Integer-Coefficient FIR Filter Sharpening for Equiripple Stopbands and Maximally Flat Passbands

(Invited Paper)

Jeffrey O. Coleman
Naval Research Laboratory, Washington, DC

***Abstract*—Linear-phase FIR filters with deep, equiripple stopbands are constructed using Chebyshev polynomials to sharpen short linear-phase subfilters. The Chebyshev recursion yields a natural implementation structure using only local interconnections between subfilter copies. Small-integer coefficients often suffice for the subfilters, because of their shallow stopbands, and power-of-two coefficients often suffice for the sharpening structure. Included are several passband-flattening techniques, most of which also use only small-integer coefficients.**

I. Introduction

Can an FIR filter with ultra-simple coefficients, like small integers or powers of two, and realizable with only local interconnect routing have a deep, equiripple stopband?

The filter in Fig. 1 is the causal (7 samples of causality shimming), pipelined (13 more samples of pipeline delays) version of a zero-phase filter with frequency response $G(f) = 2^{-20}T_7(F(f))$, where $F(f) = 2 + 2\cos(2\pi f)$ and $T_7(x) = 64x^7 - 112x^5 + 56x^3 - 7x$ is the Chebyshev polynomial of the first kind of degree seven. For $|x| \gg 1$ the first term dominates so that $T_7(x) \approx 2^6x^7$ and $G(0) = 2^{-20}T_7(4) \approx 2^{-20}2^6(2^2)^7 = 1$. Like its siblings of other degrees, this Chebyshev polynomial ripples with $|T_7(x)| \le 1$ for $|x| \le 1$ so that $G(f)$ ripples with $|G(f)| \le 2^{-20}$ for $|F(f)| \le 1$, which in turn holds when $\frac{1}{3} \le f \le \frac{2}{3}$. At those frequencies then, $G(f)$ has an equiripple stopband down about 120 dB. See Fig. 2.

Classic filter sharpening [1] using rescaled Chebyshev polynomials will allow us below to generalize on this example, most importantly by enlarging the single-point passband. Some of these ideas were previously explored for CIC filters [2]. Here, arbitrary linear-phase subfilters are used instead, signal scaling is improved, and the math is simplified. The realization topology here does not directly apply to CIC filters, however.

II. Rescaling the Chebyshev Polynomials

Chebyshev polynomials of the first kind are defined by

$$T_n(x) \triangleq \begin{cases} 1 & \text{for } n = 0, \\ x & \text{for } n = 1, \\ 2xT_{n-1}(x) - T_{n-2}(x) & \text{for } n > 1. \end{cases} \quad (1)$$

Assume $\alpha > 0$ and for $n = 1, 2, \ldots$ define

$$P_{\alpha,n}(x) \triangleq 2\alpha^{n/2}T_n(\alpha^{-1/2}x/2) \quad (2)$$

so that $P_{\alpha,1}(x) \triangleq x$ and $P_{\alpha,n}(x)$ has a leading term of x^n for $n = 1, 2, \ldots$. Then for $n = 3, 4, \ldots$, recursion (1) yields

This work was supported by the base program at the Naval Research Laboratory, where the work was done but from from which the author will retire before this conference. Author contact info at http://alum.mit.edu/www/jeffc

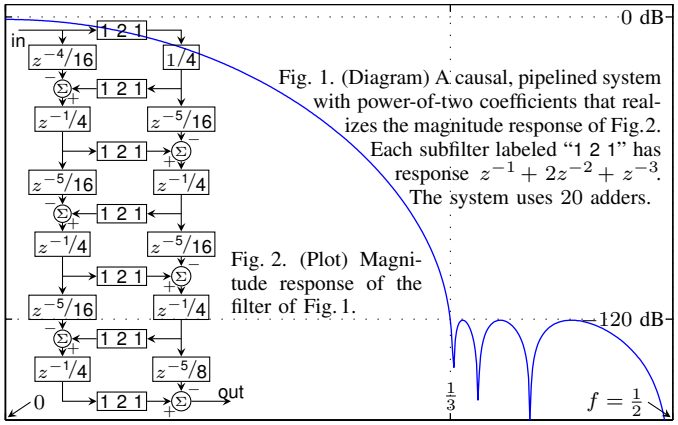

Fig. 1. (Diagram) A causal, pipelined system with power-of-two coefficients that realizes the magnitude response of Fig. 2. Each subfilter labeled "1 2 1" has response $z^{-1} + 2z^{-2} + z^{-3}$. The system uses 20 adders.

