Noname manuscript No.
(will be inserted by the editor)

Two-dimensional digital filters with sparse coefficients

Wu-Sheng Lu · **Takao Hinamoto**

Received: date / Accepted: date

Abstract Is sparsity an issue in 2-D digital filter design problems to explore and why is it important? How a 2-D filter can be designed to retain a desired coefficient sparsity for efficient implementation while achieving best possible performance subject to that sparsity constraint? These are the focus of this paper in which we present a two-phase design method for 2-D FIR digital filters in two most common design settings, namely, the least squares and minimax designs. Simulation studies are presented to illustrate each phase of the proposed design method and to evaluate the performance of the filters designed.

Keywords 2-D digital filters · Coefficient sparsity · Convex optimization

1 Introduction

Design and analysis of two-dimensional (2-D) digital filters have been a field of active research since 1970's. As a result, a good number of analysis techniques and design algorithms for 2-D digital filters have been developed and they now form an important part of multidimensional digital signal processing [1]–[4]. One of the design issues is coefficient (impulse response) sparsity. The issue is evidently of importance as it is directly related to filter implementation efficiency and cost. Several authors have investigated digital filters with sparse coefficients, especially for filters with specific system structures or specific classes of filters. These include the frequency-response masking (FRM) filters initiated by Lim [5] and extended to the 2-D case [6], narrowband 2-D fan filters and 3-D cone filters using shaped 2-D window functions proposed by Khademi and Bruton [7][8], and sparse half-band like FIR

Wu-Sheng Lu
Dept. of Electrical and Computer Eng.
University of Victoria
Victoria, BC, Canada V8W 3P6
E-mail: wslu@ece.uvic.ca

Takao Hinamoto
Graduate School of Engineering
Hiroshima University
Higashi-Hiroshima, 739-8527, Japan
E-mail: hina@crest.ocn.ne.jp

filters by Gustafsson et al [9]. In a basic FRM filter [5], a prototype filter and its complement are upsampled, yielding sparse filter coefficients and a reduced transition width. They are connected in cascade to a pair of frequency-response masking filters to approximiate a desired sharp frequency response. On the other hand, by using a parabolically-bounded shaped 2-D window and properly thresholding the magnitude of the windowed impulse response, narrowband 2-D fan filters with sparse coefficients and satisfactory performance are obtained [7][8]. In [9], a method for the design of filters close to half-band filter with slightly relaxed specifications is developed. It is shown that one can obtain filters of this type with sparse coefficients by allowing an increase in passband ripple.

Inspired by the recent development in compressive sensing [10][11][12], this paper examines the 2-D filter design problem from a sparsity-promoting perspective and addresses several design issues for 2-D FIR filters with sparse coefficients. Our objective is to develop a design methodology that applies to *general* 2-D FIR filters without assuming specific system structures or limiting to specific filter classes. In order for the design method to accommodate a wide range of filtering scenarios, we in this paper consider two most common design settings, namely, the frequency-weighted least squares and minimax designs. For each type of designs, we present a two-phase design method that accommodates a sparsity-promoting measure and yields globally optimal solutions subject to target coefficient sparsity. Illustrations of the design concept, technical details of the design algorithms, and simulation studies are presented in Sections 2–3, Sections 4–5, and Section 6, respectively. We remark that the authors have also looked into the sparsity problem for one-dimensional (1-D) digital filters and some preliminary results have recently appeared in [13]. However, the studies presented in this paper differ from that of [13] in several ways as we have to deal with issues facing the class of 2-D filters that do not exist for its 1-D counterpart.

Throughout, we use boldfaced upper-case letters to denote matrices and lower-case letters to denote vectors. Two vector norms for $\mathbf{x} = [x_1 \;\; x_2 \;\; \cdots \;\; x_n]^T$, which are frequently used in the paper, are the l_1-norm and l_2-norm that are defined as $\|\mathbf{x}\|_1 = \sum_{i=1}^{n} |x_i|$ and $\|\mathbf{x}\|_2 = \left(\sum_{i=1}^{n} x_i^2\right)^{1/2}$, respectively. For a continuous function $F(\omega_1, \omega_2)$ defined over region Ω, its L_2 and L_∞ norms are defined as

$$\|F(\omega_1,\omega_2)\|_2 = \left[\iint_{\Omega} |F(\omega_1,\omega_2)|^2 \, d\omega_1 \, d\omega_2\right]^{1/2}$$

and

$$\|F(\omega_1,\omega_2)\|_\infty = \max_{(\omega_1,\omega_2)\in\Omega} |F(\omega_1,\omega_2)|$$

respectively.

2 Linear-Phase 2-D FIR Digital Filters

2.1 *Transfer function and impulse response*

We consider the transfer function of a representative linear-phase 2-D FIR filter

$$H(z_1,z_2) = \sum_{i=0}^{N_1}\sum_{j=0}^{N_2} H_{ij} z_1^{-i} z_2^{-j} = \mathbf{z}_1^T \hat{\mathbf{H}} \mathbf{z}_2 \tag{1}$$

with N_1 and N_2 even integers, $\mathbf{z}_1 = [1\ z_1^{-1}\ \cdots\ z_1^{-N_1}]^T$, $\mathbf{z}_2 = [1\ z_2^{-1}\ \cdots\ z_2^{-N_2}]^T$. The phase-response linearity of (1) implies that $\hat{\mathbf{H}}$ is quadrantally symmetrical. Namely, if we denote

$$\hat{\mathbf{H}} = \begin{bmatrix} \mathbf{H}_{11} & \mathbf{h}_{12} & \mathbf{H}_{13} \\ \mathbf{h}_{21}^T & h_{22} & \mathbf{h}_{23}^T \\ \mathbf{H}_{31} & \mathbf{h}_{32} & \mathbf{H}_{33} \end{bmatrix} \tag{2}$$

$n_1 = N_1/2$ and $n_2 = N_2/2$, then $\mathbf{H}_{11},\ \mathbf{H}_{13},\ \mathbf{H}_{31},\ \mathbf{h}_{12}$, and $\mathbf{h}_{21}^T$ are related to the rest part of $\hat{\mathbf{H}}$ as

$$\begin{aligned} \mathbf{H}_{11} &= \texttt{flipud}(\texttt{fliplr}(\mathbf{H}_{33})) \\ \mathbf{H}_{13} &= \texttt{flipud}(\mathbf{H}_{33}),\ \mathbf{H}_{31} = \texttt{fliplr}(\mathbf{H}_{33}) \\ \mathbf{h}_{12} &= \texttt{flipud}(\mathbf{h}_{32}) \quad \text{and} \quad \mathbf{h}_{21}^T = (\texttt{fliplr}(\mathbf{h}_{23}^T)) \end{aligned}$$

where `flipud` and `fliplr` represent the operations of flipping a matrix upside down and from left to right, respectively. As a result, the frequency response of the filter is given by [14]

$$H(\omega_1, \omega_2) = e^{-j(n_1\omega_1 + n_2\omega_2)} \mathbf{c}_1^T(\omega_1)\mathbf{H}\mathbf{c}_2(\omega_2) \tag{3}$$

with $\mathbf{c}_i(\omega_i) = [1\ \cos\omega_i\ \cdots\ \cos n_i\omega_i]^T$ for $i = 1,\ 2$, and

$$\mathbf{H} = \begin{bmatrix} h_{22} & 2\mathbf{h}_{23}^T \\ 2\mathbf{h}_{32} & 4\mathbf{H}_{33} \end{bmatrix} \tag{4}$$

