Equalization of Data Acquisition Channels Using Digital Filters

I. Kollár*, R. Pintelon**, Yves Rolain** and J. Schoukens**

*Department of Measurement and Instrument Engineering
Technical University of Budapest, Hungary

**Department ELEC
Vrije Universiteit Brussel, Belgium

Received July 2, 1990.

Abstract

The possibility of the compensation of amplitude and phase errors of data acquisition channels is investigated. The complex passband error can be significantly reduced by an appropriate digital filter, as it is illustrated on the example of a dynamic signal analyzer. The aspects of IIR and FIR equalizers are studied: it turns out that for the given anti-aliasing filter (Cauer filter, order 11) an FIR filter of length 60...100 can perform as well as a 26/26 IIR one. Because of the absence of stability problems and the ease of implementation, the use of FIR filters is suggested.
Keywords: equalization, compensation, signal conditioning, data acquisition channel, digital filter design, approximation in complex domain.

1 Introduction

Digital signal processing works on sampled equivalents of continuous-time signals. In order to avoid aliasing, components over the half of the sampling frequency have to be suppressed. This is usually done by an anti-aliasing filter. A perfect anti-aliasing filter should not alter the signal in the passband, which means that the absolute value of the transfer function is constant, and that the

Published in Periodica Polytechnica Ser. Electrical Engineering, Vol. 34, No. 3, pp. 167-178, 1990.
The research has been supported by the Belgian National Fund for Scientific Research (NFWO). It was accomplished during a sabbatical leave of Dr. Kollár to the Department ELEC, Vrije Universiteit Brussel, Belgium.

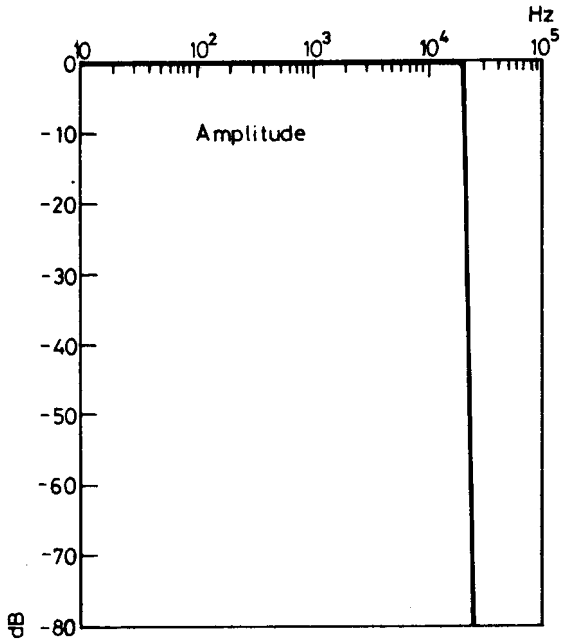

Figure 1: Transfer function of an anti-aliasing filter (Cauer filter, order: 11)

phase is linear. On the other hand, the attenuation has to be sufficiently high outside the passband. A transition band is usually tolerable between the passband and the stopband, if the maximum passband frequency is smaller than the half of the sampling frequency, thus the anti-aliasing filter need not exhibit a discontinuity at the edge of the passband.

Theoretically, such a filter is realizable with arbitrary precision, if ideal (or nearly ideal) components can be used, and the complexity is not limited. However, in practice this is not the case. A reasonably steep transition band behaviour and a good stopband attenuation (e.g. >80–100 dB) can be easily realized, but both the amplitude and phase of the filter will usually exhibit variations in the passband. Even if it were possible to design a nearly ideal filter, component tolerances would spoil the performance. In principle it is possible to adjust some elements, but this is usually very cumbersome and expensive (one adjustment step impairs the effect of others, consequently an iterative procedure is to be followed; it is often difficult to define the concrete goal of a single adjustment step). Moreover, the adjustable components may deteriorate the long-term stability of the transfer function (except if laser trimming is used), and this is undesirable again.

Because of these facts, the filter order is usually chosen not higher than 10–15, which limits the performance.

