

Performance Estimation for Lowpass Ternary Filters

Brenton Steele

Dynamic Hearing Pty, Ltd., 2 Chapel Street, Richmond 3121, Victoria, Australia
Email: bsteele@dynamichearing.com.au

Peter O'Shea

*School of Electrical and Electronic Systems Engineering, Queensland University of Technology,
GPO Box 2434, Brisbane 4001, Australia*
Email: pj.oshea@qut.edu.au

Received 16 December 2002 and in revised form 25 June 2003

Ternary filters have tap values limited to -1 , 0 , or $+1$. This restriction in tap values greatly simplifies the multipliers required by the filter, making ternary filters very well suited to hardware implementations. Because they incorporate coarse quantisation, their performance is typically limited by tap quantisation error. This paper derives formulae for estimating the achievable performance of lowpass ternary filters, thereby allowing the number of computationally intensive design iterations to be reduced. Motivated by practical communications systems requirements, the performance measure which is used is the worst-case stopband attenuation.

Keywords and phrases: ternary filter, performance estimation, FIR filter, stopband attenuation.

1. INTRODUCTION

Ternary filters typically have a transversal finite impulse response (FIR) section cascaded with a recursive filter section [1]. The transversal FIR section has an impulse response which is quantised to one of three values (-1 , 0 , and $+1$) and the recursive section is usually selected to allow for simple hardware implementation (i.e., it is designed to have no or very few multiplications). Because of the very coarse quantisation of ternary filters, it is often (although not always) necessary to sample at a higher rate than would be required if infinitely fine quantisation were used. Often, optimisation methods are used to ensure that the required sampling rates in ternary filters are kept as low as possible.

It is very important to be able to predict the achievable performance of a ternary filter before the design process begins; otherwise, key filter parameters will have to be determined by trial and error, and many iterations of the optimisation process may be required. Performance is a broad concept, and one must, therefore, clearly define what aspect of performance is important in the design of ternary filters. Previous authors such as Lockhart [2] have used mean square error (MSE) as a measure of performance, while others such as Abeysekera and Padhi [3] have used the filtered quantisation noise power. Both these performance estimates tend to give an indication of a filter's average stopband attenuation, as noise (i.e., error) is most noticeable in the stopband, where the signal-to-quantisation error ratio is the lowest. In many practical situations, what is required is not the average

stopband attenuation, but the maximum level of stopband gain (i.e., the worst-case attenuation). In scenarios where a filter is processing radio frequency (RF) signals, for example, the worst-case attenuation is very important. The purpose of such filters is to block unwanted signals from interfering with the desired signal. If an interferer happened to fall at the same frequency as the worst-case attenuation, it would not be sufficiently attenuated. It is, therefore, desirable that the achievable performance estimation method be capable of predicting the worst-case stopband attenuation in addition to the average stopband attenuation. To be useful, the accuracy of this estimate should be of the order of 10 dB or better.

In what follows, achievable performance estimation formulae based on the worst-case stopband attenuation are derived. Section 2 specifies the assumptions that are made in the derivations. The performance derivations themselves are presented in Section 3. The achievable performance estimates are compared with the actual performance of some practical filters in Section 4. Section 5 then compares the performance of ternary filters to some other common filter architectures, while Section 6 presents conclusions.

2. ASSUMPTIONS

2.1. Filter structure

It is assumed that the ternary filter structure consists of a ternary tap FIR transversal filter cascaded with an integrator,

