Design of High-Performance Filter Banks for Image Coding

Di Xu and Michael D. Adams
Dept. of Elec. and Comp. Engineering, University of Victoria
PO Box 3055, STN CSC, Victoria, BC, V8W 3P6, Canada
dixu@ece.uvic.ca and mdadams@ece.uvic.ca

***Abstract*—*An optimization-based method for the design of high-performance separable filter banks for image coding is proposed. This method yields linear-phase perfect-reconstruction systems with high coding gain, good frequency selectivity, and certain prescribed vanishing-moment properties. Several filter banks designed with the proposed method are presented and shown to work extremely well for image coding, outperforming the well-known 9/7 filter bank from the JPEG-2000 standard in most cases.**

***Keywords*—*optimal filter-bank design, wavelet transforms, lifting framework, image coding.**

I. Introduction

Filter banks have proven to be an extremely valuable tool for image coding applications [1], [2], [3]. In order to be effective in such applications, however, a filter bank must typically have a number of desirable characteristics, namely, perfect reconstruction (PR), linear phase, high coding gain [4], good frequency selectivity, and certain vanishing-moment properties. To design filter banks having all of the preceding characteristics is a challenging task. In this paper, we propose a new optimal design method based on [5] that yields high-performance separable filter banks with all of the aforementioned desirable characteristics.

The remainder of this paper is structured as follows. Section II briefly introduces some of the notation used herein. Then, Section III introduces some essential background related to filter banks. Our new design method is presented in detail in Section IV, and Section V provides some examples of filter banks produced using our proposed method and evaluates their performance for image coding. Finally, Section VI concludes the paper with a summary of our work and some closing remarks.

II. Notation and Terminology

Before proceeding further, a brief digression is in order regarding the notation used herein. The set of integers is denoted as $\mathbb{Z}$. Matrices and vectors are denoted by uppercase and lowercase boldface letters, respectively. For matrix multiplication, we define the product notation as $\prod_{k=M}^{N} \boldsymbol{A}_k \triangleq \boldsymbol{A}_N \boldsymbol{A}_{N-1} \ldots \boldsymbol{A}_{M+1} \boldsymbol{A}_M$. For a positive semi-definite matrix $\boldsymbol{A}$, we denote its square root (e.g., as defined in [6]) as $\boldsymbol{A}^{1/2}$. For a function f, its gradient and transposed gradient with respect to $\boldsymbol{x}$ are denoted as $\nabla_{\boldsymbol{x}} f$ and $\nabla_{\boldsymbol{x}}^T f$, respectively, where the subscript $\boldsymbol{x}$ may be omitted when clear from the context. The transfer function, impulse response, and frequency response of a filter H are denoted as $H(z)$, h, and $\hat{h}$, respectively.

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

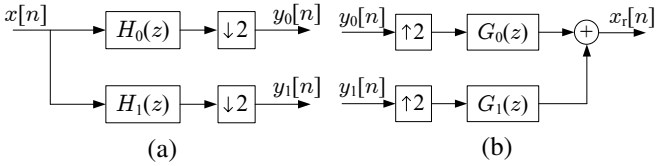

Fig. 1. The canonical form of a 1-D 2-channel filter bank. (a) Analysis side and (b) synthesis side.

III. Filter Banks

A one-dimensional (1-D) two-channel filter bank has the canonical form shown in Fig. 1. Such a filter bank consists of lowpass and highpass analysis filters H_0 and H_1, lowpass and highpass synthesis filters G_0 and G_1, as well as downsamplers and upsamplers. In our design method, rather than representing a filter bank in its canonical form, we instead use the lifting framework. The lifting realization [7] of a 1-D two-channel filter bank is shown in Fig. 2, where there are 2λ lifting filters $\{\mathrm{F}_k\}_{k=0}^{2\lambda-1}$. Without loss of generality, we assume that only $F_0(z)$ and $F_{2\lambda-1}(z)$ may be identically zero. The analysis filter transfer functions can be calculated from the lifting parameterization as

$$H_0(z) = H_{0,0}(z^2) + zH_{0,1}(z^2) \quad \text{and} \tag{1a}$$

$$H_1(z) = H_{1,0}(z^2) + zH_{1,1}(z^2), \tag{1b}$$

where $\begin{bmatrix} H_{0,0}(z) & H_{0,1}(z) \\ H_{1,0}(z) & H_{1,1}(z) \end{bmatrix} = \prod_{k=0}^{\lambda-1} \left(\begin{bmatrix} 1 & F_{2k+1}(z) \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ F_{2k}(z) & 1 \end{bmatrix} \right)$. The transfer functions of the synthesis filters can be similarly derived. Since the synthesis filters are completely determined by the analysis filters, we focus primarily on the analysis side of the filter bank in what follows.

