 lists only the positive-frequency IIR poles found from the right hand side of the even-symmetric impulse response shown in Figure 2. Only the right hand side of $g_o[n]$ is input to the Prony's method analysis. As SVD is highly sensitive to noise in individual sequence samples, this particular application is perfectly suited to the SVD because of the high 64-bit precision provided at the SVD algorithm input.

Table 1. Poles of $g_o(t)$. Complex conjugates and conjugate reciprocals not listed.

Damping d	**Frequency** θ [normalized radians]	**Amplitude** A_{dB} [dB]	**Phase** ϕ [rads]
0.8453486430642601	1.9816725942614788	-56.6452788169135600	2.1166355292310737
0.7335858735633031	1.8045568547009641	-26.6539233689077230	-2.5274509377016940
0.7211042720568758	1.5673204811852091	-16.7767035262900560	0.0161358357478311
0.4873572404090880	1.4977702191245537	-18.9146017435699920	-1.0571429978160689
0.1405612015576965	1.3147965577602470	-14.1388037100154930	-0.4834711319434999
0.7734455289939324	1.3330424477404466	-43.3808363913245220	-2.9113879721102003

The data in Table 1 can be used to reconstruct the support-truncated signal as follows:

$$\left.\begin{aligned} t[n] &= \sum_i A_i\, e^{j\phi_i}\, a_i^n\, u[n] \\ &= \sum_i A_i\, d_i^n\, e^{j(\theta_i n + \phi_i)}\, u[n] \end{aligned}\right\} \quad ;\ n < K \qquad (7)$$

where the a_i are the poles, previously discussed in (5) and (6), appearing in Table 1 in their polar form as damping (pole radius) and frequency (pole angle). The Prony's method analysis additionally provides the complex gain associated with each pole, in the form $A_i\, e^{j\phi_i}$. To construct $g_o[n]$ from $t[n]$ we must form

$$g_o[n] = \left(t[n-K] + t^*[K-1-n]\right) + \left(t^*[n-K] + t[K-1-n]\right) \qquad (8)$$

Implementation Accuracy

Regarding only the quantization errors in the individual coefficients of the sequence $g_o[n]$, the Prony's method synthesis represented by the parameters in Table 1 is a faithful reproduction to the support-truncated sampled prototype filter to within a mean-square error of -121.754146 dB. The total error in fidelity, including the support truncation and quantization, remains less than -100 dB.

Of course, the FIR implementation introduces no coefficient quantization error, so its total error in fidelity is determined by support truncation.

Computational Intensity

Through the combination of complex conjugate pole pairs, each pair is implemented in second order sections which require 6 multiplications due to the support truncation circuitry. From Equ.(8) we count 12 pairs that require implementing. Hence the computational intensity in terms of multiplications is 6x12=72.

Comparing this to the 66 multiplications required for the FIR implementation, we conclude that it is cheaper to stick with the FIR implementation.

Conclusions

We sampled a continuous prototype synthesis filter $g_o(t)$ using a shah function shifted by one-half sample. By forcing an even number of samples in the even-symmetric FIR $g_o[n]$, we insured one necessary condition for convergence of the iterated filter bank. That led to a modification of the Meyer-construction which produced an antisymmetric wavelet.

We have determined that the most efficient implementation option is an FIR approach. This came about because of the high energy-compaction of the impulse response that we achieved via $\beta(x)$ (Equ.(2)). It stands to reason that further research may find other polynomials that optimize the compaction.

References

[Mallat] Stephane Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998

[Vetterli] Martin Vetterli, *Wavelets and Subband Coding*, Prentice-Hall, 1995

[O&S] A.V. Oppenheim, R.W. Schafer, *Discrete-Time Signal Processing*, Prentice-Hall, 1989

[Dattorro] Jon Dattorro, 'The Implementation of Digital Filters for High Fidelity Audio, Part II - FIR', *Audio in Digital Times*, The Proceedings of the Audio Engineering Society 7th International Conference, Toronto, pp.168-180, 1989 May 14-17

[Andreas] Jon C. Dattorro, Albert J. Charpentier, David C. Andreas, 'Decimation Filter as for a Sigma-Delta Analog-to-Digital Converter', United States Patent No.5,027,306 June 25, 1991

[Luthra] A. Luthra, 'Extension of Parseval's Relation to Nonuniform Sampling', IEEE Transactions on Acoustics Speech & Signal Processing, vol.36, no.12, 1988 December

[Wang] Avery Li-Chun Wang, 'On Fast FIR Filters Implemented As Tail-Canceling IIR Filters', Stanford, Calif. : CCRMA, Dept. of Music, Stanford University, 1994
LOCATION: Music ML63 .S785 NO.90 (Library has c.1)

[Numerical] Press et al., *Numerical Recipes in C*, Cambridge, 1991

[Marple] S. Lawrence Marple, Jr., *Digital Spectral Analysis with Applications*, Prentice-Hall, 1987

<u>Three papers on the topic of Laplace analysis:</u>

[Laplace]

Friedman, IEEE Transactions on Acoustics Speech & Signal Processing, August 1981, pg.923

Betser, IEEE Transactions on Acoustics Speech & Signal Processing, November 1991, pg.2540

Beliczynski, IEEE Transactions on Acoustics Speech & Signal Processing, March 1992, pg.532