Signal Processing 6 (1984) 97–112
North-Holland

MULTI-DIMENSIONAL SUB-BAND CODING: SOME THEORY AND ALGORITHMS

Martin VETTERLI, member EURASIP

Laboratoire d'Informatique Technique, Ecole Polytechnique Fédérale de Lausanne, 16 Chemin de Bellerive, CH-1007 Lausanne, Switzerland

Received 28 September 1983
Revised 24 November 1983

Abstract. A system is proposed in order to split a multi-dimensional signal into N sub-bands, which are then subsampled by N. Subsequent upsampling and filtering allows the recovery of the original signal. Main features are a good bandpass characteristic of the channels, automatic aliasing cancellation and spectral invariance of the overall system. The one dimensional case, known as the quadrature mirror filter (QMF, [1]), is generalized for both the separable and the non-separable case. A parallel implementation, based on pseudo-QMF filters, is presented as an efficient way to split a signal into equal sub-bands.

Zusammenfassung. Ein System wird vorgeschlagen um ein multi-dimensionales Signals in N frequenzbänder zu unterteilen, in welchen nur einer von N Abtastwerten beibehalten wird. Von diesen Kanälen kann das originale Signal mittels Filtrierung perfekt reconstruiert werden. Wichtigste Merkmale sind eine gute Bandpass Charakteristik der Kanäle, automatische Unterdrückung der spektralen Rückfaltungen sowie spektrale Invarianz des gesamten Systems. Der 1-dimensionale Fall, bekannt als quadratur Spiegel Filter (QMF, [1]), wird für den separablen und den nicht-separablen Fall generalisiert. Eine parallele Ausführung, basierend auf pseudo-QMF Filtern, wird als eine effiziente Technik für die Teilung von einem Signal in Frequenzbändern vorgestellt.

Résumé. Un système pour la séparation d'un signal multi-dimensionnel en N sous-bandes, sous-échantillonées d'un facteur N, est présenté. Il est possible de reconstruire le signal original à partir des canaux au moyen d'un sur-échantillonage suivi d'un filtrage approprié. Les propriétés principales du système sont la charactéristique passe-bande des canaux, la suppression automatique des repliements spectraux ainsi que l'invariance spectrale du système total. Le cas mono-dimensionnel, connu sous le nom de filtres miroirs en quadrature (QMF, [1]), est généralisé pour le cas séparable ainsi que pour le cas non-séparable. Une version parallèle, basée sur des filtres pseudo-QMF, est présentée comme une technique efficace de séparation d'un signal en sous-bandes.

Keywords. Efficient sub-band coding of images, efficient filter banks for images, image processing.

1. Introduction

Since its introduction [1], the quadrature mirror filter (QMF) technique has been widely used in one-dimensional sub-band coding [2]. The basic idea is to split a signal into subsampled channels and none the less keep the following two fundamental properties:

- frequency discrimination: each channel should have a good bandpass characteristic (typically: better than the DFT).
- signal recovery: if the subsampled channels are not coded, one should be able to recover the original signal without aliasing or spectral deformation (e.g., as guaranteed by the inversion property of the DFT).

Both of these requirements are fulfilled by the QMF filters: the signal is split into a low-pass and a high-pass version and both are subsampled by 2. After upsampling by 2 and interpolation with the same filters (modulo a sign change) the original signal is perfectly recovered. A tree structure is used for repeated splitting. For a review of QMF filters and sub-band coding, see references [3] and [4], and for specific filters, see reference [5].

Since the above properties are fundamental in nature, they are likely to be important in multi-dimensional signal processing as well. To be able to split a two-dimensional signal into sub-bands with a good frequency discrimination and subsample them without any loss of information seems useful in various applications, as for example in image compression, texture analysis/classification and spectral analysis.

First, the QMF concept is generalized to the two-dimensional case. Section 2 investigates the non-separable case and Section 3 deals with the separable one. Both aliasing cancellation and spectral invariance are obtained, showing that the QMF property is a fundamental one. Generalization to higher dimensions is left to the appendix. Section 4 looks at the implementation of the QMF, especially at the computational load. That the resulting rather heavy load can be substantially reduced is shown in Section 5, where a parallel implementation (making use of an FFT-type algorithm) is presented. In the concluding remarks, the results are discussed and potential applications are suggested, with hopes that some interest for these algorithms will be awakened.

2. The two-dimensional general case

Using a system as depicted in Fig. 1, we will show that the QMF condition is maintained, which means that the aliasing terms due to the subsampling are cancelled.

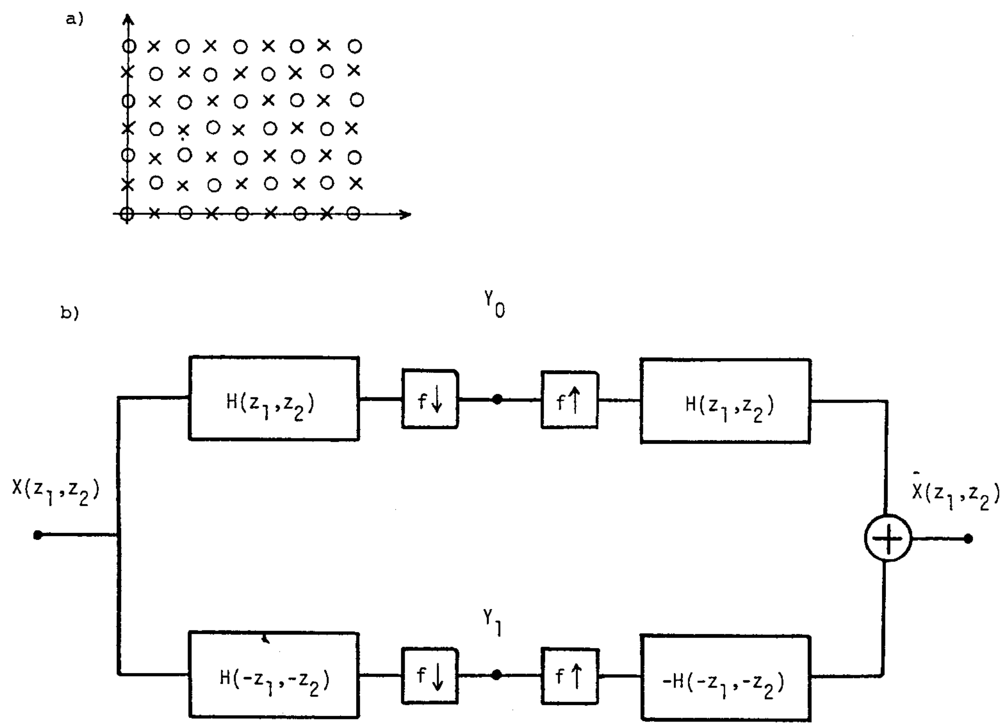

Fig. 1. General two-dimensional system. a) subsampling function, O: retained sample, ×: disregarded sample; b) system configuration.

The two-dimensional input $x(n_1, n_2)$ is filtered by $h(n_1, n_2)$ and its mirror filter $(-1)^{n_1+n_2}h(n_1, n_2)$. The resulting channels are subsampled by the function $f\downarrow$ which keeps one out of two samples, namely the samples where n_1+n_2 is even. For the reconstruction, the channels are upsampled by the reciprocal function $f\uparrow$ (which inserts zeroes where samples were dropped by $f\downarrow$), filtered and added, thus producing $\hat{x}(n_1, n_2)$.

