Naval Research Laboratory

Washington, DC 20375-5320

NRL/MR/5320--10-9239

Second-Order Cone Formulations of Mixed-Norm Error Constraints for FIR Filter Optimization

DAN P. SCHOLNIK
JEFFREY O. COLEMAN

Advanced Radar Systems Branch
Radar Division

June 25, 2010

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
25-06-2010	Memorandum Report	

4. TITLE AND SUBTITLE

Second-Order Cone Formulations of Mixed-Norm Error Constraints for FIR Filter Optimization

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
61153N

6. AUTHOR(S)

Dan P. Scholnik and Jeffrey O. Coleman

5d. PROJECT NUMBER

5e. TASK NUMBER
EW021-05-43

5f. WORK UNIT NUMBER
8574

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/MR/5320--10-9239

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

10. SPONSOR / MONITOR'S ACRONYM(S)
NRL

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

It has been previously shown that second-order cone programming (SOCP) can be used to efficiently optimize FIR filter frequency responses in the sense of L1, L2, and Linf. Although SOCP cannot be used for the more general case of Lp minimization, two techniques for combining norms provide the flexibility to minimize a continuum of norms from L1 to Linf. The recently introduced vector epsilon-norm is here extended to complex functions on an interval of the real line, and is compared with a simple convex combination of norms. The dual norms of both are found as well. Discretized approximations of these norms can be used for filter design with SOCP, as shown by several examples.

15. SUBJECT TERMS

FIR filter
Second-order cone programming
Lp norms

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dan P. Scholnik
Unclassified	Unclassified	Unclassified	UL	22	19b. TELEPHONE NUMBER *(include area code)* (202) 404-1943

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

(Short titles here differ from longer actual captions.)

1 Introduction

Today, numerical optimization is the approach of choice for all but the simplest digital filter designs. Most software used for optimization-based filter design can be roughly separated into two categories. The first is the general-purpose solver (the `fmincon` and `fminunc` functions of the ubiquitous Matlab Optimization Toolbox are prime examples), which assumes little about the problem structure and permits almost any objective or constraints. This flexibility, however, is offset by a great deal of inefficiency and is unsuitable for large problems. The second group is the custom solvers, written to solve a specific problem or related set of problems. The classic Remez-exchange-based Parks-McClellan [1] algorithm (function `remez` in Matlab) is still the premier representative of this class, used to design optimal L_∞ (minimax) filters. Another popular technique is the method of least-squares, providing filters optimal in the L_2 sense. The efficiency of these methods is tempered by their inflexibility, as they severely restrict the available problem types.

Increasingly, custom methods are being developed that provide more flexibility, offering more than just L_2 and L_∞ designs. In [2], iterative reweighted least squares (IRLS) was used to allow general L_p norm design, with different values of p applied to the filter passband and stopband. Slow convergence in some cases was later improved [3]. Another approach is to combine L_∞ and L_2 optimization methods. Peak Constrained Least Squares (PCLS) [4] is one such approach, where a least-squares design is performed subject to an L_∞ constraint. An equivalent unconstrained method is to minimize a combined norm, formed as the convex combination of the L_∞ and L_2 norms [5].

Straddling the line between general solver and custom algorithm are tools designed to interface to solvers of highly structured canonical problems. Important examples of structured problems include linear programs (LP), second-order cone programs (SOCP) [6], and semidefinite programs (SDP) [7]. LPs, SOCPs, and SDPs are all convex optimization problems that can be solved by a relatively new class of robust and highly efficient interior-point algorithms [8, 9, 10, 11]. Convexity is important, as it guarantees a global solution when a solution exists. Linear programming is limited to an objective and constraints that are all affine in the optimization variables. Even so, a great number of filter design problems can be solved using LP. Early tools designed to automatically generate linear programs from filter specifications include the FF design language [12] and the METEOR program [13]. The qdes program [14] similarly produces quadratic programs (convex quadratic objective, affine constraints) from linear-controller specifications. SOCP augments LP with second-order cone constraints, which include convex quadratic constraints. The authors' own Opt Matlab toolbox [15] provides filter-related data structures and a flexible interface to several second-order cone programming (SOCP) engines. SDP further adds linear matrix inequalities, which restrict variables or their affine combinations to the cone of positive-semidefinite matrices. LMITOOL [16], though not filter-specific, provides a friendly interface to various semidefinite-programming (SDP) engines. The restriction to a suitable one of these problem types leads more efficiently to the solution while preserving a surprising amount of flexibility through the creative combination of multiple constraints. A big advantage of this separation of high-level design and low-level numerical solution is that new constraints and new constraint types can be added to an existing design without the need to extend or retune (or deeply understand) the algorithm.

In this paper we show that SOCP can be used to optimize filters using an arbitrary number of L_∞, L_2, and L_1 constraints (including PCLS), as well as constraints on linear combinations of the L_∞, L_2, and L_1 norms. A alternate method of combining norms through linear combinations of their unit balls is also presented, and it is shown that a vector norm introduced by Burdakov and Merkulov [17], when extended to complex functions, is a special case that combines L_∞ and

Manuscript approved January 12, 2010.

L_2. This norm is of interest because it has similar characteristics to the L_p norm, which cannot be optimized using SOCP unless $p \in \{1, 2, \infty\}$. It is shown that there is a convenient duality relationship between combining norms directly and indirectly through their unit balls, with the result that two complete families of norms spanning from L_1 to L_2 to L_∞ can be optimized using SOCP. Examples are provided to demonstrate the different properties of these combined norms. A condensed version of this work was presented in [18].

2 Problem Specification

The classic approach to filter design in the frequency domain is to minimize some norm $\|E\|$ of a weighted frequency response error $E(f) = W(f)[G(f)H(f) - D(f)]$, where W is a real, nonnegative *weighting function*, G is the equivalent frequency response of other (fixed) filters in cascade, H is the frequency response of the filter being designed, and D is the desired frequency response. Both E and G are in general complex-valued periodic functions of real optimization variables $\mathbf{x}$. For flexibility here a more general filter-design formulation is chosen:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{f}^T \mathbf{x} \\ \text{s.t.} \quad & \|E_k\| \leq \delta_k, \qquad k = 1, \ldots, K. \end{aligned}$$

Both the error functions and the norms can be different for each constraint. The bounds δ_k can be constants or affine functions of the optimization variables. Other SOCP constraints can be added as needed, but will not be considered here.

This paper will restrict its attention to FIR filters, so that the optimization variables $\mathbf{x}$ represent the real and imaginary parts of the filter coefficients and any required auxiliary variables:

$$G(f) = \sum_n (x_{2n} + j x_{2n+1}) e^{-j2\pi f n}.$$

This can also be written

$$G(f) = \Psi^T(f)\mathbf{x},$$

where $\Psi(f)$ is the vector of exponentials corresponding to frequency f. No assumptions are made regarding coefficient symmetry. The restriction to FIR filters ensures that G and thus E are affine functions of the optimization variables, and that $\|E\|$ is a convex function of the optimization variables. IIR filter problems are nonconvex with multiple local minima and cannot, in general, be directly optimized using SOCP or other convex methods. Iterative methods such as Steiglitz-McBride [19] can be used to apply the methods presented here to IIR filters, as is done in [5] and [20], among others. These methods solve a particular joint design of the FIR numerator and denominator filters at each step. While they can produce reasonable filters, they do not truly optimize the IIR frequency response [21].

Since FIR (and stable IIR) frequency responses have bounded magnitudes and are periodic with period one they are effectively in $L_\infty[0,1]$, the space of bounded complex-valued functions on the interval $[0,1]$. This implies that they also belong to the space $L_p[0,1]$ for any $p \geq 1$, and in particular to the space $L_2[0,1]$ of finite-energy complex-valued functions on $[0,1]$. This largely removes the need to consider pathological functions in the following analysis.

