

ADAPTIVE DESIGN OF DIGITAL FILTERS

Bernard Widrow, Paul F. Titchener, and Richard P. Gooch

Information Systems Laboratory
Stanford University
Stanford, California 94305

ABSTRACT

In this paper, we present a novel technique for the design of FIR and IIR digital filters. The design approach begins with the specification of a discrete set of arbitrary magnitude and phase characteristics which describe a desired filter response. These frequency domain characteristics are used to create an ideal "pseudo-filter" whose impulse response is unknown and possibly non-causal, but whose input/output characteristics can be determined for a finite sum of sinusoids. Time-domain techniques common to adaptive system identification are then used to identify a realizable FIR or IIR digital filter which best matches the pseudo-filter. The advantages of this method include the ability to specify response at arbitrarily-spaced frequencies, to use arbitrary cost weighting, and to apply (possibly non-linear) constraints to the range of the filter coefficients.

I. INTRODUCTION

Many techniques are available for the design of FIR and IIR digital filters. In this paper we present a new method that we call the Pseudo Filter Design Technique (PFDT). We feel this technique is substantially different from any technique currently in the literature, for it allows a degree of flexibility in the design procedure that is unmatched by previous works.

It can be shown that the PFDT encompasses the Frequency Sampling FIR design of Rabiner et al. [1]. There is also a similarity between the statistical IIR design method of Sharf and Luby [2] and the PFDT since both solutions involve second order statistics. The statistical method involves analytical expressions for the desired frequency response and its weighting function. For arbitrary filter designs, analytical expressions may be hard to specify. The PFDT is more flexible since it involves discrete frequency and weighting specifications. Another commonly used IIR filter design method is that of Burrus and Parks [3]. There, an IIR filter is matched to a high order FIR filter that has the desired frequency response. This restricts somewhat the specification process, and in essence requires the design of two filters for each final product.

II. TECHNIQUE DESCRIPTION

The PFDT is best explained through concepts common to adaptive system identification. Figure 1 shows the basic system identification structure. The idea is to select a model and determine the model parameters which describe or approximate an unknown system (commonly referred to as the plant). In most situations, the model cannot exactly match the plant, so an optimality criterion is established. The criterion

usually involves the minimization of the squared error between the plant and model outputs (output error), or the squared error between the plant and model equations (equation error).

The connection between the above example and the PFDT is seen in Figure 1. The unknown plant is replaced by an ideal filter which represents the desired characteristics. The model is replaced by the digital filter to be designed, referred to in this paper as the modeling filter. It represents an FIR or IIR digital filter whose order is specified by the user.

Our technique involves generating the ideal pseudo filter based on a set of magnitude and phase specifications at discrete frequencies which may or may not be uniformly spaced. A filter possessing these characteristics may have an impulse response which is non-causal and therefore non-realizable. For this reason we call the ideal filter a "pseudo-filter". Even though the pseudo-filter is non-realizable, it will have input/output characteristics which are known for an input consisting of a sum of sinusoids whose frequencies are those of the specification set. The goal is to have the modeling filter (the designed filter) match the pseudo-filter at those frequencies. Typically the number of frequency specifications is large compared to the degrees of freedom in the modeling filter (number of poles and zeroes), and it will be impossible to make an exact match between the pseudo-filter and the modeling filter. In this case, the modeling filter is chosen so as to minimize the squared error between the outputs of the two filters.

In Figure 2 it is seen that the input excitation consists of a weighted sum of sinusoids at the specification frequencies. The effect of selectively increasing the cost of the weighting parameters is to cause a tighter fit between the pseudo-filter and the designed filter at the frequencies selected. Otherwise, the cost in the least squares sense of misfit at these frequencies would be high.

Figure 2 includes a block diagram of the specific structure investigated in this paper. The modeling filter would be a true IIR filter if the pseudo-filter output were replaced by the model filter output as indicated by the dashed line. The error generated would then be the output error between the IIR modeling filter and the pseudo-filter. This is not done since minimization of this output error results in a highly nonlinear problem. The configuration shown generates an error related to the difference between the model and the pseudo-filter equations and is thus referred to as an equation error method. The equation error is a linear function of the modeling filter coefficients, and therefore minimization of the squared error is a quadratic minimization problem.

This work was supported by NSF under grant No. ENG 78-08526 and by ONR under contract No. N00014-76-C-0929.

This structure is commonly used in the field of system identification, and is a basic component of many of the more elaborate structures that system identification researchers have developed [4]. One drawback is that there is no guarantee the resulting filter will be stable, i.e., will have all of its poles inside the unit circle. Methods for insuring filter stability will be developed and discussed in sections III and IV.

