Defence Research Establishment Pacific

Centre de recherches pour la défense pacifique

DEFENCE RESEARCH ESTABLISHMENT PACIFIC

CFB Esquimalt, FMO Victoria, B.C. V0S 1B0

Technical Memorandum 93-11

THE DETECTION PERFORMANCE OF A LONG-TERM SPECTRAL INTEGRATOR USING NORMALIZED AND QUANTIZED DATA

by

B.H. Maranda

February 1993

Approved By:

Chief DREP

Research and Development Branch
Department of National Defence

Canada

ABSTRACT

The effect of normalization and quantization on the detection performance of a long-term spectral integrator is examined. The performance is described by the receiver operating characteristic, which is computed numerically from the probability density function of the detection statistic. It is found that the loss in performance due to normalization is small, less than 0.2 dB for an example case. The loss due to quantization depends on the number of bits used. When the quantizer uses only one bit (two levels), the loss is approximately 1 dB. This loss can be decreased to less than 0.1 dB by the use of three bits (8 levels) in the quantizer.

RÉSUMÉ

L'effet de la normalisation et de la quantification sur la performance de détection d'un intégrateur spectral de longue durée est examiné. La performance est décrite par la caractéristique opérationelle de réception, qui est calculée numériquement au moyen de la densité de probabilité à la sortie du détecteur. La perte de performance due à la normalisation est menue, moins de 0,2 dB dans un exemple. La perte due à quantification dépend du nombre de bits utilisés. Dans le cas où le quantificateur n'utilise qu'un bit (deux niveaux), la perte est d'environ 1 dB. Cette perte peut être réduite à moins de 0,1 dB en utilisant 3 bits (8 niveaux) de quantification.

1. INTRODUCTION

The standard method for the detection of narrowband signals is to use an energy detector implemented by a fast Fourier transform (FFT) followed by a squarer.[1,2] For purposes of display and storage, the FFT spectra are normalized and then quantized; often a short-term noncoherent integration, consisting of the addition of only a few FFT spectra, will be performed before the normalization and quantization. A long-term noncoherent integration is then attained visually by stacking the quantized spectra in a frequency-time-intensity (FTI) display, referred to as a lofargram by the sonar community.[1] To minimize storage requirements it is desirable to quantize the data heavily, and studies based on human performance have shown that little is to be gained by using more than 3 bits, or 8 levels, in an FTI display.[3,4]

More recently, an algorithm has been proposed in which the long-term integration is performed not visually on a display, but internally in a computer.[5] For a beamformed system, this algorithm allows much greater flexibility in the path of integration through the sonar spectra: the integration is no longer limited to a single beam, but can pass from one beam to another. However, the algorithm was implemented in Ref. 5 by summing unnormalized spectra with floating-point arithmetic, whereas an operational system would likely be restricted to using the normalized and quantized spectra that reside in a sonar database. The small word-width of the quantized data would be an advantage in the implementation of a dedicated hardware summation device, but if the quantization were too coarse the sensitivity of the detector would be degraded. The purpose of this paper is to examine the effect of both normalization and quantization on the narrowband detection performance of a long-term spectral integrator, in order to determine whether heavily quantized spectra give adequate performance. This question is important, for the object of the long-term integration is to allow the detection of targets that cannot be detected on the short-term spectra.[5] The studies of human performance cited earlier, indicating that 3-bit quantization is sufficient, cannot be said to answer the question, since the result may be a property of the human visual system (e.g., an inability to resolve more than 8 intensities with any fidelity.) The current study, based on a complete mathematical analysis of the end-to-end processing system, provides a more concrete answer for the computerized system envisioned here. In fact, the results should provide upper bounds on the detection performance of a human operator viewing lofargrams.

2. DETECTOR STRUCTURES

The performance will be compared for the three detector structures shown schematically in Fig. 1. Detector 3 is of most interest, as it most closely models the processing done in practice, but analyzing the sequence of detectors shown in Fig. 1 allows the in-

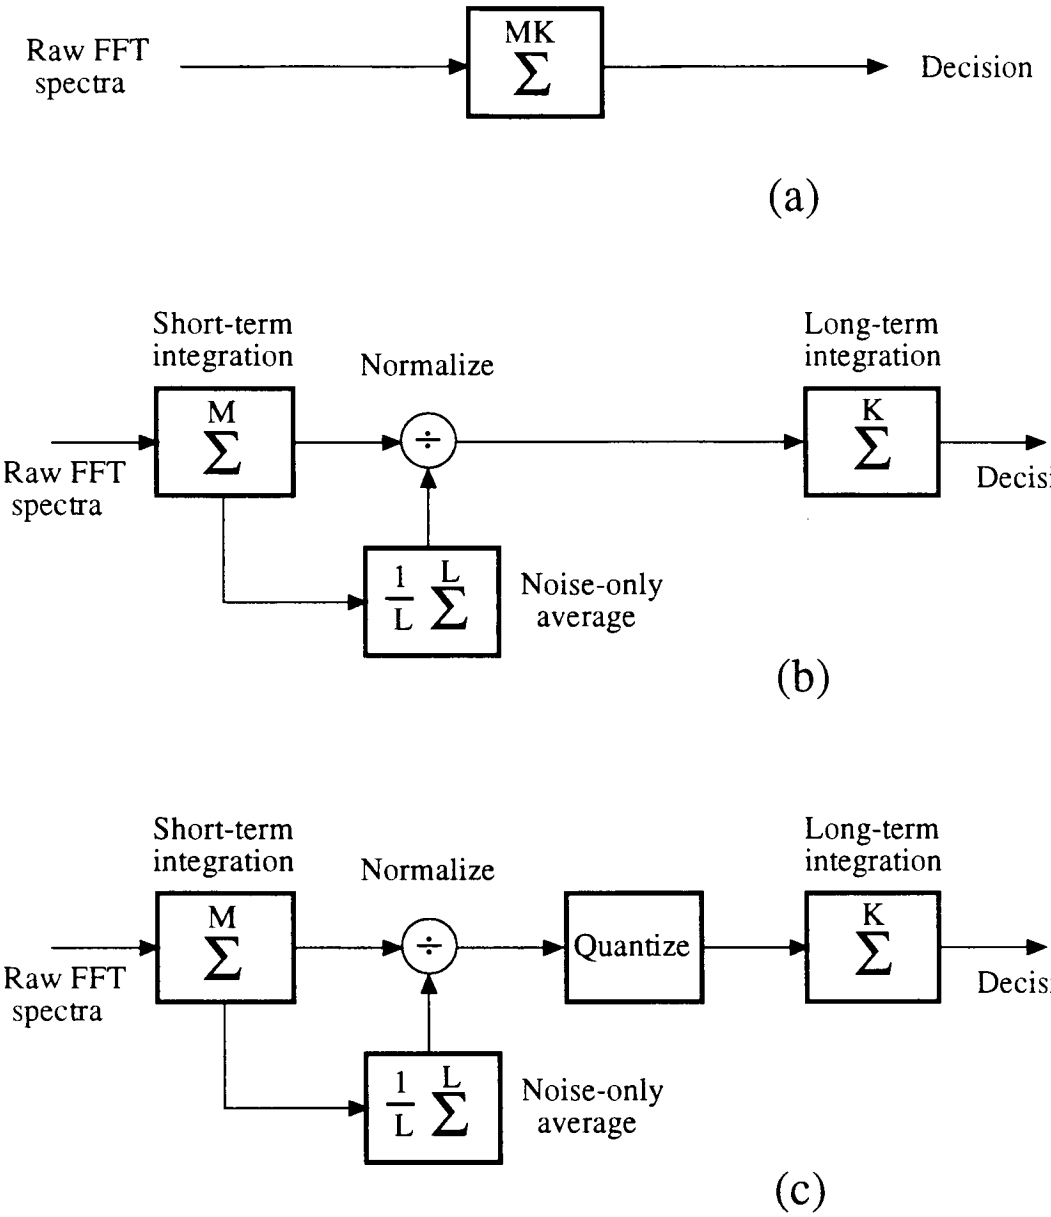

Figure 1. Schematics of the detectors used in the performance analysis. (a) Detector 1, (b) Detector 2, (c) Detector 3.

dividual sources of performance loss to be identified. Detector 1 is the baseline system, consisting simply of an integrator of unnormalized and unquantized spectra. Detector 2 considers the degradation in performance caused by normalization alone, and Detector 3 considers the further effect of quantization. Performance analysis for similar systems can be found in the literature[6–8], but the analysis provided here, and the models used for the detectors, are believed to be more complete than those available elsewhere.

A short overview of the most complex detector, Detector 3, will now be given. (It should then be obvious how the simpler detectors work, since they merely omit parts of Detector 3.) First, a block of sampled data is transformed into the frequency domain using the fast Fourier transform, and the squared magnitude of each bin is computed. A single spectrum computed in this way will be called a raw spectrum. Next, the raw spectra from M consecutive data blocks are summed to yield an integrated spectrum. It is assumed that M is small, typically less than 10, and these integrated spectra will be referred to as short-term spectra. (In the case $M = 1$, the short-term and raw spectra are the same.) The spectral bins are then normalized and quantized to discrete levels. Finally, a total of K quantized spectra are integrated, where K is large (on the order of 100), and a decision is made on this long-term integrated value. It is the K normalized and quantized short-term spectra that in a conventional sonar processor are presented in an FTI display and stored in the database.

In the next section, the signal and noise models are discussed, and the analytical and numerical techniques required for the performance analysis of the detectors are derived. Readers who are interested only in the numerical results can proceed to Section 4.

3. METHODS OF PERFORMANCE ANALYSIS

In order to determine the detection performance of the given detectors, a statistical model for the noise and the signal is required. It will be assumed that the input noise is Gaussian distributed with zero mean. The real and imaginary parts of an FFT bin are then also Gaussian distributed with zero mean, and their common variance will be denoted by σ^2. The signal will be taken to be a sinusoid located at the centre of an FFT bin, and its phase is assumed to be uniformly distributed from block to block. It will be assumed that the input data blocks are not overlapped, so that successive FFT spectra are independent. In what follows, x will be used as a generic symbol to denote a random variable appearing at any point in the three detectors just described. When the discussion requires a more specific notation, modifications will be made to the basic notation; e.g., $\tilde{x}$ will represent a normalized variable. When it is necessary to distinguish between noise-only and signal+noise, we shall adhere to the standard nomenclature that H_0 is the noise-only hypothesis and H_1 the signal+noise hypothesis; hence $p_x(\cdot|H_0)$

denotes the probability density function (pdf) of the random variable x when noise alone is present. However, for simplicity, we shall often use the variables n and s (and the forms $\tilde{n}$, $\tilde{s}$, etc.) to refer to the noise and signal+noise cases respectively. For example, we may also write $p_n(\cdot) = p_x(\cdot|H_0)$.

All of the detectors described above form an output decision variable, x, which is compared with a threshold, T. When $x \geq T$ a detection is declared, and when $x < T$ noise alone is declared. (Note that the decision variable is a continuous random variable for the first two detectors, and is discrete in the quantized system; for the discrete case, the integrals that follow are replaced by summations.) The performance of each detector will be described by the probability of false alarm, P_{fa}, and the probability of detection, P_d. Defining

$$P_j = \int_T^\infty p_x(\alpha|H_j)\, d\alpha, \tag{1}$$

we have $P_{fa} = P_0$ and $P_d = P_1$. In order to calculate the probabilities of detection and false alarm the statistics of the decision variable x must be known, and this is the primary focus of the rest of the paper. In some cases it was not possible to express the pdf for x in a closed form, and it was then necessary to proceed numerically using characteristic functions. The characteristic function of the random variable x is defined by $\phi(\xi; x) = E\{\exp(i\xi x)\}$, where $E\{\cdot\}$ denotes expectation. The inversion formula of Fourier analysis then yields

$$p_x(\alpha) = \frac{1}{2\pi} \int_{-\infty}^\infty \phi(\xi; x) e^{-i\xi\alpha}\, d\xi, \tag{2}$$

whence P_{fa} and P_d may be evaluated using Eq. (1). Fortunately, it is possible to compute P_d and P_{fa} directly from the characteristic function $\phi(\xi; x)$ without computing the pdf $p_x(\cdot)$ as an intermediate result. There are several numerical methods that one may use, but since a numerical software package[9] with highly accurate quadrature routines was available, the desired probabilities were computed by direct numerical integration of the formula[10]:

$$P_j = \frac{2}{\pi} \int_0^\infty \mathrm{Im}\{\phi(\xi; x|H_j)\} \frac{\cos \xi T}{\xi}\, d\xi. \tag{3}$$

It might appear that the integrand in Eq. (3) is $O(\xi^{-1})$ as $\xi \to 0^+$, so that the integral diverges, but it can be shown that the integrand approaches a finite limit (the mean of x) at the origin. The usefulness of characteristic functions stems from the well-known fact that the characteristic function of a sum of independent random variables is the product of the characteristic functions of the individual variables; the application details will be given in the next section.