Fig. 2. (Plot) Magnitude response of the filter of Fig. 1.

$$\begin{aligned} &P_{\alpha,n}(x) \\ &= 2\alpha^{n/2}\big(\alpha^{-1/2}xT_{n-1}(\alpha^{-1/2}x/2) - T_{n-2}(\alpha^{-1/2}x/2)\big) \\ &= 2\alpha^{(n-1)/2}xT_{n-1}(\alpha^{-1/2}x/2) - 2\alpha^{n/2}T_{n-2}(\alpha^{-1/2}x/2) \\ &= xP_{\alpha,n-1}(x) - \alpha P_{\alpha,n-2}(x). \end{aligned}$$

To include $n = 2$, set $\alpha P_{\alpha,0}(x) = xP_{\alpha,1}(x) - P_{\alpha,2}(x)$. Use (2) on the right to obtain $\alpha P_{\alpha,0}(x) = x^2 - 2\alpha T_2(\alpha^{-1/2}x/2)$. Then use $T_2(x) = 2x^2 - 1$ to obtain $\alpha P_{\alpha,0}(x) = x^2 - 2\alpha(2\alpha^{-1}x^2/4 - 1) = 2\alpha$ and thus an alternative, recursive definition

$$P_{\alpha,n}(x) \triangleq \begin{cases} 2 & \text{for } n = 0, \\ x & \text{for } n = 1, \\ xP_{\alpha,n-1}(x) - \alpha P_{\alpha,n-2}(x) & \text{for } n > 1. \end{cases} \quad (3)$$

This definition will ultimately guide our synthesis of the filters, while both it and original definition (2) will help us understand their frequency-response behavior. That understanding begins with several important properties.

1) Ripple: Since $T_n(x)$ ripples between ± 1 for $|x| \le 1$, definition (2) implies that $P_{\alpha,n}(x)$ ripples between $\pm 2\alpha^{n/2}$ for $|\alpha^{-1/2}x/2| \le 1$ or $|x| \le 2\alpha^{1/2}$.

2) Behavior for small α: Chebyshev polynomial $T_n(x) \approx 2^{n-1}x^n$ for $|x| \gg 1$, so for $x \gg 2\alpha^{1/2}$ definition (2) implies $P_{\alpha,n}(x) \approx x^n$. This implies $P_{\alpha,n}(1) \approx 1$ for small α.

3) Derivatives at unity: Likewise, for $\alpha \ll 1$, the kth derivative at unity $P^{(k)}_{\alpha,n}(1) \approx \frac{n!}{(n-k)!} = n(n-1)\cdots(n-k+1)$.

The last property extended to $k = 0$ subsumes $P_{\alpha,n}(1) \approx 1$ from the second, so the first and last properties are key. They govern stopband and passband behavior respectively.

III. Equiripple-Stopband Filters

The linear programs used in some of the example designs of this section and the next are formulated and carried out in

