Digital Filters with Sparse Coefficients

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

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

***Abstract*—Is sparsity an issue in filter design problems? and why is it important? How a digital filter can be designed to have a sparse impulse response for efficient implementation while achieving improved performance relative to its non-sparse counterpart? In an attempt to address these questions, this paper comes up with a design technique for optimal linear-phase FIR filters with sparse impulse responses.**

I. Introduction

Research in the analysis and design of digital filters has stayed active since late 1960's, as a result the field has matured to form an important part of theory and practice in digital signal processing [1]–[5]. Inspired by the recent development in compressive sensing and sparse signal processing [6]–[9], this paper takes a new look at the filter design problem and attempts to address the following questions: (i) is sparsity an issue in filter design problems to explore and why is it important? and (ii) how a digital filter can be designed to have a sparse impulse response for efficient implementation while achieving improved performance relative to its non-sparse counterpart? In brief terms, our design method can be described as a two-phase algorithm in that, for a desired frequency response and an upper bound of filter order, the locations of impulse response components that can be set to zero with minimal impact on filter performance are identified and the number of such locations are maximized in design phase 1; and an optimal FIR filter (either in least-squares or minimax sense) subject to the sparsity identified in phase 1 is then designed in phase 2. Illustrations and technical details of the design method are given in Sections 2, 3, and 4, and design examples are presented in Section 5.

II. Observations and The Design Problem

For simplicity of presentation, throughout we examine a class of linear-phase FIR filters whose transfer functions assume the form of

$$H(z) = \sum_{i=0}^{N} h_i z^{-i} \tag{1}$$

with N even and $h_i = h_{N-i}$ for $i = 0,\ 1,\ \ldots,\ N/2$.

A. Is sparsity an issue in filter design problems?

An impulse response $\{h_i,\ i = 0,\ 1,\ \ldots,\ N\}$ is said to be *sparse* if a considerable number of h_i's are exactly equal to zero. An impulse response is said to be *K-sparse* if there are only K nonzero h_i's. The primary reason we are interested in digital filters with sparse coefficients is because this sparsity implies reduced implementation complexity, hence real-time application potential and cost effectiveness. Speculating on why sparsity has not been an explicit issue of research for digital filters in the past, we mention a generic observation that the impulse of a digital filter is typically not sparse, see e.g., Fig. 1. It is also observed, on the other hand, that usually an impulse response contains h_i's of small magnitude relative to a given threshold δ, see e.g. Fig. 1. Moreover, if two filters are designed to approximate a desired frequency response, then the impulse response of the filter of higher-order contains more small-magnitude h_i's than its lower-order counterpart, see e.g. Fig. 1b.

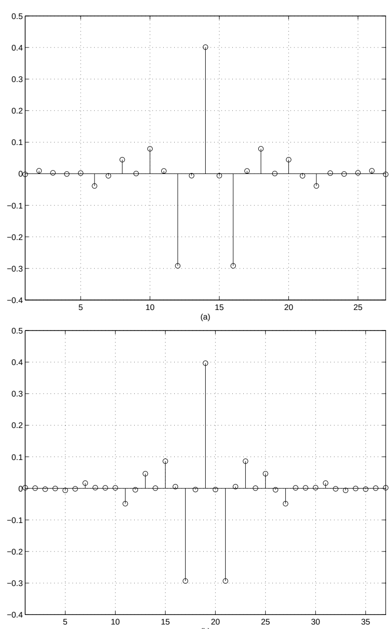

Fig. 1. (a) Impulse response of a bandpass FIR filter of order 26, no zero-valued components, there are $L = 4$ components with magnitude less than $\delta = 0.0017$; (b) impulse response of a bandpass FIR filter of order 36, no zero-valued components, there are $L = 12$ components with magnitude less than δ.

Suppose $\boldsymbol{h}$ contains L coefficients with magnitude less than δ, then a sparse $\hat{\boldsymbol{h}}$ can be constructed within a small vicinity of $\boldsymbol{h}$ by simply setting the L small coefficients to zero. This procedure is often called hard-thresholding in the literature. By viewing $\hat{\boldsymbol{h}}$ as the impulse of an FIR filter $\hat{H}(z)$, we see that $\hat{H}(z)$ has a sparse impulse response if L is not small, and the closeness between the two filters in the frequency domain is indicated by

$$\max_{\omega} |\hat{H}(e^{j\omega}) - H(e^{j\omega})| \leq L \cdot \delta \tag{2}$$

which can be verified using the Cauchy-Schwartz inequality.

