The 1999 Midwest Symposium on Circuits and Systems, New Mexico State University, Aug. 1999

Design of Nonlinear-Phase FIR Filters with Second-Order Cone Programming

Jeffrey O. Coleman
Dan P. Scholnik

jeffc@alum.mit.edu
scholnik@nrl.navy.mil

Naval Research Laboratory
Radar Division
Washington, DC

***Abstract*—We extend classic linear-programming design of linear-phase FIR filters to second-order cone-programming design of nonlinear-phase FIR filters. Real and complex FIR filters with optional analog or digital cascade filters are designed using cones on discrete frequencies to approximate Chebychev, mean-absolute, and rms errors.

I. Introduction

Linear programming has long been used for the design of linear-phase FIR filters [1], [2], especially when there are constraints on both time and frequency responses [3]–[5]. Typically a frequency-response constraint across a band is approximated using a grid of frequencies closely spaced across that band, with a constraint of the form $|H(f)-H_{\text{ideal}}(f)| \leq b$ implemented as the pair of inequalities $-b \leq H(f)-H_{\text{ideal}}(f) \leq b$. At any given grid frequency f a pair of linear constraints in the linear program results because the expression $H(f) - H_{\text{ideal}}(f)$ is affine in the design variables, which for a real filter are just the filter weights being optimized and for a complex filter are the real and imaginary components of those weights.

For a nonlinear-phase filter, however, $H(f) - H_{\text{ideal}}(f)$ is complex and must be confined by constraint $|H(f) - H_{\text{ideal}}(f)| \leq b$ to a circle of radius b in the complex plane. Such a constraint is a so-called second-order cone in the design variables (which may be augmented to include b). A second-order cone program (SOCP) combines a family of such cones with a linear objective and can be solved (optimized) with a new generation of extraordinarily efficient solvers [6]–[10] based on interior-point methods. Most of these solvers accept linear constraints as well as cones.

The next section shows how bounds on L_∞ (Chebychev), L_2 (rms), and L_1 (mean-absolute) norms can be approximated using families of second-order cones. The result is a flexible framework for the design of FIR filters based on frequency-domain specifications that can be used to design real or complex filters with linear or nonlinear phase (the latter emphasized here), mixed analog/FIR filter cascades, and "nonuniform" FIR filters corresponding to analog tapped-delay-line (TDL) filters with taps nonuniformly spaced.

This work was supported by National Science Foundation grant MIP-9896034 to the University of Maryland Baltimore County (where Coleman and Scholnik are on the adjunct faculty and in the Ph.D. program respectively) and by the Office of Naval Research through its Base Program at the Naval Research Laboratory.

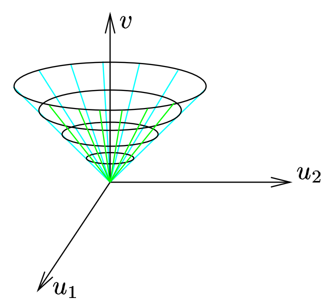

Fig. 1. Rank-two cone $||\mathbf{u}|| \leq v$.

Although it will not be explored further here, the big attraction of the SOCP approach is that once these simple ideas for using cones to constrain frequency responses are mastered, the extension to all manner of interesting and useful constraints in both time and frequency domains becomes straightforward. Simple variants of the L_1 technique discussed here, for example, parallel but improve upon a more-primitive[1] linear-programming approach to the constraint of peak intersymbol interference (ISI) and peak adjacent-channel interference (ACI) in the design of a filter for a digital-communication receiver [4]. Modification to address rms ISI and ACI instead would be straightforward using the L_2 approach of this paper.[2]

II. Theory

Second-Order Cones If $\mathbf{u} \in \mathbb{R}^M$ and $v \in \mathbb{R}$ are each affine (linear plus a constant) in optimization variables $x_1, \ldots, x_N$, then constraint $||\mathbf{u}|| \leq v$ is a *second-order cone* in those variables [6]. With $\mathbf{u}$ and v the horizontal and vertical components respectively of a three vector, the constraint surface takes the 3D cone shape shown in Fig. 1. We loosely speak of M, two in this example, as the *rank* of the cone.