3 Second-Order Cone Programming

A powerful special case of convex optimization is *second-order cone programming*, or SOCP [6]. SOCP allows the minimization of a linear objective, subject to second-order cone (SOC) constraints,

which encompass linear and convex quadratic constraints. One standard formulation of a SOCP is

$$\begin{aligned} \min_{\mathbf{x}} . \quad & \mathbf{f}^T\mathbf{x} \\ \text{s.t.} \quad & \|\mathbf{A}_k^T\mathbf{x}+\mathbf{b}_k\|_2 \le \mathbf{c}_k^T\mathbf{x}+d_k, \qquad k=1,\ldots,K. \end{aligned}$$

The number of linearly independent columns in $\mathbf{A}_k$ is R_k, the *rank* of the second-order cone. An equivalent notation for a SOC constraint is the pair of constraints

$$\begin{aligned} \|\mathbf{A}\mathbf{x}+\mathbf{b}\|_2^2 &\le (\mathbf{c}^T\mathbf{x}+d)^2 \\ 0 &\le \mathbf{c}^T\mathbf{x}+d. \end{aligned}$$

Although this form tends to be the most convenient for filter design due to the explicit inequality constraints for some applications the dual problem [6]

$$\begin{aligned} \max_{\mathbf{y}} . \quad & \textstyle\sum_k(-d_k\ \mathbf{b}_k^T)\mathbf{y}_k \\ \text{s.t.} \quad & (-\mathbf{c}_k,\ \mathbf{A}_k)\mathbf{y}_k = \mathbf{f}_k \\ & \left\|([\mathbf{y}_k]_2,\ldots,[\mathbf{y}_k]_{R_k})^T\right\|_2 \le [\mathbf{y}_k]_1 \qquad k=1,\ldots,K \end{aligned}$$

is preferred. The primal and dual problems are equivalent in that any problem that can be expressed in one form can be converted to the other via auxiliary variables. Most modern solvers iteratively solve both problems simultaneously, so the choice of problem formulation is generally based on convenience.

Although SOCP seems quite limited, with only one basic type of constraint, it can be used to solve a wide range of problems using creative combinations of multiple constraints. A linear constraint results as a special case when $\mathbf{A}_k$ and $\mathbf{b}_k$ are zero. Of specific interest here, SOCP can be used to exactly constrain l_1, l_2, and l_∞ norms of error vectors that are affine in the optimization variables. Using discretization SOCP can be used to approximately constrain L_1, L_2, and L_∞ norms of affine error functions [6, 22]. Exact L_2 formulations also exist [23] that do not require discretization. Any constraint can also be used as the objective via an auxiliary variable:

$$\begin{aligned} \min_{\mathbf{x},\delta} . \quad & \delta \\ \text{s.t.} \quad & \|\mathbf{A}_k^T\mathbf{x}+\mathbf{b}_k\|_2 \le \delta. \end{aligned}$$

The following sections briefly summarize how SOCP can be used to constrain the various norms of an error function $E : [0,1] \times \mathbb{R}^N \to \mathbb{C}$ that is affine in the optimization variables, the dependence on which is hereafter implicit.

3.1 L_∞-Norm Constraints

For a function E the L_∞ or Chebychev norm is defined as

$$\|E\|_\infty \triangleq \operatorname*{ess\,sup}_f |E(f)|.$$

Since the functions of interest (FIR transfer functions) are continuous and periodic, this reduces to

$$\|E\|_\infty \triangleq \max_f |E(f)|.$$

For a function E that is affine in the optimization variables, the L_∞-norm can be (approximately) constrained using SOCP by discretizing the interval $[0,1]$. Letting $\{0, f_1, \ldots, f_{K-1}, 1\}$ partition the interval $[0,1]$ (choosing $f_k = k/K$ is common), the set of rank-2 SOC constraints

$$\left\| \begin{array}{c} \mathrm{Re}\{E(f_k)\} \\ \mathrm{Im}\{E(f_k)\} \end{array} \right\|_2 \le \delta, \qquad k=0,\ldots,K$$

bounds the maximum value of $|E(f_k)|$ on the grid by δ, which can be a constant or an auxiliary optimization variable. This will be written more compactly in the sequel as

$$|E(f_k)| \leq \delta, \qquad k = 0, \ldots, K.$$

The spacing between the grid points must be chosen to limit violations of the bound between constraints to acceptable levels. Errors between the discretized approximation and the true norm can be made arbitrarily small at the cost of an increased number of constraints. Often a looser grid is used in the initial stages of a design to obtain rapid solutions, while a tight grid is used for the final design. For an FIR filter of length N an approximate rule of thumb for the minimum grid spacing is $1/(20N)$.

If E is a real-valued function, representing a linear-phase filter, then using SOC constraints as shown above results in a degenerate rank-1 cone. Using cones here is inefficient; instead use a pair of linear constraints per frequency:

$$-\delta \leq E(f_k) \leq \delta, \qquad k = 0, \ldots, K.$$

3.2 L_2-Norm Constraints

In contrast to the L_∞ norm, which requires one constraint per discrete frequency, the L_2 norm requires only a single SOC constraint. The L_2 norm is defined as

$$\|E\|_2 \triangleq \left(\int_0^1 |E(f)|^2 \, df \right)^{\frac{1}{2}}.$$

Since E was assumed affine in the optimization variables $\mathbf{x}$, the result can be written as a positive semidefinite quadratic form in $\mathbf{x}$:

$$\|E\|_2^2 = \mathbf{x}^T \mathbf{Q} \mathbf{x}.$$

The real, positive-semidefinite kernel $\mathbf{Q}$ can be factored as $\mathbf{Q} = \mathbf{S}^T \mathbf{S}$, reducing the right side to the desired form:

$$\begin{aligned} \|E\|_2^2 &= (\mathbf{S}\mathbf{x})^T (\mathbf{S}\mathbf{x}) \\ &= \|\mathbf{S}\mathbf{x}\|_2^2. \end{aligned}$$

Thus an L_2 constraint can always be represented exactly with a SOC, provided the integral can be calculated exactly. In many cases this can be achieved using the random process formulation of [23]. If not, the integral must be evaluated using a numerical method, perhaps the simplest being a Riemann sum over a uniform partition $\{f_k\}$ of $[0,1]$. The resulting SOC constraint can be written

$$\sum_k |E(f_k)|^2 \Delta f \leq \delta^2,$$

where δ is a nonnegative constant or auxiliary variable. Compared to the L_∞ norm it takes a tighter grid to get the same level of approximation error using a Riemann sum, a rough rule of thumb being $\Delta f \approx 1/(40N)$. The choice of a small spacing has less impact since it does not increase the size of the optimization problem itself.

3.3 L_1-Norm Constraints

The L_1 norm is the trickiest of the norms to constrain using SOCP as well as the least common. One application is in the design of communication filters to control intersymbol interference [24]. The L_1 norm is defined as

$$\|E\|_1 \triangleq \int_0^1 |E(f)|\, df.$$

Approximating the integral directly as a Riemann sum does not result in a SOC. Instead, introduce the real, nonnegative auxiliary function $\beta : [0,1] \to \mathbb{R}^+$ and consider the pair of constraints

$$|E(f)| \le \beta(f)$$
$$\int_0^1 \beta(f)\, df \le \delta.$$

Due to the nonnegativity of β, this is equivalent to

$$\int_0^1 |E(f)|\, df \le \delta.$$

To discretize, choose a uniform partition $\{f_k\}$ of $[0,1]$ and introduce the auxiliary optimization variables $\{\beta_k\}$. The following set of constraints then approximates the integral:

$$|E(f_k)| \le \beta_k, \quad k = 0, \ldots, K$$
$$\sum_k \beta_k \Delta f \le \delta.$$

As with the L_∞ norm, linear constraints rather than SOC constraints should be used when E is real-valued. Although this method introduces many new variables as well as constraints, the resulting problem is quite sparse and can be efficiently handled by most available solvers. The L_1 norm is even more sensitive than the L_2 norm to small values in its argument, and as such an accurate approximation requires an even tighter grid of frequencies. A rule of thumb is $\Delta f \approx 1/(60N)$.

