A Class of Adaptive Matched Digital Filters*

by

K. Steiglitz and J.B. Thomas

Department of Electrical Engineering

Princeton University, Princeton, N.J.

ABSTRACT

An adaptive digital matched filter structure is developed for the case where the input signal is of known form and finite duration and the input noise has a power spectral density which is all-pole. The effect of noise spectrum identification errors on system performance is investigated both theoretically and experimentally. It is shown that, when the noise is highly correlated, the adapting structure leads to significant improvement in the output signal-to-noise ratio (and hence in the detection characteristics) with relatively short measurement times. This suggests the use of switching logic to allow noise adaptation only when measurements indicate a highly correlated noise background.

1. INTRODUCTION

Most signal detection systems, particularly those which are optimum in some sense, are designed to be operated in a fixed noise environment. It is highly desirable in many situations to have available systems which are relatively insensitive to input noise statistics. Such systems are particularly useful, for example, when the noise is stationary over only a short time interval.

A reasonable approach to this problem is to consider detection schemes which have been made adaptive to a class of input noises by the addition of a noise-identification system. The increasing use of digital techniques and the general availability of digital facilities permit the easy implementation of such identification systems. Of particular simplicity are those designed to estimate parameters in an assumed all-pole noise spectral density function [1].

A wide class of detection schemes uses matched filters followed by threshold comparators [2]. In fact, for the detection of signals of known form in Gaussian noise, such detectors are optimum [2] in the Neyman-Pearson sense [i.e. for a fixed false alarm probability (level) they give the greatest probability of a hit (power)]. These matched filters have transfer functions which depend on the input noise spectrum [3]. This paper considers a class of adaptive matched digital filters; i.e. filters which operate on sampled-data and whose transfer functions are determined from estimates of the input noise power spectral density.

In the next section, the realizable transfer function of the matched digital filter is obtained. Although the derivation is a straightforward extension from the continuous case, it does not appear to be commonly available in the literature.

---

* This work was partially supported by the National Science Foundation under Grant GP-1647.

2. MATCHED DIGITAL FILTERS

The matched digital filter is shown schematically in Fig. 1, where T is the sampling interval and H(z) is the sampled transfer function of a linear time-invariant digital filter with impulsive response h(nT). The input consists of samples of a known signal s(t) and of an additive random noise n(t), assumed to be a wide-sense stationary process with zero mean and power spectral density $\phi(z)$. The filter H(z) is to be chosen to maximize the output S/N power ratio.

In complete analogy with the continuous case [3], the output noise variance $\mathcal{N}$ and the squared output signal component $\mathcal{S}$ at time t = NT can be written as

$$\mathcal{N} = \frac{1}{2\pi j} \oint H(z)H(z^{-1})\phi(z)\frac{dz}{z} \tag{1}$$

and

$$\mathcal{S} = \left[\frac{1}{2\pi j} \oint H(z)S(z)z^{N}\frac{dz}{z}\right]^2 \tag{2}$$

or

$$\mathcal{S} = \left[\frac{1}{2\pi j} \oint H(z^{-1})S(z^{-1})z^{-N}\frac{dz}{z}\right]^2 \tag{3}$$

where capital letters have been used to denote the z-transforms of lower case time functions and where integration is around the unit circle in the counter-clockwise direction.

It is desired to choose H(z) so as to maximize the ratio $\mathcal{S}/\mathcal{N}$ or, equivalently [4], the term $\mathcal{N} - \lambda\mathcal{S}$ where λ is a Lagrangian multiplier. Differentiating with respect to $H(z^{-1})$ under the integral sign and equating the result to zero yields

$$\frac{1}{2\pi j} \oint\left[\frac{H(z)\phi(z)}{z} - \lambda\frac{S(z^{-1})z^{-N}}{z}\right]dz = 0 \tag{4}$$

Hence the integrand can be set equal to a function X(z) which is analytic inside the unit circle [5]:

$$\frac{H(z)\phi(z)}{z} - \lambda\frac{S(z^{-1})z^{-N}}{z} = X(z) \tag{5}$$

The noise spectrum can be factored into two components $\phi^{+}(z)$ and $\phi^{-}(z)$ having poles and zeroes only inside and only outside the unit circle respectively. This factorization should be performed so that $\phi^{+}(z)$ does not possess a pole or zero at infinity and so that $\phi^{-}(z)$ [which is just $\phi^{+}(z^{-1})$] does not possess a pole or zero at the origin [6]. Then Eq. (5) becomes

$$\frac{H(z)\phi^{+}(z)}{z} - \lambda \frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)} = \frac{X(z)}{\phi^{-}(z)} \tag{6}$$

Now let the coefficient of λ in Eq. (6) be expressed as a partial fraction expansion of the form

$$\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)} = \left[\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)}\right]_{+} + \left[\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)}\right]_{-} \tag{7}$$

where the term $[\ ]_+$ contains only poles which are inside the unit circle and the term $[\ ]_-$ contains only poles outside the unit circle. Then Eq. (6) can be written as

$$\frac{H(z)\phi^{+}(z)}{z} - \lambda\left[\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)}\right]_{+} = \frac{X(z)}{\phi^{-}(z)} + \lambda\left[\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)}\right]_{-} \tag{8}$$