The process of down and upsampling can be viewed as a modulation by the function $f(n_1, n_2)$ given by (2.1):

$$f(n_1, n_2) = 1/2(1+\mathrm{e}^{\mathrm{j}\pi(n_1+n_2)}). \tag{2.1}$$

In the x-transform domain, with $H(z_1, z_2)$ being the transfer function of the filter, the mirror filter becomes simply $H(-z_1, -z_2)$. The modulation by $f(n_1, n_2)$ becomes a convolution in the z-transform domain. Therefore, the output $\hat{X}(z_1, z_2)$ is equal to:

$$\begin{aligned}\hat{X}(z_1, z_2) &= 1/2[H(z_1, z_2)\cdot X(z_1, z_2)+H(-z_1, -z_2)\cdot X(-z_1, -z_2)]\cdot H(z_1, z_2)\\ &\quad -1/2[H(-z_1, -z_2)\cdot X(z_1, z_2)+H(z_1, z_2)\cdot X(-z_1, -z_2)]\cdot H(-z_1, -z_2)\\ &= 1/2[H^2(z_1, z_2)-H^2(-z_1, -z_2)]\cdot X(z_1, z_2).\end{aligned} \tag{2.2}$$

We see that the aliasing term corresponding to $X(-z_1, -z_2)$ disappears automatically, whatever the transfer function $H(z_1, z_2)$ is.

Assuming that $H(z_1, z_2)$ is a two-dimensional symmetric linear phase filter of dimension $L_1\times L_2$, its transfer function can be written as in (2.3) where $H^*(\omega_1, \omega_2)$ is real [6] (we assume normalized frequency).

$$H(\mathrm{e}^{\mathrm{j}\omega_1}, \mathrm{e}^{\mathrm{j}\omega_2}) = \exp[-\mathrm{j}(((L_1-1)/2)\omega_1+((L_2-1)/2)\omega_2)]\cdot H^*(\omega_1, \omega_2). \tag{2.3}$$

Since $H(-\mathrm{e}^{\mathrm{j}\omega_1}, -\mathrm{e}^{\mathrm{j}\omega_2}) = H(\mathrm{e}^{\mathrm{j}(\omega_1+\pi)}, \mathrm{e}^{\mathrm{j}(\omega_2+\pi)})$ we see that if L_1 or L_2 are even (but not both) then eq. (2.2) becomes:

$$\begin{aligned}\hat{X}(\omega_1, \omega_2) &= 1/2\cdot\exp[-\mathrm{j}((L_1-1)\omega_1+(L_2-1)\omega_2)]\\ &\quad\times[H^{*2}(\omega_1, \omega_2)+H^{*2}(\omega_1+\pi, \omega_2+\pi)]\cdot X(\omega_1, \omega_2).\end{aligned} \tag{2.4}$$

Thus, given that the sum of the modules squared is a constant, the reconstructed signal $X(\omega_1, \omega_2)$ is a perfect replica of the input signal modulo a constant factor and shifts.

If L_1 and L_2 are both even or odd, eq. (2.4) is not true anymore. Instead, the system of Fig. 1 has to be modified by placing appropriate delays as for example in Fig. 2. Eq. (2.2) becomes:

$$\begin{aligned}\hat{X}(z_1, z_2) &= 1/2[H(z_1, z_2)\cdot X(z_1, z_2)+H(-z_1, -z_2)\cdot X(-z_1, -z_2)]\cdot z_1^{-1}\cdot H(z_1, z_2)\\ &\quad +[z_1^{-1}H(-z_1, -z_2)\cdot X(z_1, z_2)-z^{-1}\cdot H(z_1, z_2)\cdot X(-z_1, -z_2)]\cdot H(-z_1, -z_2)\\ &= 1/2\cdot z_1^{-1}\cdot[H^2(z_1, z_2)+H^2(-z_1, -z_2)]\cdot X(z_1, z_2),\end{aligned} \tag{2.5}$$

Again, the aliasing terms are automatically cancelled. Since both L_1 and L_2 are even or odd, it follows from (2.3) and (2.5) that:

$$\hat{X}(\omega_1, \omega_2) = 1/2\exp[-\mathrm{j}(L_1\cdot\omega_1+(L_2-1)\cdot\omega_1)][H^{*2}(\omega_1, \omega_2)+H^{*2}(\omega_1+\pi, \omega_2+\pi)]\cdot X(\omega_1, \omega_2). \tag{2.6}$$

The sum of squares being constant, the reconstruction is again perfect (modulo shifts and a constant multiplicative factor). We note that the delay can be arbitrarily placed on z_1 or z_2 and on either branch.

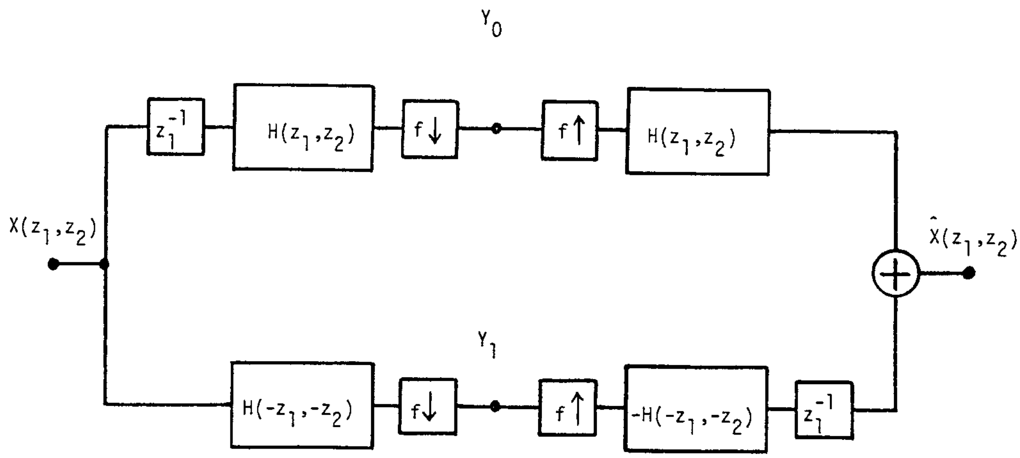

Fig. 2. General two-dimensional system with delays.

A typical filter configuration is shown in Fig. 3. Of course, the splitting can be applied again to the subsampled channels, resulting in a tree structure well known from the one-dimensional sub-band coding techniques using QMF's.

3. The two-dimensional separable case

A system as shown in figure 4 is used to split a two-dimensional signal into 4 channels, each subsampled by 4. The mirror filters are obtained from a single prototype filter modulated by $(-1)^{n_1}$, $(1)^{n_2}$ and $(-1)^{n_1+n_2}$. By generalizing the subsampling concept in [7], we see that a signal $x(n_1, n_2)$ subsampled by 4 (2 along each axis) results in a signal $y(n_1, n_2)$ with the following z-transform:

$$Y(z_1, z_2) = 1/4[X(z_1^{1/2}, z_2^{1/2}) + X(-z_1^{1/2}, z_2^{1/2}) + X(z_1^{1/2}, -z_2^{1/2}) + X(-z_1^{1/2}, -z_2^{1/2})]. \tag{3.1}$$

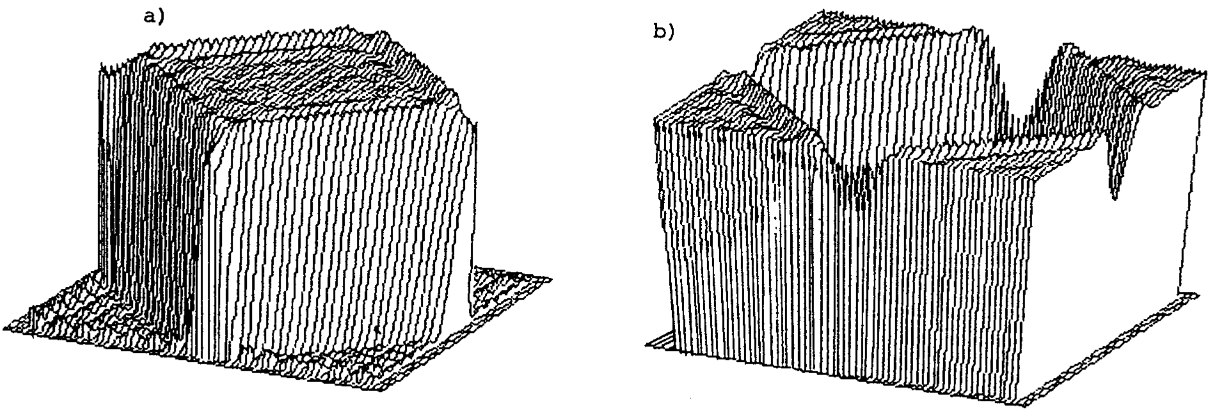

Fig. 3. Magnitude of the frequency response of typical non-separable filters. a) low-pass filter $h(n_1, n_2)$. b) mirror filter $(-1)^{n_1+n_2}h(n_1, n_2)$.

