


Enriching the Art of FIR Filter Design via Convex Optimization

[Efficient evaluation of inherent tradeoffs]

Finite impulse response (FIR) filters have played a central role in digital signal processing since its inception. As befits that role, a myriad of design techniques is available, ranging from the quite straightforward windowing and frequency-sampling techniques to some rather sophisticated optimization-based techniques; e.g., [1]–[7]. Among the most prominent optimization-based techniques is the Parks-McClellan algorithm [8] for the design of “equiripple” linear phase FIR filters. One of the key features of that technique is the efficiency of the underlying Remez exchange algorithm. However, computing resources have grown more plentiful since the Parks-McClellan algorithm was developed [9], and this has spawned the development of more flexible design methodologies. Of particular note are METEOR [10] and the peak-constrained least-squares (PCLS) approach [11], [12]. METEOR is a flexible platform for FIR filter design problems that can be formulated as the optimization of a linear objective subject to linear constraints; i.e., as a linear program. One such problem is the design of a linear-phase low-pass filter with a “ripple” constraint in the pass-band, a constraint on the stop-band level, and the constraint that the pass-band response be a concave function of frequency. The PCLS approach provides efficient constraint exchange algorithms for finding filters that minimize a “least squares” approximation error subject to linear constraints; i.e., solve a quadratic program. One example is the design of a low-pass filter that minimizes the stop-band energy subject to a bound on the stop-band level.

Linear and quadratic programs are two of the simpler forms of convex optimization problem, and effective algorithms for solving them have been available for some time. Around the time that METEOR



© BRAND X PICTURES

and the PCLS approach were developed, breakthroughs were being made in the development of algorithms for solving more general convex optimization problems, and the subsequent developments have expanded the class of optimization problems that can be efficiently solved [13]. Furthermore, general-purpose implementations of those algorithms are readily available [14], [15], and some of the recently developed interfaces to those implementations are remarkably easy to use [16], [17].

An outcome of that development is that there is a rich class of filter design problems that can be optimally solved in a manner that is reliable, efficient, and requires little programming effort. Indeed, as also observed in [18]–[20], these convex optimization tools constitute a flexible platform for FIR filter design that captures the spirit of METEOR, but encapsulates a richer class of design problems. However, the impact of convexity extends beyond the solution of individual design problems to the interactive procedure that typifies the art of filter design. Effective FIR filter design requires judicious compromises to be made between competing properties of the filter; e.g., [11]. Convex optimization enriches that design process by enabling efficient computation of some of the inherent tradeoffs in FIR filter design; that is, fundamental tradeoffs that cannot be exceeded by any design method. The resulting tradeoff curves enable designers to quantify the extent to which certain desirable properties of a filter must be compromised to improve other aspects of the filter. As an example, insight from convex-optimization-based tradeoffs underlies the somewhat simpler tradeoff that was employed in the design of a spectral shaping filter in the HomePlug AV standard for home networking over power lines [21].

Of course, there is a significant number of filter design problems that are not convex; perhaps most notably those involving quantized filter coefficients. For these problems, finding a globally optimal filter is often a computationally overwhelming task, and hence the designer's judgment typically plays an even greater role in the design process. Convex optimization can also enhance that design process, through the generation of useful starting points for local optimization techniques, the ability to evaluate locally optimal solutions against inherent tradeoffs or bounds thereon, and the development of bounds for use in branch-and-bound algorithms [22] for globally optimal filters.

FIR FILTERS

An FIR filter of length L can be represented by its impulse response, $h[n]$, $0 \leq n \leq L - 1$, or by its frequency response

$$H(e^{j\omega}) = \sum_{n=0}^{L-1} h[n] e^{-j\omega n}. \quad (1)$$

We will find it convenient to represent the impulse response using the vector \mathbf{h} , the elements of which we will index from zero so that the n th element, $[\mathbf{h}]_n$, is $h[n]$, $0 \leq n \leq L - 1$. This enables us to write the frequency response as

$$H(e^{j\omega}) = \mathbf{v}(\omega)^H \mathbf{h}, \quad (2)$$

where $[\mathbf{v}(\omega)]_n = e^{j\omega n}$, and the superscript $(\cdot)^H$ denotes the conjugate transpose. For simplicity, we will focus on filters with real-valued coefficients, $\mathbf{h} \in \mathbb{R}^L$, but much of our discussion extends to filters with complex coefficients.

A popular class of FIR filters is the class of linear-phase filters. Phase linearity can be achieved by ensuring that the impulse response is either symmetric or antisymmetric about its midpoint. Linear-phase filters can be represented in a generic way [1], [3], with $\tilde{\mathbf{h}}$ denoting one half of the impulse response, and with the frequency response taking the form $H(e^{j\omega}) = e^{j\Theta(\omega)} \tilde{H}(e^{j\omega})$, where the amplitude response

$$\tilde{H}(e^{j\omega}) = \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \quad (3)$$

is real valued, $\Theta(\omega) = \theta_0 + \mu\omega$ with θ_0 and μ depending only on the filter length and symmetry, and $\tilde{\mathbf{v}}(\omega)$ is defined implicitly. The superscript $(\cdot)^T$ denotes the transpose. For odd-length filters with even symmetry, which are often called Type I linear-phase filters, $\theta_0 = 0$, $\mu = -(L - 1)/2$, $[\tilde{\mathbf{v}}(\omega)]_0 = 1$, and for $n > 0$, $[\tilde{\mathbf{v}}(\omega)]_n = 2\cos(n\omega)$.

In addition to the frequency and amplitude responses, the power spectrum of the filter, $|H(e^{j\omega})|^2$, also arises naturally in a number of design contexts. Using (2) we can write

$$|H(e^{j\omega})|^2 = \mathbf{h}^T \mathbf{v}(\omega) \mathbf{v}(\omega)^H \mathbf{h}, \quad (4)$$

and hence for each frequency the power spectrum is a convex quadratic function of \mathbf{h} . The power spectrum can also be written as a linear function of the autocorrelation of $h[n]$,

$$r_h[k] = \sum_n h[n] h[n + k]. \quad (5)$$

In particular, since $R_h(e^{j\omega}) = |H(e^{j\omega})|^2$, and $r_h[k]$ is a Type I linear-phase filter,

$$|H(e^{j\omega})|^2 = \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h. \quad (6)$$

GENERIC FORMULATION OF FILTER DESIGN PROBLEMS

The design of an FIR filter is often formulated as a constrained optimization problem. Typically, that problem is constructed in two stages. First, the characteristics of desirable filters are described in terms of constraints. Then a performance metric is selected and we seek the best of the desirable filters according to that metric. For example, the set of desirable filters might be those for which the magnitude response lies within the spectral mask in Figure 1 and the performance metric might be the “least-squares” error between the obtained frequency response and a desired response.

The description of the characteristics of desirable filters may include both inequality and equality constraints. We will write each of the inequality constraints in the form $f_m(\mathbf{x}) \leq \xi_m$, where m is the index of the constraint, \mathbf{x} represents the vector of

design variables, which will include one of h , \tilde{h} and \tilde{r}_h , $f_m(\cdot)$ describes the characteristic of the filter, and ξ_m is the value below which $f_m(\cdot)$ must lie for the filter to be deemed desirable. (Lower-bound constraints can be written in this form by multiplying both sides by -1 .) Each of the equality constraints will be written in the form $g_q(\mathbf{x}) = \zeta_q$.

The set of vectors \mathbf{x} that satisfy the constructed constraints represents the set of filters that are desirable in the sense that they satisfy the characteristics specified by the values of ξ_m and ζ_q . Often there is a primary characteristic, say $f_{m_0}(\mathbf{x})$, that is to be optimized over the set of desirable filters, and in that case the generic design problem can be written as

$$\min_{\mathbf{x}, \gamma} \gamma \tag{7a}$$

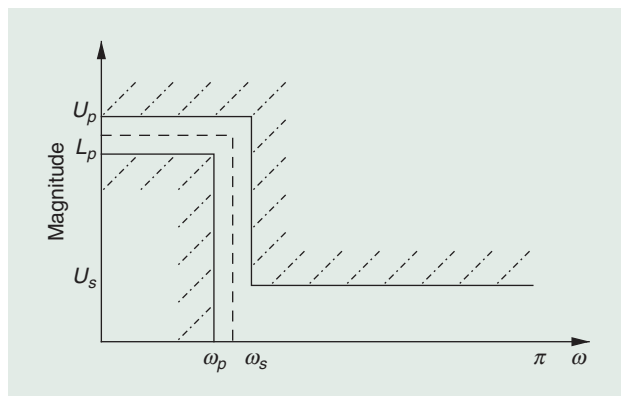
$$\text{subject to } f_{m_0}(\mathbf{x}) \leq \gamma, \tag{7b}$$

$$f_m(\mathbf{x}) \leq \xi_m, \quad m \neq m_0 \tag{7c}$$

$$g_q(\mathbf{x}) = \zeta_q, \tag{7d}$$

where we have left the ranges of the integers m and q implicit. Although the objective, $f_{m_0}(\mathbf{x})$, is also left implicit, a useful interpretation of (7) is that it seeks the tightest version of the m_0 th inequality such that there exists a filter that satisfies the constraints. This interpretation will allow us to handle some mild variations of (7) in a straightforward way; see the section “Variations on the Theme of (7).” One natural variant is the problem of optimizing a weighted sum of several characteristics.

In general, problems of the form in (7) can be quite difficult to solve. Globally optimal solutions can be obtained using branch-and-bound methods [22], but those methods are typically rather time consuming. An alternative is to seek locally optimal solutions by applying a sequential algorithm to appropriately selected starting points [23].