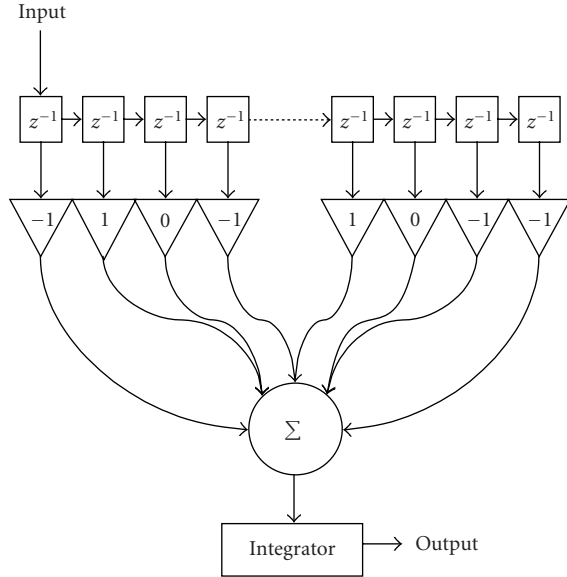


FIGURE 1: Ternary filter structure used in performance estimation derivation.

as illustrated by the diagram in Figure 1. This is a common architecture used for ternary filters, and is simple to implement.

Note that the integrator is not restricted to being a simple first-order integrator; it can have an arbitrarily high order. That is, it is a structure with a transfer function given by

$$H_{\text{int}}(z) = \frac{1}{(1 - z^{-1})^N}, \quad (1)$$

where N is the order. Note that the assumed filter structure has some distinct similarities to sigma-delta modulator (SDM) based filtering schemes. The latter typically have an FIR transversal filter section (with the impulse response being generated by an SDM) followed by a comb filter. Ternary filters thus differ from SDM based filters in that (i) their impulse responses are usually generated by an optimisation algorithm (such as dynamic programming) and (ii) the reconstruction filter is typically an integrator rather than a comb filter. These two differences can provide significant benefits [1]. The use of optimisation algorithms enables the best possible performance for a given optimisation criterion, and the use of integrators enables very efficient hardware implementations. A hardware integrator, for example, requires only a delay and summer, whereas a comb filter requires several delays, a summer, and some storage elements. This hardware advantage is particularly great in cases where the desired filter is very narrowband. If a comb filter section is used to help realise these narrowband filters, the comb section also typically needs to have a narrow bandwidth. This narrow bandwidth in turn translates to a large memory requirement.

2.2. Transversal FIR filter section and target filter

The taps of the transversal FIR filter are quantised to three levels (-1 , 0 , and $+1$). When the impulse response of the transversal FIR filter is passed through the integrator, the multibit impulse response of the complete ternary filter is obtained. The transversal filter is assumed to be designed using a dynamic programming style optimisation algorithm. It is assumed that the target (or desired) filter response is an infinite resolution, lowpass FIR filter with a flat passband and a flat stopband. The stopband attenuation of this ideal filter is infinite.

2.3. Quantisation noise model

The impulse response of the FIR transversal filter section is assumed to be given by

$$h_{\text{Transversal}}(n) = h_{\text{True}}(n) + q_{\text{Transversal}}(n), \quad (2)$$

where $h_{\text{True}}(n)$ is the true or desired transversal filter impulse response and $q_{\text{Transversal}}(n)$ is the additive quantisation noise, the samples of which are assumed to be uniformly distributed. The impulse response of the overall filter (i.e., the transversal filter cascaded with the integrator) will be

$$\begin{aligned} h(n) &= h_{\text{Transversal}}(n) * h_{\text{int}}(n) \\ &= h_{\text{True}}(n) * h_{\text{int}}(n) + q_{\text{Transversal}}(n) * h_{\text{int}}(n) \\ &= h_{\text{True}}(n) * h_{\text{int}}(n) + q(n), \end{aligned} \quad (3)$$

where $h_{\text{int}}(n)$ is the impulse response of the integrator and $q(n)$ is the additive quantisation noise in the overall filter. The samples of $q(n)$ are assumed to be white and identically distributed. The transfer function of $h(n)$ is the transfer function of the overall filter. It is specified by

$$H(\omega) = H_{\text{True}}(\omega)H_{\text{int}}(\omega) + Q(\omega) = H_{\text{Target}}(\omega) + Q(\omega), \quad (4)$$

where $H_{\text{Target}}(\omega)$ is the target filter transfer function, and $H_{\text{True}}(\omega)$, $H_{\text{int}}(\omega)$, and $Q(\omega)$ are the Fourier transforms of $h_{\text{True}}(n)$, $h_{\text{int}}(n)$, and $q(n)$, respectively.