B. A design formulation faithful to coefficient sparsity

Although (2) shows that $H(e^{j\omega})$ is well behaved for small coefficient variations, any changes made in the coefficients inevitably lead to performance degradation, making a sparse $\hat{H}(e^{j\omega})$ only suboptimal. For illustration clarity, we introduce a term *equivalent non-sparse filter*. Suppose a sparse $\hat{H}(z)$ is of order N with L zero coefficients (thus it has $K = N + 1 - L$ nonzero coefficients). An FIR $H_e(z)$ with a non-sparse impulse response is said to be *equivalent* to $\hat{H}(z)$ if $H_e(z)$ is designed to approximate the same desired frequency response (as $\hat{H}(z)$ does) and contains also K nonzero coefficients. Evidently, here the term "equivalent" is meant to require the same number of multiplications per output sample. As an example, Fig. 2a shows the amplitude response of a sparse bandpass $\hat{H}(z)$ of order $N = 36$ obtained by hard-thresholding (with $\delta = 0.0017$) an equiripple bandpass filter of order 36. There are $L = 12$ zero coefficients in $H(z)$. Fig. 2b shows the amplitude response of an equivalent non-sparse bandpass filter $H_e(z)$ of order 24 ($= N - L$). We see that the sparse $\hat{H}(z)$ fails to keep the equiripple property and its stopband attenuation is worse than its equivalent non-sparse counterpart.

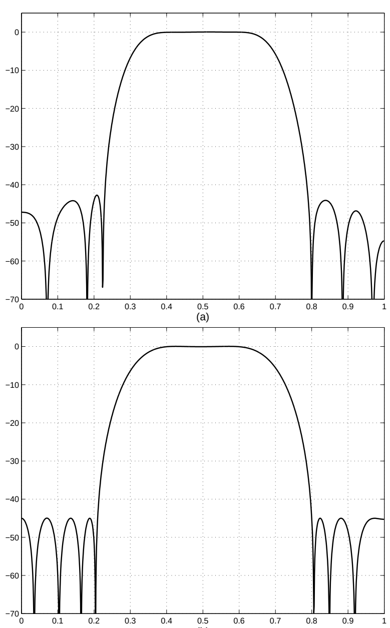

Fig. 2. (a) Amplitude response of a sparse bandpass $\hat{H}(z)$ of order 36 with $K = 25$ nonzero coefficients; (b) amplitude response of an equiripple non-sparse bandpass filter of order 24.

From the observations made above, it is quite clear that one needs a new formulation for filter design that is faithful to coefficient sparsity while maintaining optimal performance. In our problem formulation, an order upper bound $\hat{N}$ rather than a specific filter order is given. This is because sparsity is one of the design considerations and the sparsity depends heavily on filter length. We remark that the order upper bound may be determined by system's requirement such as the largest acceptable group delay, etc. The design objective is to obtain a linear-phase FIR filter with a K-sparse impulse response that optimally approximates a desired frequency response $H_d(\omega)$ in either a least-squares (LS) or a minimax sense.

III. The Design Method at a Glance

The design of an optimal filter with sparse coefficients is accomplished in two phases. The aim of the first phase is at identifying the locations where the filter coefficients should be set to zero to satisfy the sparsity requirement. Since one is dealing with an impulse response to approximate a given $H_d(\omega)$ and since one wants to enhance its sparsity, phase 1 of the design is achieved by minimizing a weighted sum:

$$\underset{\boldsymbol{a}}{\text{minimize}} \; \|H(e^{j\omega}) - H_d(\omega)\|_{2,\infty} + \mu\|\boldsymbol{a}\|_1 \tag{3}$$

where $\|\cdot\|_{2,\infty}$ denotes L_2-norm for (LS) or L_∞-norm (for minimax), vector $\boldsymbol{a}$ is related to the filter coefficients as $\boldsymbol{a} = [a_0 \;\; a_1 \;\; \cdots \;\; a_{\hat{N}/2}]^T$, and

$$H(e^{j\omega}) = \sum_{i=0}^{\hat{N}} h_i e^{-ji\omega} = e^{-j\hat{N}\omega/2} \sum_{i=0}^{\hat{N}/2} a_i \cos i\omega \tag{4}$$