4 Combining Norms

Although pure L_∞ and L_2 solutions have been used in digital-filter design for some time and remain the most popular in practice[1] in real world designs it is common to want to trade off the properties of the different designs. One way to do this is to optimize the filter with respect to a norm that lies in-between L_2 and L_∞ in some sense. The L_p norm, for example, is defined as

$$\|E\|_p \triangleq \left(\int_0^1 |E(f)|^p\, df \right)^{\frac{1}{p}}, \qquad p \ge 1,$$

and thus provides a continuum of norms from L_1 through L_∞. Although SOCP can handle only $p = \{1, 2, \infty\}$, there are two related methods that can be used to combine two norms to form a new norm. In each case a single scalar parameter controls the relative contribution from each of the component norms. The first, the convex combination of norms denoted here as the α-norm, already enjoys widespread use under various guises. The second, the convex combination of unit balls, is an extended version of the recently introduced ε-norm [17]. The former enjoys better theoretical justification in terms of typical system requirements, while the latter appears to more closely approximate the L_p norm. Both will be described in turn, followed by the derivation of their dual norms.

[1] The window method doesn't count as a filter design technique.

4.1 Convex Norm Combinations—The α-Norm

The most straightforward way to combine the L_2 and L_∞ norms is through a simple convex combination. Letting $\alpha \in [0,1]$, the *α-norm* is defined here as

$$\|E\|_\alpha \triangleq \alpha\|E\|_2 + (1-\alpha)\|E\|_\infty \tag{1}$$

This is essentially equivalent to the norm introduced in [5]. It is straightforward to verify the properties of a norm for the α-norm, as they all follow immediately from the properties of the L_2 and L_∞ norms. As noted in [5], the unconstrained minimization

$$\min_{\mathbf{x}} . \quad \|E\|_\alpha \tag{2}$$

and the constrained minimization

$$\begin{aligned} \min_{\mathbf{x}} . \quad & \|E\|_2 \\ \text{s.t.} \quad & \|E\|_\infty \le \beta \end{aligned} \tag{3}$$

have the same solution space and are thus functionally equivalent. The latter problem has been dubbed *peak-constrained least squares* (PCLS) in the literature [4] or simply constrained L_2 design [2]. An advantage of the unconstrained approach is that a solution is always obtained, as without constraints there can be no infeasible region. The constrained approach, on the other hand, allows the peak error to be explicitly fixed to meet given specifications. In any practical design there will be many iterations, and so the choice is really between iterating α and iterating the peak-error bound β. Either can be formulated as an SOCP: the α-norm can be minimized by the program

$$\begin{aligned} \min_{\mathbf{x},\delta,\beta} . \quad & \alpha\delta + (1-\alpha)\beta \\ \text{s.t.} \quad & |E(f_k)| \le \beta, \ \ k = 0,\ldots,K \\ & \sum_k |E(f_k)|^2 \Delta f \le \delta^2, \ \ \delta \ge 0, \end{aligned} \tag{4}$$

where δ and β are nonnegative auxiliary variables, while the PCLS problem can be solved as

$$\begin{aligned} \min_{\mathbf{x},\delta} . \quad & \delta \\ \text{s.t.} \quad & |E(f_k)| \le \beta, \ \ k = 0,\ldots,K \\ & \sum_k |E(f_k)|^2 \Delta f \le \delta^2, \ \ \delta \ge 0, \end{aligned}$$

with β a constant and δ an auxiliary variable.

The usual justification for using PCLS stems from the observation that a small increase in the peak error of a L_∞-optimal filter frequency response usually allows a large reduction in the L_2 norm. Generally the peak sidelobes of an L_2-optimal frequency response only exceed those of the L_∞-optimal response in a narrow spectral region near the transitions between passband and stopband. PCLS allows these sidelobes to be controlled in an otherwise L_2-optimal design. In terms of general system parameters the L_∞ bound controls the worst-case performance in the presence of unknown spectral components, while the L_2 norm determines the average performance; a natural tradeoff. In the examples, it will be seen that filters optimized according to the combined α-norm retain properties of both L_∞ and L_2-optimized filters.

Some properties of the α-norm can be derived from the definition and the properties of the L_p norms. Clearly L_2 and L_∞ result as special cases of the α-norm with $\alpha = 1,0$ respectively. The norm is monotonically nonincreasing in α, that is $\|E\|_{\alpha_1} \le \|E\|_{\alpha_2}$ for $\alpha_1 \ge \alpha_2$, and is bounded by its component norms: $\|E\|_2 \le \|E\|_\alpha \le \|E\|_\infty$. Both follow from the general $L_p[0,1]$ relation $\|E\|_p \le \|E\|_q$ for $p \le q$.

4.2 Convex Unit-Ball Combinations—The ε-Norm

In the previous section we showed that the α-norm is actually a family of new norms that can be obtained through a direct convex combination of the L_2 and L_∞ norms. Another way to combine the two norms is via their *unit balls*, defined by $B(\|\cdot\|) \triangleq \{E : \|E\| \leq 1\}$. A standard analysis result [25] holds that a norm is uniquely identified by its unit ball, and that further any set that meets certain conditions is the unit ball of a norm, implicitly defining the norm. There are various ways to state the conditions, one of which is given in the following lemma:

Lemma 1. *Let the closed bounded convex set B in a complex vector space X also satisfy*

1. $x \in B \implies e^{j\phi}x \in B,\ \forall \phi \in [0, 2\pi]$ *(symmetry)*
2. $\forall x \in X\ \exists \lambda \in \mathbb{R} : \lambda \neq 0, \lambda x \in B.$ *(absorbing)*

Then B is the unit ball of the norm defined by

$$\|x\| \triangleq \inf\{\lambda \geq 0 : x \in \lambda B\}.$$

Now, in the space $L_2[0,1]$ for $0 < \varepsilon < 1$ consider the set

$$\begin{aligned} B_\varepsilon &= \varepsilon B(\|\cdot\|_2) \oplus (1-\varepsilon)B(\|\cdot\|_\infty) \\ &\triangleq \{\varepsilon U + (1-\varepsilon)V : \|U\|_2 \leq 1, \|V\|_\infty \leq 1\} \end{aligned} \tag{5}$$

a convex Minkowski sum of the unit balls of the L_2 and L_∞ norms. It is straightforward to show that, since both unit balls necessarily meet the conditions of Lemma 1, so does their convex combination. Thus we have implicitly defined a new norm with unit ball B_ε, known as the ε-norm and denoted $\|\cdot\|_\varepsilon$. Clearly for $\varepsilon = 1$ and $\varepsilon = 0$ this reduces to the L_2 and L_∞ norms, and so the ε-norm represents a new family spanning the continuum between the two.

The ε-norm was first described in [17] as a family of vector norms on $\mathbb{R}^n$. The definition and two associated lemmas from that paper, extended here to $L_2[0,1]$, provide greater geometrical insight into the properties of the ε-norm by directly incorporating the properties of the L_∞ norm. The proofs are deferred to the appendix.

Definition: For a function $E \in L_2[0,1]$ and $0 < \varepsilon < 1$, define the ε-norm of E, denoted $\|E\|_\varepsilon$, as the unique solution $\nu^*(E)$ of

$$h(\nu, E) \triangleq \left\| \left[|E| - (1-\varepsilon)\nu \right]_+ \right\|_2 - \varepsilon\nu = 0. \tag{6}$$

The positive operator $[\cdot]_+$ is defined pointwise by $[E(f)]_+ \triangleq \max\{0, E(f)\}$.

Lemma 2. *For any function E the pair of functions*

$$E^\varepsilon = e^{j\angle E} \left[|E| - (1-\varepsilon)\|E\|_\varepsilon \right]_+ \tag{7}$$

$$E^{1-\varepsilon} = E - E^\varepsilon \tag{8}$$

define the unique decomposition $E = E^\varepsilon + E^{1-\varepsilon}$ *called the* ε-decomposition *such that*

$$\|E^\varepsilon\|_2 = \varepsilon\|E\|_\varepsilon \tag{9}$$

$$\|E^{1-\varepsilon}\|_\infty = (1-\varepsilon)\|E\|_\varepsilon. \tag{10}$$

Lemma 3. *The unit ball of the solution to* (6) *is* B_ε.