If the function H(z) is physically realizable, then the left side of Eq. (8) is analytic outside the unit circle; the right side is analytic inside the unit circle by definition; therefore both sides must be equal to a constant which can be shown to be zero [7]. Since the value of λ is arbitrary for our purposes [3], it may be set equal to unity and we have

$$H(z) = \frac{z}{\phi^{+}(z)} \left[\frac{S(z^{-1})z^{-N}}{z\phi^{-}(z)}\right]_{+} \tag{9}$$

as the transfer function of the realizable matched digital filter. If physical realizability can be ignored, then it follows directly from Eq. (4) [or from Eq. (9)] that the matched filter is [8]

$$H(z) = \frac{S(z^{-1})z^{-N}}{\phi(z)} \tag{10}$$

When the signal s(t) is of finite duration, N can always be chosen large enough so that $S(z^{-1})z^{-N}$ is analytic outside the unit circle. (This is equivalent to delaying the signal long enough so that it vanishes for negative time.) If, in addition the power spectral density $\phi(z)$ of the noise is all-pole, then N can be chosen large enough so that the term $[\ ]_+$ in Eq. (9) is analytic outside the unit circle. In this case the matched filter transfer function reduces to Eq. (10) where N ≥ [degree of $S(z^{-1})$] + [degree of $1/\phi^{-}(z)$]. Since the noise has been assumed to be all-pole, its spectral density can be written as

$$\phi(z) = 1/[D(z)D(z^{-1})] \tag{11}$$

where the polynomial D(z), given by

$$D(z) \quad = \quad 1 + \alpha_1 z^{-1} + \ . \ . \ . \ + \alpha_p z^{-p} \tag{12}$$

has zeroes only inside the unit circle. Note that any scale factors can be ignored since they do not affect the S/N ratio improvement.

For the all-pole case, the matched filter is given by

$$H(z) \quad = \quad S(z^{-1})D(z)D(z^{-1})z^{-N} \tag{13}$$

The output noise variance of Eq. (1) becomes

$$\eta = \frac{1}{2\pi j} \oint |S(z)D(z)|^2 \frac{dz}{z} \ , \tag{14}$$

and the squared output signal component $\mathcal{S}$ is given by Eq. (3) as

$$\mathcal{S} = \left[\frac{1}{2\pi j} \oint |S(z)D(z)|^2 \frac{dz}{z} \right]^2 \tag{15}$$

Thus the greatest (optimum) output S/N power ratio is

$$\rho_{opt} \quad = \quad (\mathcal{S}/\eta)_{opt} \quad = \quad \frac{1}{2\pi j} \oint |S(z)D(z)|^2 \frac{dz}{z} \tag{16}$$

For the white noise case, the matched filter of Eq. (13) becomes

$$H(z) \quad = \quad S(z^{-1})z^{-N} \tag{17}$$

When the noise spectrum is unknown, it would be reasonable to use this matched filter. Then, if the actual (but unknown) spectrum is given by Eq. (11), the output S/N power ratio would be

$$\rho_w \quad = \quad (\mathcal{S}/\eta)_w \quad = \quad \frac{\left[\frac{1}{2\pi j} \oint |S(z)|^2 \frac{dz}{z} \right]^2}{\frac{1}{2\pi j} \oint \left| \frac{S(z)}{D(z)} \right|^2 \frac{dz}{z}} \tag{18}$$

The adaptive scheme to be considered will be discussed in the next section and is based on approximating the polynomial D(z) of Eq. (12) using estimates $\hat{a}_i$ of the p parameters α_i in that equation:

$$\hat{D}(z) \quad = \quad 1 + \hat{\alpha}_i z^{-1} \ + \ \ . \ . \ . + \hat{\alpha}_p z^{-p} \tag{19}$$

The adaptive matched filter $\hat{H}(z)$ will then be given by Eq. (13) with D(z) replaced by $\hat{D}(z)$. The S/N power ratio at the output of this filter when the input noise is actually given by Eq. (11) will be

$$\rho_{adapt} = (\mathcal{S}/\eta)_{adapt} = \frac{\left[\frac{1}{2\pi j} \oint |S(z)\hat{D}(z)|^2 \frac{dz}{z}\right]^2}{\frac{1}{2\pi j} \oint \left| \frac{S(z)\,[\hat{D}(z)]^2}{D(z)} \right|^2 \frac{dz}{z}} \qquad (20)$$

3. ADAPTIVE SCHEME

The adaptive scheme is based on the estimation of the p parameters of the noise power spectral density, $\underline{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_p)$. The computed estimate $\hat{\underline{\alpha}}$ of $\underline{\alpha}$ is used to adjust the parameters of the matched filter, as shown in Fig. 2. The estimation technique has been described previously [1] and is based on least mean square regression. The basic steps in the identification computation are as follows:

1. From N sample points of the input signal, the p + 1 mean lagged products are computed:

$$f_j = \frac{1}{N-j} \sum_{i=1}^{N-j} y_i y_{i+j} \qquad j = 0, \ldots, p. \qquad (21)$$

2. The pxp matrix $F = (f_{|i-j|};\ i, j = 1, \ldots, p)$ is formed and inverted.

3. The estimate $\hat{\underline{\alpha}}$ is computed from

$$\hat{\underline{\alpha}} = -F^{-1} \begin{bmatrix} f_1 \\ \cdot \\ \cdot \\ \cdot \\ f_p \end{bmatrix} \qquad (22)$$

It has been shown [1] that the error vector $\hat{\underline{\alpha}} - \underline{\alpha}$ is asymptotically normally distributed with zero mean and covariance matrix asymptotic to $F^{-1}/(N-p)$, [with the normalization implied by $\phi(z) = 1/D(z)D(z^{-1})$]. Thus the per unit standard deviations of the errors in the coefficients are roughly $N^{-1/2}$ which gives some basis on which to investigate the effect of parameter identification errors on S/N improvement.