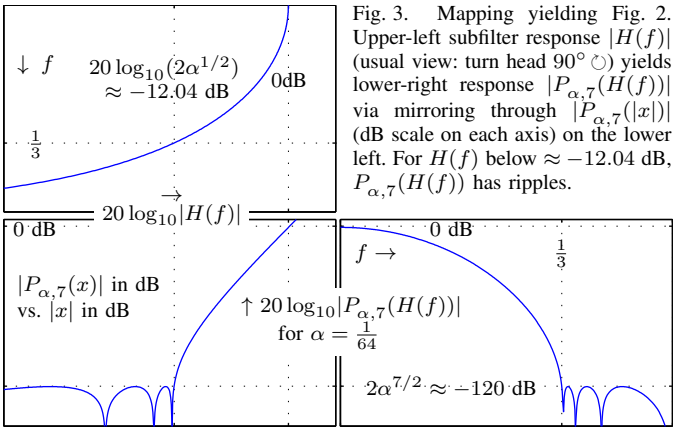

Fig. 3. Mapping yielding Fig. 2. Upper-left subfilter response $|H(f)|$ (usual view: turn head 90° ↻) yields lower-right response $|P_{\alpha,7}(H(f))|$ via mirroring through $|P_{\alpha,7}(|x|)|$ (dB scale on each axis) on the lower left. For $H(f)$ below ≈ -12.04 dB, $P_{\alpha,7}(H(f))$ has ripples.

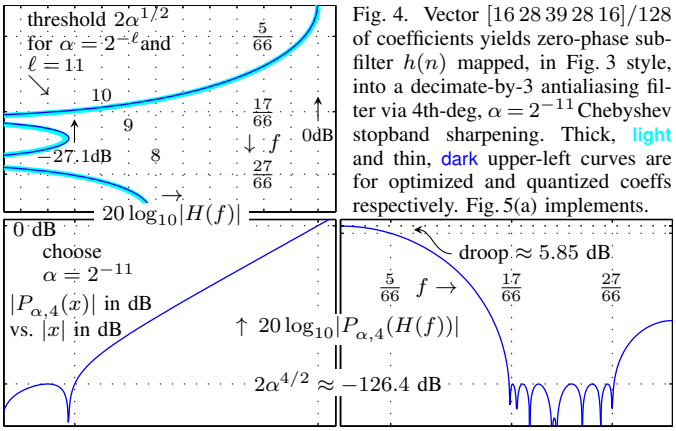

Fig. 4. Vector $[16\,28\,39\,28\,16]/128$ of coefficients yields zero-phase subfilter $h(n)$ mapped, in Fig. 3 style, into a decimate-by-3 antialiasing filter via 4th-deg, $\alpha = 2^{-11}$ Chebyshev stopband sharpening. Thick, light and thin, dark upper-left curves are for optimized and quantized coeffs respectively. Fig. 5(a) implements.

the Opt [3] toolbox of this author and his colleagues.

Properties 1 and 3 above imply that the passbands/stopbands of zero-phase FIR filter $H(f)$ become the passbands/stopbands of derived zero-phase frequency response $P_{\alpha,n}(H(f))$ when $\alpha \ll 1$. Proper design of $H(f)$ gives $P_{\alpha,n}(H(f))$ equiripple stopbands. In the Fig. 3 example, polynomial $P_{\alpha,7}(x)$ constructed with $\alpha = 1/64$ maps $H(f) = \big(1+\cos(2\pi f)\big)/2$, the Fourier transform of impulse response $h(n) = \big(\delta(n+1) + 2\delta(n) + \delta(n-1)\big)/4$, into $P_{\alpha,7}\big(H(f)\big)$, the Fig. 2 response expressed as $2^{-20}T_7(2+2\cos(2\pi f))$ in the introduction.

The $H(f)$ of Fig. 3 was uninspired. Sharpening is more useful with $H(f)$ and α chosen to give $H(f)$ at least a roughly equiripple stopband of height near but not exceeding $2\alpha^{1/2}$.

Let us walk through a design. Suppose we seek a linear-phase antialiasing filter for decimation by 3 with passband to $f = 5/66 \approx 0.076$. This centers the stopband on $f = 22/66$ with width $10/66$. Nulls at $f = 0$ or $f = 1/2$ are unneeded, so use a Type-I linear-phase subfilter. We want subfilter stopband depth ≥ 10 dB, with 15 dB or 20 dB being better, with the shortest subfilter. Impulse-response length 3 allows only one spectral null in $0 \leq f \leq 1/2$, so choose length 5 and two nulls for greater stopband depth and lower-degree sharpening.

To begin, give subfilter $H(f)$ a one-point passband at $f = 0$ (defer passband flattening to Section IV) and minimize stopband magnitude using a linear program or a Parks-McClellan approach. The $|H(f)|$ on the upper left in Fig. 4 results.