FIR/TDL Frequency Response The filter structure considered here is a tapped-delay-line (TDL) filter, which generalizes a standard FIR filter by allowing nonuniform and noninteger tap spacing. The impulse response of a TDL filter is given

[1] It approximated the unit circle by a polygon.
[2] In this case the approach of [11] is cleaner, however.

by $h(t) = \sum_k h_k \delta(t - T_k)$, where coefficient h_k is the complex weight applied to the signal from a tap at delay T_k. In this paper the delays $\{T_k\}$ are assumed fixed. Only the weights $\{h_k\}$ are optimized. The frequency response corresponding to $h(t)$ is $H(f) = \sum_k h_k e^{-j2\pi f T_k}$, which is periodic with period $1/T$ for the common FIR case $T_k = kT$. At any particular frequency this frequency response is linear in the coefficients $\{h_k\}$ and thus in their real and imaginary parts, which are generally optimization variables (or affine combinations of them).

In filter design we often wish to bound or minimize an L_∞, L_1, or L_2 norm of a frequency-response expression on a subset $\mathcal{F}$ of the real line. Below we show the expression as $H(f)$ with the understanding that it might typically be $H(f)$ in a stopband constraint but $H(f) - e^{-j2\pi f\tau}$ in a passband constraint or perhaps $A(f)H(f)$ or $A(f)H(f) - e^{-j2\pi f\tau}$ when the specifications apply to the cascade with some analog filter. Here τ is the desired net passband delay of the filter or cascade. Occasionally $\mathcal{F}$ is discrete, but more often it is not, in which case approximation by some finite set $\mathcal{F}_d$ of closely spaced frequencies is key.

The L_∞ or Chebychev norm. To approximate the Chebychev-norm bound $\max_{\mathcal{F}}\{|H(f)|\,W(f)\} \leq \delta$, where $W(f)$ is a nonnegative real weighting function, we substitute a finite grid of frequencies and use the family of constraints

$$\{|H(f)|\,W(f) \leq \delta : f \in \mathcal{F}_d\}.$$

Each of these constraints is equivalent to

$$\left\| \begin{pmatrix} \mathrm{Re}\{\mathrm{H(f)W(f)}\} \\ \mathrm{Im}\{\mathrm{H(f)W(f)}\} \end{pmatrix} \right\| \leq \delta$$

and so is a rank-two cone (rank one if $H(f)$ is purely real or imaginary).[3] The left side is minimized (rather than bounded) if δ is both an optimization variable and the objective function.

Examples in the Literature. Minimization and bounding of Chebychev norms of frequency responses in this fashion has been used in the SOCP design of simple equiripple filters [6] and, in a two-dimensional parallel, simple sensor arrays [12]. Lu [13] goes on to formulate nonlinear-phase FIR-filter designs in which an L_2 frequency-response norm is minimized subject to Chebychev constraints on both passband and stopband. He used a semidefinite-program (SDP) formulation, but if each of his δ variables is replaced with the corresponding δ^2, each of his linear matrix inequalities can be formulated instead as a second-order cone. His designs are therefore fully equivalent to SOCPs. Coleman [11] and Scholnik [14] also use SOCP-equivalent SDPs to bound/minimize a variety of frequency- and time-domain L_2 norms. A different approach to frequency-domain L_2 norms is discussed next.

[3] Since our approximating family of rank-two cones is formally $||H||_\infty \leq \delta$ in the space $L_\infty(\mathbb{R}, \mu)$, where positive measure μ is concentrated on $\mathcal{F}_d$ and assigns mass $W(f)$ to each $f \in \mathcal{F}_d$, we are technically justified in speaking here of bounding a norm.