Filtering/downsampling and upsampling/filtering as in Fig. 4 gives the following output:

$$\begin{aligned}\hat{X}(z_1, z_2) = 1/4[[&H(z_1, z_2)X(z_1, z_2) + H(-z_1, z_2)X(-z_1, z_2)\\ &+H(z_1, -z_2)X(z_1, -z_2) + H(-z_1, -z_2)X(-z_1, -z_2)]H(z_1, z_2)\\ &-[H(-z_1, z_2)X(z_1, z_2) + H(z_1, z_2)X(-z_1, z_2)\\ &+H(-z_1, -z_2)X(z_1, -z_2) + H(z_1, -z_2)X(-z_1, -z_2)]H(-z_1, z_2)\\ &-[H(z_1, -z_2)X(z_1, z_2) + H(-z_1, -z_2)X(-z_1, z_2)\\ &+H(z_1, z_2)X(z_1, -z_2) + H(-z_1, z_2)X(-z_1, -z_2)]H(z_1, -z_2)\\ &+[H(-z_1, -z_2)X(z_1, z_2) + H(z_1, -z_2)X(-z_1, z_2)\\ &+H(-z_1, z_2)X(z_1, -z_2) + H(z_1, z_2)X(-z_1, -z_2)]\cdot H(-z_1, -z_2)].\end{aligned} \tag{3.2}$$

All terms containing aliased components of the original signal $x(n_1, n_2)$ have to be cancelled. It can be seen that a necessary and sufficient condition is the separability of the filter, that is:

$$h(n_1, n_2) = h_1(n_1)h_2(n_2), \quad H(z_1, z_2) = H_1(z_1)H_2(z_2). \tag{3.3}$$

In that case, eq. (3.2) reduces to:

$$\begin{aligned}\hat{X}(z_1, z_2) = 1/4\cdot[&H_1(z_1)H_2(z_2) - H_1(-z_1)H_2(z_2) - H_1(z_1)H_2(-z_2)\\ &-H_1(-z_1)H_2(-z_2)]\cdot X(z_1, z_2).\end{aligned} \tag{3.4}$$

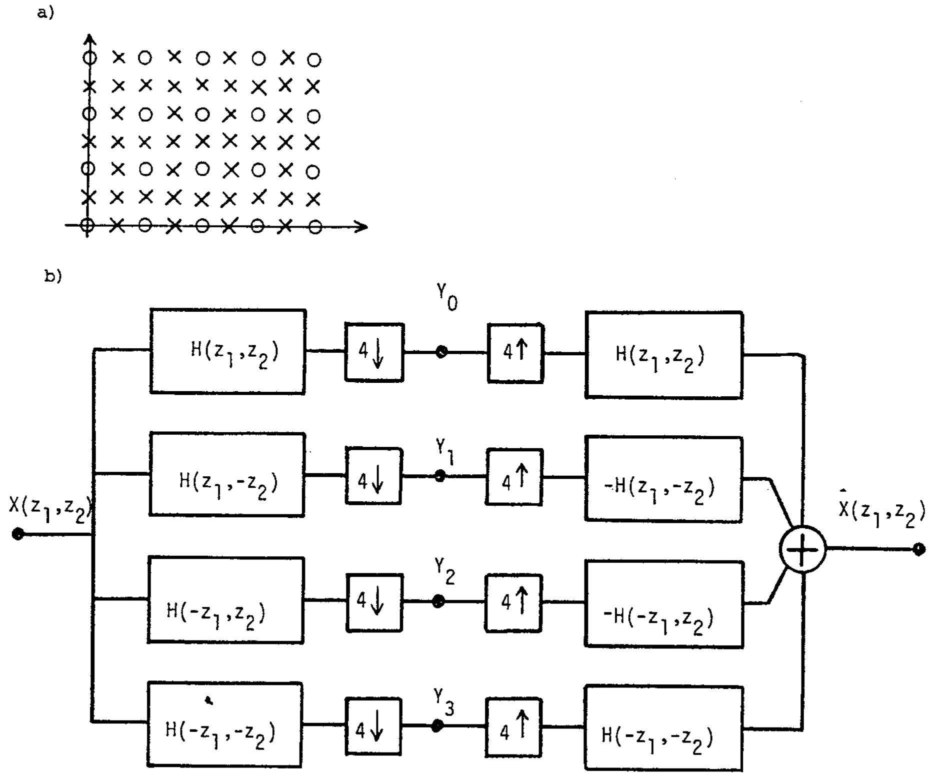

Fig. 4. Two-dimensional system with subsampling by 4. a) subsampling function, O: retained sample, ×: disregarded sample; b) system configuration.

All aliased components have therefore been cancelled. It is assumed that both $h_1(n_1)$ and $h_2(n_2)$ are even length linear phase filters. In that case, $X(\omega_1, \omega_2)$ is equal to (from (2.3)):

$$\hat{X}(\omega_1, \omega_2) = 1/4 \cdot \exp[-j((L_1-1)\omega_1 + (L_2-1)\omega_2)][H_1(\omega_1)H_2(\omega_2) + H_1(\omega_1+\pi)H_2(\omega_2) + H_1(\omega_1)H_2(\omega_2+\pi) + H_1(\omega_1+\pi)H_2(\omega_2+\pi)]X(\omega_1, \omega_2). \tag{3.5}$$

If the sum of the modules squared is equal to a constant, then the signal is perfectly reconstructed. Since the filter is separable, the sum of the one-dimensional filters has to be equal to 1 and thus conventional QMF's can be used. The computation can be performed first along one axis and then along the other. Analysis and synthesis are done as shown in Fig. 5. A typical filter configuration in the frequency domain is presented in Fig. 6.

