UNIVERSITY OF WESTMINSTER

WestminsterResearch
http://www.wmin.ac.uk/westminsterresearch

Compromises in FIR matched filter design.

G.D. Cain[1]
Anush Yardim[2]
M.A. Mughal[1]

[1]DSP Creations Limited, UK
[2]Cavendish School of Computer Science

Compromises in FIR Matched Filter Design

G.D. Cain[+], A. Yardim* and M.A. Mughal[+]

[+] DSP Creations Limited, UK

* University of Westminster, UK

Abstract- With duration-limited signals there is the opportunity for perfect matched filtering by a suitable FIR filter, providing the noise can be construed as being supplied through an all-pole coloration filter. The popular matrix solution formulation does not make it obvious just what size the optimal filter length must be, and the signal vector zero-padding mechanism needed for coaxing out the optimal coefficient vector is also unclear. Difficulties are compounded when the filter length allowable for implementation falls short of the optimal length. Worse yet, if it happens that the noise is shaped by a coloration filter which has some zeros (i.e., is MA or ARMA instead of just AR), then any FIR filter can only be an approximation to the ideal IIR matched filter. In either case, decreasing filter length requires compromise strategies that are not at all transparent.

We base our analysis approach to the FIR problem setup in terms of (time) correlations and convolutions in which the whitening filter has the central role. It is then easy to see that both "pole-only" and "some-zero" noise cases yield optimal SNR values that are exactly calculatable by a time-domain scalar product. The inevitable degradations of SNR with decreasing FIR filter lengths are, in turn, readily quantifiable. We study several compromise strategies arising from the whitening filter convolution approach and find (albeit with a very limited set of test cases) that they are not attractive when contrasted to the common matrix solution. The matrix solution itself, meanwhile, is shown to demand close attention to zero-padding patterns employed in it if best performance is to be obtained as filter length is reduced. Fortunately, exhaustive zero-padding assessment is a practical proposition, and this is our recommended procedure at this early stage of investigation.

I. INTRODUCTION

Matched filtering is of enormous importance in many application arenas, ranging from its origins in radar/sonar signal detection, through biomedical pattern recognition to Multi-User Detection in the latest CDMA communication systems.

In view of its vital nature it is surprising that it is so difficult to find useful basic theoretical expositions of this workhorse in the literature, and there are almost no useful numerical examples to guide DSP design approaches. To be sure, there is plenty of worthy and substantial coverage in older books of continuous-time matched filter properties [1]-[4] (where the main focus is on imposing causality constraints), but the bulk of "meaty" digital matched filter coverage seems to be restricted to a mere handful of texts [5]-[8]. If one were to think that the sparsity of treatment is due to the problem area being too simple, or already satisfactorily expounded, that would be a thoroughly incorrect conclusion.

The causality constraint in continuous-time coverage of the topic (never mind the duration-limited constraint we face if we are to have FIR filters) quickly leads to delta functions located at the time origin. This confronts us with severe compromises if we are to convert such unruly prototypes (perhaps by impulse-invariance?) to digital counterparts. Meanwhile, the direct FIR approaches in the few books already cited are couched in matrix terms which we believe does little to illuminate the core design issues and tradeoffs.

The approach we prefer for better understanding is based on deterministic filter manipulations involving only convolutions and (time) correlations. Thus, we view statistical correlation simply as scaled time correlation (i.e., all Power Spectral Densities result from passage of white noise through coloration filters and so are merely scaled versions of the magnitude-square of these filter gains). This avenue of attack immediately reveals that there is a crisp hierarchy of design difficulty when the signal to be match-filtered is duration-limited; everything is down to the nature of the noise coloration filter:

(1) white noise from an allpass coloration filter (standard, but trivial)
(2) pole-only (AR) noise from an all-pole coloration filter (easy, with perfect results achievable)
(3) (MA/ARMA) noise from a coloration filter having zeros [and perhaps poles] (hard; demands compromises to make an IIR requirement into an acceptable FIR solution)
(4) noise from a coloration filter with zeros on the z-plane unit circle [and poles perhaps] (diabolically difficult, with severe optimality degradation likely)

A whitening filter corresponding to a given coloration filter has the inverse of its gain magnitude, but can have any convenient phase. We care only about the time autocorrelation of a whitener filter's impulse response, $h_W(k)$:

$$R_W(k) = h_W^*(-k) * h_W(k) \qquad (1)$$

where k is the time index, $*$ represents convolution and superscript asterisk stands for complex conjugate.