III. MATHEMATICAL FORMULATION

In this section we show how to find the modeling filter coefficients which best match the pseudo-filter. Referring to Figure 2, we make the following definitions. The input to both the pseudo-filter and the modeling filter is given by

$$u_t = \sum_{k=1}^N c_k \sin w_k t \quad (1)$$

where c_k represents the cost weighting and w_k represents the radian frequencies at which the ideal filter transfer characteristics are specified. The quantities N , c_k , and w_k , ($k=1, \dots, N$) are free to be selected by the user. The output of the pseudo-filter is given by

$$y_t = \sum_{k=1}^N c_k m_k \sin(w_k t + \varphi_k) \quad (2)$$

where m_k and φ_k are the magnitude and phase specifications determining the desired filter characteristics. The modeling filter output is

$$\hat{y}_t = B_t * u_t - A_{t-1} * y_t \quad (3)$$

where B and A represent the zero and pole coefficients with order q and p , respectively. Note that if y_t is replaced by \hat{y}_t , Eq. (3) would represent a general IIR filter. The justification for writing Eq. (3) as it stands, is that \hat{y}_t and y_t should be approximately equal. The error is given by

$$e_t = y_t - \hat{y}_t \quad (4)$$

$$= (\delta_t + A_{t-1}) * y_t - B_t * u_t$$

Note that this error is linear in the A and B parameters. This is important when minimizing the squared error since the solution is found simply by solving a set of linear equations as will now be shown.

We now wish to find the coefficients of A and B which minimize the squared error criterion.

$$\underset{A, B}{\text{Minimize}} \left\{ \sum_{t=-\infty}^{\infty} e_t^2 \right\} \quad (5)$$

This can be restated in matrix notation as

$$\underset{\Theta}{\text{Min}} \left\{ \Theta^T R \Theta - 2P^T \Theta \right\} \quad (6)$$

where

$$R = \begin{bmatrix} R_{uu} & R_{uy} \\ R_{yu} & R_{yy} \end{bmatrix}, \quad P = \begin{bmatrix} P_{yu} \\ P_{yy} \end{bmatrix}, \quad \Theta = \begin{bmatrix} B \\ A \end{bmatrix} \quad (7)$$

Let (i, j) represent the row and column of the above submatrices. Then

$$R_{uu}(i, j) = \sum_{k=1}^N c_k^2 \cos[w_k(i-j)t] \quad 1 < i, j < q \quad (8)$$

$$R_{yy}(i, j) = \sum_{k=1}^N c_k^2 m_k \cos[w_k(i-j+1) - \varphi_k] \quad 1 < i < q, \quad 1 < j < p$$

$$R_{yu}(i, j) = R_{uy}(j, i) \quad 1 < i < p, \quad 1 < j < q$$

$$R_{yy}(i, j) = \sum_{k=1}^N c_k^2 m_k^2 \cos[w_k(i-j)] \quad 1 < i, j < p$$

and

$$P_{yu}(i) = \sum_{k=1}^N c_k^2 m_k \cos[w_k i - \varphi_k] \quad 1 < i < q$$

$$P_{yy}(i) = \sum_{k=1}^N c_k^2 m_k^2 \cos[w_k(i+1)] \quad 1 < i < p$$

and

$$A(i) = -a_i \quad 1 < i < q$$

$$B(i) = b_i \quad 1 < i < p$$

The solution to this minimization is found by solving the following linear equations for Θ^* .

$$R \Theta^* = P \quad (9)$$

The effect of hard linear constraints is to restrict the coefficient vector Θ to some subspace defined by the underdetermined set of linear equations $C \Theta = d$. Thus the problem becomes,

$$\underset{\Theta}{\text{Min}} \left\{ \Theta^T R \Theta - 2P^T \Theta \right\} \quad \text{subject to } C \Theta = d \quad (10)$$

The solution to this problem can be found using Lagrange multiplier techniques and is given by

$$\Theta_{hc}^* = [I - R^{-1}C(C^T R^{-1}C)^{-1}C^T]R^{-1}P \quad (11)$$

$$+ R^{-1}C[C^T R^{-1}C]^{-1}d$$

Hard constraints can be used to design linear phase FIR filters. A sufficient condition for a filter to have linear phase is that its coefficients be symmetric about their midpoint. This can be accomplished using linear constraints by letting $C = [I \ I]$, and $d = 0$ where I is a reversed identity matrix.

Similarly, soft constraints could be applied to the filter coefficients. Unlike hard constraints, the soft constraints allow the coefficients to deviate from the constraint subspace with a resulting increase in cost. This problem is written as

$$\underset{\Theta}{\text{Min}} \left\{ \Theta^T R \Theta - 2P^T \Theta + (C^T \Theta - d)^T B (C^T \Theta - d) \right\} \quad (12)$$

where the constraining subspace is $C \Theta = d$. The solution to this problem is

$$\Theta_{sc}^* = [R + CBC^T]^{-1} [P + CBD] \quad (13)$$

where B represents a weighting matrix usually chosen as γI .