Since the short-term integration of raw FFT spectra is common to two of the detectors considered above, we write down the statistics of the short-term spectra for future reference. In the noise-only case, each bin of a short-term spectrum has a chi-squared (χ^2) distribution with $2M$ degrees of freedom[11]; the density function is

$$p_n(\alpha) = \frac{\alpha^{M-1} e^{-\alpha/2\sigma^2}}{(2\sigma^2)^M \, \Gamma(M)}, \qquad \alpha \geq 0, \tag{4}$$

with corresponding characteristic function

$$\phi(\xi; n) = (1 - 2i\sigma^2\xi)^{-M}. \tag{5}$$

The signal+noise bin has a non-central χ^2 distribution[11] with pdf

$$p_s(\alpha) = \frac{1}{2\sigma^2} \left(\frac{\alpha}{M\rho^2} \right)^{(M-1)/2} \exp\left(-\frac{M\rho^2 + \alpha}{2\sigma^2} \right) I_{M-1}\left(\frac{\sqrt{\alpha \, M\rho^2}}{\sigma^2} \right), \qquad \alpha \geq 0, \tag{6}$$

and characteristic function

$$\phi(\xi; s) = e^{-M\gamma}(1 - 2i\sigma^2\xi)^{-M} \exp\left(\frac{M\gamma}{1 - 2i\sigma^2\xi} \right), \tag{7}$$

where $I_\nu(\cdot)$ is the modified Bessel function[12] of the first kind and order ν, and $M\rho^2$ is the non-centrality parameter. Here ρ^2 is the magnitude squared of the signal amplitude, and the signal-to-noise ratio (SNR) has been defined by $\gamma = \rho^2/2\sigma^2$. This is the output SNR in the bin containing the signal for a raw spectrum.

3.1. Detector 1: Integration of raw spectra

The expressions for the density functions obtained after integrating $N \equiv MK$ raw FFT spectra follow immediately from the above, since the χ^2 distribution exhibits the *reproducing* property: the sum of χ^2 variables remains χ^2, but with an increased number of degrees of freedom. Therefore, it is only necessary to replace M by N in Eqs. (4) and (6) to obtain the required densities. The probabilities of false alarm and detection are then given by

$$P_{fa} = Q(N, T/2\sigma^2), \tag{8}$$

$$P_d = e^{-N\gamma} \sum_{k=0}^{\infty} \frac{(N\gamma)^k}{k!} Q(N + k, T/2\sigma^2), \tag{9}$$

where Q is the complement of the incomplete Gamma function, defined as

$$Q(N, x) = \frac{1}{\Gamma(N)} \int_x^\infty \alpha^{N-1} e^{-\alpha} d\alpha. \tag{10}$$

The series for P_d in Eq. (9) is derived by inserting the power series for the Bessel function of Eq. (6) into the integral (1) and then interchanging the order of summation and integration. The probabilities P_{fa} and P_d depend on the parameters T, σ^2 and ρ^2 only through the ratios T/σ^2 and $\gamma = \rho^2/2\sigma^2$. Note that, given a value of T/σ^2 that yields a pre-assigned value of P_{fa}, it is necessary to know the noise level σ^2 in order to adjust the threshold T to attain this pre-assigned value. That is, the detector as given here is for the case in which the noise level σ^2 is known; this ideal case provides the basis for comparison with the more practical detectors that follow.

A routine for computing Q is available in the NAG library of numerical algorithms[9], and Eq. (8) was used directly for numerical evaluation. The infinite series in Eq. (9) is more difficult to use for numerical purposes, since in the parameter ranges of interest the terms increase for perhaps the first hundred indices k before decreasing, and the dynamic range required is considerable. However, by careful scaling of the intermediate results, Eq. (9) was used successfully for rapid and accurate computation of P_d. The results from Eq. (9) were checked using another NAG routine, which computes lower tail probabilities for a non-central χ^2 variate. (These probabilities are converted to P_d by subtracting from unity, thus giving reduced accuracy for small values of P_d.) The results generally agreed to more than 8 significant figures.*

3.2. Detector 2: Integration of normalized spectra

The purpose of the normalizer is to remove the dependence on the unknown noise variance σ^2 by estimating it from the data itself. (This estimate is made indirectly, as the normalizer actually divides by an estimate $\hat{\mu}_n$ of the true mean μ_n of a noise-only spectral bin.) In practice, the normalizer has another task that is just as important: reducing the dynamic range of the short-term spectra. This would be unnecessary if the input noise were indeed white, as a single estimate $\hat{\mu}_n$ would characterize the noise background across the whole frequency band. In actuality, at low frequencies ocean ambient noise can vary in level by over 10 dB in a band 100 Hz wide.[13] For this reason the normalizer is designed to estimate only a local mean, under the assumption that that noise spectrum is locally white (i.e., that it varies little over 20–30 adjacent bins). The estimated local mean can then be used to reduce the often large dynamic range of the FFT spectra before quantization, in this manner allowing fewer bits to be used.

A complete statistical analysis of the normalization schemes used in practice does not seem to be tractable, as such schemes may involve multiple passes over the spectral data,

* All numerical calculations were done in double-precision arithmetic on a VAX computer, unless otherwise noted.

comparisons and replacements, etc.[14] However, the complications of most normalization techniques result from their attempt to handle departures in the data from an ideal model (such as caused by unwanted narrowband interference), and if it is assumed that these more complicated procedures handle these departures well, then the output of the normalizer is approximately the same as a simple normalizer working with ideal data. Thus we shall assume that the estimate $\hat{\mu}_n$ is equivalent to the sum of L independent noise-only bins, divided by L. The effect is to increase the degrees of freedom in the χ^2 statistics, and from Eq. (4) we can immediately write down the pdf for $\hat{\mu}_n$:

$$p_{\hat{\mu}_n}(\alpha) = \left(\frac{L}{2\sigma^2}\right)^{ML} \frac{\alpha^{ML-1}}{\Gamma(ML)} \exp\left(-\frac{\alpha L}{2\sigma^2}\right), \qquad \alpha \geq 0. \tag{11}$$

Next, we desire to find the density functions for the normalized random variables $\tilde{n} \equiv n/\hat{\mu}_n$ and $\tilde{s} \equiv s/\hat{\mu}_n$, denoted generically by $\tilde{x} = x/\hat{\mu}_n$. The density function for the ratio of two independent random variables can be expressed in integral form[15], and we obtain

$$p_{\tilde{x}}(\alpha) = \int_0^\infty w\, p_{\hat{\mu}_n}(w)\, p_x(\alpha w)\, dw, \qquad \alpha \geq 0. \tag{12}$$

Substituting Eqs. (4) and (11) into the integral of Eq. (12) yields the expression

$$p_{\tilde{n}}(\alpha) = \frac{1}{B(M, ML)} \frac{(\alpha/L)^{M-1}}{L\,(1+\alpha/L)^{ML+M}}, \qquad \alpha \geq 0, \tag{13}$$

where $B(\cdot,\cdot)$ is the Beta function[12]. The density in Eq. (13) corresponds to the F–distribution, or Beta distribution of the second kind. As stated above, the purpose of the normalizer is to remove the dependence on the unknown variance σ^2 of the noise, and it may be noted that the density $p_{\tilde{n}}$ is independent of σ^2. It follows that the probability of false alarm can be set to a desired value with no knowledge of σ^2; such a detector is known as a constant false-alarm rate (CFAR) detector.[16] For the signal+noise case, the integral in Eq. (12) leads to the pdf

$$p_{\tilde{s}}(\alpha) = \frac{e^{-M\gamma}(\alpha/L)^{M-1}}{L\,(1+\alpha/L)^{ML+M}} \sum_{k=0}^{\infty} \frac{1}{k!\, B(M+k, ML)} \left(\frac{\alpha M\gamma}{L+\alpha}\right)^k, \qquad \alpha \geq 0. \tag{14}$$

This density function corresponds to the non-central F–distribution[16] and depends only on the SNR, γ, and not on the absolute levels of the signal or noise. Note that Eq. (14) correctly reduces to the noise-only pdf of Eq. (13) when $\gamma = 0$. The distributions in Eqs. (13) and (14) can also be found using the Mellin transform.[17]

The next stage in the analysis is to consider the summation of K normalized spectra. However, the distributions in this case do not enjoy the reproducing property of the χ^2 distribution. We shall proceed in the usual way for sums of independent variates, using characteristic functions. Let $\Sigma_{\tilde{x}}$ denote the sum of K normalized FFT bins. Since the bins from successive FFTs are independent and identically distributed, the characteristic function of $\Sigma_{\tilde{x}}$ is given by

$$\phi(\xi; \Sigma_{\tilde{x}}) = \phi(\xi; \tilde{x})^K. \tag{15}$$

The specific characteristic functions for $\Sigma_{\tilde{n}}$ and $\Sigma_{\tilde{s}}$ were evaluated from Eq. (15) and then used in the integral in Eq. (3) for the numerical evaluation of P_{fa} and P_d. Efficient numerical evaluation of $\phi(\xi; \tilde{x})$ is therefore important, because it typically must be evaluated for hundreds of values of ξ during the quadrature of Eq. (3). Accuracy is also important, since raising to the power K amplifies the error by the same factor. To illustrate this, let $\hat{a}$ denote the computed value of a true value a; then we may write $\hat{a} = (1 + \delta)a$, where δ is the relative error. Assuming that $|\delta| \ll 1$ we have from the binomial expansion that $\hat{a}^K \cong (1 + K\delta)a^K$, showing that the relative error has been magnified by the factor K, i.e., by a factor of up to 400 for the examples considered subsequently.

Unfortunately, even the more basic characteristic functions $\phi(\xi; \tilde{n})$ and $\phi(\xi; \tilde{s})$, required to form the characteristic functions for $\Sigma_{\tilde{n}}$ and $\Sigma_{\tilde{s}}$ via Eq. (15), are not expressible in terms of simple functions[18], and some effort must be expended in determining an efficient method for their accurate evaluation. We first consider the computation of the characteristic function of $\tilde{n}$. From Eq. (13) we may write

$$\phi(\xi; \tilde{n}) = \frac{1}{B(M, ML)} \int_0^\infty \frac{w^{M-1}}{(1+w)^{ML+M}} e^{i\xi L w} dw. \tag{16}$$

A recurrence relation has been developed for $\phi(\xi; \tilde{n})$ from this integral[18], but it is necessary to start the recursion with a sine or cosine integral, which in turn must be numerically evaluated in some fashion. Furthermore, the numerical stability of the recursion would have to be examined. However, Eq. (16) is suitable for direct numerical quadrature, due to the rapid decay of the rational function in the integrand (here $ML \gg M$).

We turn now to the computation of $\phi(\xi; \tilde{s})$. The direct approach of taking the Fourier transform of Eq. (14) does not lead to an easy numerical problem, since the function in Eq. (14) is itself difficult to evaluate. Thus a completely different approach was used for the numerical evaluation of $\phi(\xi; \tilde{s})$; in the end, this different approach was also used for the evaluation of $\phi(\xi; \tilde{n})$ instead of Eq. (16), since then a single numerical procedure could be used. The equation on which this formulation is based comes from taking the

Fourier transform of Eq. (12) to find that[15]

$$\phi(\xi;\tilde{x}) = \int_0^\infty p_{\hat{\mu}_n}(w)\,\phi\left(\frac{\xi}{w};x\right)dw,$$

which gives, after inserting the expression for $p_{\hat{\mu}_n}$ from Eq. (11) and changing variables,

$$\phi(\xi;\tilde{x}) = \frac{1}{\Gamma(ML)}\int_0^\infty y^{ML-1}e^{-y}\,\phi\left(\frac{L\xi}{2\sigma^2 y};x\right)dy. \tag{17}$$

This integral can be efficiently evaluated using Gauss-Laguerre quadrature, which was the method eventually selected for evaluating the characteristic functions. Details on the computer implementation of the Gauss-Laguerre quadrature are presented in Appendix A.

3.3. Detector 3: Integration of normalized and quantized spectra

To store the normalized spectra in the database of a sonar system, they must first be quantized.* In order to store as many spectra as possible, it is desired to make the quantization as coarse as possible, while at the same time incurring only a small loss in performance. It has been found that three bits (8 levels) are sufficient to give good performance in an FTI display[3,4], and a three-bit quantizer is implemented in the experimental passive sonar system being developed at DREP and described in Ref. 19. We shall examine the effect of different quantizers to ascertain whether three bits remain sufficient when the long-term integration is performed by computer instead of visually.

The quantizer is described by a set $\{x_1, x_2, \ldots, x_{B-1}\}$ of $B-1$ breakpoints that partition the quantizer input into intervals, and a set $\{y_0, y_1, \ldots, y_{B-1}\}$ of B output levels onto which the intervals are mapped. It is assumed that $x_j < x_{j+1}$ for $j = 1, \ldots, B-2$, and the interval $[x_1, x_{B-1}]$ will be referred to as the quantization interval. The input/output relation of the quantizer can be described most succinctly by appending two breakpoints $x_0 = -\infty$ and $x_B = \infty$ to the set of breakpoints above; then the following relation holds:

$$q[x] = y_j \quad \text{when} \quad x_j < x < x_{j+1}.$$

The quantizer output can be assigned arbitrarily when x is exactly equal to a breakpoint, since this happens with probability zero. It will be assumed that the set of output levels is $\{0, 1, \ldots, B-1\}$, i.e., that $y_j = j$. As noted elsewhere[7], choosing output levels of the

* For the purpose of analysis, the high-precision floating-point arithmetic usually used for the signal processing prior to quantization can be viewed as infinite-precision arithmetic.

form $y_j = h_0 + j\,h$ for any constants h_0 and h would yield identical detection performance. Furthermore, although in general B can be any natural number, it will be chosen to be of the form $B = 2^b$, where b is the number of bits used to store the quantized data. This choice maximizes the number of the levels that can be encoded in b bits; for example, the level $y_j = j$ can be directly encoded by the binary representation of the integer j.