The use of the lifting realization has a number of advantages over the canonical form. The key benefit, however, is that the PR condition is automatically satisfied. Additionally, the linear phase requirement can be easily met by choosing the lifting filters to have certain symmetries, as we shall see shortly. Since the PR and linear-phase conditions can be imposed via the lifting framework, there is no need for explicit optimization constraints to ensure that these conditions hold. This greatly reduces the complexity of the subsequent optimal-design problem.

As suggested above, the linear-phase condition can be easily imposed through a clever choice of the lifting filters $\{\mathrm{F}_k\}_{k=0}^{2\lambda-1}$.

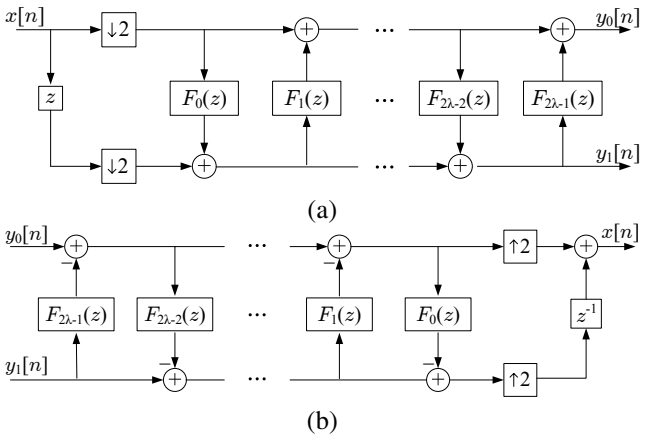

Fig. 2. The lifting realization of a 1-D two-channel filter bank. (a) Analysis side and (b) synthesis side.

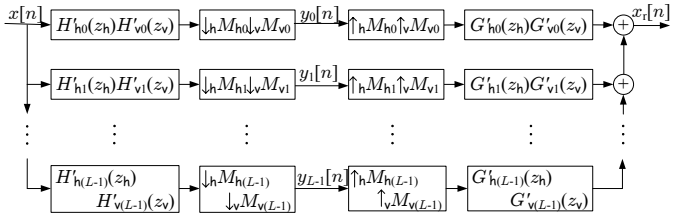

Fig. 3. The equivalent L-channel nonuniform filter bank associated with the N-level tree-structured filter bank.

It can be shown [8] that if the $\{F_k(z)\}_{k=0}^{2\lambda-1}$ are chosen to be of either of the following two forms, then the resulting filter bank will have linear phase:

$$F_k(z) = \begin{cases} \sum_{i=0}^{(L_k-2)/2} a_{k,i}(z^{-i} + z^{i+1}) & \text{for even } k \\ \sum_{i=0}^{(L_k-2)/2} a_{k,i}(z^{-i-1} + z^{i}) & \text{for odd } k \end{cases} \tag{2a}$$

or

$$F_k(z) = \begin{cases} -1 & \text{for } k = 0 \\ \frac{1}{2} + \tilde{F}_1(z) & \text{for } k = 1 \\ \tilde{F}_k(z) & \text{for } k \geq 2, \end{cases} \tag{2b}$$

where $\tilde{F}_k(z) = \sum_{i=1}^{(L_k-1)/2} \tilde{a}_{k,i}(z^{-i} + z^{i})$, L_k is the length of the lifting filter F_k, and the $\{L_k\}$ are all even and all odd in (2a) and (2b), respectively.

To construct a 2-D filter bank from a 1-D filter bank, we simply apply the 1-D filter bank in each of the two dimensions in succession. This results in a separable four-channel 2-D filter bank. Furthermore, in practice, we usually apply the 2-D filter bank in an N-level tree structure, decomposing the lowest-frequency band at each level in the tree. The resulting N-level tree-structured filter bank can be equivalently expressed (via the noble identities [9]) in the form of an L-channel nonuniform filter bank as shown in Fig. 3, where $L = 3N + 1$. In the diagram, the $\{\mathrm{H}'_{\mathsf{h}k}\}$, $\{\mathrm{H}'_{\mathsf{v}k}\}$, $\{\mathrm{G}'_{\mathsf{h}k}\}$, and $\{\mathrm{G}'_{\mathsf{v}k}\}$ denote the equivalent horizontal analysis, vertical analysis, horizontal synthesis, and vertical synthesis filters, respectively, and the $\{M_{\mathsf{h}k}\}$ and $\{M_{\mathsf{v}k}\}$ denote horizontal and vertical upsampling/downsampling factors, respectively. As a matter of notation, the subscripts "h" and "v" on $\downarrow$, $\uparrow$, and z indicate correspondences with the horizontal and vertical directions, respectively.