Property 1 of Section II makes $2\alpha^{1/2}$ a threshold. Where $|H(f)|$ is below it (esp. if it covers a range of amplitudes) sharpened response $P_{\alpha,n}\big(H(f)\big)$ exhibits stopband ripples. In the upper-left Fig. 4 plot, grid lines show these threshold values in dB for power-of-two choices of α, the choices that will be most convenient in implementation later. To maximize the ultimate sharpened-filter stopband depth per degree of sharpening, we want the smallest α, so we choose the lowest threshold that upper bounds the subfilter stopband response.

We generally lose little by quantizing $h(n)$ so that the original bounding threshold remains the threshold after quantization. In Fig. 4, for example, 2^{-11} is the smallest power-of-two α that has its threshold everywhere above the stopband response of the optimal subfilter. Here 7-bit quantization of the subfilter impulse-response samples is the coarsest that preserves this relationship. (Hand tuning quantized values can sometimes improve matters, but it did not seem so this time.)

Once α has been fixed, we choose n to give sharpened response $P_{\alpha,n}\big(H(f)\big)$ an adequate final stopband depth $2\alpha^{n/2}$ as per Section II's Property 1. Here $n = 2$, 3, 4, and 5 yield depths of -60.2, -93.3, -126.4, and -159.5 dB respectively, and $n = 4$ is chosen for the final design.

The Fig. 5(a) transposed-form implementation of the sharpened filter follows recursion (3). Reversing arrows creates a direct-form structure. Each subfilter $H(f)$ is in zero-phase form: its impulse response is centered on the origin. To replace a zero-phase subfilter with a causal one, cut the system vertically so that the subfilter is cut and insert identical delays wherever branches are cut. Choose the delay amount to make the subfilter causal. Pipeline-register delays can be added in a similar way, such as after each addition node shown. Of course it was these steps that gave the Fig. 1 filter the form shown.

In the lower-right plot of Fig. 4, passband droop is fairly severe at ≈ 5.85 dB. The next section addresses possible ways to reduce passband droop (without resorting to a cascaded compensator filter, which of course is always an option).

IV. Passband Flattening

A. Flatten the Subfilter Passband

The most obvious way to flatten the passband of sharpened filter $P_{\alpha,n}(H(f))$ is simply to flatten the passband of subfilter $H(f)$. This is illustrated in Fig. 6. Relative to the Fig. 4 design, here a slightly longer subfilter is required, and because less subfilter stopband depth is obtained, higher-order sharpening is needed to obtain roughly similar stopband depth in the sharpened filter. These tradeoffs are of course unsurprising.

B. Zero a Derivative of the Sharpening Polynomial

To obtain stronger passband flattening, alter the sharpening polynomial. Let $n < m$ and define new sharpening polynomial

$$P_{\alpha,n,m}(x) \triangleq \frac{1}{m-n}\big(mP_{\alpha,n}(x) - nP_{\alpha,m}(x)\big). \qquad (4)$$

Then use Section II's Properties 2 and 3 to see that at $x = 1$, where passband mapping behavior is determined, $P_{\alpha,n,m}(x) \approx 1$ with derivative $P'_{\alpha,n,m}(x) \approx 0$.