For the convenience of subsequent development, the zero-phase frequency response, i.e.

$$A(\omega_1, \omega_2) = \mathbf{c}_1^T(\omega_1)\mathbf{H}\mathbf{c}_2(\omega_2) \tag{5}$$

needs to be reformulated as an inner product of a frequency-dependent variable vector with a design vector that collects the entries of $\mathbf{H}$ in (4). This can be done by using a property of matrix trace: $\text{trace}(\mathbf{AB}) = \text{trace}(\mathbf{BA})$. It follows that

$$\begin{aligned} A(\omega_1, \omega_2) &= \mathbf{c}_1^T(\omega_1)\mathbf{H}\mathbf{c}_2(\omega_2) = \text{trace}[\mathbf{c}_1^T(\omega_1)\mathbf{H}\mathbf{c}_2(\omega_2)] \\ &= \text{trace}[\mathbf{c}_2(\omega_2)\mathbf{c}_1^T(\omega_1)\mathbf{H}] = \text{trace}[\mathbf{C}(\omega_1, \omega_2)\mathbf{H}] \\ &= \mathbf{c}^T(\omega_1, \omega_2)\mathbf{h} \end{aligned} \tag{6}$$

where $\mathbf{C}(\omega_1, \omega_2) = \mathbf{c}_2(\omega_2)\mathbf{c}_1^T(\omega_1)$, thus $\mathbf{c}(\omega_1, \omega_2)$ is a vector of length $n = (n_1 + 1)(n_2 + 1)$ generated by stacking the rows of $\mathbf{C}(\omega_1, \omega_2)$ and then transposing it, and $\mathbf{h}$ is a vector obtained by stacking the columns of $\mathbf{H}$ from left to right.

2.2 *Is the impulse response of a typical 2-D FIR filter sparse?*

A discrete sequence (vector) $\mathbf{x}$ of length n is said to be K-sparse if $\mathbf{x}$ contains K nonzero entries with $K \ll n$. In the context of filter design, we are interested in developing design methodologies for digital filters with satisfactory performance *and* sparse impulse responses because filters such as these require considerably less number of multiplications and additions for implementation. Such development is especially significant for 2-D filters as the number of coefficients involved in a 2-D filter of size $N_1 \times N_2$ is in the order of $O(N_1N_2)$ compared with $O(N)$ coefficients for 1-D filters of length N.

There are two questions immediately arising from this coefficient sparsity issue: Are the impulse responses of typical 2-D FIR filters sparse? and, if not, are there simple techniques to approximate a nonsparse impulse response by a sparse one without sacrificing much of the performance?

To address the first question, we examined a variety of FIR filters in terms of types and sizes, these include circularly symmetric lowpass, highpass, bandpass, and bandstop filters, fan filters, and diamond-shaped filters with sizes varying from 5×5 to 45×45. For each type of filters of a given size, designs with various passband and stopband edges were examined. Among all tested, no impulse response having at least one zero coefficient was found. The examination is therefore indicative that impulse responses of typical 2-D FIR filters are generically nonsparse. On the other hand, however, a typical impulse response contains considerable number (relative to the filter size) of coefficients whose magnitudes are small. As an example, Fig. 1 depicts the impulse response of a diamond-shaped linear phase FIR filter of size 15×15 with normalized passband edge $\omega_p = 0.5\pi$ and stopband edge $\omega_a = \pi$ that was designed by minimizing a least-squares (LS) criterion, namely

$$e_2 = \iint_{\Omega} |H(\omega_1, \omega_2) - H_d(\omega_1, \omega_2)|^2 d\omega_1 d\omega_2 \tag{7}$$

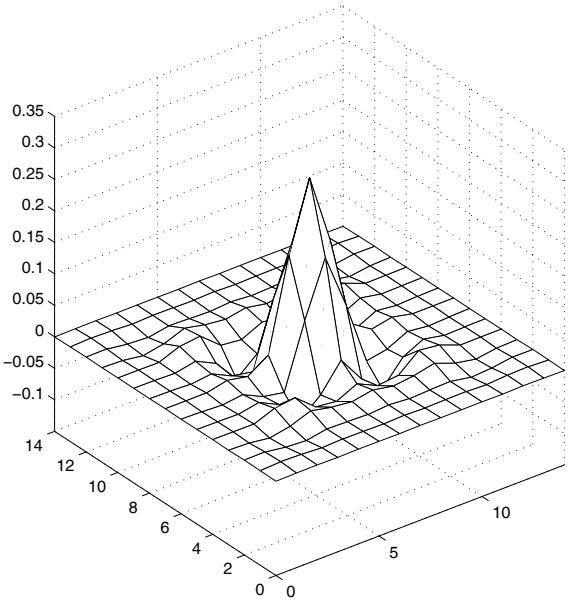

Fig. 1 Impulse response of a diamond-shaped FIR filter of size 15×15.

where Ω denotes the union of filter's passband and stopband. Among its 225 coefficients, there are 104 coefficients with magnitudes less than 10^{-3}. Taking both sides of the above analysis into account, we see that a typical 2-D impulse response is not sparse by definition, but is nearly sparse as many of its entries are fairly close to zero. Unfortunately, implementing a 2-D filter with a nearly sparse impulse response does not lead to substantial

complexity/cost reduction as the number of multipliers and adders practically remain the same.

To address the second question, we start with a seemingly reasonable approach to generating a sparse impulse response, that is to nullify the entries of a nonsparse impulse response, whose magnitudes fall below a prescribed tolerance ε_t. For illustration, the technique was applied to the diamond-shaped filter discussed above with $\varepsilon_t = 10^{-3}$ to generate an impulse response with 104 zero entries. Denoting the frequency response of the nonsparse diamond-shaped filter and its sparse counterpart by $H_{15}(\omega_1, \omega_2)$ and $H_{15}^{(s)}(\omega_1, \omega_2)$, respectively, the e_2 error defined in (7) was found to be $e_2[H_{15}(\omega_1, \omega_2)] = 0.6640 \times 10^{-5}$ and $e_2[H_{15}^{(s)}(\omega_1, \omega_2)] = 0.1430 \times 10^{-3}$. We see that the performance of the sparse filter is considerably degraded compared with its nonsparse counterpart.