The numerical procedure for analyzing this case is somewhat simpler than for Detector 2, where no quantization is done. The reason is that the probability density function after the quantizer is discrete, and hence the characteristic function is easily calculated with the fast Fourier transform; algorithmic details can be found in the literature.[7] In the special case of one-bit quantization, the output of the quantizer obeys the binomial distribution and a simpler analysis than the FFT method can be used. The only remaining detail is the method used to compute the discrete pdf at the quantizer output. Let $F_{\tilde{x}}(\alpha) = \Pr\{\tilde{x} \le \alpha\}$ denote the cumulative distribution function of a normalized spectral bin $\tilde{x}$ at the input to the quantizer; the probability that the output will be assigned to level y_j is then given by $F_{\tilde{x}}(x_{j+1}) - F_{\tilde{x}}(x_j)$. From Eqs. (13) and (14) we know that $F_{\tilde{x}}$ is either the F–distribution or the non-central F–distribution, depending on whether the spectral bin contains noise alone ($\tilde{x} = \tilde{n}$) or signal+noise ($\tilde{x} = \tilde{s}$). A routine from the NAG library[9] was used to evaluate the required cumulative distributions. As a check, we also note that integration of Eq. (12) yields the expression

$$\begin{aligned} F_{\tilde{x}}(\alpha) &= \int_0^\infty p_{\hat{\mu}_n}(w) F_x(\alpha w)\, dw \\ &= \frac{1}{\Gamma(ML)} \int_0^\infty y^{ML-1} e^{-y}\, F_x(2\sigma^2 \alpha y / L)\, dy, \end{aligned} \tag{18}$$

where F_x is the distribution function for an unnormalized bin x from a short-term spectrum, i.e., χ^2 or non-central χ^2. The integral (18) has a form similar to the one in Eq. (17) and like that integral can be evaluated numerically using Gauss-Laguerre quadrature. The distribution function F_x can be evaluated as described in Sec. 3.1 with appropriate modifications (such as replacing N by M). The results obtained from (18) evaluated in this way agreed very well with those produced by the NAG routine.

4. NUMERICAL EXAMPLES OF DETECTOR PERFORMANCE

Based on the theory of the preceding section, computer programs were written to evaluate the probabilities of detection and false alarm for the three detector structures of Fig. 1. Some examples are now given; however, in order to prevent the number of examples from growing very large, the parameters have been chosen to represent a typical case (similar to that in Ref. 5). The total number of raw spectra is taken to be

$N = 400$, so that an FFT length of 9 sec would yield an observation time of one hour. For correct comparison with Detector 1, the parameters M and K in Detectors 2 and 3 must be chosen to have the product $N = MK$, where $M \ll K$. Unless otherwise stated, the selected values are $M = 4$ and $K = 100$.

4.1. Detector 1: Integration of raw spectra

Figure 2 shows the probability of detection for the baseline system as a function of SNR for several values of the false alarm probability. The value of P_{fa} for a given curve is given by its intercept with the left axis, which corresponds to the noise-only case. (In other words, P_d reduces to P_{fa} along the left axis, where the signal is absent). Fig. 3 shows the receiver operating characteristic (ROC) curves for the same detector. ROC plots, which show the probabilities of detection and false alarm for various SNRs, will be used in this memorandum as the standard format for displaying detector performance.

4.2. Detector 2: Integration of normalized spectra

Figure 4 shows the ROC curves for the detector when normalization, but not quantization, is implemented. The number of noise-only bins used in the normalizer was $L = 20$. As expected, the performance has been slightly degraded; i.e., the curve for a given SNR in Fig. 4 lies below the corresponding curve in Fig. 3. The performance loss due to normalization is measured by computing the extra SNR required by Detector 2 to obtain the same values of P_{fa} and P_d as Detector 1. A contour plot of the loss is shown in Fig. 5, in which the curves are labelled with the loss in dB. The compilation of this figure required two-dimensional interpolation of the data in Figs. 3 and 4, followed by contouring. The loss varies little over the ROC plot, and has a typical value of 0.156 dB for the range of detection and false alarm probabilities shown in the figure. This means that an increase of about 0.156 dB in SNR is sufficient to overcome the effect of normalization over this range.

The above case was for $M = 4$ and $K = 100$. Results were also produced for $M = 1$, 2, and 8, with K being adjusted to keep the product MK equal to 400. The curves for different values of M are plotted in Fig. 6. The plot has been restricted to a small range of P_d to allow a better view of the fairly small performance differences. As might be expected intuitively, the greatest loss is incurred by normalizing raw spectra ($M = 1$). The biggest improvement is obtained by averaging just two FFTs before normalization, with the incremental improvement decreasing as M is increased. Aside from the case $M = 1$ the performance differences are quite small, and only the case $M = 4$ will be considered below.

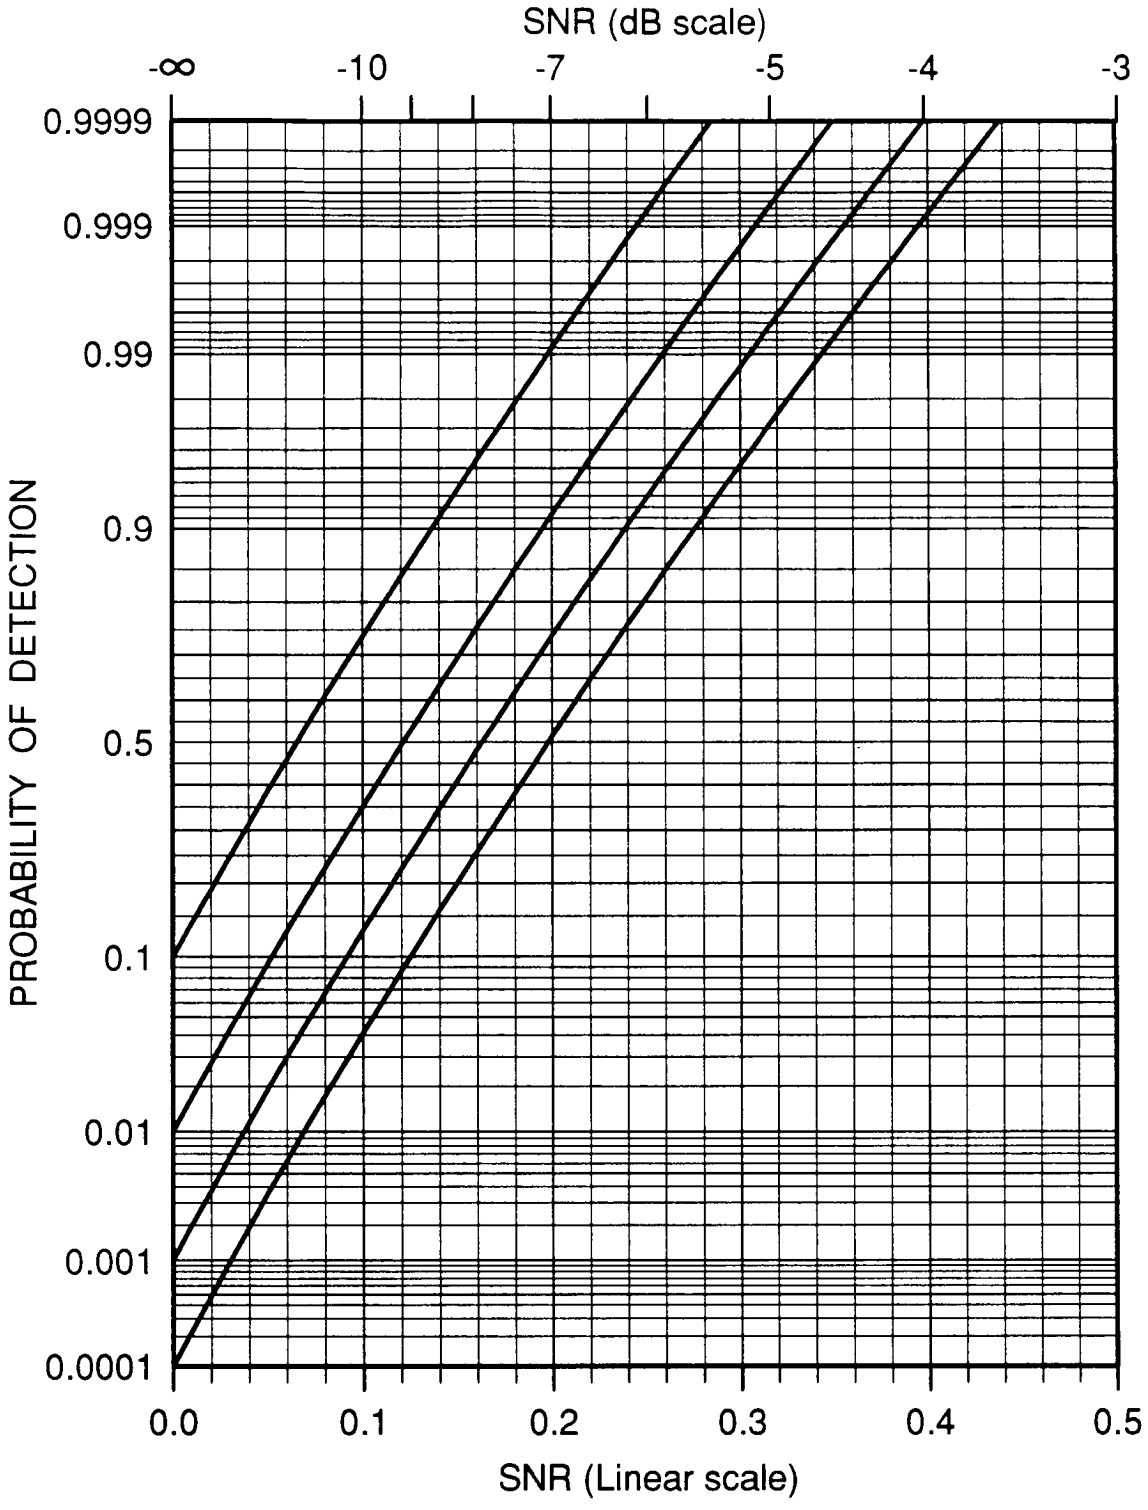

Figure 2. Probability of detection for Detector 1 with $N = 400$. The probability of false alarm for each curve is given by the value of the curve's intercept with the left axis.

