Signal Processing 24 (1991) 117-126
Elsevier

Interpolating and Nyquist filters with constraints at certain frequencies*

Peter Kabal
Department of Electrical Engineering, McGill University, 3480 University, Montreal, Quebec, Canada H3A 2A7, and INRS-Telecommunications, Université du Quebec, 3 Place du Commerce, Verdun, Quebec, Canada H3E 1H6

Eric Dubois
INRS-Telecommunications, Université du Québec, 3 Place du Commerce, Verdun, Quebec, Canada H3E 1H6

Received 6 November 1989
Revised 26 November 1990

Abstract. This paper presents modifications to procedures for designing FIR interpolating filters and FIR Nyquist filters that allow the output of the filter to be error-free for certain input frequencies. For a minimum mean-square error interpolator, the modifications result in an expanded set of linear equations which include the constraints. For minimax stop-band Nyquist filters, the modifications involve factoring out a filter which implements the constraints on the transfer function. Approximate techniques for implementing the constraints are also discussed.

Zusammenfassung. Dieser Beitrag stellt Modifikationen von Entwurfsverfahren für interpolierende FIR Filter und Nyquist FIR Filter die das Ausgangssignal des Filters bei gewissen Eingangsfrequenzen fehlerfrei machen vor. Für einen interpolator nach dem mittleren Fehlerquadratskriterium führen diese Modifikationen auf ein erweitertes System linearer Gleichungen, die die Einschränkungen enthalten. Für Nyquistfilter mit im tschebyscheffschen Sinne entworfenem Sperrband führen die Modifikationen auf einen Filterfaktor, der die Einschränkungen auf die Übertragungsfunktion verwirklicht. Es werden approximative Verfahren zur Verwirklichung der Einschränkungen diskutiert.

Résumé. Cet article présente des modifications à des procédures connues pour la conception de filtres d'interpolation RIF et de filtres de Nyquist RIF, afin d'inclure des contraintes assurant la reconstruction parfaite pour certaines fréquences. Pour un interpolateur minimisant l'erreur quadratique moyenne, les modifications impliquent la résolution d'un système d'équations linéaires, augmenté pour inclure les contraintes. Pour des filtres Nyquist avec bande d'arrêt 'minimax', la procédure modifiée impose un facteur qui réalise les contraintes à la fonction de transfert. Des techniques approximatives pour réaliser les contraintes sont également discutées.

Keywords. Interpolating filter, Nyquist filter, filter design.

1. Introduction

Interpolating filters and Nyquist filters are commonly used for sampling rate conversion operations. In this paper we discuss some problems associated with the traditional methods used to design finite impulse response (FIR) interpolating and Nyquist filters. Specifically, the residual aliasing at some frequencies (such as dc) can render the filters unsuitable for practical applications. We show how the design procedures can be modified to assure error-free reproduction for inputs at certain frequencies.

The work for this paper was motivated by the observation that conventionally designed interpolating filters can perform badly in certain applica-

* This work was supported by the Natural Sciences and Engineering Research Council of Canada.

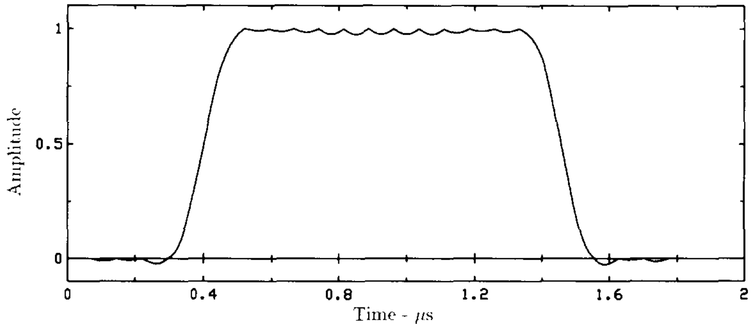

Fig. 1. Response of an unconstrained interpolation filter to an input T-step pulse (125 ns rise and fall times) ($N = 35$, 279 coefficients, sampling rate 472.5 MHz, designed for a raised-cosine power spectrum with transition between 5 MHz and 6 MHz).

tions. A particular application involved an interpolating filter for video transcoding (converting between a colour/luminance representation and a composite representation with different sampling frequencies). In constant parts of the image, noticeable periodic components were inserted by the interpolating filter. The cause was traced to the fact that the interpolating filter did not interpolate a constant input to a constant value. Figure 1 shows the response of such a filter to an input pulse (design details are given later). Note the ripple at the top of the pulse response. As another example, Nyquist filters are required in certain NTSC compatible extended definition television transmission systems [1]. Similar problems with a constant input are encountered for this application.