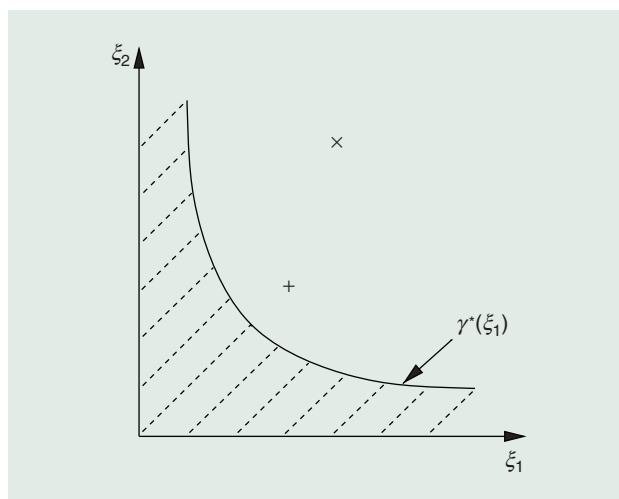
[FIG1] A low-pass filter mask (solid) and a desired response (dash-dot).

EFFECTIVE FIR FILTER DESIGN REQUIRES JUDICIOUS COMPROMISES TO BE MADE BETWEEN COMPETING PROPERTIES OF THE FILTER.

The situation is quite different in the case that each $f_m(\mathbf{x})$ in (7) is a convex function of \mathbf{x} and each $g_q(\mathbf{x})$ is linear. In that case, the problem in (7) is convex [13], and any locally optimal solution is globally

optimal. For a large class of such problems there are some analytical results that provide insight into the nature of the optimal solution [13], and recently developed general-purpose software tools [14]–[17] enable globally optimal solutions to be efficiently obtained with little programming effort. These tools also provide reliable detection of infeasibility; i.e., when there is no filter of the chosen length that satisfies all the constraints. That enables us to accommodate performance metrics that are only quasi convex [13]; see the section “Variations on the Theme of (7).” Many filter design problems have rather sparse constraints, and some of these general-purpose tools are able to exploit that sparsity for computational and numerical advantage. Although some filter design problems possess additional structure that can be exploited by custom designed software, with contemporary computing technology the approach that we will highlight typically enables interactive design of filters with several hundred taps.

When the problem in (7) is convex, it also provides an efficient, flexible framework for examining inherent tradeoffs in FIR filter design; that is, fundamental tradeoffs that cannot be exceeded by any design method. As an example, consider the simple case in which there are just two inequality constraints. As illustrated in Figure 2, there is a curve that partitions the (ξ_1, ξ_2) plane into a region containing pairs for which there exists a filter that satisfies both constraints, and pairs for which there is no feasible filter. As is implicit in the notation $\gamma^*(\xi_1)$, each point on this curve can be obtained by solving the problem in (7) with $f_2(\mathbf{x})$ chosen as the objective (i.e., $m_0 = 2$) and ξ_1



[FIG2] The inherent tradeoff between ξ_1 and ξ_2 for a generic convex problem. All pairs (ξ_1, ξ_2) on or above the curve can be achieved using a filter of the given length, and no pairs below the curve can be achieved.

fixed at the desired value. When (7) is convex, each of those problems can be efficiently solved. In contrast, when the problem in (7) is not convex, the potential for multiple locally optimal solutions typically means that the tradeoff curve that can be obtained with reasonable computational effort is the tradeoff curve achieved by the chosen design method, rather than the inherent tradeoff.

The potential for multiple locally optimal solutions means that the design of filters subject to nonconvex constraints typically involves considerable judgment on behalf of the designer. For example, if the optimization approach involves the application of a sequential algorithm to a number of starting points, the designer must eventually decide when the best of the local solutions that has been obtained so far is “good enough.” As illustrated in “Using an Inherent Tradeoff in a Nonconvex Design,” in some cases insight from inherent tradeoffs generated by convex design problems can enhance that decision process.

Now that we have identified some of the advantages of filter design problems in which each constraint yields a convex feasible set, a natural question to ask is how rich is the corresponding class of filter design criteria. We will provide some examples in the following two sections, but at this point it is worth pointing out that bounds on linear functions of \mathbf{x} yield convex feasible sets, as do upper bounds on convex quadratic functions of \mathbf{x} , but lower bounds on convex quadratic functions of \mathbf{x} , upper bounds on nonconvex quadratic constraints, and quadratic equality constraints do not.

Although many filter design criteria can be written as linear or convex quadratic constraints, the class of convex criteria is far more diverse; e.g., [13] and [17]. In some cases, that

THE POTENTIAL FOR MULTIPLE LOCALLY OPTIMAL SOLUTIONS MEANS THAT THE DESIGN OF FILTERS SUBJECT TO NONCONVEX CONSTRAINTS TYPICALLY INVOLVES CONSIDERABLE JUDGMENT ON BEHALF OF THE DESIGNER.

convexity is not immediately apparent from the natural parameterization of the design problem in terms of the impulse response, h . Indeed, there are several interesting problems in which the desirable characteristics of the filter result in nonconvex constraints

on h that can be rewritten as convex constraints on the autocorrelation vector, $\tilde{\mathbf{r}}_h$. Examples include filter design problems with constraints on the magnitude response [24], [25], root-Nyquist filters (e.g., [26] and [27]), and wavelets and multirate filter banks [28]–[31]. However, before we solve the corresponding version of the problem in (7), we must ensure that $\tilde{\mathbf{r}}_h$ corresponds to a valid autocorrelation. Since $R_h(e^{j\omega}) = |H(e^{j\omega})|^2$, this can be ensured by requiring that

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h \geq 0 \text{ for all } \omega \in [0, \pi]. \quad (8)$$

Although this constraint is linear in $\tilde{\mathbf{r}}_h$, it is semi-infinite, in the sense that there is one constraint for each $\omega \in [0, \pi]$. Some techniques for tackling such constraints are discussed in “Finite Representations of Spectral Mask Constraints.” Once the optimal autocorrelation has been found, a filter that generates that autocorrelation can be extracted using standard spectral factorization techniques [25], [32]. Although the minimum phase spectral factor is often a reasonable choice, if a secondary objective is available, it can be used to select the most appropriate spectral factor [33].

In the following two sections we will provide some examples of common filter design criteria that yield convex constraints on at least one of h , \tilde{h} , and $\tilde{\mathbf{r}}_h$. These examples illustrate the flexibility provided by the convex design platform: For any collection of constraints that are convex in the chosen design variable, the problem in (7) is convex and is amenable to the powerful algorithms and analytical tools discussed above. In a later section, we will provide a design example in which we take advantage of that flexibility, and we will provide a simple implementation using MATLAB and CVX [17].

FREQUENCY-DOMAIN CRITERIA

We now provide some examples of frequency-domain design criteria that yield convex constraints on h , \tilde{h} , or $\tilde{\mathbf{r}}_h$.

SPECTRAL MASKS

In many applications it is desirable that the filter satisfy a “spectral mask” of the form

$$L(e^{j\omega}) \leq |H(e^{j\omega})| \leq U(e^{j\omega}) \text{ for all } \omega \in [0, \pi], \quad (9)$$

for some given (nonnegative) bounds $L(e^{j\omega})$ and $U(e^{j\omega})$, respectively. Although the semi-infinite nature of this constraint can be tackled using the techniques described in “Finite Representations of Spectral Mask Constraints,” this constraint is

USING AN INHERENT TRADEOFF IN A NONCONVEX DESIGN

Consider a generic design problem in which the goal is to find a filter that achieves small values of the functions $f_1(\mathbf{x})$ and $f_2(\mathbf{x})$, which are convex, and also satisfies a third, nonconvex, constraint. This problem is somewhat representative of the problem of designing a low-pass filter with small stop-band energy, a low stop-band level, and quantized coefficients. Although the corresponding version of the problem in (7) is not convex, if that problem is “relaxed” by removing the nonconvex constraint, the resulting problem becomes convex and the inherent tradeoff between ξ_1 and ξ_2 can be efficiently obtained, as illustrated in Figure 2. If the chosen algorithm for the original nonconvex problem achieves the point in Figure 2 marked by the + then it might be deemed to be “good enough,” because it lies close to the “knee” point of the inherent tradeoff for filters that are not required to satisfy the nonconvex constraint. This would enable the design process to be terminated at that point. In contrast, a filter that achieves the × in Figure 2 is less likely to be deemed good enough, at least in the early phases of the design process.

awkward to deal with because the feasible set for the lower bound is not convex, unless $L(e^{j\omega})$ is identically zero. To make that more apparent, we can rewrite (9) as

$$L(e^{j\omega})^2 \leq \mathbf{h}^T \mathbf{v}(\omega) \mathbf{v}(\omega)^H \mathbf{h} \leq U(e^{j\omega})^2 \quad \text{for all } \omega \in [0, \pi]. \quad (10)$$

At each frequency, this corresponds to upper and lower bounds on a convex quadratic function of \mathbf{h} . The upper bound yields a convex feasible set, whereas the lower bound does not; cf. [13].

In the case of linear-phase filters, the mask is usually applied to the amplitude response: $\tilde{L}(e^{j\omega}) \leq \tilde{H}(e^{j\omega}) \leq \tilde{U}(e^{j\omega})$ for all $\omega \in [0, \pi]$, where $\tilde{L}(e^{j\omega})$ and $\tilde{U}(e^{j\omega})$ are not restricted to being nonnegative. Since $\tilde{H}(e^{j\omega}) = \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}}$, this mask generates two semi-infinite linear constraints on $\tilde{\mathbf{h}}$; e.g., [10].

In the case of autocorrelation-based designs, we can use the fact that $|H(e^{j\omega})|^2 = R(e^{j\omega})$ to rewrite (10) as two semi-infinite linear constraints on $\tilde{\mathbf{r}}_h$, namely,

$$L(e^{j\omega})^2 \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h \leq U(e^{j\omega})^2 \quad \text{for all } \omega \in [0, \pi]. \quad (11)$$

The spectral mask in (9) is absolute, and hence it imposes implicit scaling constraints on the filter. In some applications, different scalings, such as energy normalization, are required. In those cases, it is often more appropriate to impose a relative mask of the form $L(e^{j\omega}) = \beta \bar{L}(e^{j\omega})$ and $U(e^{j\omega}) = \beta \bar{U}(e^{j\omega})$ for

given bounds $\bar{L}(e^{j\omega})$ and $\bar{U}(e^{j\omega})$, where $\beta > 0$ is a design variable that will be incorporated into \mathbf{x} along with one of \mathbf{h} , $\tilde{\mathbf{h}}$, or $\tilde{\mathbf{r}}_h$. These relative masks can be handled in an analogous way, and they facilitate some interesting variants of some conventional design problems; e.g., [34, Ex. 1].