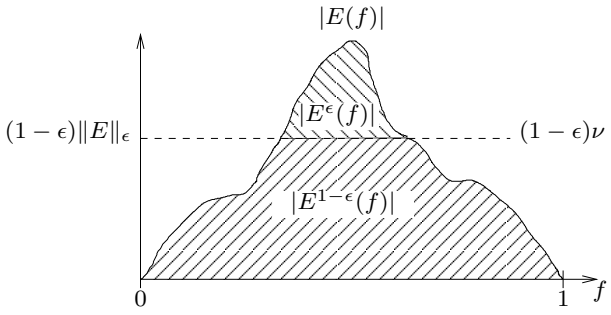

Figure 1: Graphical representation of the ε-decomposition.

This last lemma is included simply to verify that the ε-norm of [17] is indeed the same as the norm implicitly defined in (5). Lemma 2 offers a convenient geometric interpretation of the definition of the ε-norm as illustrated in Fig. 1, the plot of a representative function $|E|$. The dashed line represents the value of $(1-\varepsilon)\nu$ in (6). The portion of $|E|$ below the dashed line is a flat-topped function with L_∞ norm $(1-e)\nu$, while the portion above the line is the function inside the L_2 norm in (6). When $\nu = \|E\|_\varepsilon$ the value of the L_2 norm is $\varepsilon\|E\|_\varepsilon$, the function above the line is $|E^\varepsilon|$, and the function below the line is $|E^{1-\varepsilon}|$. Both E^ε and $E^{1-\varepsilon}$ have the same phase as E and add constructively. The relative "size" of E^ε and $E^{1-\varepsilon}$ varies with ε; for $\varepsilon = 1$, $E^{1-\varepsilon}$ vanishes, as ε is decreased, $E^{1-\varepsilon}$ grows at the expense of E^ε.

For filter design purposes, we need a way to constrain the ε-norm of affine error functions. The following set of constraints use the construction of the norm to together effectively bound $\|E\|_\varepsilon$ by ν:

$$\begin{aligned} \|U\|_2 &\leq \varepsilon\nu \\ \|E-U\|_\infty &\leq (1-\varepsilon)\nu, \end{aligned}$$

where U represents an auxiliary function implemented as a set of complex optimization variables on the discretized set of frequencies. The ε-norm of the error can be approximately minimized via the SOCP

$$\begin{aligned} \min_{\mathbf{x},\nu,\mathbf{U}} \ & \nu \\ \text{s.t.} \ & \sum_k |U_k|^2 \Delta f \leq (\varepsilon\nu)^2 \\ & |E(f_k) - U_k| \leq (1-\varepsilon)\nu, \ \ k = 0,\ldots,K. \end{aligned} \tag{11}$$

4.3 The Dual Norms of the α- and ε-Norms

The α-norm and ε-norm both span a continuum from L_2 to L_∞. To extend both through L_1, recall that the dual norm of L_p, $2 \leq p < \infty$ is L_q, where $1 < q = p/(p-1) \leq 2$, and that L_1 and L_∞ are dual norms. This suggests that the duals of the ε- and α-norms will similarly span a continuum from L_2 to L_1. This is verified by the following lemma.

Lemma 4. *The dual norm of the ε-norm is given by*

$$\|F\|_\varepsilon^D = \varepsilon\|F\|_2 + (1-\varepsilon)\|F\|_1. \tag{12}$$

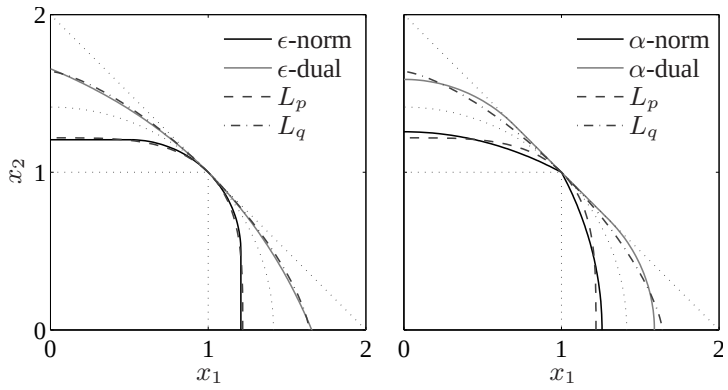

Figure 2: A graphical comparison of the unit balls of the L_p, ε-, and α-norms in the first quadrant of $\mathbb{R}^2$.

The formal definition of the dual norm and the proof of this lemma are found in the appendix.

The lemma reveals that the ε-norm dual is a convex combination of the L_2 and L_1 norms. This suggests a rather satisfying symmetry that is completed with the dual of the α-norm, $\|\cdot\|_\alpha^D$, whose unit ball is a convex combination of the unit balls of the L_2 and L_1 norms:

$$B(\|\cdot\|_\alpha^D) = \alpha B(\|\cdot\|_2) \oplus (1-\alpha) B(\|\cdot\|_1).$$

This norm cannot be defined in the form of (6), which depends on the properties of the L_∞ norm, but we can constrain it similarly with the pair of constraints

$$\|U\|_2 \leq \alpha \nu$$
$$\|E-U\|_1 \leq (1-\alpha)\nu.$$

The preceding shows that the two methods for combining norms are intimately connected: an ε-norm-style convex combination of unit balls of two norms has as its dual norm an α-norm-style convex combination of the duals of the two norms. Together the ε-norm and its dual form a continuum of norms connecting L_1 to L_2 to L_∞, as does the α-norm and its dual.

Although it is difficult to visualize the unit ball of an infinite-dimensional norm, it is instructive to consider the unit balls of the various norms in $\mathbb{R}^2$, representing a (very) coarsely discretized problem. In Fig. 2, the unit balls for the ε-norm and dual and the α-norm and dual are compared in turn to the unit ball of a coarse approximation of L_p,

$$\left(\frac{1}{2}|x_1|^p + \frac{1}{2}|x_2|^p\right)^{1/p}.$$

Here $p = 3.5$, $q = 1.4$, $\varepsilon = 0.5$, and $\alpha = 0.7$, chosen empirically to best match corresponding norms visibly. In both plots the diagonal line, quarter circle, and square, representing the first quadrant of the unit balls of L_1, L_2, and L_∞, respectively, are shown for reference. Inspection reveals the ε-norm to be a closer visual match to L_p than is the α-norm. This observation will be reinforced in the examples.

4.4 Two More Combined Norms

So far the two combined norm/dual norm pairs $\{\|\cdot\|_\varepsilon^D, \|\cdot\|_\varepsilon\}$ and $\{\|\cdot\|_\alpha^D, \|\cdot\|_\alpha\}$ have been formed by combining the L_1, L_2, and L_∞ norms. We can form two more families of norms by forming the

convex combination of the unit balls of L_1 and L_∞, and by convex combination of the two norms themselves. These two new families form a dual pair, and both span the whole range from L_1 to L_∞. These norms do not resemble L_p in general nor do they include L_2 as a special case. However, they do have the advantage that they can be implemented using only linear programming when E is real-valued.

Although limited to pairwise combinations of norms easily constrained using SOCP, the analysis in this paper can be trivially extended to the convex combination of an arbitrary (but finite) number of norms or unit balls, as well as to norms other than L_p.

5 Examples

In this section we present two examples. The first, a standard lowpass filter with reduced group delay, is repeatedly optimized to illustrate the characteristic error responses corresponding to the various norms. The second, a multiband filter, demonstrates the flexibility of the SOCP approach to subject different regions of the frequency response of a filter to different norm constraints.

5.1 Example 1: A Lowpass Filter

A simple lowpass filter was optimized using the various combined-norm SOCP techniques presented in the previous section and compared to L_1, L_2, and L_∞ optimizations as well as the general L_p solution. The filter has 35 real coefficients with no symmetry and a reduced passband delay (relative to a linear phase filter) of 10 samples, resulting in a complex frequency response. The passband and stopband cover the intervals $[0, 0.1]$ and $[0.15, 0.5]$, with the remainder of the frequency-response period defined by symmetry. An error weighting function W of one in the passband, four in the stopband, and zero in the transitions was applied.