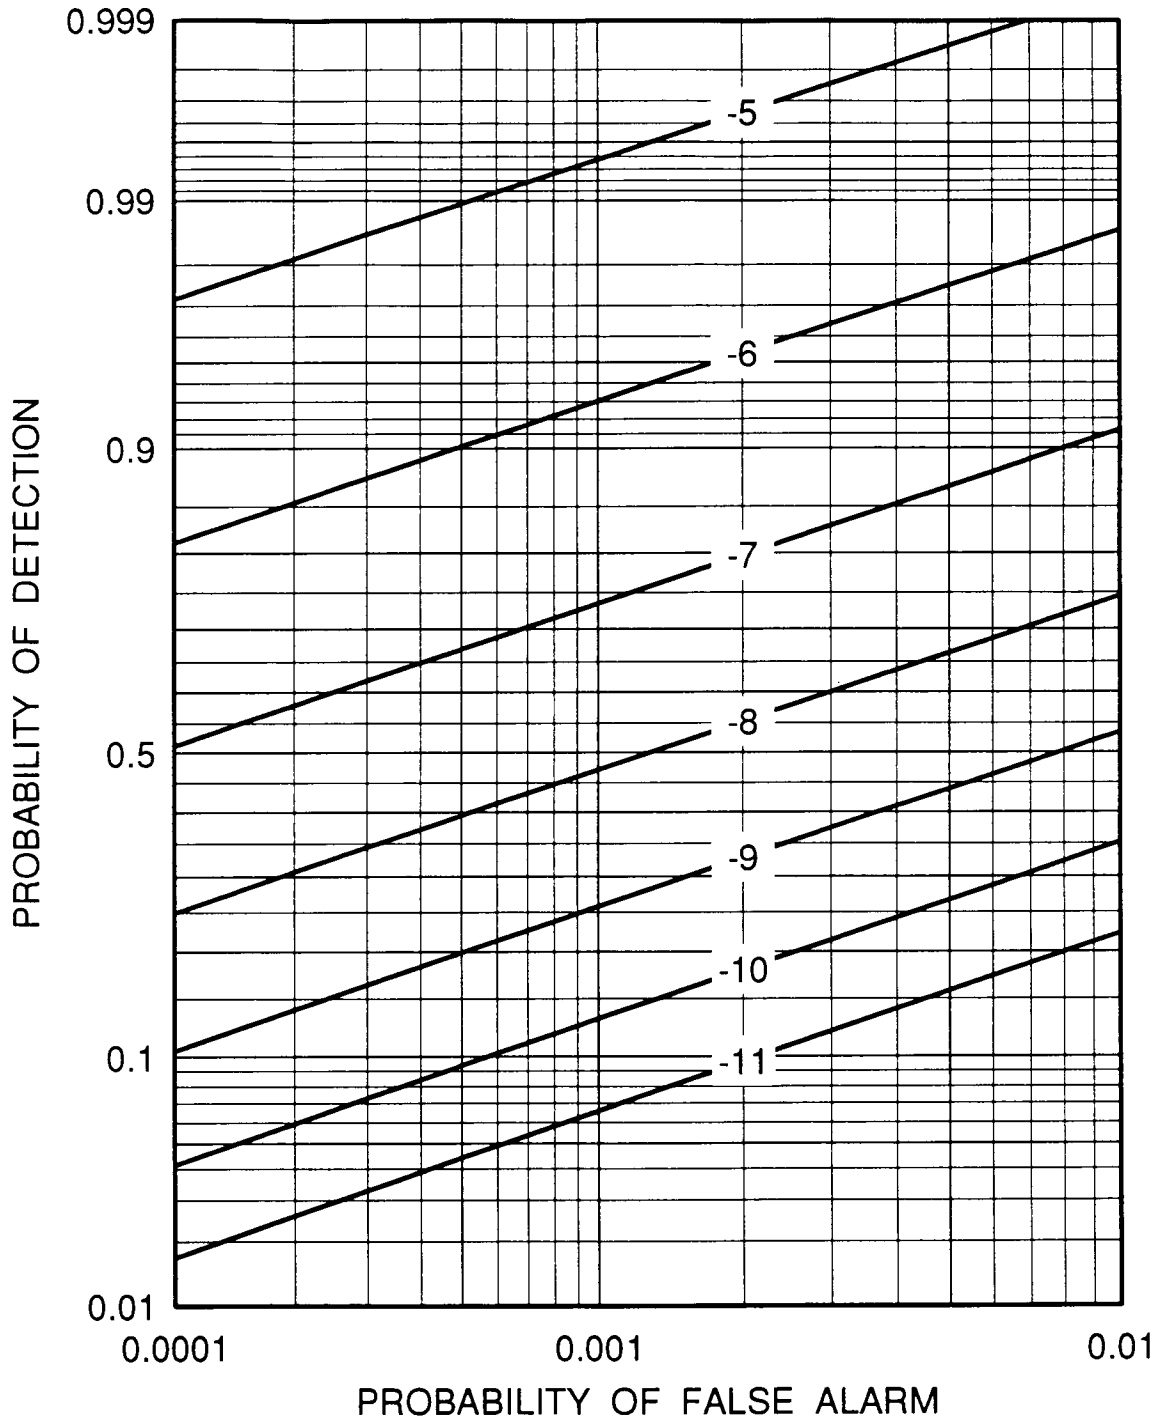

Figure 3. ROC curves for Detector 1 with $N = 400$. The curves are labelled with the SNR in dB.

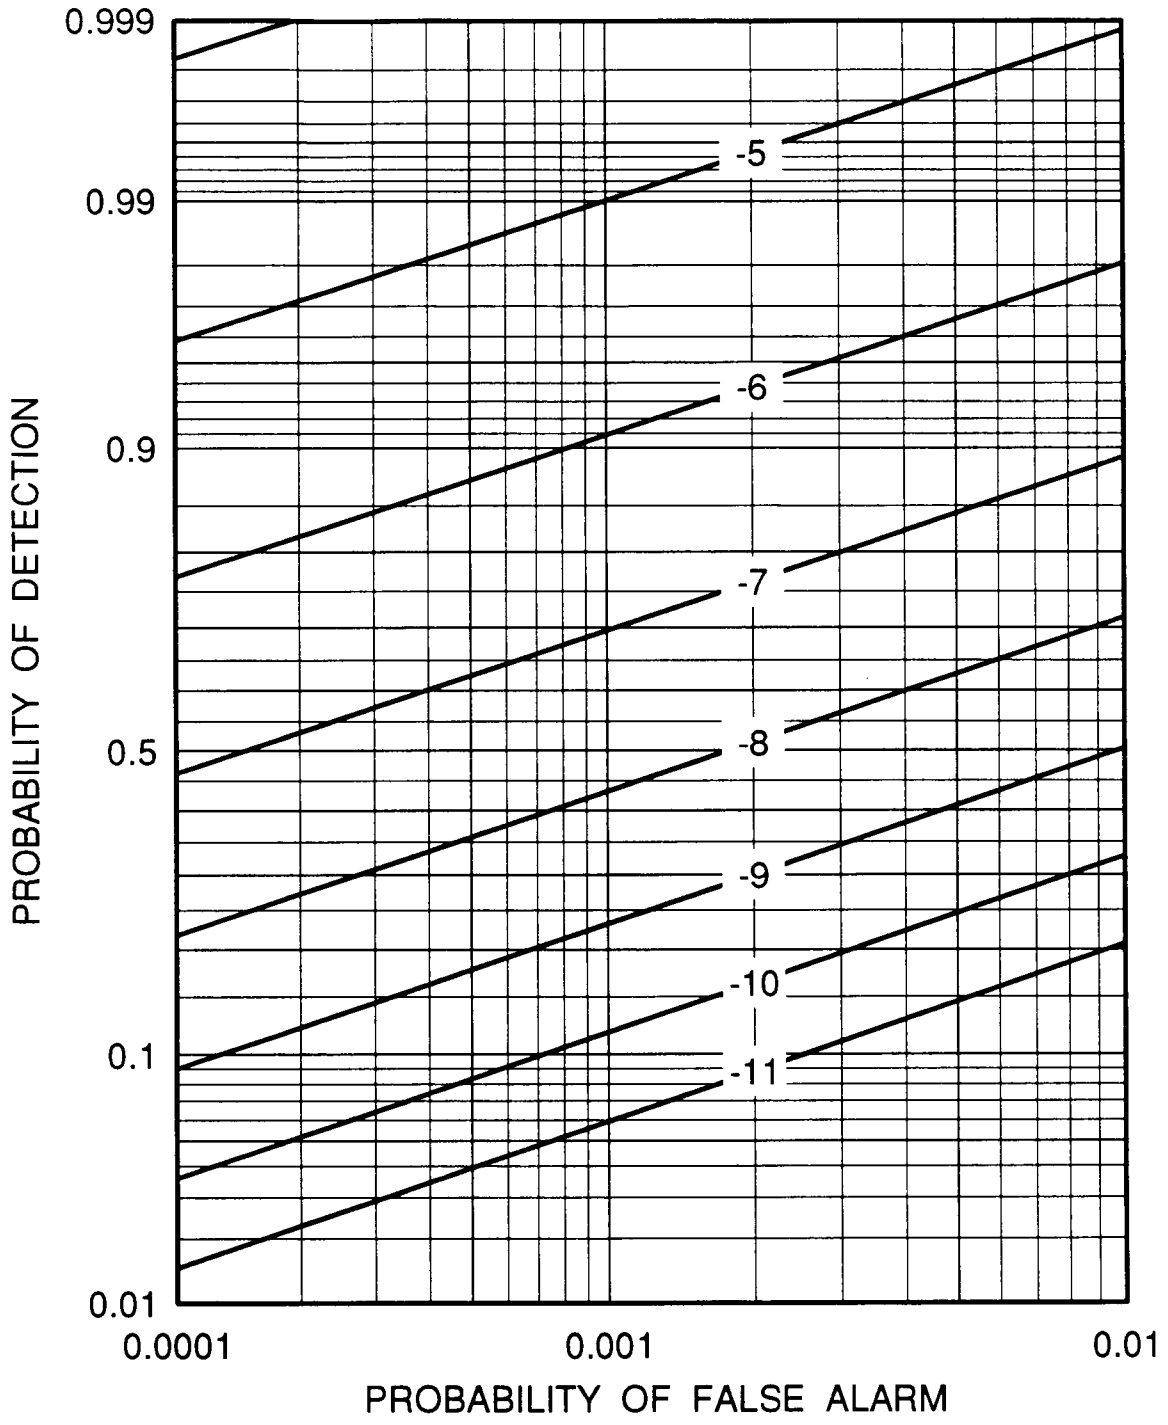

Figure 4. ROC curves for Detector 2 with $M = 4$, $K = 100$, and $L = 20$. The curves are labelled with the SNR in dB.

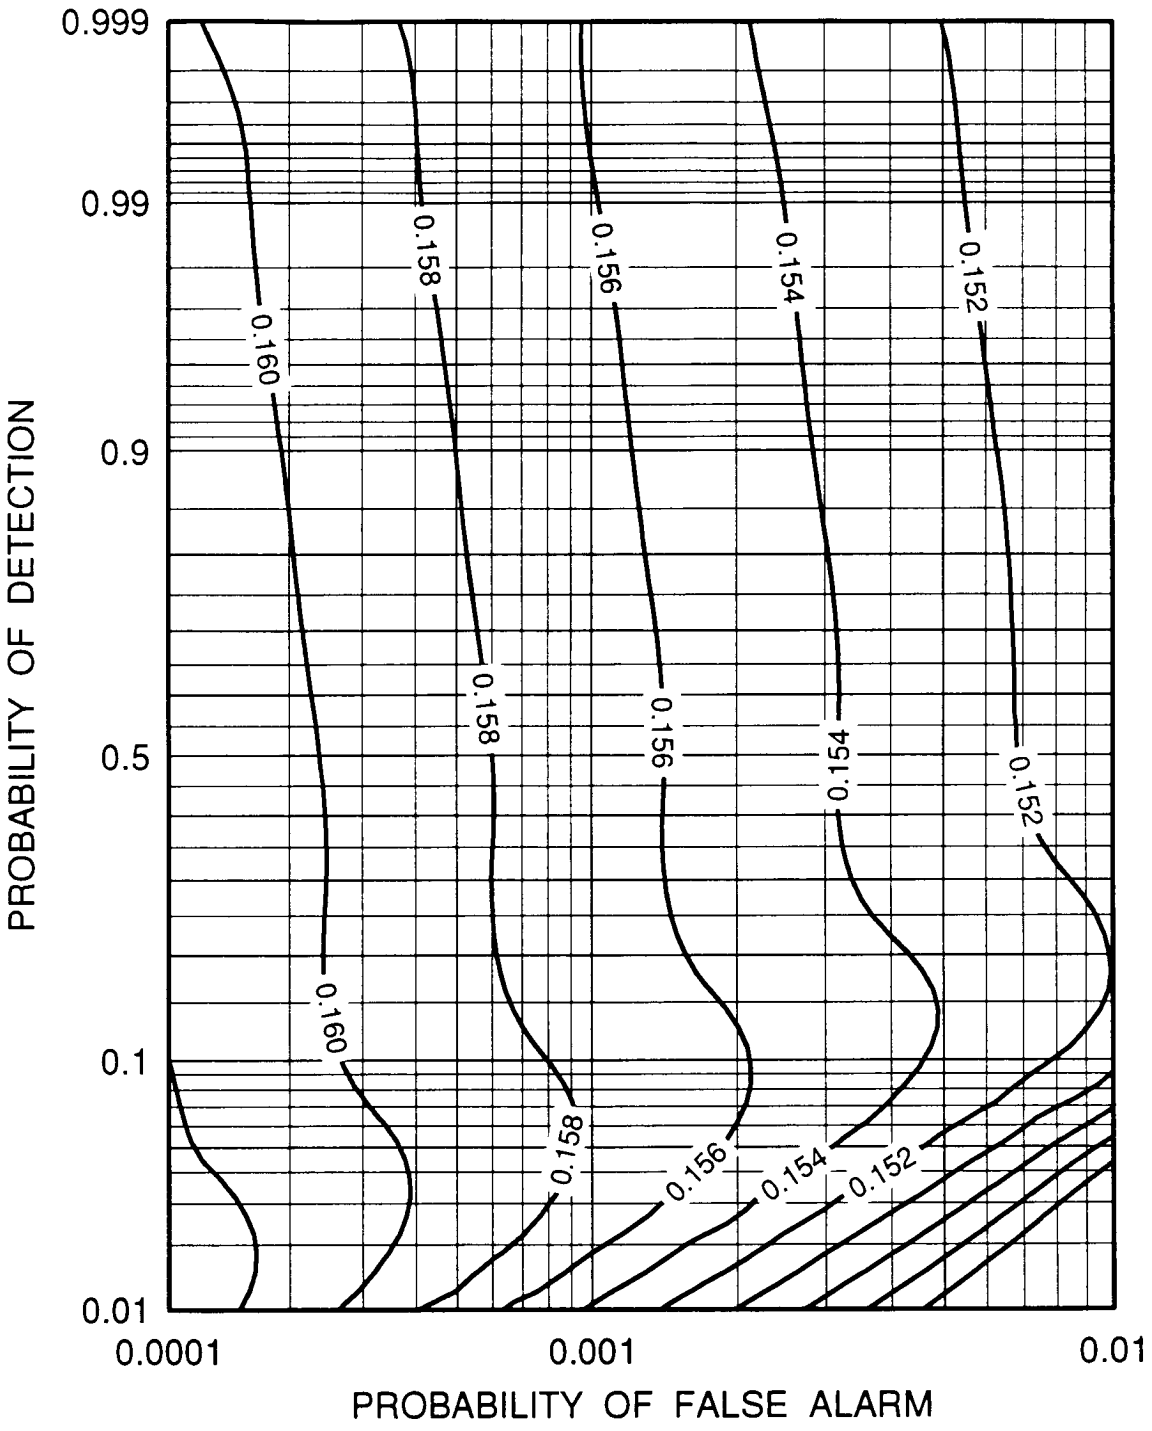

Figure 5. Loss due to normalization (dB).