Several interpretations of the dc interpolation issue are possible. In an interpolator, subsampled filters are used in round-robin fashion. If the responses of these subfilters to a constant input are not the same, the resultant interpolated values will have a periodic variation. From a frequency domain point of view, the aliased dc components appear at frequencies $2\pi l/N$, where N is the interpolation factor. If the response of the interpolating filter to the aliased components is not exactly zero, the resulting interpolated values will not be constant. In fact, the stopband lobes of the interpolation filter may peak in the neighborhood of some of the aliased frequency values, exacerbating the effect.

This suggests that the design procedures for interpolation and Nyquist filters should be modified to allow for control of the interpolated dc values. More generally, exact reconstruction can be stipulated at any given frequency. For instance in a video application, control of the response at the colour subcarrier frequency will ensure that no unwanted amplitude modulation of the subcarrier occurs.

In the following we consider constraining the response at one or more frequencies. Two generic filter design strategies are examined. For interpolation filters, the technique described by Oetken et al. [5] minimizes the mean-square interpolation error for a signal with a given correlation function (or equivalently given power spectral density). For Nyquist filters, the technique described by Ramachandran and Kabal [6] designs Nyquist filters with a minimax stopband. Modifications to these design techniques to contrain the response at certain frequencies are developed. Examples are given for filters designed with and without frequency constraints.

2. Interpolating and Nyquist filters

Interpolating filters and Nyquist filters perform similar functions. For an interpolating filter, the primary concern is accurate interpolation. For the design procedure that we consider, this means minimizing the mean-square interpolation error for an input signal of given power spectral density. The corresponding frequency response of the interpolation filter will show characteristic peaks in the stopband response at frequencies corresponding to regions in which the aliased input frequencies have little or no energy (see the responses in [5]).

Nyquist filters are used in data transmission systems to prevent intersymbol interference. For these filters, the primary concern is the suppression of out-of-band energy, while ensuring regularly space zero crossings in the impulse response. The Nyquist filters considered in this paper have a minimax stopband behavior. For a well attenuated stopband, the passband will be flat and the transition region will have the appropriate symmetry to guarantee the zero crossing property.

For an input sequence $s(n)$ and a filter response $h(n)$, the output can be written in terms of a convolution,

$$y(n) = \sum_k h(k) s(n-k). \tag{1}$$

In this paper, FIR filters with real coefficients will be considered. In that case, the summation is over the non-zero elements of $h(\cdot)$. In addition, the filter will be considered to be non-causal to simplify coefficient indexing.

Interpolation and Nyquist filters ideally satisfy the property that every Nth sample of the output sequence reproduces the input data values. The requirement that $y(lN) = s(lN)$ for general $s(\cdot)$ implies that

$$h(lN) = \begin{cases} 1 & \text{for } l = 0, \\ 0 & \text{for } l \neq 0. \end{cases} \tag{2}$$

This is the zero crossing property referred to earlier. This relationship also indicates that the range of coefficient indices must include zero.[1] Normally the input to the filter is a sequence with $N-1$ zeros interposed between sample values. For interpolation, the sequence is a rate-increased version of the sequence to be interpolated. For Nyquist filters, the sequence is the rate-increased version of the data sequence to be transmitted. In either case, n can be written as $mN+q$ with $0 \leq q \leq N-1$ and the convolution sum can be expressed as

$$y(mN+q) = \sum_i h_q(i) s((m-i)N). \tag{3}$$

The coefficient $h_q(i)$ is an element of the qth subfilter $h(iN+q)$. When continuously processing input samples, the subfilters are used in round-robin fashion.

The frequency response of an interpolating filter can be written in polyphase form,

$$\begin{aligned} H(e^{j\omega}) &= \sum_{q=0}^{N-1} \sum_i h_q(i)\, e^{-j\omega iN}\, e^{-j\omega q} \\ &= \sum_{q=0}^{N-1} H_q(e^{j\omega N})\, e^{-j\omega q}, \end{aligned} \tag{4}$$

where

$$H_q(z) \triangleq \sum_i h_q(i) z^{-i}. \tag{5}$$

The constraint that a sinusoid at frequency[2] ω_c be reproduced with no distortion is equivalent to the condition that the frequency response of subfilter q satisfies

$$H_q(e^{j\omega_c N}) = e^{j\omega_c q}, \quad 0 \leq q \leq N-1. \tag{6}$$

This result can be obtained directly from (3) by requiring that $y(n) = s(n) = \cos(\omega_c n + \theta)$ for arbitrary θ.