It can be shown that the optimal filter impulse response matched to a given duration-limited signal $s(k)$ having length N_s samples has an impulse response

$$h_{opt}(k) = s^*(-k) * R_W(k - k_p) \qquad (2)$$

where k_p is the particular time at which the peak signal instantaneous power-to-average noise power filter output has been optimized.

In the benign case where the noise coloration filter is all-pole, then $h_W(k)$ is FIR (of length N_W) and consequently $R_W(k)$ has finite duration. Therefore there is no difficulty in choosing

$$k_p = (N_S - 1) + (N_W - 1) \qquad (3)$$

and arriving at an FIR filter length N_{filt} which achieves its optimal length, N_{opt} :

$$N_{filt} = N_{opt} = N_S + 2(N_W - 1) \qquad (4)$$

This simple expression gives us some good insight. With white noise (our welcome trivial case), N_w is one and the best filter has its length equal to the signal. Indeed, we will set N_s as the smallest allowable compromise collapse of N_{filt} when reducing length below N_{opt}.

With the filter resulting from the convolution set out in (2) we are guaranteed to have the best possible SNR (taken in the usual matched filter sense by adopting length N_{opt} for N_{filt}'s value). This is given by a scalar product:

$$SNR_{opt} = [R_S(k) * R_W(k)]_{k=0} \qquad (5)$$

II. CONSTRAINING N_{filt} BELOW ITS OPTIMAL VALUE

But what happens if the filter length demanded by (4) cannot be accommodated? (This might happen, say, if N_s is itself already large - maybe if it is a long chirp sequence - and hardware limitations then preclude large N_{filt}). Inevitably we will have to restrict N_{filt} below the optimal length indicated by (4) (so that we have a suboptimal filter which we will just call h(k)) and the achievable SNR will degrade below that value given by (5), throwing open the question of the best k_p value (i.e. equation (4) need not be in effect, and it will be necessary to locate the time at which $s(k)*h(k)$ takes on its peak absolute value).

Cadzow [5, p. 449, p. 460] presents the only FIR specimen problem with pole-only noise that we have been able to find in the literature, so - despite its very small signal length (N_s=2) - it merits brief exposure. It allows us to introduce symbology (linked to the padding options used for its matrix solution) that will be crucial to recognize in all our subsequent plots of results. He takes a single-pole filter with pole located at z=0, so

N_W=2. On page 447 Cadzow advocates only a single padding strategy (see "Prepending" below), but we find that two other cases are essential to incorporate and survey:

TABLE A
CADZOW PROBLEM SUMMARY

	Padded Signal Vector	N_{filt}	Achieved SNR (times 3)	Symbol
Original	[3 -1]	2	52	
Prepend 1	[0 3 -1]	3	61	□
Append 1	[3 -1 0]	3	53	Δ
Prepend 2	[0 0 3 -1]	4	61	□
Append 2	[3 -1 0 0]	4	53	Δ
Sympend 2	[0 3 -1 0]	4	62 (optimum)	✱

Obviously we should start at the bottom of Table A and move upward if we are considering our startpoint as the optimal, full-length solution. There are 2 suboptimal solutions which possess the same length, and decreasing N_{filt} takes us up the table, encountering different padding mechanics as we go. Our best percent of perfect SNR with N_{filt}=3 is 100*(61/62)=98.4% and with N_{filt}=2 is 100*(52/62)=83.9%. This monotonic decrease of SNR effectiveness with tightening filter length constraints seems natural; it is therefore all the more surprising to find aberrant situations in our subsequent measurements, with decreased N_{filt} values occasionally outperforming their larger neighbors!

Proceeding to a somewhat larger specimen problem of our own choosing, we devise the all-pole coloration filter of Fig. 1 having 7 poles (so N_w=8) deployed in such a way as to give an oscillatory $R_w(k)$ and a matched filter pre-disposed to be bandpass:
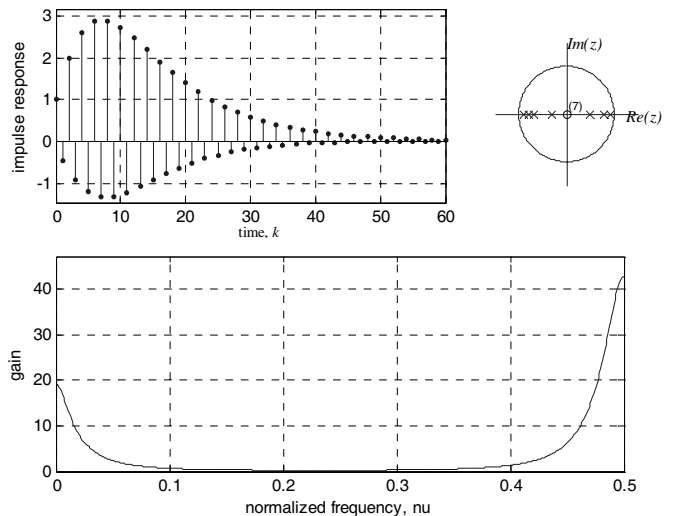

