UNIVERSITY OF WESTMINSTER

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

Dynamic matched filtering: animating the action.

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

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

Copyright © [2004] IEEE. Reprinted from 2004 IEEE 11th Digital Signal Processing Workshop & 2nd IEEE Signal Processing Education Workshop, pp.73-78.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Westminster's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners. Users are permitted to download and/or print one copy for non-commercial private study or research. Further distribution and any use of material from within this archive for profit-making enterprises or for commercial gain is strictly forbidden.

Whilst further distribution of specific materials from within this archive is forbidden, you may freely distribute the URL of WestminsterResearch. (http://www.wmin.ac.uk/westminsterresearch).

In case of abuse or copyright appearing without permission e-mail wattsn@wmin.ac.uk.

2004 IEEE 11th Digital Signal Processing Workshop
& IEEE Signal Processing Education Workshop

DYNAMIC MATCHED FILTERING - ANIMATING THE ACTION

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

+ DSP Creations Limited, UK
* University of Westminster, UK

ABSTRACT

A Simulink-based block diagram modelling environment is described which makes investigation of demanding DSP concepts such as FIR matched filtering easy and fun. Users interact with the experiment to a remarkable degree, watching scope displays while tuning parameter values or moving sliders to effect model changes during run time in a dynamic fashion. Instrumentation for achieved signal-to-noise ratio sits alongside displays advising the experimenter of theoretically optimal SNR for the current parameter settings. A small example problem using a single-pole noise-shaping filter is seen to be very enlightening, especially since a variable-coefficient matched filter block is employed which is self-designing in response to the prevailing pole radius and resonant frequency selection.

1. INTRODUCTION

Simulink is the block diagram tool in the MATLAB family of software products. It permits nearly effortless visual assembly of complicated signal generators, processing system elements and measuring "instruments" into runnable behavioural system simulation configurations while still retaining access to the MATLAB Workspace, renowned for its ease of ready calculations, friendly graph plotting and rapid development capability for high-level algorithms. Thus Simulink's potential for accelerating learning of topics in signals and systems areas is unparalleled.

We have produced a collection of specialized Simulink blocks which are especially conducive to ready understanding of many common DSP and Communication System concepts. Matched filtering is an excellent case in point. First off, though there are many sophisticated treatments of spectral factorization issues, relatively few textbooks treat digital matched filters at a basic level, and very few indeed which provide good example problems that can be studied and easily extended. Particularly with FIR matched filters there is a tendency to cast all expressions weightily in matrix form, which sometimes shrouds the underlying simplicity of the calculations that are actually needed, as well as tradeoffs available.

The best examples we have found are in Cadzow [1, p. 445], although some useful problems are also found in [2]-[5]. Still, the flow of calculations fails to lead the student to firmly grasp much of this topic that is of such vital importance in radar, sonar, biomedicine, and modern multi-carrier modulation systems. We felt this was a clear case where Simulink could shed some life, light, and learning motivation!

2. THE HIERARCHY OF DIFFICULTY

In dealing with FIR matched filtering it is important to first understand what makes some problems easy, some difficult, and some enormously difficult. We have found it better to gain this feel by abandoning the matrix-based approach which is central in the texts already cited.

The approach we prefer 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 colouration 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 colouration filter:

(1) white noise from an allpass colouration filter (standard, but trivial)
(2) pole-only (AR) noise from an all-pole colouration filter (easy, with perfect results achievable)
(3) (MA/ARMA) noise from a colouration filter having zeros [and perhaps poles] (hard; demands compromises to make an IIR requirement into an acceptable FIR solution)

0-7803-8434-2/04/$20.00©2004 IEEE

(4) noise from a colouration filter with zeros on the z-plane unit circle [and poles perhaps] (diabolically difficult, with severe optimality degradation likely)

In this paper we concentrate on difficulty level (2) above, since our introductory laboratory exercise "matchlab1" is totally devoted to that situation, and it harbours plenty of interesting detail. Later laboratory work assails difficulty levels (3) and (4).

3. CALCULATIONAL FRAMEWORK

Before having Learners model and experiment with matched filtering action it is necessary to clearly plant the idea of the optimality criterion being imposed (a curious SNR which is peak instantaneous output power due to the incoming signal - to -average output power due to the inbound noise), and also to develop the set of resulting design equations which must routinely be used.

Learners first have a one-hour lecture session devoted to these issues before any laboratory exposure. During that session they will have worked two elementary matched filter problems: an easy white noise case from [5] and a first taste of Cadzow's two-sample signal problem [1, p.445] where he limits consideration to the suboptimal filter size of two coefficients. One of our objectives in the laboratory is to see that experience broadened to the ideal size of four coefficients. (All modest sizes, to be sure; but sufficiently illustrative of all the vital concepts).