and $\mu > 0$ is a scalar weight. The L_1-penalty term in (3) helps produce an impulse response that tends to be more sparse. This is based on a recent discovery that under certain conditions the sparsest solution of a underdetermined linear system $\boldsymbol{Ax} = \boldsymbol{b}$ can be found by minimizing L_1-norm $\|\boldsymbol{x}\|_1$ subject to $\boldsymbol{Ax} = \boldsymbol{b}$ [10]. Problem (3) is a convex problem which admits a unique solution $\boldsymbol{a}$ (hence $\boldsymbol{h}$) that can be found using an efficient solver such as SeDuMi. Hard-thresholding with an appropriate δ is then applied to $\boldsymbol{h}$ to yield an $\hat{\boldsymbol{h}}$ with the desired sparsity. Note that depending on the value of δ used, the dimension of $\hat{\boldsymbol{h}}$, N, may or may not be equal to $\hat{N}$, i.e., $N \leq \hat{N}$. In phase 2 of the design, a filter $H_s(z)$ of order N that optimally approximates $H_d(\omega)$ subject to the coefficient sparsity identified in phase 1 is designed. Suppose $\hat{\boldsymbol{h}}$ (obtained in phase 1) contains L zeros, thus the solution vector $\hat{\boldsymbol{a}}$ in phase 1 contains $L/2$ zeros. Let the locations of the zeros in $\hat{\boldsymbol{a}}$ be i_k for $k = 1, 2, \ldots, L/2$. Filter $H_s(z)$ is designed by solving the constrained problem

$$\underset{\boldsymbol{a}}{\text{minimize}} \quad \|H(e^{j\omega}) - H_d(\omega)\|_{2,\infty} \tag{5a}$$

$$\text{subject to:} \quad a_{i_k} = 0 \quad \text{for } k = 1, 2, \ldots, L/2 \tag{5b}$$

Note that in (5a) the L_1-penalty term has been dropped so that the filter is genuinely optimal while the constraints in (5b) ensure the coefficient sparsity. Also note that (5) is a convex problem. Details in solving (3) and (5) are given next.

IV. Algorithms for LS and Minimax Designs

A. An algorithm for weighted LS designs

Phase 1: Given an order upper bound $\hat{N}$, a desired $H_d(\omega) = e^{-j\hat{N}\omega/2}A_d(\omega)$ and desired sparsity K, and let $H(z)$ assume the form of (4). We can write the objective function in (3) as

$$J_2(\boldsymbol{a}) = \int_{-\pi}^{\pi} W(\omega)[\boldsymbol{a}^T\boldsymbol{c}(\omega) - A_d(\omega)]^2\, d\omega + \mu\|\boldsymbol{a}\|_1 \quad (6)$$

where $\boldsymbol{c}(\omega) = [1 \quad \cos\omega \quad \cdots \quad \cos\hat{N}\omega/2]^T$ and $W(\omega) \geq 0$ is a weighting function. The first term in (6) is quadratic in $\boldsymbol{a}$, hence up to a constant $J_2(\boldsymbol{a})$ can be written as

$$J_2(\boldsymbol{a}) = \boldsymbol{a}^T\boldsymbol{Q}\boldsymbol{a} + \boldsymbol{a}^T\boldsymbol{p} + \mu\|\boldsymbol{a}\|_1 \quad (7)$$

If we place a bound for each component of $\boldsymbol{a}$, i.e. $|a_i| \leq d_i$ and treat d_i's as auxiliary variables, then minimizing $J_2(\boldsymbol{a})$ in (7) can be formulated as

$$\begin{aligned} &\text{minimize} \quad && \boldsymbol{x}^T\hat{\boldsymbol{Q}}\boldsymbol{x} + \boldsymbol{x}^T\hat{\boldsymbol{p}} && (8a)\\ &\text{subject to:} \quad && \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{0} && (8b)\end{aligned}$$

where $\boldsymbol{x} = [a_0 \quad \cdots \quad a_{\hat{N}/2} \quad d_0 \quad \cdots \quad d_{\hat{N}/2}]^T$,

$$\hat{\boldsymbol{Q}} = \begin{bmatrix} \boldsymbol{Q} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \end{bmatrix},\ \hat{\boldsymbol{p}} = \begin{bmatrix} \boldsymbol{p} \\ \mu\boldsymbol{e} \end{bmatrix} \quad \text{and} \quad \boldsymbol{A} = \begin{bmatrix} \boldsymbol{I} & \boldsymbol{I} \\ -\boldsymbol{I} & \boldsymbol{I} \end{bmatrix}$$

