Design of Optimum Multi-Dimensional Energy Compaction Filters

N. Damera-Venkata
Imaging Technology Dept.
Hewlett-Packard Labs
Palo Alto, CA

J. Tuqan
Dept. of ECE
Univ. of California
Davis, CA

B. L. Evans
Dept. of ECE
Univ. of Texas
Austin, TX

Abstract

We discuss the design of optimum signal-adapted multi-dimensional energy compaction filters. As in the one-dimensional (1-D) case, the energy compaction problem is linear in the auto-correlation coefficients of the compaction filter which must also satisfy the multi-dimensional (m-D) equivalent of the Nyquist-(M) condition. If a minimum-phase spectral factor exists the optimum compaction filter is recovered using the m-D Discrete Hilbert Transform (DHT). If a minimum phase spectral factor does not exist we propose an iterative algorithm based on multi-objective goal attainment. We try to enforce the Nyquist-$\mathbf{M}$ *condition while simultaneously forcing the autocorrelation coefficients of the compaction filter to be as close as possible to the coefficients of the product filter and the compaction gain of the optimum compaction filter to be close to the compaction gain produced by using the optimum product filter.*

1 Introduction

The problem of maximizing the function

$$\max_{H(e^{j\omega})} \frac{1}{2\pi} \int_{-\pi}^{\pi} |H(e^{j\omega})|^2 W(e^{j\omega}) \qquad (1)$$

subject to

$$\frac{1}{M} \sum_{k=0}^{M-1} |H(e^{j(\omega - 2\pi k/M)})|^2 = 1 \qquad (2)$$

has received significant attention in the past due to its occurrence in different signal processing problems including the design of optimal transmit and receive filters for data modems, design of orthonormal filter-banks and wavelets, identification of time-varying systems, echo cancellation, in the design of optimal quantization for a class of non-bandlimited signals and the design of optimal energy compaction filters (see [1][2] for references pertaining to these applications). Each of these applications essentially involves variations on the function chosen for $W(e^{j\omega})$. Fig. 1 shows the block diagram of the energy compaction problem. The goal is to design $H(e^{j\omega})$ so that the output variance is maximized subject to the constraint (2). Assuming that the input is a wide-sense-stationary (WSS) stochastic process we see that the energy compaction problem is a special case of the problem stated in (1) and (2) with $W(e^{j\omega}) = S_{xx}(e^{j\omega})$ representing the power spectrum of the input signal (note that decimation does not change the output variance). The constraint of (2) is called the Nyquist-M condition [3]. It can readily be seen that the optimum solution is function of the *product filter* $|H(e^{j\omega})|^2$ only. For one-dimensional signals once the optimum product filter is obtained, one-dimensional spectral factorization can be carried out to obtain an optimal solution for the compaction filter. Kirac and Vaidyanathan [1] give several methods for obtaining the optimal product filter coefficients, such as analytical methods, linear programming and a window based approach. Tuqan and Vaidyanathan [2]use a state space approach to find the globally optimal solution for the product filter and minimum phase spectral factor of the product filter without explicit spectral factorization. Thus $|H(z)|^2 = H_{min}(z)H_{max}(z)$ where $H_{min}(z)$ and $H_{max}(z)$ represent the minimum and maximum phase spectral factors respectively. In this paper we are interested in the solution of the multi-dimensional (m-D) version of the optimum energy compaction problem. Although, similar to the one-dimensional (1-D) case, the optimal solution may be regarded as a linear programming problem in the coefficients of the product filter, the subsequent spectral factorization is the hard part. The lack of a factorization algorithm for m-D polynomials makes this an ill-posed task. Thus obtaining an optimal product filter does not necessarily imply the solution of the energy compaction problem. In this paper we explore a couple of alternatives. First, if the product filter does indeed have a minimum phase spectral factor, we use the m-D version of the Discrete Hilbert Transform [4][5] (DHT) to obtain the optimum solution. This is an extension of the 1-D minimum phase filter de-