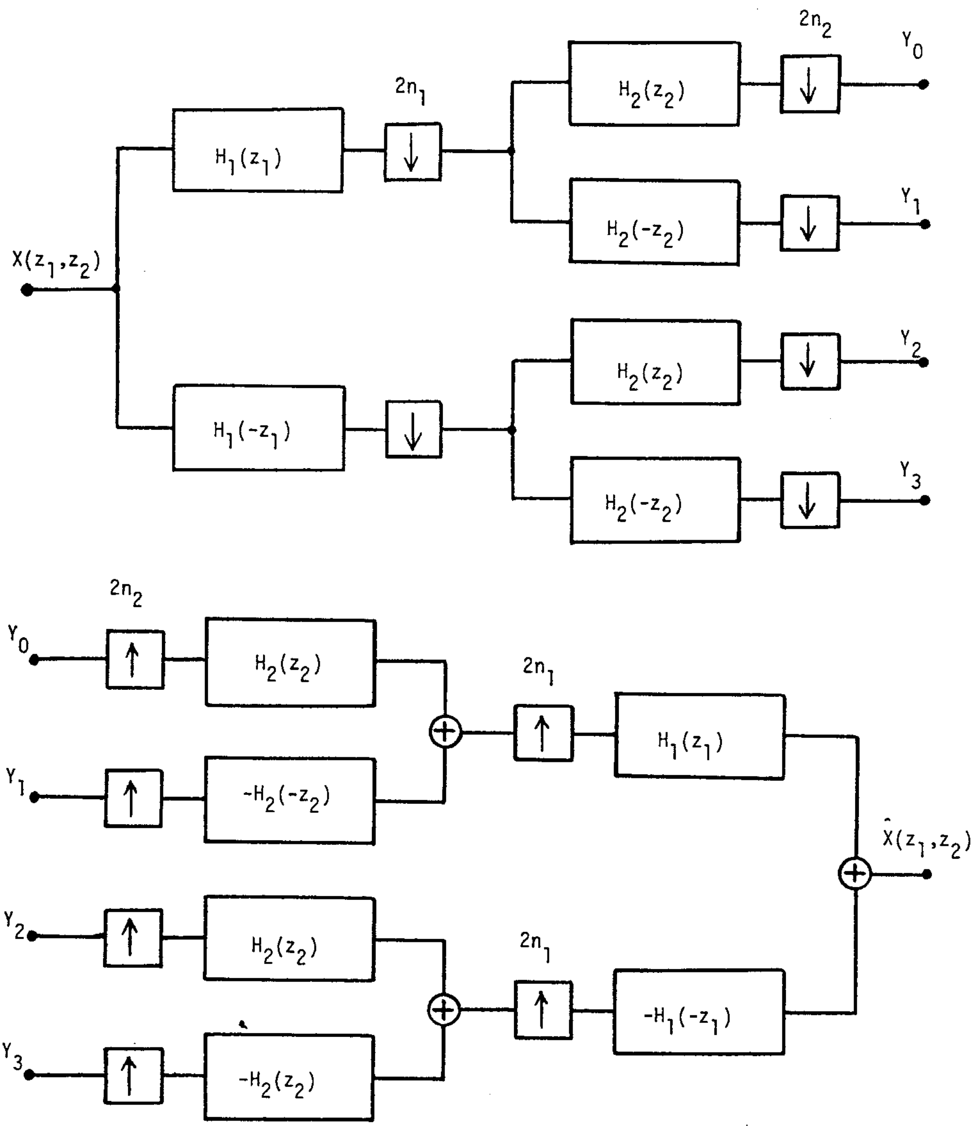

Fig. 5. Analysis/synthesis with separable filters.

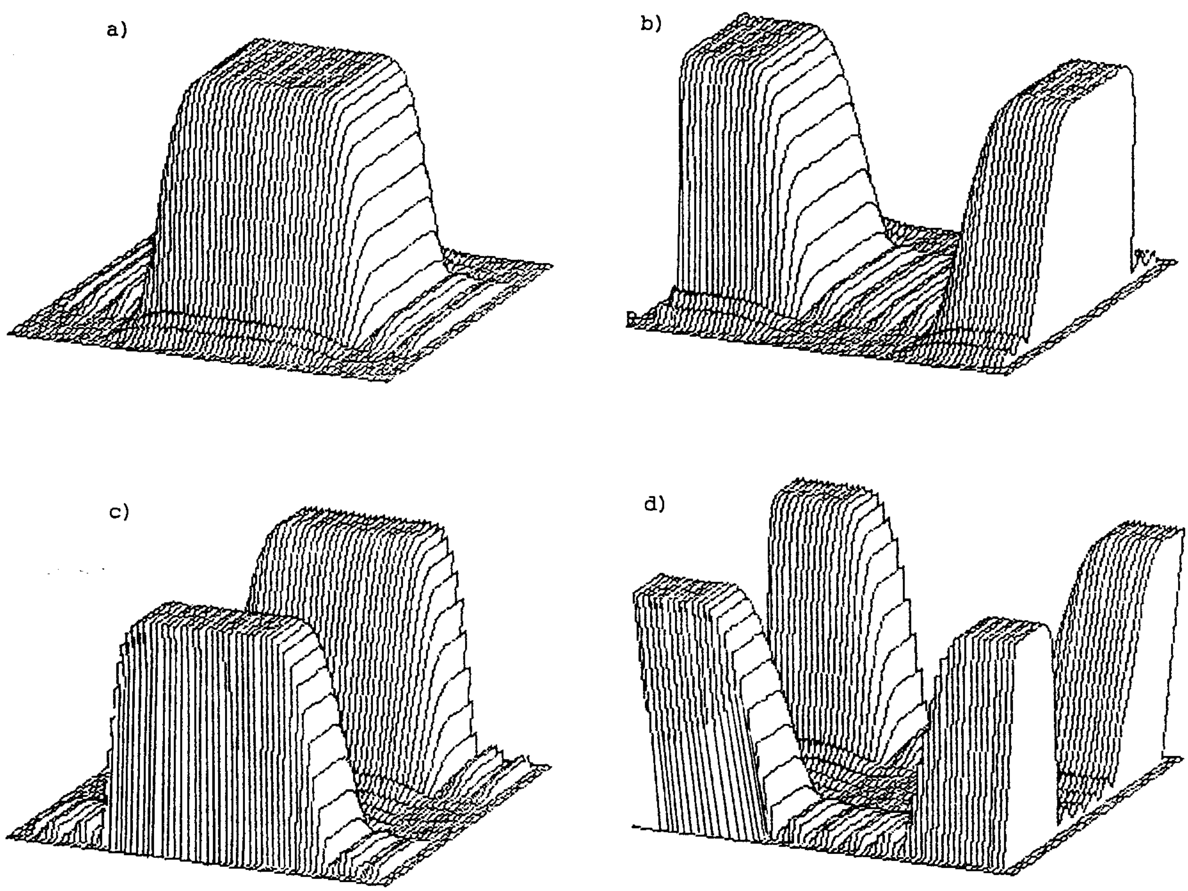

Fig. 6. magnitude of the frequency response of typical separable filters. a) low-pass filter $h(n_1, n_2)$, b) mirror filter $(-1)^{n_1}h(n_1, n_2)$, c) mirror filter $(-1)^{n_2}h(n_1, n_2)$, d) mirror filter $(-1)^{n_1+n_2}h(n_1, n_2)$.

Note that the only perfect prototype filter is the 2×2 unit filter $1/2\cdot(1+z_1^{-1}+z_2^{-1}+z_1^{-1}z_2^{-1})$ which performs, together with its mirror filters, a 2×2 Hadamard transform. A tree structure using this basic filter will perform a Hadamard transform on the input [9].

4. QMF Implementation

Looking first at the non-separable QMF's of Section 2, we note that even length linear phase filters (even symmetry) are not realizable for our problem. The constraint in (2.6) implies that the prototype filter has a 3 dB attenuation on the symmetry axis to its mirror filter. When the filter is of even length and symmetric, it is easy to see that this constraint is violated at the points $\omega_1=0$, $\omega_2=0$ and $\omega_1=\pi$, $\omega_2=0$ where the prototype filter has zeroes. The same holds for its mirror image as well, and the overall system has therefore zero transmission at these two points. Mixed even and odd length filters as in (2.4) will produce one of the two zeroes. Of course, these zeroes are of no consequence when working with real life signals since they occur at half sampling frequency, a location where the signals have low energy.