Since we have elected to use a lifting parameterization for our later optimization problem formulation, we need to relate the analysis-filter frequency responses, vanishing-moment properties, and coding gain to the lifting-filter coefficients. In the case of the frequency responses and vanishing-moment properties, these relationships can be derived in a straightforward manner via (1) and (2). In the case of the coding gain, the relationship can be derived as follows. First, we determine the filters of the equivalent nonuniform filter bank (as shown in Fig. 3) via (1), (2), and the noble identities. Then, one can show [4] that the coding gain G_{SBC} is given by

$$G_{\mathsf{SBC}} = \prod_{k=0}^{L-1} \left(\frac{\alpha_k}{A_k B_k}\right)^{\alpha_k}, \quad \text{where} \tag{3a}$$

$$A_k = \sum_{m\in\mathbb{Z}} h'_{\mathsf{h}k}(m) \sum_{n\in\mathbb{Z}} h'_{\mathsf{v}k}(n) \sum_{p\in\mathbb{Z}} h'_{\mathsf{h}k}(p) \sum_{q\in\mathbb{Z}} h'_{\mathsf{v}k}(q) r(m-p, n-q),$$

$B_k = \alpha_k \sum_{m\in\mathbb{Z}} g'^2_{\mathsf{h}k}(m) \sum_{n\in\mathbb{Z}} g'^2_{\mathsf{v}k}(n)$, $\alpha_k = (M_{\mathsf{h}k} M_{\mathsf{v}k})^{-1}$, and r is the normalized autocorrelation function (NACF) of the input image. In practice, the NACF is chosen, depending on the most appropriate image model, as follows:

$$r(x,y) = \begin{cases} \rho^{|x|+|y|} & \text{for separable model} \\ \rho^{\sqrt{x^2+y^2}} & \text{for isotropic model,} \end{cases} \tag{3b}$$

where ρ is the correlation coefficient (typically, $0.9 \leq \rho \leq 0.95$). In all of our subsequent work, we assume $\rho = 0.95$ for both models.

IV. Design Method

As indicated earlier, we seek to design filter banks having numerous desirable characteristics, namely, PR, linear phase, high coding gain, good frequency selectivity, and certain prescribed vanishing-moment properties. Since the PR and linear-phase properties are structurally imposed via the lifting framework, we need not consider them further. Thus, the design problem at hand reduces to one explicitly involving only coding gain, frequency selectivity, and vanishing-moment properties.

Now, let us consider the formulation of the design problem as an optimization. Let $\boldsymbol{x}$ denote the vector of independent lifting-filter coefficients. We choose G, a measure related to coding gain, as our objective function to maximize. Let G_{sep} and G_{iso} denote the coding gain (in dB) obtained from (3) using the separable and isotropic models, respectively. In our work, we consider three possible choices for G as given by

$$G(\boldsymbol{x}) = \begin{cases} G_{\mathsf{sep}}(\boldsymbol{x}) & \text{separable only} \quad (4a) \\ G_{\mathsf{iso}}(\boldsymbol{x}) & \text{isotropic only} \quad (4b) \\ \min\{G_{\mathsf{sep}}(\boldsymbol{x}), G_{\mathsf{iso}}(\boldsymbol{x})\} & \text{joint.} \quad (4c) \end{cases}$$

That is, we consider the maximization of each of the separable and isotropic coding gains individually as well as the maximization of the minimum of both coding gains together. This

last case is motivated by the observation that many images are nonstationary, exhibiting both separable and isotropic behavior in different regions. Thus, we might suspect that there is an advantage to having both coding gains high.

The remaining filter bank properties are handled as constraints. To quantify the frequency selectivity of the analysis filters, we employ a stopband-energy measure. In particular, we define the stopband energy of the analysis filter H_k as

$$b_k(\boldsymbol{x}) \triangleq \int_{S_k} |\hat{h}_k(\omega, \boldsymbol{x})|^2 d\omega, \quad k \in \{0, 1\}, \tag{5}$$

where $S_0 = [\pi - \omega_b, \pi]$, $S_1 = [0, \omega_b]$, $\hat{h}_k$ denotes the frequency response of H_k, and ω_b denotes the stopband width of the analysis filters. In passing, we note that a choice of $\omega_b = \frac{3\pi}{8}$ is made in our work. To facilitate the introduction of moment constraints, we define the moment-constraint functions

$$c_k(\boldsymbol{x}) \triangleq \|\boldsymbol{m}_k(\boldsymbol{x})\|, \quad k \in \{1, 2, \ldots, n\}, \tag{6}$$

where $\boldsymbol{m}_k$ is a ν_k-dimensional vector function with its elements corresponding to moments of interest (i.e., moments that are to be constrained). Each $\boldsymbol{m}_k$ may contain only one moment (i.e., $\nu_k = 1$) or a group of several moments (i.e., $\nu_k > 1$). In this way, moments can controlled either individually or jointly.