Fig. 1. An All-Pole Coloration Filter

The signal chosen is a 7-sample linear-phase (for visual convenience) triangle sequence described by this vector: [1.75

3.5 5.25 7 5.25 3.5 1.75]. We make a set of experimental measurements as N_{filt} ranges from $N_s=7$ to the optimal length of 21:
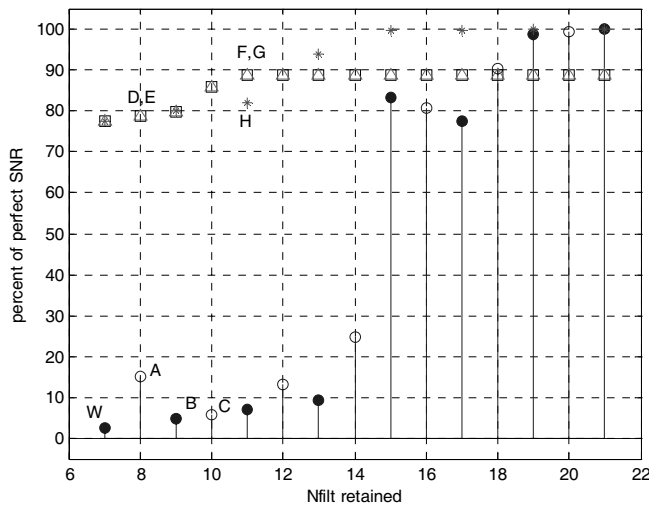

Fig. 2. Variation of Optimal SNR Percentage Achievement versus Allowed Matched Filter Size (N_s Odd, Linear-Phase Signal) with Impulse Response Truncation

Clearly the non-monotonic variation evident in Fig. 2 is a bit unexpected, since - as mentioned earlier - it would seem intuitively that larger filter sizing should relentlessly improve performance. There are two prime categories of design being shown in this diagram. First, the disembodied symbols near the top (values D, E, H for instance) are the result of matrix calculations of the sort we have met in Table A, and the symbols signify the same zero-padding strategies already introduced. Notice that prepending and appending of clustered zeros can be done for every N_{filt} selected, while symmetric padding is only possible here for odd N_{filt} values. Note also that these asterisk symbols are often better than prepending or appending, but NOT ALWAYS (see points F and G out-performing point H)!

The stemplotted values result from a different design strategy and are (until N_{filt} grows to 18 and bigger anyway) disappointingly poor contenders. The way that the stemplot designs arise is by taking the optimal design (where the asterisked matrix design coefficient set perfectly coincides with that obtained by (2), at $N_{filt}=21$) and progressively truncating their length as we move leftward along the horizontal axis. Observe that the stems come in two varieties; the filled circles represent the symmetric truncation obtainable only at odd N_{filt} settings, while the unfilled circles represent the better of a left-skewed and a right-skewed truncation for every even filter length.

At first it seems counter-intuitive that skewed impulse responses such as those contributing to points A and C should be better than the symmetric response that gave us point B!

And to further surprise us, the filter for point C is worse than the smaller one at A! But this curious ambivalence is no more than the edge effects familiar from flat windowing; we have a highly oscillatory $R_w(k)$ convolved with a non-negative signal via (2), so we expect an oscillatory $h_{opt}(k)$ which, when truncated, will be very dramatically affected as positive-to-negative outrider samples are lopped off.