The result in (6) means that the frequency response $H_q(e^{j\omega_c N})$ must correspond to a shift of q samples. Note that the right-hand side has a magnitude of one. This is necessary to allow the $q=0$ case to be consistent with the interpolation property (2). In addition, note that if the response satisfies (6) at frequency ω_c, the response also satisfies this condition at frequency $-\omega_c$.

Using the constraint given by (6), it can be shown from the polyphase form of the frequency response

[1] For Nyquist filters, the zeroth sample is often set to N instead of 1. This gives a passband gain of one.

[2] Frequencies are normalized with respect to the sampling frequency of the interpolated sequence.

that

$$H(e^{j(\pm\omega_c+2\pi l/N)}) = \begin{cases} N, & l = mN, \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

This indicates that the z-transform $H(z)$ has roots on the unit circle. For general ω_c, the factor containing these unit circle roots can be written as

$$\begin{aligned} H_c(z) &= \prod_{l=1}^{N-1} (z - e^{j(\omega_c+2\pi l/N)})(z - e^{-j(\omega_c+2\pi l/N)}) \\ &= \frac{(z^N - e^{j\omega_c N})(z^N - e^{-j\omega_c N})}{(z - e^{j\omega_c})(z - e^{-j\omega_c})}. \end{aligned} \tag{8}$$

This filter serves to suppress aliased versions of the sinusoid. For ω_c of the form $\pi m/N$, the root factor differs from the general form since the constraints of (6) become real for some or all of the subfilters. Specifically, for a dc constraint ($m = 0$), the root factor becomes

$$\begin{aligned} H_c(z) &= \prod_{l=1}^{N-1} (z - e^{j2\pi l/N}) \\ &= \frac{z^N - 1}{z - 1}. \end{aligned} \tag{9}$$

3. Interpolating filter design

The design of unconstrained minimum mean-square error interpolating filters is discussed in Appendix A. For each frequency ω_c, the complex constraint of (6) can be expressed as a pair of real constraints which need to be considered for each subfilter. Consider the case of a single frequency. The constraints can be written as

$$\boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{c} = \cos(\omega_c q) \quad \text{and} \quad \boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{s} = -\sin(\omega_c q), \tag{10}$$

where the elements of $\boldsymbol{h}_q^{\mathrm{T}}$ are $h_q(i)$ and the elements of $\boldsymbol{c}$ and $\boldsymbol{s}$ are $c_i = \cos(\omega_c iN)$ and $s_i = \sin(\omega_c iN)$, respectively. For ω_c of the form $\pi m/N$, one of the pair of real constraints is automatically satisfied for some or all values of q. These real constraints for the general case can be incorporated into the design procedure using Lagrange multipliers. The mean-square error in (A.4) is augmented with two terms:

$$\varepsilon_q = r_{xx}(0) - 2\boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{\alpha}_q + \boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{R}\boldsymbol{h}_q + 2\lambda_c\boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{c} + 2\lambda_s\boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{s}. \tag{11}$$

The constraints on the response can be combined with the error-minimization equations to give the following linear equations for the constrained minimization:

$$\left[\begin{array}{c|cc} \boldsymbol{R} & \boldsymbol{c} & \boldsymbol{s} \\ \hline \boldsymbol{c}^{\mathrm{T}} & 0 & 0 \\ \boldsymbol{s}^{\mathrm{T}} & 0 & 0 \end{array}\right] \left[\begin{array}{c} \boldsymbol{h}_q \\ \hline \lambda_c \\ \lambda_s \end{array}\right] = \left[\begin{array}{c} \boldsymbol{\alpha}_q \\ \hline \cos(\omega_c q) \\ -\sin(\omega_c q) \end{array}\right]. \tag{12}$$

In contrast to the unconstrained optimization, the system of equations is no longer Toeplitz. This means that some efficiency of solution must be sacrificed to enforce the constraints. The changes to include additional constraint frequencies or a constraint frequency of the form $\pi m/N$ are straight-forward.

4. Nyquist filter design

The linear phase FIR Nyquist filter design procedure of [6] can be modified to allow exact reproduction of certain frequencies. The focus of that paper is the design of factorable filters. That aspect will be stressed less here, though the modified Nyquist filter can be forced to be factorable and the factorization techniques discussed in [6] still apply.

The constraint of a correct response at frequency ω_c (assumed to be in the passband for clarity of exposition) requires that the responses at the stopband frequencies $2\pi l/N \pm \omega_c$, $l = 1, \ldots, N-1$ be exactly zero. This can be accomplished by forcing the overall Nyquist filter to have a factor with zeros at those frequencies. This factor is the filter $H_c(z)$ of (8).