In the one-pole case when the signal is a single pulse, S(z) = 1, the S/N improvement can be found directly in terms of the pole position α and the error in the coefficient measurement ϵ. It is

$$\frac{\rho_{adapt}}{\rho_w} = \frac{1+\alpha^2}{1-\alpha^2}\left(1 - \frac{2\epsilon^2}{1+\alpha^2} \cdot \frac{[(1-\alpha^2)(1-(\alpha+\epsilon)^2)] + \epsilon^2}{[(1-\alpha^2)(1+(\alpha+2\epsilon)^2)] + \epsilon^4}\right), \tag{23}$$

with

$$\frac{\rho_{opt}}{\rho_w} = \frac{1+\alpha^2}{1-\alpha^2} . \tag{24}$$

Figure 3 shows the average S/N improvement as a function of pole position for various errors which correspond to record lengths of roughly 200-1000 points. It is seen that when the noise is highly correlated ($|\alpha|$ near 1) the effects of measurement errors are not great and allow significant improvement over a single fixed filter. Only the case when $S(z) = 1$ will be considered in the remainder of this paper, since the introduction of signal parameters tend to obscure the results. The use of longer signals allows integration of signal energy and will usually lead to performance improvement. It is of interest to note that when the noise is nearly white ($|\alpha|$ near zero), the measurement errors result in a filter which is less effective than the fixed filter of Eq. (17). Thus, if the noise is relatively uncorrelated, it may not be profitable to attempt noise adaptation. This suggests that switching logic be used to allow adaptation only when measurements indicate a highly correlated noise background. The switching criteria would depend on the record length, the signal-to-noise ratios, and the relative costs of false alarms and misses.

When the noise spectral density has more than one pole, similar results obtain, although the calculation of average S/N improvement for all ranges of the parameters becomes inconvenient. Figure 4 shows the results for the two-pole case as a function of α_2 when $\alpha_1 = 0.5$. It is seen that for nominal errors in the coefficient identification the average S/N improvement is again a relatively insensitive function of measurement errors, and that good performance can be obtained when the noise is highly correlated. It is expected that this conclusion is valid in higher order cases.

4. EXPERIMENTAL RESULTS

In order to investigate the behavior of the adaptive matched filter under realistic conditions, the proposed adaptive scheme was implemented on an IBM 7094 computer. The correlated noise was obtained by filtering rectangularly distributed independent random numbers. Record lengths of 1000 points were used, and 10 single pulse signals were added to the noise sample at distributed points in time. Of importance are the actual distribution of the output S/N for various input S/N; the detection possibilities based on the output amplitude density distribution; and the behavior of the system when the noise is not all-pole (the structure problem).

Figures 5 and 6 show measured output S/N improvement in the one-pole case for low and high input S/N ranges, respectively. It is seen that the average S/N improvement agrees with the theoretical results of Fig. 3, but that the statistical spread depends heavily on the signal level. For small input S/N, the S/N improvement is rather unstable. This is to be expected since some signals in this case tend to become completely submerged in the noise. For larger signals the operation of the adaptive filter becomes more consistent, as would be expected intuitively. Fig. 7 shows a plot of the measured mean and variance of the S/N improvement for a single fixed pole at 0.9 as a function of the input S/N. The decrease in the variance of S/N improvement is quite pronounced. The average S/N improvement decreases with increasing signal strength because the signals present interfere with the noise spectrum identification and degrade the estimate of α_1.

To gain more insight into the operation of the filter under different signal level and noise conditions, the input and output amplitude densities for a 1000 point record and a one-pole noise spectral density were tabulated Figs. 8 and 9 show these densities under adverse conditions [relatively low correlation ($\alpha_1 = .7$) and relatively low signal strength (input S/N = 1.5)]. It is seen that in this case the signals that are not buried in the noise are recovered by the matched filter. The detection probability in this case is typically increased from 3/10 to 5/10 with a much lower test level (false alarm probability). Figs. 10 and 11 show the input and output amplitude densities when the signal strength is higher (input S/N = 2.0) and the noise more correlated ($\alpha = 0.9$). Here the improvement in detection capability is more pronounced, and the detection probability is raised from about 3/10 to very nearly 1 with a very small level. The "T" and "U" symbols at the top of Figs. 9 and 11 indicate the amplitude of the matched filter transient response at points adjacent to the signal. Thus, the false signals in these plots are associated with real signals and are not properly false alarms. This indicates a trade-off between detection capability and determination of time of occurrence.

Another important consideration is the behavior of the filter when different order approximants are used in the noise spectral density model. When the noise is actually all-pole it is important that the order of the approximant be at least as great as the actual order of the noise spectral density. This is illustrated in Fig. 12 which shows the measured mean and variance of S/N improvement for ten single pulse signals in 1000 record points as a function of the number of poles used in the noise model. The noise in this case was all-pole with 3 poles at $\pm j.8$ and $-.5$. For first and second order approximants to the noise spectral density, the performance is relatively poor; for a three-pole model the performance is relatively good; and for larger numbers of poles the performance is practically unaffected by the order of the noise model.

This is in contrast to the behavior of the filter when the noise is all zero. Figure 13 shows the measured mean and variance of the S/N improvement for ten signals when the noise background has the spectral density

$$\phi(z) = (1 - .9z^{-1})^2(1 - .9z)^2 \qquad (25)$$

In this case the performance improves steadily with the number of poles in the noise model and finally stabilizes for $p = 6$. Thus when the noise is not all-pole it is important to use a large number of parameters in the all-pole model. It is noteworthy that significant S/N improvement can be obtained in this case, despite the fact that the model chosen is inappropriate.