SPECTRAL FLATNESS

The optimization of filters subject to spectral mask constraints often leads to “ripples” in the response. Less oscillatory spectra can be obtained by constraining the curvature of the response with respect to frequency.

For linear-phase filters, we can consider the amplitude response $\tilde{H}(e^{j\omega})$. Since $d^p \tilde{H}(e^{j\omega})/d\omega^p = \tilde{\mathbf{v}}^{(p)}(\omega)^T \tilde{\mathbf{h}}$, where $\tilde{\mathbf{v}}^{(p)}(\omega) = d^p \tilde{\mathbf{v}}(\omega)/d\omega^p$, bounds on the derivatives of $\tilde{H}(e^{j\omega})$ result in semi-infinite linear constraints on $\tilde{\mathbf{h}}$. For example, the constraint that $\tilde{H}(e^{j\omega})$ is concave in frequency over a pass-band \mathcal{P} corresponds to $\tilde{\mathbf{v}}^{(2)}(\omega)^T \tilde{\mathbf{h}} \leq 0$ for all $\omega \in \mathcal{P}$, [10]. In autocorrelation-based designs one can impose curvature constraints on the power spectrum $|H(e^{j\omega})|^2 = \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h$ in an analogous way. Furthermore, in certain cases, bounds on the derivatives of $|H(e^{j\omega})|$ can be handled directly [25].

The behavior of a filter’s response can also be modified by forcing the derivatives of the response to be zero at certain frequencies. These “flatness” constraints correspond to linear equality constraints on $\tilde{\mathbf{h}}$ or $\tilde{\mathbf{r}}_h$; e.g., [4], [20], [30], and [31].

FINITE REPRESENTATIONS OF SPECTRAL MASK CONSTRAINTS

By their very nature, spectral masks, such as that in (9), yield semi-infinite constraints, in the sense that each value of ω generates different constraints. To incorporate these constraints into a conventional optimization framework, they must be represented in a finite manner. Although some rather sophisticated discretization techniques available [35], a popular approach is to approximate each semi-infinite constraint by sampling it uniformly in frequency. The sampled constraints can also be “tightened” in a number of ways to account for the behavior of the response in between the sample points. In the case of sampling (11) at $N + 1$ uniformly spaced values of ω and uniform tightening of the mask in an additive sense, this approximation results in the $2N + 2$ linear constraints

$$L(e^{j\omega_i})^2 + \epsilon \leq \tilde{\mathbf{v}}(\omega_i)^T \tilde{\mathbf{r}}_h \leq U(e^{j\omega_i})^2 - \epsilon \quad (12)$$

for $\omega_i = i\pi/N$, $i = 0, 1, 2, \dots, N$, where $\epsilon \geq 0$ tightens the mask. Although there are cases in which the pair (N, ϵ) can be chosen so that satisfaction of (12) guarantees satisfaction of (11) (e.g., [25]), a popular rule of thumb is to choose $N \approx 15L$ and to allow ϵ to be quite small. If, as is often the case, there are discontinuities in $L(e^{j\omega})$ and $U(e^{j\omega})$ that do not lie on the uniform sampling grid and then imposing additional constraints corresponding to those points of discontinuity can significantly improve the quality of the sampled approximation.

The uniform sampling approach in (12) is convenient because it yields constraints that can be handled directly by general-purpose convex optimization solvers. A popular alternative is to employ a constraint-exchange approach; e.g., [1], [8], [11], and

[12]. In each iteration of such algorithms, a set of active constraints is postulated based on information regarding the current iterate, and the next iterate is generated by solving the simpler optimization problem obtained by considering only those active constraints. Although the selection of a subset of the constraints usually means that each iteration of an exchange algorithm incurs a relatively low computational cost, considerable effort may be required to guarantee convergence.

The discussion here raises a more general question of whether it is possible to precisely represent the mask constraints for linear-phase or autocorrelation-based designs in a finite convex manner. Dual parameterization methods [36] yield precise finite representations, but they do not necessarily yield convex problems. In contrast, by developing generalizations of the Positive-Real Lemma and Bounded-Real Lemma of system theory, for a large class of functions $L(e^{j\omega})$ and $U(e^{j\omega})$ the two semi-infinite constraints in (11) (or the corresponding constraints for the linear-phase case) can be precisely represented in a finite and convex way using linear matrix inequalities; e.g., [34]–[39]. Perhaps the simplest example of this approach concerns the nonnegativity constraint in (8). That semi-infinite linear constraint is equivalent to the existence of an $L \times L$ positive semidefinite symmetric matrix \mathbf{X} for which $\text{trace}(\mathbf{X}) = [\tilde{\mathbf{r}}_h]_0$ and the sum of the elements of the k th off-diagonal of \mathbf{X} equals $[\tilde{\mathbf{r}}_h]_{kr}$, $k = 1, 2, \dots, L - 1$; e.g., [26], [29], and [40]. Since the set of positive semidefinite matrices is convex [13] and the equality constraints are linear, we have a finite convex representation.

**SOME FILTER DESIGN PROBLEMS
REQUIRE THE USE OF A
NORMALIZATION CRITERION,
TYPICALLY WITH RESPECT TO THE
ENERGY OF THE IMPULSE RESPONSE
OR THE DC GAIN.**

The above constraints allow us to control the slope of the response on a linear-linear plot. However, we may also wish to constrain it on a log-linear (decibel) plot or on a log-log plot (i.e., the Bode diagram). For linear phase filters and for autocorrelation-based designs, such constraints also yield semi-infinite linear constraints on $\tilde{\mathbf{h}}$ and $\tilde{\mathbf{r}}_h$, respectively, [25].

BOUNDS ON THE PHASE RESPONSE

In the direct design of a general FIR filter, one may wish to impose bounds on the phase response of $H(e^{j\omega})$ that are analogous to those for the magnitude response; i.e., $F(e^{j\omega}) \leq \angle H(e^{j\omega}) \leq G(e^{j\omega})$. Bounds of this form can be written as semi-infinite linear constraints on \mathbf{h} [41], [42].

WEIGHTED INTEGRAL SQUARED APPROXIMATION ERROR

One of the classical approaches to FIR filter design is to pose the design problem as the approximation of a desired frequency response $D(e^{j\omega})$, and to minimize a measure of the approximation error; see “Approximation Errors and Induced Norms.” In this section, we will consider the “least-squares” approximation error, and in the following section we will consider the peak error.

For a given real-valued weighing function $W(\omega) \geq 0$, the constraint that the weighted integral of the squared approximation error be less than ξ is

$$\int_{-\pi}^{\pi} W(\omega) |D(e^{j\omega}) - H(e^{j\omega})|^2 d\omega \leq \xi. \quad (13)$$

This constraint can be written as

$$\mathbf{h}^T \mathbf{Q} \mathbf{h} - 2\mathbf{b}^T \mathbf{h} + c \leq \xi, \quad (14)$$

where $\mathbf{Q} = \int_{-\pi}^{\pi} W(\omega) \mathbf{v}(\omega) \mathbf{v}(\omega)^H d\omega$, $\mathbf{b} = \int_{-\pi}^{\pi} W(\omega) \text{Re}(D(e^{j\omega})^* \mathbf{v}(\omega)^*) d\omega$, $c = \int_{-\pi}^{\pi} W(\omega) |D(e^{j\omega})|^2 d\omega$, and $\text{Re}(\cdot)$ denotes the real part. (In some cases, these integrals can be found in closed form.) The matrix \mathbf{Q} is positive semidefinite and hence the constraint in (14) is convex.

For linear phase filters, we can consider the error between the amplitude response $\tilde{H}(e^{j\omega})$ and a real-valued desired response, say $\tilde{D}(e^{j\omega})$. The weighted integral squared error between these responses can be written as a convex quadratic function of $\tilde{\mathbf{h}}$ that is analogous to that in (14).

In (13), the approximation error is measured in the “complex domain” [1], and hence the magnitude error and the phase error are intertwined. However, in some cases it might be appropriate to impose a corresponding constraint on the error between the power spectra of $D(e^{j\omega})$ and $H(e^{j\omega})$; i.e.,

$$\int_{-\pi}^{\pi} W(\omega) ||D(e^{j\omega})|^2 - |H(e^{j\omega})|^2|^2 d\omega \leq \xi. \quad (15)$$

The expression on the left-hand side is a quartic function of \mathbf{h} and hence is not necessarily convex. However, by recalling that $|H(e^{j\omega})|^2 = R_h(e^{j\omega})$, we can reparameterize this constraint as a convex quadratic constraint on $\tilde{\mathbf{r}}_h$ analogous to (14).

WEIGHTED PEAK APPROXIMATION ERROR

A constraint on the weighted peak error takes the form

$$\max_{\omega} W(\omega) |D(e^{j\omega}) - H(e^{j\omega})| \leq \xi. \quad (16)$$

This “Chebyshev” constraint is equivalent to the semi-infinite convex quadratic constraint [6], [18]–[20], [46]

$$W(\omega) (\mathbf{h}^T \mathbf{v}(\omega) \mathbf{v}(\omega)^H \mathbf{h} - 2\text{Re}(D(e^{j\omega})^* \mathbf{v}(\omega)^H) \mathbf{h} + |D(e^{j\omega})|^2) \leq \xi \quad \text{for all } \omega \in [0, \pi]. \quad (17)$$

(An alternative formulation appears in [47].) In the linear-phase case, a constraint on the weighted peak error between $\tilde{H}(e^{j\omega})$ and a real-valued desired response $\tilde{D}(e^{j\omega})$ (cf. [8]) can be written as the following pair of semi-infinite linear constraints, e.g., [10]

$$W(\omega) \tilde{D}(e^{j\omega}) - \xi \leq W(\omega) \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq W(\omega) \tilde{D}(e^{j\omega}) + \xi. \quad (18)$$

A constraint on the weighted peak error of the power spectra,

$$\max_{\omega} W(\omega) ||D(e^{j\omega})|^2 - |H(e^{j\omega})|^2| \leq \xi, \quad (19)$$

yields nonconvex quadratic constraints on \mathbf{h} but semi-infinite linear constraints on $\tilde{\mathbf{r}}_h$, analogous to those in (18).