For reference, the filter was designed by minimizing the L_1, L_2, and L_∞ norms of the error using SOCP as described in Section 3. The magnitude and weighted error responses of the resulting filters are shown in Fig. 3. Here we clearly see the classic characteristic responses of the three norms. The L_∞ design is equiripple across the entire band. The L_2 design error peaks fall monotonically away from the transition band, and only exceed the L_∞ error close to the transition band. This is a primary reason many argue against pure L_∞ designs. The L_1-optimized error falls off even faster, at a cost of larger errors at the band edges. Qualitatively, L_1 designs put emphasis on lowering the lowest errors, L_∞ designs put emphasis on lowering the highest errors, and L_2 designs occupy a middle ground.

Now consider intermediate designs between the L_2 and L_∞ designs. Three methods have been presented: a true L_p design using a general-purpose or custom solver, and minimizing the α and ε combined norms using the SOCPs in (4) and (11). Each was performed for three intermediate values and are shown in Fig. 4. The values of α and ε are given in the plots for reference only; there is no known convenient mapping between these parameters and the closest approximating value of p. The values were chosen empirically to achieve a good visual match between the various design methods for comparison purposes. As one might expect, the L_p designs provide a smooth transition from L_2 to L_∞, with the slope of the error rolloff gradually decreasing as the overall error level increases. The peak error is still at the transition, but the gap between the highest and lowest error peaks shrinks as p increases. Comparing the L_p plots to the corresponding α-norm designs, we see that the error peaks of the latter no longer fall off monotonically. Instead we have an equiripple region starting at the transition that grows with decreasing α. The equivalent constrained formulation of (3) suggests why: since the largest error ripples typically occur near the transition, the peak-error constraint is active in this region. Away from the transition, the peak constraint is inactive, and the response

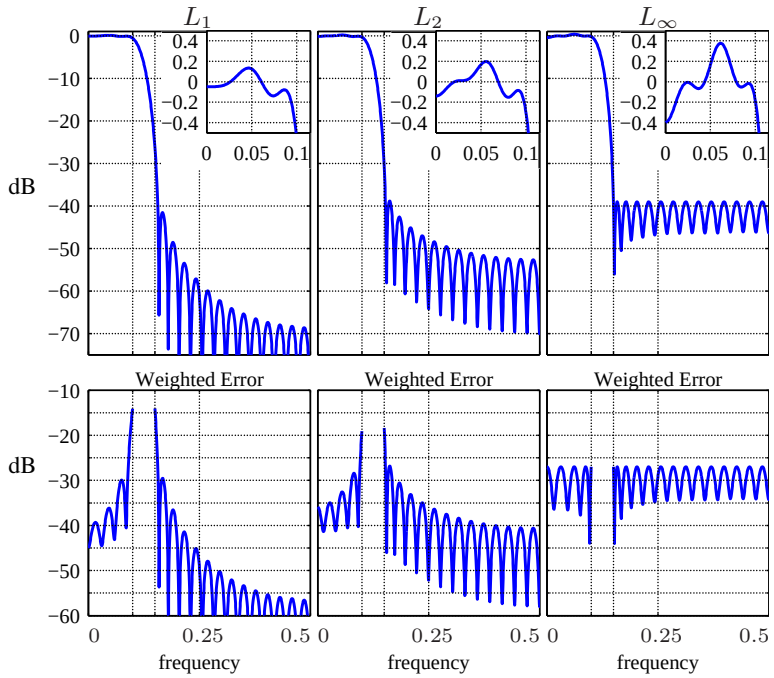

Figure 3: L_1, L_2, and L_∞-optimized FIR filter responses.

resembles a pure L_2 response. Thus the effect of combining the L_2 and L_∞ norms in this way is to partition the frequency axis into L_2 and L_∞ regions. The ε-norm, on the other hand, like the L_p norm, smoothly transitions from equiripple errors to L_2-type errors everywhere on the frequency axis as ε is increased from 0 to 1, and generally has no equiripple errors for $\varepsilon > 0$ (although the lowest error peaks are very nearly equal in height). Visually, at least, the ε-norm designs provide a good approximation of true L_p designs.

To compare norms between L_1 and L_2, the same design was also optimized using the duals of the norms from the previous example. A representative design using each of the three dual norms is shown in Fig. 5. Although the comparison is less obvious here, the ε-norm dual provides a closer approximation to the L_p norm for $1 < p < 2$. Again the α-norm shows a band-partitioning effect: near the transition the peak error is higher and falls off faster like an L_1 design, while near the Nyquist frequency the slope is much flatter like an L_2 design.

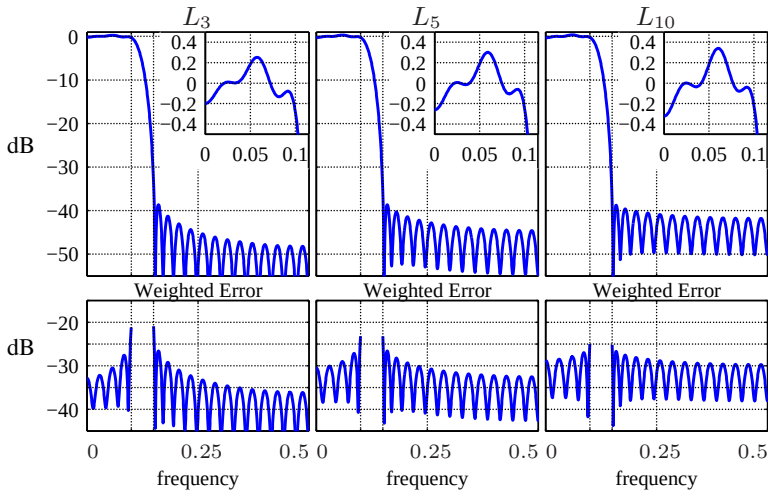

(a) L_p-optimized FIR filter responses, $2 < p < \infty$.

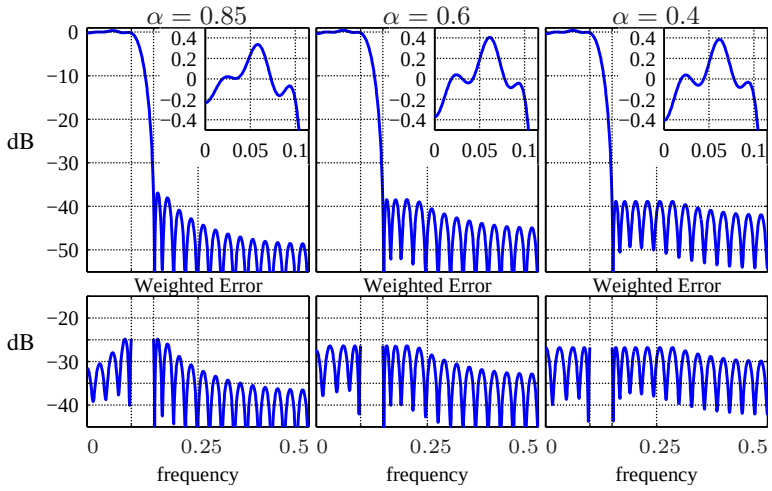

(b) α-norm optimized filters.

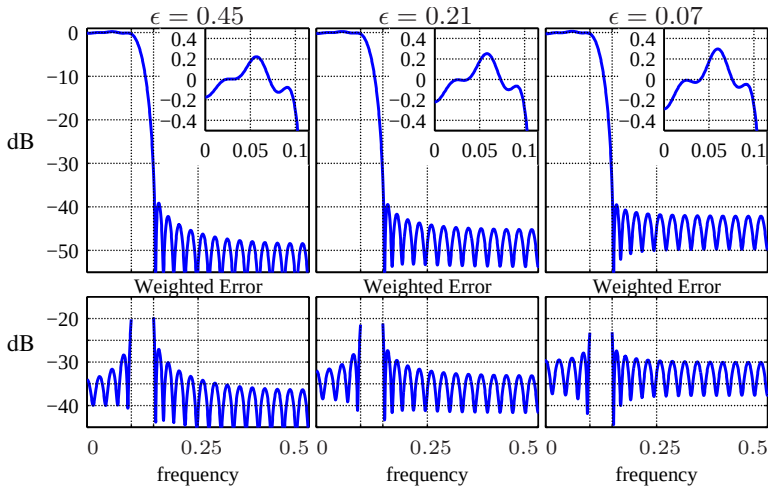

(c) ε-norm optimized filters.