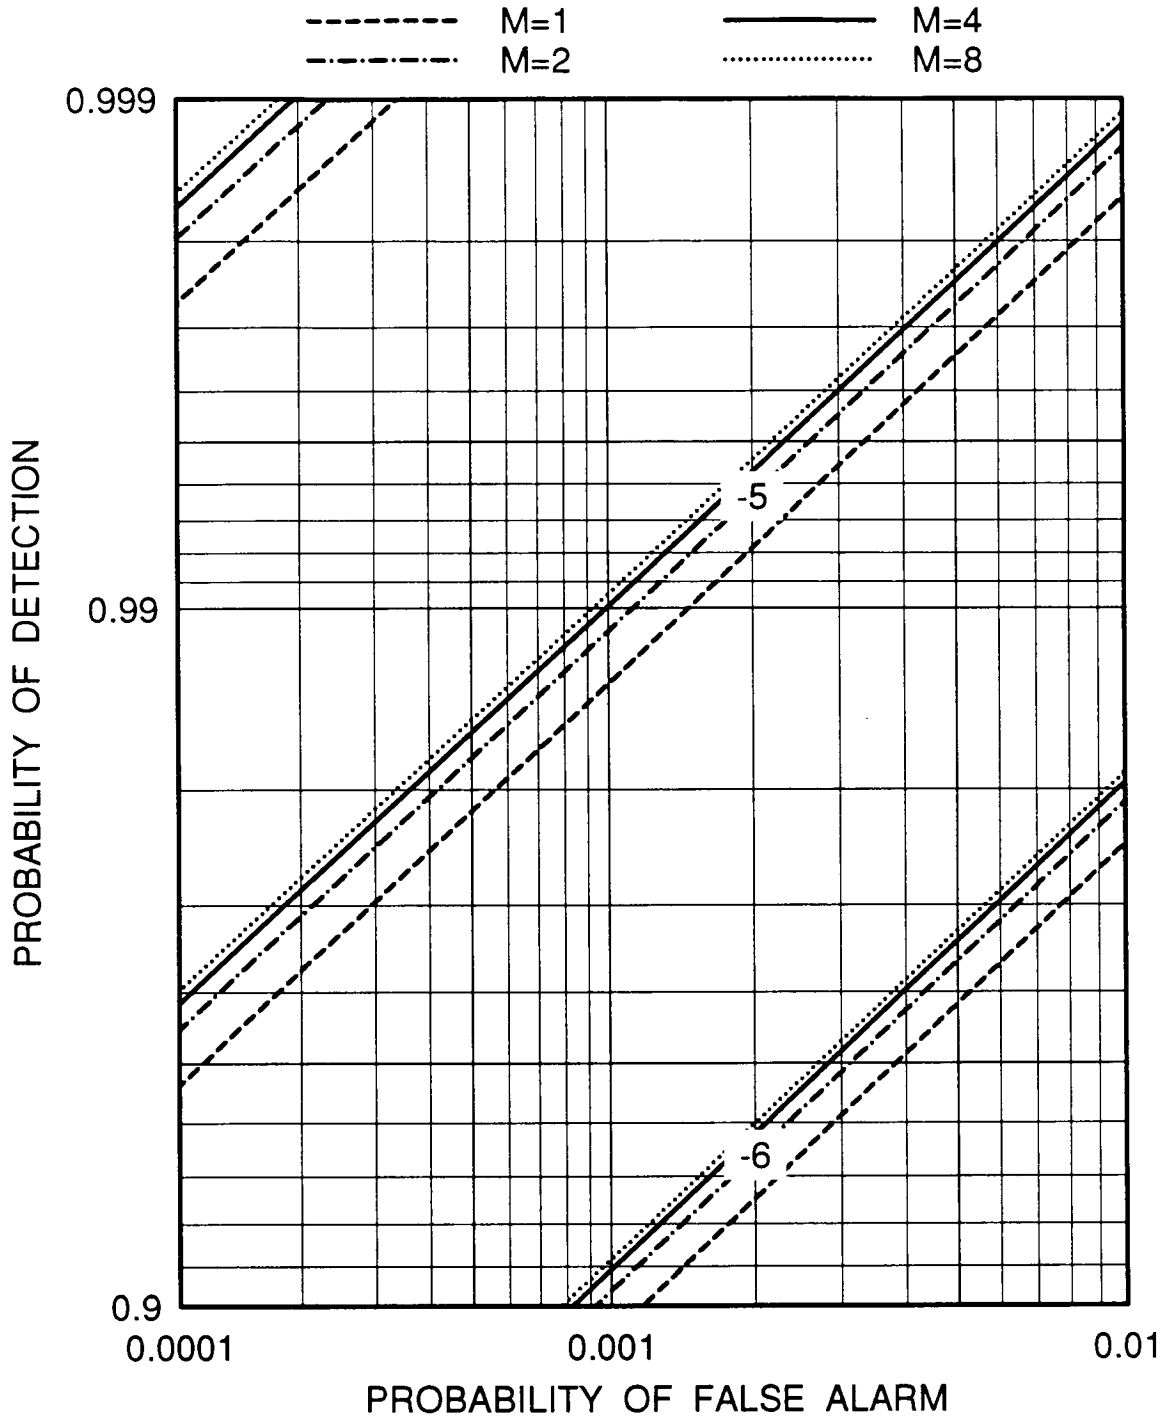

Figure 6. ROC curves for Detector 2 with different values of M, the parameter K being adjusted to maintain $MK = 400$. Each group of four curves is labelled with the SNR in dB.

4.3. Detector 3: Integration of normalized and quantized spectra

We now come to the most realistic detector model, in which the spectra are both normalized and quantized. The number of possible cases for analysis grows significantly, as both the number of levels and the breakpoints of the quantizer can be adjusted. The analysis will start with a one-bit quantizer, and then consider the effect of increasing the number of bits one at a time. The effect of changing the quantization interval and the position of the breakpoints within the interval will also be considered. However, there is no simple procedure for finding the optimal position for the breakpoints, and in some cases trial and error was used. Also, it should be noted that the ROC now consists of discrete points, since the decision variable at the detector output is discrete. That is, for unquantized data it is possible to sweep out a given ROC curve by continuously varying the threshold, whereas for quantized data the threshold is an integer. However, the discrete results will be plotted below as continuous curves for ease in interpretation.

The statistics of the quantizer input should be taken into account when setting the breakpoints. Here the input is a normalized short-term spectrum, and from the density function in Eq. (13) it can be shown that a noise-only bin, $\tilde{n}$, has a mean and standard deviation given by

$$\begin{aligned} \mu_{\tilde{n}} &= \frac{ML}{ML-1}, \\ \sigma_{\tilde{n}} &= \frac{1}{\sqrt{M}}\{1+O(L^{-1})\} \cong \frac{1}{\sqrt{M}}, \end{aligned} \tag{19}$$

where L is the number of bins used in the normalizer, assumed to be sufficiently large that terms of order $O(L^{-1})$ are small. The mean and standard deviation of the normalized bin are independent of the unknown noise variance σ^2, as expected. In all cases the mean is approximately unity, and the standard deviation decreases as $1/\sqrt{M}$. Thus as M increases there is less statistical fluctuation in the short-term spectra. The implication is that a fixed quantizer cannot be made to provide optimal performance for different values of M: the length of the quantization interval should be decreased as M is increased in order to maintain the same visual quality of an FTI display (e.g., in a sonar system in which an operator can alter M). However, as stated above, only the case $M = 4$ will be examined here.

4.3.1. ONE-BIT QUANTIZATION

The simplest quantizer encodes a normalized spectral bin into a single bit, 0 or 1. Since there is only one breakpoint, a good choice for it can be determined by trial and error without too much difficulty. It was found that a breakpoint at 1.3 gives the best performance in this example case. From Eq. (19) the mean of a normalized noise-only

bin is $ML/(ML-1) \cong 1.0127$, so the breakpoint is slightly above the mean of the distribution. The ROC curves for this breakpoint are shown in Fig. 7. A comparison of Figs. 4 and 7 shows that the loss due to quantization is very close to 1 dB. This remark is confirmed by Fig. 8, which shows contours of quantization loss. This is the loss due to quantization only, the comparison again being made with the results given in Fig. 4 for normalized data. The loss has remarkably little variation over the entire plot, and slightly exceeds 1 dB in the regions of interest. The fact that the loss is almost constant makes the notion of a "good" breakpoint a simple one, since there is a single choice of the breakpoint that achieves good detection performance simultaneously in all parts of the receiver operating characteristic.

Let p_{fa} (as opposed to P_{fa}) denote the probability of false alarm when a single short-term spectrum is thresholded; i.e., before the long-term integration is performed. Delisle and Kroenert[20] modelled the two-level FTI display by assuming that, in practice, the breakpoint is adjusted so that the value of p_{fa} is approximately 0.5 for each bin on the display. In other words, the breakpoint is set to the median of the noise-only pdf at the input to the quantizer. The noise background on the display then has the "salt and pepper" appearance that is characteristic of such displays. Some numerical computation shows that a breakpoint of about 0.9219 sets $p_{fa} = 0.5$ for a normalized short-term spectrum with $M = 4$. The quantization loss for the entire system, including the integration of K quantized spectra, is shown in Fig. 9. The loss is 0.3–0.4 dB greater than when the breakpoint is set to 1.3, and so the model assumed in Ref. 20 does not provide optimal performance among two-level systems, as least in the example considered here. The effect of setting the breakpoint to 1.3 is to decrease p_{fa} below 0.5, which would make an FTI display appear visually somewhat lighter than when $p_{fa} = 0.5$.

The sub-optimality of choosing the median as the breakpoint may be a consequence of the asymmetry of the pdf at the input to the quantizer. This asymmetry is indicated by the fact that the median of the noise-only pdf, 0.9219, is less than its mean value of 1.0127. It would be a topic for future study to see whether the median is the optimal breakpoint for the Gaussian pdf, a symmetric pdf with a common value for its mean and median. If this last statement were true, it would show that approximating the pdf at the input to the quantizer by a Gaussian pdf (as is done in some studies) yields misleading results.

4.3.2. TWO-BIT QUANTIZATION

In this case there are four levels and three breakpoints. The limited number of breakpoints made it feasible to optimize their position by a brute-force search. Again, the fact that the loss varies only slightly over most of the ROC was important, as the