Since there are only 121 nonzero coefficients in $H_{15}^{(s)}(z_1, z_2)$, from an implementation point of view a fair comparison should be made between $H_{15}^{(s)}(z_1, z_2)$ and a nonsparse 2-D filter with 121 coefficients. To this end, an LS-optimal diamond-shaped filter of size 11×11 with the same design specifications as that for $H_{15}(z_1, z_2)$ was designed and the transfer function obtained is denoted as $H_{11}(z_1, z_2)$. The e_2 error for $H_{11}(z_1, z_2)$ was found to be $e_2[H_{11}(z_1, z_2)] = 0.0708 \times 10^{-3}$. The amplitude responses of $H_{15}^{(s)}(z_1, z_2)$ and $H_{11}(z_1, z_2)$ are shown in Fig. 2. We note that the performance of the sparse $H_{15}^{(s)}(z_1, z_2)$ is not as good as that of the equivalent nonsparse $H_{11}(z_1, z_2)$. The above example is indicative that simply nullifying small-magnitude coefficients will not produce a sparse 2-D filter with satisfactory performance. A focus point of this paper is to develop algorithms for the design of 2-D FIR filters with sparse coefficients that outperform their equivalent nonsparse counterparts.

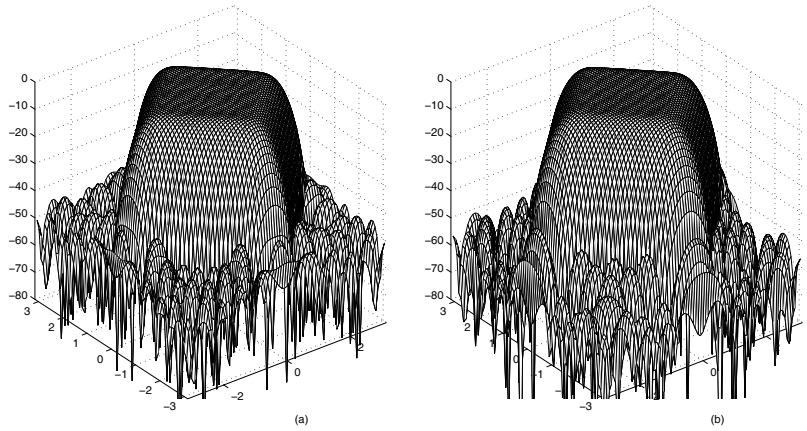

Fig. 2 Amplitude responses of (a) $H_{15}^{(s)}(z_1, z_2)$ and (b) $H_{11}(z_1, z_2)$.

3 The Design Method at A Glance

The design of an optimal linear-phase 2-D filter with sparse coefficients involves two distinct phases. We regard the impulse response of a 2-D FIR filter as matrix $\mathbf{H}$ (see (3)) and the indices $(i,\ j)$ of its entries h_{ij} as *locations*. The aim of the first design phase is at identifying

the locations where the filter coefficients will be set to zero to satisfy a sparsity requirement. This is followed by a second design phase in that the remaining nonzero entries of the impulse response are optimally determined by minimizing a certain error measure of how close the filter being designed is to a desired frequency response.

3.1 *Design Phase I*

In brief terms, phase 1 of the design method is accomplished by optimizing a fidelity term that measures the closeness between the frequency response of the filter and a given (desired) one, in combination with a sparsity-promoting measure. In the context of filter design, the fidelity term, in accordance with the two most common design settings, namely the least squares and minimax designs, is the L_2 or L_∞-norm of approximation error $H(\omega_1,\omega_2) - H_d(\omega_1,\omega_2)$ over a region of interest in the frequency baseband $-\pi \le \omega_1,\ \omega_2 \le \pi$. A recent discovery in the theory of compressive sensing [10][11][12] is that under certain conditions the sparsest solution of a underdetermined linear system $\mathbf{Ax} = \mathbf{b}$ can be found by minimizing l_1-norm $\|\mathbf{x}\|_1$ subject to $\mathbf{Ax} = \mathbf{b}$. It is important to note that the l_1-norm $\|\mathbf{x}\|_1$ is a continuous and globally convex function of $\mathbf{x}$ (it is not differentiable however). Summarizing, phase 1 of the design amounts to minimizing a weighted sum of these two terms with respect to vectorized impulse response $\mathbf{h}$:

$$\underset{\mathbf{h}}{\text{minimize}}[\|H(\omega_1,\omega_2) - H_d(\omega_1,\omega_2)\|_{2,\infty} + \mu\|\mathbf{h}\|_1] \tag{8}$$

where $\mu > 0$ is a scalar weight that balances filter's sparsity with its fidelity. We remark that the objective function in (8) may be modified to include a frequency-weighting function $W(\omega_1,\omega_2)$ to help discriminate frequency region of interest over the baseband, details will follow. Once the solution $\mathbf{h}$ of problem (8) is obtained, an index set of the most appropriate locations for the impulse response to be set to zero to satisfy a given sparsity constraint can be identified by hard-thresholding the entries of $\mathbf{h}$ with an appropriate threshold ε_t. It is this index set that is the essential outcome of design phase 1.

3.2 *Design Phase II*

As observed from the example in Sec. 2, sparsifying an impulse response inevitably degrades the filter's performance. This necessitates a second phase of the design in that the filter is optimized against an L_2 or L_∞ error measure subject to filter's sparsity identified in phase 1. Let $\mathscr{I}$ be the index set produced in phase 1, that collects the locations of the impulse response that are most adequate to be set to zero, phase 2 is essentially a design step that solves the constrained problem

$$\underset{\mathbf{h}}{\text{minimize}} \quad \|H(\omega_1,\omega_2) - H_d(\omega_1,\omega_2)\|_{2,\infty} \tag{9a}$$

$$\text{subject to:} \quad h_i = 0 \qquad \text{for} \qquad i \in \mathscr{I} \tag{9b}$$

As a result, a 2-D FIR filter with desired coefficient sparsity and best possible performance in accordance with a given error measure is obtained. A noteworthy feature of the design method is that both the problems in (8) and (9) can be cast as low-order (linear or quadratic) *convex programming* (CP) problems [15], hence the solutions are globally optimal, and can be calculated using efficient and reliable CP solvers such as SeDuMi [16] and CVX [17]. In the rest of the paper, we present algorithmic details addressing the problems in (8) and (9) as well as several design results.