It is important to note that the assumptions listed in the previous paragraph may not always be fully satisfied in practice. Firstly, the assumption that $q(n)$ is white implies that the stopband (which should be dominated by the error signal) is flat. This will only be the case if the filter design process creates the desired flat passband, and this will not always be the case. If, for example, time domain MSE criteria are used for the optimisation criterion, the stop bands are often not flat. Secondly, the actual impulse response will differ from the ideal target impulse response not only because of the quantisation noise, but also the limitations of the finite length of the transversal FIR filter. For these reasons, the performance estimates derived in this paper can be considered estimates of “achievable performance.” They can be achieved if the errors due to the finite length of the transversal filter impulse response are insignificant compared with the quantisation noise effects and if the ternary filter design algorithm spreads the quantisation noise out evenly over the entire band. These issues are explored more fully in [4].

3. DERIVATION

3.1. Average stopband attenuation

The ternary filter response is assumed to be given by

$$H(\omega) = \begin{cases} G_{\text{PB}}, & 0 < \omega < 2\pi \frac{F_C}{F_S}, \\ G_{\text{SB}}, & 2\pi \frac{F_C}{F_S} < \omega < \pi, \end{cases} \quad (5)$$

where G_{PB} and G_{SB} are the passband and stopband gains, respectively. The response in (5) is depicted graphically in Figure 2. The performance analysis of a ternary filter can be simplified by separating the filter into the transversal FIR and integrator parts. The overall response can then be defined as

$$H(\omega) = H_{\text{Transversal}}(\omega) \times H_{\text{int}}(\omega), \quad (6)$$

where $H_{\text{Transversal}}(\omega)$ is the transversal FIR filter's response and $H_{\text{int}}(\omega)$ is the integrator's response. The amplitude response for an arbitrary N th-order integrator can easily be shown to be

$$|H_{\text{int}}(\omega)| = \frac{1}{(2 \sin \omega/2)^N}, \quad 0 < \omega < 2\pi. \quad (7)$$

Substituting (7) into (6) and combining with (5) yield the following transversal filter amplitude response:

$$|H_{\text{Transversal}}(\omega)| = \begin{cases} G_{\text{PB}} \left(2 \sin \frac{\omega}{2}\right)^N, & 0 < \omega < 2\pi \frac{F_C}{F_S}, \\ G_{\text{SB}} \left(2 \sin \frac{\omega}{2}\right)^N, & 2\pi \frac{F_C}{F_S} < \omega < \pi. \end{cases} \quad (8)$$

The amplitude response in (8) is graphically illustrated in Figure 3. In accordance with the assumptions in Section 2, the quantisation noise level in the transversal filter section will largely dictate the achievable stopband level of the overall filter. Assuming that the design algorithm spreads the quantisation error of the overall ternary filter evenly across the entire band, the expected value of the quantisation error amplitude spectrum will be

$$|H_{\text{NoiseComp}}(\omega)| = G_{\text{SB}}, \quad 0 \leq \omega \leq \pi. \quad (9)$$

The quantisation noise transfer function for the transversal section of the filter will thus be

$$|H_{\text{TransversalNoiseComp}}(\omega)| = G_{\text{SB}} \left(2 \sin \frac{\omega}{2}\right)^N, \quad 0 \leq \omega \leq \pi. \quad (10)$$

Assume for the moment that one wants to determine the SNR (which is effectively the passband-to-stopband ratio) in an ideal lowpass transversal FIR filter impulse response in which (i) the quantisation error is uniformly distributed, (ii) there are ternary coefficients, (iii) the cutoff frequency is F_C , (iv) the sampling frequency is F_S , and (v) there is *no* noise

shaping. The standard quantisation noise equation may be used to calculate such an SNR [5, page 37]. It is given by

$$\text{SNR}_{\text{no-shaping}} = 1.76 + 20 \log 3 - 10 \log \left(\frac{F_C}{F_S/2} \right) \quad (11)$$