The L_2 or rms norm. Reference [11] discusses the systematic reduction of a fairly general L_2 bound of the form

$$\left(\int_{\mathcal{F}} |H(f)|^2\,W(f)\,df\right)^{\frac{1}{2}} \leq \delta$$

first to a second-order cone (unidentified as such) and then to a linear matrix inequality. As presented, however, that formulation cannot handle an analog cascade filter except in special cases in which it can be absorbed into weight function $W(f)$. Here we use a more-general but less-elegant approach and replace the exact bound with the (generally high-rank) cone

$$\left(\sum_{f \in \mathcal{F}_d} |H(f)|^2\,W(f)\,\Delta f\right)^{\frac{1}{2}} \leq \delta.$$

The L_1 or mean-absolute norm. Augment the optimization variables with $\{\delta_f\}$ and enforce rank-two cone constraints $\{|H(f)|\,W(f)\Delta f \leq \delta_f : f \in \mathcal{F}_d\}$ and linear constraint $\sum_{f\in\mathcal{F}_d} \delta_f \leq \delta$ to obtain

$$\int_{\mathcal{F}} |H(f)|\,W(f)\,df \approx \sum_{f\in\mathcal{F}_d} |H(f)|\,W(f)\,\Delta f \leq \sum_{f\in\mathcal{F}_d} \delta_f \leq \delta,$$

which bounds an approximation to the L_1 norm on the left.

III. Examples

Two examples below spark discussions for options for stopband and passband control. The second also illustrates optimization of an embedded filter to system specifications.

Example—stopband error measures. Figure 2 shows the magnitude response $20\log_{10}|H(f)|$ versus fT of a nonlinear-phase FIR filter $H(f)$ with 50 (uniformly spaced) complex weights optimized to create an analytic filter with a stopband for $-0.45 \leq fT \leq -0.03$, a Chebychev passband for $0.03 \leq fT \leq 0.45$, and with net (arbitrarily chosen) passband delay of $\tau = 25T$. In each of 526 rank-two second-order cone constraints spaced uniformly across the passband, the complex passband error is fixed at $1 - 10^{-0.1/20}$, corresponding to approximately 0.2 dB peak-to-peak ripple. A like number of rank-two second-order cone constraints spaced uniformly across the stopband are used in minimization of stopband error using each of the three norms. The optimization was set up in `matlab` for solution by SeDuMi [9], an interior-point solver that can accept both second-order cone and linear constraints.

Asymmetric L_2 and L_1 stopbands result from transition regions of differing widths. As usual, going from a Chebychev stopband to an L_2 stopband costs some attenuation close to the transition band but gains substantially greater attenuation elsewhere. This effect is more pronounced in the L_1 case, because there is more minimization pressure on the higher sidelobes

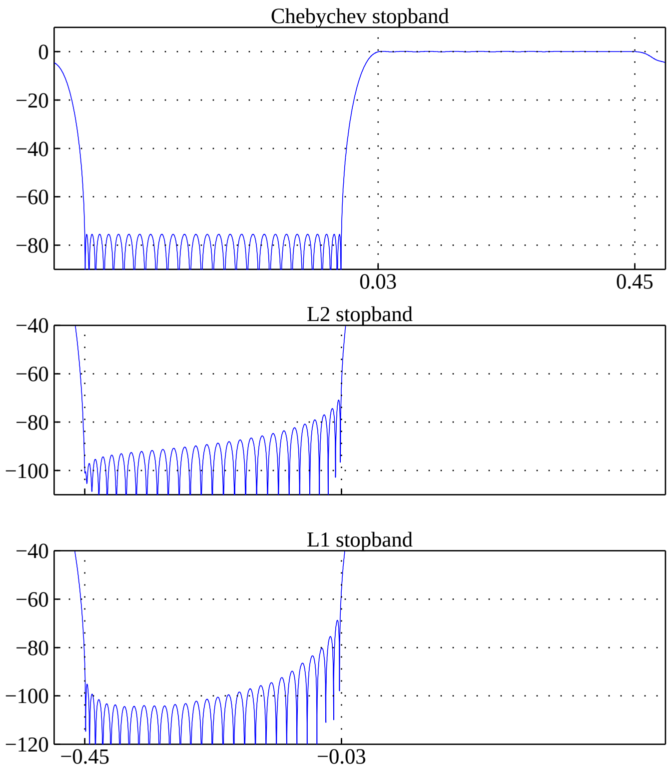

Fig. 2. Gain (dB) vs. fT for a 50 tap nonlinear-phase complex FIR analytic filter with a Chebychev passband and an optimal stopband (three norms).