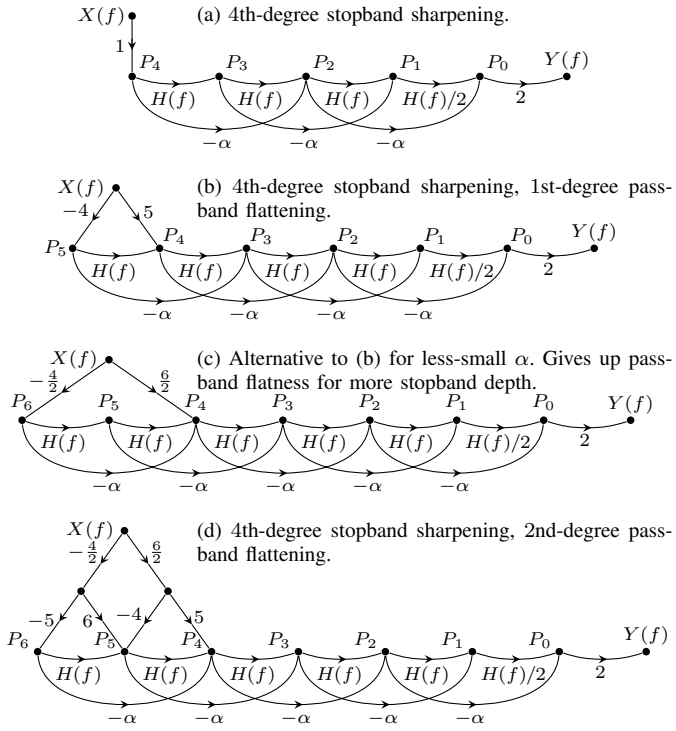

Fig. 5. The transfer function from node P_n to output $Y(f)$ is $P_{\alpha,n}(H(f))$ and has stopband ripple wherever $H(f) \approx 0$. Passband flattening from input $X(f)$ linearly combines two or more of those transfer functions in proportions that zero one or more derivatives of $P_{\alpha,n}(x)$ at $x = 1$. See Figs. 7 and 8.

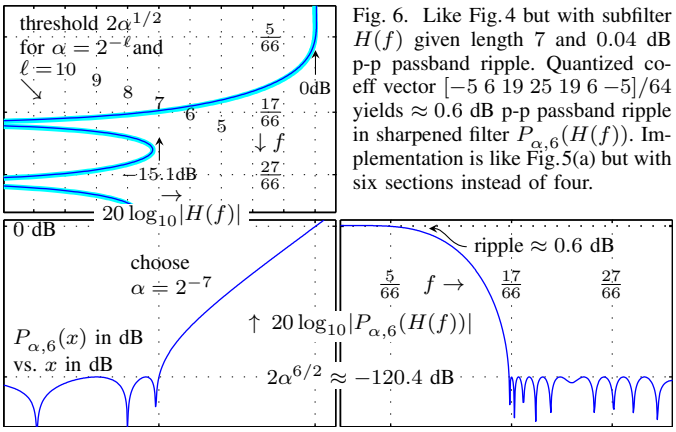

Fig. 6. Like Fig. 4 but with subfilter $H(f)$ given length 7 and 0.04 dB p-p passband ripple. Quantized coeff vector $[-5\ 6\ 19\ 25\ 19\ 6\ -5]/64$ yields ≈ 0.6 dB p-p passband ripple in sharpened filter $P_{\alpha,6}(H(f))$. Implementation is like Fig. 5(a) but with six sections instead of four.

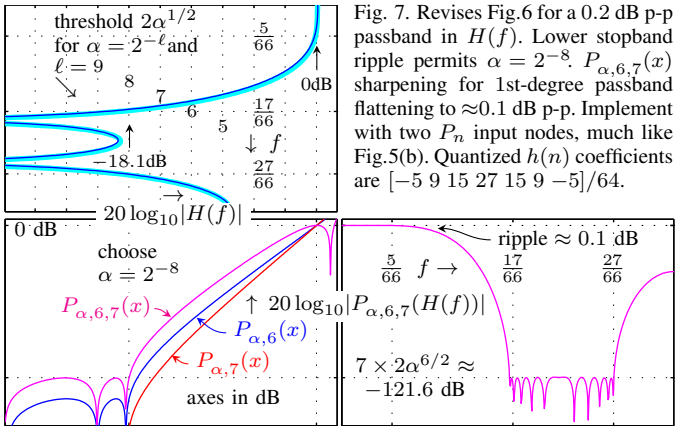

Fig. 7. Revises Fig. 6 for a 0.2 dB p-p passband in $H(f)$. Lower stopband ripple permits $\alpha = 2^{-8}$. $P_{\alpha,6,7}(x)$ sharpening for 1st-degree passband flattening to ≈0.1 dB p-p. Implement with two P_n input nodes, much like Fig. 5(b). Quantized $h(n)$ coefficients are $[-5\ 9\ 15\ 27\ 15\ 9\ -5]/64$.

