

Turbocharging Interpolated FIR Filters

Interpolated finite impulse response (IFIR) filters drastically improve the computational efficiency of low-pass finite impulse response (FIR) filters designed using the traditional Parks-McClellan design method [1]. This article presents a trick to further improve the computational efficiency of traditional IFIR filters.

TRADITIONAL IFIR FILTERS

Recall that a traditional IFIR filter comprises a band-edge shaping subfilter followed by a low-pass masking subfilter, where both subfilters are traditionally implemented as linear-phase tapped-delay FIR filters. The band-edge shaping subfilter has a sparse impulse response, with all but every M th sample being zero, that shapes the final IFIR filter's passband, transition band, and stopband responses. (Integer M is called the "expansion" or "stretch" factor of the band-edge shaping subfilter.) Because the band-edge shaping subfilter's frequency response contains $M-1$ unwanted periodic passband images, the masking subfilter is used to attenuate those images and can be implemented with few arithmetic operations.

Although design curves are available for estimating the optimum value for expansion factor M which minimizes filter computational workload [1], the fol-

"DSP Tips and Tricks" introduces practical design and implementation signal processing algorithms that you may wish to incorporate into your designs. We welcome readers to submit their contributions to Associate Editors Rick Lyons (r.lyons@ieee.org) or Britt Rorabaugh (dspboss@aol.com).

lowing expression enables the optimum M value to be computed directly [2]

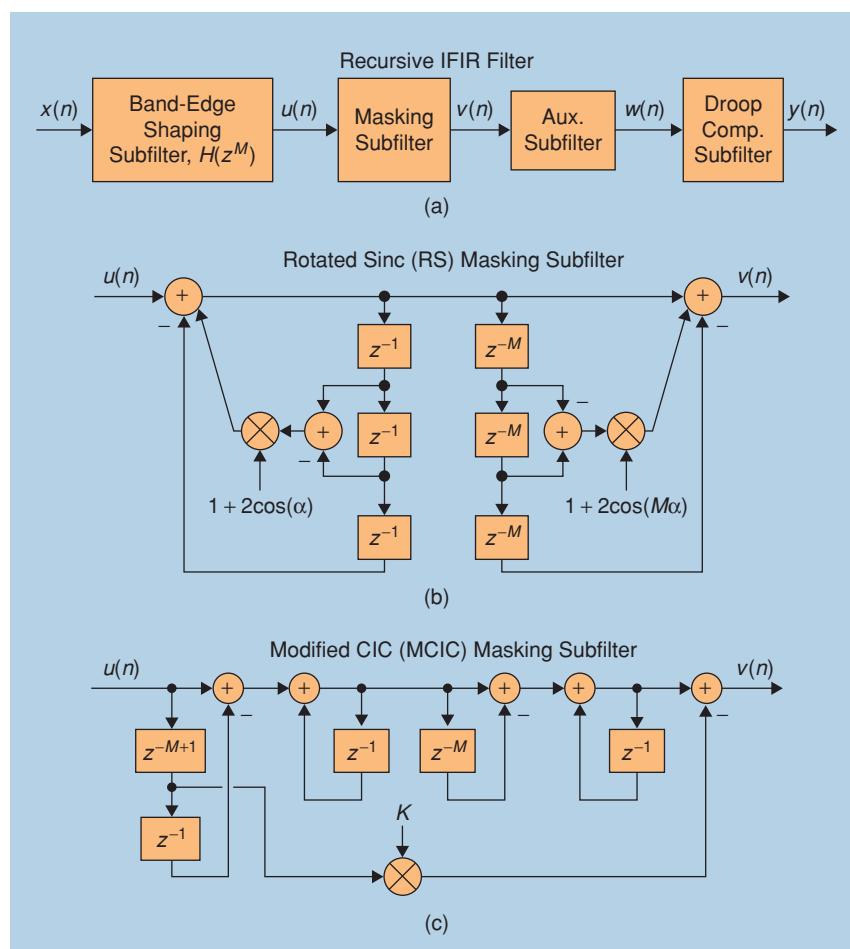
$$M_{\text{opt}} = \frac{1}{f_{\text{pass}} + f_{\text{stop}} + \sqrt{f_{\text{stop}} - f_{\text{pass}}}}. \quad (1)$$

The f_{pass} and f_{stop} frequency values are normalized to the filter input sample rate, f_s , in Hz. For example, $f_{\text{pass}} = 0.1$ is equivalent to a cyclic frequency of $f_{\text{pass}} = f_s/10$ Hz. The value of M_{opt} ,

computed using (1), is rounded to the nearest integer.