Since the performance of the adaptive matched filter depends largely on the noise being highly correlated, it is interesting to investigate the filter's behavior with a periodic, deterministic noise background and single pulse signals. This was done with 200 point records of repetitive sine, square, and triangle waves with 3 scattered single pulse signals. A third order noise model was assumed. The amplitude of the noise background was decreased by factors of 435, 748, and 2610 respectively, without decreasing the signal amplitude. This indicates a marked ability of the adaptive scheme to mimic highly regular structures and suggests the use of such schemes in pattern recognition problems.

5. CONCLUSIONS

This paper indicates that useful noise adaptation can be achieved for signals in correlated noise with an all-pole model for the noise spectral density. The qualitative performance of the adaptive loop of Fig. 2 can be described as follows:

1. Performance is relatively insensitive to identification errors, thus allowing short measurement times and fast adaptation in quasi-stationary environments.

2. The adaptive matched filters can be useful in conjunction with threshold detection if the noise is highly correlated.

3. The use of high order all-pole models is justified for a wide variety of noise spectral densities, provided that the time required to measure p + 1 mean lagged products and invert a pxp matrix is tolerable.

4. The adaptive matched filter is very effective in removing highly regular deterministic patterns from random signals.

REFERENCES

1. K. Steiglitz, "Power Spectrum Identification for Adaptive Systems," IEEE Trans. on Applications and Industry, Vol. 83, No. 72, pp. 195-97; May, 1964.

2. C.W. Helstrom, Statistical Theory of Signal Detection, Pergamon Press, Inc., New York, N.Y.; 1960, Chap. IV.

3. W.B. Davenport and W.L. Roots, An Introduction to the Theory of Random Signals and Noise, McGraw-Hill Book Company, Inc., New York, N.Y.; 1958, Chap. II.

4. L.A. Zadeh and J.R. Ragazzini, "An Extension of Wiener's Theory of Prediction," Jour. Appl. Physics, Vol. 21, pp. 645-655; July, 1950.

5. S.S.L. Chang, "Statistical Design Theory for Strictly Digital Sampled-Data Systems," Trans. AIEE, Communication and Electronics, Vol. 76, pp. 702-709; 1957.

6. S.S.L. Chang, "Statistical Design Theory for Digital Controlled Continuous Systems," Trans. AIEE, Applications and Industry, Vol. 77, pp. 191-201; 1958.

7. S.S.L. Chang, Synthesis of Optimum Control Systems, McGraw-Hill Book Company, Inc., New York, N.Y.; 1961, Chap. 6, p. 139.

8. L.A. Wainstein and V.D. Zubakov, Extraction of Signals from Noise, Prentice-Hall, Inc., Englewood Cliffs, N.J.; 1962, Sec. 25.

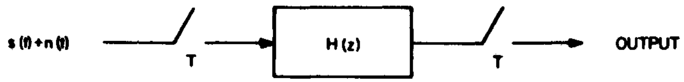

Fig. 1 - Block diagram of digital matched filter.

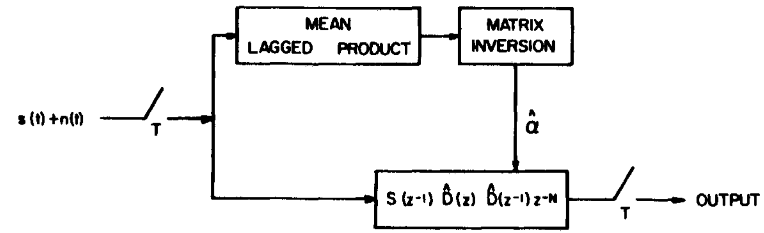

Fig. 2 - Block diagram of adaptive digital matched filter.

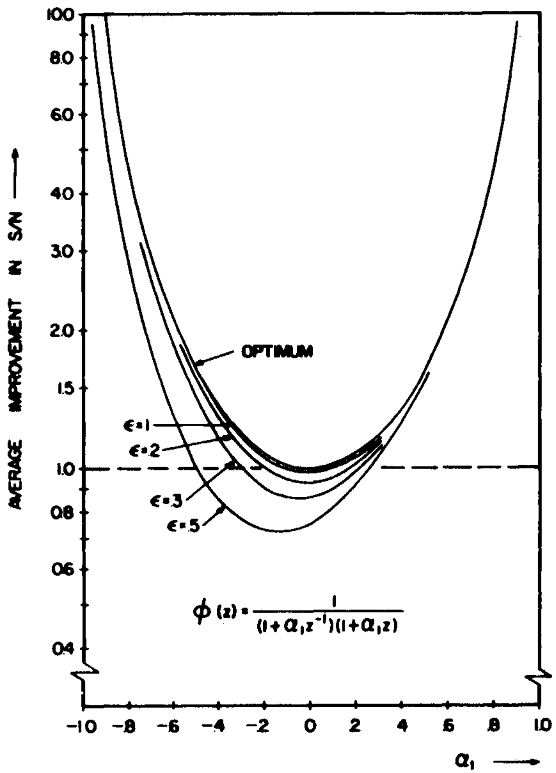

Fig. 3 -
Average improvement in S/N vs. pole position with no measurement error and with typical measurement errors for a one-pole noise spectral density.