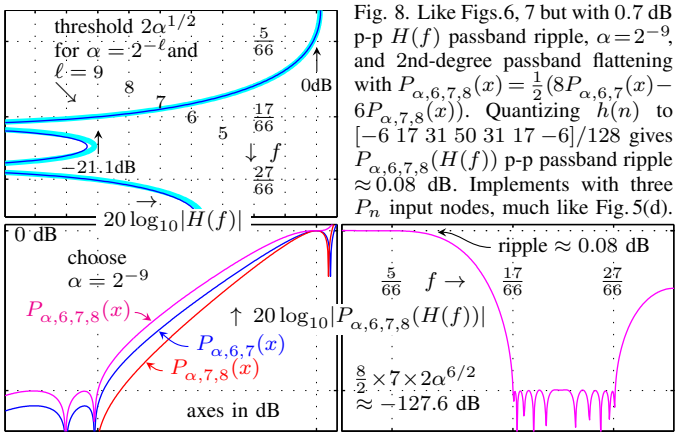

Fig. 8. Like Figs. 6, 7 but with 0.7 dB p-p $H(f)$ passband ripple, $\alpha = 2^{-9}$, and 2nd-degree passband flattening with $P_{\alpha,6,7,8}(x) = \frac{1}{2}(8P_{\alpha,6,7}(x) - 6P_{\alpha,7,8}(x))$. Quantizing $h(n)$ to $[-6\ 17\ 31\ 50\ 31\ 17\ -6]/128$ gives $P_{\alpha,6,7,8}(H(f))$ p-p passband ripple ≈ 0.08 dB. Implements with three P_n input nodes, much like Fig. 5(d).

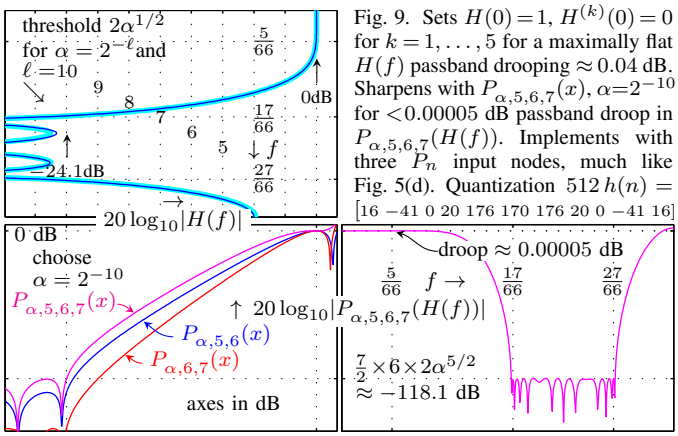

Fig. 9. Sets $H(0) = 1$, $H^{(k)}(0) = 0$ for $k = 1, \ldots, 5$ for a maximally flat $H(f)$ passband drooping ≈ 0.04 dB. Sharpens with $P_{\alpha,5,6,7}(x)$, $\alpha = 2^{-10}$ for < 0.00005 dB passband droop in $P_{\alpha,5,6,7}(H(f))$. Implements with three P_n input nodes, much like Fig. 5(d). Quantization $512\,h(n) = [16\ -41\ 0\ 20\ 176\ 170\ 176\ 20\ 0\ -41\ 16]$

For $|x| \ll 1$ the behavior of $P_{\alpha,n,m}(x)$ is roughly that of its dominant first term from definition (4), so that term determines stopband behavior. When α is very small, the usual case $P_{\alpha,n,n+1}(x)$ can be used, and implementation is as in Fig. 5(b). A design example appears in Fig. 7.

If α is too large for the suppression of the second term to ensure an adequate approximation to an equiripple stopband, $P_{\alpha,n,n+2}(x)$ can be used instead, as pictured in Fig. 5(c). (Space restrictions preclude showing a design example.)