4 An Algorithm for LS Designs

4.1 CP Formulation for Design Phase 1

There are two possible formulations for design phase 1: one takes the form of (8) with L_2 norm in its first term, and the other formulates the problem as

$$\begin{aligned} &\text{minimize} && \|\mathbf{h}\|_1 && \text{(10a)} \\ &\text{subject to:} && \|H(\omega_1,\omega_2) - H_d(\omega_1,\omega_2)\|_2 \le \delta && \text{(10b)} \end{aligned}$$

with δ a prescribed upper bound for the L_2 approximation error. These two formulations are equivalent because the objective function in (8) may be interpreted as the Lagrangian of problem (10) up to a constant [15]. Therefore our attention in phase 1 will be focused on problem (8) with a slight modification so that the objective function becomes frequency weighted. To be specific, we consider designing a linear-phase FIR filter $H(z_1,z_2)$ in (1) with N_1 and N_2 even that approximates a desired frequency response

$$H_d(\omega_1,\omega_2) = e^{-j(n_1\omega_1+n_2\omega_2)}A_d(\omega_1,\omega_2) \tag{11}$$

in a weighted least squares sense subject to coefficient sparsity no greater than K. With (6), (8) and (11), phase 1 of the design aims at solving the problem of minimizing function $J_2(\mathbf{h})$ with

$$J_2(\mathbf{h}) = \left\{ \iint_{\Omega} W(\omega_1,\omega_2)[\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)]^2 d\omega_1 d\omega_2 \right\}^{1/2} + \mu\|\mathbf{h}\|_1 \tag{12}$$

where the weighting functin assumes the form (in our simulation studies)

$$W(\omega_1,\omega_2) = \begin{cases} 1 & (\omega_1,\omega_2) \in \text{ passbands} \\ w & (\omega_1,\omega_2) \in \text{ stopbands} \\ 0 & \text{elsewhere} \end{cases} \tag{13}$$

hence the region Ω where the integration in the fidelity term of (12) is carried out is simply the union of all passbands and stopbands within the baseband $[-\pi,\pi]\times[-\pi,\pi]$. A straightforward numerical implementation of function $J_2(\mathbf{h})$ is given by

$$J_2(\mathbf{h}) \approx \left\{ \sum_{i=1}^{M} W(\omega_i)[\mathbf{h}^T\mathbf{c}(\omega_i) - A_d(\omega_i)]^2 \right\}^{1/2} + \mu\|\mathbf{h}\|_1 \tag{14}$$

where $\Omega_d = \{\omega_i = (\omega_1^{(i)}, \omega_2^{(i)}),\ 1 \le i \le M\}$ is a dense set of M frequency grids uniformly placed over Ω. Let

$$\hat{\mathbf{A}} = \begin{bmatrix} W^{1/2}(\omega_1)\mathbf{c}^T(\omega_1) \\ W^{1/2}(\omega_2)\mathbf{c}^T(\omega_2) \\ \vdots \\ W^{1/2}(\omega_M)\mathbf{c}^T(\omega_M) \end{bmatrix},\ \mathbf{a} = \begin{bmatrix} W^{1/2}(\omega_1)A_d(\omega_1) \\ W^{1/2}(\omega_2)A_d(\omega_2) \\ \vdots \\ W^{1/2}(\omega_M)A_d(\omega_M) \end{bmatrix}$$

then (14) becomes

$$J_2(\mathbf{h}) \approx \|\hat{\mathbf{A}}\mathbf{h} - \mathbf{a}\|_2 + \mu\|\mathbf{h}\|_1 \tag{15}$$

To deal with the non-differentiability of $\|\mathbf{h}\|_1$, we place an upper bound for each entry of $\mathbf{h}$, i.e.,

$$|h_i| \le d_i \qquad \text{for } 1 \le i \le n$$

with $n = (n_1+1)(n_2+1)$. An upper bound β is also introduced for the fidelity term in (15), in this way a constrained problem equivalent to miniming $J_2(\mathbf{h})$ in (15) is obtained as

$$\begin{aligned} \text{minimize} \quad & \beta + \mu \sum_{i=1}^{n} d_i && \text{(16a)} \\ \text{subject to:} \quad & \|\hat{\mathbf{A}}\mathbf{h} - \mathbf{a}\|_2 \le \beta && \text{(16b)} \\ & |h_i| \le d_i \qquad \text{for } 1 \le i \le n && \text{(16c)} \end{aligned}$$

By treating the bounds β and d_i $(1 \le i \le n)$ as auxiliary design variables and defining

$$\mathbf{d} = \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_n \end{bmatrix}, \ \mathbf{x} = \begin{bmatrix} \beta \\ \mathbf{d} \\ \mathbf{h} \end{bmatrix}, \ \mathbf{A}_1 = [\mathbf{0}\ \hat{\mathbf{A}}]_{M\times(2n+1)}, \ \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}_{(2n+1)\times 1},$$

$$\mathbf{f} = \begin{bmatrix} 1 \\ \mu\mathbf{e}_n \\ \mathbf{0} \end{bmatrix}_{(2n+1)\times 1}, \ \mathbf{e}_n = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}_{n\times 1}, \ \mathbf{A}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{I}_n & \mathbf{I}_n \\ \mathbf{0} & \mathbf{I}_n & -\mathbf{I}_n \end{bmatrix}_{2n\times(2n+1)}$$

the problem in (16) becomes

$$\begin{aligned} \text{minimize} \quad & \mathbf{f}^T\mathbf{x} && \text{(17a)} \\ \text{subject to:} \quad & \|\mathbf{A}_1\mathbf{x} - \mathbf{a}\|_2 \le \mathbf{b}^T\mathbf{x} && \text{(17b)} \\ & \mathbf{A}_2\mathbf{x} \ge \mathbf{0} && \text{(17c)} \end{aligned}$$

which is a standard second-order cone programming (SOCP) problem [15]. From the unique and global solution $\mathbf{x}^*$ of (17), $\mathbf{h}^*$ can be obtained as the last n entrices of $\mathbf{x}^*$. Because of the use of the sparsity-promoting term $\mu\|\mathbf{h}\|_1$ in (15) (hence in (16) and (17)), $\mathbf{h}^*$ is indicative of the index locations at which the impulse response can adequately be set to zero. A target impulse-response sparsity K can be readily satisfied by hard-thresholding $\mathbf{h}^*$ with an appropriate threshold ε_t. The set of indices at which the impulse response is set to zero is denoted by $\mathscr{I}_2^*$ and, by definition, it is characterized by

$$\mathscr{I}_2^* = \{i : 1 \le i \le n,\ |h_i^*| \le \varepsilon_t\} \tag{18}$$

It is this set $\mathscr{I}_2^*$ that is the essential result produced by design phase 1.