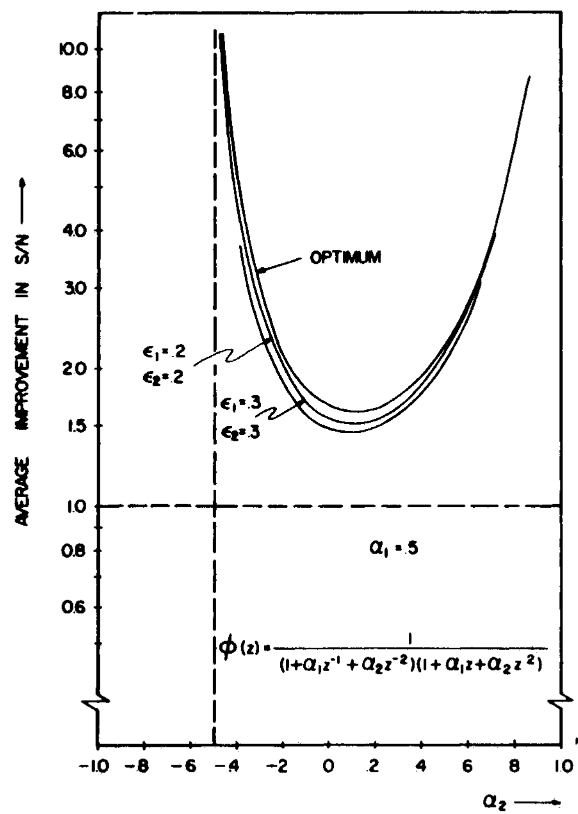

Fig. 4 -
Average improvement in S/N vs. α_2 with α_1=0.5 with no measurement errors and with typical measurement errors for a two-pole noise spectral density.

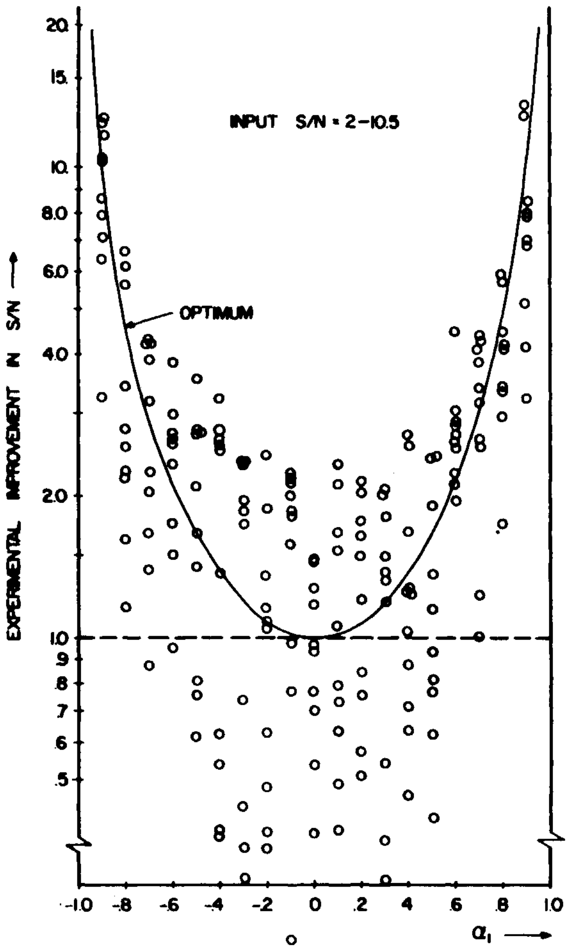

Fig. 5 -
Measured S/N improvement vs. pole position for a one-pole noise spectral density. Each experimental point represents one single pulse signal in a record of 1000 points. Input S/N= 2 - 10.5 .

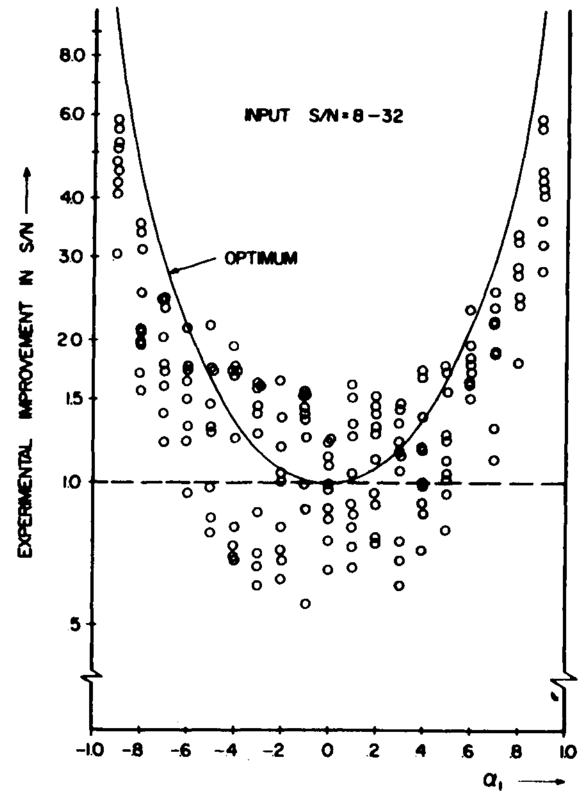

Fig. 6 - Same as Fig. 5 except input S/N= 8 - 32 .

