Polyphase Filter Design with Reduced Phase Non-Linearity

DR. ARTUR KRUKOWSKI and DR. IZZET KALE
Applied DSP and VLSI Research Group,
University of Westminster,
115 New Cavendish Street, London,
UNITED KINGDOM

Abstract: - The perfect linear phase requirement can be achieved only by direct design with an FIR filter at the expense of high filter order and hence high computational burden for the given specification. Alternatively, IIR filters can be designed with much smaller orders than their FIR counterparts, but at the expense of the non-linear phase. However, it is possible to design almost linear-phase IIR filters for a given specification of the allowed phase or group delay ripples. This paper presents an improvement to the existing algorithms for the design of almost linear–phase, arbitrary-band IIR filters and gives a different viewpoint on these methods, showing a considerable improvement in maximum ripples of the magnitude and group delay responses. The idea of decreasing the phase non-linearity of the N-path polyphase IIR filters with a combination of allpass sections of different order and coefficient polarity is presented. The effect of coefficient quantization, in view of the fixed-point implementation, on the performance of the resulting filter is also shown.

Key-Words: - Filter design, Phase-linearity, Digital filter, Polyphase structure, Multi-rate system.

1 Introduction

For systems, which cater for wide-band signals, it is important to ensure the same time differences between signal spectrum components before and after filtering. This requirement is met if the phase response of the filter is linear and has no constant factor, or when the group delay function is constant:

$$\tau(\nu) = -\frac{1}{2\pi}\frac{\partial\phi(\nu)}{\partial\nu} = const \qquad (1)$$

A flatness of the group delay function is enough to assure the phase response linearity for the polyphase filter [1]-[3], shown in Fig. 1, due to its monotonic phase response crossing the zero at DC.

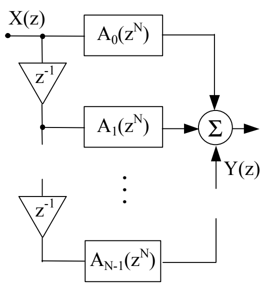

$$H(z) = \frac{1}{N}\sum_{n=0}^{N} A_n(z) z^{-n}$$

$$A_n(z) = \prod_{k=1}^{K_n} A_{n,k}(z) \qquad (2)$$

$$A_{n,k}(z) = \frac{\alpha_{n,k} + z^{-N}}{1 + \alpha_{n,k} z^{-N}}$$

Fig. 1 The general N-path polyphase structure.

The group delay of the structure is given by:

$$\tau(\nu) = \frac{N}{2}\left(\frac{N-1}{2} + \sum_{k=1}^{K} \frac{1-\alpha_\kappa^2}{1+\alpha_\kappa^2 + 2\alpha_\kappa^2 \cos 4\pi\nu}\right) + \sum \delta(\nu_z) \qquad (3)$$

The shape of the overall group delay function is an average of group delays of all allpass subfilters. The second term in (3) is responsible for the phase jumps due to zeros of the filter on the init circle. The example shapes of the group delay for different coefficients and section orders are shown in Fig. 2.

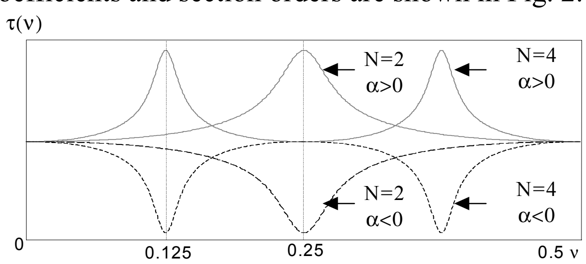

Fig. 2 Example group delay of the polyphase filter.

The phase non-linearity can be decreased in two ways, either by designing the corrector with the phase response opposite to the filter's one or by designing the filter directly for almost flat group delay. It can be noticed from Fig. 2 that for negative coefficients the group delay is symmetric against the Y-axis to the one for the positive coefficient. Therefore the correction can be achieved with a combination of allpass filters having negative coefficients. A cascade of different order allpass sections can be used to achieve a better correction.