A typical example is shown in *Fig. 1.* This is the transfer function of an audio

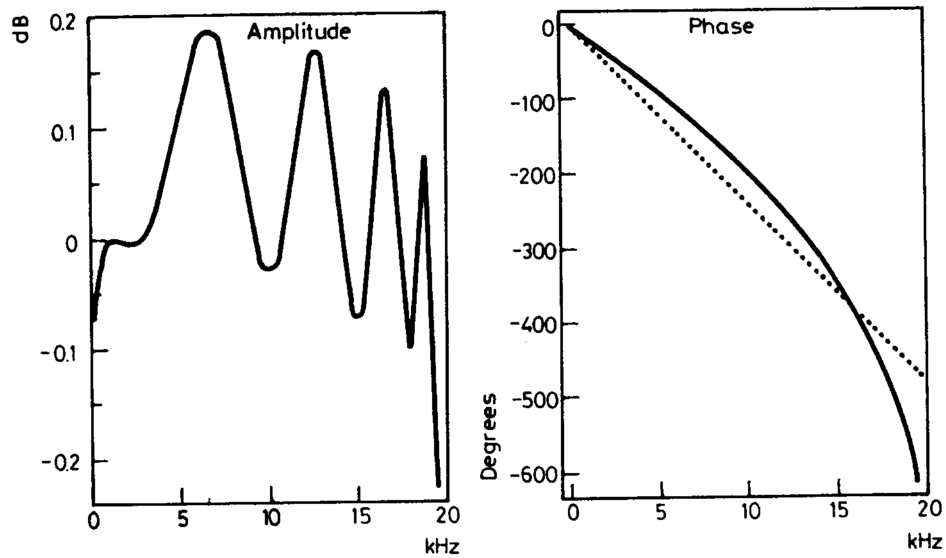

Figure 2: Passband of the anti-aliasing filter

band Cauer filter of order 11, produced with laser trimmed thick film technology. The shape seems to be perfect, however, a closer look at the passband (*Fig. 2*) shows that, while the transition band is very steep, the passband frequency response exhibits substantial errors (about $\pm\, 0.2\,\text{dB}$ ripple in the magnitude, the maximum phase deviation from an LS-fitted straight line is about 50° in the middle, and much worse near 20 kHz).

If the non-ideal behaviour of the filter is known, we can take this into consideration when designing the digital data processing algorithm. But the speed of modern digital signal processors offers another, even more attractive solution: the data series can also be corrected by a digital filter in real time *before* any other processing. In the following sections we are going to deal with the aspects of this correction and with the methods that can be applied in the design of the digital filters. As an example, equalization of the input channel of a dynamic signal analyzer will be treated (DSA710, for details see Pintelon et al, 1990a).

2 Design of the Correction Filters

The requirement is clear: the digital filter should correct the anti-aliasing filter for gain and phase in the passband, without introducing too large gains elsewhere and, last but not least, it must be realizable. The latter means that it must be safely stable, and possibly of low order, to be implemented at low cost. For the design procedure an effective algorithm is needed which does not demand too much computing time, and can be performed possibly without human intervention.

The first task of the procedure is the accurate measurement of the transfer function: the measurement error must be reasonably smaller than that required after compensation. This measurement is often to be done as a one-channel one (with calibrated input signal). We will not dwell here much on this question, a good solution is given by Pintelon et al (1990a), using a low crest factor multisine as an excitation signal.

Once the transfer function is known on a sufficiently dense grid, the design can be started. We will try to reduce the relative complex error:

$$e_{\mathrm{dB}} = 20 \cdot \log\big((|tf| + |e|)/|tf|\big) ,$$

where tf is the transfer function and e is the complex error of the fit at a given frequency.