As an alternative to (19) one can consider the weighted peak error in decibels. Given a set of frequencies \mathcal{W} for which $W(\omega) > 0$ and $D(e^{j\omega}) \neq 0$, that error can be written as

$$\max_{\omega \in \mathcal{W}} W(\omega) |20 \log_{10}(|D(e^{j\omega})|) - 20 \log_{10}(|H(e^{j\omega})|)|. \quad (20)$$

In the case of autocorrelation-based designs, the constraint that this error is bounded by ξ is equivalent to the following pair of semi-infinite linear constraints on $\tilde{\mathbf{r}}_h$, [25],

$$|D(e^{j\omega})|^2 / \eta(\omega) \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_g \leq \eta(\omega) |D(e^{j\omega})|^2 \quad (21)$$

for all $\omega \in \mathcal{W}$, where $\eta(\omega) = 10^{\xi/(10W(\omega))}$.

WEIGHTED L₁ APPROXIMATION ERROR

In addition to the weighted L_2 and L_{∞} norms of the error $E(e^{j\omega})$ that were considered above, other weighted L_p errors may be of interest [48]. In particular, we may be interested in a bound on the weighted L_1 norm of the error

$$\int_{-\pi}^{\pi} W(\omega) |D(e^{j\omega}) - H(e^{j\omega})| d\omega \leq \xi.$$

In the linear-phase case, a convex formulation of the version of this constraint for the amplitude response can be obtained (e.g., [18] and [49]), and a specialized algorithm was developed in [50].

TIME-DOMAIN CRITERIA

Now we provide some examples of time-domain design criteria that yield convex constraints on h or \tilde{r}_h . For linear phase filters, h is a (simple) linear function of \tilde{h} , and hence the discussion below regarding h also applies to \tilde{h} .

SIMPLE NORMALIZATION CRITERIA

Some filter design problems require the use of a normalization criterion, typically with respect to the energy of the impulse response or the DC gain. Setting the DC gain to K corresponds to the linear equality constraint $\sum_{n=0}^{L-1} h[n] = K$, but specifying the filter energy results in a quadratic equality constraint on h , which is not convex. In contrast, for autocorrelation-based designs, setting the filter energy to \mathcal{E} corresponds to $r_h[0] = \mathcal{E}$, which is linear, and hence convex. Setting the DC gain to K in such designs corresponds to $r_h[0] + 2 \sum_{k=1}^{L-1} r_h[k] = K^2$, which is also linear.

ENVELOPE CONSTRAINTS

In some applications, the filter design problem may include a constraint that the time-domain response to a particular input lies within a prescribed envelope [51]. For a “sampled-data” implementation using an FIR filter, this envelope constraint leads to two semi-infinite linear constraints on h ; one pair of constraints for each instant of time [51].

INTERPOLATION

In the conversion of a signal of a given sampling rate to a corresponding signal at a higher sampling rate, it may be desirable to impose an interpolation constraint that ensures that the samples from the original signal appear unchanged in the “up-converted” signal. For sampling rate changes by an integer factor M , this

constraint corresponds to the linear equality constraints $h[0] = 1$ and $h[iM] = 0, i = 1, 2, \dots, \lfloor (L-1)/M \rfloor$.

ℓ_1 APPROXIMATION ERROR

As discussed in “Approximation Errors and Induced Norms,” the ℓ_1 norm of $d[n] - h[n]$, where $d[n]$ is the desired impulse response, is an induced norm of the error system. A bound on this approximation error, $\sum_{n=0}^{L-1} |d[n] - h[n]| \leq \xi$, can be written as the following $2L + 1$ linear constraints on h and the additional variable $t \in \mathbb{R}^L$ [13], $d[n] - t[n] \leq h[n] \leq d[n] + t[n]$, and $\sum_{n=0}^{L-1} t[n] \leq \xi$. When $d[n] = 0$, this corresponds to the constraint $\|h\|_1 \leq \xi$, which is a popular proxy for a sparsity constraint on h ; e.g., [52].

SELF-ORTHOGONALITY

In a number of applications, including the design of pulse shapes for communications and the design of wavelets and multirate filter banks, it is desirable that the impulse response be orthogonal, or almost orthogonal, to translated versions of itself at integer multiples of, say, M . That is, we would like to have $\sum_n h[n] h[n + iM]$ to be zero, or small, for nonzero integers i . Unfortunately, this leads to nonconvex quadratic constraints on h ; e.g., [53]. However, the orthogonality constraint yields linear equality constraints on \tilde{r}_h , namely $r_h[iM] = 0$. If small deviations from self-orthogonality are permissible, we can collect the terms $r_h[iM]$, $i = 1, 2, \dots, \lfloor (L-1)/M \rfloor$, in a vector \check{r}_h and impose $\|\check{r}_h\|_1 \leq \xi$, $\|\check{r}_h\|_2 \leq \xi$, or $\|\check{r}_h\|_\infty \leq \xi$, all of which can be written as convex constraints on \tilde{r}_h and perhaps some additional variables; e.g., [27] and [54].

VARIATIONS ON THE THEME OF (7)

As mentioned in the section “Generic Formulation of Filter Design Problems,” the problem in (7) can be interpreted as finding the tightest version of the m_0 th constraint such that there is a filter that satisfies all the constraints. As we now illustrate, variations on this interpretation yield a number of interesting efficiently solvable design problems that do not quite fit into the framework of (7) but are very closely related.

APPROXIMATION ERRORS AND INDUCED NORMS

A useful interpretation of some of the common measures used in approximation-based filter design can be obtained by asking how different the outputs of the desired and designed systems can be for a given class of bounded inputs; e.g., [43]. This corresponds to considering the approximation error $E(e^{j\omega}) = D(e^{j\omega}) - H(e^{j\omega})$ as the frequency response of a stable system and asking how large the output of that system can be for the given class of inputs. Different choices for the notions of size of the inputs and outputs yield different “induced norms” (e.g., [44, Ch. 5] and [45, Ch. 2]) on the error system and hence different measures of the approximation error.

When the input class is the set of bounded energy signals and the output is measured in terms of its energy, the induced norm is the peak value of $|E(e^{j\omega})|$; i.e., the L_∞ norm of $E(e^{j\omega})$. For cases

in which $D(e^{j\omega})$ has an inverse discrete-time Fourier transform (of possibly infinite length), denoted by $d[n]$, the induced norm from the energy of the input to the peak absolute value of the time-domain output is the square-root of the unweighted integral squared error; i.e., the L_2 norm of $E(e^{j\omega})$ on $[-\pi, \pi]$. Similarly, the induced norm from the peak absolute value of the input to the peak absolute value of the output is the ℓ_1 norm of $e[n] = d[n] - h[n]$, namely $\sum_n |d[n] - e[n]|$; e.g., [45, Ch. 2].

In the main text we consider weighted approximation errors, which correspond to the induced norms of appropriately pre-filtered error systems. The weighted integral squared error is also 2π times the mean square value of the output of the error system when the input is a stationary random process with power spectral density $W(\omega)$; e.g., [44, Ch. 5].

MINIMUM LENGTH DESIGN

Consider the problem of finding the shortest filter that satisfies the constraints; i.e., $\min_{\mathbf{x} \in \mathbb{R}^L, L \in \mathbb{Z}_+} L$, subject to $f_m(\mathbf{x}) \leq \xi_m$ and $g_p(\mathbf{x}) = \zeta_p$. Since L is an integer, the objective is not convex. However, when the feasible sets of all the constraints are convex, this problem is quasi convex (cf. [13]) and can be efficiently solved. In particular, given lengths L_p and L_s for which the problem is known to be feasible and infeasible, respectively, we can perform a bisection-based search on $[L_s, L_p]$ for the minimum feasible length. At each step, the question of whether there exists a filter of the given length that satisfies all the constraints is resolved. When the feasible set is convex, infeasibility can be reliably detected and that feasibility problem can be efficiently solved. Since the bisection search algorithm is also efficient, a filter of minimum length can be efficiently found. In the absence of other information, the initial value of L_s can be set to zero, and one can obtain an initial value for L_p by iteratively doubling a postulated value until a feasible problem is found; see, e.g., [26] and [54].

TIGHTEST SPECTRAL MASK DESIGN

Another variation is to find the tightest spectral mask such that a feasible filter of a given length exists, subject to other convex constraints; e.g., [25] and [27]. For simplicity, let us consider linear phase filters and the low-pass spectral mask in Figure 1. The problem of minimizing the stop-band level of this mask can be written as: $\min_{\tilde{\mathbf{h}}, \gamma} \gamma$ subject to

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq U_p \quad \text{for all } \omega \in [0, \omega_s], \quad (22a)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \geq L_p \quad \text{for all } \omega \in [0, \omega_p], \quad (22b)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \geq U_p \quad \text{for all } \omega \in [\omega_p, \omega_s],$$

$$-\gamma \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq \gamma \quad \text{for all } \omega \in [\omega_s, \pi]. \quad (22c)$$

This problem is convex (e.g., [10]), as is the problem of minimizing the pass-band ripple (on a linear scale), which takes a similar form, but with the constraints

$$C_p - \gamma \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq C_p + \gamma \quad \text{for all } \omega \in [0, \omega_p], \quad (23a)$$

$$C_p + \gamma \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq C_p + \gamma \quad \text{for all } \omega \in (\omega_p, \omega_s), \quad (23b)$$

$$-U_s \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq U_s \quad \text{for all } \omega \in [\omega_s, \pi], \quad (23c)$$

where $C_p = (L_p + U_p)/2$. To minimize the ripple on a decibel scale, we define $\check{C}_p = \sqrt{L_p U_p}$ and replace (23a) by

$$\check{C}_p/\gamma \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq \check{C}_p\gamma \quad \text{for all } \omega \in [0, \omega_p], \quad (24)$$

and (23b) by $\check{C}_p/\gamma \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq \check{C}_p\gamma$, for all $\omega \in (\omega_p, \omega_s)$. In (24) the upper bound is linear, but the lower bound is not. However, the lower bound can be rewritten as $\check{C}_p \leq \gamma \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}}$, and, at each frequency, this constraint can be transformed [25] into a (convex) second-order cone constraint; cf. [13, p. 197].