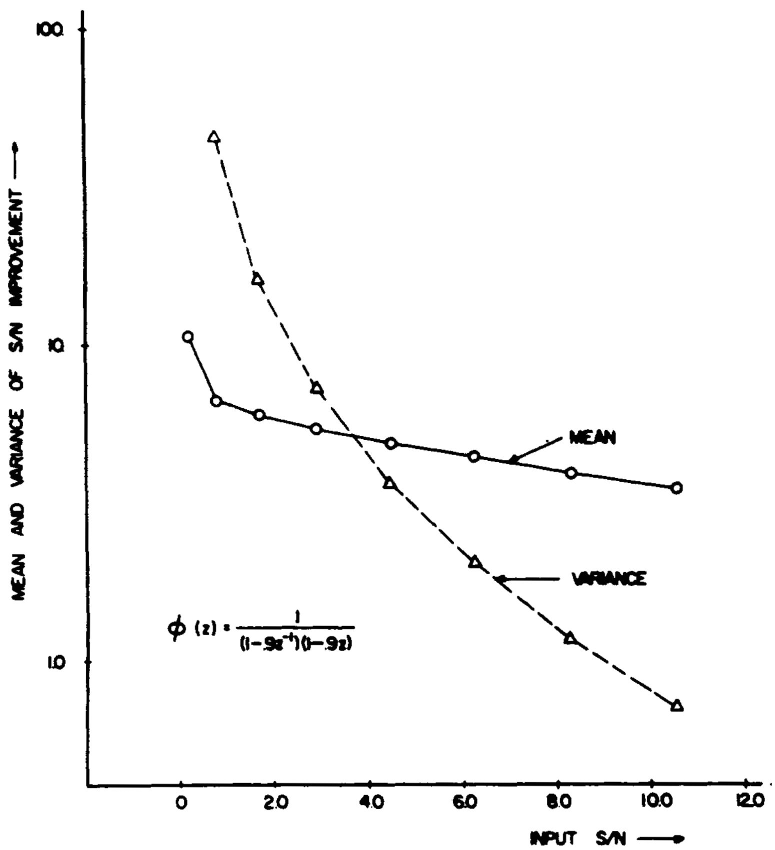

Fig. 7 - Measured mean and variance of S/N improvement for 10 single pulse signals in a 1000 point record vs. input S/N for fixed one-pole noise.