C. Zero Multiple Derivatives

Even better passband performance or an improved passband/stopband tradeoff can be obtained by sharpening with a polynomial with two zero derivatives at $x = 1$. Let $n < m < \ell$ and define new sharpening polynomial

$$P_{\alpha,n,m,\ell}(x) \triangleq \frac{1}{\ell - n}\big(\ell P_{\alpha,n,m}(x) - n P_{\alpha,m,\ell}(x)\big).$$

Equation (4) yields $P''_{\alpha,n,m}(1) \approx -nm$ and so the last of these:

$$P_{\alpha,n,m,\ell}(1) \approx \frac{1}{\ell - n}\big(\ell - n\big) \approx 1$$

$$P'_{\alpha,n,m,\ell}(1) = \frac{1}{\ell - n}\big(\ell P'_{\alpha,n,m}(1) - n P'_{\alpha,m,\ell}(1)\big) \approx 0$$

$$P''_{\alpha,n,m,\ell}(1) \approx \frac{-\ell nm + nm\ell}{\ell - n} = 0.$$

Unit increments between n, m, and ℓ are the norm when α is sufficiently small. The design example in Fig. 8 illustrates.

This approach extends easily to zero even more derivatives.

D. Maximally Flat Passband

Instead of giving zero-phase subfilter $H(f)$ passband ripples, we can simply zero some number of its $f=0$ derivatives. We zero that sample of the kth derivative by zeroing kth moment $\sum_n h(n)n^k$, taking care of course not to confuse time index n with our use elsewhere of n as a polynomial degree. We are constructing $H(f)$ to be real and even, so this is automatic for odd k. We can focus on even k.

Here's one approach. First, construct some short real, even impulse responses—here it is done by hand—with small-integer coefficients and some number of zero even moments. These impulse responses are zero beyond the n values shown:

n	$=-5$	-4	-3	-2	-1	0	1	2	3	4	5
$h_0(n)=$						1					
$h_1(n)=$			1	-6	15	0	15	-6	1		
$h_2(n)=$		1	-4	4	4	0	4	4	-4	1	
$h_3(n)=$	1	-2	-3	8	2	0	2	8	-3	-2	1

Each has zero 2nd and 4th moments and therefore a zero kth moment for $k\in\{1,2,3,4,5\}$. The same is therefore true of $h(n) \triangleq x_0h_0(n) + x_1h_1(n) + x_2h_2(n) + x_3h_3(n)$. We can choose the x_i to be small integers that give lowpass-filter $h(n)$ the stopband behavior we need, and the passband will be relatively flat out to some bandwidth. How far? Let us experiment using the same stopband specifications as before.

Solving a linear program to optimize x_i to minimize δ subject to $1 \le H(0)$ and $|H(f)| \le \delta$ for each f in a grid of 51 frequencies spaced across the stopband yields $[x_0,\ldots,x_3] \approx [0.3315, 0.0236, -0.0178, 0.0313]$. Quantizing to $[x_0,\ldots,x_3]=[170,12,-9,16]/512$ yields the Fig. 9 design.

E. Optimized (Noninteger Coefficient) Sharpening Polynomial

The final design example, shown in Fig. 10, both revisits the last technique and introduces another. The last subfilter design approach is used but with fewer and shorter basis sequences:

n	$=-4$	-3	-2	-1	0	1	2	3	4
$h_0(n)=$					1				
$h_1(n)=$			-1	4	0	4	-1		
$h_2(n)=$		-1	2	1	0	1	2	-1	

Each has a zero 2nd moment and so, by even symmetry, zero kth moments for $k\in\{1,2,3\}$. Variables x_1, x_2, and x_3 in $h(n)=x_0h(1)+x_1h(1)+x_2h(2)$ were first optimized to minimize a stopband bound subject to requirement $H(0)=1$ and then rounded to $[x_0,x_1,x_2]=[14\ 1\ 3]/32$. Setting $\alpha=2^{-7}$ and optimizing a_7, a_8, and a_9 to minimize passband ripple of sharpening polynomial $P_\alpha(x) \triangleq a_7P_{\alpha,7}(x) - a_8P_{\alpha,8}(x) + a_9P_{\alpha,9}(x)$ subject to a -111 dB stopband bound yielded a sharpened filter with some 37% of the dB passband ripple obtained with reference polynomial $P_{\alpha,7,8,9}$ (and 0.4 dB better stopband ripple as well). This modest improvement relative to that reference comes at the cost of giving up integer coefficients for optimal ones requiring precision implementation. The optimal approach is unlikely to be worth such a sacrifice.