RECURSIVE IFIR FILTERS

Our IFIR filter optimization trick that improves the filter's computational efficiency replaces the traditional single-stage masking subfilter with a cascade of subfilters, as shown in Figure 1(a). We present two detailed structures of the masking subfilter, a rotated sinc (RS)-masking subfilter and a "modified"



[FIG1] Recursive IFIR filter: (a) subfilters, (b) RS-masking subfilter structure, (c) MCIC-masking subfilter.

cascaded integrator-comb (MCIC)-masking subfilter, which are shown in Figure 1(b)–(c), respectively.

THE RS-MASKING SUBFILTER

The RS-masking subfilter is called a “rotated sinc” filter and was originally proposed for use with sigma-delta analog-to-digital converters [3]. Factor α is the angular positions (in radians) of the subfilter’s poles/zeros on the z -plane near $z = 1$ and M is the band-edge shaping subfilter’s integer expansion factor. The z -domain transfer function of the RS-masking subfilter is given by

$$H_{RS}(z) = \frac{1 - Az^{-M} + Az^{-2M} - z^{-3M}}{1 - Bz^{-1} + Bz^{-2} - z^{-3}}, \quad (2)$$

where $A = 1 + 2\cos(M\alpha)$ and $B = 1 + 2\cos(\alpha)$. There are multiple ways to implement (2); however, the structure in Figure 1(b) is the most computationally efficient.

An RS-masking subfilter has M triplets of zeros (three closely spaced zeros) even-

ly distributed around the z -plane’s unit circle. The triplets of zeros create periodic response stopbands centered exactly at the frequencies where we require band-edge shaping subfilter passband image attenuation. The triplet of zeros near $z = 1$ are overlayed with three poles where pole-zero cancellation makes our masking subfilter a low-pass filter. The expansion factor M in (2) determines how many RS-masking subfilter response stopbands (triplets of zeros) are distributed around the unit circle, and angle α determines the width of those stopbands. For proper operation, angle α must be less than π/M . As it turns out, thankfully, the RS-masking subfilter’s impulse response is both finite in duration and symmetrical, so the subfilter exhibits linear phase.

IMPROVED MASKING ATTENUATION WITH THE RS-MASKING FILTER

Because the RS-masking subfilter may not provide sufficient stopband attenuation, we can improve that attenuation, at a minimal computational cost, by applying the masking subfilter’s $v(n)$ output to

an auxiliary subfilter as shown in Figure 1(a). There are two candidate auxiliary subfilters that should be considered.

■ Auxiliary subfilter 1 defined by

$$H_{Aux1}(z) = \frac{1 - z^{-M}}{1 - z^{-1}} \quad (3)$$

can achieve an additional 15–25 dB of stopband attenuation. When two’s complement fixed-point arithmetic is used, this subfilter is a guaranteed-stable cascaded integrator-comb (CIC) low-pass filter that places a single additional z -domain zero at the center of each RS-masking subfilter stopband.

■ Auxiliary subfilter 2 defined by

$$H_{Aux2}(z) = \frac{1 - 2\cos(M\alpha/2)z^{-M} + z^{-2M}}{1 - 2\cos(\alpha/2)z^{-1} + z^{-2}} \quad (4)$$

can achieve a masking stopband attenuation greater than 90 dB. This low-pass subfilter places a pair of z -domain zeros at the center of each RS-masking subfilter stopband.

The RS-masking and auxiliary subfilters have some droop in their passband magnitude responses. If their cascaded passband droop is intolerable in your application, then some sort of passband droop compensation must be employed as indicated in Figure 1(a). We suggest using a