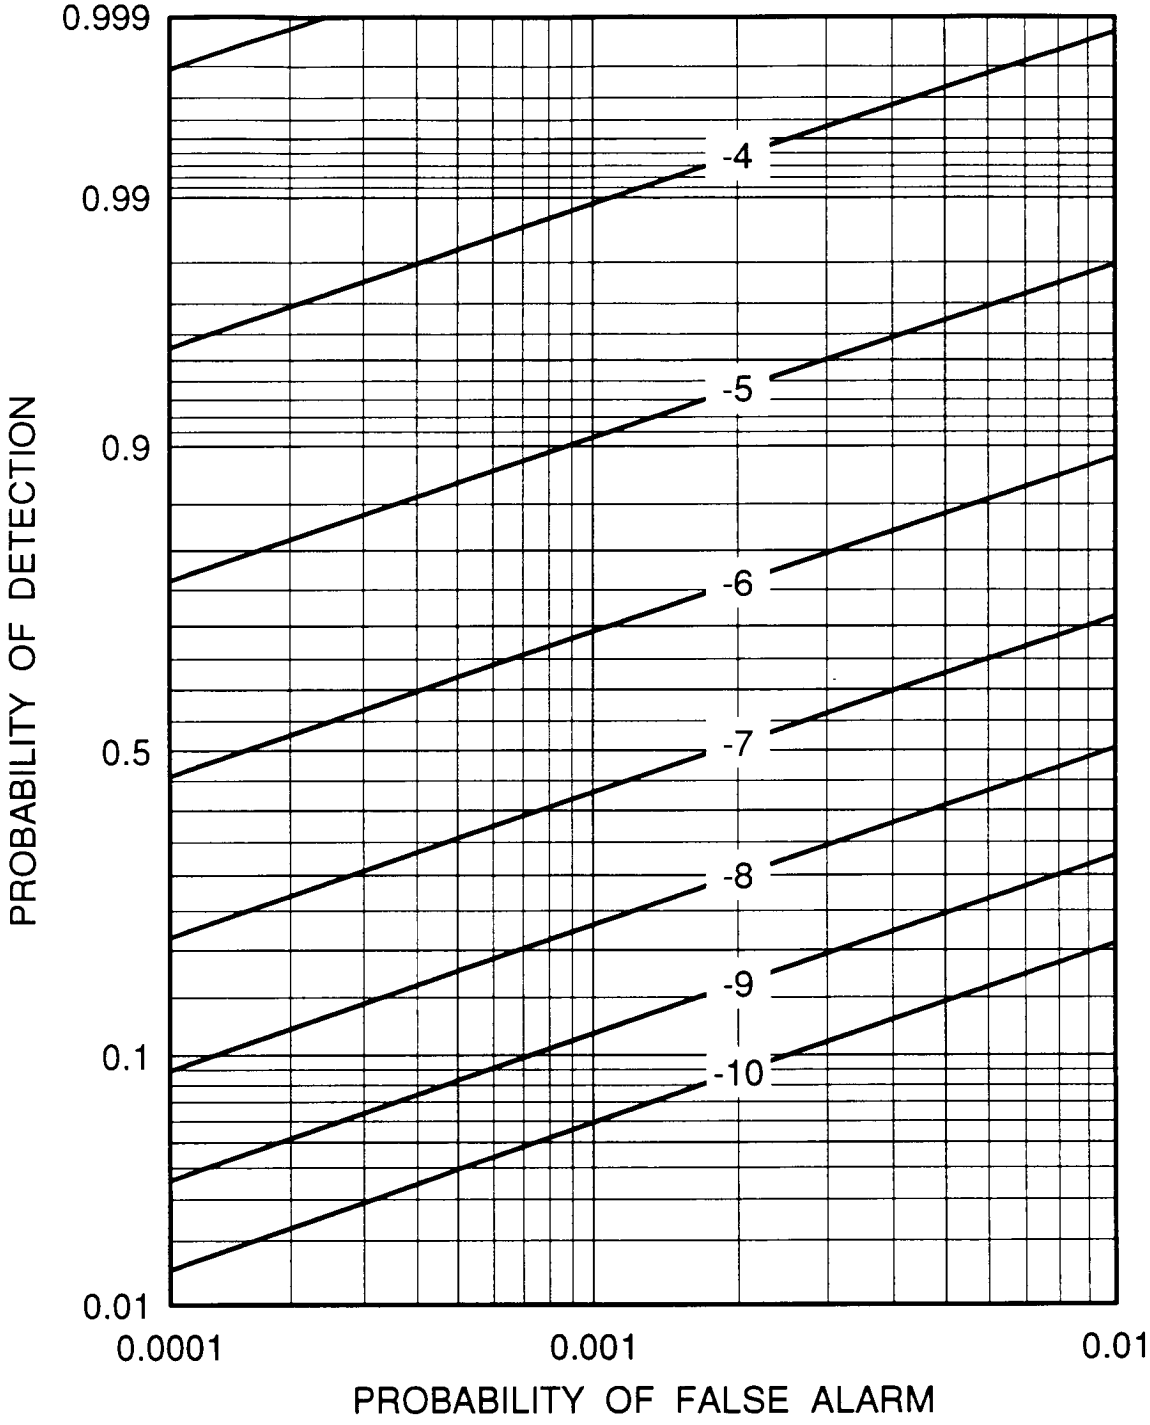

Figure 7. ROC curves for Detector 3 with one-bit quantization. The curves are labelled with the SNR in dB.

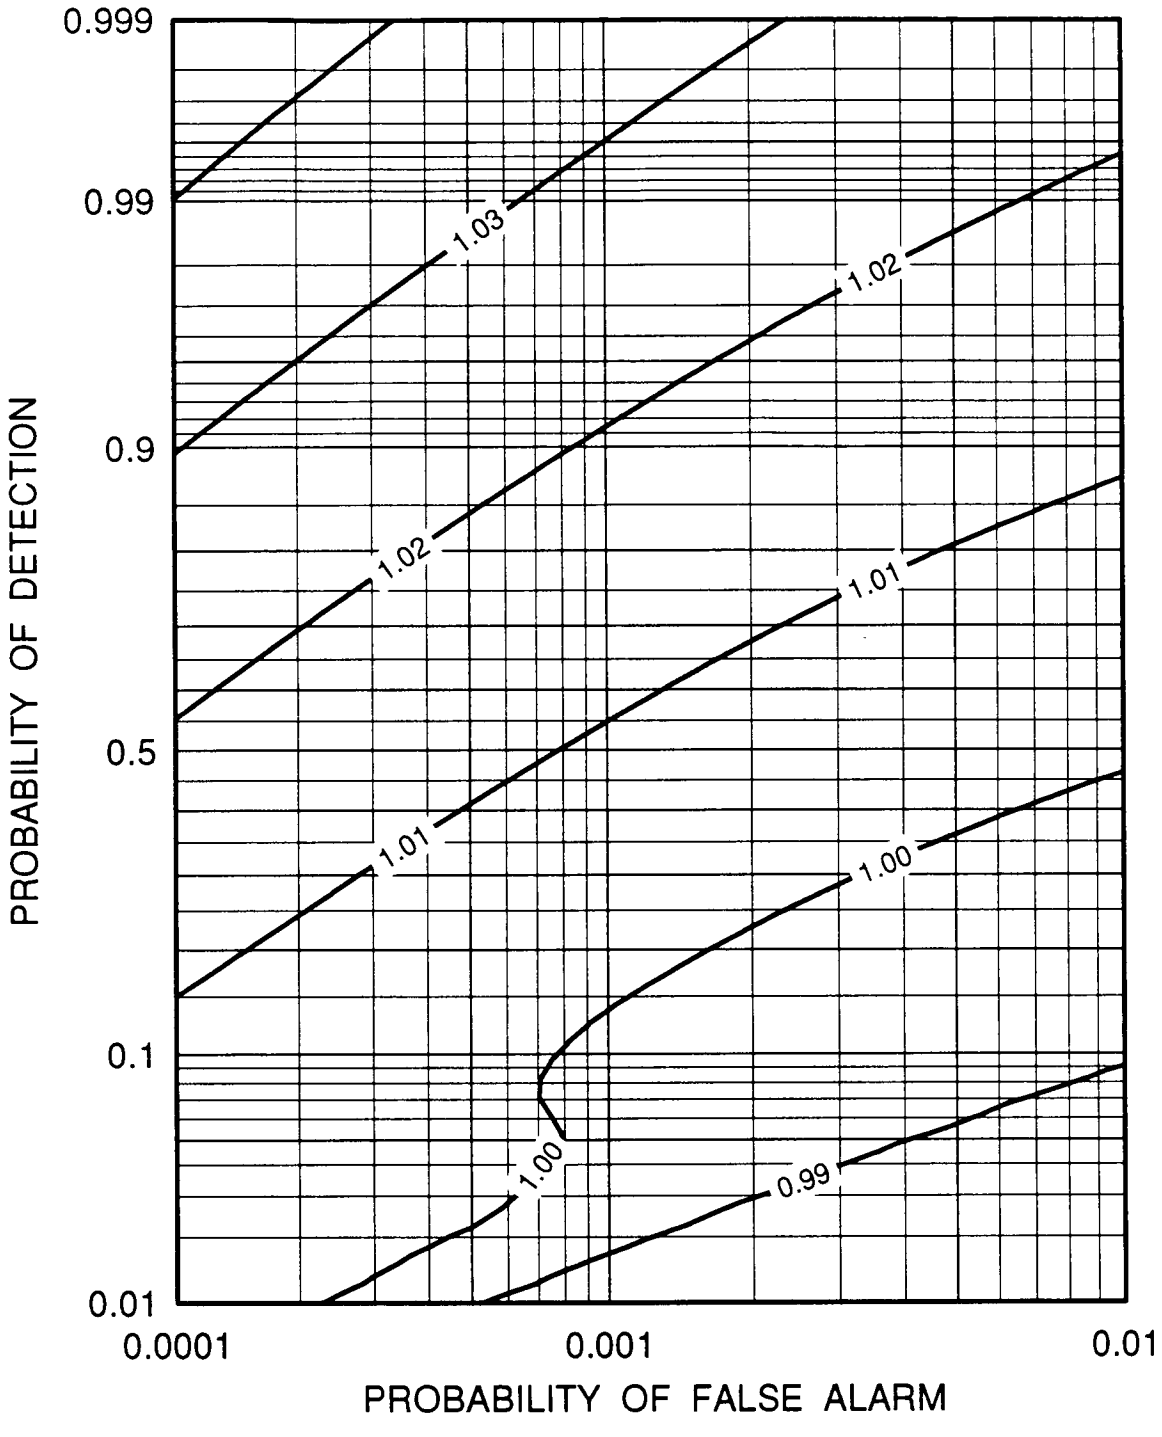

Figure 8. Loss resulting from one-bit quantization (dB).

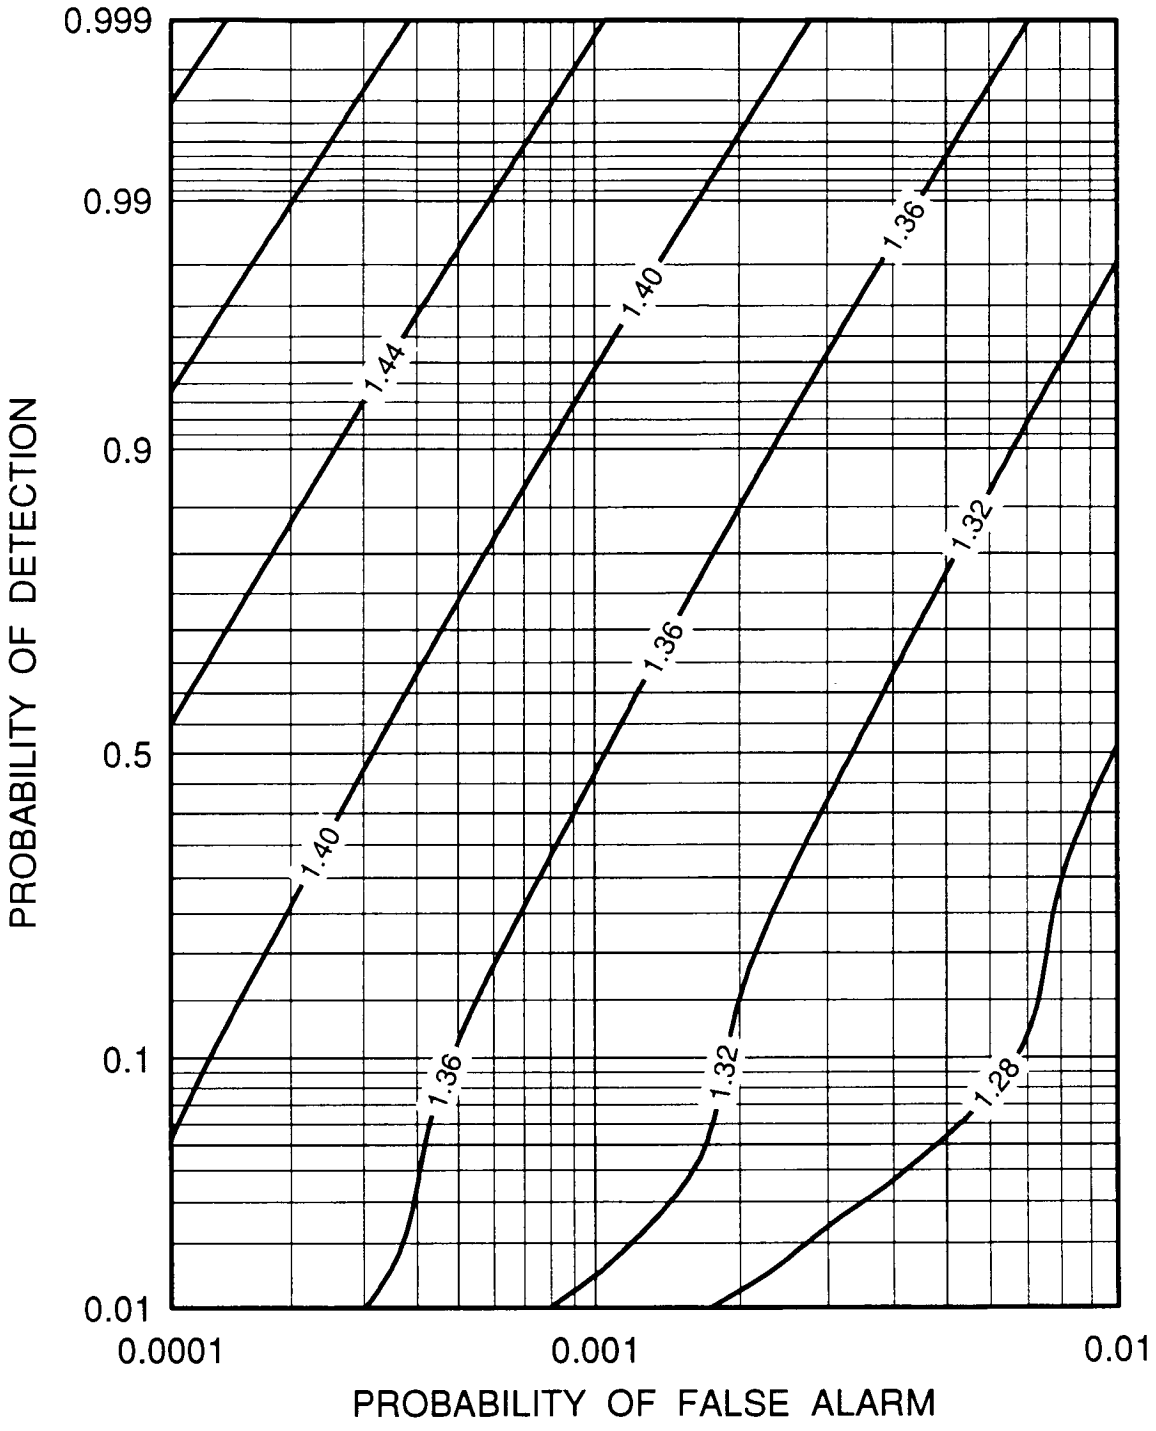

Figure 9. Loss resulting from one-bit quantization (dB) when the breakpoint is set to the median of the noise-only pdf.

detection performance could be optimized at a single point on the characteristic. The breakpoints that optimize the performance at the single point then provide close-to-optimal performance over the rest of the ROC as well.