norm	*iterations*	*CPU seconds*
Chebychev	22	58.6
L_1	24	515.4
L_2	25	33.6

TABLE I
Optimization runtimes for Fig. 2 using a 266 MHz Pentium II under Linux.

in the L_2 optimization than in the L_1 optimization due to the squaring in the norm in the former. This comparison suggests (and we believe) that Chebychev stopbands are seldom justified. Table I shows the CPU time and the number of interior-point iterations required for the optimization by SeDuMi. The L_1 optimization requires more time because of the large number of auxiliary optimization variables $\{\delta_f\}$.

Example—an embedded filter. Figure 3 illustrates the design of a complex nonlinear-phase FIR filter for use with a D/A hold function and analog postfiltering. The cascade is an analytic filter suitable for a modulation application [15], [16]. The cascade passband ranges over $0.85 \leq fT \leq 1.65$ with the center $fT = 1.25$ as the desired carrier frequency. Sixth-order (in the lowpass prototype) Chebychev bandpass filter $A(f)$, with 1 dB ripple and corner frequencies fT of 0.7 and 1.7, provides close-in stopband attenuation at the expense of passband distortion. FIR filter $H(f)$, with 31 taps spaced at the hold-function width of $T/5$, provides passband equalization and

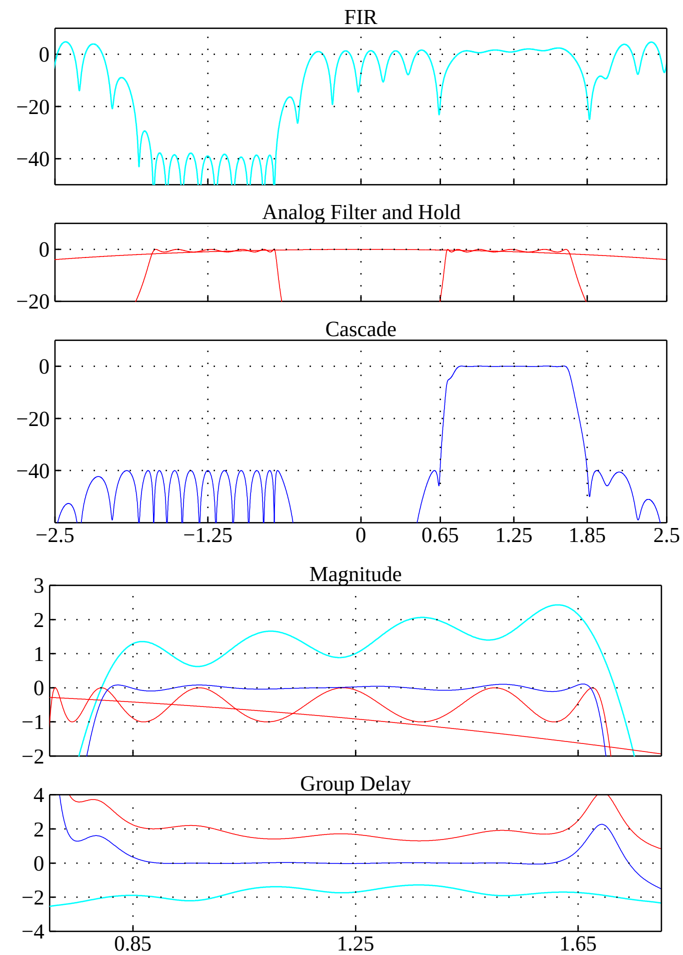

Fig. 3. The cascade of this optimized FIR filter with an analog filter and D/A hold function has a Chebychev stopband and a minimum rms-error passband.