The heart of the procedure in [6] involves decomposing the Nyquist filter into two parts – one controlling the stopband and one controlling the passband. Interaction between the passband

and stopband filters requires an iteration of the designs, i.e., the stopband filter is redesigned taking into account the response of the passband filter.

The modified procedure to include the effect of a constrained response at frequency ω_c can be summarized as follows. Let the passband be the frequency range $[0, \omega_p]$. The stopband is then the frequency range $[2\pi/N - \omega_p, \pi]$. Let the passband linear phase filter be $H_0(z)$ and let the stopband linear phase filter be $H_1(z)$. The overall filter is $H_0(z)H_c(z)H_1(z)$. The weighting function for the stopband region is $W(\omega)$.

0. Initialization. Let $H_0(z) = 1$.
1. Stopband design. The desired response of the filter $H_1(z)$ has the value one at $\omega = 0$ and is zero in the stopband. The stopband weighting is

$$W'(\omega) = W(\omega)|H_0(e^{j\omega})H_c(e^{j\omega})|.$$

   The Remez exchange algorithm operating on a discretized frequency grid can be used to design a stopband filter which approximates the desired response [2].
2. Passband design. $H_0(z)$ is designed to force zero crossings into the impulse response of the overall filter. This requires the solution of a set of linear equations to determine the coefficients of $H_0(z)$. The linear equations involve the impulse response coefficients of the combined filter $H_c(z)H_1(z)$.
3. Iterate the design, repeating steps 1 and 2 until the responses no longer change significantly.

For certain combinations of filter length and interpolating ratio, $H_1(z)$ has an even number of coefficients. In that case, since $H_1(z)$ is linear phase, there is a null in the frequency response at $\omega = \pi$. This null obviates the need for $H_c(z)$ to include a zero at $z = -1$.

Factorability of the Nyquist filter can be assured by using $H_1^2(z)$ as the stopband filter (with appropriate modification of the weighting function) and using $H_c^2(z)$ as the constraint filter. In the factorable design, the factor $H_c(z)$ appears in both the minimum phase and the maximum phase parts of the filter. The presence of the term $H_c^2(z)$ can reduce the number of coefficients available for the design of the passband and stopband filters significantly. However, in Nyquist filter applications, the stopband suppression requirement usually applies to each of the minimum and maximum phase factors individually. This means that the number of coefficients for each of these factors is sufficiently large to allow for the inclusion of at least some frequency constraints.

Although the focus of this paper is the constraining of the response for certain input frequencies, similar principles can be applied to force a null into the response of the filter at some particular stopband frequency. In data transmission applications, such a null can be used for pilot tone insertion.

5. Approximate solutions

The discussion above has involved exact constraints on the response for input sinusoids. In this section, alternate approximate solutions are investigated.

5.1. Interpolating filter

The formulation for a constrained response changes the form of the equations from that for the unconstrained mean-square minimization. If the input signal model is modified to include a sinusoidal component, the mean-square error will include the effect of this sinusoid. By increasing the amplitude of the sinusoidal component relative to the other components of the signal, the error at the frequency of the sinusoidal component will receive more weighting.

The sinusoidal component is a signal of the form

$$s(n) = A_c \cos(\omega_c n + \theta), \tag{13}$$

where θ is a uniformly distributed random phase. This randomized phase component is appropriate to force the solution to be independent of phase. In addition, the resulting autocorrelation is stationary. The autocorrelation of the signal plus sinusoid is

$$r'_{xx}(k) = r_{xx}(k) + \tfrac{1}{2}A_c^2 \cos(\omega_c k). \quad (14)$$

The coefficients of the interpolating filter can be found by using the modified correlation function in the equations for the unconstrained minimization.

A potential problem with this approach is that the correlation matrix becomes increasingly ill-conditioned as A_c increases. The rank of the matrix approaches 2 in the limit of large A_c. For moderate values of A_c, the effect of the modified correlation matrix will shift the solution so as to reduce the mean-square interpolation error for inputs of frequency ω_c.

5.2. Nyquist filter

The standard Nyquist design procedure can be modified to include constraints at a certain frequency by using a constrained filter design procedure for designing the stopband portion of the filter. As an alternative, an increased weighting can be given to the aliased frequencies of a given input sinusoid. A large weight at the aliased frequencies will tend to force zeros to occur at those frequencies. Such strategies are appealing since the standard Nyquist design procedure needs to be only slightly modified to constrain or weight the response for sinusoidal inputs of a given frequency.

6. Design examples

Several filters were designed to illustrate the procedure described in this paper.