Combining (4), (5), and (6), we obtain the following abstract optimization problem to be solved:

$$\begin{aligned} \text{maximize} \quad & G(\boldsymbol{x}) && \text{(7a)} \\ \text{subject to:} \quad & b_k(\boldsymbol{x}) \leq \varepsilon_k, \ \ k \in \{0, 1\} \ \text{ and} && \text{(7b)} \\ & c_k(\boldsymbol{x}) \leq \gamma_k, \ \ k \in \{1, 2, \ldots, n\}, && \text{(7c)} \end{aligned}$$

where the $\{\varepsilon_k\}$ and $\{\gamma_k\}$ are strictly positive tolerances for the stopband-energy and moment constraints. Since the $\{\gamma_k\}$ are strictly positive, we do not impose exact vanishing-moment conditions; rather, we only ensure that the moments in question are very small. In a practical sense, there is no significant disadvantage to allowing moments to deviate slightly from zero, as true vanishing-moment properties are usually lost during implementation anyhow, due to finite-precision effects. This more relaxed form of moment constraint is, in fact, advantageous, resulting in a larger feasible region and therefore potentially better designs.

Unfortunately, the abstract optimization problem given by (7) is highly nonlinear and somewhat difficult to solve directly. For this reason, rather than directly solving this problem, we instead employ an approach based on the iterative solution of a reduced-order problem.

To simplify our initial discussion, we consider our optimization-solution technique for the first two possibilities for the objective function given by (4a) and (4b). In these cases, we use a Taylor series in the variable $\boldsymbol{\delta}$ to approximate the G, $\{b_k\}$, and $\{c_k\}$ about a fixed operating point $\boldsymbol{x}$, such that the approximation of G is linear, and the approximations of the $\{b_k\}$ and $\{c_k\}$ are at most quadratic. This reduced-order problem can then be expressed as a second-order cone programming (SOCP) problem [10], as explained shortly, and solved relatively easily using software such as SeDuMi [11] (which is employed in our work). Since the various Taylor series involved are only good approximations in the neighbourhood of the operating point $\boldsymbol{x}$ (i.e., for small $\boldsymbol{\delta}$), we must impose an additional constraint on $\boldsymbol{\delta}$ to force a solution to be in this neighbourhood. As the neighbourhood of the solution is limited, we then update our operating point to $\boldsymbol{x} + \boldsymbol{\delta}$ and repeat this process in order to permit solutions farther away from the current operating point. The optimization terminates when either the increase in G or the change in the operating point becomes sufficiently small.

Let us now consider the formulation of the SOCP problem mentioned above. A SOCP problem has the general form:

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^T \boldsymbol{x} && \text{(8a)} \\ \text{subject to:} \quad & \|\boldsymbol{A}_k^T \boldsymbol{x} + \boldsymbol{c}_k\| \leq \boldsymbol{b}_k^T \boldsymbol{x} + d_k, \quad k \in \{1, 2, \ldots, K\}. && \text{(8b)} \end{aligned}$$

It can be shown that the abstract optimization problem in (7) can be converted to the following form:

$$\begin{aligned} \text{maximize} \quad & \nabla^T G(\boldsymbol{x}) \boldsymbol{\delta} && \text{(9a)} \\ \text{subject to:} \quad & \|\boldsymbol{Q}_k^{1/2}(\boldsymbol{x})\boldsymbol{\delta} + \boldsymbol{q}_k(\boldsymbol{x})\| \leq \varepsilon_k - b_k(\boldsymbol{x}) \\ & \qquad + \boldsymbol{q}_k^T(\boldsymbol{x})\boldsymbol{q}_k(\boldsymbol{x}), \ \ k \in \{0, 1\}, && \text{(9b)} \\ & \|\nabla^T \boldsymbol{m}_k(\boldsymbol{x})\boldsymbol{\delta} + \boldsymbol{m}_k(\boldsymbol{x})\| \leq \gamma_k, \\ & \qquad k \in \{1, 2, \ldots, n\}, \ \text{and} && \text{(9c)} \\ & \|\boldsymbol{\delta}\| \leq \beta, && \text{(9d)} \end{aligned}$$

where

$$\boldsymbol{Q}_k(\boldsymbol{x}) = \int_{S_k} \nabla_{\boldsymbol{x}} \hat{h}_k(\omega, \boldsymbol{x}) \nabla_{\boldsymbol{x}}^T \hat{h}_k(\omega, \boldsymbol{x}) d\omega,$$

$$\boldsymbol{q}_k(\boldsymbol{x}) = \boldsymbol{Q}_k^{-1/2}(\boldsymbol{x}) \int_{S_k} \hat{h}_k(\omega, \boldsymbol{x}) \nabla_{\boldsymbol{x}}^T \hat{h}_k(\omega, \boldsymbol{x}) d\omega,$$

and $\boldsymbol{\delta}$ is a small perturbation from the operating point $\boldsymbol{x}$. Clearly, the preceding problem is of the form of (8), where the optimization is performed with respect to the variable $\boldsymbol{\delta}$.