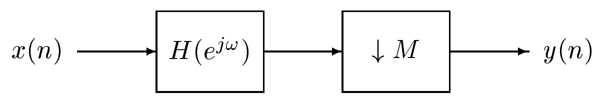

Figure 1: Block diagram illustrating the optimum energy compaction problem.

sign procedure using the 1-D DHT [6]. However in several cases a minimum phase spectral factor does not exist[7]. Application of the DHT to such a product filter results in a filter whose magnitude square response is neither Nyquist nor matches the desired product filter response. For such cases we use an iterative algorithm to solve for the energy compaction filter that is a close approximation to the actual magnitude response of the desired spectral factor. Section 2 sets up the m-D optimum energy compaction filter design problem and obtains a solution for the optimum product filter. Section 3 explores the use of the DHT in obtaining an optimum solution by finding the minimum phase spectral factor. Section 4 uses an iterative method to obtain an approximate solution in the case when the designed product filter does not have a minimum phase spectral factor. Finally Section 5 concludes the paper by summarizing the results. Without loss of generality we discuss the design problem for two-dimensions (2-D). The m-D problem is a straightforward extension of the 2-D case.

2 Optimum 2-D product filter design

In this section we set up the problem of finding an optimum 2-D product filter as a linear programming problem in the coefficients of the product filter. First, we convert the formulation to the spatial domain. The output variance σ_y^2 in the spatial domain (assuming a real stochastic input process $x(n_1, n_2)$ in terms of the product filter coefficients is

$$\begin{aligned} \sigma_y^2 &= \sum_{n_1=-N_1}^{N_1} \sum_{n_2=-N_2}^{N_2} p(n_1,n_2) r_{xx}(n_1,n_2) \qquad (3) \\ &= r_{xx}(0,0) + 2\sum_{n_2=1}^{N2} p(0,n_2) \\ &\quad +2 \sum_{n_1=1}^{N_1} \sum_{n_2=-N_2}^{N_2} p(n_1,n_2) r_{xx}(n_1,n_2) \end{aligned}$$

where we have used Parseval's identity [8] and the fact that $p(n_1, n_2) = p(-n_1, -n_2)$ and $r_{xx}(n_1, n_2) = r_{xx}(-n_1, -n_2)$ which represent the symmetry in the coefficients of the product filter and the signal autocorrelation function respectively. We have also assumed without loss of generality that $p(0,0) = 1$ [2]. We must maximize σ_y^2 subject to the following constraints on the product filter

$$p(\mathbf{M}[n_1 \ n_2]^T) = \delta(n_1, n_2) \qquad (4)$$

$$\begin{aligned} P(\omega_1,\omega_2) &= 1 + \sum_{n_2=1}^{N2} p(0,n_2)\cos(\omega_2 n_2) \qquad (5) \\ &+ 2\sum_{n_1=1}^{N_1} \sum_{n_2=-N_2}^{N_2} p(n_1,n_2) \cdot \\ &\quad \cos(\omega_1 n_1 + \omega_2 n_2) \\ &\geq 0 \ \forall \omega_1, \omega_2 \end{aligned}$$

The constraint (4) represents the 2-D Nyquist-$\mathbf{M}$ constraint. The matrix $\mathbf{M}$ is the two dimensional sampling matrix. The constraint (5) enforces the fact that the product filter frequency response represents a squared magnitude response and hence is constrained to be positive. Thus the optimum product filter may be obtained by solving the linear programming problem

$$\max \ \mathbf{r}_{\downarrow\mathbf{M}}^T \mathbf{p}_{\downarrow\mathbf{M}} \qquad (6)$$

subject to

$$\mathbf{c}_{\downarrow\mathbf{M}}(\omega_1,\omega_2)^T \mathbf{p}_{\downarrow\mathbf{M}} \geq -0.5 \qquad (7)$$