Interpolating filters

First, we return to the application that motivated the search for means to constrain the response. Consider the interpolating filter needed to convert a CCIR luminance signal (sampling rate 13.5 MHz) to an NTSC luminance signal (sampling rate 14.4 MHz). A sampling rate conversion factor of 35/33 is used. This corresponds to interpolating by a factor of 35 and then subsampling by a factor of 33. Here we focus on the interpolation step. An interpolation filter was designed ($N = 35$) with 279 taps (8 non-zero taps for each output sample), and with the input power spectrum modelled as a raised-cosine response with transition between 5 MHz and 6 MHz. Applying a sampled test pulse (rise and fall times 125 ns, measured between the 10% and 90% points) gives the time response shown in Fig. 1. The peak ripple due to dc reconstruction errors is 2.3%.

The interpolation filter was redesigned with a dc constraint. The corresponding pulse response appears in Fig. 2. The top of response is now flat. The change in impulse response needed to accomplish this change is very small - the maximum change in any filter coefficient is 0.0034.

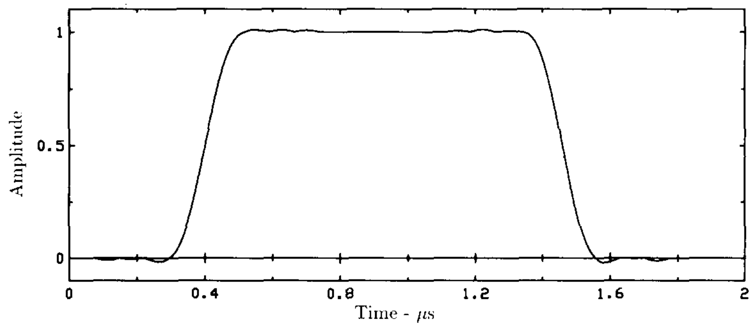

Fig. 2. Response of a constrained interpolation filter to an input T-step pulse (125 ns rise and fall times) ($N = 35$, 279 coefficients, sampling rate 472.5 MHz, designed for a raised-cosine power spectrum with transition between 5 MHz and 6 MHz, constrained at dc).

Consider an example to illustrate a constrained design for a short filter. The interpolating ratio is $N = 5$. The input power spectrum is assumed to be a raised-cosine, flat to $0.8\pi/N$ and then rolling off to become zero at π/N. Each subfilter has 6 coefficients, giving a filter with a total length of 29 coefficients (24 non-zero, non-unity coefficients). The overall filter is linear phase. Figures 3 and 4 compare the frequency response of the unconstrained design with that for a filter which requires exact reproduction of dc and a frequency of $0.5\pi/N$. Three degrees of freedom for each subfilter are used up in meeting these constraints.

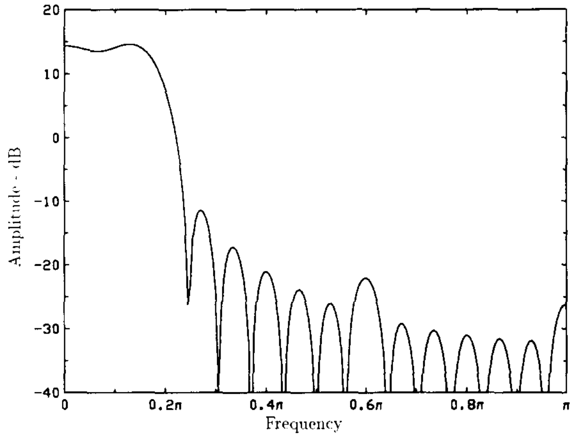

Fig. 3. Frequency response of an unconstrained interpolation filter ($N = 5$, designed for a raised-cosine power spectrum with transition between 0.16π and 0.2π).

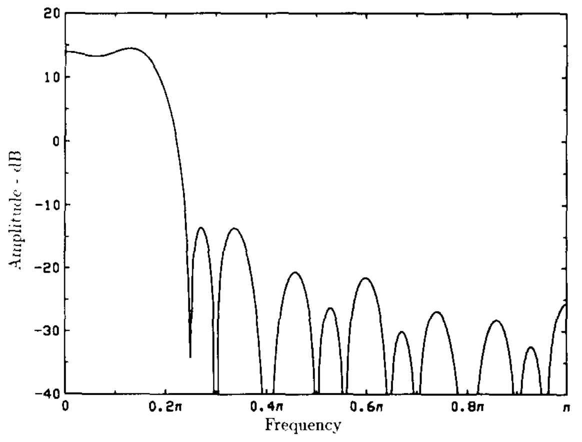

Fig. 4. Frequency response of a constrained interpolation filter ($N = 5$, designed for a raised-cosine power spectrum with transition between 0.16π and 0.2π, constrained at dc and $\omega_c = 0.1\pi$). Note the spectral nulls at 0.3π, 0.4π, 0.5π, 0.7π, 0.8π and 0.9π.