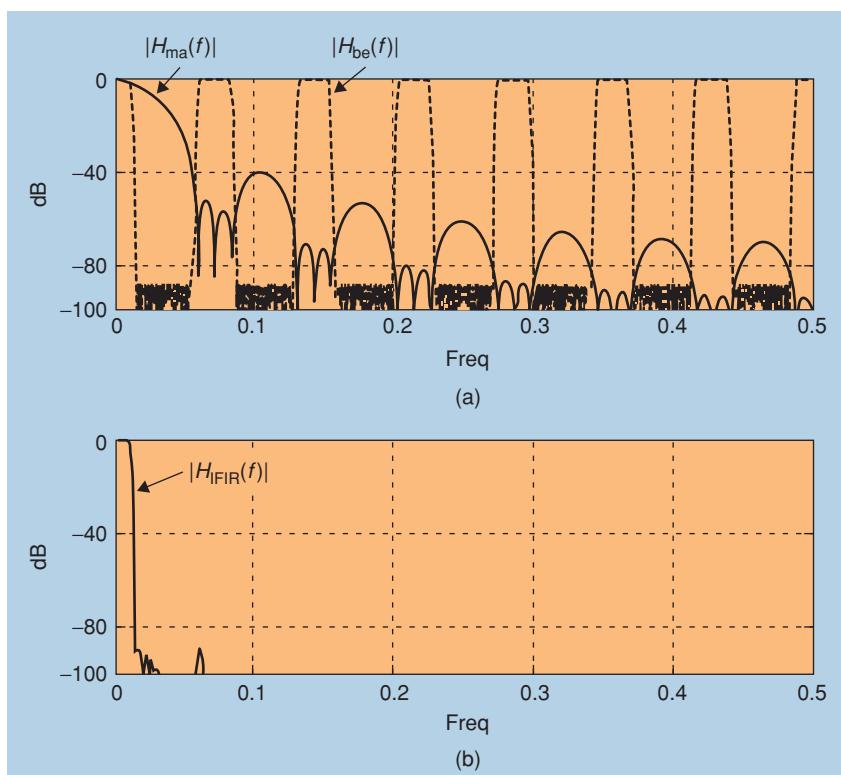
■ Compensation subfilter defined by

$$H_{CS}(z) = 1 - Cz^{-M} + z^{-2M}, \quad (5)$$

which was originally proposed for use with high-order cascaded integrator comb (CIC) low-pass filters [4]. That compensation subfilter has a monotonically rising magnitude response beginning at 0 Hz—just what we need for passband droop compensation. The coefficient C in (5) is typically in the range of 4–10 and is determined empirically.

THE MCIC-MASKING SUBFILTER

An alternate structure of the masking subfilter in Figure 1(a) is the “modified CIC” (MCIC) subfilter shown in Figure 1(c). The MCIC subfilter was inspired by, and it is an optimized version of, a CIC filter proposed in [5].



[FIG2] Recursive IFIR filter performance: (a) band-edge shaping and RS-masking subfilter responses, (b) response using cascaded auxiliary subfilter 2 and compensation subfilters.

The MCIC-masking subfilter is a guaranteed-stable linear-phase low-pass filter whose frequency response (like that of the RS-masking subfilter) has periodic stopbands located at the center of each band-edge shaping subfilter passband image. The coefficient K controls the width of the MCIC-masking subfilter stopbands. While the MCIC-masking subfilter requires one fewer multiplier and fewer delay elements than the RS-masking subfilter, sadly the number of delay elements of an MCIC-masking subfilter is not reduced in decimation applications as with an RS-masking subfilter, as we will discuss later.