Nevertheless, only odd length filters will be considered in the following, especially since they are also computationally more efficient. We assume a linear phase $L\times L$ filter $h(n_1, n_2)$ where L is odd. The linear phase constraint together with the fact that the transfer function is identical along the ω_1 and ω_2 axis leads to the following properties of the filter (and mirror filter) impulse response matrix:

- each line and column is symmetric.
- the ith line is identical with the ith column.

This leaves only $1/4(L^2-1)+1$ out of the $2L^2$ possible values for the filter and mirror filter taps. The evaluation of one output of the filter and its mirror filter requires each:

$$\begin{aligned} 1/4\cdot(L^2-1)+1 &\quad \text{mults,} \\ L^2-1 &\quad \text{adds.} \end{aligned} \tag{4.1}$$

This result is obtained when taking advantage of the symmetries. Neglecting the border effects but noting that only one out of two outputs has to be computed (due to the subsequent subsampling), we see that a $P\times P$ picture can be split into two subpictures with a load of:

$$\begin{aligned} (1/4\cdot(L^2-1)+1)\cdot P^2 &\quad \text{mults,} \\ (L^2-1)\cdot P^2 &\quad \text{adds.} \end{aligned} \tag{4.2}$$

It should not be overlooked that, because of the delay on one branch, the filter and its mirror image do not work on the same input, thus not allowing further reductions. For a repeated, k-stage splitting ($k=3$ is depicted in Fig. 7), we note that the filter surface can be divided by 2 at each stage [8] as $L_1^2=L^2$, $L_2^2=L^2/2$, $L_3^2=L^2/4$, where L_i is the filter length at the ith stage. The channels are also subsampled by 2 but this saving is cancelled by the fact that the number of channels doubles. Therefore, the load for k stages (or equivalently 2^k channels) is roughly:

$$\begin{aligned} 1/2\cdot(1-1/2^k)\cdot L^2\cdot P^2 &\quad \text{mults,} \\ 2\cdot(1-1/2^k)\cdot L^2\cdot P^2 &\quad \text{adds.} \end{aligned} \tag{4.3}$$

Turning to the separable case, we note that all simplifications using symmetries, common factors and zero-crossings of the impulse response can be taken over from the one-dimensional case. The prototype filter is obtained as the product of a single length-L linear phase QMF filter (L even). One stage, or 4 channels, requires [3]:

$$\begin{aligned} L\cdot P^2 &\quad \text{mults,} \\ L\cdot P^2 &\quad \text{adds.} \end{aligned} \tag{4.4}$$

If the splitting is repeated, the filter length can be reduced as $L_1=L$, $L_2=L/2$, $L_3=L/4$, where L_i is the filter length at the ith stage. The subsampling is counterbalanced by the increasing number of channels, and the load for k stages or 4^k channels is:

$$\begin{aligned} 2\cdot(1-1/2^k)\cdot L\cdot P^2 &\quad \text{mults,} \\ 2\cdot(1-1/2^k)\cdot L\cdot P^2 &\quad \text{adds.} \end{aligned} \tag{4.5}$$

A two stage splitting is shown in Fig. 8. Further reductions can be obtained by using odd length filters [3]. Note also that the original signal can be recovered with the same amount of computation.

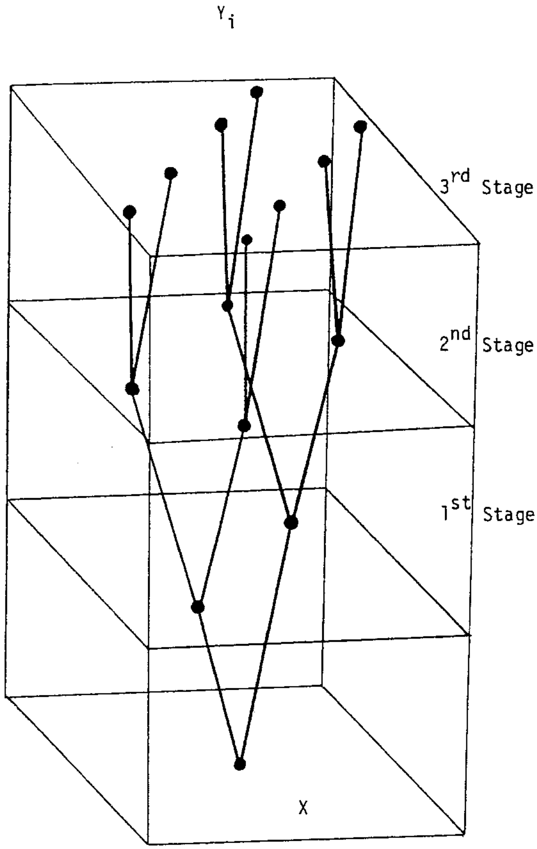

Fig. 7. Non-separable QMF decomposition with 3 stages or 8 channels.

5. Parallel implementation

In this section, we first generalize a result on filter banks where the various filters are derived by uniform modulation from a single prototype filter. This technique, known from transmultiplexers [10], replaces N filters by reduced filters (or polyphase network) and a FFT-like transform (for example a modified DCT as in [11]).

Assume $N_1 \times N_2$ filters uniformly distributed over the frequency space and derived from a prototype low-pass $h(n_1, n_2)$ (see Fig. 9).

$$h_{i_1,i_2}(n_1, n_2) = h(n_1, n_2) \cdot \cos\left(2\pi n_1 \frac{(2i_1+1)}{4N_1}\right) \cdot \cos\left(2\pi n_2 \frac{(2i_2+1)}{4N_2}\right). \tag{5.1}$$

A further assumption is that the prototype filter is a square FIR filter of dimension $L_1 \times L_2 = 2N_1K_1 \times 2N_2K_2$. The output of the filter h_{i_1,i_2} is therefore:

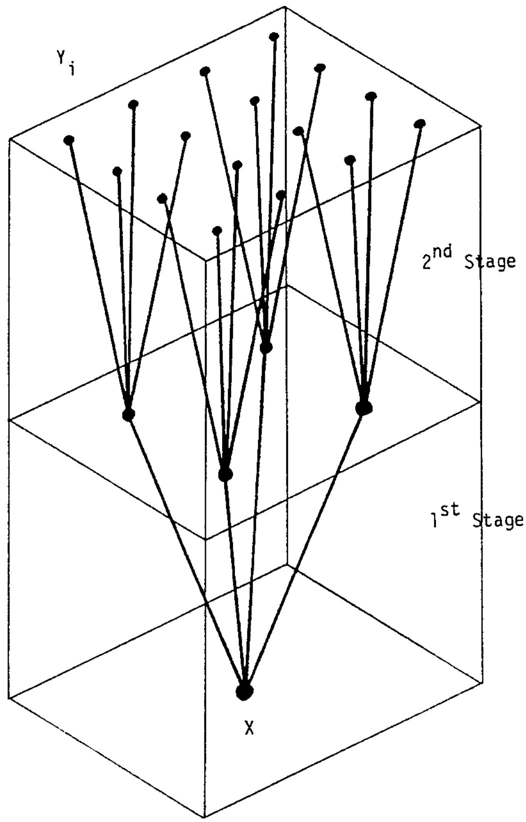

Fig. 8. Separable QMF decomposition with 2 stages or 16 channels.

$$
\begin{aligned}
y_{i_1,i_2}(m_1, m_2) &= \sum_{n_1=0}^{2N_1K_1-1} \sum_{n_2=0}^{2N_2K_2-1} x(m_1-n_1, m_2-n_2)\cdot h(n_1, n_2) \\
&\quad \cdot \cos\left(2\pi n_1 \frac{(2i_1+1)}{4N_1}\right)\cdot \cos\left(2\pi n_2 \frac{(2i_2+1)}{4N_2}\right) \\
&= \sum_{k_1=0}^{K_1-1} \sum_{k_2=0}^{K_2-1} \sum_{l_1=0}^{2N_1-1} \sum_{l_2=0}^{2N_2-1} x(m_1-2N_1k_1-l_1, m_2-2N_2k_2-l_2) \\
&\quad \cdot h(2N_1k_1+l_1, 2N_2k_2+l_2) \\
&\quad \cdot (-1)^{k_1}\cdot \cos\left(2\pi l_1 \frac{(2i_1+1)}{4N_1}\right)\cdot (-1)^{k_2}\cdot \cos\left(2\pi l_2 \frac{(2i_2+1)}{4N_2}\right).
\end{aligned} \tag{5.2}
$$