The original unconstrained filter has a peak dc ripple of 7.2% and a peak amplitude ripple of 2.5% for a sinusoidal input at $0.5\pi/N$. The unconstrained filter has normalized mean-square interpolation errors of 0, 6.63×10^{-3}, 17.12×10^{-3}, 17.12×10^{-3} and 6.63×10^{-3} for sample phases $q =$ 0, 1, 2, 3 and 4, respectively. The signal-to-noise ratio (SNR) corresponding to the mean-square error averaged over sample phases is 20.22 dB. For the constrained filter, the corresponding normalized mean-square interpolation errors are 0, 7.29×10^{-3}, 18.79×10^{-3}, 18.79×10^{-3} and 7.29×10^{-3}. These mean-square errors are calculated using the same raised-cosine power spectrum model for the input signal. The SNR for the constrained case is 19.82 dB. The penalty paid for the constraints is small as measured by the mean-square interpolation error. The constrained filter has zeros at the appropriate frequencies to avoid aliasing of the constraint frequencies. Since these zeros are in the stopband, the stopband attenuation of the filter is comparable to that of the unconstrained filter.

The approximate design technique using an autocorrelation function with added sinusoidal components was tested for this design case. With the relative amplitude of the autocorrelation of the dc and sinusoidal components set to 1000 times that of the original autocorrelation function at $k = 0$, the results are very close to that of the contrained optimization. The resulting filter exhibits ripples of less than 0.0002% at the constrained frequencies. No numerical difficulties were encountered with this case (double precision arithmetic) unless the relative amplitudes of the autocorrelation terms for the dc and sinusoidal components were set to be extremely large (more than 10^{12}).

Subfilters of different lengths

In the previous example, each subfilter was of length 6. A second example, using the same constraint frequencies, shows the effect of using subfilters of different lengths, again with a total of 24 non-zero, non-unity coefficients. However, now the subfilter lengths are distributed as 5 ($q=1$), 7 ($q=2$), 7 ($q=3$) and 5 ($q=4$). As shown in the appendix, the design of interpolating filters accommodates an overall FIR filter $h(n)$ with arbitrary length subfilters. However, if we use a design procedure which generates equal length subfilters, we can design several $h(n)$ of different lengths and extract the needed subfilters. The length 5 subfilters were extracted from a 23 tap filter, while the length 7 subfilters were extracted from a 35 tap filter. Decreasing the length of the subfilters for $q=1$ and $q=4$ from 6 to 5 increases the mean-square error at those sampling phases, while increasing the length of the subfilters from 6 to 7 for $q=2$ and $q=3$ decreases the mean-square error for those sampling phases. The resulting filter has normalized mean-square errors of 0, 12.03×10^{-3}, 13.13×10^{-3}, 13.13×10^{-3} and 12.03×10^{-3}. There is a small net benefit with different length subfilters - the corresponding SNR improves to 19.97 dB. The frequency response of this filter is shown in Fig. 5. Note that in this case the passband is much flatter, but the stopband suppression is poorer.

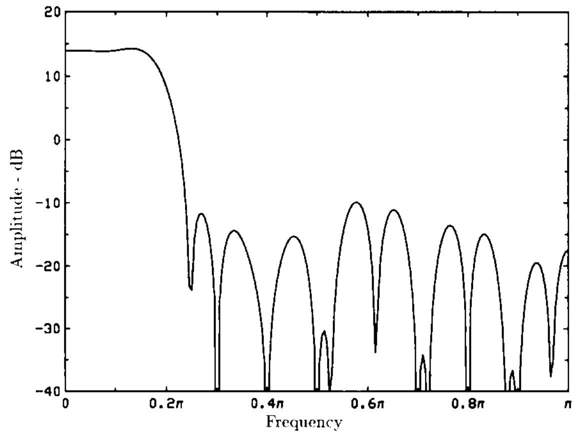

Fig. 5. Frequency response of a constrained interpolation filter with different length subfilters ($N=5$, designed for a raised-cosine power spectrum with transition between 0.16π and 0.2π, constrained at dc and $\omega_c=0.1\pi$). Note the spectral nulls at 0.3π, 0.4π, 0.5π, 0.7π, 0.8π and 0.9π.