$$= 11.3 - 10 \log \left(\frac{F_C}{F_S/2} \right). \quad (12)$$

Now, the transversal FIR filter section under consideration in this paper differs from that described in the previous paragraph in that there is noise shaping. Accordingly, the expression for the SNR of the transversal filter section in this paper can be obtained by appropriately amending the expression in (12). The improvement in SNR at F_C (i.e., at the edge of the passband) due to noise shaping will be given by the ratio of the average noise power spectral density to the noise power spectral density at F_C . Using (10), this ratio is

$$\text{SNR}_{\text{improvement}}(F_C) = 10 \log \left[\frac{(1/\pi) \int_0^\pi [G_{\text{SB}} (2 \sin \omega/2)^N]^2 d\omega}{[G_{\text{SB}} (2 \sin \pi F_C/F_S)^N]^2} \right]. \quad (13)$$

The SNR at F_C can thus be obtained by adding the right-hand sides of (12) and (13), giving

$$\begin{aligned} \text{SNR}_{F_C} &= 11.3 - 10 \log \left(\frac{F_C}{F_S/2} \right) \\ &+ 10 \log \left[\frac{1}{\pi} \int_0^\pi \left(2 \sin \frac{\omega}{2}\right)^{2N} d\omega \right] \\ &- 20N \log \left[2 \sin \left(\frac{\pi F_C}{F_S} \right) \right]. \end{aligned} \quad (14)$$

Now, in conventional lowpass SDM systems, the signal transfer function in the passband is typically flat, while the noise transfer function increases monotonically with frequency. For SDM systems, therefore, the SNR varies as a function of frequency within the passband. For the system considered in this paper, however, the signal and noise transfer functions have (up to a constant scaling factor) the *same* spectral shape. See, for example, (8) and (10). As a result, the SNR in the passband is constant. The SNR in (14), therefore, represents the SNR across all frequencies within the passband. It may also be thought of as the “average stopband attenuation.”

The second last term of the right-hand side of (14), is independent of the filter bandwidth, and can be calculated in advance to simplify predictions. Values of this term (defined as k) for integrators of first- to fourth-order are given in Table 1.

3.2. Worst-case stopband attenuation

Now that the “average stopband attenuation” is available via (14), the worst-case stopband attenuation must be calculated. As stated in Section 2.3, the quantisation noise in the impulse response for the complete filter is assumed to be




FIGURE 2: Desired ternary filter response.




FIGURE 3: Transversal FIR filter response.

white and identically distributed. The transfer function of the overall filter is obtained by taking the discrete Fourier transform of this impulse response. The real and imaginary components of the samples in the transfer function can, therefore, be well approximated as white and Gaussian, using a central limit theorem argument.

To find the worst-case stopband attenuation, the highest amplitude component in the stopband region (i.e., the highest amplitude “frequency bin”) must be determined. Since the real and imaginary components of the stopband samples are assumed Gaussian, the amplitude will be Rayleigh [6]. This type of distribution can be described by (see [6, pages 59, 442]) the probability density function (pdf)

$$f_X(x) = \begin{cases} \frac{2}{b}(x-a)e^{-(x-a)^2/b}, & x \geq a, \\ 0, & x < a, \end{cases} \quad (15)$$

 TABLE 1: Constant k versus integrator order.

Integrator order	k (dB)
First	3.0
Second	7.8
Third	13.0
Fourth	18.5

and with the cumulative distribution function (cdf)

$$F_X(x) = \begin{cases} 1 - e^{-(x-a)^2/b}, & x \geq a, \\ 0, & x < a. \end{cases} \quad (16)$$