The above equations can be written as:

$$
\begin{aligned}
y_{i_1,i_2}(m_1, m_2) &= \sum_{l_1=0}^{2N_1-1} \sum_{l_2=0}^{2N_2-1} a(l_1, l_2)\cdot \cos\left(2\pi l_1 \frac{(2i_1+1}{4N_1}\right)\cdot \cos\left(2\pi l_2 \frac{(2i_2+1)}{4N_2}\right), \\
a(l_1, l_2) &= \sum_{k_1=0}^{K_1-1} \sum_{k_2=0}^{K_2-1} x(m_1-2N_1k_1-l_1, m_2-2N_2k_2-l_2) \\
&\quad \cdot h(2N_1k_1+l_1, 2N_2k_2+l_2)\cdot(-1)^{k_1}\cdot(-1)^{k_2}.
\end{aligned} \tag{5.3}
$$

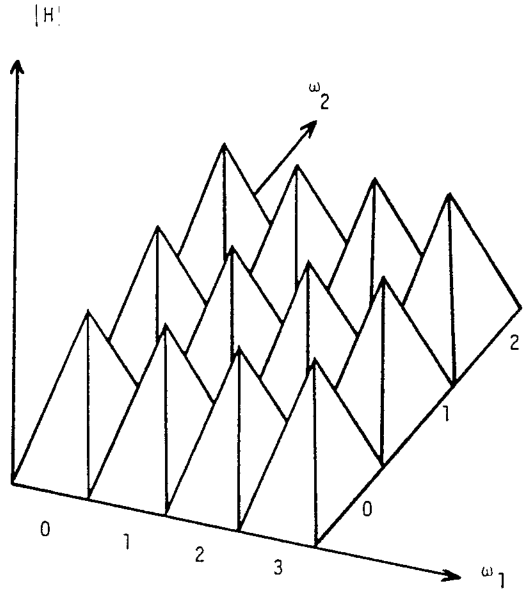

Fig. 9. Uniformly distributed filters with $N_1 = 4$ and $N_2 = 3$.

Therefore, the evaluation of $N_1 \times N_2$ dimension $2N_1K_1 \times 2N_2K_2$ filters has been reduced to $2N_1 \times 2N_2$ filters of dimension $K_1 \times K_2$ and a DCT of dimension $2N_1 \times 2N_2$ where only a quarter of the outputs are computed. This can lead to substantial savings in the computational load.

This rather general result can be used for a special case of filters which allow perfect reconstruction after subsampling, the so-called pseudo-QMF filters [12], [13]. Without going into details, we restate the conditions for pseudo-QMF filterbanks (where N filters are used in parallel in order to filter a signal into N subbands, subsampled by N) derived from a prototype FIR low-pass in the one-dimensional case (see Fig. 10).

– anti-aliasing constraint:

$$|H(\omega)| = 0, \qquad |\omega| \geq 2\pi/2N. \tag{5.4}$$

The magnitude of the transfer function has zero transmission twice above the cut-off frequency.

– flatness constraint:

$$|H(\omega)|^2 + |H(\omega - 2\pi/2N)|^2 = 1, \qquad 0 \leq \omega \leq 2\pi/2N. \tag{5.5}$$

The sum of the squares of the magnitude has to be one over the whole spectrum.

– modulation constraint:

$$h_i(n) = h(n) \cdot \cos(2\pi(2i+1)(2n-N)/8N), \qquad 0 \leq i \leq N. \tag{5.6}$$

This is one of the possible modulation schemes (there is some freedom in the choice of the phase), but this leads to an efficient implementation.

Using filters of this type allows one to filter a signal into N channels, to subsample by N, then to upsample by N, filter by a similar set of filters, sum the outputs and recover the original signal (see Fig. 11). The aliasing components produced by the down/upsampling process are automatically cancelled if (5.4) is respected, and no spectral deformation occurs if (5.5) holds.

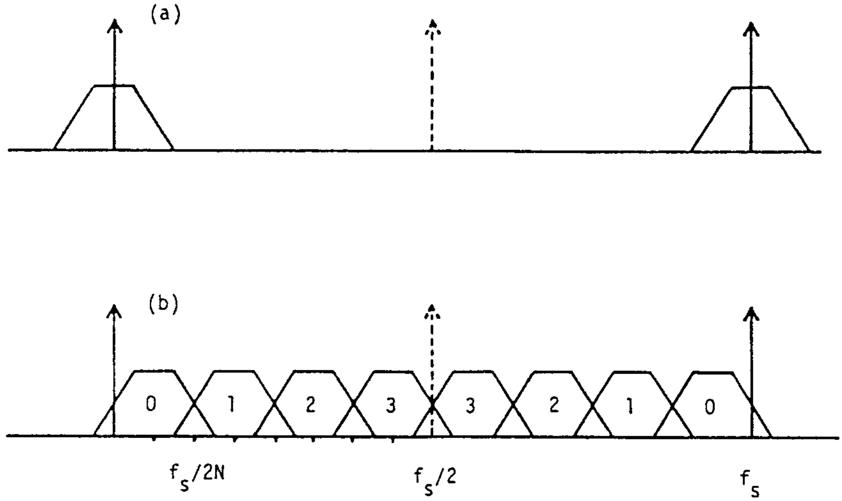

Fig. 10. Pseudo-QMF filters. a) prototype low-pass filter, b) modulated filters.

The performance of the system defined by (5.4–5.6) has been investigated elsewhere [14], our purpose is only to show its computational performance when applied to 2-dimensional signal processing.

Since desirable bandsplitting schemes are separable and that the pseudo-QMF conditions lead to separable filters as well, we will only consider this case. The two dimensional signal is filtered by N pseudo-QMF filters in one dimension, subsampled by N, then filtered by N pseudo-QMF's in the other dimension and subsampled by N in that dimension as well.

Efficient ways to implement such a system are investigated in [14]. Drawing from this work, we use, as an example, the radix-4 version which requires $(3L+9)/4$ multiplications and $L+42$ adds for each set of outputs in one dimension (L is equal to $2NK+1$). We note that since a FFT-like algorithm is used, the efficiency grows with the radix. Assume now $L=25$ and a 256×256 picture to be split into 16 subbands. In Table 1, the various implementations are roughly compared (negecting the borders).

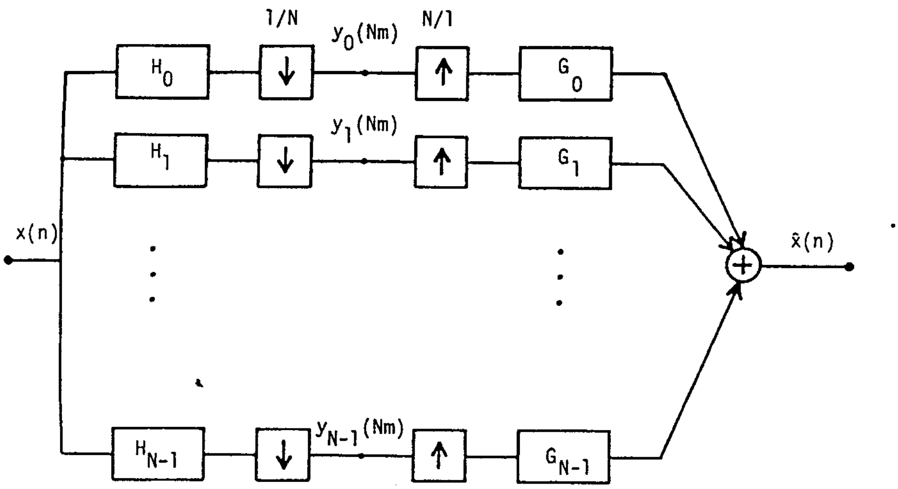

Fig. 11. Parallel decomposition with pseudo-QMF filters.