2.1 A direct approach (IIR filter)

The most straightforward approach is to identify the transfer function in a parametric form, and invert the z–domain equivalent (Pintelon et al, 1990a). Since usually both the input and the output measurements are noisy, an algorithm prepared for this situation is to be used (see e.g. Pintelon and Schoukens, 1990b). The correct solution is to make an s–domain identification, and use the obtained smooth transfer function for the further digital filter design steps. However, in the investigated case we found that the identification may as well be done directly in the z–domain. A good fit has been achieved using an IIR filter of order 14/14, which is reasonable, since the identification extends also to the signal conditioning amplifier, thus the order 11 of the filter itself should be increased, and the resting 1–2 poles can be explained by the s–domain/z–domain mapping (Pintelon et al, 1990a). The relative complex error of the fit was everywhere less than 6.1 mdB (*Fig. 3*). This error is the sum of the random errors of the measured data and the modelling error. The detailed analysis of the modelling error needs a more thorough analysis. For the purpose of this study we assumed that the modelling error is less than 3 mdB (see Pintelon et al, 1990a), and demanded from the further fittings not to introduce more additional complex error than 7 mdB, with respect to the 14/14 fit. The worst-case total compensation error (10 mdB) corresponds to 0.12% amplitude error or 0.066° phase error.

In this procedure, two difficulties were encountered:

a) since a complex function is to be fitted, the passband behaviour contains information concerning the stopband, too, especially because near 20 kHz the filter already starts to attenuate. This means that the correction filter will try to invert also a part of the transition band (and maybe of the stopband), having a large stopband gain. Fortunately, in our very case the edge of the passband is close to the half of the sampling frequency (25.6 kHz), which resulted in a limitation of the gain: in the transition band the correction filter begins to compensate the rolloff, but its maximal gain is not larger than about 30 dB (see *Fig. 4*).

