JOINT OPTIMIZATION OF MULTIDIMENSIONAL DECIMATION AND INTERPOLATION FILTERS

Pierre Faubert and Eric Dubois
INRS-Télécommunications
Ile-des-Soeurs, Québec, Canada. H3E IH6

Abstract

The joint optimization of prefilters and interpolators in a subsampling and interpolation system is addressed. An objective error criterion which differentially weights filtering and aliasing error is developed, and a procedure is presented for joint design of the filters to minimize this error measure. The technique is extended to the multidimensional case. The method has been succesfully applied in the subsampling and interpolation of time-varying imagery.

1. Introduction

The efficient source coding of digital signals for bandwidth conservation is a major area of research in telecommunications. One very important bit rate reduction technique is subsampling (also called decimation). This technique, along with its dual operation of interpolation, is the main focus of this paper. The question that will be addressed is: What is the optimum way to subsample and then interpolate back a digital signal? Implicit here is the choice of a fidelity criterion, otherwise the meaning of the word "optimum" is far from precise. That will be discussed later. But fundamentally, answering this question amounts to the design of two low-pass filters called the decimation and interpolation filters (or pre- and post-filters) [1].

Previous work on this problem has mainly concentrated on optimizing either the prefilter or the interpolator. A good example of this is the work of Oetken on optimizing the interpolator subject to a mean square error criterion [2]. In one instance, an attempt was made to jointly optimize a pair of Gaussian prefilter and interpolator through the use of a series of subjective tests [3]. This paper presents a technique for jointly optimizing the decimation and interpolation filters using an objective fidelity criterion. Both filters are constrained to have the characteristics of finite impulse response (FIR) and linear phase. The objective function is the mean square error with an added differential weighting of the two basic components of the error introduced by the decimation and interpolation system (DIS), namely loss of resolution and aliasing.

Section 2 of the paper presents the one-dimensional DIS and analyses it mathematically. The input/output relation is derived, and the two-component nature of the error is studied. Section 3 introduces the objective function chosen and deals with the details of filter design. Section 4 extends the previous developments to multidimensional signals. Finally section 5 analyses the filter characteristics and discusses a few aspects of their performance.

2. The one-dimensional DIS

Fig. 1(a) presents the block diagram of the decimation and interpolation system. The input digital signal u is first filtered by the prefilter of impulse response h_1, yielding sequence w (block 1). Then w is subsampled by an integer factor S, retaining only every S^{th} sample (block 2). At this point the decimation operation is completed, and the operations of channel coding and transmission take place. At the other end of the DIS, all sampling points that were previously omitted are restored in the signal and given the value zero (block 3) to form sequence r. Finally the interpolation is completed by convolving r with the postfilter of impulse response h_2, to obtain the output sequence v (block 4).

Blocks 2 and 3 can be replaced by the multiplication of sequence w with a periodic unit impulse train of period S, as illustrated in Fig. 1(b). Calling this new sequence s, the input/output relationship becomes:

$$\begin{aligned} v(n) &= h_2(n) * [s(n) \times [h_1(n) * u(n)]] \\ &= h_2(n) * [\frac{1}{S}\sum_{k=0}^{S-1} \exp(j\frac{2\pi}{S}kn) \times [h_1(n) * u(n)]]. \end{aligned} \tag{1}$$

where the second line is obtained by replacing s by its discrete Fourier series representation. Before going any further, the mathematical nature of the signal u must be known. The hypothesis made in this paper is that u is a wide-sense stationary stochastic process with zero mean. Hence w is also a zero mean wide-sense stationary stochastic process. The autocorrélation functions of u and w are related by:

$$\begin{aligned} R_w(n) &= E[w(m)w(m+n)] \\ &= h_1(n) * h_1(-n) * E[u(m)u(m+n)] \\ &= h_1(n) * h_1(-n) * R_u(n) \end{aligned} \tag{2}$$

Since process r is obtained by a repetitive and periodic processing (multiplication by s), it is cyclostationary [4]. This means that its first and second order statistics are periodic. The usual procedure is then to average its autocorrelation function over one period S in order to get a function of a single argument. Averaging and taking into account the postfiltering, the input/output relation of the DIS becomes:

$$\begin{aligned} R_v(n) = \frac{1}{S^2} h_2(n) * h_2(-n) * [\sum_{k=0}^{S-1} \exp(j\frac{2\pi}{S}kn) \times \\ [h_1(n) * h_1(-n) * R_u(n)]] \end{aligned} \tag{3}$$

85CH2238-4 REG7/0186 $01.00 © 1985 IEEE

The power spectral density (PSD) of the output process is given by the Fourier transform of R_v:

$$\Phi_v(f) = \frac{1}{S^2}|H_2(f)|^2 \sum_{k=0}^{S-1} \Phi_u(f - \frac{k}{S})|H_1(f - \frac{k}{S})|^2 \quad (4)$$