with $\boldsymbol{I}$ the identity matrix of dimension $1 + \hat{N}/2$ and $\boldsymbol{e}$ the all-one vector of dimension $1 + \hat{N}/2$. Since $\boldsymbol{Q}$ is positive semidefinite, (8) is a standard convex quadratic programming problem which can be solved by efficient solvers including MathWorks' optimization toolbox and SeDuMi. The locations in the impulse response where zeros should be placed are identified by applying hard thresholding with an appropriate δ to the solution vector $\boldsymbol{a}$ (the first $1 + \hat{N}/2$ components of $\boldsymbol{x}$). The value of δ is set to yield L zeros in $\boldsymbol{h}$ so that the required sparsity $K = N + 1 - L$ is met.

Phase 2: Let the indices of zero components in $\boldsymbol{a}$ be given by i_k for $k = 1,\ 2,\ \ldots,\ L/2$. We now solve the constrained problem

$$\begin{aligned} &\text{minimize} \ \ e_2(\boldsymbol{a}) = \int_{-\pi}^{\pi} W(\omega)[\boldsymbol{a}^T\boldsymbol{c}(\omega) - A_d(\omega)]^2\, d\omega && (9a)\\ &\text{subject to:}\ a_{i_k} = 0 \quad \text{for } k = 1,\ 2,\ \ldots,\ L/2 && (9b)\end{aligned}$$

The constraints in (9b) can be eliminated by substituting (9b) into (9a). This leads to an unconstrained convex problem

$$\text{minimize}\, e_2(a) = \int_{-\pi}^{\pi} W(\omega)[\tilde{\boldsymbol{a}}^T\tilde{\boldsymbol{c}}(\omega) - A_d(\omega)]^2\, d\omega \quad (10)$$

where $\tilde{\boldsymbol{a}}$ is a "compressed" version of $\boldsymbol{a}$, generated by deleting its zero components, and $\tilde{\boldsymbol{c}}(\omega)$ is a vector with corresponding cosine functions. Evidently, (10) is quadratic and can be expressed (up to a constant) as

$$e_2(\tilde{\boldsymbol{a}}) = \tilde{\boldsymbol{a}}^T\tilde{\boldsymbol{Q}}\tilde{\boldsymbol{a}} + \tilde{\boldsymbol{a}}^T\tilde{\boldsymbol{p}} \quad (11)$$

where $\tilde{\boldsymbol{Q}} = \int_{-\pi}^{\pi} W(\omega)\tilde{\boldsymbol{c}}(\omega)\tilde{\boldsymbol{c}}^T(\omega)\, d\omega$ is positive definite. The unique minimizer of (11) is given by $\tilde{\boldsymbol{a}} = -\frac{1}{2}\tilde{\boldsymbol{Q}}^{-1}\tilde{\boldsymbol{p}}$ and the optimal sparse $\boldsymbol{a}$ is obtained by inserting $L/2$ zeros back into $\tilde{\boldsymbol{a}}$ at indices i_k for $k = 1,\ 2,\ \ldots,\ L/2$.

B. An algorithm for weighted minimax designs

Phase 1: The main difference from the LS design is that here we solve a minimax problem with a weighted L_1-penalty term:

$$\underset{\boldsymbol{a}}{\text{minimize}}[\max_{\omega} W(\omega)|\boldsymbol{a}^T\boldsymbol{c}(\omega) - A_d(\omega)| + \mu\|\boldsymbol{a}\|_1] \quad (12)$$

By introducing an upper bound η for the first term over a set of frequency grids Ω in the frequency region of interest, we convert (12) into

$$\begin{aligned} &\text{minimize} \quad && \eta + \mu\|\boldsymbol{a}\|_1 && (13a)\\ &\text{subject to:} && W(\omega)|\boldsymbol{a}^T\boldsymbol{c}(\omega) - A_d(\omega)| \leq \eta \text{ for } \omega \in \Omega && (13b)\end{aligned}$$

