

Improved Narrowband Low-Pass IIR Filters in Fixed-Point Systems

"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 the Associate Editors, Rick Lyons (r.lyons@ieee.org) or Britt Rorabaugh (dspboss@aol.com).

Due to their resistance to quantized-coefficient errors, traditional second-order infinite impulse response (IIR) filters are the fundamental building blocks in computationally efficient high-order IIR digital filter implementations. However, when used in fixed-point number systems, the inherent properties of quantized-coefficient second-order IIR filters do not readily permit their use in narrowband low-pass filtering applications. Narrowband low-pass IIR filters have traditionally had a bad reputation; for example, MATLAB's Signal Processing Toolbox documentation warns: "All classical IIR low-pass filters are ill conditioned for extremely low cutoff frequencies."

This article presents a neat trick that can be used to overcome the shortcomings of narrowband second-order low-pass IIR filters, with no increase in filter coefficient bit widths and no increase in the number of filter multiplies per output sample.

MOTIVATION

Narrowband low-pass IIR filters are difficult to implement because of intrinsic limitations on their z -plane pole locations. Let's quickly review the

restrictions on the z -plane pole locations of a standard second-order IIR filter whose structure is shown in Figure 1(a).

Such an IIR filter, having a transfer function given by

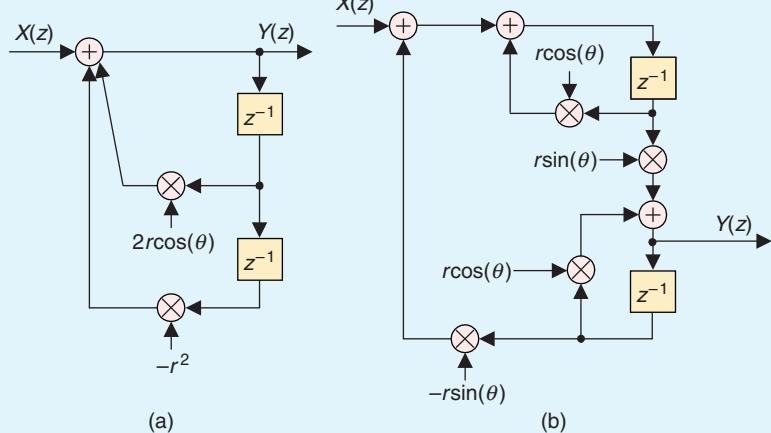
$$H(z) = \frac{1}{1 - 2r\cos(\theta)z^{-1} + r^2z^{-2}} = \frac{1}{(1 - re^{j\theta}z^{-1})(1 - re^{-j\theta}z^{-1})} \quad (1)$$

has a pair of conjugate poles located at radii of r , at angles of $\pm\theta$ radians. In fixed-point implementations, quantizing the $2r\cos(\theta)$ and r^2 coefficients restricts the possible pole locations [1], [2]. On the z -plane, a pole can only reside at the intersection of a vertical line defined by the quantized value of $2r\cos(\theta)$ and a concentric circle whose radius is defined by the square root of the quantized value of r^2 . For example, Figure 2 shows the first quadrant of possible z -plane pole locations when five magnitude bits are used to represent the filter's two coefficients. Notice the irregular spacing of