We can look at a couple of snapshot designs lifted out of the family of filters which underlie the symbols in Fig. 2. Fig. 3 shows the complete signalling spectral environment for the best SNR case achievable (with $N_{filt}=21$) contrasted with the most extremely truncated case (point W) where $N_{filt}=7$:
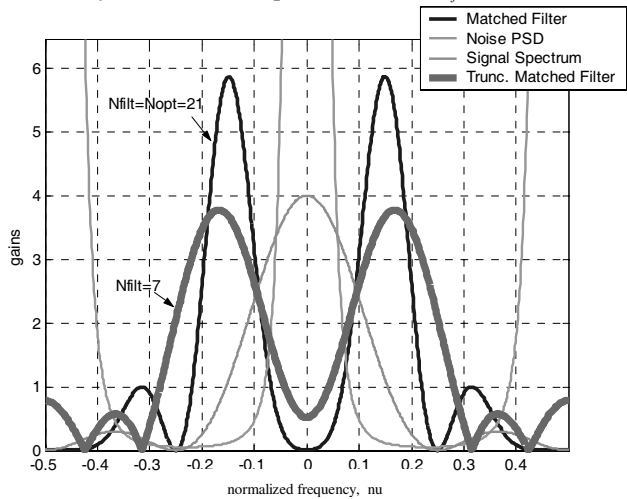

Fig. 3. Spectral Plots Corresponding to the Two Extreme Filters Among Those Comprising the Stemplots in Fig. 2

The heavy $N_{filt}=7$ plot displays the indelicate smearing we anticipate when truncation of an impulse response takes place. The damage to performance this entails is catastrophic, taking us down from 100% effectiveness to under 3% when we only retain the central 7 coefficients from the centre of the optimal impulse response!

It is worth pointing out that Frequency Sampling design was carried out as an alternative to impulse response truncation, and that similarly smeared gain results led to stemplots reminiscent of the disappointing ones in Fig. 2.

III.   CHANGING THE SIGNAL'S SIZE AND SYMMETRY

We naturally want to expand our tiny base of experience. So the first thing new to explore is the whether N_s even acts the same erratic way as N_s odd. A change to a linear-phase $N_s=6$ triangle signal vector reveals behaviour broadly similar to Fig. 2 (but a bit more orderly, as Fig. 4 shows).
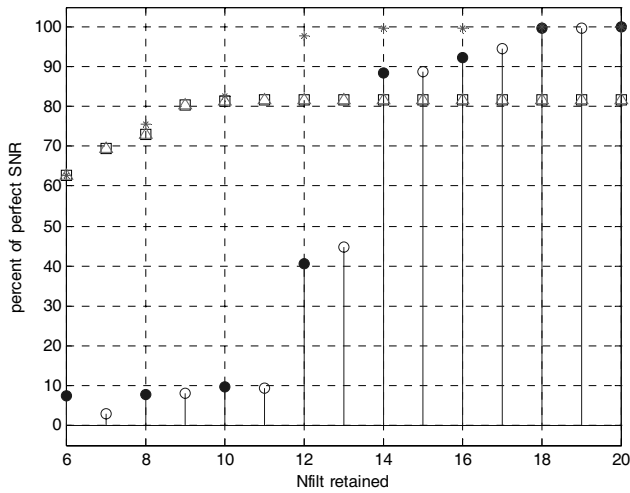
Fig. 4. Variation of Optimal SNR Percentage Achievement versus Allowed Matched Filter Size (N_s Even, Linear-Phase Signal) with Impulse Response Truncation

Far more interesting is a switch from linear-phaseness (our triangular signal above has all its z-plane zeros on the unit circle) to minimum-phaseness for another N_s=7 signal (we chose the signal created by the MATLAB vector [7:-1:1]). The impact is startling:
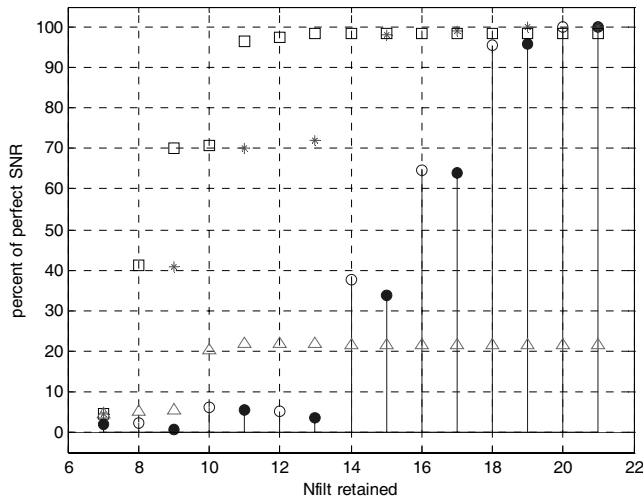