Nyquist filter

The design example for a linear phase Nyquist filter has 29 coefficients. The design uses $N=4$ and a stopband edge at 0.375π. The filter is factorable into a minimum and maximum phase part, each with 15 coefficients. The overall Nyquist filter shows a peak dc error of 0.15%. This filter was redesigned using the factorable Nyquist approach with the dc response constrained. The frequency responses of the unconstrained and constrained designs are shown in Figs. 6 and 7. The constraint

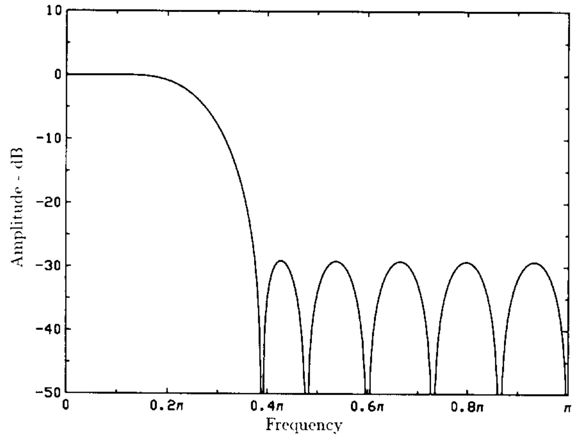

Fig. 6. Frequency response of an unconstrained minimum phase factor of a Nyquist filter ($N=4$, stopband edge 0.375π).

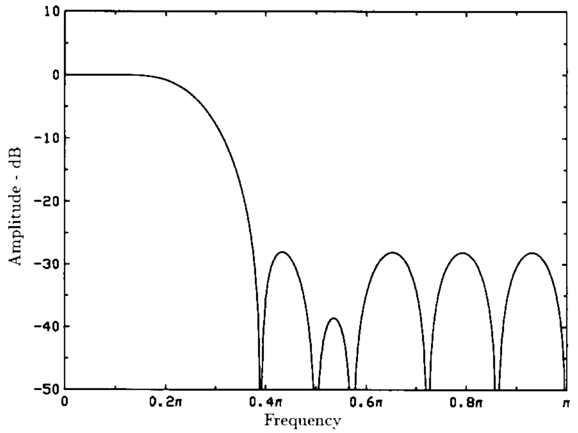

Fig. 7. Frequency response of a constrained minimum phase factor of a Nyquist filter ($N=4$, stopband edge 0.375π, constrained at dc). Note the spectral null at 0.5π.

reduces the stopband suppression by 2.1 dB. The alternate approach using a large weighting at the aliased constraint frequencies generates a filter very close to that for the exact constraint method.

7. Summary and conclusions

This paper has introduced the concept of constraining the response of an interpolating or Nyquist filter to give error-free response at certain frequencies. The ability to design filters with such properties is important in many practical applications. The results show that the presence of constraints need not significantly degrade the performance of the filter as measured in terms of interpolation error or stopband suppression. In addition, approximate techniques can be applied to the filter design process. These approaches have the advantage of requiring only minor modifications to existing design procedures.

Appendix A. Minimum mean-square error interpolator

The system under consideration is shown in Fig. A.1. The input signal is assumed to be a stationary process with autocorrelation function $r_{xx}(k)$. The input signal is subsampled and then sample rate increased to produce the signal $u(n)$. The signal $u(n)$ contains every Nth sample of $x(n)$ with zero samples in between.

The output signal $y(n)$ can be written as

$$y(n) = \sum_k h(k) u(n-k), \tag{A.1}$$

or in vector form

$$y(n) = \boldsymbol{h}^{\mathrm{T}} \boldsymbol{u}^{(n)}. \tag{A.2}$$

The limits of the sum (and hence the range of indices in the vectors) are left unspecified in order to simplify the notation. The mean-square interpolation error is

$$\begin{aligned}\varepsilon(n) &= \mathrm{E}[x^2(n)] - 2\boldsymbol{h}^{\mathrm{T}} E[x(n)\boldsymbol{u}^{(n)}] \\ &\quad + \boldsymbol{h}^{\mathrm{T}} E[\boldsymbol{u}^{(n)}\boldsymbol{u}^{(n)\mathrm{T}}]\boldsymbol{h} \\ &= r_{xx}(0) - 2\boldsymbol{h}^{\mathrm{T}}\boldsymbol{\alpha}^{(n)} + \boldsymbol{h}^{\mathrm{T}}\boldsymbol{R}^{(n)}\boldsymbol{h},\end{aligned} \tag{A.3}$$

where the matrix $\boldsymbol{R}^{(n)}$ is the matrix of correlations for $u(n)$ and the vector $\boldsymbol{\alpha}^{(n)}$ is the cross-correlation between $x(n)$ and $\boldsymbol{u}^{(n)}$. The nonzero elements of $\boldsymbol{R}^{(n)}$ and $\boldsymbol{\alpha}^{(n)}$ can be expressed in terms of the correlation of $x(n)$. Many rows and columns of $\boldsymbol{R}^{(n)}$ and elements of $\boldsymbol{\alpha}^{(n)}$ are zero due to the fact that $N-1$ out of every N samples of $u(n)$ are zero.