Nonlinear constraints may also be applied to the filter coefficients. However, when nonlinear constraints are applied, a closed form solution may not exist and iterative methods for finding a solution must be used.

The unconstrained problem involves the solution of a set of linear equations with a matrix that has a block Toeplitz structure. The solution can be found efficiently by the use of the Levinson algorithm. The equations can be solved with relative ease since only $O_r[(p+q)^2]$ computations and $O_r[(p+q)]$ memory locations are required.

In some cases when designing a linearly constrained high order filter on a small computer, it may be easier to implement a recursive method for finding the filter coefficients. The LMS adaptive algorithm [5] requires $O_r[p+q]$ memory and can often be easily constrained and applied to these problems. Computational

time may be greater, but programming efforts and memory requirements will be relaxed.

IV. FILTER DESIGN IMPLEMENTATION

One starts the design process by forming a discrete set of psuedo-filter magnitude and phase specifications that describe the response of the filter to be designed. The number and location of these specifications can be freely chosen by the designer to best describe the filter response desired. A cost weighting is associated with each magnitude specification. Typically this cost is initially set to one.

The FIR case corresponds to the designer selecting the number of poles to be zero ($p=0$), and choosing the number of zeros ($q-1$). The solution then gives the set of coefficients for the filter B . After viewing the response of the filter B , the designer can increase the cost weighting at certain specification points where a tighter fit is desired. Additional specification points may also be introduced to further improve the designed filter response, and a final design is arrived at in an interactive manner.

In the IIR design case, the problem of unstable poles must be handled. If the designer is only concerned with the magnitude response of the final filter (phase response is unimportant), the pole reflection technique can be used. It is commonly known that the poles of an unstable filter can be reflected inside the unit circle, yielding a stable filter with identical magnitude response, but altered phase response. Therefore to stabilize the designed filter, the unstable poles are reflected inside the unit circle.

In the IIR design case when the phase of the designed filter is important, the designer has two alternatives. The first is to make the chosen phase specification points reasonable for a stable IIR filter that has the chosen magnitude specification points. Depending on the filter to be designed, this may or may not be easy to do. For the standard filter types such as lowpass and highpass, the characteristic phase responses of stable filters are well known and can be used as guidelines for the chosen phase specifications. When designing filters that have a more general magnitude response, some thought and experimentation may be necessary to pick phase specifications that will result in stability of the designed filter, and a final stability check is needed.

The second alternative, when strict adherence to design phase characteristics is important, is to place constraints on the A filter coefficients. Since the polynomial that represents the denominator of the transfer function of the designed filter can, without loss of generality, begin with a unit constant coefficient, constraining its higher order coefficients to be increasingly smaller will tend to pull the poles of the filter in towards the origin. Thus a soft linear constraint on the A filter coefficients can be used to "persuade" the designed filter to be stable. Alternatively, a nonlinear constraint on the A filter coefficients that forces the filter poles to remain inside the unit circle can be implemented. This requires the use of some recursive solution technique, such as a suitably constrained version of the LMS adaptive algorithm.

V. SIMULATION RESULTS

Figures 3a,b and 4 are the results of some filter design simulations. The cross marks in the figures

represent the magnitude, phase and cost weighting points specified for the desired (psuedo) filter, and the continuous curves represent the response of the resulting filter design. Figures 3a,b demonstrate the design of a 50 weight FIR filter that has a magnitude response that initially increases and then sharply cuts off. A linear phase response was specified. The virtuosity of the technique is clear. This is no mere Butterworth or Chebyshev design. Figure 3b shows how increasing the cost weighting at some of the specification points can tighten the fit (in both magnitude and phase) of the designed filter.

Figure 4 shows a 9 pole, 9 zero IIR lowpass filter design. For this simple example, the phase specification points were chosen so as to result in a stable filter. The design provides flat pass band response, sharp cutoff, and close adherence to a -50 db gain in the stop band.

VI. CONCLUSION

In this paper we have developed a new technique for the design of FIR or IIR digital filters. It has been shown that the technique is easy to implement and possesses flexibility in the design process. These two traits make the PFDT a desirable addition to any digital filter design package.

The general system identification framework on which this technique is based suggest the possibility of utilizing more elaborate identification structures. Methods for implementing nonlinear constraints on the filter poles are currently undergoing development. One such method involves the use of either second order sections or a lattice filter in order to maintain filter stability. Techniques for implementing an adaptive cost weighting are being investigated. The application of output error minimization techniques are also being researched.