Figure 4: Mixed L_2/L_∞-optimized filters.

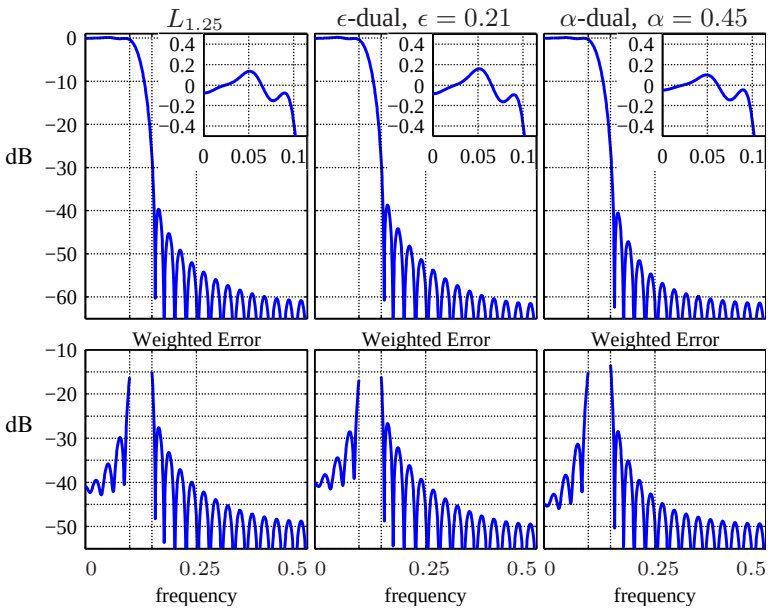

Figure 5: Mixed L_1/L_2-optimized filter responses.

5.2 Example 2: A Multiband Filter

To demonstrate the flexibility of SOCP-based design, a 75-complex-tap nonlinear-phase FIR filter was designed with four passbands and four stopbands. Each passband was designed to approximate a different delay, while the error in each stopband was constrained using a different mixed norm:

$$\begin{aligned} \min_{\mathbf{x}} . \quad & \|E_{\text{pb}}\|_2 \\ \text{s.t.} \quad & \|E_{\text{sb1}}\|_\alpha \leq 10^{-40/20} \\ & \|E_{\text{sb2}}\|_\varepsilon \leq 10^{-40/20} \\ & \|E_{\text{sb3}}\|_\alpha^D \leq 10^{-50/20} \\ & \|E_{\text{sb4}}\|_\varepsilon^D \leq 10^{-50/20}. \end{aligned}$$

The passbands occupy the intervals $[0,0.05]$, $[0.25,0.3]$, $[0.5,0.55]$, and $[0.75,0.8]$ with a single passband weighting function of unity on the intervals and zero elsewhere. The desired response on the passbands is unity gain with delays of 37, 36, 35, and 34. The stopbands occupy the intervals $[0.07,0.23]$, $[0.32,0.48]$, $[0.57,0.73]$, and $[0.82,0.98]$, each with a weighting function of unity on the corresponding interval and zero elsewhere. This was solved using the following SOCP program

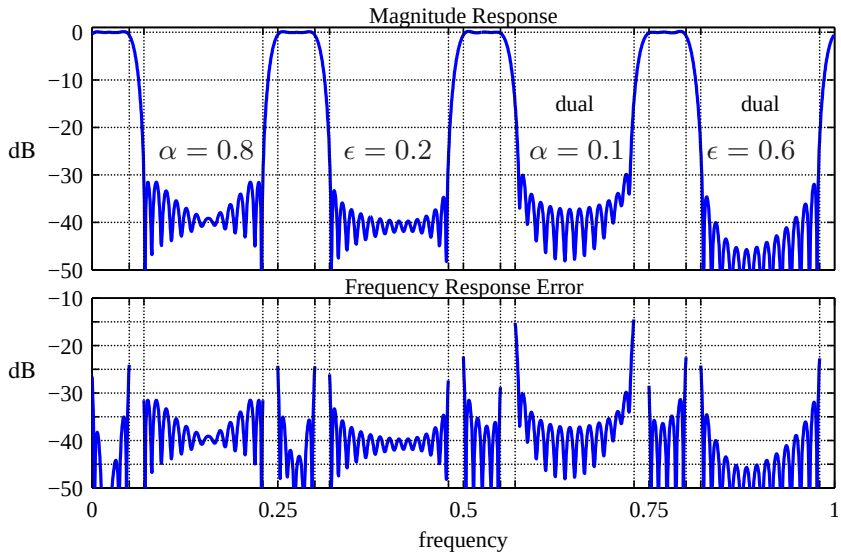

Figure 6: Multiband filter response.

(the horizontal lines delineate the different norm constraints):

$$
\begin{aligned}
\min_{\mathbf{x},\delta,\delta_1,\gamma_1,\{U_{2,k}\},\{\beta_{3,k}\},\delta_4,\{\beta_{4,k}\},\{U_{4,k}\}} \quad & \delta \\
\text{s.t.} \quad & \textstyle\sum_k |E_{\text{pb}}(f_k)|^2 \Delta f \le \delta^2, && \delta \ge 0 \\
\hline
& \textstyle\sum_k |E_{\text{sb1}}(f_k)|^2 \Delta f \le \delta_1^2, && \delta_1 \ge 0 \\
& |E_{\text{sb1}}(f_k)| \le \gamma_1, && k = 0,\ldots,K \\
& \alpha\delta_1 + (1-\alpha)\gamma_1 \le 10^{-40/20} && \\
\hline
& \textstyle\sum_k |U_{2,k}|^2 \Delta f \le \varepsilon^2 10^{-40/10} && \\
& |E_{\text{sb2}}(f_k) - U_{2,k}| \le (1-\varepsilon)10^{-40/20}, && k = 0,\ldots,K \\
\hline
& \textstyle\sum_k |U_{3,k}|^2 \Delta f \le \alpha_D^2 10^{-50/10} && \\
& |E_{\text{sb3}}(f_k) - U_{3,k}| \le \beta_{3,k}, && k = 0,\ldots,K \\
& \textstyle\sum_k \beta_{3,k} \Delta f \le (1-\alpha_D)10^{-50/20} && \\
\hline
& \textstyle\sum_k |E_{\text{sb4}}(f_k)|^2 \Delta f \le \delta_4^2, && \delta_4 \ge 0 \\
& |E_{\text{sb4}}(f_k)| \le \beta_{4,k}, && k = 0,\ldots,K \\
& \varepsilon_D\delta_4 + (1-\varepsilon_D)\textstyle\sum_k \beta_{4,k}\Delta f \le 10^{-50/20}. &&
\end{aligned}
$$

Here ε, α, ε_D, and α_D are constants, and δ, δ_1, γ_1, $\{U_{2,k}\}$, $\{\beta_{3,k}\}$, δ_4, $\{\beta_{4,k}\}$, and $\{U_{4,k}\}$ are auxiliary variables. Although this is a much more complicated specification than one is likely to encounter in practice, it translates to only about 35 lines of Matlab code to solve using the Opt toolbox [15]. The resulting magnitude responses of the filter and the overall error are shown in Fig. 6. The same characteristics seen in the previous example can be seen here in the multiple stopbands.