In addition to tightening the stop-band level and pass-band ripple constraints, we may also be interested in increasing their extent; that is, finding the smallest stop-band edge, or the largest pass-band edge such that there is a feasible filter of the given length. For the stop-band edge case, we have $\min_{\tilde{\mathbf{h}}, \gamma} \gamma$ subject to

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq U_p \quad \text{for all } \omega \in [0, \gamma), \quad (25a)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \geq L_p \quad \text{for all } \omega \in [0, \omega_p], \quad (25b)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \geq U_p \quad \text{for all } \omega \in [\omega_p, \omega_s],$$

$$-U_s \leq \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{h}} \leq U_s \quad \text{for all } \omega \in [\gamma, \pi]. \quad (25c)$$

Here, γ enters the constraints in a somewhat different manner from that in (7), and the above tightest mask problems. However, for each value of γ the feasible set is convex, and if $\gamma_1 > \gamma_2$ then the feasible set for $\gamma = \gamma_1$ contains that for $\gamma = \gamma_2$. Therefore, the problem is quasi convex and can be efficiently solved using a bisection-based search analogous to that for the minimum length design; see, e.g., [27].

DESIGN EXAMPLE

We now pursue a simple design example that encapsulates some of the principles outlined above. We consider the design of the prototype filter for an over-sampled near-perfect reconstruction generalized discrete Fourier transform (NPR-GDFT) filter bank (e.g., [54]) with M subchannels and down-sampling factor $K < M$. (This design also has an application in filtered multitone modulation [55].) The prototype filter is a low-pass filter, and its desirable properties include a small energy in the stop band $[\pi/K, \pi]$, a low stop-band level, a bounded response in the pass band $[0, \pi/M]$ and the transition band, and that it comes close to satisfying a self-orthogonality constraint [54]. With the energy normalization $\sum_n |h[n]|^2 = K/M$, these constraints become

$$\mathbf{h}^T \mathbf{Q}_s \mathbf{h} \leq \xi_1, \quad (26a)$$

$$|\mathbf{v}(\omega)^H \mathbf{h}| \leq \xi_2 \quad \text{for all } \omega \in [\pi/K, \pi], \quad (26b)$$

$$|\mathbf{v}(\omega)^H \mathbf{h}| \leq \xi_3 \quad \text{for all } \omega \in [0, \pi/K], \quad (26c)$$

$$\sum_{i \neq 0} \left| \sum_n h[n] h[n + iM] \right|^2 \leq \xi_4^2, \quad (26d)$$

respectively. In (26a), $\mathbf{Q}_s = (1/\pi) \int_{\pi/K}^{\pi} \mathbf{v}(\omega) \mathbf{v}(\omega)^H d\omega$ and can be found analytically. Although (26a) is convex and (26b) and (26c) can be made convex [cf. (10)], (26d) is a quartic constraint on \mathbf{h} and the energy normalization is a quadratic equality constraint, neither of which is convex. Fortunately, all four constraints and the normalization can be written as convex constraints on the autocorrelation $r_h[k]$, and the problem of minimizing the stop-band energy in (26a) subject to the other constraints can be written as

$$\min_{\tilde{\mathbf{h}}} \mathbf{g}^T \tilde{\mathbf{r}}_h \quad (27a)$$

$$\text{subject to } \tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h \leq \xi_2^2 \quad \text{for all } \omega \in [\pi/K, \pi], \quad (27b)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h \leq \xi_3^2 \quad \text{for all } \omega \in [0, \pi/K], \quad (27c)$$

$$\|S\tilde{\mathbf{r}}_h\|_2 \leq \xi_4/\sqrt{2}, \quad (27d)$$

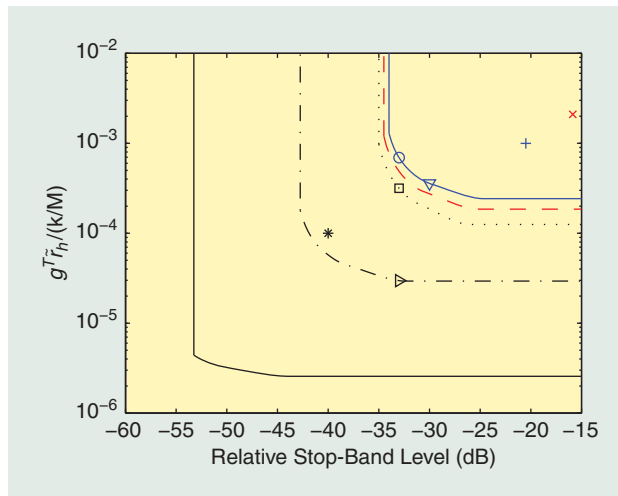
$$[\tilde{\mathbf{r}}_h]_0 = K/M, \quad (27e)$$

$$\tilde{\mathbf{v}}(\omega)^T \tilde{\mathbf{r}}_h \geq 0 \quad \text{for all } \omega \in [0, \pi], \quad (27f)$$

where the last constraint ensures that $\tilde{\mathbf{r}}_h$ represents a valid autocorrelation. The elements of \mathbf{g} are $[\mathbf{g}]_0 = 1 - 1/K$ and $[\mathbf{g}]_k = -2\sin(\pi k/K)/(\pi k)$, $k \geq 1$. The constraint in (27d) represents $\sum_{i>0} r_h[iM]^2 \leq \xi_4^2/2$, and hence \mathbf{S} consists of rows $iM + 1$, $i > 0$, of the $L \times L$ identity matrix. The problem in (27) has a linear objective, three semi-infinite linear constraints, (27b), (27c), and (27f), a convex quadratic constraint (27d), and a linear equality constraint (27e). Hence, it is convex. A simple MATLAB implementation using CVX [17] is provided in “A Simple MATLAB/CVX Implementation of (27).” In that implementation, (27b) and (27c) have been approximated by sampling and uniform additive tightening, and (27f) has been precisely transformed into linear equality constraints on a positive semi-definite symmetric matrix; see “Finite Representations of Spectral Mask Constraints.” Once the optimal autocorrelation has been found, an optimal prototype filter can be obtained by spectral factorization. In this particular application, a factor that is close to having linear phase will often be preferred.

STOP-BAND TRADEOFFS

Now that we have a convex formulation, we can efficiently explore some of the inherent design tradeoffs. To complement the tradeoffs considered in [54], in Figure 3 we provide the inherent tradeoff between the stop-band energy $\mathbf{g}^T \tilde{\mathbf{r}}_h$ and the relative stop-band level ξ_2^2/ξ_3^2 for different values of the self-orthogonality bound $\xi_4 = \alpha K/M$, for an NPR-GDFT filter bank with $M = 8$, $K = 6$, and $L = 48$. As in [54], we set $\xi_3^2 = 10^{0.1}K$. For each curve, all points on or above the curve are achievable with a length 48 filter that satisfies the corresponding self-orthogonality bound, and no point below the curve can be



[FIG3] Tradeoff between the stop-band energy and the stop-band level for length-48 filters with different self-orthogonality bounds $\xi_4 = \alpha K/M$; from outer to inner: no self-orthogonality constraint, $\alpha = 10^{-1}$, 10^{-2} , 5.34×10^{-3} , 10^{-3} .

achieved. The spectra of two filters that achieve the inherent tradeoff for $\alpha = 10^{-3}$ are plotted in Figure 4(a) and (b).

In the absence of a constraint on the self-orthogonality metric, the stop-band tradeoff is a conventional peak-constrained least squares tradeoff [11], but in GDFT filter banks the self-orthogonality constraint captures the important performance criterion of distortion. Figure 3 shows that relatively

A SIMPLE MATLAB/CVX IMPLEMENTATION OF (27)