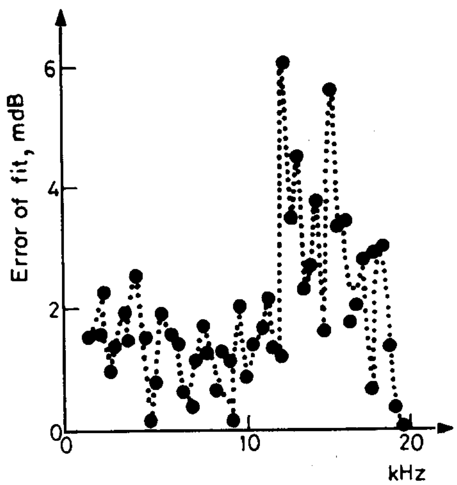

Figure 3: Complex relative error of the 14/14 fit

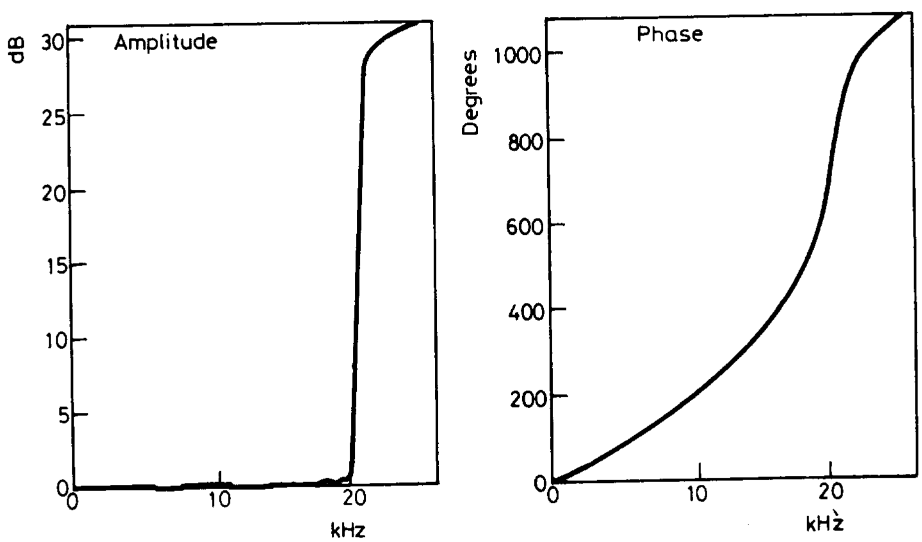

Figure 4: Magnitude and phase response of the inverted z–domain filter

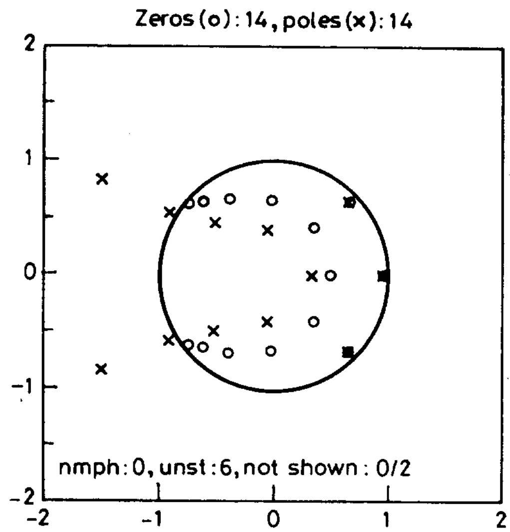

Figure 5: Pole-zero pattern of the inverted z–domain filter

b) Although the Cauer filter is theoretically minimum in phase, the z–domain equivalent not necessarily preserves this property (or at least not at high frequencies), thus the inverse may be unstable (*Fig. 5*).

Pintelon et al (1990a) solved the problem of instability by inversion of the unstable poles. This procedure does not affect the amplitude response, but spoils the phase fit. The phase can then be corrected by using additional allpass sections, which do not alter the amplitude response, either. By using a phase corrector of order 20/20 (fitted in complex minimax sense), the complex error could be reduced below 7 mdB.

Though the design of the correction filter is not trivial (we have to minimize the phase error by a *stable* allpass filter), with some experience and experimentation an appropriate value of the allowed delay can usually be chosen, which results in a stable fit (Pintelon, 1990).

2.2 Decreasing the order

It is obvious that the above described very straightforward method is not globally optimal, since the two separate steps had different goals. This inspired us to look for a stable, good fit with an order lower than 34 (Kollár et al, 1991). However, our attempts to find such a fit using the available passband information failed, because even by using the same delay as for the phase corrector, some of the poles were always unstable (usually near to the edge of the passband). We also observed very different stopband curves, which is not surprising since these are extrapolated from the noisy passband data. Thus, we decided

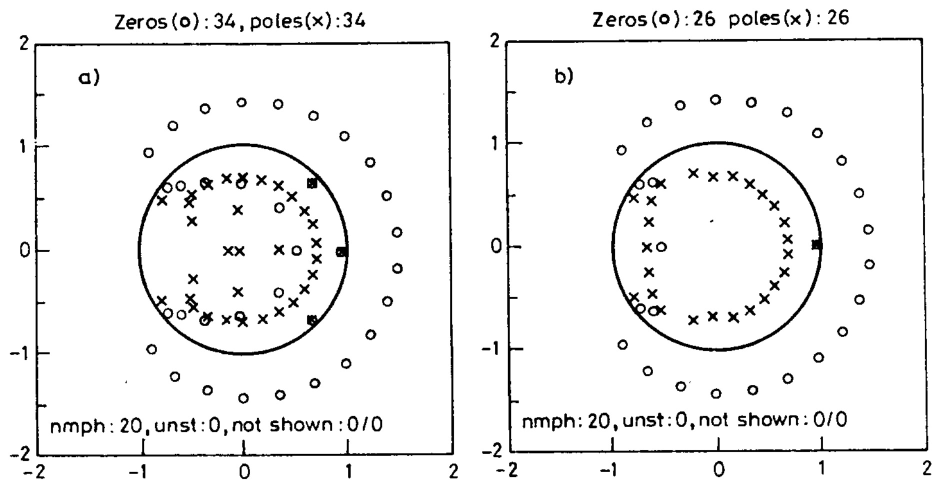

Figure 6: Pole-zero patterns of the compensating filters: *a)* The 34/34 filter (stabilized 14/14 fit, cascaded with a 20/20 phase equalizer); *b)* The reduced-order filter (26/26)