Next, we write $\boldsymbol{a} = \boldsymbol{u} - \boldsymbol{v}$ with $\boldsymbol{u} \geq \boldsymbol{0}$ and $\boldsymbol{v} \geq \boldsymbol{0}$. Vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ can be set as $\boldsymbol{u} = \max\{\boldsymbol{a},\ \boldsymbol{0}\}$ and $\boldsymbol{v} = \max\{-\boldsymbol{a},\ \boldsymbol{0}\}$. In this way, the L_1-norm of $\boldsymbol{a}$ can be expressed as $\|\boldsymbol{a}\|_1 = \boldsymbol{e}^T\boldsymbol{u} + \boldsymbol{e}^T\boldsymbol{v}$ and (13) becomes

$$\begin{aligned} &\text{minimize} \quad && \boldsymbol{c}^T\boldsymbol{x} && (14a)\\ &\text{subject to:} && |(\boldsymbol{u} - \boldsymbol{v})^T\boldsymbol{c}(\omega) - A_d(\omega)| \leq \eta \text{ for } \omega \in \Omega && (14b)\\ & && \boldsymbol{u} \geq \boldsymbol{0} \quad \text{and} \quad \boldsymbol{v} \geq \boldsymbol{0} && (14c)\end{aligned}$$

where $\boldsymbol{x} = [\eta \ \ \boldsymbol{u}^T \ \ \boldsymbol{v}^T]^T$ and $\boldsymbol{c} = [1 \ \ \mu\boldsymbol{e}^T \ \ \mu\boldsymbol{e}^T]^T$. We see that the objective function as well as the constraints in (14) are all linear, hence (14) is a linear programming (LP) problem which can be solved using e.g. MathWorks' optimization toolbox. The rest of phase 1 is identical to the counterpart of the LS algorithm described in Sec. 4A.

Phase 2: Here one solves the constrained minimax problem

$$\begin{aligned} &\underset{\boldsymbol{a}}{\text{minimize}} \quad && \max_{\omega} W(\omega)|\boldsymbol{a}^T\boldsymbol{c}(\omega) - A_d(\omega)| && (15a)\\ &\text{subject to:} && a_{i_k} = 0 \qquad \text{for } k = 1,\ 2,\ \ldots,\ L/2 && (15b)\end{aligned}$$

Like the LS algorithm, substituting (15b) into (15a) leads the above problem to an unconstraint minimax problem as

$$\underset{\boldsymbol{a}}{\text{minimize}} \max_{\omega} W(\omega)|\tilde{\boldsymbol{a}}^T\tilde{\boldsymbol{c}}(\omega) - A_d(\omega)| \quad (16)$$

The problem in (16) can in turn be converted to

$$\begin{aligned} &\text{minimize} \quad && \eta && (17a)\\ &\text{subject to:} && W(\omega)|\tilde{\boldsymbol{a}}^T\tilde{\boldsymbol{c}}(\omega) - A_d(\omega)| \leq \eta \text{ for } \omega \in \Omega && (17b)\end{aligned}$$

which is an LP problem when the upper bound η is treated as an auxiliary variable. The rest of phase 2 is identical to that of the LS algorithm in Sec. 4A.

V. Design Examples

The algorithms described in Sec. 4 were applied to design linear-phase FIR filters with sparse coefficients. Presented below are three examples for illustrating the proposed algorithms. In all examples, the weight $W(\omega)$ was set to one in the frequency bands of interest and zero elsewhere.

Example 1 An LS lowpass filter with normalized passband edge $\omega_p = 0.5$, stopband edge $\omega_a = 0.6$, sparsity $K = 23$, and order upper bound $\hat{N} = 32$ was designed by applying the algorithm in Sec. 4A. With $\mu = 0.004$ and $\delta = 0.006$, the algorithm yielded a filter of order $N = 28$ whose response contains $L = 6$ zeros, thus the sparsity requirement $K = N +$

$1 - L = 23$ is met. The L_2 approximation error achieved was 0.0005. For comparison, with the same design specifications (without sparsity) an equivalent non-sparse LS lowpass filter of order 22 was designed. The non-sparse filter yielded an L_2 approximation error 0.0012.

Example 2 A minimax lowpass filter with $\omega_p = 0.55$, $\omega_a = 0.6$, $K = 53$, and $\hat{N} = 64$ was designed by applying the algorithm in Sec. 4B. With $\mu = 0.07$, $\delta = 0.003$, and 190 equally spaced frequency grids over the passbnad and stopband, the algorithm yielded a filter of order $N = 64$ with $L = 12$ zero coefficients, thus satisfying $K = N + 1 - L = 53$. The L_∞ (maximum) error was found to be 2.1332×10^{-2}. An equivalent non-sparse equiripple FIR filter of order 52 with the same design specifications (without sparsity) was designed using the Parks-McClellan algorithm. The L_∞ error it achieved was 3.4657×10^{-2}. The amplitude responses of these filters are shown in Fig. 3.