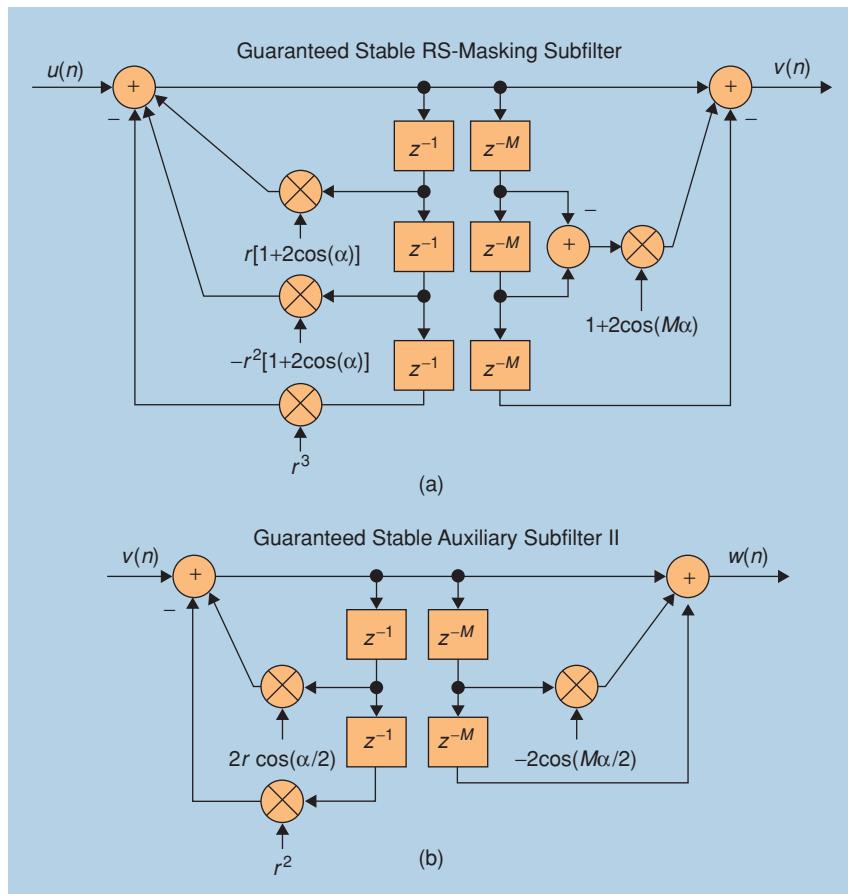
RECURSIVE IFIR FILTER DESIGN

To summarize, designing a recursive IFIR filter using an RS-masking subfilter comprises the following steps:

- 1) Based on the desired low-pass filter's f_{pass} and f_{stop} frequencies, use (1) to determine a preliminary value for the band-edge shaping filter's integer expansion factor M .
- 2) Choose an initial value for the RS-masking subfilter's α using $\alpha = 2\pi f_{\text{pass}}$. Adjust α to maximize the attenuation of the band-edge shaping subfilter's passband images.
- 3) If the RS-masking subfilter does not provide sufficient passband image attenuation, employ one of the auxiliary filters defined by (3) and (4).
- 4) Choose an initial value for C (starting with $4 < C < 10$) for the compensation subfilter. Adjust C to provide the desired passband flatness.
- 5) Continue by increasing M by one (larger values of M yield lower-order band-edge shaping subfilters) and repeat steps 2 through 4 until either the RS-masking/auxiliary subfilter combination no longer supplies sufficient passband image attenuation or the compensation subfilter no longer can achieve acceptable passband flatness.
- 6) Sit back and enjoy a job well done.

RECURSIVE IFIR FILTER DESIGN EXAMPLE

We can further understand the recursive IFIR filter design by considering the scenario where we desire a narrowband low-pass IFIR filter with $f_{\text{pass}} = 0.01fs$, a



[FIG3] Guaranteeing stability: (a) stable RS-masking subfilter and (b) stable auxiliary subfilter 2.

peak-peak passband ripple of 0.2 dB, a transition region bandwidth of $f_{\text{trans}} = 0.005fs$, and 90 dB of stopband attenuation. A traditional Parks-McClellan-designed FIR low-pass filter satisfying these very demanding design specifications would require a 728-tap filter. Designing a recursive IFIR filter to meet our design requirements yields an $M = 14$ band-edge shaping subfilter having 52 taps and an RS-masking subfilter with $\alpha = \pi/42$ rad. When used in place of the RS-masking subfilter, a $K = 6$ MCIC-masking subfilter also meets the stringent requirements of this design example.

The dashed curve in Figure 2(a) shows the design example band-edge shaping subfilter's $|H_{\text{be}}(f)|$ frequency magnitude response with its periodically spaced image passbands which must be removed by the RS-masking subfilter. Selecting $\alpha = \pi/42$ for the RS-masking subfilter yields the $|H_{\text{ma}}(f)|$ magnitude response shown by the solid curve. The magnitude

response of our final recursive IFIR filter, using an auxiliary subfilter 2 and a $C = 6$ compensation subfilter, is shown in Figure 2(b). The final filter's passband peak-peak ripple is less than 0.2 dB.

For comparison purposes, Table 1 lists the design example computational workload, per filter output sample, for the various filter implementation options discussed above. The “PM FIR” table entry means a single Parks-McClellan-designed, tapped-delay line FIR filter. The “two-stage-masking IFIR” table entry refers to the implementation of a traditional IFIR filter's masking subfilter as a separate IFIR filter in order to reduce the overall computations workload, as proposed in [2]. The “efficiency gain” column indicates the percent reduction in additions plus multiplications with respect to a standard IFIR filter. As these results show, for cutting the computational workload of traditional IFIR filters, the recursive IFIR filter is indeed a sharp knife.

[TABLE 1] RECURSIVE IFIR FILTER DESIGN EXAMPLE COMPUTATIONAL WORKLOAD.