To complete our discussion, we must now consider the third possibility for the objective function (4c). In this case, the reduced-order optimization problem can be expressed as

$$\begin{aligned} \text{maximize} \quad & t \\ \text{subject to:} \quad & G_{\mathsf{sep}}(\boldsymbol{x}) + \nabla^T G_{\mathsf{sep}}(\boldsymbol{x})\boldsymbol{\delta} \geq t, \\ & G_{\mathsf{iso}}(\boldsymbol{x}) + \nabla^T G_{\mathsf{iso}}(\boldsymbol{x})\boldsymbol{\delta} \geq t, \ \ \text{and} \\ & \text{(9b-9d)}, \end{aligned}$$

where t is an auxiliary variable. By introducing the new augmented vector variable $\tilde{\boldsymbol{\delta}} \triangleq \begin{bmatrix} t & \boldsymbol{\delta} \end{bmatrix}^T$, the problem at hand can be trivially cast as a SOCP problem with respect to $\tilde{\boldsymbol{\delta}}$.

A. Design Parameter Selection

Having introduced our design method, we now comment briefly on the selection of several key design parameters. In the case of the frequency-selectivity constraints, an appropriate choice of tolerances $\{\varepsilon_k\}$ is critical to achieving a good design. Based on experimentation, for a stopband width of $\omega_b = \frac{3\pi}{8}$,

Table 1. Characteristics of a subset of the test images

Image	Size, Precision†	Model‡	Description
gold	720×576, 8	separable	houses and countryside
target	512×512, 8	—	patterns and textures
sar2	800×800, 12	isotropic	synthetic aperture radar

‡best-fitting image model (e.g., separable or isotropic)
†precision in bits per sample

a choice of $\varepsilon_k \in [0.02, 0.14]$ is highly effective. The selection of the moment constraints (i.e., the $\{\boldsymbol{m}_k\}$ and $\{\gamma_k\}$) is also quite important. Since filter banks satisfying (2a) and (2b), respectively, have all of their even and all of their odd moments automatically vanish, only the remaining moments need to be considered. Again, based on our experiments, good designs satisfying (2a) are obtained when the moments are constrained such that the norm of the vector of zeroth dual and primal moments is less than $2 \cdot 10^{-5}$. Lastly, since the optimal solution obtained depends on the choice of the initial point $\boldsymbol{x}$, the quality of the design can be improved by finding multiple solutions and then selecting the best one. As a practical matter, we found that an effective strategy in this regard is to consider many initial points with lifting-filter coefficients of magnitude 2 or less, as the best designs typically have coefficients in this range.

V. Experimental Results

Having introduced our design method, we would now like to present some experimental results related to our method. Before proceeding, however, we first introduce some details concerning our test environment. In all of our experiments, for test data, all of the twenty-six reasonably-sized continuous-tone grayscale images from the JPEG-2000 test set [12] were employed. Often, in what follows, we will focus our attention on a representative subset of these images, namely, the gold, sar2, and target images, whose characteristics are given in Table 1. To evaluate the performance of filter banks for image coding, the EZW [2], SPIHT [3], and MIC [13] codecs were employed, all of which utilize reversible integer-to-integer transforms. Since similar results were obtained with all three codecs, however, only results from the MIC codec are presented herein. In all of our coding experiments, a six-level wavelet decomposition was employed. In what follows, we use the notation l_0/l_1 to indicate that a filter bank has lowpass and highpass analysis filters of lengths l_0 and l_1, respectively. We also identify the lengths $\{L_k\}$ of the lifting filters (i.e., the lifting configuration) using the notation $\{L_0, L_1, \ldots\}$.

A. Choice of Objective Function

As explained earlier, our design method allows for three possible objective functions as given by (4). We also suggested that, of these possibilities, (4c) might be most desirable. Now, we experimentally show this to be the case.

To begin, for each of several different lifting configurations, we used our method to design three filter banks, one for each of the three choices of objective function in (4). In so doing, we were able to make an interesting observation. As it turns out, in all of our tests, optimizing with respect to either of the objective functions (4b) or (4c) always led to the same optimal design. This is due to the fact that, for filter banks with good frequency selectivity, G_{sep} always tends to be greater than G_{iso}. Therefore, maximizing (4c) is equivalent (in a practical sense) to maximizing (4b). With the above observation in mind, we combine the cases of (4b) and (4c) in the remainder of this discussion.

Table 2. Characteristics of the filter banks designed using different objective functions for the 9/7 and 6/14 cases

Transform	$G^{\ddagger}_{\mathsf{sep}}$	$G^{\ddagger}_{\mathsf{iso}}$	$b^{\dagger}_0$	$b^{\dagger}_1$	Van. Mom.*
9/7-ref	14.9734	12.1781	0.0628	0.0347	4, 9.560
9/7-sep	14.9735	12.1781	0.0628	0.0347	2, 5.79E-5
9/7-iso/jnt	14.9326	12.1809	0.0570	0.0351	2, 0.0041
6/14-ref	14.8801	12.0457	0.0340	0.0241	3, 3.000
6/14-sep	15.0912	11.9285	0.0252	0.0131	1, 0.0483
6/14-iso/jnt	14.9766	12.0738	0.0212	0.0229	1, 0.0643

*index and magnitude of the first dual moment with magnitude greater than $2 \cdot 10^{-5}$
‡coding gain (in dB) for a six-level decomposition
†stopband energy as defined by (5)

Herein, we consider two sets of optimal designs with analysis filter lengths of 9/7 and 6/14 (which correspond to lifting configurations of $\{2, 2, 2, 2\}$ and $\{1, 3, 5\}$, respectively). The characteristics of these filter banks are shown in Table 2, where the suffixes "sep", "iso", and "jnt" are used to designate the optimal designs obtained using the objective functions (4a), (4b), and (4c), respectively. For comparison purposes, in the case of each of the two sets, we also consider a filter bank (having an identical lifting configuration) that is already known to be effective for image coding. In the table, the reference filter bank is designated by the suffix "ref".