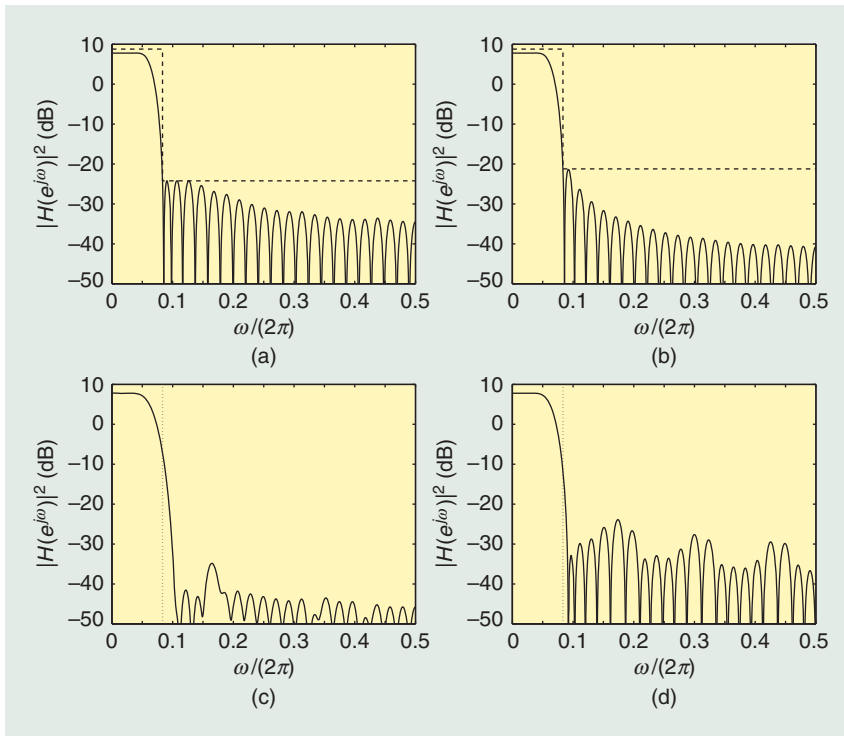
```
L=48;M=8;K=6;rel_sb_level_dB=-33;
alpha=10^(-3);

xi_3_sqr=10^(0.1)*K; xi_4=alpha*K/M;
xi_2_sqr=xi_3_sqr*10^(rel_sb_level_
dB/10);
sf=15; eps_sf=1e-6;

omega_s=pi/K; k_vecT=1:(L-1);
Ns_pt=ceil(L*sf*omega_s/pi);
Ns_s=ceil(L*sf*(1-omega_s/pi));
omega_ptb=linspace(0,omega_s,Ns_pt+1)';
omega_ptb=omega_ptb(1:end-1);
omega_sb=linspace(omega_s,pi,Ns_s)';
VtT_ptb=[ones(Ns_pt,1),
2*cos(omega_ptb*k_vecT)];
VtT_sb=[ones(Ns_s,1),2*cos(omega_sb*k_
vecT)];

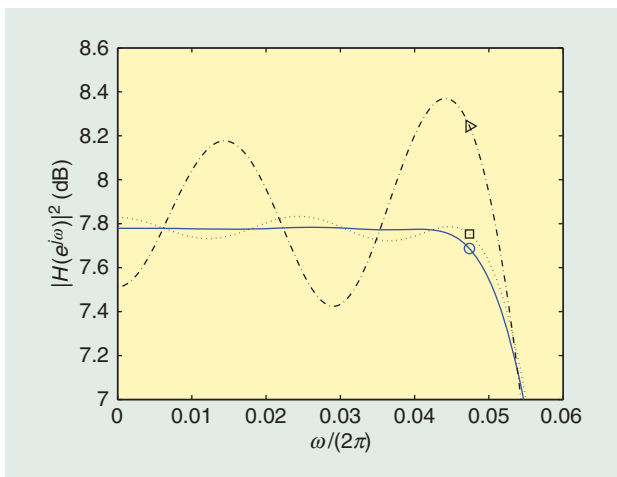
N_isi=floor((L-1)/M); k_isi_vec
=M*(1:N_isi)';
S=eye(L); S=S(k_isi_vec+1,:);

gT=[1-1/K, -2*sin(pi*k_vecT/K)./
(pi*k_vecT)];
cvx_begin
cvx_solver sedumi; cvx_precision
high;
variable rtilde_h(L);
variable X(L,L) symmetric;
minimize (gT*rtilde_h)
subject to
VtT_sb*rtilde_h <= xi_2_
sqr-eps_sf;
VtT_ptb*rtilde_h <= xi_3_sqr_
eps_sf;
norm(S*rtilde_h) <= xi_4/sqrt(2);
[1,zeros(1,L-1)]*rtilde_h == K/M;
X==semidefinite(L);
trace(X)==rtilde_h(1);
for n=1:L-1,
sum(diag(X,n))==rtilde_h(n+1);
end;
cvx_end
```



[FIG4] (a)–(d) Power spectra of filters achieving the points a) \circ : stop-band 33 dB down; b) ∇ : stop-band 30 dB down; c) \times : Interpolated symlet; d) $+$: Harteneck et al. [56] in Figure 3. The stop-band edge is indicated by the mask (dashed) or a dotted line.

mild constraints on the self-orthogonality have a considerable impact on the inherent tradeoff; the achievable region for $\alpha = 10^{-1}$ is significantly smaller than that in the absence of the self-orthogonality constraint. The impact of the self-orthogonality constraint on the pass-band spectrum of the filter is illustrated in Figure 5. As is implicit in Figure 3, the gross spectral features of these filters are similar to those in Figure 4(a), but they satisfy different self-orthogonality bounds. Figure 5 shows that the filters that are closer to being



[FIG5] Pass-band spectra of filters achieving the indicated points in Figure 3.

self-orthogonal have substantially flatter pass-band spectra; see also [27].

To illustrate the role that the tradeoffs in Figure 3 can play in evaluating filters designed using other techniques, consider the filter with the spectrum in Figure 4(c). This filter was obtained by interpolating the length-12 symlet filter to length 48; see [54]. The self-orthogonality metric for this filter is $5.34 \times 10^{-3} K/M$. Even though its side-lobes decay quite rapidly, this filter has a rather large response at the stop-band edge and this is reflected in the position of the stop-band tradeoff achieved by this filter (the \times in Figure 3) with respect to the corresponding inherent tradeoff (the red dashed curve). In particular, Figure 3 shows that for the attained self-orthogonality bound, significantly lower stop-band levels and stop-band energies can be achieved.

Another approach to prototype design for NPR-GDFT filter banks is to consider linear-phase prototypes and to minimize a weighted sum of the stop-band energy and a self-orthogonality metric analogous to

(26d), [56]. That problem is not convex, but good locally optimal filters can often be found using an iterative reweighted least-squares technique; cf. [48]. The stop-band tradeoff achieved by the filter designed using that method that has a self-orthogonality metric of $10^{-3} K/M$ is marked with the $+$ in Figure 3, and its gross power spectrum is provided in Figure 4(d). The gap between the $+$ and the corresponding inherent tradeoff (the blue solid curve in Figure 3) is due to the linear phase constraint and, possibly, to the fact that the algorithm only generates locally optimal solutions. That said, phase linearity is sometimes highly desirable, and linear-phase filters are easier to implement; see also [57].

To close this example, we find the shortest filter that can achieve the point marked with the $*$ in Figure 3, in the absence of the self-orthogonality constraint and for self-orthogonality bounds $\xi_4 = \alpha K/M$, with $\alpha = 10^{-1}$, 10^{-2} , and 10^{-3} . Using the approach in the section “Minimum Length Design,” these lengths were found to be 36, 47, 55, and 59, respectively. As predicted by Figure 3, when $\alpha = 10^{-2}$ and 10^{-3} filters longer than 48 are required.

NONCONVEX DESIGN PROBLEMS

Our focus so far has been on FIR filter design problems that are, or can be transformed into, convex or quasi-convex optimization problems. Of course, there are many interesting design problems that are not convex, including the design of filters with lower bounds on the magnitude response and bounds on the phase response; the design of linear phase filters with self-orthogonality constraints; the design of filters with quantized

coefficients; and a variant of the design example that involves two prototype filters and group delay specifications [58]. There is a plethora of approaches to generating “good,” or even optimal, solutions to such problems, but the effectiveness of a given approach depends quite strongly on the nature of the problem at hand. In lieu of a comprehensive review of such approaches, in this section we will try to place a few of the more prominent general-purpose approaches in the context of our earlier discussion.

For nonconvex design problems with continuously valued coefficients, one reasonable approach is to adopt a sequential approximation algorithm [23]. In each step of such algorithms, we first construct a local approximation of the nonconvex design problem around the current iterate. Often that approximation is chosen so that the resulting local subproblem is convex, and sometimes the subproblem has a closed-form solution. The next iterate is generated by (coarsely) solving the local subproblem problem, and the algorithm continues until appropriate convergence criteria are satisfied. The sequential quadratic programming algorithm [23], which underlies MATLAB’s `fmincon` function, is an example of this approach, and related approaches were applied to some FIR filter design problems in [59] and [60]. (Although it is somewhat different, the iterative reweighted least squares approach [48] also involves sequential approximation.) Sequential algorithms require a starting point, and given the potential for a locally optimal solution to be significantly suboptimal in a global sense, these algorithms are typically run from several starting points. As is implicit in “Using an Inherent Tradeoff in a Nonconvex Design,” there is considerable value in choosing “good” starting points. One approach to that task is to globally approximate the nonconvex problem by a convex one, and to use the solution of that convex problem, or a modified version thereof, as one of the starting points.

Two classes of global approximations that can also assist the designer in the evaluation of locally optimal solutions are relaxation and restriction. Simply put, relaxation involves loosening or removing some of the constraints in the generic design problem in (7), and hence the feasible set expands. When that expanded feasible set is convex, the relaxed problem can be efficiently solved. Although the resulting solution might not satisfy the constraints of the original problem, it does provide a lower bound on the globally optimal value of the original problem. (Relaxation can also be used to generate lower bounds for branch-and-bound approaches [22] to globally solving the original problem.) The lower bounds provided by relaxation can be used to construct an outer bound on the region of tradeoff points that can be achieved; i.e., points below a tradeoff curve for the relaxed problem cannot be achieved. This was implicit in “Using an Inherent Tradeoff in a Nonconvex Design.”

TWO CLASSES OF GLOBAL APPROXIMATIONS THAT CAN ALSO ASSIST THE DESIGNER IN THE EVALUATION OF LOCALLY OPTIMAL SOLUTIONS ARE RELAXATION AND RESTRICTION.

Although points on the tradeoff curve for the relaxed problem might not be achievable while satisfying the nonconvex constraints, if we obtain a solution to the original problem that is “close” to that curve, it is likely to be deemed a “good” solution.

That judgment can be enhanced by using restriction, in which some of the constraints are tightened, or new constraints are added, and the feasible set shrinks. When that shrunk feasible set is convex, the restricted problem can be efficiently solved. The restricted problem provides an upper bound on the globally optimal value of the original problem and an inner bound on the region of achievable tradeoff points; i.e., points on or above the tradeoff curve for the restricted problem are achievable. An example of restriction arises in problems with a lower bound constraint on the magnitude response. This is a nonconvex constraint, but if the constraint that the filter have linear phase is added, the feasible set becomes convex.

The design of filters with quantized coefficients has long been associated with the field of discrete optimization; e.g., [61] and [62]. A conventional approach to finding optimal discrete coefficient filters is to perform a branch-and-bound search [22] in which the lower bounds are generated by solving a relaxed problem in which the free variables no longer required to take on discrete values. If the coefficient quantization constraint is the only nonconvex constraint in the design problem, then the relaxed problem can be efficiently solved and a lower bound efficiently obtained. In that case, it might be tempting to remove the coefficient quantization constraint all together, solve the convex relaxed problem, and then round that solution to the nearest quantized point. Although that approach is known to perform quite poorly in some cases, the quantized solution of the completely relaxed problem is a candidate for the starting point of a local search algorithm for a “good” quantized solution. Furthermore, that point can also be used to specify a randomized rounding procedure that can often be quite effective; e.g., [63].