Note that v inherits the cyclostationarity of r. It is now clear that the subsampling operation introduces $S-1$ replicas of the PSD of w in the frequency domain, equispaced along the frequency axis. The role of the interpolation filter is to eliminate these extraneous replicas in order to keep only those that were originally present in the PSD of w. The role of the prefilter is to reduce the bandwidth of the PSD of u in order to minimize the presence of aliasing in the closely packed PSD of process r. Both filters must clearly perform low-pass operations. The pair must also have an in-band gain equal to the subsampling factor S.

We now turn our attention to the error that the DIS introduces. First we form the error process $x_e(n) = u(n) - v(n)$. It is also cyclostationary. Averaging over a period, its autocorrelation function is found to be:

$$\begin{aligned} R_{x_e}(n) &= E[(u(m) - v(m))(u(n+m) - v(n+m))] \\ &= R_u(n) + \\ &\frac{1}{S^2}h_2(n) * h_2(-n) * [[h_1(n) * h_1(-n) * R_u(n)]s(n)] \\ &- \frac{2}{S^2}h_2(n) * h_1(n) * R_u \end{aligned} \quad (5)$$

which becomes, when transformed to the frequency domain:

$$\begin{aligned} \Phi_{x_e}(f) = &\Phi_u(f)[1 - \frac{2}{S}|H_1(f)||H_2(f)| \\ &+ \frac{1}{S^2}|H_1(f)|^2|H_2(f)|^2] \\ &+ \frac{1}{S^2}|H_2(f)|^2 \sum_{k=0}^{S-1} \Phi_u(f - \frac{k}{S})|H_1(f - \frac{k}{S})|^2 \end{aligned} \quad (6)$$

The PSD of the error process is broken up into a sum of two terms. These correspond to its two basic components. The first one (within the large bracket) is associated with loss of resolution (denoted $B(f)$ from now on). It represents the effect of the cascade filtering of the prefilter and the postfilter on the baseband spectrum. The degradations it introduces in the signal are the loss of fine spatial detail caused by the attenuation of energy in the higher part of the baseband spectrum. The second component is associated with the aliasing implied by the replicas of the PSD of w introduced by the subsampling (it is denoted $A(f)$ from here on). It shows up in the output signal as masking of small spatial details, this being caused by the intrusion in the passband of v of unattenuated energy belonging to the extraneous PSD replicas and Moiré patterns in periodic structures in the image. Both components of the error are a function of the frequency response of the prefilter and the postfilter, so there exists a basic interdependence between the two filters. It is only natural then to jointly optimize their performance. It is crucial to note that there usually exists a trade-off between the elimination of the loss of resolution and the elimination of aliasing (it is always true when subsampling is severe enough to cause aliasing; we will limit ourselves to that realistic situation). The case of the prefilter demonstrates clearly this fact. On one hand, we want its passband to be narrow enough to substantially reduce aliasing; on the other hand we want its passband to be wide enough so as not to introduce too much loss of resolution.

3. Objective function and filter design

To our knowledge, this work is the first attempt at jointly optimizing the decimation and interpolation filters according to an objective fidelity criterion. It seems logical then to choose an objective function that is analytically tractable. That is why the mean square error (MSE) was chosen.

In many systems, and in particular in image processing, the final element is the human user. It is well known that minimizing the MSE does not necessarily mean that the subjective satisfaction is maximized. Moreover the two components of the error identified above do not generally have the same subjective impact. In an attempt to devise a better and more flexible fidelity criterion, a differential weighting of the two components of the error has been added. Hence loss of resolution and aliasing can be given two distinct "cost functions". Denoting them respectively $W_{lr}(f)$ and $W_a(f)$, we form the following PSD:

$$\Phi_{x_{dw}}(f) = |W_{lr}(f)|^2 \times B(f) + |W_a(f)|^2 \times A(f) \quad (7)$$

The spatial error process that corresponds to this PSD will be called the doubly-weighted error sequence ($x_{dw}(n)$). To obtain its explicit expression, we define first the following two processes ($\delta(n)$ is the unit impulse):

$$b(n) = u(n) * [\delta(n) - \frac{1}{S}h_2(n) * h_1(n)] \quad (8)$$

$$a(n) = \frac{1}{S}h_2(n) * [\sum_{k=1}^{S-1} \exp(j\frac{2\pi}{S}kn)[h_1(n) * u(n)]] \quad (9)$$