The second way of designing approximately linear-phase IIR filters is to modify the structure from Fig. 1 by replacing the allpass filter in one of the branches with a bulk delayor, z^{-K}, (K being the order of the allpass filter) as in Figure 3. The allpass filter A(z) represents the transfer function of allpass filters in other branches of the structure.

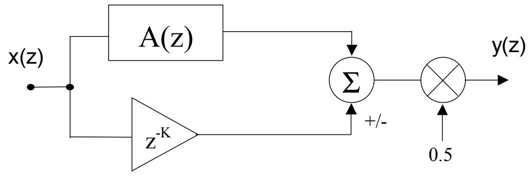

Fig. 3 The structure of an approximately linear-phase polyphase two-path filter.

In order to achieve lowpass filter response the allpass filter A(z) has to be designed in-phase with the bulk delay at the low frequencies and distant by π at the frequencies close to Nyquist (or other way round for the highpass filter). This way A(z) and hence the overall filter have an almost linear-phase. The idea as such is not new and was suggested and used by Curtis [1][4] and employed in designs published by Lawson [5] and Lu [6]. The current design methods are based on the standard idea of composing two identical IIR (non-linear phase) filters to achieve an approximately linear-phase characteristic [7] or applying iterative quadrature programming methods [6]. Such an approach does not allow much flexibility, limiting the number of points of freedom to half of what would be available when standard IIR filters are used. Thus the resulting filters posses larger stopband ripples as well as larger group delay ripples than what the structure is really capable of achieving. The additional advantage of the presented structure is that it is capable of providing a complimentary set of lowpass and highpass filters, simply by replacing the adder with a subtractor [1]-[3]. This feature is very attractive for applying such filters in almost linear-phase IIR quadrature filter banks.

2 Phase Compensation with Cascade of Allpass Sections

Correcting the group delay (phase linearity) of the polyphase filters with allpass sections $A_{n,k}(z)$ from (2) having negative coefficients is especially applicable within frequency bands, which are the power of two divisions of the Nyquist frequency, i.e. 0.25, 0.125 etc. As it was shown before in Fig. 2, the shape of the allpass phase response having a negative coefficient is opposite to the one having the positive one. As the total group delay of the polyphase structure is equal to the average of its allpass components, it has the same bell-like shape, peaking at half-Nyquist for second-order allpasses, quarter-Nyquist for forth-order ones, etc. Here the case of two-path structure (N=2) is discussed. For structures with more paths the compensation is very similar.

The simplest corrector is a single second-order allpass section. The best correction is achieved when the group delay of the final filter at DC is equal to the group delay at its cutoff frequency, ν_c, i.e.:

$$\tau(\nu_c) = \tau(0)$$

$$\tau_o(\nu_c) + \frac{2\left(1-\alpha_{c0}^2\right)}{1+\alpha_{c0}^2+2\alpha_{c0}^2\cos 4\pi\nu_c} = \tau_o(0) + 2\frac{1-\alpha_{k0}}{1+\alpha_{c0}} \quad (4)$$

where α_{ck} are corrector coefficients. The effect of such correction can be seen in Fig. 4 for a two-coefficient (α_0=0.125, α_1=0.525) two-path (N=2) filter. This single-coefficient corrector give a 6.5 times decrease of group delay peak-to-peak error in the signal band up to $\nu_c=0.125$.

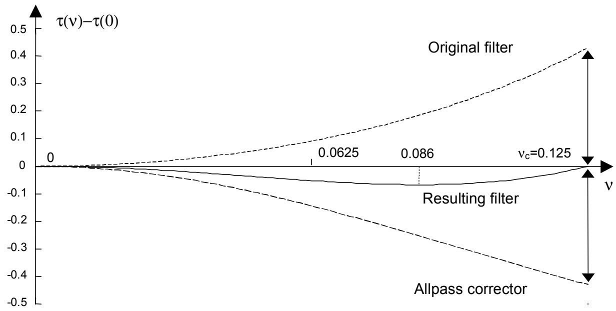

Fig. 4 The group delay correction with a single allpass section for the example two-coefficient filter up to $\nu_c=0.125$.