for the optimal product filter coefficient vector $\mathbf{p}_{\downarrow\mathbf{M}}$. The vector $\mathbf{r}_{\downarrow\mathbf{M}}$ is obtained by deleting the elements corresponding to the sampling lattice of the matrix $\mathbf{M}$ from the vector $\mathbf{r} = [r_{xx}(0,1)\cdots r_{xx}(0,N_2), r_{xx}(1,-N_2)\cdots r_{xx}(N_1,N_2)]^T$. The vector $\mathbf{c}_{\downarrow\mathbf{M}}(\omega_1,\omega_2)$ is given by deleting the elements corresponding to the sampling lattice of $\mathbf{M}$ from the vector $\mathbf{c}(\omega_1,\omega_2) = [\cos(\omega_2)\cdots\cos(\omega_2 N_2), \cos(\omega_1 - \omega_2 N_2)\cdots\cos(\omega_1 N_1 + \omega_2 N_2)]^T$. The constraints of (7) may be discretized by substituting $(\omega_1,\omega_2) = (e^{j\frac{2\pi k_1}{L_1}}, e^{j\frac{2\pi k_2}{L_2}})$, $k_1 = 0, 1, \ldots L_1 = 1, k_2 = 0, 1, \ldots L_2$ resulting in $L_1 L_2$ linear constraint equations. $\mathbf{p}_{\downarrow\mathbf{M}}$ represents the vector of product filter coefficients by deleting the elements corresponding to the sampling lattice of $\mathbf{M}$ from the vector $\mathbf{p} = [p(0,1)\cdots p(0,N_2), p(1,-N_2)\cdots p(N_1,N_2)]^T$. The product filter $p(n_1,n_2)$ is completely determined by $\mathbf{p}_{\downarrow\mathbf{M}}$ since it may be generated from $\mathbf{p}_{\downarrow\mathbf{M}}$ using the symmetry of the product filter and upsampling using the matrix $\mathbf{M}$.

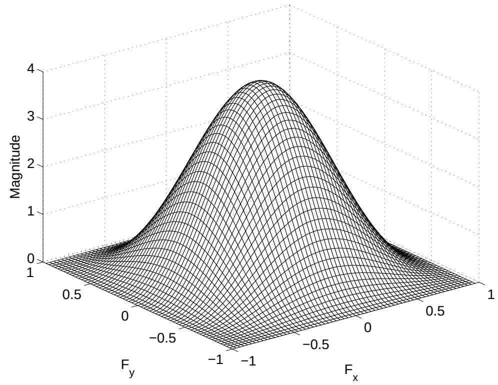

(a)

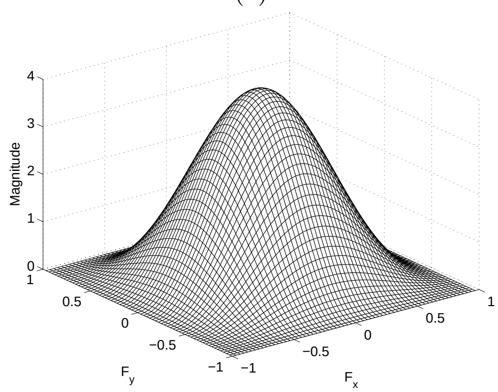

(b)

Figure 2: Frequency responses of (a) optimum product filter and (b) the product filter corresponding to the optimum compaction filter of example 1, designed using the DHT.

3 Optimum 2-D compaction filter design using the Discrete Hilbert Transform

The magnitude response of the product filter is the squared magnitude response of the optimal compaction filter. If a minimum phase spectral factor exists, its magnitude response may be obtained from the magnitude response of the product filter since

$$|H_{min}[k_1,k_2]| = \sqrt{|P[k_1,k_2]|} \qquad (8)$$

The 2-D Discrete Hilbert Transform (DHT) relates the magnitude spectrum of a 2-D sequence to its minimum phase spectrum if it exists. The minimum phase spectrum exists if the complex cepstrum is causal. This is the assumption we make when we apply the DHT. Once the minimum phase spectrum is determined we reconstruct the minimum phase polynomial by combining the desired magnitude and phase responses, and take the inverse FFT. We use the 2-D DHT developed by Read an Treitel [4]. Extension to higher dimensions is given in [5].