The breakpoints were varied in the range from 0.5 to 3.5 in steps of 0.1, with the usual constraint that $x_1 < x_2 < x_3$, and the best choice was found to be $\{0.8, 1.4, 2.1\}$. The additional improvement obtainable by refining the breakpoints, say by using a grid of step size 0.01, is negligible. Contours of quantization loss (again relative to Fig. 4) are shown in Fig. 10. The loss is decreased dramatically by going from two to four output levels in the quantizer, but at slightly less than 0.3 dB it is still greater than the normalization loss.

4.3.3. THREE-BIT QUANTIZATION

The number of breakpoints is now 7, making a brute-force search computationally infeasible. To reduce the number of parameters, the breakpoints were assigned via a μ–law characteristic described in Appendix B. When $\mu = 0$ the breakpoints are assigned linearly within the quantization interval, whereas for $\mu > 0$ the lowest breakpoints cluster together more. There are thus three parameters to vary: the extreme breakpoints x_1 and x_{B-1} which define the quantization interval and μ. There is no guarantee that the breakpoints found are optimal, since they are constrained to obey the μ–law. The final choice for the breakpoints was $\{0.55, 0.84, 1.15, 1.49, 1.86, 2.26, 2.70\}$, corresponding to a value $\mu = 0.8$. Thus the points are skewed slightly downward from a linear spacing. The resulting loss, shown in Fig. 11, is now under 0.1 dB and is approximately half that due to normalization.

4.3.4. FOUR-BIT QUANTIZATION

It was found that, with 4-bit quantization, it is possible to reduce the loss to a small value simply by choosing breakpoints with a linear spacing. For example, linearly spacing the 15 breakpoints from 0.5 and 3.0 yields a loss of less than 0.03 dB over most of the ROC plot. The plot is omitted.

5. CONCLUSIONS

A numerical study has been conducted to determine the effect of spectral normalization and quantization on the performance of a system that uses long-term spectral integration to detect narrowband signals in noise. The study was based on a complete mathematical representation of the processing that is used in practical sonars, with a realistic signal and noise model. Some conclusions are:

1. The loss due to normalization is small, less than 0.2 dB for the cases examined. The loss can be reduced slightly by integrating at least two FFT spectra before the normalizer.

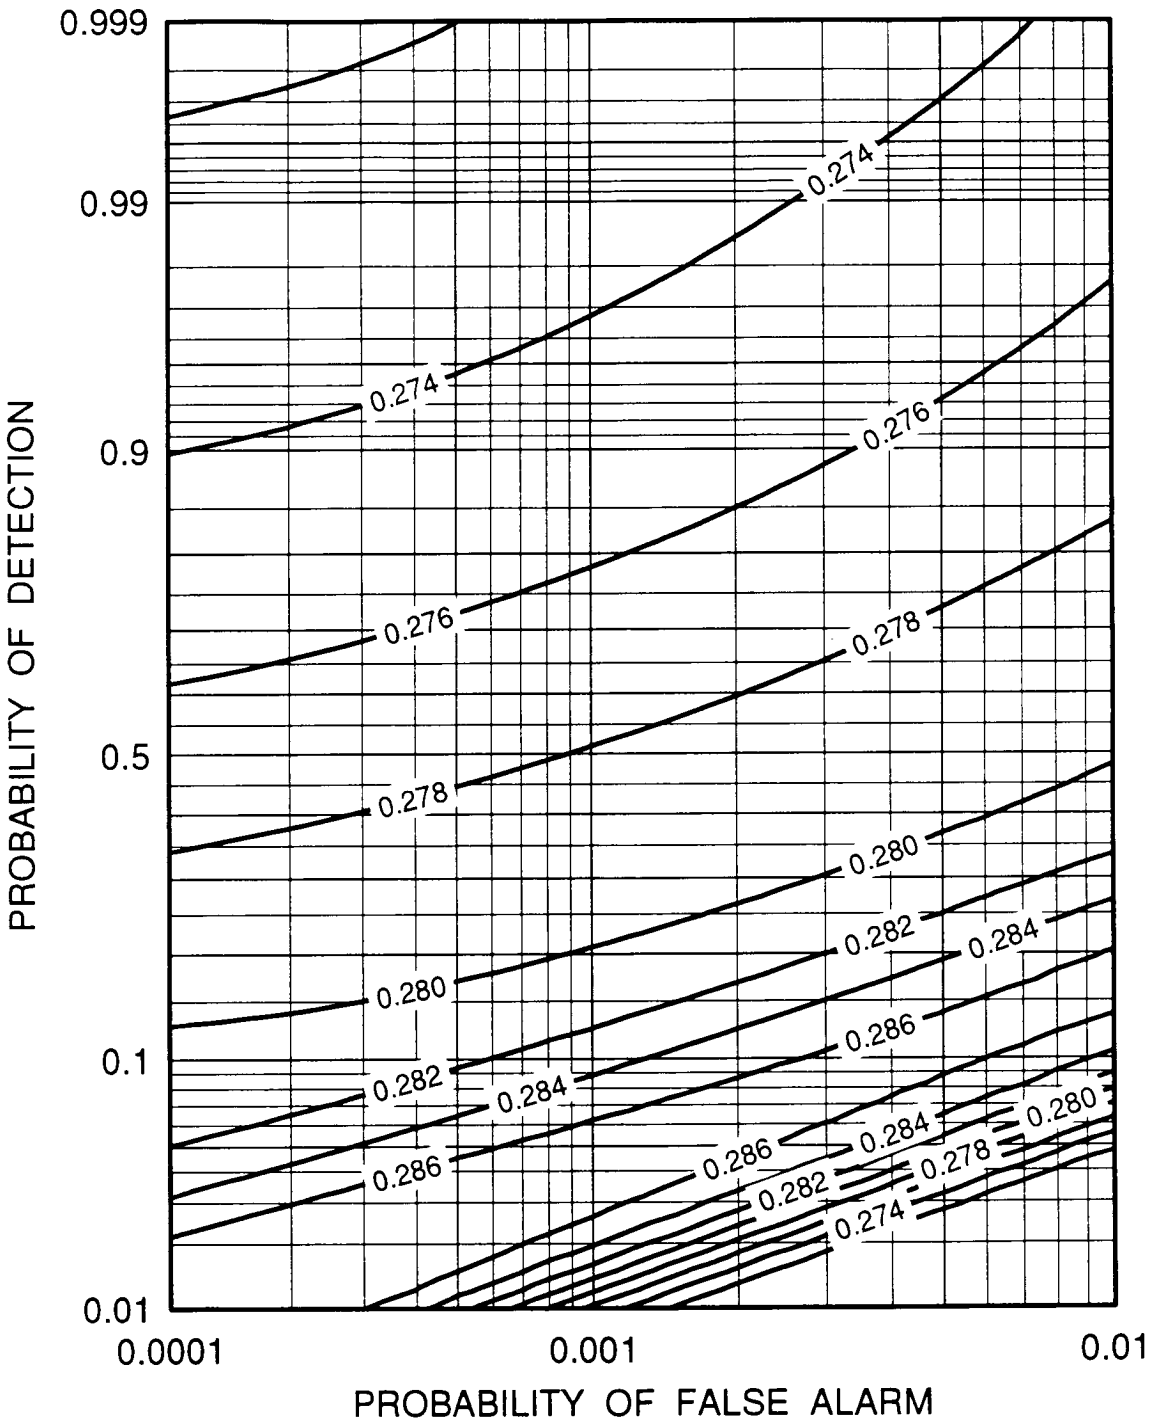

Figure 10. Loss resulting from two-bit quantization (dB).

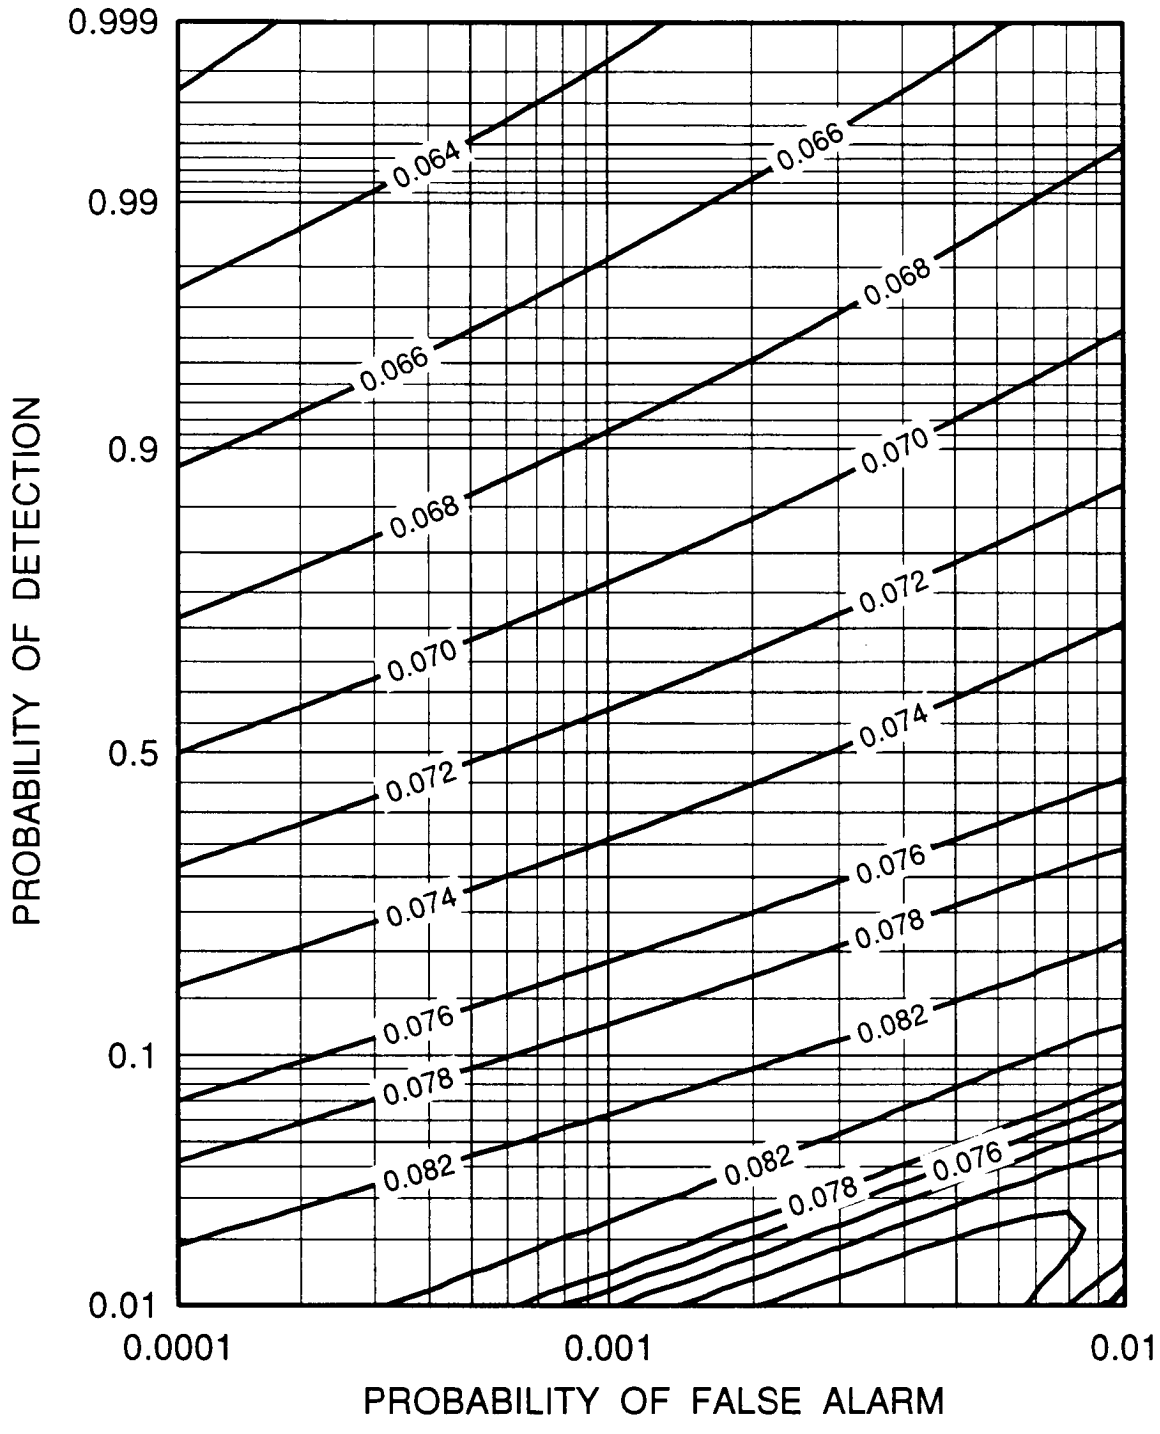

Figure 11. Loss resulting from three-bit quantization (dB).