OTHER FILTER ARCHITECTURES

Although the focus of this article is on the design of FIR filters with real-valued coefficients, many of the principles are quite generic and can be extended to other filter architectures. Perhaps the most straightforward extension is to FIR filters with complex-valued coefficients. In that case, the frequency response and power spectrum remain linear functions of the impulse response and autocorrelation, respectively, and most of the convex design criteria that we have considered remain convex. This extension enables the application of some of our discussion to the design of narrow-band antenna arrays; e.g., [34] and [64]. Another extension is to filters with “FIR-like” architectures, in which each delay element in the FIR filter is

replaced by a transfer function; e.g., [65]. These include discrete-time Laguerre networks (e.g., [66]) and “variable” filters with Farrow structure; e.g., [5, p. 45].

Convex optimization also has an important role to play in the design of infinite impulse response (IIR) filters. For design problems based on a spectral mask, convex design formulations can be obtained [37]. However, design criteria that involve both the magnitude and phase responses, and the requirement that the filter be stable, generally lead to nonconvex constraints; e.g., [67]. When a sequential approximation approach is applied to such problems, there are several techniques for obtaining a convex local approximation (e.g., [4], [19], [67]–[71]), and the stability constraint can be tackled by restricting the denominator coefficients to a set that is guaranteed to yield stable filters and has a convex description [68], [70]. Convex optimization can also play a role in the selection of the starting points. One candidate starting point is an FIR filter of the same order as the numerator, and an IIR starting point can be generated by first designing an FIR filter of higher order than the IIR filter and then employing a model reduction procedure; e.g., [4]. In the multistage approach to IIR filter design [70], different IIR design methods are applied in sequence, with the starting points being the local solutions obtained by the previous method.

The design of multidimensional filters is somewhat more involved than that of one-dimensional filters, but some aspects of our discussion extend to that case, too; e.g., [19].

CONCLUSION

The effective design of FIR filters requires judicious compromises to be made between competing properties of the filter. As we have argued herein, when those properties lie in a rich class of design criteria that are convex in the design variables, the inherent tradeoffs between these properties, and filters that attain these tradeoffs, can be reliably obtained. Furthermore, some recently developed software [16], [17] provides a platform that enables these tradeoffs to be obtained with little programming effort. In our examples we only considered pair-wise tradeoffs, but the extension to tradeoff surfaces in higher dimensions is straightforward, even though they tend to be more difficult to visualize. As we illustrated, many common FIR filter design criteria are convex, but the class of convex design criteria is much richer than these examples indicate; cf. [13] and [17]. Moreover, the general-purpose convex design platform provides the flexibility to combine convex design criteria in a variety of ways. This platform also offers the opportunity to develop customized algorithms that exploit the structure of particular classes of FIR filter design problems.

Making judicious design compromises is substantially more difficult when one or more of the criteria is not convex. However, as we have argued that convex optimization can enhance the design process for such filters in several

ways, including: the provision of inner and outer bounds on the design tradeoffs, which can be used to evaluate locally optimal filters; the development of convex local approximations of the design problem for use in sequential approximation algorithms; the provision of starting points for such algorithms; and the development of bounds for use in branch-and-bound algorithms for globally optimal filters.

Although our focus has been on highlighting ways in which convex optimization enriches the established art of FIR filter design, the reliability and efficiency of convex optimization techniques also gives rise to new opportunities, including the embedding [13] of flexible FIR filter design algorithms into autonomous, or semiautonomous, systems.

ACKNOWLEDGMENTS

The author would like to thank Ramy Gohary, of the Communication Research Centre, Ottawa, Ontario, Canada, for his insightful comments on an early draft. This work was supported in part by the Canada Research Chairs Program.

AUTHOR

Timothy N. Davidson (davidson@mcmaster.ca) received the B.Eng. (Hons. I) degree from the University of Western Australia in 1991 and the D.Phil. degree from the University of Oxford in 1995. He is a professor in the Department of Electrical and Computer Engineering at McMaster University, Hamilton, Ontario, Canada, where he holds the (Tier II) Canada Research Chair in Communication Systems and is currently associate director of the School of Computational Engineering and Science. He was awarded the 1991 Rhodes Scholarship for Western Australia.

REFERENCES

- [1] T. W. Parks and C. S. Burrus, *Digital Filter Design*. New York: Wiley, 1987.
- [2] A. Antoniou, *Digital Filters. Analysis, Design, and Applications*, 2nd ed. New York: McGraw-Hill, 1993.
- [3] T. Saramäki, “Finite impulse response filter design,” in *Handbook for Digital Signal Processing*, S. K. Mitra and J. F. Kaiser, Eds. New York: Wiley, 1993, ch. 4, pp. 155–277.
- [4] L. J. Karam, J. H. McClellan, I. W. Selesnick, and C. S. Burrus, “Digital filtering,” in *The Digital Signal Processing Handbook*, V. K. Madisetti and D. B. Williams, Eds. Boca Raton, FL: CRC, 1998, ch. 11, pp. 1–86.
- [5] T. I. Laakso, V. Välimäki, M. Karjalainen, and U. K. Laine, “Splitting the unit delay,” *IEEE Signal Processing Mag.*, vol. 13, pp. 30–60, Jan. 1996.
- [6] W.-S. Lu, “Design of digital filters and filter banks by optimization: A state of the art review,” in *Proc. European Signal Processing Conf.*, Tampere, Finland, Sept. 2000, pp. 351–354.
- [7] A. Tkachenko, P. P. Vaidyanathan, and T. Q. Nguyen, “On the eigenfilter design method and its applications: A tutorial,” *IEEE Trans. Circuits Syst. II*, vol. 50, pp. 497–517, Sept. 2003.
- [8] T. W. Parks and J. H. McClellan, “Chebyshev approximation for nonrecursive digital filters with linear phase,” *IEEE Trans. Circuit Theory*, vol. CT-19, no. 2, pp. 189–194, 1972.
- [9] J. H. McClellan and T. W. Parks, “A personal history of the Parks–McClellan algorithm,” *IEEE Signal Processing Mag.*, vol. 22, pp. 82–86, Mar. 2005.
- [10] K. Steiglitz, T. W. Parks, and J. F. Kaiser, “METEOR: A constraint-based FIR filter design program,” *IEEE Trans. Signal Processing*, vol. 40, pp. 1901–1909, Aug. 1992.
- [11] J. W. Adams and J. L. Sullivan, “Peak-constrained least-squares optimization,” *IEEE Trans. Signal Processing*, vol. 46, pp. 306–321, Feb. 1998.
- [12] I. W. Selesnick, M. Lang, and C. S. Burrus, “Constrained least square design for FIR filters without specified transition bands,” *IEEE Trans. Signal Processing*, vol. 44, pp. 1879–1892, Aug. 1996.
- [13] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