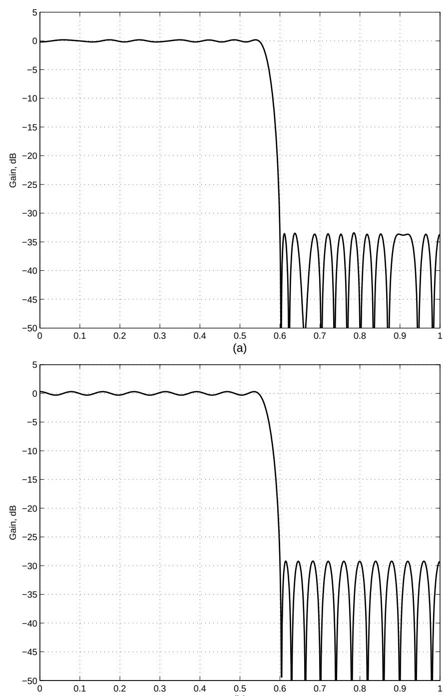

Fig. 3. Amplitude response of (a) the sparse minimax lowpass filter and (b) the equivalent non-sparse minimax lowpass filter in Example 2.

Example 3 A minimax bandpass filter with $\omega_{a1} = 0.265$, $\omega_{p1} = 0.4$, $\omega_{p2} = 0.6$, $\omega_{a2} = 0.73$, $K = 35$ and $\hat{N} = 64$ was designed using the algorithm in Sec. 4B. With $\mu = 0.04$, $\delta = 4 \times 10^{-4}$, and 220 equally spaced frequency grids over the passbnad and stopbands, the algorithm produced a filter of order $N = 64$ with $L = 20$ zero coefficients, thus meeting $K = N + 1 - L = 35$. The L_∞ error was found to be 2.5404×10^{-4}. An equivalent non-sparse equiripple bandpass filter of order 34 with the same design specifications (without sparsity) was designed using the Parks-McClellan algorithm. The L_∞ error it yielded was 6.0742×10^{-3}. The amplitude responses of these filters are depicted in Fig. 4.

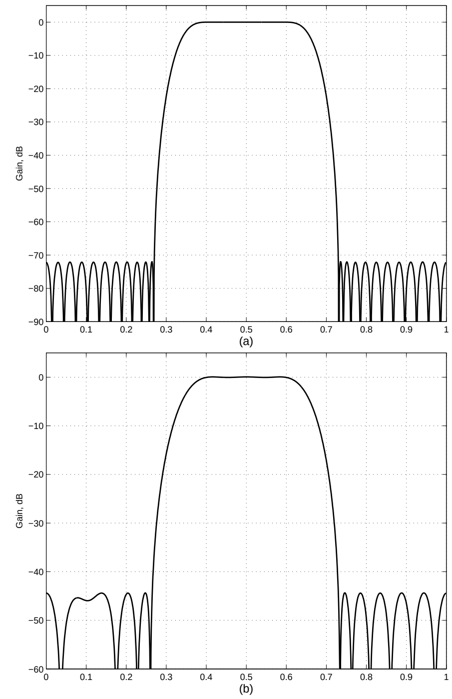

Fig. 4. Amplitude response of (a) the sparse minimax bandpass filter and (b) the equivalent non-sparse minimax bandpass filter in Example 3.

VI. Concluding Remarks

After this work was done, [11] was brought to the authors' attention, where sparsity of half-band like FIR filters was examined. We also remark that there exist several techniques for efficient implementation of FIR filters, e.g., the frequency response masking technique, and it shall be interesting to examine these techniques with a comparative study.

References

[1] B. Gold and C. M. Rader, *Digital Processing of Signals*, McGraw-Hill, New York, 1969.

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

[3] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, Prentice-Hall, Englewood Clliffs, NJ, 1989.

[4] S. K. Mitra, *Digital Signal Processing*, McGraw-Hill, New York, 1998.

[5] A. Antoniou, *Digital Signal Processing: Signals, Systems, and Filters*, McGraw-Hill, New York, 2005.

[6] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[7] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[8] E. J. Candes and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[9] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, Mar. 2008.

[10] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, 2009.

[11] O. Gustafsson, L. S. DeBrunner, V. DeBrunner, and H. Johansson, "On the design of sparse half-band like FIR filters," *41st Asilomar Conf.*, pp. 1098-1102, Nov. 2007.