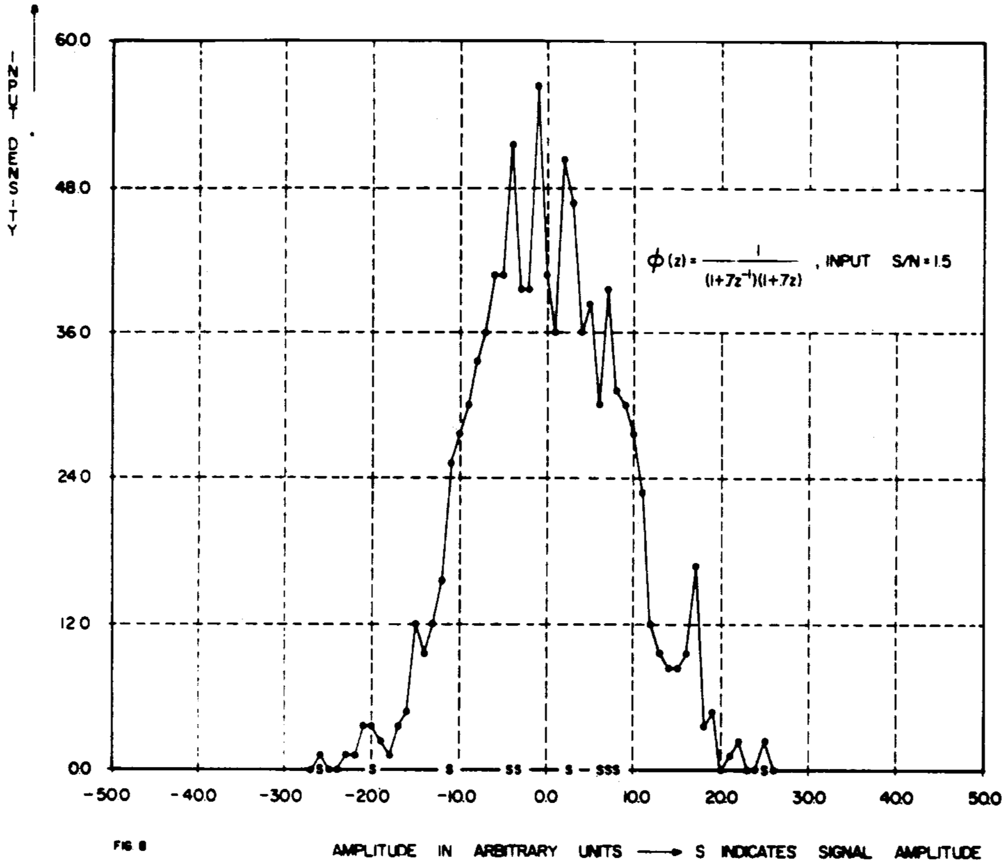

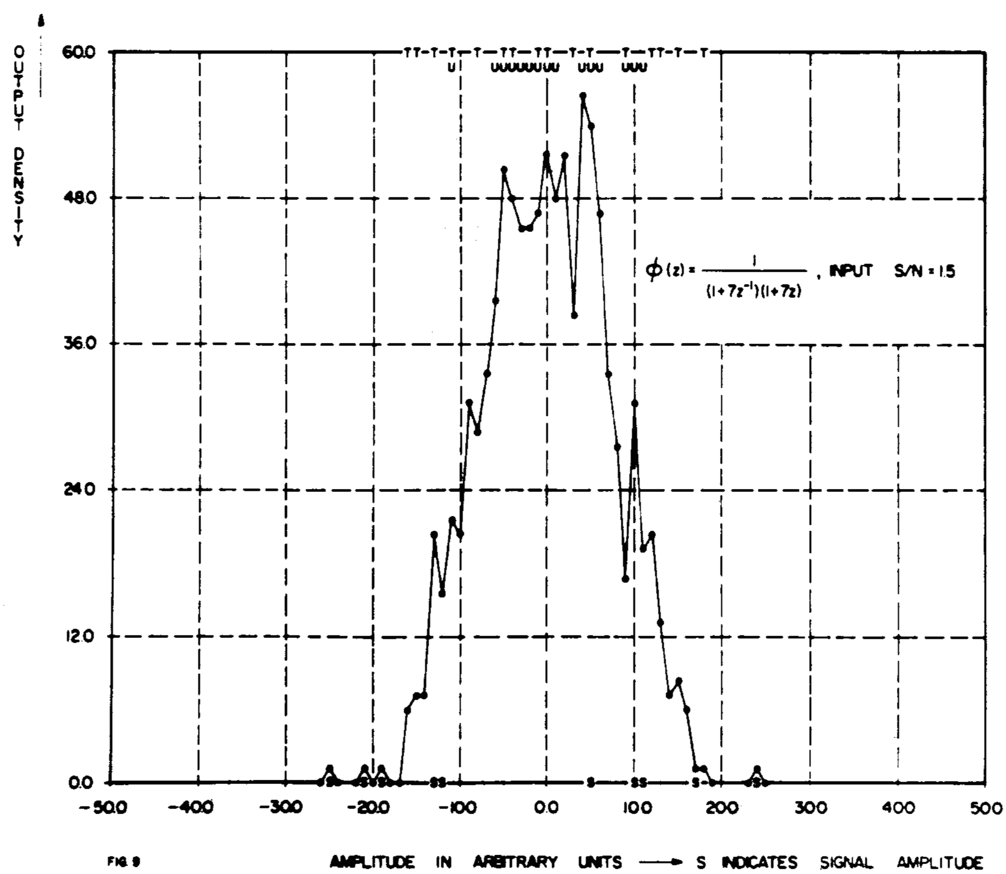

Figs. 8 and 9 - Measured input and output amplitude density for a 1000 point record with relatively uncorrelated one-pole noise and input S/N = 1.5 .