suppression of the image spectrum. The cascade's magnitude response $|H(f)D(f)A(f)|$ is constrained to a -40 dB level across a Chebychev stopband excluding only $.65 < fT < 1.85$ (an L_1 or L_2 stopband might typically be more appropriate), and the L_2 norm of passband error $H(f)D(f)A(f) - e^{-j2\pi f\tau}$ (for a suitable τ) is minimized. Rank-two cones are used for constraints in the passband, stopband, and transition bands (to prevent huge gains), and the ratio of FIR frequency-response period to grid spacing is fixed at thirty times the number of taps. (Five or ten times is adequate for preliminary results.) The optimization requires 30 seconds evenly split between the `matlab` problem setup and 23 iterations in SeDuMi.

Alternatives for controlling passband error. There are a variety of alternatives for controlling mean-square passband error. At any particular frequency f, a passband error like $H(f)D(f)A(f) - e^{-j2\pi f\tau}$ is affine in frequency response $H(f)$, which in turn is affine in the filter weights whose real and imaginary parts are some subset (generally) of the optimization variables. Therefore, the real and imaginary parts of

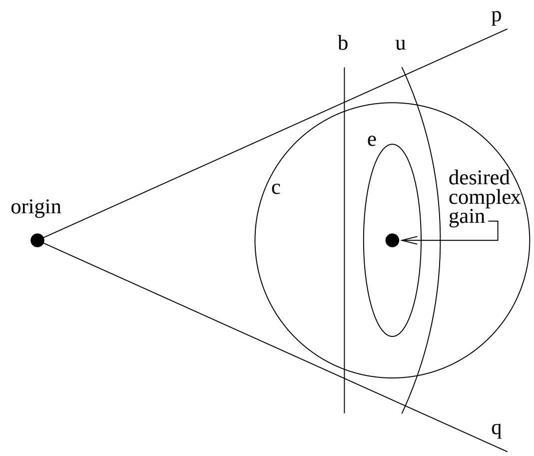

Fig. 4. Some convex error bounds in the complex plane that can be used in passband constraint or optimization.

the passband error are affine in those optimization variables. Straight-line bounds on that error (or the response itself) in the complex plane become (so-called) linear constraints (they are actually affine) on the variables, and elliptical or circular upper bounds on the passband error (or the response itself) become simple second-order cones in those variables.

Some of these passband-error bounds are illustrated in Fig. 4, which shows the complex plane of the passband frequency response, $H(f)D(f)A(f)$ in our example, at some fixed frequency f. The desired complex gain at frequency f, here $e^{-j2\pi f\tau}$, appears on the right as a point. At the given frequency, bounding this gain between lines p and q would bound the phase response, and confining the error to circle c would bound the complex frequency-response error. Ellipse e could be used instead if magnitude should be more tightly constrained than phase. Because second-order cones are necessarily convex constraints, using them to exactly bound magnitude error is generally not possible, because the annulus to which the frequency response would have to be confined is nonconvex and therefore not representable as the intersection of convex regions. But when phase is narrowly constrained, confining the response between some line b and some circle u can be used to approximately bound magnitude.

IV. Conclusions

Second-order cone programming is now available as a reasonably approachable tool for constrained optimization of filters of substantial complexity in reasonable times on modest computers. The approach allows much more design flexibility than the classical minimax, least-squares, and eigenfilter approaches, and the speed advantages of these classic methods are quickly becoming irrelevant in all but special cases like huge filters or applications requiring on-the-fly design for real-time applications. For the past year we have exclusively used linear-programming (occasionally) and SOCP (usually) when designing FIR filters. We have used SOCP approaches to design linear- and nonlinear-phase filters; real and complex filters; lowpass, highpass, bandpass, analytic, and time-delay filters; standalone filters and filters with specifications on a cascade; single-rate and multi-rate filters; halfband and other Nyquist filters; one-dimensional and two-dimensional filters; and filters with over 600 real coefficients. Based on this experience, we propose that SOCP deserves a central place in the tool collections of DSP engineers.