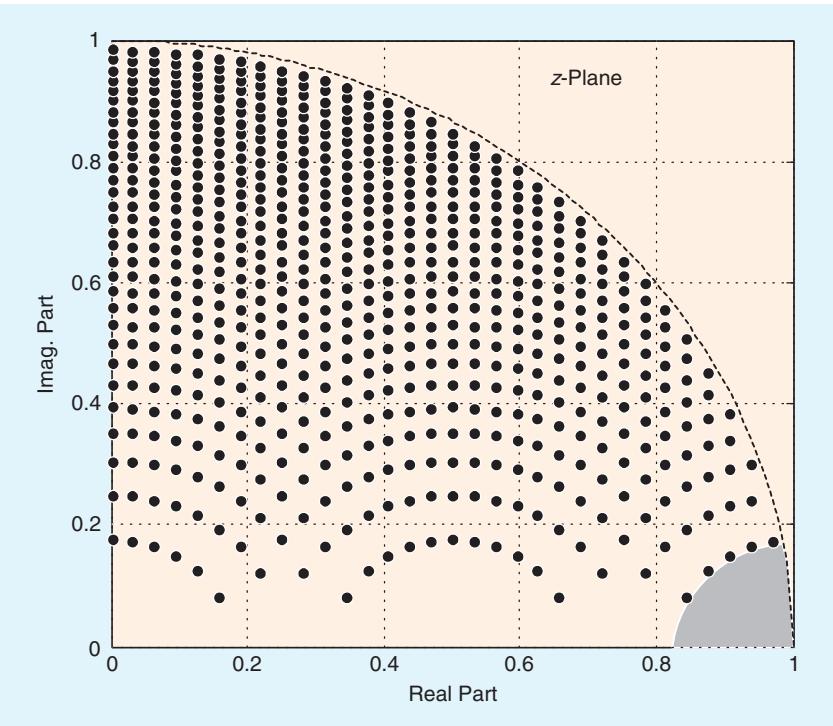
those permissible pole locations. (Due to trigonometric symmetry, the pole locations in the other three quadrants of the z -plane are mirror images of those shown in Figure 2.)

THE PROBLEM

So here's the problem: if we use floating-point software to design a very narrowband low-pass IIR filter (implemented as cascaded second-order filters) having poles residing in the shaded area near $z = 1$, subsequent quantizing of the designed filter coefficients to five magnitude bits will make the poles shift to one of the locations shown by the dots on the border of the shaded region in Figure 2. Unfortunately, that pole shifting, inherent in the Figure 1(a) IIR filter implementation due to coefficient quantization, prevents us from realizing the desired narrowband low-pass filter. We can always reduce the size of the shaded forbidden zone near $z = 1$ in Figure 2 by increasing the number of bits used to represent the second-order



[FIG1] Second-order IIR filters: (a) standard form; (b) coupled form.



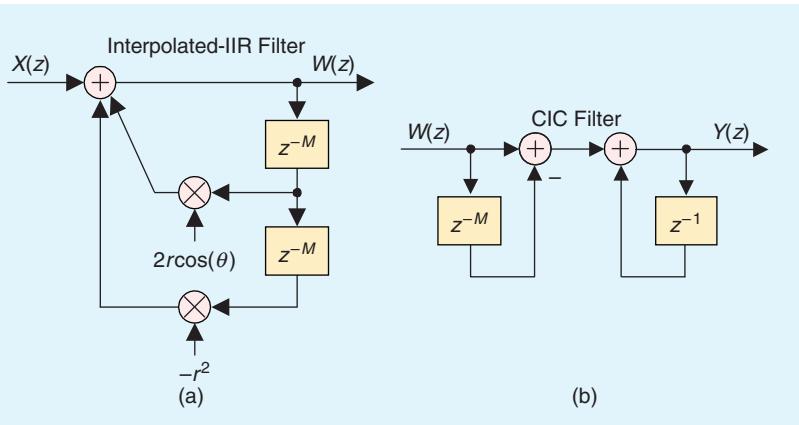
[FIG2] Possible pole locations for five-magnitude-bit coefficient quantization.

filters' coefficients. In some filter implementation scenarios, however, increasing coefficient word bit widths may not be a viable option.

One solution to the above problem is to use a so-called coupled-form IIR filter (also called the Gold-Rader filter [3]) structure, shown in Figure 1(b), having a transfer function given by

$$H_{cf}(z) = \frac{r \sin(\theta) z^{-1}}{1 - 2r \cos(\theta) z^{-1} + r^2 z^{-2}}. \quad (2)$$

Because the coupled-form filter's quantized coefficients in Figure 1(b) are linear in $r \cos(\theta)$ and $r \sin(\theta)$, its possible pole locations are on a regularly spaced grid on the z -plane defined by $z = r \cos(\theta) + j r \sin(\theta)$. This enables us to build second-order narrowband low-pass IIR filters with poles in the desired shaded region of Figure 2. Unfortunately, the coupled-form filter requires twice the number of multiplies needed by the standard second-order IIR filter in Figure 1(a).