Now it can be shown that $b(n)$ and $a(n)$ have respectively $B(f)$ and $A(f)$ as PSD. It can also be shown that these two processes are uncorrelated, taking into account the cyclostationarity of $a(n)$. Denoting $w_{lr}(n)$ as the inverse transform of $W_{lr}(f)$ and $w_a(n)$ as the inverse transform of $W_a(f)$, we finally get the expression of the doubly-weighted error sequence:

$$x_{dw}(n) = w_{lr}(n) * b(n) + w_a(n) * a(n) \quad (10)$$

Figure 2 illustrates a system that could be used to form the error process x_{dw}, comprising every element of the DIS. Note that the weighting functions appear now as digital filters. Stated explicitly, the objective function, denoted MSE_{dw}, is:

$$\text{MSE}_{dw}(n) = E[||x_{dw}(n)||^2] \quad (11)$$

The optimum decimation and interpolation filters that we seek are those which minimize MSE_{dw}. In order to eliminate phase distortion, both filters are restricted to be zero phase FIR filters, so that

$$h_i(n) = h_i(-n), i = 1, 2. \quad (12)$$

The two weighting filters are also constrained to have these characteristics. A relatively straightforward but tedious development then leads to the explicit expression of MSE_{dw} in terms of the autocorrelation R_u, the weighting filter impulse responses and the impulse responses of the prefilter and the postfilter. The coefficients of the latter two are the unknowns of the expression. The final expression will not be printed here because it is too lengthy.

The MSE_{dw} is a biquadratic function of both filter impulse responses, each one being present up to the second degree. Two minimization algorithms have been developed to obtain the values of the finite number of coefficients allowed to the two filters. One was the use of the quasi-Newton algorithm applied to the nonlinear function MSE_{dw}. The other is the iterative minimization of the quadratic function obtained in terms of the coefficients of either the prefilter or the postfilter. This method starts with an initial set of reasonable values for the coefficients of the prefilter, and continues by replacing the obtained values in the quadratic equations for the other filter. This continues up to the point where the marginal gain in MSE from one iteration to the next goes below a certain threshold. Both methods have been used a great number of times in largely varying conditions (filter orders,subsampling factors,etc.) and they have always been found to give the same solutions, independently of starting point.

4. The multidimensional case

Uniform sampling in the multidimensional case is done according to a sampling lattice. It is a structured set of points in space that are spanned by all integer linear combinations of a finite set of real vectors, called the basis of the lattice. Let the process u be a multidimensional signal of dimension D. Let the sampling lattice according to which it is sampled be Λ. Then Λ's basis contains D vectors. Let M_Λ be the $D \times D$ matrix whose columns are the D basis vectors. It is called a sampling matrix of the lattice. Then the sampling operation can be expressed as [5]:

$$u(\mathbf{n}) = u_c(M_\Lambda \mathbf{n}), \mathbf{n} \in Z^D \quad (13)$$

where the bold character $\mathbf{n}$ represents a multidimensional vector of integers. In the present context, the signal is subsampled from lattice Λ to sub-lattice Γ. The former is a lower density lattice every point of which also belongs to Λ. Then, as in section 2, the subsampling factor is an integer. It is expressed in terms of the determinants of the sampling matrices:

$$S = \frac{|\det M_\Gamma|}{|\det M_\Lambda|} \quad (14)$$

The extension to the multidimensionnal case is easily made making use of the relation 13. All arguments of each expression of section 2 simply must be interpreted as multidimensional integer vectors.

5. Results

Figure 3 shows one example of filter characteristics obtained with this method. It shows the frequency responses a pair of one-dimensional prefilter and postfilter for $S = 10$ and an exponential autocorrelation model with a parameter equal to 0.9. The orders of the filters are 29 and 49 respectively. First, we can see that both filters have a passband that is well adjusted to the subsampling factor (cutoff frequency=.05 on the normalized axis), although the prefilter frequency response is slightly inferior to 0 dB in the upper part. However, it is very interesting to note that the postfilter boosts those frequencies, such that the cascade gain is almost equal to $10\log(10)$ over all the passband. The interpolator also aligns transmission minima on the center on the replicas introduced by the subsampling.

A study of the gain to be expected from prefiltering has been carried out for various exponential autocorrelation models, for signals of dimension 1, 2 and 3. It has been found to be in the range of 1.2 to 2.0 dB. A comparative study of their performance versus Gaussian and maximally flat pairs of filters has also been carried out, yielding average gains of 1.1 and 0.5 dB respectively. Identification of pairs of filters that offer good performance for the latter two types of filter stresses a strong point of the joint optimization method: the ease with which a well adapted pair of prefilter and postfilter can be obtained automatically.