The value of a and b can be related to the mean, \bar{X} , of the Rayleigh distributed random variable through

$$\bar{X} = a + \sqrt{\frac{\pi b}{4}}. \quad (17)$$

In the scenario discussed here, $a = 0$, so b can be defined in terms of the mean via

$$b = \frac{4}{\pi} \bar{X}^2. \quad (18)$$

For a given level of confidence, the peak value of a single frequency bin in the stopband can be deduced from (16) and (18). The following equation gives the amplitude (relative to the mean) below which the amplitude in a stopband bin will fall with (100P%) confidence:

$$\frac{x_{\max}}{\bar{X}} = \sqrt{-\frac{4}{\pi} \ln(1-P)}. \quad (19)$$

Now, the stopband bin values, being the result of discrete Fourier transformation of the impulse response, are all independent. Since there will be $N_{\text{bins}} = N_{\text{taps}}(0.5 - F_C/F_S)$ independent samples in the stopband, the amplitude (relative to the mean) below which all stopband bin values will fall with (100P%) confidence is

$$\left(\frac{x_{\max}}{\bar{X}}\right) = \sqrt{-\frac{4}{\pi} \ln(1 - P^{1/N_{\text{bins}}})}. \quad (20)$$

The maximum relative deviation in (20) can also be expressed in dB by the equation

$$20 \log\left(\frac{x_{\max}}{\bar{X}}\right) = 20 \log\sqrt{-\frac{4}{\pi} \ln(1 - P^{1/N_{\text{bins}}})}. \quad (21)$$

Equation (21) can be applied to the problem of estimating the worst-case stopband attenuation for the ternary filter specified in Sections 2 and 3. The mean, \bar{X} , for the case under consideration is the mean value of G_{SB} . The right-hand side of (21) gives the amount (in dB) by which the average stopband attenuation exceeds the worst-case stopband

attenuation, with $P\%$ confidence. By combining (14) with (21), an estimate of the worst-case stopband attenuation can be derived

$$\begin{aligned} \text{Atten}_{\text{Worst_case}} = & 11.3 - 10 \log \left(\frac{F_C}{F_S/2} \right) \\ & - 20N \log \left[2 \sin \left(\frac{F_C}{F_S} \pi \right) \right] \\ & + k - 20 \log \sqrt{-\frac{4}{\pi} \ln(1 - P^{1/N_{\text{bins}}})} \text{ dB}, \end{aligned} \quad (22)$$

where k is as in Table 1.

4. SIMULATION RESULTS

The dynamic programming design method in [4] was used to design ternary filters with a range of normalised cutoff frequencies (F_C/F_S) and integrator orders. The performance of these filters was compared with the achievable performance predicted by (14) and (22).

The graphs in Figures 4 and 5 show the comparison between measured performance and achievable performance. The circles, triangles, squares, and diamonds indicate measurements for first-order, second-order, third-order, and fourth-order integrators, respectively. The achievable performance is indicated with a solid line. The performance measure considered in Figure 4 is average stopband attenuation, while that considered in Figure 5 is the worst-case stopband attenuation, assuming a confidence level of 95%. This comparison is performed for ternary filters with first- to fourth-order integrators, and for normalised cutoff frequencies ranging from 1/40 to 1/8.

It can be seen from Figures 4 and 5 that for first-, second-, and third-order filters, the dynamic programming design method in [4] yields performance which is comparatively close to the achievable limits predicted by (14) and (22). That is, the stopband attenuation for these filters is within about 10 dB of that predicted by (14) and (22). For *fourth-order* filters, however, the dynamic programming design method in [4] gives rise to filters with comparatively poor stopband attenuation. The reason for the comparatively poor performance in the fourth-order case is due to the fact that the dynamic programming method in [4] is an MSE-based method which does not guarantee the optimum solution in terms of the worst-case stopband attenuation. Rather it only gives solutions which tend to minimise the *average* stopband attenuation. The fourth-order integrator based filters have been observed to have nonflat passbands, which have relatively low average stopband attenuations, but comparatively high worst-case stopband attenuations. Note that the existence of these nonflat passbands violates one of the key assumptions in Section 2.3, namely that the error samples are white and identically distributed. This suboptimal performance suggests the need for alternative ternary filter design methods. An alternative method that uses a minimax criterion is presented in [4].




FIGURE 4: Average stopband attenuation versus 1/normalised cutoff frequency (F_S/F_C).