6 Conclusions

In this paper we demonstrated the ability of SOCP to constrain more interesting FIR filter error norms than the usual L_∞, L_2, and L_1. In particular, the vector ε-norm was extended and shown to be related to a direct norm combination, and it was shown that these combining techniques in fact generate three different norm/dual norm pairs that can be constrained with SOCP. The examples provide a qualitative comparison of the resulting filter responses to true L_p designs. The real strength of SOCP is not its ability to solve these particular unconstrained optimizations, although it can do so efficiently. Rather, it lies in the tremendous flexibility it provides, allowing an arbitrary number of constraints of the sort presented here while at the same time enforcing a uniform structure that allows for efficient solution.

References

[1] J. H. McClellan, T. W. Parks, and L. R. Rabiner, "A computer program for designing optimum FIR linear phase digital filters," *IEEE Trans. Audio Electroacoust.*, vol. AU-21, no. 6, Dec. 1973.

[2] C. S. Burrus, J. A. Barreto, and I. W. Selesnick, "Iterative reweighted least-squares design of FIR filters," *IEEE Trans. Signal Processing*, vol. 42, no. 11, pp. 2926–2936, Nov. 1994.

[3] R. A. Vargas and C. S. Burrus, "Adaptive iterative reweighted least squares design of Lp FIR filters," in *Proc. IEEE Int. Conf. Acoustic, Speech, and Signal Processing*, vol. 3, Mar. 1999, pp. 1129–1132.

[4] J. Adams and J. Sullivan, "Peak-constrained least-squares optimization," *IEEE Trans. Signal Processing*, vol. 46, no. 2, pp. 306–321, Feb. 1998.

[5] W. Lertniphonphun and J. McClellan, "Unified design algorithm for complex FIR and IIR filters," in *Proc. IEEE Int. Conf. Acoustic, Speech, and Signal Processing*, Salt Lake City, UT, May 2001.

[6] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret, "Applications of second-order cone programming," *Linear Algebra and its Applications*, vol. 284, pp. 193–228, Nov. 1998.

[7] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM Review*, vol. 38, no. 1, pp. 49–95, Mar. 1996.

[8] Jos F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," in *Optimization Methods and Software*, 1999, vol. 11–12, pp. 625–653, special issue on Interior Point Methods.

[9] R. Vanderbei, "LOQO: An interior point code for quadratic programming," *Optimization Methods and Software*, vol. 12, pp. 451–484, 1999.

[10] K. Toh, M. Todd, and R. Tutuncu, "SDPT3—a Matlab software package for semidefinite programming," *Optimization Methods and Software*, vol. 11, pp. 545–581, 1999.

[11] E. D. Andersen and K. D. Andersen, "The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm," in *High Performance Optimization*. Kluwer Academic Publishers, 2000, pp. 197–232.

[12] J. O. Coleman, "The use of the FF design language for the linear-programming design of finite-impulse-response digital filters for digital-communication and other applications," Ph.D. dissertation, University of Washington, Seattle, WA, Dec. 1991.

[13] K. Steiglitz, T. W. Parks, and J. F. Kaiser, "METEOR: A constraint-based FIR filter design program," *IEEE Trans. Signal Processing*, vol. 40, no. 8, pp. 1901–1909, Aug. 1992.

[14] S. Boyd, V. Balakrishnan, C. Barratt, N. Khraishi, X. Li, D. Meyer, and S. Norman, "A new CAD method and associated architectures for linear controllers," *IEEE Trans. Automat. Contr.*, vol. 33, no. 3, pp. 268–283, Mar. 1988.

[15] J. O. Coleman, D. P. Scholnik, and J. J. Brandriss, "A specification language for the optimal design of exotic FIR filters with second-order cone programs," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Pacific Grove, CA, Nov. 2002.

[16] L. E. Ghaoui, R. Nikoukhah, and F. Delebecque, "LMITOOL: a package for LMI optimization," in *Proc. IEEE Conf. on Decision and Control*, Dec. 1995.

[17] O. Burdakov and B. Merkulov, "On a new norm for data fitting and optimization problems," Linköping University, Linköping, Sweden, Tech. Rep. LiTH-MAT-R-2001-29, 2001.

[18] D. P. Scholnik, "Mixed-norm FIR filter optimization using second-order cone programming," in *Proc. IEEE Int. Conf. Acoustic, Speech, and Signal Processing*, Orlando, FL, May 2002.

[19] K. Steiglitz and L. E. McBride, "A technique for the identification of linear systems," *IEEE Trans. Automat. Contr.*, vol. 10, pp. 461–464, Oct. 1965.

[20] R. A. Vargas and C. S. Burrus, "On the design of Lp IIR filters with arbitrary frequency response," in *Proc. IEEE Int. Conf. Acoustic, Speech, and Signal Processing*, vol. 6, May 2001, pp. 3829–3832.

[21] P. Regalia, "Comments on 'A weighted least-squares method for the design of stable 1-D and 2-D IIR digital filters'," *IEEE Trans. Signal Processing*, vol. 47, no. 7, pp. 2063–2065, July 1999.

[22] J. O. Coleman and D. P. Scholnik, "Design of nonlinear-phase FIR filters with second-order cone programming," in *Proc. 1999 Midwest Symp. on Circuits and Systems (MWSCAS '99)*, Las Cruces NM, Aug. 1999.

[23] J. O. Coleman, "Systematic mapping of quadratic constraints on embedded FIR filters to linear matrix inequalities," in *Proc. 1998 Conf. on Information Sciences and Systems (CISS '98)*, Princeton, NJ, Mar. 1998.

[24] J. O. Coleman and D. W. Lytle, "Linear-programming techniques for the control of intersymbol interference with hybrid FIR/analog pulse shaping," in *Proc. IEEE Int'l Conf. on Communications*, June 1992.

[25] R. T. Rockafellar, *Convex Analysis*. Princeton University Press, 1970.

[26] H. L. Royden, *Real Analysis*, 3rd ed. Prentice Hall, 1988.

A The ε-Norm

Definition: For a function $E \in L_2[0,1]$ and $0 < \varepsilon < 1$, define the ε-norm of E, denoted $\|E\|_\varepsilon$, as the unique solution $\nu^*(E)$ of

$$h(\nu, E) \triangleq \left\| \left[|E| - (1-\varepsilon)\nu \right]_+ \right\|_2 - \varepsilon\nu = 0. \tag{13}$$

Proof. We must show that $\nu^*(E)$ exists, is unique and possesses the properties of a norm, namely

1. $\|E\|_\varepsilon \geq 0$ (Positivity)
2. $\|E\|_\varepsilon = 0 \iff E = 0$ (Positive definiteness)
3. $\|\gamma E\|_\varepsilon = |\gamma| \|E\|_\varepsilon$ (Homogeneity)
4. $\|E+F\|_\varepsilon \leq \|E\|_\varepsilon + \|F\|_\varepsilon$. (Triangle Inequality)

To show existence and uniqueness, first note that $h(\nu, E)$ is a continuous, monotonically decreasing function of ν. For any $\nu > \|E\|_\infty/(1-\varepsilon)$, we see that $h(\nu, E) = -\varepsilon\nu < 0$, and since $h(0,E) = \|E\|_2 \geq 0$ as well, there must exist a unique solution $\|E\|_\varepsilon \triangleq \nu^*(E)$. The positivity of $\nu^*(E)$ is guaranteed by the positivity of the L_2 norm. Positive definiteness (in the sense of L_p) follows directly from (13) and the positive definiteness of the L_2 norm. Homogeneity follows from the relation $h(|\alpha|\nu, \alpha E) = |\alpha| h(\nu, E)$, which implies that $h(|\alpha|\nu^*(E), \alpha E) = 0$ and $\nu^*(\alpha E) = |\alpha|\nu^*(E)$. To show the triangle inequality we use the triangle inequality for the L_2 norm and the fact that $[E+F]_+ \leq [E]_+ + [F]_+$ (in a pointwise sense):

$$\begin{aligned}
h(\nu^*(E)+\nu^*(F), E+F) + \varepsilon(\nu^*(E)+\nu^*(F)) & \\
&= \left\| \left[|E+F| - (1-\varepsilon)(\nu^*(E)+\nu^*(F)) \right]_+ \right\|_2 \\
&\leq \left\| \left[|E| - (1-\varepsilon)\nu^*(E) \right]_+ + \left[|F| - (1-\varepsilon)\nu^*(F) \right]_+ \right\|_2 \\
&\leq \left\| \left[|E| - (1-\varepsilon)\nu^*(E) \right]_+ \right\|_2 + \left\| \left[|F| - (1-\varepsilon)\nu^*(F) \right]_+ \right\|_2 \\
&= h(\nu^*(E), E) + h(\nu^*(F), F) + \varepsilon(\nu^*(E)+\nu^*(F)) \\
&= \varepsilon(\nu^*(E)+\nu^*(F)).
\end{aligned}$$