Having produced several sets of optimal filter banks as described above, we then proceeded to compare the coding performance of the filter banks in each set. For each of the filter banks in each set, we compressed all of the twenty-six test images in a lossy manner at several bit rates. In each case, we measured the relative difference between the distortions (in PSNR) obtained with our design and the appropriate reference (i.e., "ref") filter bank. The results are summarized in statistical form in Table 3(a). In particular, we provide the mean and median relative differences (with positive values corresponding to our designs outperforming the reference filter bank). As well, we indicate the percentage of cases in which our filter bank outperforms the reference filter bank. From Table 3(a), we can see that, in both the 9/7 and 6/14 cases, designs based on the (joint) objective function (4c) (designated by the suffix "jnt") have better performance relative to the reference filter bank than the designs based on the (separable-only) objective function (4a). In Table 3(b), we provide the actual distortions obtained for three representative images. Here, we can see that the filter banks with the jointly-highest coding gains perform better overall for all three images, in spite of the images having significantly different autocorrelation models. Clearly,

Table 3. Lossy compression results for the filter banks designed using different objective functions. (a) Summary statistical results over all twenty-six test images and five bit rates. (b) Specific results for three images.

(a)

Transform	Mean (%)	Median (%)	Outperform (%)
9/7-sep	-0.0049	-0.0001	46.15
9/7-iso/jnt	0.1488	0.1070	87.69
6/14-sep	-0.5848	-0.5112	20.77
6/14-iso/jnt	0.0331	0.0279	61.54

(b)

Image	CR†	PSNR (dB)			
		9/7-sep	9/7-iso/jnt	6/14-sep	6/14-iso/jnt
gold	8	36.76	36.88	36.77	36.96
	16	33.77	33.84	33.50	33.78
	32	31.26	31.27	31.07	31.15
	64	29.15	29.15	28.82	28.83
	128	27.33	27.32	27.06	27.26
target	8	41.48	41.59	40.94	41.68
	16	33.55	33.55	32.79	32.91
	32	27.08	27.19	27.37	27.00
	64	22.70	22.84	22.39	22.43
	128	19.13	19.14	18.49	18.26
sar2	8	30.33	30.35	30.24	30.39
	16	26.61	26.62	26.49	26.49
	32	24.69	24.70	24.64	24.78
	64	23.55	23.55	23.48	23.46
	128	22.73	22.73	22.59	22.68

†compression ratio

Table 5. Lossy compression results for the various filter banks. (a) Summary statistical results over all twenty-six test images and five bit rates. (b) Specific results for three images.

(a)

Transform	Mean (%)	Median (%)	Outperform (%)
9/7	0.1488	0.1070	87.69
9/11	0.5371	0.0554	59.23
13/11	0.1863	0.0872	74.62
17/11	0.5804	0.2422	77.69
13/15	0.5936	0.1633	68.46

(b)

Image	CR†	PSNR (dB)					
		9/7-J	9/7	9/11	13/11	17/11	13/15
gold	8	36.75	36.88	37.34	36.85	37.17	**37.39**
	16	33.75	33.84	**34.00**	33.76	33.91	33.95
	32	31.23	31.27	**31.35**	31.24	31.32	31.27
	64	29.16	29.15	29.24	29.17	29.17	**29.25**
	128	27.32	27.32	**27.39**	27.35	27.37	27.34
target	8	41.46	41.59	42.92	42.12	42.81	**43.11**
	16	33.54	33.55	33.47	33.83	**34.00**	33.45
	32	27.07	27.19	26.65	27.84	**27.88**	26.84
	64	22.70	22.84	22.35	23.11	**23.19**	22.42
	128	19.16	19.14	18.86	19.35	**19.46**	18.94
sar2	8	30.32	30.35	30.30	30.33	30.33	**30.37**
	16	26.61	26.62	26.59	**26.62**	26.60	26.57
	32	24.69	**24.70**	24.65	24.69	24.69	24.70
	64	23.55	**23.55**	23.52	23.54	23.53	23.54
	128	22.73	22.73	22.70	**22.74**	22.74	22.67

†compression ratio

the above experiment demonstrates that there is a potential benefit to jointly optimizing both of the separable and isotropic coding gains (i.e., using the objective function (4c)).