Finally, our optimum filters were applied to the problem of subsampling and interpolating image sequences. Excellent results were obtained on black and white moving pictures for subsampling factors of up to 8 with originals sampled at $4f_{sc}$ and quantized with eight bits. The trade-off between loss of resolution and aliasing has also been briefly explored with the preliminary result that for large subsampling factors it is preferable to give higher weight to the aliasing component.

6. Conclusion

A new design method for jointly optimizing decimation and interpolation filters has been presented. The filters present excellent performance. The method has the great advantage of automatically yielding pairs of filters well-adapted to any DIS. The next step is to explore in depth the differential-weighting of the two components of the error from a subjective standpoint.

References

[1] R.E. Crochiere and L.R. Rabiner, " Interpolation and decimation of digital signals - A tutorial review", Proc. of the IEEE, vol. 69, No. 3, March 1981.

[2] G. Oetken, T.M. Parks and H.S. Schüssler, "New results in the design of digital interpolators", IEEE Trans. on ASSP, vol ASSP-23, No. 3, June 1975.

[3] J.N. Ratzel, The Discrete Representation of Spatially Continuous Images, doctorate thesis, Massachusetts Institute of Technology, August 1980.

[4] L.E. Franks, Signal Theory, Englewood Cliffs, NJ: Prentice-Hall, 1969.

[5] R.M. Mersereau and T.C. Speake, "The processing of periodically sampled signals", IEEE Trans. on ASSP. vol. ASSP-31. No. 1, pp 188-194, February 1983.

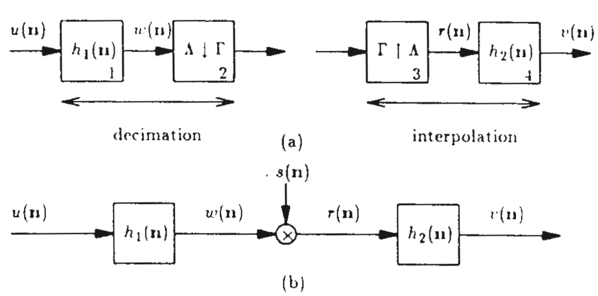

Fig. 1 The decimation and interpolation system (DIS).

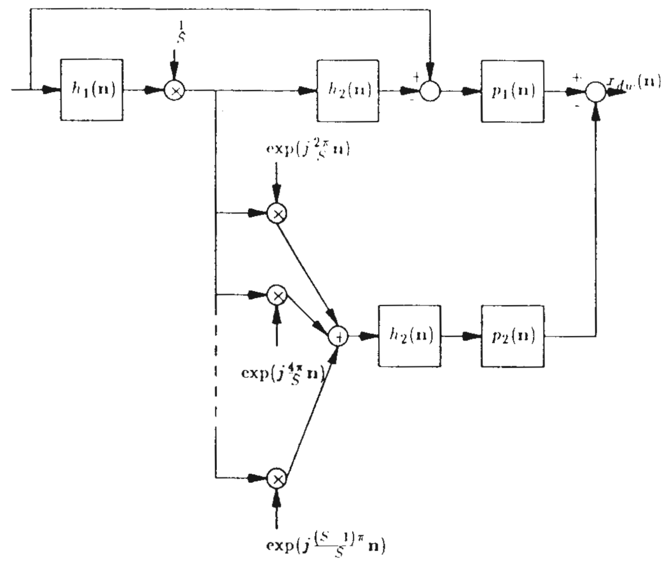

Fig. 2 Doubly-weighted error sequence formation.

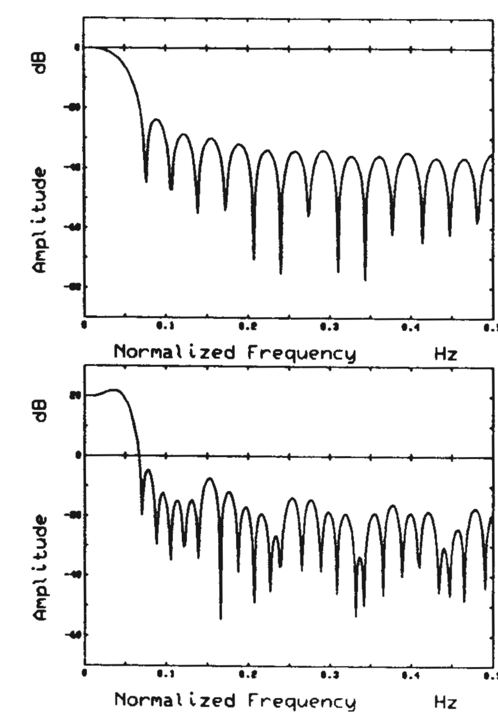

Fig. 3 Pre-filter and postfilter frequency responses for $S = 10$ and $R(n) = (.9)^n$, the orders being 29 and 49 respectively.