The correction can be improved by cascading additional higher order allpass sections. In general arbitrary orders can be chosen. However, the best results are achieved by for doubling of the consecutive orders of the allpass sections. Apart from symmetry of compensation, this gives an additional advantage for multirate systems [8], [9]. The order of the second compensator section should be chosen to be $0.5/\nu_c$, i.e. the overall K-section compensator transfer function becomes:

$$H_c(\nu) = \frac{\alpha_{c0}+z^{-2}}{1+\alpha_{c0}z^{-2}} \cdot \sum_{m=0}^{K-2} \frac{\alpha_{c0}+z^{-\frac{2^m}{\nu_c}}}{1+\alpha_{c0}z^{-\frac{2^m}{\nu_c}}} \quad (5)$$

Higher order compensators can be designed by minimizing the sum of the peak differences of the overall group delay squared with respect to an average value, as in Fig. 5, i.e.:

$$C = \sum_{k=1}^{K} \left|\tau(\nu_k) - \tau_{med}\right|^2 \quad , \quad \tau_{med} = \frac{1}{K}\sum_{m=1}^{K} \tau(\nu_m) \quad (6)$$

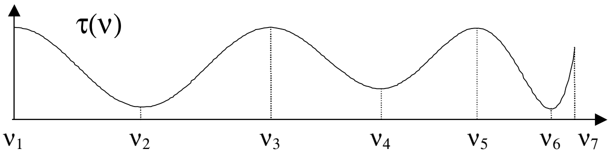

Fig. 5 The non-linear optimisation cost function.

It was noticed for the number of corrector sections greater than two that the absolute values of coefficients for consecutive sections approximately follow a geometric series. Therefore finding the first two section coefficients allows estimating a good starting point for the multi-section corrector optimisation. The calculation of the second section coefficient is explained in Fig. 6.

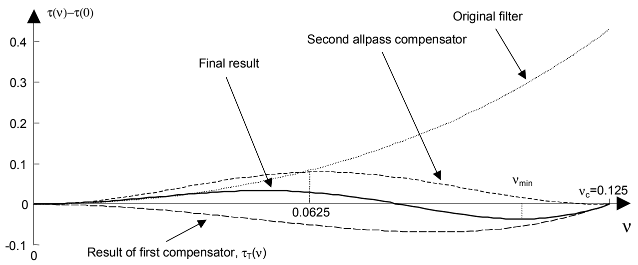

Fig. 6 The correction with two allpass sections.

The coefficient α_{c1} is calculated to satisfy (7):

$$\begin{aligned}&\tau_T\left(\frac{\nu_c}{2}\right)-\tau_T(0)+\tau_c\left(\alpha_{c1},\frac{\nu_c}{2}\right)-\tau_c(0)=\\&=\tau_T(\nu_c)-\tau_T(\nu_{\min})+\tau_c(\alpha_{c1},\nu_c)-\tau_c(\alpha_{c1},\nu_{\min})\end{aligned}\tag{7}$$

where τ_T is the group delay of the filter compensated with a single allpass section. Then, the initial set of compensator coefficients for an optimisation is:

$$\alpha_{init}=\left\{\alpha_{c0},\alpha_{c1},\frac{\alpha_{c1}^2}{\alpha_{c0}},\ldots,\frac{\alpha_{c1}^{K-1}}{\alpha_{c0}^{K-2}}\right\}\tag{8}$$

The group delay correction up to $\nu_c=0.125$ for the same example two-coefficient filter with four allpass sections is shown in Fig. 8.