The mean-square interpolation error varies with the sampling time n. Further examination shows that the mean-square error is cyclo-stationary. Writing $n = mN + q$, the mean-square error depends only on the sample phase q. For a given sample phase q, only a subset (every Nth coefficient) of the coefficients of $h(n)$ is used in producing the output sample. Then one can find a reduced set of equations for sample phase q,

$$\varepsilon_q = r_{xx}(0) - 2\boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{\alpha}_q + \boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{R}\boldsymbol{h}_q, \tag{A.4}$$

where

$$\boldsymbol{R} = \begin{bmatrix} r_{xx}(0) & r_{xx}(N) & \cdots & \\ r_{xx}(-N) & r_{xx}(0) & \cdots & \\ \vdots & \vdots & & \vdots \\ & & \cdots & r_{xx}(0) \end{bmatrix}, \tag{A.5}$$

and

$$\boldsymbol{h}_q = \begin{bmatrix} \vdots \\ h(-N+q) \\ h(q) \\ h(N+q) \\ \vdots \end{bmatrix}$$

and

$$\boldsymbol{\alpha}_q = \begin{bmatrix} \vdots \\ r_{xx}(-N+q) \\ r_{xx}(q) \\ r_{xx}(N+q) \\ \vdots \end{bmatrix}. \tag{A.6}$$

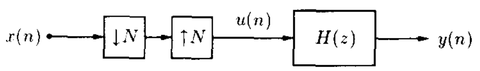

Fig. A.1. Block diagram for interpolation model.

The matrix $\boldsymbol{R}$ is Toeplitz and does not depend on q.

The minimum mean-square error is found by differentiating the above equation with respect to the filter coefficients. This leads to the discrete Wiener-Hopf equation,

$$\boldsymbol{R}\boldsymbol{h}_q = \boldsymbol{\alpha}_q. \tag{A.7}$$

The resulting mean-square error for sample phase q is

$$\varepsilon_q = r_{xx}(0) - \boldsymbol{h}_q^{\mathrm{T}}\boldsymbol{\alpha}_q. \tag{A.8}$$

This value depends on q. With the appropriate range of filter coefficients ($h(0)$ included) the error is exactly zero for $q = 0$.

For equal length subfilters, the mean-square error increases from 0 for sample phase zero to a maximum for a sample phase at or near $N/2$ (see [4, 5]). This suggests that increasing the lengths of subfilters for sample phases with large mean-square errors relative to other subfilters would be beneficial. In some cases, for a fixed computational effort (averaged over all subfilters), the *average* mean-square error can be decreased over the case in which all subfilters are of the same length.

Acknowledgment

The concept of constraining the dc response of an interpolating filter using a Lagrange multiplier was conceived originally by H. L. Nguyen of INRS-Telecommunications for a video transcoding application [3].

References

[1] E. Dubois and R. O'Shaughnessey, "NTSC compatible HDTV transmission methods based on subampling", *Proc. Third International Workshop on HDTV*, Vol. 1, Torino, Italy, 30 August-1 September 1989.

[2] J.H. McClellan, T.W. Parks and L.R. Rabiner, "A computer program for designing optimum FIR linear phase digital filters", *IEEE Trans. Audio Electroacoust.*, Vol. AU-21, December 1973, pp. 506-526.

[3] H.L. Nguyen, "Conversion de fréquence d'échantillonnage: étude théorique et architecture des filtres d'interpolation", M.Sc. thesis, INRS-Telecommunications, University of Quebec, October 1990.

[4] G. Oetken, "A new approach for the design of digital interpolating filters", *IEEE Trans. Acoust. Speech Signal Process.*, Vol. ASSP-27, December 1979, pp. 637-643.

[5] G. Oetken, T.W. Parks and H.W. Schüssler, "New results in the design of digital interpolators", *IEEE Trans. Acoust. Speech Signal Process.*, Vol. ASSP-23, June 1975, pp. 301-309.

[6] R.P. Ramachandran and P. Kabal, "Minimax design of factorable Nyquist filters for data transmission systems", *Signal Processing*, Vol. 18, No. 3, November 1989, pp. 327-339.