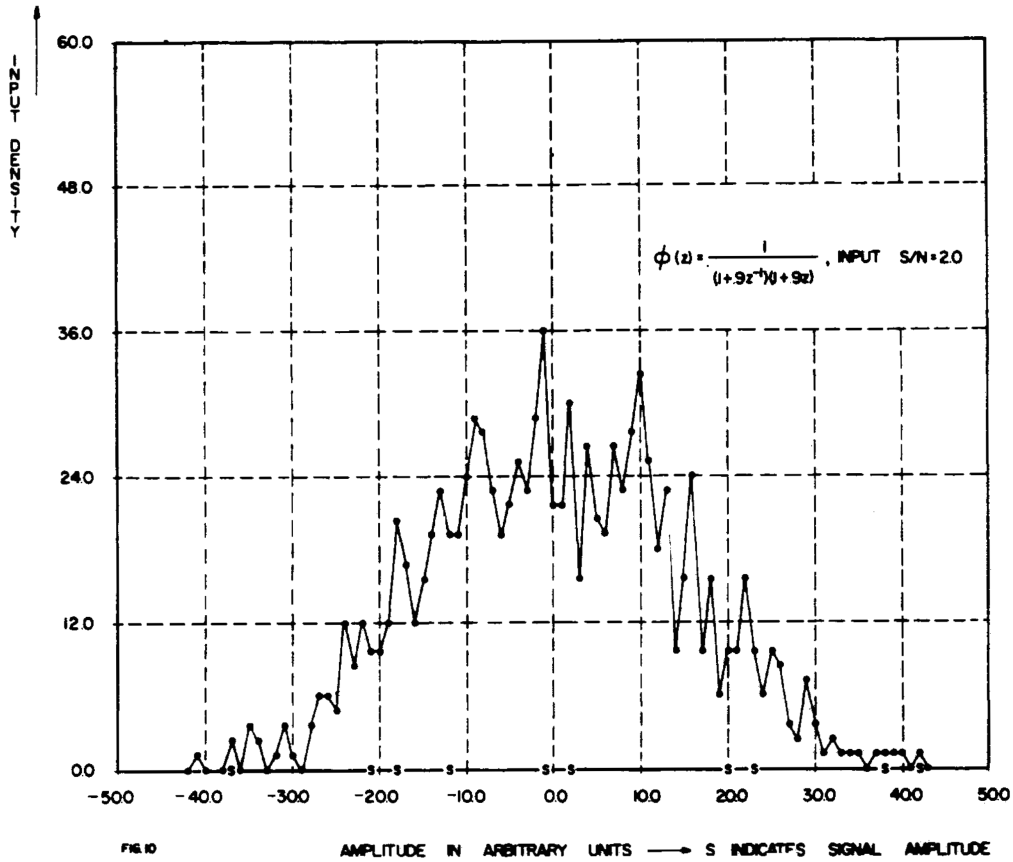

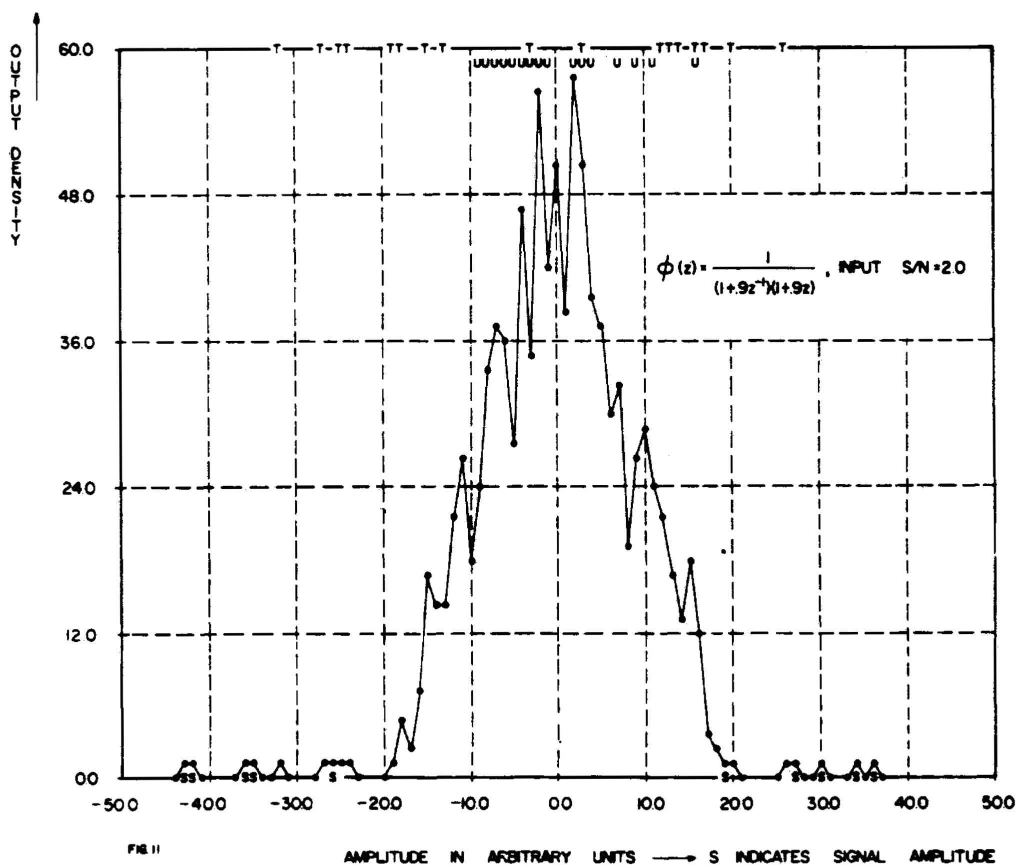

Figs. 10 and 11 - Same as Figs. 8 and 9 except the noise is more correlated and input S/N = 2.0

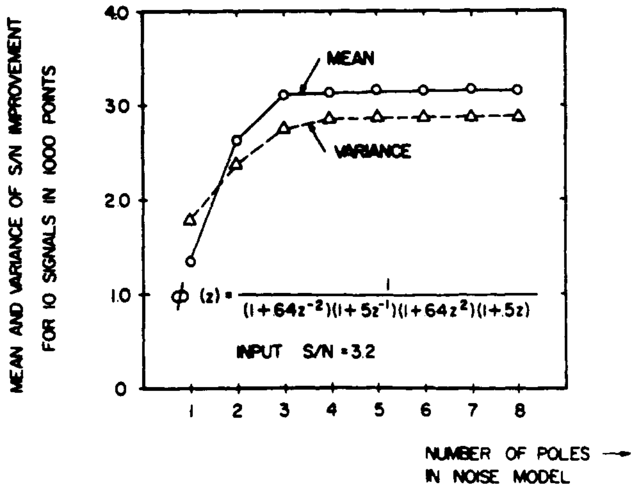

Fig. 12 - Measured mean and variance of S/N improvement for 10 single pulse signals in a 1000 point record vs. number of poles in the noise model.Noise spectral density is all-pole.

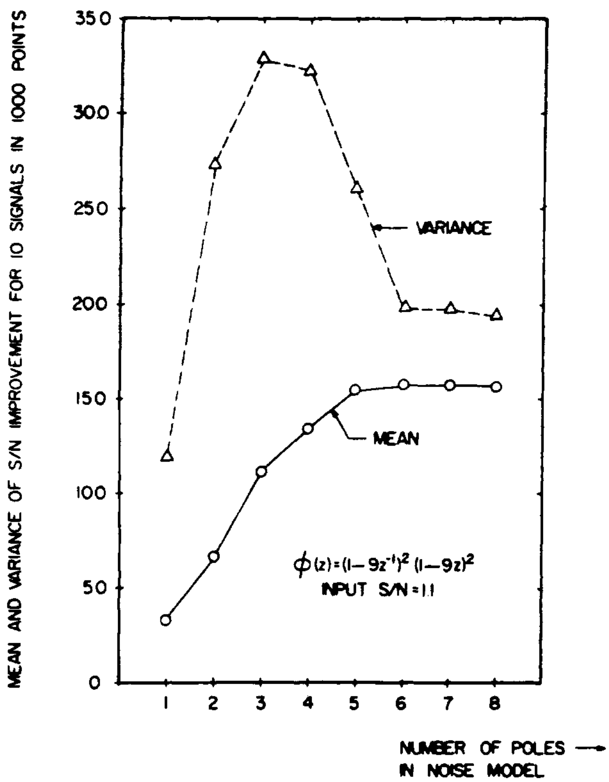

Fig. 13 - Measured mean and variance of S/N improvement for 10 single pulse signals in a 1000 point record vs. number of poles in the noise model. Noise spectral density is all-zero.