Fig. 5. Variation of Optimal SNR Percentage Achievement versus Allowed Matched Filter Size (N_s Odd, Minimum-Phase Signal) with Impulse Response Truncation

As Fig. 5 demonstrates, prepending our zero-padding makes very much more sense (up until N_{filt}=17 and above) than symmetric padding. Appending the zeros relegates matrix solutions to the sort of poor performance seen for truncation results back in Fig. 2. Meanwhile, these truncation results themselves are even worse than seen in Fig. 2. Unsurprisingly, flipping to use a maximum-phase signal causes the positions of the triangles and squares of Fig. 5 to interchange! Clearly, the

symmetry or asymmetry of the signal to be matched can have a profound effect on the way preparations must be made to (suboptimally) "match"-filter it.

Thus far, we can conclude that truncation of an optimal impulse response is risky and is probably likely to under-perform dramatically if more than a very few outlying samples are expunged. (This, by the way, bodes ill for the truncation procedures advocated in the harsher situation of IIR matched filters by several of our References).

IV.   REDUCED ALL-POLE COLORATION MODELLING

Another strategy for compromise as filter length decreases is (assuming we are still at difficulty level (2), having pole-only noise) to progressively reduce the number of poles utilized, so that each N_{filt} size represents a pseudo-optimal situation; that is, recalibrate coloration to an approximate model that allows each N_{filt} size to seem to be optimal filter length. We can use the all-pole modelling procedure described in Hayes [9, p.181], operating on the original (full) all-pole coloration impulse response to down-size the model smoothly.

Of course our noise model gets progressively more outrageous as N_{filt} decreases, until we ultimately have the fiction of white noise when N_{filt} bottoms out at the value N_s. Despite its obvious drawbacks, this mythical outlook gives noteworthy results:
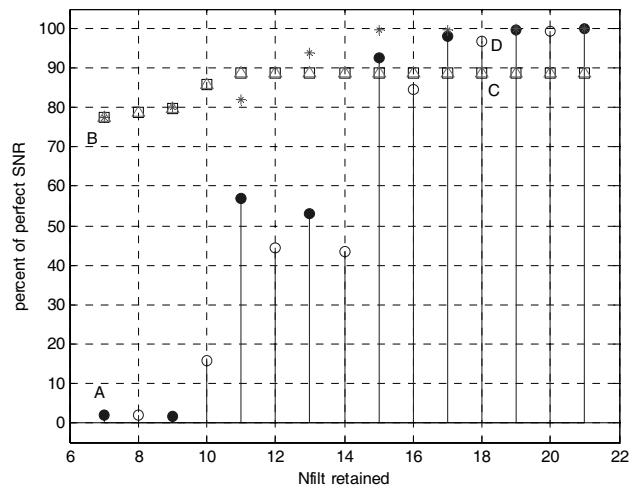

Fig. 6. Variation of Optimal SNR Percentage Achievement versus Allowed Matched Filter Size (N_s Odd, Linear-Phase Signal) with Reducing Pole-Only Coloration Modelling

In Fig. 6 we have the same matrix solution results as seen in Fig. 2. The new element here is the substantially-improved stemplot which treats the effects of progressively-modified all-pole model. If we look at point A we see the white noise solution (that is, h_{opt} is merely our $s(k)$, since we have returned to our linear-phase creation as in Fig. 2). Still there is a

massive shortfall compared to the matrix solution shown as point B. (Clearly, taking account of noise shaping <u>matters</u>).

The stems manage to close a lot of the gap to the matrix solutions and, near to optimal vector size, do better (see points C and D). While still short of the matrix solution effectiveness, reducing all-pole sizing offers something when we encounter greater levels of problem difficulty, stages (3) and (4).

The story then gets gloomy: the coloration filter has transmission zeros. Of course these become poles of the whitening filter, and we must decide how to drop from the now-infinite dimensionality actually required, to an acceptable FIR compromise. Since we are free to make our whitening filter have any phase we want, we can at least guarantee whitener stability by making the coloration filter become minimum-phase. Then we can again return to the issue of trying to satisfactorily model the true coloration situation by the reducing all-pole procedure just entertained above.