to provide *stopband information.* We created a new data set by taking samples of the frequency response of the well fitted (but unstable) 14th order filter, and tried to fit this by filters of order less than 34/34, but with the same delay, that is, by the form $z^d H(z)$. After some playing with the weights (the stop-band need not be fitted very exactly, moreover with uniform weighting the error around zero was somewhat larger than elsewhere), a 26/26 filter was obtained with practically the same fitting error as the 34/34 one. Even a 24/24 filter could be designed with less than 10 mdB fitting error, and a 22/22 with 25 mdB fitting error. It may be interesting to throw a glance onto the pole/zero patterns (*Fig. 6*): the poles far from the unit circle and two of the three pole/zero pairs (almost coinciding pole and zero) disappeared, and their effect was substituted by a slight rearrangement of the rest. The pole/zero pairs were also present already in the original 14/14 filter, but there the order could not be decreased without a significant increase in the error: this filter has not enough freedom to 'simulate' the effect of these poles and zeros, if they are missing. In the 24/24 fit the last pole/zero pair also disappears, but this is already at the cost of a slight increase in the error.

With the above technique we did not succeed in further reducing the order without a significant increase in the error. However, it still remains an open question whether even lower order stable filters exist with similar performance. We incline to think that such filters exist, but at present we do not know any reasonable method to find them.

2.3 FIR filter

Although the above filters were implemented and worked well, the whole procedure is rather intricate and time consuming, furthermore it needs human interaction, first of all during the allpass design. If an FIR filter could be used, there would be no stability problem at all, and what is more, the implementation would also be much simpler. The widespread opinion (or rather: prejudice, as we are going to notice, see Kollár and Rolain, 1993) is that FIR filters which have the same performance as a good IIR filter, have a very large order, thus they are of no practical use. But the absence of stability problems is so attractive, that it is worth a trial.

Before starting the design, three important aspects have to be discussed. First, since we are going to use a high-order FIR filter, the degrees of freedom will increase significantly. This means that we need a dense enough grid of the desired passband, to avoid obtaining a good fit on a coarse grid, while in between the fit is poor. For the above fits we used about 50 complex amplitudes in the total audio band, providing $2 \cdot 50 = 100$ equations. For the allpass design the point density was already increased by a factor of 2, and for the FIR fit by a factor of 4.

Second, the increased degrees of freedom may give rise also to large stopband gains. This can be avoided by artificially introducing zero amplitudes in the stopband (the grid must be dense enough again, otherwise the algorithm will just put zeros to each zero amplitude), with a reasonably small weight, to allow large deviations from zero, compared to the desired small passband error.

Third, the fit is to be performed in the complex domain. Unfortunately enough, the well-known, very effective Parks–McClellan (Remez exchange) algorithm (McClellan, Parks and Rabiner, 1979) works for the real (or imaginary) case only, producing a linear phase FIR filter, fitted in minimax sense. For the complex case there is still no standard algorithm. Ellacott and Williams (1976) used the reweighted LS technique (RWLS). Chen and Parks (1987) transformed the complex domain problem into an approximately equivalent linear programming one, and solved the latter using Algorithm 495. Leeb and Henk (1989) used Padé approximant techniques in an iterative scheme, based on the Remez exchange algorithm. Nagy (1990) used the Remez exchange algorithm to fit the real and imaginary parts of the transfer function separately, combined them to a nonlinear phase FIR filter (this brought his error under $\sqrt{2}$ times the optimum one), and used an iterative scheme to reach the minimax fit. It is also possible to use the minimum p–error method of Deczky (1972) for the nonrecursive case.

The different algorithms should be compared systematically with respect to their speed, memory need, convergence for high orders etc. before making a choice. However, since we had ready programs for the WLS fit (the ELiS algorithm, described in (Pintelon, 1990b), can also be used for this purpose), we used a modified version of the RWLS technique. It turned out that an FIR filter of order 56 (57 taps) has fitting errors smaller than 7 mdB, and a filter of length 49 had a smaller error than 10 mdB. The filters are easily realizable (see *Fig. 7*), with acceptable stopband gain. It is interesting to observe that, because