Given a sampled magnitude spectrum $|\ P[k_1,k_2]\ |$ where $k_1 = 0,1,\ldots K_1$ and $k_2 = 0,1,\ldots K_2$ computed with a $K_1 \times K_2$ FFT , we compute the corresponding minimum phase spectral factor in two steps. First, we compute the sampled minimum phase spectrum

$$\begin{aligned}\theta[k_1,k_2] \quad &= \quad -j\text{DFT}\{\ s[m_1,m_2] \bullet \\ &\qquad \text{IDFT}\{\ln(\sqrt{|P[k_1,k_2]|} + \delta)\}\}\end{aligned} \qquad (9)$$

where $\bullet$ represents pointwise matrix multiplication and

$$s[m_1,m_2] = \begin{cases} 1 & m_1 = 0, 0 < m_2 < \frac{K_2}{2} \\ 1 & m_2 = 0, 0 < m_1 < \frac{K_1}{2} \\ -1 & m_1 = 0, 0 < m_2 < \frac{K_2}{2} \\ -1 & m_2 = 0, 0 < m_1 < \frac{K_1}{2} \\ 0 & else \end{cases} \qquad (10)$$

δ is a small number used to avoid zero as an argument to the logarithm. Second, we form $|\ H_{min}[k_1,k_2]\ |\ e^{j\,\theta[k_1,k_2]}$. The optimum energy compaction filter $h[n_1,n_2]$ may be constructed by

1. computing $|\ H_{min}[k_1,k_2]\ |$ which is the sampled magnitude spectrum of $h[n_1,n_2]$,
2. calculating $\theta[k_1,k_2]$ by using (9),
3. constructing the $K_1 \times K_2$ FFT of the minimum phase sequence as $|\ H_{min}[k_1,k_2]\ |\ e^{j\,\theta[k_1,k_2]}$,
4. taking the inverse FFT transform to obtain a minimum phase sequence, and
5. truncating the resulting sequence to the desired filter impulse response to $N_1 + 1 \times N_2 + 1$.

Example 1 : We use a 5×5 signal autocorrelation function given by $r_{xx}[n_1,n_2] = e^{-0.5\sqrt{(n_1^2+n_2^2)}}$ for all examples in this paper. Fig. 2(a) shows the optimum product filter response obtained assuming $\mathbf{M} = [[2\ 0]^T\ [0\ 2]^T]$ (rectangular sampling) and a 100×100 grid sampling. The 3×3 minimum phase spectral factor obtained using 512×512 DHT and truncating had a squared magnitude response given in Fig. 2(b) showing that a optimum compaction filter was indeed obtained. The compaction gain given by $G = (\sum r_{xx}(n_1,n_2)p(n_1,n_2))/(\sum r_{xx}(n_1,n_2)^2)$ for both cases was 0.5330.

Example 2 : Fig. 3(a) shows the optimum product filter response obtained assuming $\mathbf{M} = [[1\ 1]^T\ [1\ -1]^T]$ (quincunx sampling) and a 100×100 grid sampling. The optimum product filter had a compaction gain of $G = 0.3320$. The 3×3 minimum phase spectral

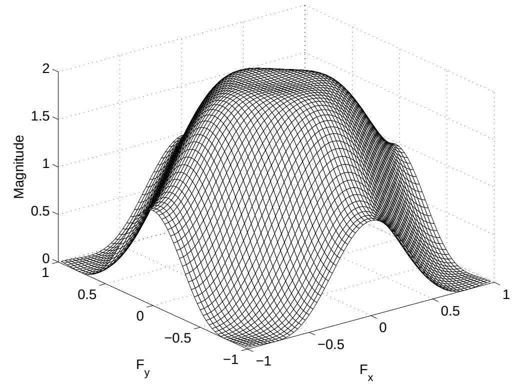

(a)

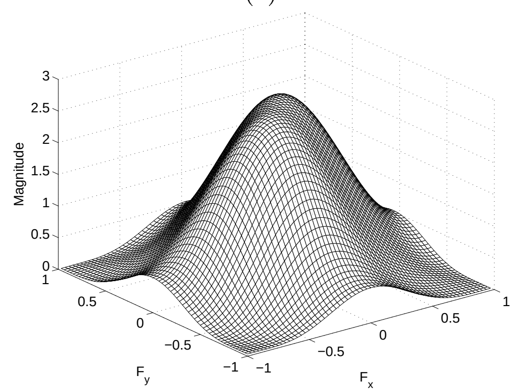

(b)