Learners have a set of slide images which expound the theory and show step-by-step the systematic flow of calculations necessary, emphasizing those for dealing with the pole-only noise case. The key equations are rehearsed here.

All practical work starts from the punchline of the optimization derivation which is the frequency-domain equation

$$\mathcal{H}_m(\nu) = \frac{\mathcal{S}^*(\nu)}{\Phi_{\mathcal{N}}(\nu)} e^{-j2\pi\nu k_p} \qquad (1)$$

This normalized frequency (ν) equation has the (DTFT) spectrum of the time (k) signal $s(k)$ which is to be matched, the noise power spectral density $\Phi_{\mathcal{N}}(\nu)$, and a complex exponential phase factor involving the particular time instant k_p specified for the sampling of the filter output. The result is the transfer function of the optimal matched filter (before any considerations of scaling).

As already mentioned we always adopt the modeller's view that the noise we have has been issued by a colouration filter acting on a white noise source with power σ_i^2 :

$$\Phi_{\mathcal{N}}(\nu) = \sigma_i^2 \left|\mathcal{H}_c(\nu)\right|^2 \qquad (2)$$

But our analysis procedure usually focuses much more on whitener filters. The whitening filter corresponding to a given colouration 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, $h_w(k)$:

$$R_w(k) = h_w^*(-k) * h_w(k) \qquad (3)$$

where again 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 N_s-long signal $s(k)$ has an impulse response

$$h_m(k) = s^*(-k) * R_w(k-k_p) \qquad (4)$$

The output delivered at time k_p by this (unscaled) filter is given by a similarly simple scalar product:

$$y_p = \sum_{k=-\infty}^{\infty} \mathcal{R}_s(k)\, \mathcal{R}_w^*(k) \qquad (5)$$

The optimal SNR achievable from the matched filter is

$$SNR_{opt} = |y_p| / \sigma_i^2 \qquad (6)$$

Finally, a word about filter scaling. Learners are pleasantly surprised when it is pointed out that - of course - the filter coefficients can be scaled in any way desired, without affecting the SNR. But this does not mean that no attention should be paid to coefficient scaling. We prefer to cast our final declaration of our optimal filter as

$$h_{opt}(k) = h_m(k) / y_p \qquad (7)$$

This ensures unit-voltage output at time k_p and makes it easy to compare performance with non-matched filters which are similarly scaled.

4. THE FOUNDATION EXPERIMENT

We successfully relied upon these few equations when we field-tested several matched filter experiments in about a dozen laboratory sessions, involving both industrial short course participants and postgraduate students. We initially utilize a simple two-sample signal being received in additive noise - first white (of course) and later pink noise. We used exactly the flagship problem in Cadzow and several extensions. Since the noise-shaping filter is a particularly simple first-order pole-only filter, we know that an exact FIR optimal solution exists and that we can make all necessary calculations using the above equations for the optimal coefficients and the optimal Signal-to-Noise delivered - by hand. Thus, all simulation results can be readily cross-checked for confidence re-enforcement.

Figure 1 shows the splash panel that a Learner sees; it conducts him through the colour-coded laboratory Tasks:

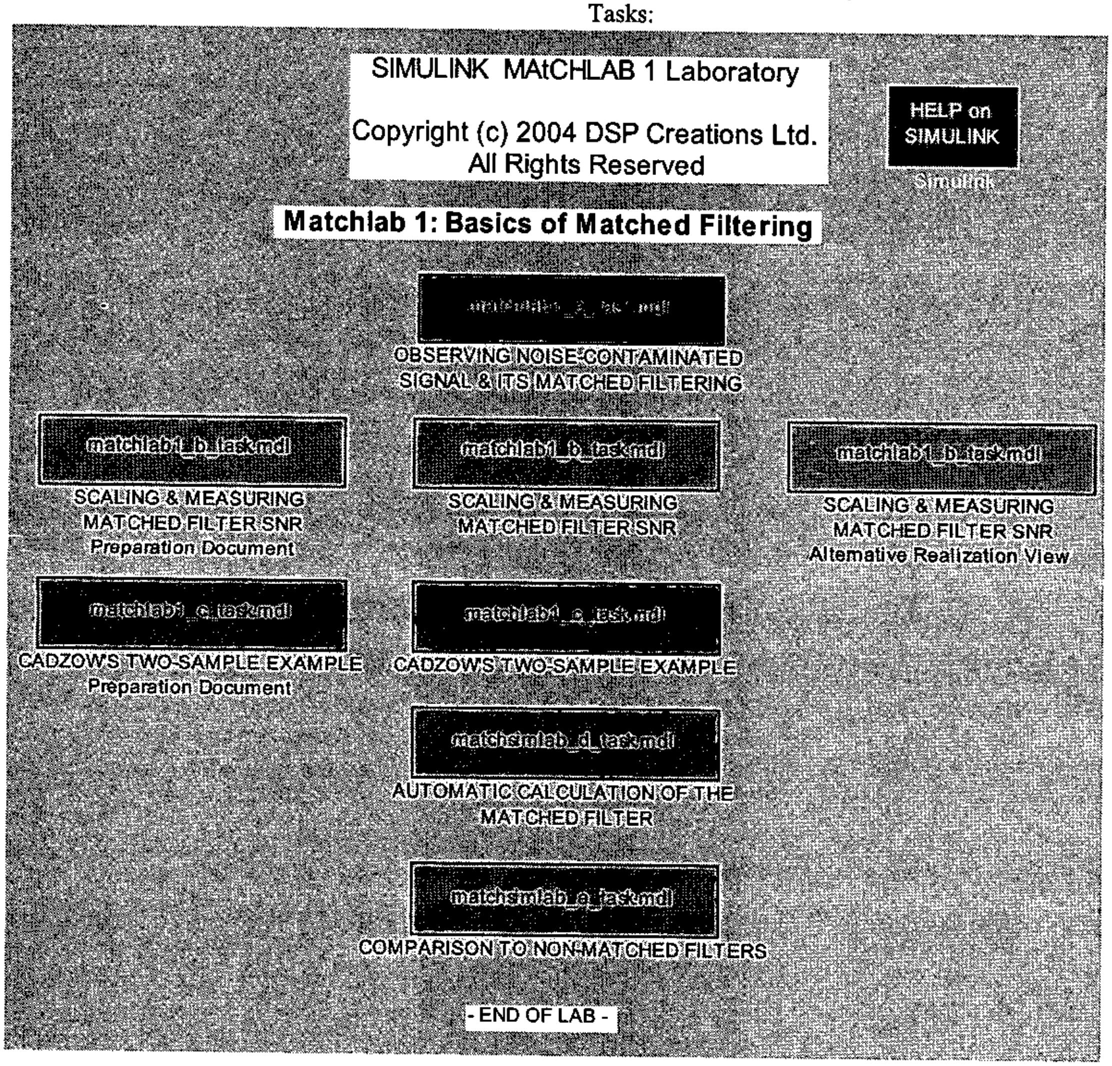

Figure 1 Laboratory Task Splashpanel

Of course the Instructor has, as well as this Task sheet, the similar Solution splashpanel, giving useful guidance on what Learners should gain from each Task.

Figure 2 shows the model being employed midway through the exercise, Task c:

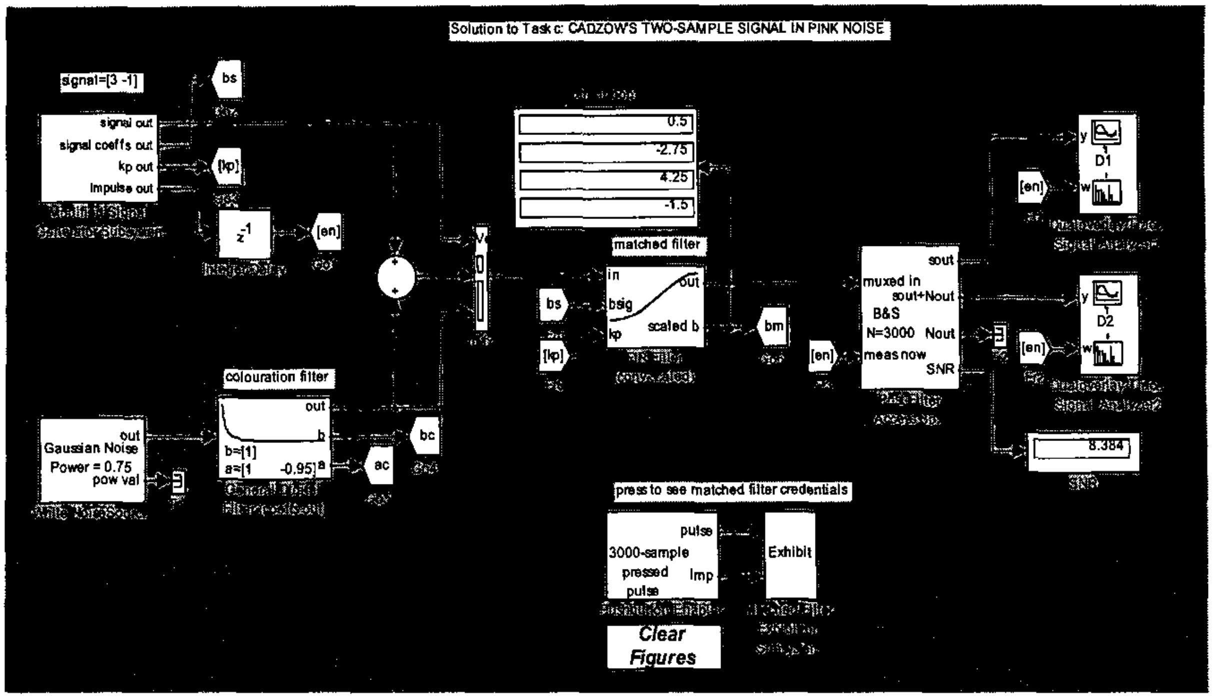

<u>Figure 2</u> Task c- Modelling Cadzow's Two-Sample Signal Example

By the time Learners reach this point in their work plan they have been "warmed up" to smoothly use the "equipment" provided, can make changes to block parameters and can turn attention to the animation of the calculations just outlined.

At upper left position in the model can be seen the signal generation apparatus, while the white noise feeding the colouration filter (causing the noise to become "pinked") is located at the lower left. The matched filter is situated in the model's middle, and has a coefficient display box above it and a cluster of design result visualization blocks below it. At the extreme right comes the Signal Analyzers for viewing the noisy waveforms and the SNR measurement facilities.

With this testbench Learners immediately see active waveforms confirming their correct calculations (done separately in the MATLAB command window and broadcasted directly into the appropriate Simulink blocks) and can see how close measured SNR comes to their theoretical expectation.

The grand finale of this experiment is the model of Task e, which compares the suboptimal SNR performance of two non-matched filters with a matched filter. This model is much more elaborate than any previously seen and would overwhelm the Learner if it were engaged before the gentler buildup of Tasks a-d.

Figure 3 shows this intricate model, which affords enormously flexible animation of all the ideas which have gone before. At the left edge are a ramp signal source (a somewhat larger signal vector than heretofore) and a white noise source that is "resonated" by a second-order resonator colouration filter. This filter is **slider-adjustable**, so that different pole radius and pole frequency shifts can be manually varied at simulation run time. Such variability makes for a **dynamic** style of experimentation that greatly enhances its appeal!

Three different filters are subjected to testing: the upper two processing branches have arbitrarily-chosen filter coefficient sets alongside (in the lower branch) the truly optimal matched filter. This truly optimal filter is being **designed automatically**, with no intervention by the student whatsoever. What's more, as slider settings on the noise-shaping filter are manipulated, the new matched filter required is instantly enacted and put to work processing the endless flow of noise-contaminated signal pulses!

At the right edge of the model window is instrumentation for exhibiting <u>measured</u> SNRs alongside the theoretically predicted optimal SNR. On various 3-d scopes connected it is apparent just which contributions signal and noise are making to the composite real-world noisy signal and just how critical it is to pick the unique sampling instant for eliciting best detection performance