Although we need to be mindful of the difficulties of all-pole design where abrupt gain transitions are found [10, p.338], we can often arrive at an acceptable approximation to the true coloration and hence successfully employ our whitener correlation analysis approach. Fig. 7 shows a respectable 9-pole/9-zero ARMA noise coloration situation:
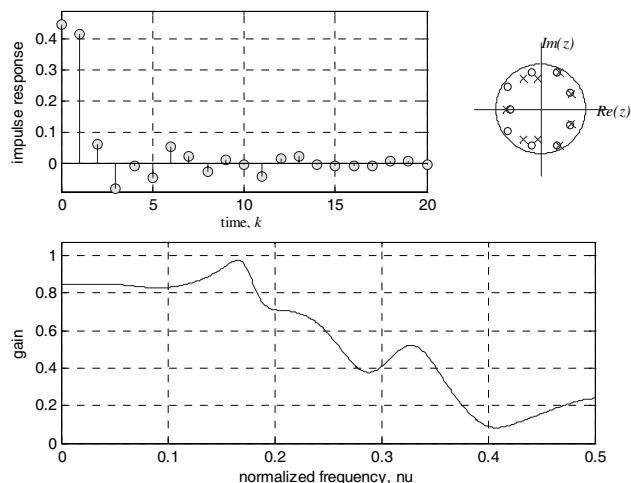

Fig. 7. An ARMA Coloration Filter

Following the same all-pole approximating procedure which yielded an erratic Fig. 6 even when we started out all-pole, we find this time that the matrix solutions are much more closely approached by the correlation-designed filters. Despite not delivering the IIR matching that theory demands, most filters are satisfactory, near the 90% effectiveness level.

This serves to emphasize that, depending heavily on the coloration severity, there are occasions when very gratifying results are obtainable by <u>all</u> our methods.
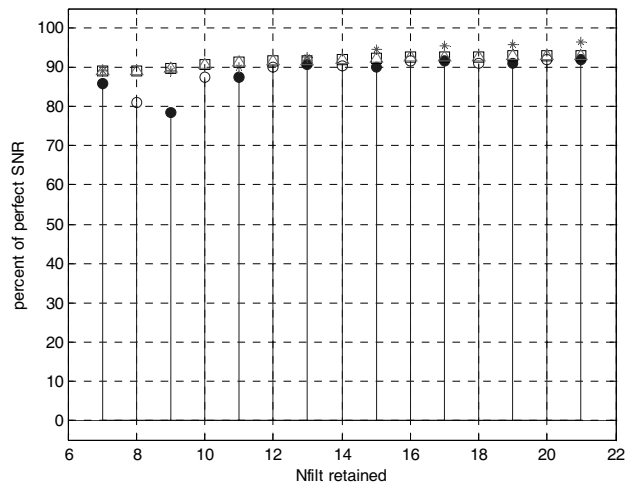

Fig. 8. Variation of Optimal SNR Percentage Achievement versus Allowed Matched Filter Size (N_s Odd, Linear-Phase Signal) with Reducing Pole-Only Modelling of an ARMA Coloration Filter

V. RECOMMENDATIONS FOR PRACTITIONERS

Obtain precisely the best SNR that can ever be obtained with the finite duration signal you are given as well as the optimal filter length and use these to inform your decision among compromise alternatives. If the number of allowable coefficients is markedly below the optimal length, our experience thus far suggests that a matrix solution may well be preferred. However, zero-padding strategies must be exhaustively evaluated to extract the best matrix solution for that value of constrained filter length. Along with this, the location of the required k_p must be determined by convolution and response peak location.

REFERENCES

[1]   DiFranco, J.V. and W.L. Rubin, *Radar Detection*, Artech, 1980.
[2]   Franks, L.E., *Signal Theory*, Prentice-Hall, 1969.
[3]   McDonough, R.N. and A. D. Whalen, *Detection of Signals in Noise*, Second Edition, Academic Press, 1995.
[4]   Thomas, J.B., *Introduction to Statistical Communication Theory*, Wiley, 1969.
[5]   Cadzow, J.A., *Foundations of Digital Signal Processing and Data Analysis*, Macmillan, 1987.
[6]   Therrien, C.W., *Discrete Random Signals and Statistical Signal Processing*, Prentice-Hall, 1992.
[7]   Thomas, J.B., *An Introduction to Communication Theory and Systems*, Springer-Verlag, 1988.
[8]   Mix, D.F., *Random Signal Processing*, Prentice-Hall, 1985.
[9]   Hayes, M.H., *Statistical Digital Signal Processing and Modelling*, Wiley, 1996.
[10] Jackson, L.B., *Digital Filters and Signal Processing*, Third Edition, Kluwer, 1996.