Figure 3: Frequency responses of (a) optimum product filter and (b) the product filter corresponding to the "optimum compaction filter" of example 2, designed using the DHT .

factor obtained using 512×512 DHT and truncating had a squared magnitude response shown in Fig. 3(b). In this case the DHT failed to produce an optimum compaction filter. In fact it did not produce a valid compaction filter at all (since it resulted in a product filter with significant coefficient values at the lattice locations of $\mathbf{M}$ of the order of 10^{-1}). In Section 4 we will obtain an iterative solution to this problem and find a valid optimum compaction filter with magnitude squared response close to the response of the desired product filter.

4 Iterative compaction filter design

In this section we consider the case when the designed product filter does not have a minimum phase spectral factor. In this case, as seen in section 3 the DHT method fails to produce an optimum compaction filter in general. We note that the coefficients of the product filter should be equal to the autocorrelation coefficients of the desired optimum compaction filter $h(n_1, n_2)$. This results in the following $S = 4N_1N_2 + 2N_1 + 2N_2 + 1$ nonlinear equations in the coefficients of the compaction filter.

$$\sum_{(m_1,m_2)\in\mathcal{S}} h(m_1, m_2)h(m_1 + n_1, m_2 + n_2) = p(n_1, n_2) \tag{11}$$

where $n_1 = -N_1 \ldots N_1$ and $n_2 = -N_2 \ldots N_2$ and $\mathcal{S}$ represents the support of the filter $h(n_1, n_2)$. These equations may be solved using the Levenberg-Marquardt algorithm [9]. Several initial guesses must be tried to avoid local minima. Further if an exact solution may not exist. A least squares solution is obtained. In this case the designed filter would not satisfy the Nyquist-$\mathbf{M}$ condition exactly so would not be a valid compaction filter. We formulate the optimum design problem as a multi-objective goal attainment problem [10]. This involves expressing the set of autocorrelation equations in the form of a vector of objective functions with the right hand side as the design goals. Thus we form the vectors $f(\mathbf{h}) = [f_0(\mathbf{h})f_1(\mathbf{h})\cdots f_{S-1}(\mathbf{h})]$ and $f^* = [f_0^* f_1^* \cdots f_{S-1}^*]$ for a total of S objective functions and design goals. Since the product filter is Nyquist-$\mathbf{M}$ the corresponding elements of f^* are zero. We add the objective function $f_S(\mathbf{h}) = f_{\downarrow\mathbf{M}}^T(\mathbf{h})\mathbf{r}_{\downarrow\mathbf{M}}$ and the corresponding design goal $f_S^* = \mathbf{r}_{\downarrow\mathbf{M}}^T\mathbf{p}_{\downarrow\mathbf{M}}$. This objective expresses the energy compaction of the optimum product filter in terms of the filter coefficients of the energy compaction filter to be designed. The energy compaction filter design may then be expressed as

$$\min_{\gamma,\mathbf{h}} \gamma \tag{12}$$

subject to:

$$f_i(\mathbf{h}) - w_i\,\gamma = f_i^*,\ i = 0,\ 1,\ \ldots S \tag{13}$$

The problem formulation allows the objectives to be under or over achieved enabling the designer to be relatively imprecise about initial design goals. The relative degree of under or over-achievement of the goals is controlled by a vector of weighting coefficients $\mathbf{w} = [w_0 w_2 \ldots w_S]^T$. The term γ introduces an element of slackness into the problem, which otherwise imposes that the goals be rigidly met. The weighting vector, $\mathbf{w}$, enables the designer to express a measure of the relative trade-offs between the objectives. For the optimum energy compaction filter design problem, we use unit weights for the objectives corresponding to the zero coefficients of the product filter. The other weights are set to relatively high values (we use $100 \times |p(n_1, n_2)|$ to have the same percentage under or over attainment of the goals at the non-zero coefficients of $p(n_1, n_2)$). Thus this formulation yields valid