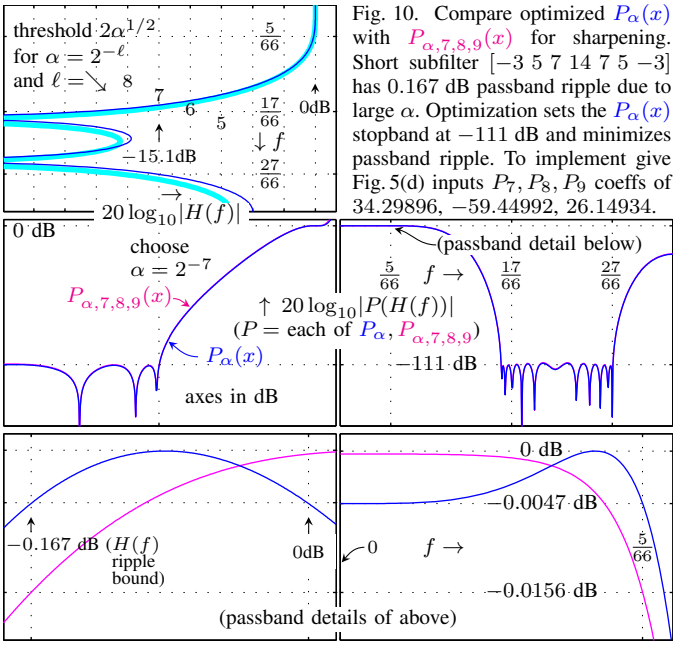

Fig. 10. Compare optimized $P_\alpha(x)$ with $P_{\alpha,7,8,9}(x)$ for sharpening. Short subfilter $[-3\ 5\ 7\ 14\ 7\ 5\ -3]$ has 0.167 dB passband ripple due to large α. Optimization sets the $P_\alpha(x)$ stopband at -111 dB and minimizes passband ripple. To implement give Fig. 5(d) inputs P_7, P_8, P_9 coeffs of 34.29896, -59.44992, 26.14934.

V. Conclusions

Copies of a short linear-phase FIR subfilter with simple coefficients can be arranged using the Fig. 5 topologies to obtain linear-phase FIR filters with deep, equiripple stopbands using only power-of-two feedforward coefficients α and, if desired, extremely flat passbands as well. The key strength of the approach is unquestionably robustness to subfilter-coefficient quantization. Some design examples here use five-bit coefficients when quantization of actual impulse-response samples would require roughly 20 bits.

This quick study leaves many questions open. How should design parameters be traded off? Could better subfilter design approaches simplify coefficients further? How does implementation complexity compare to standard approaches, both ordinary and multiplierless? Is this approach's locality of connection useful, either in VLSI or in FPGA? What about IIR subfilters? What about sharpened filters as subfilters?

References

[1] J. Kaiser and R. Hamming, "Sharpening the response of a symmetric nonrecursive filter by multiple use of the same filter," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 25, no. 5, pp. 415–422, Oct. 1977.

[2] J. O. Coleman, "Chebyshev stopbands for CIC Decimation Filters and CIC-implemented array tapers in 1D and 2D," *IEEE Trans. Circuits Syst. I*, vol. 59, no. 12, pp. 2956–2968, Dec. 2012.

[3] J. Coleman, D. Scholnik, and J. Brandriss, "A specification language for the optimal design of exotic FIR filters with second-order cone programs," in *Proc. IEEE Asilomar Conf. on Signals, Systems and Computers*, vol. 1, Nov. 2002, pp. 341–345, the IEEEXplore file is bad, so find a preprint.

[4] Y. Voronenko and M. Püschel, "Multiplierless multiple constant multiplication," *ACM Trans. Algorithms*, vol. 3, May 2007, http://doi.acm.org/10.1145/1240233.1240234.