B. Design Examples

To demonstrate the effectiveness of our design method, we now consider the lossy and lossless coding performance of five filter banks constructed with our method. For all of these designs, the objective function (4c) was employed, as this was earlier determined to be the best choice. For comparison purposes, we also consider the well-known 9/7 filter bank from JPEG 2000 [1], which we henceforth refer to by the name 9/7-J (in order to distinguish it from other filter banks having the same analysis-filter lengths). Several characteristics of our five optimal designs as well as the 9/7-J filter bank are shown in Table 4. Due to space constraints, the lifting-filter coefficients for our optimal designs are not presented here, but can be obtained from [14].

First, we consider the lossy coding performance of the various filter banks. To evaluate lossy coding performance, each filter bank was used to compress all of the test images at several bit rates. Then, we measured the relative difference between the distortions (in PSNR) obtained with each of our optimal designs and the reference 9/7-J filter bank. The results are summarized in statistical form in Table 5(a). In particular, we provide the mean and median relative differences (with positive values corresponding to our designs outperforming the reference 9/7-J filter bank). As well, we indicate the percentage of cases in which our optimal design outperforms the 9/7-J filter bank. From these results, it is clear that all of our optimal designs outperform the 9/7-J filter bank in the majority of cases. For example, our 9/7 optimal design outperforms the 9/7-J filter bank 87.69% of the time. Our four other designs outperform the 9/7-J filter bank by margins ranging from about 59 to 78%. The above results are extremely encouraging, given that the 9/7-J filter bank is well known for its exceptional lossy-coding performance. In Table 5(b), we provide the actual PSNR results obtained for a representative subset of the test images, where the best result is highlighted in each case. From the table, we can see that, even for images with different statistical properties (such as the three considered here), our optimal designs outperform the 9/7-J filter bank, sometimes by a margin as high as 1.65 dB. Lastly, we would like to note that our optimal designs also lead to good subjective image quality, comparable to that of the 9/7-J filter bank. In Fig. 4, we provide an example of lossy image reconstructions obtained with the various filter banks. One can see that the quality of the reconstructions produced by our optimal designs is comparable to that obtained with the 9/7-J filter bank.

As an aside, we would like to note that, as can be seen from Table 4, our 9/7 design and the 9/7-J filter bank have the same lifting configuration. While the 9/7-J filter bank has four dual and four primal vanishing moments, our 9/7 design has only two dual and two primal near-vanishing moments and slightly higher isotropic coding gain as well. With our design approach, by reducing the number of constrained moments and relaxing the requirement that moments be exactly zero, we are able to gain additional freedom, which ultimately allows higher-performance filter banks to be constructed.

In our work, we also evaluated the lossless coding performance of the various filter banks. Each of the filter banks was

Table 4. Characteristics of various filter banks

Transform	$\{L_k\}$	$G_{\mathsf{sep}}^{\ddagger}$	$G_{\mathsf{iso}}^{\ddagger}$	$b_0^{\dagger}$	$b_1^{\dagger}$	Van. Mom.*	$\hat{h}_0(0)$	$\hat{h}_0(\pi)$	$\hat{h}_1(0)$	$\hat{h}_1(\pi)$
9/7-J	{2,2,2,2}	14.9734	12.1781	0.0628	0.0347	4, 9.560	1.2302	2.31E-9	3.94E-9	1.6258
9/7	{2,2,2,2}	14.9326	12.1809	0.0570	0.0351	2, 0.0041	1.2504	2.11E-6	-9.07E-6	1.5995
9/11	{4,2,2}	14.9283	12.1118	0.1110	0.0430	2, 0.2755	1.2394	-5.62E-6	-9.43E-6	1.6137
13/11	{4,2,2,2}	15.0406	12.2059	0.0299	0.0268	2, 0.0680	1.1986	2.15E-5	-1.93E-5	1.6687
17/11	{2,2,4,4}	15.1174	12.2177	0.0306	0.0275	2, 0.3367	1.1869	9.85E-6	4.87E-6	1.6851
13/15	{6,2,2}	14.6409	12.0741	0.0940	0.0354	2, 0.1685	1.3465	-5.41E-6	-9.42E-6	1.4854

*index and magnitude of the first dual moment with magnitude greater than $2 \cdot 10^{-5}$
‡coding gain (in dB) for a six-level decomposition
†stopband energy as defined by (5)