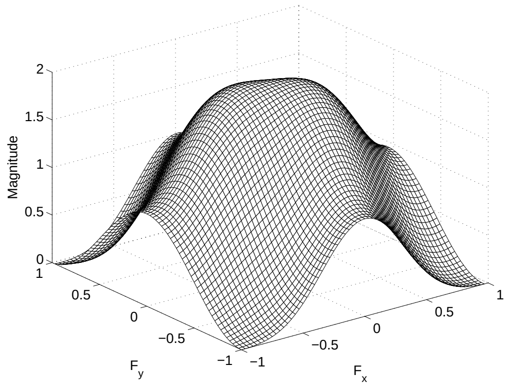

Figure 4: Frequency response of the product filter corresponding to the optimum compaction filter of example 3 designed using multi-objective optimization.

compaction filters satisfying the Nyquist-**M** condition with compaction gain close as possible to that of the optimum product filter.

Example 3 : We attempt to solve the optimum compaction filter design problem posed in example 2. Fig. 4 shows the magnitude squared response of the optimum compaction filter designed using the multi-objective optimization approach presented in this section. A least squares solution of the autocorrelation equations was used as the initial guess. This solution did not satisfy the Nyquist-**M** condition (since it resulted in a product filter with significant coefficient values at the lattice locations of **M** of the order of 10^{-1}). The multi-objective optimization resulted in a valid compaction filter with relatively small coefficient values at the lattice locations of **M** (of the order of 10^{-4}). The compaction gain of the resulting compaction filter was 0.3224, which is close to the compaction gain of the optimum product filter of 0.3320.

5 Conclusions

We present two methods for the design of optimum signal-adapted multi-dimensional energy compaction filters. First we solve for the optimum product filter using linear programming. The compaction filter must then be obtained from the product filter by spectral factorization. If a minimum-phase spectral factor exists the optimum compaction filter is recovered using the m-D Discrete Hilbert Transform (DHT). An iterative algorithm is proposed when a minimum-phase spectral factor does not exist. However the iterative algorithm only achieves a locally optimal solution.

References

[1] A. Kirac and P. P. Vaidyanathan, “Theory and design of optimum FIR compaction filters”, *IEEE Transactions on Signal Processing*, vol. 46, pp. 903–919, April 1998.

[2] J. Tuqan and P. P. Vaidyanathan, “A state space approach to the design of globally optimal FIR energy compaction filters”, *IEEE Transactions on Signal Processing*, vol. 48, pp. 2822–2838, October 2000.

[3] P. P. Vaidyanathan, *Multirate systems and filter banks*, Englewood Cliffs, NJ: Prentice-Hall, 1993.

[4] R. R. Read and S. Treitel, “The stabilization of two-dimensional recursive filters via the Discrete Hilbert Transform,” *IEEE Transactions on Geoscience Electronics*, vol. GE-11, pp. 153–60, Jul. 1973.

[5] N. Damera-Venkata, M. Venkataraman, M. S. Hrishikesh, and P. S. Reddy, “ A new transform for the stabilization and stability testing of multidimensional recursive digital filters”, *IEEE Transactions on Circuits and Systems II*, accepted for publication.

[6] N. Damera-Venkata, B. L. Evans and S. R. McCaslin, “Design of optimal minimum phase digital FIR filters using the discrete Hilbert transform”, *IEEE Transactions on Signal Processing*, vol. 48, no. 5, pp. 1491-1495, May 2000.

[7] J. W. Woods, “Comments on the stabilization of two dimensional recursive digital filters via the Hilbert transform”, *IEEE Transactions on Geoscience Electronics*, pp.104-105. July 1974.

[8] D. Dudgeon and R. Mersereau, *Multidimensional Digital Signal Processing*, Englewood Cliffs, New Jersey, Prentice Hall, 1984.

[9] D. M. Bates and D. G. Watts, *Nonlinear Regression and its Applications*, New York: Wiley, 1988

[10] F. W. Gembicki, *Vector Optimization for Control with Performance and Parameter Sensitivity Indices*, Ph.D. Dissertation, Case Western Reserve Univ., Cleveland, Ohio, 1974.