[FIG3] Filters: (a) single-section interpolated-IIR filter; (b) first-order CIC filter.

The remainder of this article describes a novel narrowband low-pass IIR filter structure, proposed by Harris and Loudermilk, having poles residing in the shaded region of Figure 2 with no additional multiplication operations beyond those needed for a standard second-order IIR filter [4].

AN IMPROVED NARROWBAND LOW-PASS IIR FILTER

The improved low-pass IIR filter is created by replacing each unit-delay element in a standard second-order IIR filter with multiple unit-delay elements, as shown in Figure 3(a). The astute reader will notice that this scheme is reminiscent of interpolated finite impulse response (IFIR) filters, where each delay element in an FIR filter's tapped-delay line is replaced by multiple delay elements [5]. For this reason we will call the filter in Figure 3(a) an interpolated-IIR filter. The M -length delay lines, where M is a positive integer, in the interpolated-IIR filter shift a standard IIR filter's conjugate poles, originally located at $z = r e^{\pm j\theta}$ (in the vicinity of $z = 1$), to the new locations of

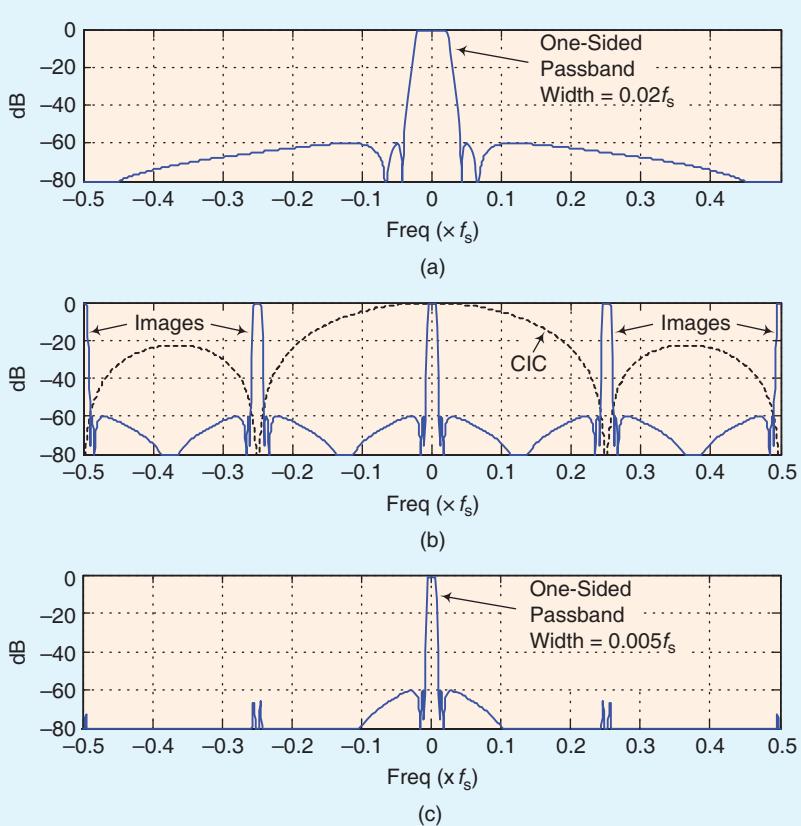
$$z_{\text{new}} = \sqrt[M]{r} \cdot e^{\pm j\theta/M}. \quad (3)$$

That is, the new conjugate pole locations are at radii of the M th root of r , at angles of $\pm\theta/M$ radians. Those interpolated-IIR filter pole locations reside in the desired shaded region of Figure 2 without using more bits to represent the original IIR filter's coefficients.

If the original Figure 1(a) second-order IIR filter contains feedforward coefficients, those coefficients are also delayed by M -length delay lines.