4.2 *CP Formulation for Design Phase 2*

As demonstrated in Sec. 2, nullifying small-magnitude impulse-response entries inevitably deteriorates the filter's performance even if a fairly small threshold is used in the nullification process. In phase 2 of the design, this problem is addressed by optimizing the rest of (nonzero) impulse response entries with respect to an LS criterion subject to the sparsity constraint. This leads to a constrained problem of the type in (9) with an L_2-norm objective

function. By introducing a weighting function in (9a) and using (6) and (11), the problem under consideration is formulated as

$$\text{minimize} \quad \iint_{\Omega} W(\omega_1,\omega_2)[\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)]^2 d\omega_1 d\omega_2 \tag{19a}$$

$$\text{subject to:} \quad h_i = 0 \qquad \text{for } i \in \mathscr{I}_2^* \tag{19b}$$

which is evidently a convex quadratic programming problem. The solution of (19) can be calculated by first eliminating the constrains in (19b) by substituting them into (19a), resulting in a unconstrained convex problem, and then solving it analytically. To this end, let $\hat{\mathbf{h}}$ and $\hat{\mathbf{c}}(\omega_1,\omega_2)$ be the vectors generated from $\mathbf{h}$ and $\mathbf{c}(\omega_1,\omega_2)$, respectively, by deleting the entries whose indices belong to set $\mathscr{I}_2^*$. Subject to the constraints in (19b), we have $\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) = \hat{\mathbf{h}}^T\hat{\mathbf{c}}(\omega_1,\omega_2)$, hence (19) is simplified to the unconstrained problem

$$\text{minimize } \hat{\mathbf{h}}^T\hat{\mathbf{Q}}\hat{\mathbf{h}} - 2\hat{\mathbf{h}}^T\hat{\mathbf{q}} + \kappa \tag{20}$$

with

$$\hat{\mathbf{Q}} = \iint_{\Omega} W(\omega_1,\omega_2)\,\hat{\mathbf{c}}(\omega_1,\omega_2)\hat{\mathbf{c}}^T(\omega_1,\omega_2)\,d\omega_1\,d\omega_2$$

$$\hat{\mathbf{q}} = \iint_{\Omega} W(\omega_1,\omega_2)\,\hat{\mathbf{c}}(\omega_1,\omega_2)\,A_d(\omega_1,\omega_2)\,d\omega_1\,d\omega_2$$

$$\kappa = \iint_{\Omega} W(\omega_1,\omega_2)\,A_d^2(\omega_1,\omega_2)\,d\omega_1\,d\omega_2$$

Because $\hat{\mathbf{Q}}$ is positive definite, the global minimum of (20) is reached by

$$\hat{\mathbf{h}}^* = \hat{\mathbf{Q}}^{-1}\hat{\mathbf{q}} \tag{21}$$

The full-scale solution of dimension n for problem (19) can now be obtained by inserting zero entries into $\hat{\mathbf{h}}^*$ so as to satisfy (19b). The solution so constructed is denoted by $\mathbf{h}_{LS}$.

5 An Algorithm for Minimax Designs

5.1 *LP Formulation for Design Phase 1*

With (6) and (11), a frequency-weighted version of formulation (8) with L_∞ norm in its first term becomes

$$\underset{\mathbf{h}}{\text{minimize}}[\underset{(\omega_1,\omega_2)\in\Omega}{\text{maximize}}\ W(\omega_1,\omega_2)|\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)| + \mu\|\mathbf{h}\|_1] \tag{22}$$

By introducing an upper bound β for $W(\omega_1,\omega_2)|\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)|$ over a finite set of frequency grids $\Omega_d = \{\omega_i = (\omega_1^{(i)}, \omega_2^{(i)}),\ 1 \le i \le M\} \subseteq \Omega$, (22) is reduced to a tractable constrained problem as

$$\text{minimize} \quad \beta + \mu\|\mathbf{h}\|_1 \tag{23a}$$

$$\text{subject to:} \quad W(\omega_i)|\mathbf{h}^T\mathbf{c}(\omega_i) - A_d(\omega_i)| \le \beta, \quad 1 \le i \le M \tag{23b}$$

In the same way as in the LS design, term $\mu\|\mathbf{h}\|_1$ in (23a) is replaced by the sum of upper bounds for the entries of $\mathbf{h}$. This leads (23) to

$$\begin{aligned} \text{minimize} \quad & \beta + \mu \sum_{i=1}^{n} d_i && \text{(24a)} \\ \text{subject to:} \quad & W(\omega_i)|\mathbf{h}^T\mathbf{c}(\omega_i) - A_d(\omega_i)| \le \beta, \quad 1 \le i \le M && \text{(24b)} \\ & |h_i| \le d_i \quad 1 \le i \le n && \text{(24c)} \end{aligned}$$

By treating the bounds β and d_i as auxiliary design variables, (24) becomes a linear programming (LP) problem that can be expressed as

$$\begin{aligned} \text{minimize} \quad & \mathbf{f}^T\mathbf{x} && \text{(25a)} \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} \ge \mathbf{b} && \text{(25b)} \end{aligned}$$

where

$$\mathbf{x} = \begin{bmatrix} \beta \\ \mathbf{d} \\ \mathbf{h} \end{bmatrix}, \ \mathbf{f} = \begin{bmatrix} 1 \\ \mu\mathbf{e}_n \\ \mathbf{0} \end{bmatrix}, \ \mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}, \ \mathbf{b} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{0} \end{bmatrix}$$

with

$$\mathbf{A}_1 = \begin{bmatrix} 1 & \mathbf{0} & W(\omega_1)\mathbf{c}^T(\omega_1) \\ \vdots & \vdots & \vdots \\ 1 & \mathbf{0} & W(\omega_M)\mathbf{c}^T(\omega_M) \\ 1 & \mathbf{0} & -W(\omega_1)\mathbf{c}^T(\omega_1) \\ \vdots & \vdots & \vdots \\ 1 & \mathbf{0} & -W(\omega_M)\mathbf{c}^T(\omega_M) \end{bmatrix}, \ \mathbf{A}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{I}_n & \mathbf{I}_n \\ \mathbf{0} & \mathbf{I}_n & -\mathbf{I}_n \end{bmatrix}, \ \mathbf{b}_1 = \begin{bmatrix} W(\omega_1)A_d(\omega_1) \\ \vdots \\ W(\omega_M)A_d(\omega_M) \\ -W(\omega_1)A_d(\omega_1) \\ \vdots \\ -W(\omega_M)A_d(\omega_M) \end{bmatrix}$$

The unique and globally optimal solution $\mathbf{h}^*$ of problem (23) can be obtained by collecting the last n components of the solution $\mathbf{x}^*$ of problem (25). Just like the LS design in Sec. 4, a target impulse response sparsity K can be satisfied by hard-thresholding $\mathbf{h}^*$ with an appropriate threshold ε_t. This yields an index set

$$\mathscr{I}_\infty^* = \{i : 1 \le i \le n, \ |h_i^*| \le \varepsilon_t\} \tag{26}$$

5.2 *LP Formulation for Design Phase 2*

With $\mathscr{I}_\infty^*$ identified in phase 1, phase 2 of the design is performed by solving the constrained minimax problem

$$\begin{aligned} \underset{\mathbf{h}}{\text{minimize}} \ \underset{(\omega_1,\omega_2)\in\Omega}{\text{maximize}} \ & W(\omega_1,\omega_2)|\mathbf{h}^T\mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)| && \text{(27a)} \\ \text{subject to:} \quad & h_i = 0 \qquad \text{for } i \in \mathscr{I}_\infty^* && \text{(27b)} \end{aligned}$$