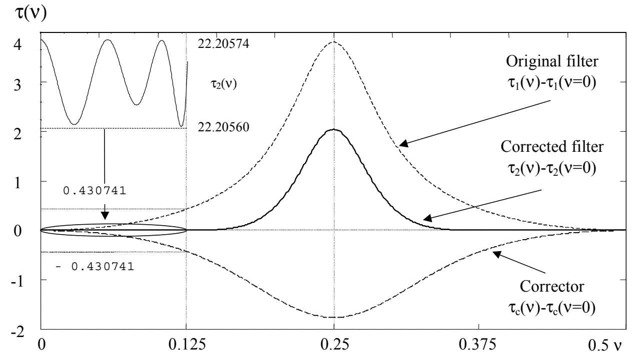

Fig. 8 Four coefficient group delay correction for the example polyphase LPF up to $\nu_{cut}=0.125$.

The group delay peak-to-peak difference was decreased from 0.43 to 1.3e-4 - over three thousand times! The correction results with the number of sections between one are four are shown in Table 1 for both floating-point (FP) and fixed-point (FX) corrector coefficients.

Sections	1	2	3	4
FP	6.50	14.74	215.78	3250
4-bit FX	2.85	2.85	2.85	2.85
8-bit FX	6.21	13.47	13.47	13.47
16-bit FX	6.34	14.73	214.47	2582

Table 1 Decrease of the peak-to-peak group delay ripples using different number of allpasses.

The factor K is a ratio between the group delay peak-to-peak error before and after the correction:

$$K=\left.\frac{\tau_1(\nu)_{\max}-\tau_1(\nu)_{\min}}{\tau_2(\nu)_{\max}-\tau_2(\nu)_{\min}}\right|_{\nu<\nu_c}\tag{9}$$

As polyphase structures are often used in high speed filtering application and most of the times are implemented in fixed-point arithmetic, it is required for the corrector coefficients to be constraint to short bit lengths. The constraint corrector coefficients were calculated from the floating-point ones by truncating them to four, eight and sixteen bits and re-optimising. He results are shown in Table 1. The performance of the correction is very much dependant on the bit length of the coefficients. It decreases quickly when shortening the wordlength of the coefficients because their values in consecutive allpass sections follow a geometric series, quickly converging to zero. For small bit lengths some of the coefficients are too small for the given number of bits and the number of corrector coefficients decreases. If the correction performance becomes unsatisfactory, an alternative has to be sought.

3 Almost Linear-Phase IIR Filter Design Algorithm

The alternative way to achieve almost linear-phase IIR filter is to use the structure from Fig. 3. The implemented design routine was based on the discrete filter least squares fit to the frequency response data ('invfreqz') fro Matlab to approximate the phase of the filter in the passband to be z^{-K} and $z^{-K-0.5}$ in the filter stopband. The important factor was the choice of the weighting function. Choosing constant weights for the passband and stopband led to stopband ripples decreasing monotonically with frequency with passband ripples monotonically increasing. Therefore an iterative update of weights was done at every iteration step of the optimisation, which was changing them according to the shape of the envelope of the magnitude response and group delay ripples in the passband and in the stopband. In general case the allpass filter used in the top branch of the structure from Fig. 3, does not necessarily have to be symmetric against ν=0.25. Also in order to ensure that both small passband ripples and high stopband attenuation is achieved, it is important to monitor the group delay ripples both in the filter passband and its stopband.

The design algorithm can be described as follows:

1. Specify the required complex magnitude response shape to equal the response of the z^{-K} delayor within the filter passband and equal the response of a $z^{-K-0.5}$ delayor in the filter stopband. Also specify the frequency grid in a logarithmic scale to be denser close to the transition band.
2. Choose the initial weights, $W(\nu)$, for the optimization equal to unity at all frequencies.
3. Perform weighted least-squares fit to the frequency response data.
4. Calculate overall group delay of the filter, $\tau(\nu)$.
5. Calculate the maxima of the modulus of the group delay function, $|\tau(\nu)|$, and interpolate the new function, $\tau^{*}(\nu)$, approximating $\tau(\nu)$ through these points.
6. Update the weights using:
$$W(\nu) = [1 + W(\nu)] \cdot [1 + \tau^{*}(\nu)] - 1 \qquad (10)$$
7. Normalize the maximum value of the weights to unity and scale the rest of them accordingly.
8. If the iteration number is less than the maximum, proceed to point three, otherwise deliver the answer vector.