References

[1] S.-P. Wu, S. Boyd, and L. Vandenberghe, *Applied Computational Control, Signals and Circuits*, chapter entitled “FIR filter design via spectral factorization and convex optimization”, pp. 51–81, Birkhauser Boston, 1997.

[2] A. Said and J. B. Anderson, “Bandwidth-efficient coded modulation with optimized linear partial-response signals,” *IEEE Trans. Inform. Theory*, vol. 44, no. 2, pp. 701–713, Mar. 1998.

[3] J. O. Coleman and Dean W. Lytle, “Linear-programming techniques for the control of intersymbol interference with hybrid FIR/analog pulse shaping,” in *Conf. Record, Int’l Conf. on Communications*, Chicago, June 1992, pp. 329.2.1–329.2.6.

[4] J. O. Coleman, “Linear-programming design of hybrid analog/FIR receive filters to minimize worst-case adjacent-channel interference,” in *Proc. 1993 Conference Information Sciences and Systems*, Mar. 1993.

[5] S. Boyd, L. Vandenberghe, and M. Grant, “Efficient convex optimization for engineering design,” in *Proc. IFAC Symposium on Robust Control Design*, Rio de Janeiro, Brazil, Sept. 1994.

[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] R. Vanderbei, “LOQO: An interior point code for quadradic programming,” Technical Report SOR 94-15, Princeton University, 1994, revised 11/30/98, to appear in *Optimization Methods and Software*.

[8] R.J. Vanderbei and D.F. Shanno, “An interior-point algorithm for nonconvex nonlinear programming,” Tech. Rep. SOR-97-21, Statistics and Operations Research, Princeton University, 1997, To appear in *Computational Optimization and Applications*.

[9] J. F. Sturm, “SeDuMi,” `matlab` toolbox, May 1998, http://www.unimaas.nl/~sturm/software/sedumi.html.

[10] F. Alizadeh, J.-P. A. Haeberly, M. V. Nayakkankuppam, M. L. Overton, and S. Schmieta, “SDPpack User’s Guide (Version 0.9 BETA for Matlab 5.0),” NYU Computer Science Dept. Technical Report, June 1997, http://www.cs.nyu.edu/faculty/overton/sdppack/sdppack.html.

[11] J. O. Coleman, “Systematic mapping of quadratic constraints on embedded FIR filters to linear matrix inequalities,” in *Proc. 1998 Conference on Information Sciences and Systems (CISS ’98)*, Mar. 1998.

[12] H. Lebret and S. Boyd, “Antenna array pattern synthesis via convex optimization,” *IEEE Trans. on Signal Processing*, vol. 45, no. 3, pp. 526–32, Mar. 1997.

[13] Wu-Sheng Lu, “Design of nonlinear-phase FIR digital filters: A semidefinite programming approach,” in *Proc. 1999 Int’l Symp. Circuits and Systems (ISCAS ’99)*, June 1999.

[14] D. P. Scholnik and J. O. Coleman, “Constrained quadratic design of FIR data-communication filters with linear matrix inequalities,” in *Proc. 1998 Conf. on Information Sciences and Systems (CISS ’98)*, Mar. 1998.

[15] Dan P. Scholnik and J. O. Coleman, “Simple, exact models of sample-interleaving demodulators/modulators for quadrature bandpass sampling/reconstruction,” in *Proc. 1997 Conf. on Information Sciences and Systems (CISS ’97)*, Baltimore, MD, Mar. 1997.

[16] Dan P. Scholnik and J. O. Coleman, “Nonuniformly offset polyphase synthesis of a bandpass signal from complex-envelope samples,” in *Proc. 1999 Int’l Symp. Circuits and Systems (ISCAS ’99)*, June 1999.