The constrains in (27b) ensures the desired sparsity for the impulse response while minimizing the largest weighted approximation error over the entire region of interest, Ω. Realistically (27) is carried out over a dense but finite set of frequency grids $\Omega_d = \{\omega_i =$

$(\omega_1^{(i)}, \omega_2^{(i)}),\ 1 \le i \le M\} \subseteq \Omega$ and the problem is solved by minimizing an upper bound β for the weighted approximation error over Ω_d subject to (27b), and it can be formulated as

$$\begin{aligned} \text{minimize} \quad & \beta && \text{(28a)} \\ \text{subject to:} \quad & W(\omega_i)|\mathbf{h}^T\mathbf{c}(\omega_i) - A_d(\omega_i)| \le \beta \quad 1 \le i \le M && \text{(28b)} \\ & h_i = 0 \qquad \text{for } i \in \mathscr{I}_\infty^* && \text{(28c)} \end{aligned}$$

Here we follow the treatment of the sparsity constraint in Sec. 4.2 to define vectors $\hat{\mathbf{h}}$ and $\hat{\mathbf{c}}(\omega)$ by deleting those entries of $\mathbf{h}$ and $\mathbf{c}(\omega)$ whose indices belong to set $\mathscr{I}_\infty^*$. Subject to constraints in (28c), we have $\mathbf{h}^T\mathbf{c}(\omega) = \hat{\mathbf{h}}^T\hat{\mathbf{c}}(\omega)$, hence (28) is simplified to

$$\begin{aligned} \text{minimize} \quad & \beta && \text{(29a)} \\ \text{subject to:} \quad & W(\omega_i)|\hat{\mathbf{h}}^T\hat{\mathbf{c}}(\omega_i) - A_c(\omega_i)| \le \beta, \quad 1 \le i \le M && \text{(29b)} \end{aligned}$$

With β as an auxiliary variable, problem (29) becomes an LP problem of the form

$$\begin{aligned} \text{minimize} \quad & \mathbf{f}^T\mathbf{x} && \text{(30a)} \\ \text{subject to:} \quad & \mathbf{A}\mathbf{x} \ge \mathbf{b} && \text{(30b)} \end{aligned}$$

with

$$\mathbf{x} = \begin{bmatrix} \beta \\ \hat{\mathbf{h}} \end{bmatrix},\ \mathbf{f} = \begin{bmatrix} 1 \\ \mathbf{0} \end{bmatrix},\ \mathbf{A} = \begin{bmatrix} 1 & W(\omega_1)\hat{\mathbf{c}}^T(\omega_1) \\ \vdots & \vdots \\ 1 & W(\omega_M)\hat{\mathbf{c}}^T(\omega_M) \\ 1 & -W(\omega_1)\hat{\mathbf{c}}^T(\omega_1) \\ \vdots & \vdots \\ 1 & -W(\omega_M)\hat{\mathbf{c}}^T(\omega_M) \end{bmatrix},\ \text{and } \mathbf{b} = \begin{bmatrix} W(\omega_1)A_d(\omega_1) \\ \vdots \\ W(\omega_M)A_d(\omega_M) \\ -W(\omega_1)A_d(\omega_1) \\ \vdots \\ -W(\omega_M)A_d(\omega_M) \end{bmatrix}$$

The optimal $\hat{\mathbf{h}}^*$ is obtained by collecting the last n entries of the solution $\mathbf{x}^*$ of problem (30), and the solution of problem (28) is then constructed by inserting zero entries into $\hat{\mathbf{h}}^*$ so as to satisfy (28c). The solution so obtained is denoted by $\mathbf{h}_{minimax}$.

6 Simulation Studies

6.1 *Effect of Using a Sparsity-Promoting Measure*

The objective functions used in phase 1 of LS and minimax designs are sparsity promoting as they both contain the term $\mu\|\mathbf{h}\|_1$ (see (14) and (23)). In this section, we provide numerical evidences to illustrate the effect of using a sparsity-promoting measure in the designs. To this end, a circularly symmetric lowpass filter of size 15×15 with $\omega_p = 0.5\pi$, $\omega_a = 0.8\pi$, and $w = 1$ (see (13)) was obtained by minimizing the objective function in (14) with $\mu = 0.5$. Since the filter is circularly symmetric, the frequency baseband can be reduced to $\{(\omega_1, \omega_2) : 0 \le \omega_1,\ \omega_2 \le \pi\}$. The set Ω_d includes 188 grid points in the passband and 446 grid points in the stopband, thus $M = 634$. By thresholding its solution with $\varepsilon_t = 0.00149$, an index set $\mathscr{I}_2^*$ of 104 locations was generated. For comparison, a circularly symmetric LS lowpass filter with the same filter size, ω_p, ω_a, and w was also designed by minimizing an objective function with the same fidelity term as in (14) but dropping the term $\mu\|\mathbf{h}\|_1$ there. By thresholding the impulse response obtained with $\varepsilon_t = 0.002$, an index set $\mathscr{I}_2$ of

104 locations was obtained. On comparing the two index sets, it was found that there were 12 locations that were in $\mathscr{I}_2^*$ but not in $\mathscr{I}_2$, which implies that there were 12 locations that were in $\mathscr{I}_2$ but not in $\mathscr{I}_2^*$. The locations associated with these sets are shown in Fig . 3.

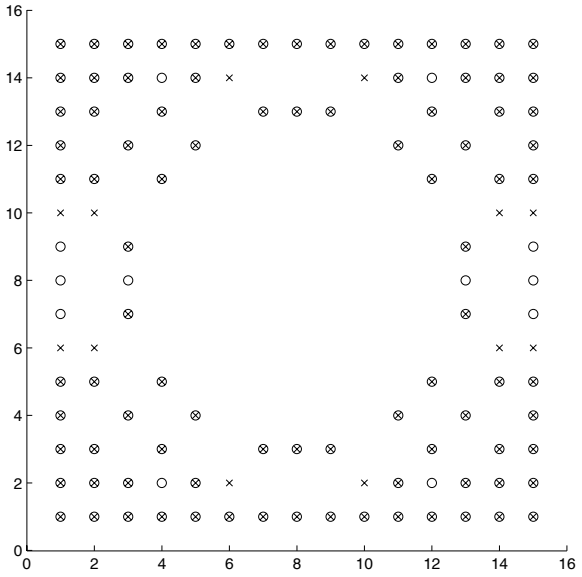

Fig. 3 Locations of zeros generated by (a) minimizing (14) ('$\circ$' in the figure) and (b) minimizing (14) without term $\mu\|\mathbf{h}\|_1$ ('$\times$' in the figure).