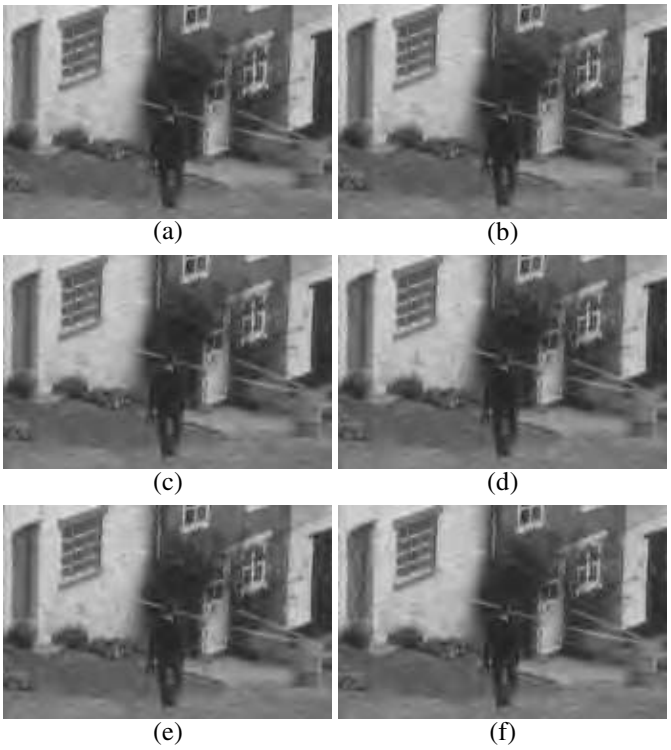

Fig. 4. Parts of the lossy reconstructions obtained after coding the gold image at a compression ratio of 32 using the (a) 9/7-J, (b) 9/7, (c) 9/11, (d) 13/11, (e) 17/11, and (f) 13/15 filter banks.

Table 6. Lossless compression results for the various filter banks

Image	NBR					
	9/7-J	9/7	9/11	13/11	17/11	13/15
gold	0.5673	0.5666	**0.5630**	0.5652	0.5657	0.5644
target	0.3173	0.3086	**0.2949**	0.2964	0.3130	0.2975
sar2	0.6350	0.6351	**0.6349**	0.6353	0.6353	0.6352
Mean†	0.4880	0.4874	**0.4771**	0.4787	0.4870	0.4828

†mean taken over all twenty-six test images

used to losslessly compress all of the twenty-six test images. The results are shown in Table 6. In particular, we provide the normalized bit rate (NBR) (i.e., the reciprocal of compression ratio) for three images as well as the mean taken over all of the twenty-six test images. Evidently, all of our filter banks perform better overall than the 9/7-J for lossless coding, with the 9/11 design being the best.

VI. Conclusions

In this paper, we proposed a method for the design of filter banks with properties desirable for image coding (i.e., PR, linear phase, high coding gain, good frequency selectivity, and certain prescribed moment properties). Several examples of filter banks constructed using our method were presented, and shown to be highly effective for image coding. In particular, our optimal designs outperformed the well-known 9/7-J filter bank (from the JPEG-2000 standard) for both lossy and lossless compression, an impressive feat given that the 9/7-J filter bank is known for its exceptional lossy coding performance.

References

[1] *ISO/IEC 15444-1: Information technology—JPEG 2000 image coding system—Part 1: Core coding system*, 2000.

[2] J. M. Shapiro, "Embedded image coding using zerotrees of wavelet coefficients," *IEEE Trans. on Signal Processing*, vol. 41, no. 12, pp. 3445–3462, Dec. 1993.

[3] A. Said and W. A. Pearlman, "A new fast and efficient image codec based on set partitioning in hierarchical trees," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–250, June 1996.

[4] J. Katto and Y. Yasuda, "Performance evaluation of subband coding and optimization of its filter coefficients," in *Proc. of SPIE Visual Communications and Image Processing*, vol. 1605, Nov. 1991, pp. 95–106.

[5] Y. Chen, M. D. Adams, and W.-S. Lu, "Design of optimal quincunx filter banks for image coding," in *Proc. of IEEE International Symposium on Circuits and Systems*, May 2006, pp. 2041–2044.

[6] P. Lancaster and M. Tismenetsky, *The Theory of Matrices: with Applications*, 2nd ed. San Diego, CA, USA: Academic Press, 1985.

[7] W. Sweldens, "The lifting scheme: A custom-design construction of biorthogonal wavelets," *Applied and Computational Harmonic Analysis*, vol. 3, no. 2, pp. 186–200, 1996.

[8] M. D. Adams and R. K. Ward, "Symmetric-extension-compatible reversible integer-to-integer wavelet transforms," *IEEE Trans. on Signal Processing*, vol. 51, no. 10, pp. 2624–2636, Oct. 2003.

[9] P. P. Vaidyanathan, *Multirate Systems and Filter Banks*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.

[10] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Applications of second-order cone programming," *Linear Algebra and its Applications*, vol. 248, pp. 193–228, Nov. 1998.

[11] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11–12, pp. 625–653, 1999.

[12] *JPEG-2000 test images*, ISO/IEC JTC 1/SC 29/WG 1 N 545, July 1997.

[13] M. D. Adams, "ELEC 545 project: A wavelet-based lossy/lossless image compression system," Dept. of Elec. and Comp. Engineering, University of British Columbia, Vancouver, BC, Canada, Apr. 1999.

[14] "Michael Adams' home page," http://www.ece.uvic.ca/~mdadams, 2006.