REFERENCES

1. L.R. Rabiner, B. Gold, and C.A. McGonegal, "An Approach to the Approximation Problem for Nonrecursive Digital Filters," *IEEE Trans. Audio Electroacoust.*, Vol. AU-18, pp. 83-106, June 1970
2. L.L. Sharf and J.C. Luby, "Statistical Design of ARMA Digital Filters", *IEEE Trans. ASSP*, Vol. 27, pp.240-247, June 1979
3. C.S. Burrus and T.W. Parks, "Time domain design of recursive digital filters," *IEEE Trans. Audio Electroacoust.*, Vol. AU-18, pp. 137-141, June 1970
4. T. Soderstrom, L. Ljung, and I. Gustausson, "A Theoretical Analysis of Recursive Identification Methods", *Automatica*, Vol. 14, pp. 231-244, 1978
5. B. Widrow et al., "Stationary and Nonstationary Learning Characteristics of the LMS Adaptive Filter," *Proc. IEEE*, Vol. 64, No. 8, pp. 1151-1162, August 1976




Fig. 1 Basic Design Approach

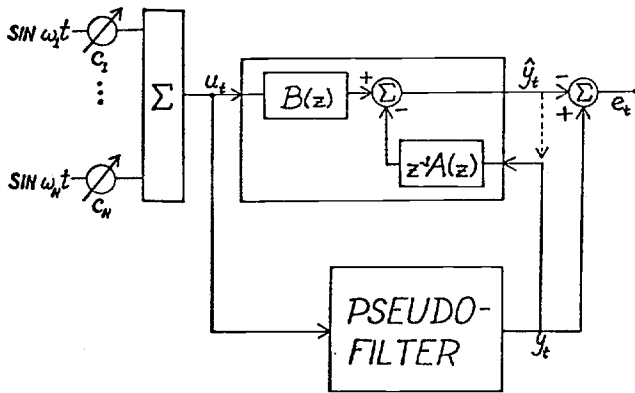


Fig. 2 Detailed Design Structure

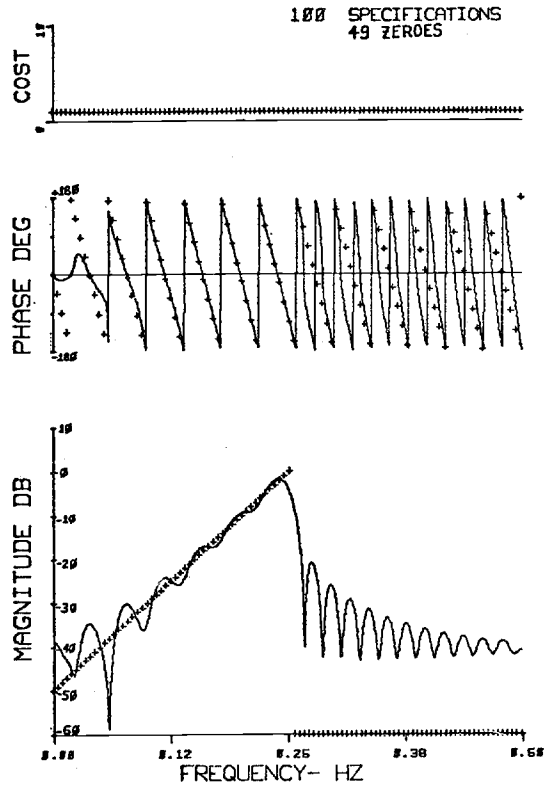


Fig. 3a Initial FIR design




Fig. 4 IIR Lowpass design

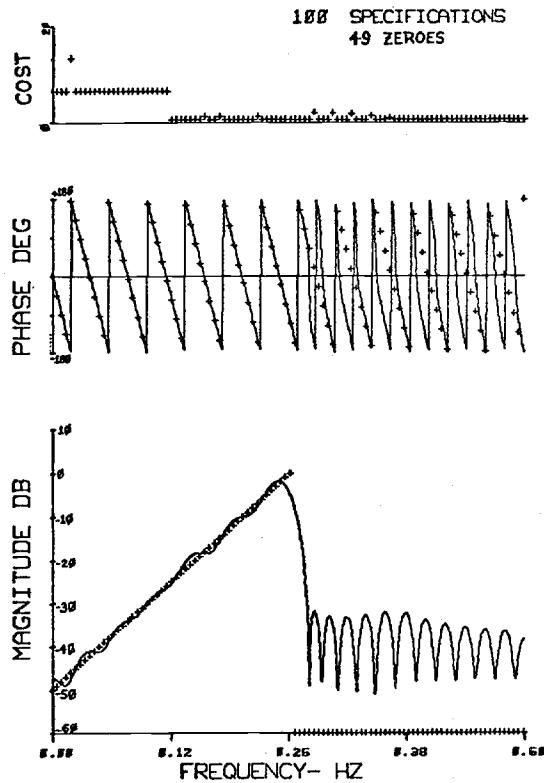


Fig. 3b Cost Weighted FIR design