Table 1

Number of multiples for the coding of a 256 by 256 picture into 16 channels

	Number of multiplies
16 general non-separable filters of dimension 25×25	$41\cdot10^6$
4 stage non-separable QMF with filters of dimensions 17×17, 13×13, 9×9 and 7×7	$9.6\cdot10^6$
16 general separables filters of dimension 25×25	$3.3\cdot10^6$
two stage separable QMF in each dimension with filter length 16 and 8	$1.57\cdot10^6$
radix-4 pseudo-QMF of length 25 in each dimension	$0.688\cdot10^6$

Obviously, substantial gains are attained over the conventional techniques, especially when using the parallel implementation.

6. Concluding remarks

New techniques were proposed for multi-dimensional signal processing. The non-separable case of the QMF filter is a new development, and even if the separable case is only an immediate generalization of the one-dimensional QMF, its introduction into a new field seems useful. The concept of sub-bands with guaranteed reconstruction from the subsampled channels is fundamental enough to justify its analysis in the context of new applications. As an example, the sub-band coding proposed for image compression in [15] forsees information reductions of the order of 5 to 10 without any adaptive coding. Using more sophisticated algorithms for the coding of the channels (or subimages) should lead to much higher compression factors.

When evaluating a coding system, one needs, on the one hand, an objective criteria like quantization noise energy. It was shown in [4] that sub-band coding attains the same theoretical limits as predictive and transform coding. On the other hand, a subjective criteria is needed as well. Results from perceptual psychology indicate a logarithmic sensitivity in the spacial frequency domain [16]. Thus, using coarser quantization for higher frequency bands seems to be a reasonable approach, and investigations are under way in order to prove this hypothesis.

Our emphasis has been on the theoretical side (proof of the generalizations) as well as on the algorithmic side (proposal of fast algorithms for the reversible sub-band coding). Especially the pseudo-QMF approach becomes very attractive, since the computational gain from the one-dimensional case is fully made to contribution. Note that the modulated filter approach (which allows a fast algorithm) leads automatically to the pseudo-QMF (unlike in the one-dimensional case, where the covering of the spectrum is not always necessary as for example in transmultiplexers), since this is the only known class of filters (beside the ideal bandpass) that combines:

1) spectral invariance of the overall system
2) automatic aliasing suppression
3) computational efficiency thanks to a FFT-like approach.

Our hope is that the proposed algorithms will be used in the various fields where they can find applications, both because of the nice properties of the resulting system and because of the reasonable computational load.

Acknowledgements

The author wishes to thank prof. H.J. Nussbaumer for stimulating discussions, the people of the Laboratoire de Traitements des Signaux (EPFL) for comments and help as well as Mrs. A. Schlageter for editing assistance. Special thanks to the 'Fond National Suisse de la Recherche Scientifique' which supported this research and to Prof. B. Wandell for comments on the perceptual aspects.

Appendix

For the general m-dimensional case, we use a convenient vectorial notation. The following definitions are introduced:

$$\boldsymbol{k} = (k_1, k_2, \ldots, k_m) \qquad \boldsymbol{k} \in \mathbb{Z}^m,$$

$$\boldsymbol{\omega} = (\omega_1, \omega_2, \ldots, \omega_m) \qquad \boldsymbol{\omega} \in \mathbb{R}^m,$$

$$\boldsymbol{z} = (z_1, z_2, \ldots, z_m) \qquad \boldsymbol{z} \in \mathbb{C}^m.$$

The unit and zero vectors are defined as $\mathbf{1} = (1, 1, \ldots, 1)$ and $\mathbf{0} = (0, 0, \ldots, 0)$. The following operations will be used:

scalar product: $\mathrm{v}^m \times \mathrm{v}^m \to \mathrm{v}$,

$$\boldsymbol{k} \cdot \boldsymbol{l} = k_1 l_1 + k_2 l_2 + \cdots + k_m l_m.$$

Schur product: $\mathrm{v}^m \times \mathrm{v}^m \to \mathrm{v}^m$,

$$\boldsymbol{k} \circ \boldsymbol{l} = (k_1 l_1, k_2 l_2, \ldots, k_m l_m).$$

General case

Assume a filter $h(\boldsymbol{n})$ with transfer function $H(\boldsymbol{z})$. Its mirror filter follows as $(-1)^{\boldsymbol{n}\cdot\mathbf{1}} \cdot h(\boldsymbol{n})$ with transfer function $H(-\boldsymbol{z})$. The same system as in Fig. 1 is used by replacing 2-dimensional signals with m-dimensional ones. Subsampling and upsampling by a factor 2 corresponds to a modulation by $f(\boldsymbol{n})$ given as:

$$f(\boldsymbol{n}) = 1/2 \cdot (1 + \mathrm{e}^{\mathrm{j}\pi(\boldsymbol{n}\cdot\mathbf{1})}). \tag{A1}$$

Therefore, the reconstructed signal $x(\boldsymbol{n})$ has the following z-transform:

$$\begin{aligned} \hat{X}(\boldsymbol{z}) &= [H(\boldsymbol{z})X(\boldsymbol{z}) + H(-\boldsymbol{z})X(-\boldsymbol{z})] \cdot H(\boldsymbol{z}) - [H(-\boldsymbol{z})X(\boldsymbol{z}) + H(\boldsymbol{z})X(-\boldsymbol{z})] \cdot H(-\boldsymbol{z}) \\ &= [H^2(\boldsymbol{z}) - H^2(-\boldsymbol{z})] \cdot X(\boldsymbol{z}). \end{aligned} \tag{A2}$$

Taking linear phase filters of appropriate length or with a delay will produce an output of the form:

$$\hat{X}(\omega) = d_m \cdot [H^{*2}(\boldsymbol{\omega}) + H^{*2}(\boldsymbol{\omega} + \pi \cdot \mathbf{1}] \cdot X(\boldsymbol{\omega}), \tag{A3}$$

where d_m is a pure m-dimensional delay. Thus, given that the quadrature relation is maintained, the reconstruction is perfect.