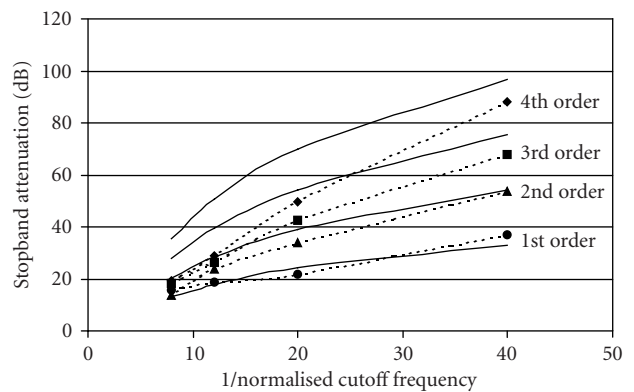


FIGURE 5: Worst-case stopband attenuation versus 1/normalised cutoff frequency (F_S/F_C).

5. COMPARISON WITH OTHER FILTERS

The graph in Figure 5 provides some insight into the effectiveness of ternary filters as compared to traditional FIR filters which use high precision filter taps. Suppose, for example, that one requires a filter with a normalised cutoff of 1/16 and a minimum stopband attenuation of 40 dB. Figure 5 indicates that such a filter can be readily implemented as a fourth-order ternary filter. The performance of such a ternary filter would be virtually the same as that of a conventional filter employing full precision [1]. The hardware for the ternary filter, however, would be much simpler than that of a full precision filter.

Now, suppose that one requires a filter different to that described in the previous paragraph, with a normalised cutoff frequency of 1/16 and a minimum stopband attenuation of 80 dB. Figure 5 indicates that 80 dB stopband attenuation can only be achieved for ternary filters with normalised cutoff frequencies of less than 1/36 (assuming that the ternary filter order is restricted to 4). To implement the required ternary filter, therefore, one would have to artificially increase the sampling rate of the input data stream by 36/16 (i.e., by 2.25). In such a case it is recommended that the

“interpolated” data stream be created with an SDM which produces reduced precision data samples. Despite the increase in the input signal sampling rate of 2.25, the ternary filter implementation would still be highly efficient. This is so because the filter taps not only would be very coarsely quantised, but the SDM derived data samples would be reduced in precision as well. Therefore, the multiplications in the filter would be trivial, and the additions for the transversal filter section would be low precision additions.

Note that while ternary filters are useful in many situations they are particularly useful for filtering sigma delta modulated data streams. The latter are found in many analog to digital converter systems and, more recently, in many digital amplification systems. SDM data streams are usually represented in single-bit format and tend to have small normalised bandwidths. If single-bit SDM data is fed into ternary filters, the low precision in both the input and the filter taps enables the multiplications and the additions in the ternary filter’s transversal filter section to be implemented with very simple hardware. Furthermore, the low normalised bandwidths of SDM derived data streams tend to lend themselves to ternary filtering with quite high stopband attenuations. In fact, it can be a useful filtering strategy to first transform the input into an SDM data stream and then apply a ternary filter. A field programmable gate array (FPGA) filter based on this approach is described in [7]. One can even use canonical signed digit (CSD) filter taps to increase performance even further. In [7], for example, it was found that by using filter taps quantised to ± 2 , ± 1 , and 0 rather than ternary coefficients, an extra 7 dB of attenuation could be obtained. This additional performance came at the cost of a 20% increase in FPGA resources.