LOW-PASS FILTER	ADDS	MULTS	EFFICIENCY GAIN
PM FIR [ORDER = 737]	727	728	—
STANDARD IFIR [$M = 10$]	129	130	—
TWO-STAGE-MASKING IFIR	95	96	26%
RECURSIVE, RS-MASKING IFIR [$M = 14$]	63	57	54%

RECURSIVE IFIR FILTER IMPLEMENTATION

The gain of a recursive IFIR filter can be very large (in the hundreds or thousands), particularly for large M and small α , depending on which auxiliary subfilter is used. As such, the recursive IFIR filter scheme is best suited for floating-point numerical implementations. Two options exist that may enable a fixed-point recursive IFIR filter implementation: 1) when filter gain scaling methods are employed and 2) swapping the internal feedback and feed-forward sections of a subfilter to minimize data word growth.

The RS-masking subfilter, and the auxiliary subfilter 1, have poles lying on the z -plane's unit circle. Such filters run the risk of being unstable should our finite-precision filter coefficients cause a pole to lie just slightly outside the unit circle. If we find that our quantized-coefficient filter implementation is unstable, at the expense of a few additional multiplications we can use the subfilters shown in Figure 3 to guarantee stability while maintaining linear phase. The stability factor r is a constant whose value is as close to, but less than, one as our number format allows. The MCIC-masking subfilter is always stable as long as two's complement fixed-point arithmetic is used.

RECURSIVE IFIR FILTERS FOR DECIMATION

If our low-pass filtering application requires the $y(n)$ output to be decimated, fortunately the RS-masking subfilter lends itself well to such a sample rate change process. To decimate $y(n)$ by M , we merely rearrange the order of the subfilters' elements so that all feed-forward paths and the band-edge shaping subfilter follow the downsample-by- M process. Doing so has two advantages: First, the zero-valued coefficients in the band-edge shaping subfilter are eliminated, reducing that subfilter's order by a factor of M .

Second, the z^{-M} delay lines in the other subfilters become z^{-1} unit-delays, which reduce signal data storage requirements.

In this decimation scenario, the multiplies by the $2 \cos(\alpha)$ and $2 \cos(\alpha/2)$ feedback coefficients must be performed at the high input sample rate. The replacement of those coefficients with $2 - 2^k$ to implement a multiply with simple high-speed binary shifts and a subtraction is discussed in [3].

ACKNOWLEDGEMENT

Appreciation is extended to Ricardo Losada for his review of, and suggestions regarding, this material.

AUTHOR

Richard Lyons (R.Lyons@ieee.org) is a consulting systems engineer and lecturer with Besser Associates in Mountain View, California. He is the author of *Understanding Digital Signal Processing* 2/E (Prentice-Hall, 2004).

REFERENCES

- [1] R. Lyons, "Interpolated narrowband lowpass FIR filters," *IEEE Signal Processing Mag.*, vol. 20, no. 1, pp. 50–57, Jan. 2003.
- [2] A. Mehrnia and A. Willson Jr., "On optimal IFIR filter design," in *Proc. 2004 Int. Symp. Circuits and Systems (ISCAS)*, May 23–26, 2004, vol. 3, pp. 133–136.
- [3] L. Lo Presti, "Efficient modified-sinc filters for sigma-delta A/D converters," *IEEE Trans. Circuits Syst. II*, vol. 47, no. 11, pp. 1204–1213, Nov. 2000.
- [4] H. Oh, S. Kim, G. Choi, and Y. Lee, "On the use of interpolated second-order polynomials for efficient filter design in programmable downconversion," *IEEE J. Selected Areas Commun.*, vol. 17, no. 4, pp. 551–560, Apr. 1999.
- [5] T. Saramäki, Y. Neuvo, and S. Mitra, "Design of computationally efficient interpolated FIR filters," *IEEE Trans. Circuits Syst.*, vol. 35, no. 1, pp. 70–88, Jan. 1988.



MOVING?
You don't want to miss
any issue of this magazine!

CHANGE YOUR ADDRESS

BY PHONE: +1 800 678 IEEE (4333)
in the U.S.A. or
+1 732 981 0060 outside the U.S.A.

BY FAX: +1 732 562 5445

BY E-MAIL: address-change@ieee.org

ONLINE: www.ieee.org,
click on quick links,
change contact info

**Be sure to have
your member number
available.**

©2000 IMAGE 100 LTD