- [14] J. F. Sturm. (1999). Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optim. Methods and Software* vol. 11–12, pp. 625–653 [Online]. Available: <http://sedumi.ie.lehigh.edu>
- [15] K. C. Toh, M. J. Todd, and R. H. Tutuncu. (1999). SDPT3—A Matlab software package for semidefinite programming, version 2.1. *Optim. Methods and Software*, vol. 11–12, pp. 545–581 [Online]. Available: <http://www.math.nus.edu.sg/~matttohkc/sdpt3.html>
- [16] J. Löfberg. (2004, Sept.). YALMIP: A toolbox for modeling and optimization in MATLAB. *Proc. IEEE Int. Symp. Computer-Aided Control System Design*, Taipei [Online]. Available: <http://control.ee.ethz.ch/~joloef/yalmip.php>
- [17] M. Grant and S. Boyd. (2008). *Graph implementations for nonsmooth convex programs. Recent Advances in Learning and Control*, V. Blondel, S. Boyd, and H. Kimura, Eds. New York: Springer, pp. 95–110 [Online]. Available: <http://stanford.edu/~boyd/cvx>
- [18] J. O. Coleman and D. P. Scholnik. “Design of nonlinear-phase FIR filters with second-order cone programming,” in *Proc. Midwest Symp. Circuits Systems*, Las Cruces, NM, Aug. 1999, pp. 409–412.
- [19] W.-S. Lu. “A unified approach for the design of 2-D digital filters via semidefinite programming,” *IEEE Trans. Circuits Syst. I*, vol. 49, pp. 814–826, June 2002.
- [20] K. M. Tsui, S. C. Chan, and K. S. Yeung. “Design of FIR digital filters with prescribed flatness and peak error constraints using second-order cone programming,” *IEEE Trans. Circuits Syst. II*, vol. 52, pp. 601–605, Sept. 2005.
- [21] K. H. Afkhamie, H. Latchman, L. Yonge, T. N. Davidson, and R. Newman. “Joint optimization of transmit pulse shaping, guard interval length, and receiver side narrow-band interference mitigation in the HomePlugAV OFDM system,” in *Proc. IEEE Workshop Signal Processing Advances Wireless Communications*, New York, June 2005, pp. 996–1000.
- [22] E. L. Lawler and D. E. Wood. “Branch-and-bound methods: A survey,” *Oper. Res.*, vol. 14, no. 4, pp. 699–719, 1966.
- [23] J. Nocedal and S. J. Wright, *Numerical Optimization*. New York: Springer, 1999.
- [24] O. Herrmann and H. W. Schüssler. “Design of non-recursive digital filters with minimum phase,” *Electron. Lett.*, vol. 6, pp. 329–330, May 1970.
- [25] S.-P. Wu, S. Boyd, and L. Vandenberghe. “FIR filter design via spectral factorization and convex optimization,” in *Applied and Computational Control, Signals, and Circuits*, vol. 1, B. Datta, Ed. Cambridge, MA: Birkhauser, May 1999, ch. 5, pp. 215–245.
- [26] T. N. Davidson, Z.-Q. Luo, and K. M. Wong. “Design of orthogonal pulse shapes for communications via semidefinite programming,” *IEEE Trans. Signal Processing*, vol. 48, pp. 1433–1445, May 2000.
- [27] T. N. Davidson. “Efficient design of waveforms for robust pulse amplitude modulation,” *IEEE Trans. Signal Processing*, vol. 49, pp. 3098–3111, Dec. 2001.
- [28] P. Moulin, M. Anitescu, K. O. Kortanek, and F. A. Potra. “The role of linear semi-infinite programming in signal adapted QMF bank design,” *IEEE Trans. Signal Processing*, vol. 45, pp. 2160–2174, Sept. 1997.
- [29] J. Tuqan and P. P. Vaidyanathan. “A state space approach to the design of globally optimal FIR energy compaction filters,” *IEEE Trans. Signal Processing*, vol. 48, pp. 2822–2838, Oct. 2000.
- [30] B. Dumitrescu and C. Popeea. “Accurate computation of compaction filters with high regularity,” *IEEE Signal Processing Lett.*, vol. 9, pp. 278–281, Sept. 2002.
- [31] J.-K. Zhang, T. N. Davidson, and K. M. Wong. “Efficient design of orthonormal wavelet bases for signal representation,” *IEEE Trans. Signal Processing*, vol. 52, pp. 1983–1996, July 2004.
- [32] T. N. T. Goodman, C. A. Micchelli, G. Rodriguez, and S. Seatzu. “Spectral factorization of Laurent polynomials,” *Advances Comput. Math.*, vol. 7, no. 4, pp. 429–454, 1997.
- [33] C. Taswell. “Constraint-selected and search-optimized families of Daubechies wavelet filters computable by spectral factorization,” *J. Comput. Appl. Math.*, vol. 121, pp. 179–195, Sept. 2000.
- [34] T. N. Davidson, Z.-Q. Luo, and J. F. Sturm. “Linear matrix inequality formulation of spectral mask constraints with applications to FIR filter design,” *IEEE Trans. Signal Processing*, vol. 50, pp. 2702–2715, Nov. 2002.
- [35] R. Hettich and K. O. Kortanek. “Semi-infinite programming: Theory, methods and applications,” *SIAM Rev.*, vol. 35, pp. 380–429, Sept. 1993.
- [36] H. H. Dam, K. L. Teo, S. Nordebo, and A. Cantoni. “The dual parameterization approach to optimal least square FIR filter design subject to maximum error constraints,” *IEEE Trans. Signal Processing*, vol. 48, pp. 2314–2320, Aug. 2000.
- [37] B. Alkire and L. Vandenberghe. “Convex optimization problems involving finite autocorrelation sequences,” *Math. Program.*, vol. 93, pp. 331–359, Dec. 2002.
- [38] B. Dumitrescu, *Positive Trigonometric Polynomials and Signal Processing Applications*. New York: Springer, 2007.
- [39] H. D. Tuan, T. T. Son, B.-N. Vo, and T. Q. Nguyen. “Efficient large-scale filter/bank design via LMI characterization of trigonometric curves,” *IEEE Trans. Signal Processing*, vol. 55, pp. 4394–4404, Sept. 2007.
- [40] B. Dumitrescu, I. Tabus, and P. Stoica. “On the parameterization of positive real sequences and MA parameter estimation,” *IEEE Trans. Signal Processing*, vol. 49, pp. 2630–2639, Nov. 2001.
- [41] M. C. Lang. “Design of nonlinear phase FIR digital filters using quadratic programming,” in *Proc. Int. Conf. Acoustics, Speech, Signal Processing*, Munich, Apr. 1997, pp. 2169–2172.
- [42] X. Lai. “Optimal design of non-linear-phase FIR filters with prescribed phase error,” *IEEE Trans. Signal Processing*, vol. 57, pp. 3399–3410, Sept. 2009.
- [43] B. A. Weisburn, T. W. Parks, and R. G. Shenoy. “Error criteria for filter design,” in *Proc. Int. Conf. Acoustics, Speech, Signal Processing*, Adelaide, Apr. 1994, vol. 3, pp. 565–568.
- [44] S. Boyd and C. Barratt, *Linear Controller Design: Limits of Performance*. Englewood Cliffs, NJ: Prentice-Hall, 1991.
- [45] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum, *Feedback Control Theory*. New York: MacMillan, 1992.
- [46] A. Potchinkov and R. Reemsten. “The design of FIR filters in the complex plane by convex optimization,” *Signal Process.*, vol. 46, pp. 127–146, Oct. 1995.
- [47] D. Burnside and T. W. Parks. “Optimal design of FIR filters with the complex Chebyshev error criteria,” *IEEE Trans. Signal Processing*, vol. 43, pp. 605–616, Mar. 1995.
- [48] C. S. Burrus, J. A. Barreto, and I. W. Selesnick. “Iterative reweighted least-squares design of FIR filters,” *IEEE Trans. Signal Processing*, vol. 42, pp. 2926–2936, Nov. 1994.
- [49] W. S. Yu, L. K. Fong, and K. C. Chang. “An ℓ_1 -approximation based method for synthesizing FIR filters,” *IEEE Trans. Circuits Syst. II*, vol. 39, pp. 578–581, Aug. 1992.
- [50] L. D. Grossmann and Y. C. Eldar. “An ℓ_1 -method for the design of linear-phase FIR digital filters,” *IEEE Trans. Signal Processing*, vol. 55, pp. 5253–5266, Nov. 2007.
- [51] B.-N. Vo, A. Cantoni, and K. L. Teo, *Filter Design with Time Domain Mask Constraints: Theory and Applications*. New York: Springer, 2001.
- [52] T. Baran, D. Wei, and A. V. Oppenheim. “Linear programming algorithms for sparse filter design,” *IEEE Trans. Signal Processing*, vol. 58, no. 3, pp. 1605–1617, Mar. 2010.
- [53] T. Q. Nguyen. “Digital filter bank design quadratic-constrained formulation,” *IEEE Trans. Signal Processing*, vol. 43, pp. 2103–2108, Sept. 1995.
- [54] M. R. Wilbur, T. N. Davidson, and J. P. Reilly. “Efficient design of over-sampled NPR GDFT filter banks,” *IEEE Trans. Signal Processing*, vol. 52, pp. 1947–1963, July 2004.
- [55] B. Borna and T. N. Davidson. “Efficient design of FMT systems,” *IEEE Trans. Commun.*, vol. 54, pp. 794–797, May 2006.
- [56] M. Harteneck, S. Weiss, and R. W. Stewart. “Design of near perfect reconstruction oversampled filter banks for subband adaptive filters,” *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 1081–1085, Aug. 1999.
- [57] B. Farhang-Boroujeny. “A square-root Nyquist (M) filter design for digital communication systems,” *IEEE Trans. Signal Processing*, vol. 56, pp. 2127–2132, May 2008.
- [58] H. H. Dam, S. Nordholm, and A. Cantoni. “Uniform FIR filterbank optimization with group delay specifications,” *IEEE Trans. Signal Processing*, vol. 53, pp. 4249–4260, Nov. 2005.
- [59] W.-S. Lu, T. Saramäki, and R. Bregović. “Design of practically perfect-reconstruction cosine-modulated filter banks: A second-order cone programming approach,” *IEEE Trans. Circuits Syst. I*, vol. 51, pp. 552–563, Mar. 2004.
- [60] Z. Lin and Y. Liu. “Design of arbitrary complex coefficient WLS FIR filters with group delay constraints,” *IEEE Trans. Signal Processing*, vol. 57, pp. 3274–3279, Aug. 2009.
- [61] D. Kodek and K. Steiglitz. “Comparison of optimal and local search methods for designing finite wordlength FIR digital filters,” *IEEE Trans. Circuits Syst.*, vol. CAS–28, pp. 28–32, Jan. 1981.
- [62] Y. C. Lim and S. R. Parker. “Discrete coefficient FIR digital filter design based upon an LMS criteria,” *IEEE Trans. Circuits Syst.*, vol. CAS–30, pp. 723–739, Oct. 1983.
- [63] U. Heute, A. Srivastav, V. Sauerland, and J. Klierer. “Fixed-point-coefficient FIR filters and filter banks: Improved design by randomized quantizations,” in *Proc. Int. Symp. Signal Processing and Applications*, Sharjah, Feb. 2007.
- [64] H. Leuret and S. Boyd. “Antenna array pattern synthesis via convex optimization,” *IEEE Trans. Signal Processing*, vol. 45, pp. 526–532, Mar. 1997.
- [65] S. Nordebo and Z. Zang. “Semi-infinite linear programming: A unified approach to digital filter design with time- and frequency-domain specifications,” *IEEE Trans. Circuits Syst. II*, vol. 46, pp. 765–775, June 1999.
- [66] H. H. Dam, A. Cantoni, S. Nordholm, and K. L. Teo. “Digital Laguerre filter design with maximum passband-to-stopband energy ratio subject to peak and group delay constraints,” *IEEE Trans. Circuits Syst. I*, vol. 53, pp. 1108–1118, May 2006.
- [67] A. Tarczyński, G. D. Cain, E. Hermanowicz, and M. Rojewski. “A WISE method for designing IIR filters,” *IEEE Trans. Signal Processing*, vol. 49, pp. 1421–1432, July 2001.
- [68] M. C. Lang. “Least-squares design of IIR filters with prescribed magnitude and phase responses and a pole radius constraint,” *IEEE Trans. Signal Processing*, vol. 48, pp. 3109–3121, Nov. 2000.
- [69] W.-S. Lu and T. Hinamoto. “Optimal design of IIR digital filters with robust stability using conic-quadratic-programming updates,” *IEEE Trans. Signal Processing*, vol. 51, pp. 1581–1592, June 1998.
- [70] B. Dumitrescu and R. Niemistö. “Multistage IIR filter design using convex stability domains defined by positive realness,” *IEEE Trans. Signal Processing*, vol. 52, pp. 962–974, Apr. 2004.
- [71] A. Jiang and H. K. Kwan. “Minimax design of IIR digital filters using iterative SOCP,” *IEEE Trans. Circuits Syst. I*, to be published. SP