It is instructive to compare ternary filter performance with that of frequency response masking (FRM) filters. The latter tends to be more suited to wideband filters, while the former tends to be suited to narrowband ones. The FRM design example in [8] will be used for comparison as it seems to be a particularly efficient filter. Being comparatively wideband, it is also the kind of example which is likely to favour an FRM implementation, and will thus be an exacting benchmark for a ternary filter. This filter (with a fractional bandwidth of 0.3 and stopband attenuation of 80 dB) required only 164 multipliers using the FRM method, whereas a traditionally designed FIR filter would have required 3177 taps. To achieve 80 dB stopband attenuation, a ternary filter with a fourth-order integrator and fractional bandwidth of 1/36 would be required (i.e., the traditionally designed filter response would have to be oversampled about 11 times). The ternary filter would, therefore, be expected to have around 35000 taps, with about 23000 of these being nonzero (i.e., ± 1) taps. Now, assuming that the input data has been interpolated with a 1-bit SDM, all of the required 23000 multiplications could be implemented with single gates. By contrast, the FRM would typically require 164 sixteen-bit multiplications, with each multiplication requiring many gates. If each 16-bit multiply requires more than 140 gates, then the ternary filter would be arguably the more efficient implementation.

6. CONCLUSIONS

A method for predicting the performance of ternary filters has been developed. The newly derived performance prediction formulae enable the designer to predict a priori the order of the recursive filter component, and in doing so, reduce the number of required design iterations.

ACKNOWLEDGMENT

The work of Brenton Steele was undertaken at the School of Electrical and Computer Engineering, RMIT University.

REFERENCES

- [1] N. Benvenuto, L. E. Franks, and F. S. Hill Jr., “Dynamic programming methods for designing FIR filters using coefficients -1 , 0 , and $+1$,” *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 34, no. 4, pp. 785–792, 1986.
- [2] G. B. Lockhart, “Digital encoding and filtering using Delta modulation,” *The Radio and Electronic Engineer*, vol. 42, no. 12, pp. 547–551, 1972.
- [3] S. S. Abeyssekera and K. P. Padhi, “Design of multiplier free FIR filters using a LADF Sigma-Delta (Σ - Δ) modulator,” in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS '00)*, vol. 2, pp. 65–68, Geneva, Switzerland, May 2000.
- [4] B. R. Steele, *Efficient signal processing through the use of Sigma Delta modulation and ternary filters*, Ph.D. thesis, RMIT University, Melbourne, Australia, 2003.
- [5] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1996.
- [6] Z. P. Peyton Jr., *Probability, Random Variables and Random Signal Principles*, McGraw-Hill, NY, USA, 4th edition, 2001.
- [7] J. Cao and P. O’Shea, “Efficient implementation of short wordlength FIR digital filters in FPGA,” in *Proc. 4th Australasian Workshop on Signal Processing and Applications (WoSPA '02)*, Stamford Plaza, Brisbane, Australia, December 2002.
- [8] L. Zhang, Y. Lian, and C. C. Ko, “A new approach for design sharp FIR filters using frequency response masking technique,” in *Proc. 9th IEEE DSP Workshop and 1st IEEE Signal Processing Education Workshop*, Hunt, Tex, USA, October 2000.

Brenton Steele received the B.Eng. degree in electrical and computer systems engineering (with first class honours) from Monash University in 1992. Then he worked for telecommunications carrier Telstra, automotive manufacturer General Motors Holden, and radio modem manufacturer Clyrcom. In March 2001, he joined the School of Electrical and Computer Engineering at RMIT University as a Research Fellow, working on a range of signal processing research projects. He completed his Ph.D. at RMIT University in 2003, undertaking research in the areas of sigma-delta modulators and ternary filters. In February 2003, he joined Dynamic Hearing Pty, Ltd. where he researches and develops signal processing algorithms for use in hearing aids. His research interests are in signal processing for communications and audio applications.



Peter O'Shea received the B.E, Dip.Ed., and Ph.D. degrees in 1978, 1983, and 1991, respectively, all from the Queensland University of Technology (QUT). He worked as an engineer at the Overseas Telecommunications Commission for three years, at the University of Queensland's Department of Electrical Engineering for four years, at QUT's School of Electrical and Electronic Systems Engineering for three years, and in the School of Electrical and Computer Systems Engineering, RMIT, for seven years. While at RMIT he received the Faculty of Engineering and President's Awards for Student Centred Teaching. He has recently rejoined the staff of the School of Electrical and Electronic Systems Engineering at QUT. His interests are in signal processing for communications and power systems, reconfigurable computing, and the use of multimedia in engineering education.