3.1 Experimental Results

During the experiments, it was found that a maximum of four iterations was required to achieve the final result within 1% difference with regard to the result obtainable if iterations were to continue for infinite number of iterations. In order to measure the performance of the method it was compared to the similar approaches suggested by Lu [6] and Lawson [5]. The example filters were designed according to the specifications given in these publications. Comparative results showing the stopband attenuation, the magnitude response ripples and the group delay deviations for the given passband and stopband cut-off frequencies, ν_p and ν_s respectively, are given in Table 2.

Design	N	ν_p, ν_s [-]	ε_p [dB]	A [dB]	$(\tau_{max} - \tau_{min})/2$ [samples]
Lu	15	0.2 0.28	0.0942	31.84	1.25 %
WLSq.	14	0.2 0.28	0.0002	53.2	0.07/0.225
Lawson	15	0.25 0.298	0.1	32.5	0.9
WLSq.	14	0.25 0.298	0.0004	38.6	0.37/1.25

Table 2 Comparison of the weighted least squared approach to approximating the linear phase of the polyphase LPF to the Lu's [6] one and the Lawson's [5] one.

Plots of the filters designed through our method are shown in Fig. 9 and Fig. 10. It can be clearly seen from these plots that the suggested weighted least square method of is advantageous when compared both to [5] and [6].

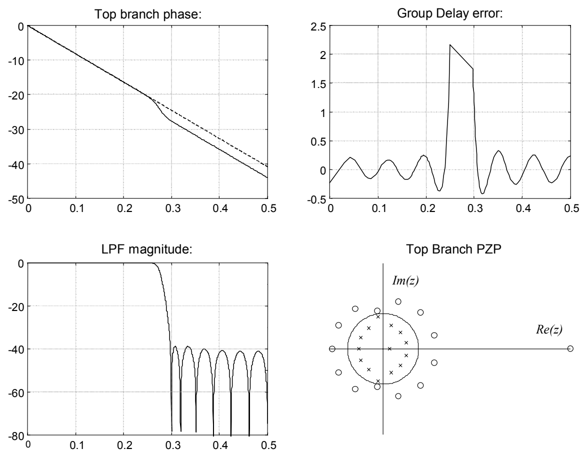

Fig. 9 Example filter designed to the same specification as Lawson's one [5].

In both cases the filter order was chosen to be one less than the ones in the competitive designs and even so the suggested approach delivered much better magnitude and group delay ripples both in the passband and in the stopband. There are two values given for the group delay ripples in Table 2. The first one is the maximum ripple value to within 96% of the bandwidth and the second one for the full bandwidth. It can be clearly seen from Fig. 9 and Fig 10 that the stopband deviations are not equiripple. The purpose was to make a compromise between achieving maximum attenuation and minimum group delay ripples in the passband.

However, it was not possible to achieve both equiripple group delay and equiripple maximum stopband attenuation [7]. It was noticed that increasing the requirements for the group delay ripples led to degradation in the stopband performance and vice versa.

Decreasing the group delay ripples by a few percent was causing degradation in the stopband attenuation by a few dB. The suggested design method was tested on a number of examples for different cut-off frequencies and transition bandwidths. The best performance for the given filter order and transition band specification were achievable for the case of the halfband filter. In such a case the design took advantage of the symmetric allpass filter response, which was easier to achieve.