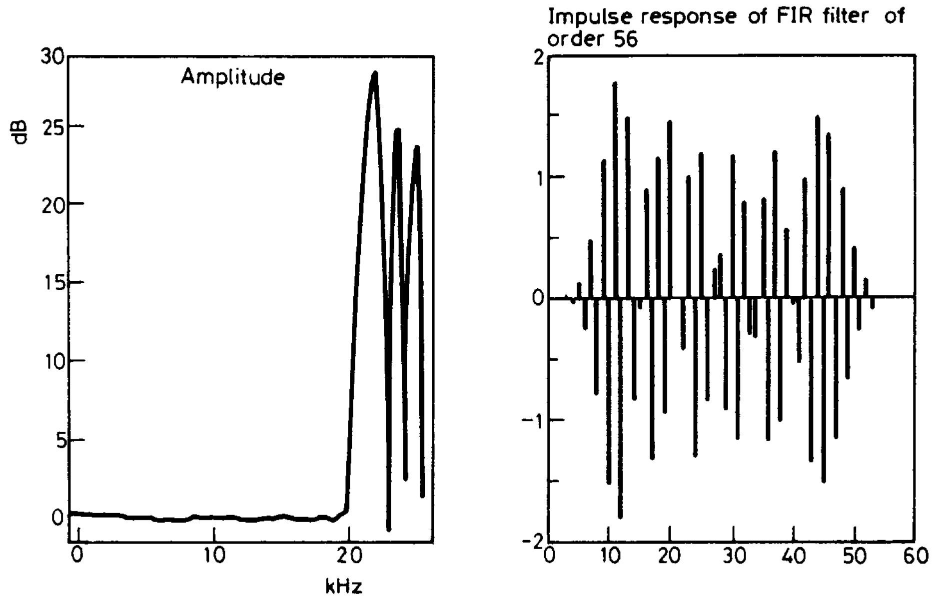

Figure 7: Amplitude response and impulse response of the length 56 FIR filter

of the 30 dB stopband gain, the impulse response has a highpass character.

By the experimentation with other anti-aliasing filter boards of the same type it turned out that though their global behaviour is very similar, sometimes the FIR order, necessary for 7 mdB fitting error, is much larger (100). Fortunately, this number is still acceptable to be realized in real time ($f_s = 51.2\,\text{kHz}$). We observed that in these cases the initial fit is definitely worse at low frequencies. This may be due to the different behaviour below 1 kHz (*Fig. 8*), which is in our experience the consequence of an impedance mismatch at the input of the anti-aliasing filter. Nevertheless, the IIR method still works fine with the same order (its error is not more than 7 mdB), thus we speculated that the worse FIR fit is due to a pole very close to the unit circle. Indeed, the pole/zero pair near $z = 1$ in the inverse filter (see *Fig. 5*) is in the first case $p_1 = 0.946855$, $z_1 = 0.946290$, while in the second case $p_2 = 0.967758$, $z_2 = 0.968282$, that is, the second pole is significantly closer to the unit circle than the first one:

$$(1 - p_1)/(1 - p_2) \approx 0.053/0.032 = 1.65\,.$$

The closer a pole to the unit circle, the more zeros are necessary to define a similar transfer function.This seems to put a limit to the applicability of the FIR filters for compensation. However, if necessary, also a 'mixed' strategy can be applied: when the resulting FIR order is too high, and there are stable poles close to the unit circle, one should realize the corresponding pole/zero pairs in first-order IIR sections, and do the rest of the compensation by an FIR filter.