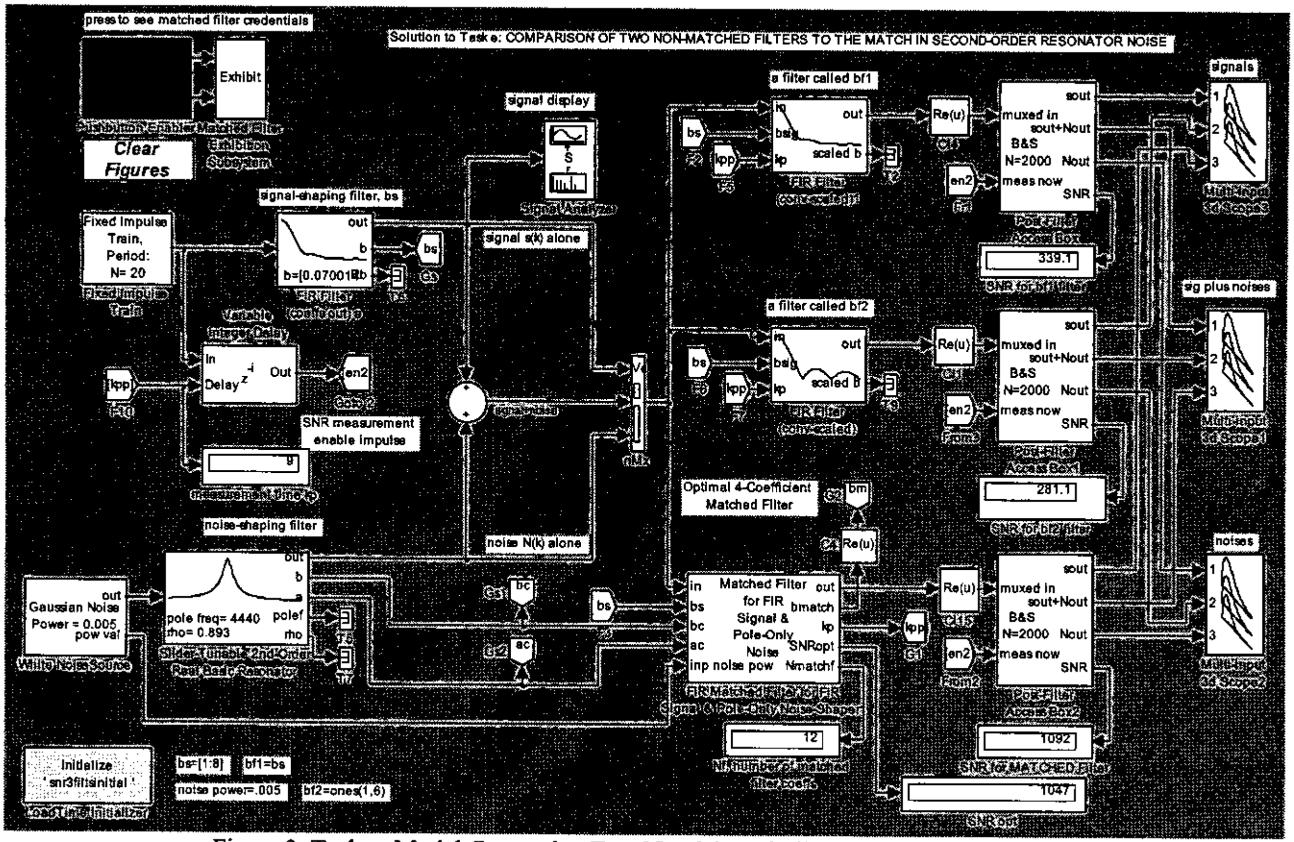

Figure 3 Task e -Model Comparing Two Non-Matced Filters to the Matched Filter

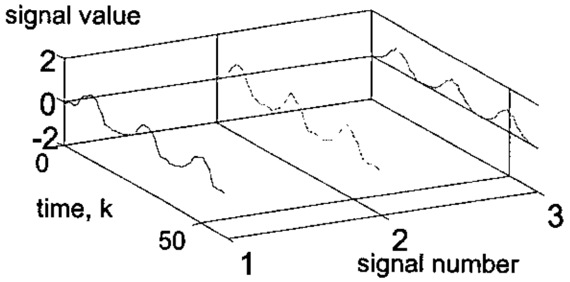

Figure 4 Post-Filter 3-d Scope Display
(Signal + Noise traces for 3 filters)

Learners find it interesting to see the SNR deterioration that accompanies pole rotations and radius decreases. After having the ease of this automatic delivery of results (verified by comparative SNR measurements) impressed upon them, students then are permitted to "take the lid off" the Simulink blocks to see inside how the calculations for the hidden noise-matching and signal template subfilters are organized. They make spot checks on the numerical validity of the results being furnished automatically by simple hand/MATLAB calculations. Thereafter, they are ready to tackle more sizable examples and to pursue - with massively-boosted confidence - more demanding signal detection and decision tasks on a subsequent occasion in "matchlab2".

The Learner has a real chance to mould the exploration plan by painless parameter changes which yield gratifying, instant pictorial responses on a range of scopes, while numerical readouts give precision wherever needed. Views on this style of highly-interactive experiments have been extremely enthusiastic, both on the part of Instructors and participants in the study programmes. Learners have voiced considerable satisfaction with "matchlab1" and the two following experiments which build on it. and we are actively extending from these to a broader range of DSP subject areas.

REFERENCES

[1] Cadzow, J.A., *Foundations of Digital Signal Processing and Data Analysis*, Macmillan,1987.

[2] Therrien, C.W., *Discrete Random Signals and Statistical Signal Processing*, Prentice-Hall, 1992.

[3]Thomas, J.B., *An Introduction to Communication Theory and Systems*, Springer-Verlag, 1988.

[4] Mix, D.F., *Random Signal Processing*, Prentice-Hall, 1985.

[5] O'Flynn, M., *Probabilities, Random Variables, and Random Processes*, Harper & Row, 1982.