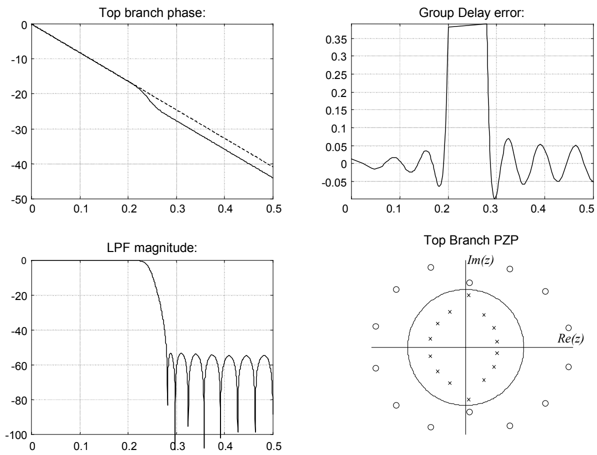

Fig. 10 Example filter designed to the same specification as Lu's one [2].

Design	N [-]	ν_p, ν_s [-]	ε_p [μdB]	A [dB]	$(\tau_{max}-\tau_{min})/2$ [samples]
New Lu	14	0.21 0.29	19.0	54.4	0.035/0.073
New Lawson	14	0.226 0.274	420.0	38.5	0.18/0.512

Table 3 Performance of the example filters with original Lu's and Lawson's specifications with the cutoff frequency set at $\nu=0.25$.

In order to demonstrate this idea the filter specification from the previous example has been modified to save its cutoff frequency at $\nu=0.25$. The performance of these example filters, which were forced to be symmetric, is presented in Table 3. It can be seen that the stopband attenuation does not increase much when the filter is symmetric against $\nu=0.25$ improving only by 1.5dB for Lu's modified specification and not at all for Lawson's modified filter. This is because the cutoff frequency in the original specification was very close to $\nu=0.25$. However, the phase response is much more dependent on the symmetry of the transfer function of an allpass, which caused a significant difference in the group delay ripples. The ripples of the group delay decreased by the factor of two for both filters.

3.2 Constraint coefficient case

It has been shown in a number of publications [8], [9] that the polyphase IIR structure as given in [1]-[3] is very attractive for constraining of its coefficients. This is true also for the special case of the polyphase IIR structure as presented in this paper. Additionally, for the case of the half-band filter the allpass in the top branch of Figure 1 can be easily decomposed into the cascade of second-order one-coefficient allpass sections. Then for the implementation of such a filter, the required number of multiplications would be equal to half of the filter order. The coefficients of these filters can be subsequently constraint to a limited wordlength and re-optimised for the best performance using, for example, the bit-flipping algorithm [10]. The effect of constraining the coefficients of the filter having Lu's specification [6] to 4, 8 and 12 bits is presented in Figure 11.

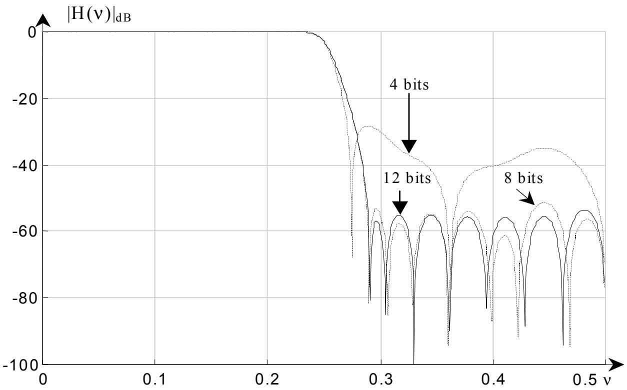

Fig. 11 The magnitude response of the filter designed to Lu's specification [6] with limited coefficient wordlengths.

Number of bits	ε_p [μdB]	A [dB]	$(\tau_{max}-\tau_{min})/2$ [samples]
4	6391.0	28.2	0.11435
6	130.03	45.2	0.07103
8	31.9	51.3	0.07109
12	19.0	53.6	0.07152
16	19.0	53.6	0.07165

Table 4 Performance of the filter designed to Lu's specification with constrained coefficients.