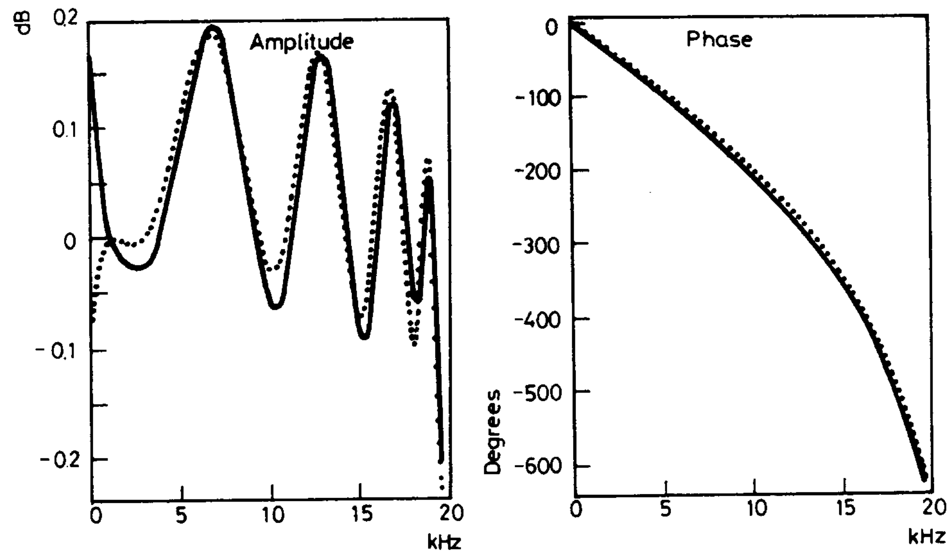

Figure 8: Comparison of the passband behaviour of two anti-aliasing filters

3 Conclusions

It has been illustrated on the example of a dynamic signal analyzer that it is possible to improve the quality of data acquisition channels by means of digital correction filters. The complex error has been significantly reduced: from 200 mdB magnitude error and $50-100^{\circ}$ phase error (see *Fig. 2*), to 10 mdB complex error (max. 10 mdB magnitude error or max. 0.066° phase error). The possibilities of the use of IIR and FIR equalizers have also been investigated: it turned out that the FIR filter is preferable because of the absence of stability problems, and the ease of the implementation, at least if the desired error reduction can be achieved by an acceptable filter length.

References

Chen, X. – Parks, T. W. (1987): Design of FIR Filters in the Complex Domain. *IEEE Transactions on Acoustics, Speech and Signal Processing*, Vol. 35, No. 2, pp. 144–153. Feb. 1987.

Deczky, A. G. (1972): Synthesis of Recursive Digital Filters Using the Minimum p–Error Criterion. *IEEE Trans. on Audio and Electroacoustics*, Vol. 20, pp. 257–263. Oct. 1976.

Ellacott, S. – Williams, J. (1976): Linear Chebyshev Approximation in the Complex Plane Using Lawson's Algorithm. *Mathematics of Computation*, Vol. 30, No. 133, pp. 35–44. Jan. 1976.

Kollár, I. – Pintelon, R. – Rolain, Y. – Schoukens, J. (1991): Another Step Towards an Ideal Data Acquisition Channel. *IEEE Transactions on Instrumentation and Measurement*, Vol. 40, No. 3, pp. 659–660. June 1991.

Kollár, I. – Rolain, Y. (1993): Design of FIR Equalizer Filters for Complex Correction of Data Acquisition Channels. *IEEE Transactions on Instrumentation and Measurement*, Vol. 42, No. 5, Oct. 1993. pp. 920-24.

Leeb, F. – Henk, T. (1989): Simultaneous Amplitude and Phase Approximation for FIR Filters. *International Journal of Circuit Theory and Applications*, Vol. 17, pp. 363–374.

McClellan, J. H. – Parks, T. W. – Rabiner, L. R. (1979): FIR Filter Design and Synthesis. In: *Programs for Digital Signal Processing*, New York, IEEE Press, pp. 5.1-1–5.1-13.

Nagy, F. (1990): Signal Conditioning and Reconstruction in Telemetry Systems by Means of Digital Filtering. *Proceedings of the European Telemetry Conference, ETC'90*, Garmisch-Partenkirchen, May 14–17, 1990. pp. 132–139.

Pintelon, R. (1990): Phase Correction of Linear Time Invariant Systems with Digital Allpass Filters. *IEEE Transactions on Instrumentation and Measurement*, Vol. IM–39, No. 2, pp. 324–330. April 1990.

Pintelon, R. – Rolain, Y. – Vanden Bossche, M. – Schoukens, J. (1990a): Towards an Ideal Data Acquisition Channel. *IEEE Transactions on Instrumentation and Measurement*, Vol. IM–39, No. 1, pp. 116–120. Feb. 1990.

Pintelon, R. – Schoukens, J. (1990b): Robust Identification of Transfer Functions in s– and z–Domain. *IEEE Transactions on Instrumentation and Measurement*, Vol. IM–39, No. 4, pp. 565–573. Aug. 1990.