To examine the effect of working with index set $\mathscr{I}_2^*$ instead of index set $\mathscr{I}_2$, two circularly symmetric lowpass LS filters of size 15×15 with the same $\omega_p,\ \omega_a$, and w given above subject to sparsity $K = 104$, specified by sets $\mathscr{I}_2^*$ and $\mathscr{I}_2$, respectively, were designed. The L_2 approximation errors of these filters are shown in Table 1. For comparison, an equivalent nonsparse LS-optimal lowpass filter of size 11×11 was designed and its L_2 error is also included in Table 1.

Table 1 Performance comparison for L_2 designs

Filter size	# of zero coeff.	index set of zero locations	L_2 error
11×11	0	nonsparse	8.0207×10^{-4}
15×15	104	$\mathscr{I}_2$	5.2914×10^{-4}
15×15	104	$\mathscr{I}_2^*$	4.7088×10^{-4}

From Table 1, the positive effect of using a sparsity-promoting measure on filter's performance is observed.

6.2 *Examples of LS Filters with Sparse Coefficients*

Four circularly symmetric lowpass (with $\omega_p = 0.5\pi,\ \omega_a = 0.7\pi$) and four diamond-shaped lowpass (with $\omega_p = 0.6\pi,\ \omega_a = \pi$) LS filters of sizes $11\times 11,\ 17\times 17,\ 23\times 23$, and 29×29 were designed by applying the algorithm described in Sec. 4, where parameters μ (see (14)) and threshold ε_t (see (18)) were chosen such that the number of zero coefficients in each design were significant relative to the filter size. The number of frequency grids placed in the passband and stopband was $M = 719$ for all four designs of circularly symmetric lowpass filter and was $M = 615$ for all four designs of diamond-shaped lowpass filters. The performance of the filters designed was evaluated in terms of L_2 approximation error e_2 (see (7)) and was compared with the equivalent LS-optimal nonsparse filters that contain the same number of nonzero coefficients as in their sparse counterparts, see the details in Tables 2 and 3.

Table 2 LS designs of circularly symmetric lowpass filters

Filter size	11×11	17×17	23×23	29×29
# of zero coeff.	72	120	304	480
M	719	719	719	719
μ	0.5	0.5	1	0.2
ε_t	0.009	0.002	0.1893×10^{-3}	0.6800×10^{-3}
Approximation error: Optimal sparse	0.0218	0.0015	8.9175×10^{-4}	1.3832×10^{-4}
Approximation error: Equi. optimal nonsparse	0.0341	0.0024	0.0014	2.5516×10^{-4}
Relative reduction in approximation error	36.07%	37.50%	36.30%	45.79%

Table 3 LS designs of diamond-shaped lowpass filters

Filter size	11×11	17×17	23×23	29×29
# of zero coeff.	72	120	304	480
M	615	615	615	615
μ	0.2	0.2	0.1	0.04
ε_t	0.007	3.91×10^{-4}	2.30×10^{-4}	9.0×10^{-5}
Approximation error: Optimal sparse	0.0055	4.0193×10^{-5}	1.0819×10^{-5}	2.0983×10^{-6}
Approximation error: Equi. optimal nonsparse	0.0083	1.2258×10^{-4}	3.3936×10^{-5}	4.6271×10^{-6}
Relative reduction in approximation error	33.73%	67.21%	68.12%	54.65%

From Tables 2 and 3, it is observed that the sparse FIR filters outperform their nonsparse counterparts. The performance improvement of the sparse FIR filters over the corresponding nonsparse filters in terms of relative reduction of L_2 approximation error is shown in the last

row of the tables. As a representative design, the magnitude responses of the optimized sparse circularly symmetric lowpass filter of size 29×29 with 480 zero coefficients and its nonsparse counterpart (of size 19×19) are depicted in Fig. 4.

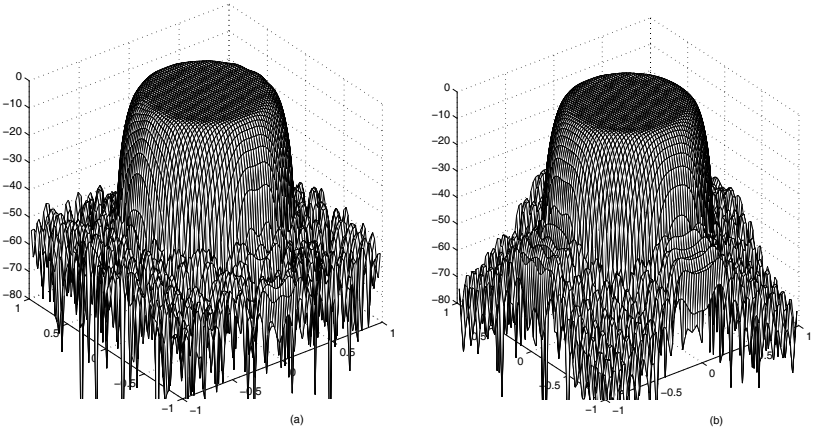

Fig. 4 Magnitude responses of (a) L_2-optimal circularly symmetric lowpass filter of size 29×29 with 480 zero coefficients and (b) equivalent L_2-optimal nonsparse filter of size 19×19 with no zero coefficients.

6.3 *Examples of Minimax FIR Filters with Sparse Coefficients*

Four circular symmetric lowpass (with $\omega_p = 0.5\pi,\ \omega_a = 0.7\pi$) and four diamond-shaped lowpass (with $\omega_p = 0.6\pi,\ \omega_a = \pi$) lowpass minimax filters of sizes $11 \times 11,\ 17 \times 17,\ 23 \times 23$, and 29×29 were designed by applying the algorithm described in Sec. 5, where parameters μ (see (23)) and ε_t (see (26)) were chosen to produce filters with sufficient coefficient sparsity relative to their size. The performance of these filters was evaluated in term of L_∞ approximation error, i.e.,

$$e_\infty = \max_{(\omega_1,\omega_2)\in\Omega} |\mathbf{h}^T \mathbf{c}(\omega_1,\omega_2) - A_d(\omega_1,\omega_2)|$$

and was compared with the equivalent minimax nonsparse filters that contain the same number of nonzero coefficients as in their sparse counterparts, see Tables 4 and 5 for details. The number M of frequency grids used in each design is also included in the tables. Note that in the tables the L_∞ approximation error are given in terms of largest ripple over the passband and minimum attenuation (in dB) over the stopband.

From Tables 4 and 5, we observe that the sparse FIR filters outperform their nonsparse counterparts in terms of the L_∞ approximation error in the passband and stopband. The performance improvement of the sparse FIR filters over the corresponding nonsparse filters in terms of relative reduction of L_∞ approximation error is shown in the last row of the tables. As a representative design, the magnitude responses of the minimax sparse diamond-shape lowpass filter of size 29×29 with 480 zero coefficients and its nonsparse counterpart (of size 19×19) are shown in Fig. 5.