Separable case

A m-dimensional input is filtered by 2^m filters and subsampled by 2^m. We call $\boldsymbol{i}_a$ the vector whose components are the base 2 representation of i, thus $i \in [0 \dots 2^m - 1]$ and $\boldsymbol{i}_b$ the vector whose components are (-1) at the power of the corresponding $\boldsymbol{i}_a$ component. The 2^m filters are defined as follows, where $H(\boldsymbol{z})$ is a prototype filter:

$$H_i(\boldsymbol{z}) = H(\boldsymbol{i}_b \circ \boldsymbol{z}), \qquad i = 0, 1, \dots, 2^m - 1. \tag{A4}$$

Following the subsampling by 2^m, the reconstruction after upsampling is performed with interpolation filters $F_i(\boldsymbol{z})$ as:

$$F_i(\boldsymbol{z}) = (-1)^{\boldsymbol{i}_a \cdot \boldsymbol{1}} H(\boldsymbol{z}). \tag{A5}$$

The modulation function is described by:

$$f(\boldsymbol{n}) = 1/2 \left[\sum_{i=0}^{2^m - 1} (-1)^{(\boldsymbol{i}_a \circ \boldsymbol{n}) \cdot \boldsymbol{1}} \right]. \tag{A6}$$

Using the shorthand $\boldsymbol{i}_b \circ \boldsymbol{z} = \boldsymbol{z}_i$, the output $X(\boldsymbol{z})$ equals:

$$\hat{X}(\boldsymbol{z}) = 1/2^m \left[\sum_{i=0}^{2^m - 1} (-1)^{\boldsymbol{i}_a \cdot \boldsymbol{1}} H_i(\boldsymbol{z}) \cdot \sum_{j=0}^{2^m - 1} H_i(\boldsymbol{z}_j) X(\boldsymbol{z}_j) \right]. \tag{A7}$$

But we have:

$$H_i(\boldsymbol{z}) = H(\boldsymbol{z}_i), \quad H_i(\boldsymbol{z}_j) = H(\boldsymbol{i}_b \circ \boldsymbol{j}_b \circ \boldsymbol{z}) = H(\boldsymbol{z}_{i \circ j}) \tag{A8}$$

Thus:

$$\hat{X}(z) = 1/2^m \sum_{j=0}^{2^m - 1} \cdot \sum_{i=0}^{2^m - 1} (-1)^{\boldsymbol{i}_a \cdot \boldsymbol{1}} \cdot H(\boldsymbol{z}_i) \cdot H(\boldsymbol{z}_{i \circ j}) \cdot X(\boldsymbol{z}_j). \tag{A9}$$

Except when $j = 0$, the sum over i should be equal to zero. When $(-1)^{\boldsymbol{j}_a \cdot \boldsymbol{1}}$ is equal to -1, we see that the following products will cancel:

$$H(\boldsymbol{z}_i) \cdot H(\boldsymbol{z}_{i \circ j})$$

$$H(\boldsymbol{z}_{i \circ j}) H(\boldsymbol{z}_{i \circ j \circ j}) = H(\boldsymbol{z}_{i \circ j}) H(\boldsymbol{z}_i). \tag{A10}$$

But for $(-1)^{\boldsymbol{j}_a \cdot \boldsymbol{1}} = 1$, there is no cancellation possible. We show now that separability of the filter is a necessary and sufficient condition. We use the following notation for separable filters:

$$H(\boldsymbol{z}_i) = H_1(z_1) H_2(z_2) \dots H_m(z_m) = \Pi H(\boldsymbol{z}_i). \tag{A11}$$

For j odd, the annulation was shown already in the non-separable case. For j even, there are several ways to show the cancellation. Introducing the vector $\boldsymbol{q}$ which has all ones but a -1 at a position where $\boldsymbol{j}_b$ has also a -1 (the position is arbitrary) we note that:

$$\Pi H(\boldsymbol{z}_i) \cdot \Pi H(\boldsymbol{z}_{i \circ j}) \quad \text{and} \quad \Pi H(\boldsymbol{z}_{i \circ q}) \cdot \Pi H(\boldsymbol{z}_{i \circ q \circ j}), \tag{A12}$$

are equal but of opposite sign in (A9), thus cancelling out. Therefore, the separability is a necessary and sufficient condition for aliasing cancellation in a system with 2^m QMF's and subsampling by 2^m. Note that separability is obviously a sufficient condition, but the purpose of the development following (A9) is to show that it is also a necessary one.

References

[1] A. Croisier, D. Esteban and C. Galand, "Perfect channel splitting by use of interpolation, decimation, tree decomposition techniques", *Int. Conf. on Information Sciences/Systems*, Patras, Aug. 1976, pp. 443–446.

[2] D. Esteban and C. Galand, "Application of quadrature mirror filters to split-band coding", *Int. Conf. on ASSP*, ICASSP 77, Hartford, May 1977, pp. 191–195.

[3] C.R. Galand and H.J. Nussbaumer, "New quadrature mirror filter structures", submitted to ASSP, 1982.

[4] C. Galand, "Codage en sous-bandes: théorie et application à la compression numérique du signal de parole", Thèse, Université de Nice, 1983.

[5] J.D. Johnston, "A filter family designed for use in quadrature mirror filter banks", *Int. Conf. on ASSP*, ICASSP 80, Denver, 1980, pp. 291–294.

[6] L.R. Rabiner and B. Gold, *Theory and Application of Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, 1975.

[7] R.W. Shafer and L.R. Rabiner, "A digital signal processing approach to interpolation", *Proc. IEEE*, Vol. 61, No. 6, June 1973, pp. 692–702.

[8] L.R. Rabiner, J.F. Kaiser, O. Hermann and M.T. Dolan, "Some comparisons between FIR and IIR digital filters", *Bell System Technical Journal*, Vol. 53, No. 2, Feb. 1974, pp. 305–331.

[9] M. Vetterli, "Tree structures for orthogonal transforms and application to the Hadamard transform", *Signal Processing*, Vol. 5, No. 6, Nov. 1983, pp. 473–484.

[10] M.G. Bellanger and J.L. Daguet, "TDM-FDM transmultiplexer: Digital polyphase and FFT", *IEEE Trans. on Communications*, Vol. COM-22, No. 9, Sept. 1974, pp. 1199–1204.

[11] M.J. Narasimha and A.M. Peterson, "Design of a 24-channel transmultiplexer", *IEEE Trans. on ASSP*, Vol. ASSP-27, No. 6, Dec. 1979, pp. 752–762.

[12] H.J. Nussbaumer, "Pseudo QMF filter bank", *IBM Technical Disclosure Bulletin*, Vol. 24, No. 6, Nov. 1981, pp. 3081–3087.

[13] J.H. Rothweiler, "Polyphase quadrature filters—A new subband coding technique", *Proc. 1983 Int. IEEE Conf. on ASSP*, pp. 1280–1283.

[14] H.J. Nussbaumer and M. Vetterli, "Computationally efficient QMF filter banks", to be presented at the 1984 Int. IEEE Conf. on ASSP.

[15] M. Kunt, "Picture coding with the general operator processor", Int. Report, Dept. of Electrical Engineering, Linköping University, Linkoping, Sweden.

[16] J. Hirsch and R. Hylton, "Limits of spacial-frequency discrimination as evidence of neural interpolation", *J. Opt. Soc. Am.*, Vol. 72, No. 10, Oct. 1982, pp. 1367–1374.

Signal Processing 11 (1986) 191
North-Holland

ERRATUM

MULTIDIMENSIONAL SUB-BAND CODING: SOME THEORY AND ALGORITHMS
Signal Processing Vol. 6, No. 2, April 1984, pp. 97–112

Martin VETTERLI (Member EURASIP)
Laboratoire d'Informatique Technique, École Polytechnique Fédérale de Lausanne, 16, Chemin de Bellerive, CH-1007 Lausanne, Switzerland

Received 28 September 1983
Revised 24 November 1983
Erratum received May 1986

This paper contains an erroneous statement on page 101. In the section on two-dimensional separable QMF's, it is stated that, in the case of the separable sub-sampling of Fig. 4(a), "separable filters are necessary and sufficient for aliasing cancellation". What was really meant was that equation (3.2) and similar equations obtained with different signs for the reconstruction filters lead always to the following condition for aliasing cancellation:

$$H(z_1, z_2)H(-z_1, -z_2) - H(z_1, -z_2)H(-z_1, z_2) = 0. \tag{1}$$

It turns out that condition (1) is weaker than separability, as shown by Woods [1] whom we thank for bringing this fact to our attention.

Reference

[1] J.W. Woods and S. O'Neil, "Sub-band coding of images", *IEEE Trans. Acoust. Speech Signal Process.*, to appear.