The filter attenuation, magnitude and group delay ripples in the passband are summarized in Table 3. For these tests all the coefficients were truncated to the required wordlength without re-optimization. Even so, the results were very close to the floating-point version. It should be noted here that these are not coefficients of the transfer function, but of the second and fourth-order sections into which the filter has been decomposed. Constraining the coefficients of the filter with arbitrary cut-off frequency makes the problem similar to constraining the taps of the general transfer function of an IIR. Hence the performance is much more sensitive to shortening of the coefficient wordlength. The simulation result showed that even for short coefficient wordlength, such as 6-bits, the responses were very close to the floating-point one. These results further suggest that for the halfband case implementing the allpass filter from Figure 1 as the cascade of smaller order allpass sections makes the coefficients much less sensitive to limiting their wordlengths.

4 Conclusion

We have presented in this paper two methods of designing lowpass/highpass filters based on the polyphase IIR structure achieving almost linear phase response.

One was to compensate the non-linear phase of the standard polyphase filter with a combination of one-coefficient allpass filters with different order. This type of correction of the phase linearity allows achieving very high improvement ratios. For the case of fractional passbands the even orders of the allpass correction sections make them suitable for operating at a lower rate in multi-rate systems.

The second approach was to use Curtis [1][4] flat group delay structure. An allpass filter in one branch was designed to follow the linear phase response of the bulk delay placed in the other branch. It has been achieved with weighted least squares algorithm with custom weight optimisation. This approach was compared to similar work by Lu and Lawson, and proved advantageous to them in terms of the required filter order, stopband attenuation and achieved group delay flatness. It was also shown that the implemented structure allows very efficient coefficient quantization, especially for the case of the halfband filter. This makes it very applicable for fixed-point implementation in very high speed filtering applications like ΣΔ-decimation filters and other applications, which necessitate a high level of phase or group delay flatness.

References:

[1] Renfors, M. and T. Saramäki, "Recursive Nth-band digital Filters - Part II: Design of multistage decimators and interpolators," *IEEE Transactions on Circuits and Systems*, vol. CAS-34, pp. 40-51, January 1987.

[2] harris, f., M. d'Oreye de Lantremange and A. G. Constantinides, "Digital signal processing with efficient polyphase recursive all-pass filters", *International Conference on Signal Processing*, Florence, 4-6 September 1991.

[3] Valenzuela, R. A. and A. G. Constantinides, "Digital signal processing schemes for efficient interpolation and decimation", *IEE Proceedings*, Pt. G., vol. 130, no. 6, pp. 225-235, December 1983.

[4] Curtis, T. E. and A. B. Webb, "High performance signal acquisition systems for sonar applications", *International Conference on A/D and D/A Conversion*, Swansea, 17-19 September 1991.

[5] Lawson, S., "On design techniques for approximately linear phase recursive digital filters", *Proceedings International Conference on Circuits and Systems (ISCAS'97),* pp. 2212-2215, June 1997.

[6] Lu, W. S., "Design of stable IIR digital filters with equiripple passbands and peak-constrained least squares stopbands", *Proceedings International Conference on Circuits and Systems (ISCAS'97),* pp. 2192-2195, June 1997.

[7] Lawson, S. S., "Direct approach to design of PCAS filters with combined gain and phase specification", *IEEE Proceedings on Visual Image Signal Processing*, vol. 141, pp. 161-167, 1994.

[8] Kale, I., R. C. S. Morling and A. Krukowski, "A high-fidelity decimator chip for the measurement of Sigma-Delta modulator performance", *IEEE Transactions on Instrumentation and Measurement*, vol. 44, no. 5, October 1995.

[9] Krukowski, A. and I. Kale, "Constrained coefficient variable cut-off polyphase decimation filters for band-pass Sigma-Delta data conversion" *IMEKO Workshop on ADC Modeling*, House of Scientists, Smolenice Castle, 7-9 May 1996.

[10] Krukowski, A., I. Kale, K. Hejn and G. D. Cain, "A bit-flipping approach to multistage two-path decimation filter design", *Second International Symposium on DSP for Communication Systems*, SPRI, 26-29 April 1994.