INTERPOLATED-IIR FILTER EXAMPLE

Here we show an example of an interpolated-IIR filter in action. With f_s representing a filter's input signal sample rate in hertz, assume we wish to implement a recursive low-pass filter whose one-sided passband width is $0.005f_s$ with a stopband attenuation greater than 60 dB. If we choose to set $M = 4$, then we start our interpolated-IIR filter design by first designing a



[FIG4] Frequency magnitude responses: (a) original IIR prototype filter; (b) zero-packed interpolated-IIR filter and CIC filters (dashed); (c) final narrowband 12-b coefficient filter.

standard IIR filter having a one-sided passband width of $M \cdot 0.005 f_s = 0.02 f_s$. Using our favorite IIR filter design software (for an elliptic IIR filter, in this example), we obtain a fifth-order prototype IIR filter. Partitioning that fifth-order prototype IIR filter into two second-order and one single-order IIR filter sections, all in cascade and having coefficients represented by 12-b words, yields the frequency magnitude response shown in Figure 4(a).

Next, replacing the unit-delay elements in the filter sections with $M = 4$ unit-delay elements results in the frequency magnitude response shown in Figure 4(b). There we see the multiple narrowband passband images induced by the $M = 4$ length delay lines of the interpolated-IIR filter. Our final job is to attenuate those unwanted passband images. We can do so by following the

cascaded increased-delay IIR filter sections with a cascaded integrator-comb (CIC) filter, whose structure is shown in Figure 3(b) [6], [7]. (The CIC filter is computationally advantageous because it requires no multiplications.) To satisfy our desired 60-dB stopband attenuation requirement, we use a second-order CIC filter—two first-order CIC filters in cascade—to attenuate the passband images in Figure 4(b). The result of our design is an interpolated-IIR and CIC filter combination whose composite frequency magnitude response meets our filter requirements; it is shown Figure 4(c).

CONCLUSIONS

We discussed the limitations encountered when using traditional second-order quantized-coefficient IIR filters to perform narrowband low-pass filtering. Next, we showed how Gold and

Rader resolved the problem, albeit with an increased computational cost of doubling the number of multiplies per filter output sample. Finally, we described and then demonstrated the interpolated-IIR filter proposed by Harris and Loudermilk to overcome the shortcomings of traditional low-pass IIR filters. The interpolated-IIR filter provides improved low-pass filter performance while requiring no increase in filter coefficient bit widths and no additional multiply operations beyond a traditional IIR filter. When it comes to narrowband IIR filters, there's a new sheriff in town.

ACKNOWLEDGMENT

Thanks go to Clay S. Turner for cogent suggestions that greatly improved the content of this article.

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) and editor of (and a contributor to) *Streamlining Digital Signal Processing, A Tricks of the Trade Guidebook* (IEEE Press/Wiley, 2007).

REFERENCES

- [1] J. Proakis, and D. Manolakis, *Digital Signal Processing-Principles, Algorithms, and Applications*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall, 1996, pp. 572–576.
- [2] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1989, pp. 382–386.
- [3] B. Gold and C. Rader, “Effects of parameter quantization on the poles of a digital filter,” *Proc. IEEE*, vol. 55, pp. 688–689, May 1967.
- [4] F. Harris and W. Lowdermilk, “Implementing recursive filters with large ratio of sample rate to bandwidth,” in *Conf. Rec. 41st Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, Nov. 4–7, 2007, pp. 1149–1153.
- [5] R. Lyons, “Interpolated narrowband low-pass FIR filters,” *IEEE Signal Processing Mag.*, vol. 20, no. 1, pp. 50–57, Jan. 2003.
- [6] E. Hogenauer, “An economical class of digital filters for decimation and interpolation,” *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-29, pp. 155–162, Apr. 1981.
- [7] R. Lyons, “Understanding cascaded integrator-comb filters,” *Embedded Syst. Program. Mag.*, vol. 18, no. 4, pp. 14–27, Apr. 2005 [Online]. Available: <http://www.embedded.com/showArticle.jhtml?articleID=160400592>