Table 4 Minimax designs of circulary symmetric lowpass filters

Filter size	11×11	17×17	23×23	29×29
# of zero coeff.	72	120	304	480
M	719	719	786	996
μ	0.2	0.1	0.01	0.002
ε_t	0.012	1.40×10^{-3}	1.32×10^{-3}	7.50×10^{-4}
Approximation error: Optimal sparse	0.1238/ 16.5891	0.0332/ 25.7330	0.0244/ 27.9328	0.0158 32.4586
Approximation error: Equi. optimal nonsparse	0.1641/ 14.5156	0.0420/ 23.7501	0.0325/ 24.9060	0.0184/ 31.4316
Relative reduction in approximation error	24.56%/ 14.28%	20.95%/ 8.35%	24.92%/ 12.15%	14.13%/ 3.27%

Table 5 Minimax designs of diamond-shaped lowpass filters

Filter size	11×11	17×17	23×23	29×29
# of zero coeff.	72	120	304	480
M	615	615	674	794
μ	0.1	0.05	0.02	0.001
ε_t	0.00650	0.00030	0.00031	0.00012
Approximation error: Optimal sparse	0.0775/ 20.7722	0.0089/ 35.3006	0.0054/ 39.0412	0.0009 49.4370
Approximation error: Equi. optimal nonsparse	0.0902/ 18.9966	0.0133/ 32.5474	0.0062/ 37.0627	0.0025/ 43.6187
Relative reduction in approximation error	14.08%/ 9.35%	33.08%/ 8.46%	12.90%/ 5.34%	64%/ 13.34%

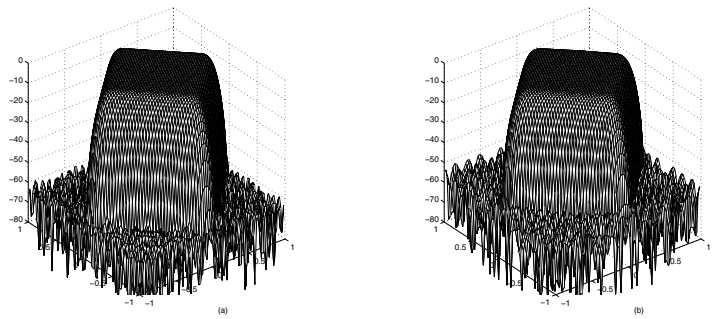

Fig. 5 Magnitude responses of (a) minimax diamond-shaped lowpass filter of size 29×29 with 480 zero coefficients and (b) equivalent minimax nonsparse filter of size 19×19 with no zero coefficients.

7 Concluding Remarks

By using a sparsity-promoting criterion in a two-phase convex optimization framework, we have developed a methodology for the design of 2-D FIR filters with sparse coefficients. Simulation studies are presented to demonstrate that with appropriate choices of design parameters, especially the values of μ and ε_t, optimal LS and minimax filters subject to a target coefficient sparsity K can be obtained to outperform their equivalent nonsparse counterparts. A drawback of the class of sparse filters studied in this paper is their larger group delay relative to their nonsparse counterparts. This should motivate studies on sparse digital filters with low group delay. It is also noticed that although optimized sparse filters offer improved performance over their nonsparse counterparts, the amount of improvement varies from case to case. This issue renders it necessary to study the interplay between the parameters μ, ε_t, and sparsity K, and how the choices of μ and ε_t effect filter's performance subject to a desired coefficient sparsity. Also, one may look into the implementation issues for digital filters with sparse coefficients in order to take full advantages of having many zero entries in impulse response. The design concept and techniques developed here are in general applicable to other types of digital filters and filter banks. Investigation of some of these issues are under way and will be reported elsewhere.

References

1. Dudgeon, D.E. & Merseau, R.M.: M*ultidimensional Digital Signal Processing*, Prentice Hall, Englewood Cliffs, N.J. (1984)
2. Lu, W.S. & Antoniou, A.: *Two-Dimensional Digital Filters*, Marcel Dekker, New York (1992)
3. Charoenlarpnopparut, C. & Bose, N.K.: Multidimensional FIR filter bank design using Groebner bases, 46 (12), 1475-1486, *IEEE Trans. Circuits Syst.* (1999)
4. Charoenlarpnopparut, C. & Bose, N.K.: Groebner Bases for Problem Solving in Multidimensional Systems, invited paper in Special Issue on Application of Groebner Bases to Multidimentsional Systems and Signal Processing (eds. by Zhiping Lin and Li Xu), 12, 365-576, *Int. J. Multidimentsional Systems and Signal Processing* (2001)
5. Lim, Y.C.: Frequency-response masking approach for the synthesis of sharp linear phase digital filters, 33(4), 357-364, *IEEE Trans. Circuits Syst*. (1986)
6. Lim, Y.C. & Lian, Y.: The optimum design of one- and two-dimensional FIR filters using the frequency response masking technique, 40(2), 88-95, *IEEE Trans. Circuits Syst. II*, (1993)
7. Khademi, L. & Bruton, L.T.: Reducing the computational complexity of narrowband 2D fan filters using shaped 2D window functions, *Int. Symp. Circuits Syst.* vol. III, 702-705 (2003)
8. Khademi, L.: Reducing the Computational Complexity of FIR Two-Dimensional Fan and Three-Dimensional Cone Filters, M.Sc. Thesis, University of Calgary, Canada, 2004, 111 pages, AAT MQ97561
9. Gustafsson, O., DeBrunner, L.S., DeBrunner, V. & Johansson, H.: On the design of sparse half-band like FIR filters, *41st Asilomar Conf.*, 1098-1102 (2007)
10. Candes, E.J., Romberg, J. & Tao, T.,: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, 52(2), 489-589, *IEEE Trans. Inform. Theory* (2006)
11. Donoho, D.L.: Compressed sensomg. 52(12), 5406-5425, *IEEE Trans. Inform. Theory* (2006)
12. Candes, E.J. & Tao, T.: Near-optimal signal recovery from random projections: Universal encoding strategies? 52(12), 5406-5425, *IEEE Trans. Inform. Theory* (2006)
13. Lu, W.S. & Hinamoto, T.: Digital filters with sparse coefficients, *Proc. International Symposium on Circuits and Systems*, 169–172, Paris, May 30-June 2, (2010)
14. Lu, W.S.: A unified approach for the design of 2-D digital filters via semidefinite programming, 49(6), 814–826, *IEEE Trans. Circuits Syst., I* (2002)
15. Antoniou, A. & Lu, W.S.: *Practical Optimization: Algorithms and Engineering Applications*, Springer, New York (2007)
16. Sturm, J.: Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones, 11–12, 625–653, *Optimization Methods and Software* (1999)
17. Grant, M., Boyd, S. & Ye, Y.: Disciplined convex programming, in *Global Optimization: From Theory to Implementation*, L. Liberti and N. Maculan (eds.) 155–210, Springer, New York (2006)