Thus, $h(\nu^*(E)+\nu^*(F), E+F) \leq 0$, and from the monotonicity of $h(\nu, E+F)$ in ν we have $\nu^*(E+F) \leq \nu^*(E)+\nu^*(F)$ as desired. □

Lemma 2. *For any function E the pair of functions*

$$E^\varepsilon = e^{j\angle E} \left[|E| - (1-\varepsilon)\|E\|_\varepsilon \right]_+ \tag{14}$$

$$E^{1-\varepsilon} = E - E^\varepsilon \tag{15}$$

define the unique decomposition

$$E = E^\varepsilon + E^{1-\varepsilon} \tag{16}$$

called the ε-decomposition *such that*

$$\|E^\varepsilon\|_2 = \varepsilon \|E\|_\varepsilon \tag{17}$$

$$\|E^{1-\varepsilon}\|_\infty = (1-\varepsilon)\|E\|_\varepsilon. \tag{18}$$

Proof. The support of E^ε is restricted to the set $I_* = I_*(E) \triangleq \{f : |E(f)| > (1-\varepsilon)\|E\|_\varepsilon\}$. By construction we have (16), (17), and (18), and thus the ε-decomposition exists. When $E = 0$ uniqueness is clear, so let $E \neq 0$ and choose a function $V \neq E^{1-\varepsilon}$ such that $\|V\|_\infty \leq (1-\varepsilon)\|E\|_\varepsilon$. Uniqueness then follows by showing that for $U = E - V$ we have $\|U\|_2 > \varepsilon\|E\|_\varepsilon$, violating (17). Using the inner-product notation $\|U\|_2^2 = (U,U)$ we have

$$\begin{aligned}\|U\|_2^2 = \|E-V\|_2^2 &= \|E^\varepsilon + E^{1-\varepsilon} - V\|_2^2 \\ &= \|E^\varepsilon\|_2^2 + 2\Re\big\{(E^\varepsilon, E^{1-\varepsilon} - V)\big\} + \|E^{1-\varepsilon} - V\|_2^2 \\ &> \varepsilon^2\|E\|_\varepsilon^2 + 2\Re\big\{(E^\varepsilon, E^{1-\varepsilon} - V)\big\}\end{aligned}$$

since $\|E^{1-\varepsilon} - V\|_2^2 > 0$ results from positive definiteness. Now consider the remaining inner-product term

$$\begin{aligned}2\Re\big\{(E^\varepsilon, E^{1-\varepsilon} - V)\big\} &= \int_{t\in I^*} 2\Re\big\{E^\varepsilon(f)\big(E^{1-\varepsilon}(f) - V(f)\big)^*\big\}\,dt \\ &= \int_{t\in I^*} 2\big(|E(f)| - (1-\varepsilon)\|E\|_\varepsilon\big)\big((1-\varepsilon)\|E\|_\varepsilon - \Re\big\{V^*(f)e^{j\angle E(f)}\big\}\big)\,dt.\end{aligned}$$

The first factor in the integrand is just $2|E^\varepsilon(f)|$ and thus is positive on I^* (and zero elsewhere). The second factor is nonnegative on I^* from the restriction on V. Since the integrand is everywhere nonnegative the integral is also, thus $\|U\|_2^2 > \varepsilon^2\|E\|_\varepsilon^2$ and the decomposition is unique. □

Lemma 3. *The unit ball of the ε-norm is a convex Minkowski sum of the unit balls of the L_2 and L_∞ norms:*

$$\begin{aligned}\{E : \|E\|_\varepsilon \leq 1\} &= \varepsilon B(\|\cdot\|_2) \oplus (1-\varepsilon)B(\|\cdot\|_\infty) \\ &\triangleq \{\varepsilon U + (1-\varepsilon)V : \|U\|_2 \leq 1,\ \|V\|_\infty \leq 1\}.\end{aligned} \tag{19}$$

Proof. By Lemma 2 a suitable $\varepsilon U = E^\varepsilon$ and $(1-\varepsilon)V = E^{1-\varepsilon}$ will always exist given $\|E\|_\varepsilon \leq 1$. It remains to show that for any U and V such that $\|U\|_2 \leq 1$ and $\|V\|_\infty \leq 1$, we have $\|\varepsilon U + (1-\varepsilon)V\|_\varepsilon \leq 1$. Apply the ε-decomposition to get $E = \varepsilon U + (1-\varepsilon)V = E^\varepsilon + E^{1-\varepsilon}$, then using an argument similar to the proof of Lemma 2 we have

$$\begin{aligned}\varepsilon^2 \geq \|\varepsilon U\|_2^2 &= \|E^\varepsilon + E^{1-\varepsilon} - (1-\varepsilon)V\|_2^2 \\ &= \|E^\varepsilon\|_2^2 + \|E^{1-\varepsilon} - (1-\varepsilon)V\|_2^2 + 2\Re\big\{(E^\varepsilon, E^{1-\varepsilon} - (1-\varepsilon)V)\big\} \\ &\geq \varepsilon^2\|E\|_\varepsilon^2\end{aligned}$$

which completes the proof. □

Lemma 4. *The dual norm of the ε-norm is given by*

$$\|F\|_\varepsilon^D = \varepsilon\|F\|_2 + (1-\varepsilon)\|F\|_1. \tag{20}$$

Proof. The dual space X^* of a complex Banach space X is the space of all bounded complex linear functionals on X. The dual space is also a Banach space, with norm [26]

$$\|f\|^D \triangleq \sup_{\|E\|\leq 1} |f(E)|. \tag{21}$$

The space of all complex-valued square-integrable functions on $[0,1]$ forms a Banach space with either the L_2 or ε-norms. We use the definition of the ε-norm unit ball to write the ε-norm dual as

$$\|f\|_\varepsilon^D = \sup_{\substack{\|U\|_2\leq 1 \\ \|V\|_\infty \leq 1}} \big|f\big(\varepsilon U + (1-\varepsilon)V\big)\big|.$$

Since f is linear and both the L_2 and L_∞ norms depend only on the magnitude of their argument, suitable U and V functions can always be found such that

$$\left|f\big(\varepsilon U + (1-\varepsilon)V\big)\right| = \varepsilon|f(U)| + (1-\varepsilon)|f(V)|$$

yielding

$$\|f\|_\varepsilon^D = \varepsilon \sup_{\|U\|_2\leq 1} |f(U)| + (1-\varepsilon) \sup_{\|V\|_\infty\leq 1} |f(V)|. \tag{22}$$

The ε-dual norm is thus the convex combination of the dual norms of L_2 and L_∞. By the Riesz Representation theorem, all bounded linear functionals on the space $L_p[0,1]$ can be represented as the inner product $f(E) = (E,F)$, where $E \in L_p[0,1]$ and $F \in L_q[0,1]$ for $q = p/(p-1)$. The functional f is thus fully represented by the $L_q[0,1]$ function F. Hölder's inequality further tells us that $|(E,F)| \leq \|E\|_p\|F\|_q$, so that the dual norm of L_p is L_q. Rewriting (22) in terms of F yields

$$\begin{aligned}\|F\|_\varepsilon^D &= \varepsilon \sup_{\|U\|_2\leq 1} |(U,F)| + (1-\varepsilon) \sup_{\|V\|_\infty\leq 1} |(V,F)| \\ &= \varepsilon\|F\|_2^D + (1-\varepsilon)\|F\|_\infty^D\end{aligned}$$

and applying the known result for L_p completes the proof. □