2. The loss due to one-bit quantization of the normalized spectra exceeds 1 dB. The use of a three-bit quantizer, typical for operational sonar systems, provides detection performance that is near optimal, and is suitable for long-term integration.

3. The position of the quantizer breakpoints is important in obtaining the full detection performance possible for a given number of levels. Some previous studies on quantization (e.g., Ref. 8) set the lower breakpoint to the mean of the noise-only pdf, but the results of this memorandum show that the lower breakpoint should be chosen below the mean. The reason for this is not hard to explain. We wish to detect a change in the pdf of a decision variable, and must do so by making inferences about the pdf on the basis of the stochastic samples being observed. The stochastic samples we observe on the lower tail of the pdf provide information about the shape of the pdf, as do the samples on the upper tail. Thus, if the samples on the lower tail are not quantized sufficiently, this information is lost. It should be noted, however, that Ref. 8 simply assumes a Gaussian pdf at the input to the quantizer, instead of the more realistic F-distribution considered here. In view of the above explanation of why quantization of the lower tail is also important, it is likely that the best lower breakpoint should be below the mean even in the Gaussian case.

4. The performance ranking of the quantizers is consistent with that obtained by conducting experiments in which the long-term integration is done by human subjects viewing an FTI display. This does not mean that the human operators match the theoretical performance, only that the ranking is the same: recent experiments conducted with human subjects viewing FTI displays with one-bit quantization indicate a loss of 3–5 dB from the theoretical.[21] (The numbers just quoted must be taken with a grain of salt, since most human-performance studies use an artificial SNR that is hard to identify with any stage in a sonar processor. The SNR values given in this paper come from a realistic signal and noise model combined with a complete model of the sonar signal processor — hence, they are much more meaningful.) Nevertheless, the more accurate theoretical results derived here demonstrate that the rapid roll-off in human performance with the number of quantizer levels is not merely a property of the human visual system.

APPENDIX A

Several integrals that occur in the body of the paper are of the form

$$\frac{1}{\Gamma(\beta)} \int_0^\infty y^{\beta-1} e^{-y} f(y)\, dy, \tag{20}$$

where f is a known function. It is assumed that the integration cannot be performed analytically and therefore must be done numerically. The integral in Eq. (20) has the

form of a Laguerre integral, and hence can be evaluated efficiently by Gauss-Laguerre quadrature. That is, if $f(y)$ is sufficiently smooth we may write

$$\frac{1}{\Gamma(\beta)} \int_0^\infty y^{\beta-1} e^{-y} f(y)\, dy \cong \sum_{j=1}^{J} w_j f(y_j), \tag{21}$$

where the abscissas y_j are the zeros of the Laguerre polynomial $L_J^{(\beta)}(\cdot)$ of degree J, and the w_j are appropriately selected weights. Here J is chosen sufficiently large to ensure good accuracy in the quadrature rule; by increasing J, the accuracy is generally improved (although round-off error can cause deteriorating accuracy with increasing J in a computer implementation). This does not immediately solve all the numerical problems, however, since the value of β in Eq. (21), equal to ML for the specific integrals (17) and (18), has a maximum value of 160 for the cases examined in this paper. The computation of the Gauss-Laguerre abscissas and weights requires that the computer be able to represent numbers on the order of $\Gamma(\beta)$ without overflow, limiting the double-precision NAG library[9] running on a VAX computer to $\beta = 34$ [since $\Gamma(34) \cong 8.7 \times 10^{36}$, just under the largest value representable in VAX double precision]. This problem was overcome by implementing a FORTRAN code for the computation of Gauss-Laguerre abscissas and weights[22] in quadruple-precision arithmetic; the added dynamic range allowed the successful computation of the desired quadrature rules. Furthermore, it was found that by scaling the weights they could be stored in the limited dynamic range of VAX double-precision arithmetic, and all applications of the quadrature rule could then be performed using double precision. That is, the quadrature formula was re-written in the form

$$\sum_{j=1}^{J} w_j f(y_j) = \kappa^{-1} \sum_{j=1}^{J} (\kappa w_j) f(y_j),$$

where κ is a scaling factor chosen such that the scaled weights κw_j fit within the approximate range $10^{\pm 38}$ of VAX double-precision arithmetic. Here κ was chosen to be a power of 2, in order to lose no accuracy in the scaling process itself. Note that the quadrature rule need be computed only once for given values of β and J.

This method turned out to be an efficient way to evaluate the characteristic function $\phi(\xi; \tilde{x})$ in Eq. (17): the weighting function $y^{\beta-1}e^{-y}/\Gamma(\beta)$ is "built into" the quadrature rule, and so only the function f must be evaluated at the quadrature abscissas. Moreover, it was found that very low-order quadrature rules sufficed to give 12 or more significant digits in the compution of $\phi(\xi; \tilde{x})$. For example, the results of the 24-point and 32-point Gauss-Laguerre rules usually agreed to 12 digits, and it was only rarely that the 48-point rule had to be employed.

APPENDIX B

Although it is possible to try to optimize the position of the quantizer breakpoints, there is no easy method for doing so. Thus we consider a quantizer with a simple analytic form.[23] This quantizer has been used in the implementation of the passive sonar processor described in Ref. 19. First, the input variable x is mapped continuously onto an output interval $[b_1, b_2]$ via a μ–law function:

$$y = \begin{cases} b_1 & \text{for } x < a_1, \\ b_1 + \dfrac{b_2 - b_1}{\ln(1+\mu)} \ln\left(1 + \mu \dfrac{x - a_1}{a_2 - a_1}\right) & \text{for } a_1 \le x \le a_2 \text{ and } \mu > 0, \\ b_1 + (x - a_1)\dfrac{b_2 - b_1}{a_2 - a_1} & \text{for } a_1 \le x \le a_2 \text{ and } \mu = 0, \\ b_2 & \text{for } x > a_2. \end{cases} \tag{22}$$

The output y of the μ–law is quantized by taking $\lfloor y \rfloor$; i.e., the largest integer less than or equal to y. By varying the value of μ, the quantizer can take on different forms. When $\mu = 0$, the quantizer break points are linearly spaced, and as μ increases they cluster near the low end of the quantization interval.

We now consider the choice of the four parameters a_1, a_2, b_1, and b_2 in Eq. (22). For a quantizer with b bits, or $B = 2^b$ levels, the output range will be chosen to be $[b_1, b_2] = [0, B-1]$. For example, a 3-bit quantizer will use $b_1 = 0$ and $b_2 = 7$. The output of the quantizer is thus an integer in the set $\{0, 1, \ldots, B-1\}$, as assumed in Sec. 3.3. To find the breakpoints x_j of the quantizer, the μ–law in Eq. (22) is first inverted. Assuming that $\mu > 0$, the result is

$$\begin{aligned} x &= a_1 + \frac{(a_2 - a_1)}{\mu}\left[(1+\mu)^{(y-b_1)/(b_2-b_1)} - 1\right] \\ &= a_1 + \frac{(a_2 - a_1)}{\mu}\left[(1+\mu)^{y/(B-1)} - 1\right] \end{aligned} \tag{23}$$

Replacing y in Eq. (23) by the output level $y_j = j$ yields the breakpoint x_j for $j = 1, \ldots, B-1$. From Eq. (23) it is seen that $a_2 = x_{B-1}$, i.e., a_2 coincides with the uppermost breakpoint. On the other hand, a_1 does not coincide with the first breakpoint x_1. Thus, for a given quantization interval $[x_1, x_{B-1}]$, a_1 was adjusted to yield a first breakpoint of x_1 and a_2 was set to x_2.

REFERENCES

1. R.O. Nielson, *Sonar Signal Processing.* Boston: Artech House, 1991.
2. R.S. Walker, "The detection performance of FFT processors for narrowband signals," DREA Technical Memorandum 82/A, Defence Research Establishment Atlantic, Dartmouth, NS, February 1982.

3. F. Kingdom and B. Moulden, "Factors affecting the detection of signals in simulated FTI displays," Displays, pp. 127–128, July 1987.

4. S.M. McFadden and K. Swansen, "Mapping of amplitude onto CRT luminance in frequency-time-intensity displays," in Proceedings of the Human Factors Society, 30th Annual Meeting, held 29 Sept–3 Oct 1986, Dayton, Ohio, pp. 164–167.

5. B.H. Maranda and J.A. Fawcett, "Detection and localization of weak targets by space-time integration," IEEE J. Oceanic Eng., vol. 16, no. 2, pp. 189–194, April 1991.

6. V.G. Hansen, "Optimization and performance of multilevel quantization in automatic detectors," IEEE Trans. Aerospace and Electronic Systems, vol. AES-10, no. 2, pp. 274–280, March 1974.

7. A.H. Nuttall, "Detection performance characteristics for a system with quantizers, OR-ing, and accumulator," J. Acoust. Soc. Am., vol. 75, pp. 1631–1642, May 1983.

8. W.A. Struzinski, "Optimizing the performance of a quantizer," J. Acoust. Soc. Am., vol. 75, pp. 1643–1647, May 1983.

9. NAG Fortran Library, Mark 14, Numerical Algorithms Group, 1400 Opus Place, Downers Grove, IL, 1990.

10. A.H. Nuttall, "Alternate forms for numerical evaluation of cumulative probability distributions directly from characteristic functions," Proceedings of the IEEE, vol. 58, no. 11, pp. 1872–1873, November 1970.

11. A.D. Whalen, *Detection of Signals in Noise.* New York: Academic Press, 1971. (Ch. 4)

12. M. Abramowitz and I.A. Stegun (eds.), *Handbook of Mathematical Functions.* New York: Dover Publications, 1965.

13. R.J. Urick, *Principles of Underwater Sound*, 3rd edition. New York: McGraw-Hill, 1983.

14. W.A. Struzinski and E.D. Lowe, "A performance comparison of four noise background normalization schemes proposed for signal detection systems," J. Acoust. Soc. Am., vol. 76, pp. 1738–1742, December 1984.

15. M. Kendall and A. Stuart, *The Advanced Theory of Statistics*, 4th edition. London: Charles Griffin, 1977. (Ch. 11)

16. L.L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis.* Reading, Massachusetts: Addison-Wesley, 1991.

17. S. Kotz and R. Srinivasan, "Distribution of product and quotient of Bessel function variates," Annals of the Inst. of Stat. Math., vol. 21, pp. 201–210, 1969.

18. A.F. Ifram, "On the characteristic functions of the F and t distributions," Sankhyā, Series A, vol. 32, pp. 350–352, 1970.

19. D.W. Craig, "Signal processing, storage and display for towed array data," in Proc. Oceans '91, held 1–3 October 1991, Honolulu, HI, pp. 1223–1227.

20. R.B. Delisle and J.T. Kroenert, "An analytical model for the detection performance of multiple channel time history display formats," J. Acoust. Soc. Am., vol. 73, pp. 2065–2070, June 1983.

21. J.C. DiVita and T.E. Hanna, "Human efficiency for visual detection of targets on cathode ray tube displays using a two-level multiple-channel time history format." J. Acoust. Soc. Am., vol. 91, pp. 1552–1564, March 1992.

22. A.H. Stroud and D.H. Secrest, *Gaussian Quadrature Formulas.* Englewood Cliffs, NJ: Prentice-Hall, 1966.

23. K.J. Hahn, "Quantization of merged intensity displays," in Proc. of IEEE Pac. Rim Conf. on Communications, Computers, and Signal Proc., 4–5 June 1987, pp. 461–464.

REPORT NO.:	DREP Technical Memorandum 93-11
TITLE:	The Detection Performance of a Long-Term Spectral Integrator Using Normalized and Quantized Data
AUTHOR(S)	Brian H. Maranda
DATE:	February 1993

1 - DSIS

1 - DRDM

1 - DMCS-3

1 - DMOR

1 - DCIEM

2 - DREA

8 - DREP

United States

1 - NUWC Attn: A.H. Nuttall

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence Research Establishment Pacific
FMO Victoria, B.C.
Canada V0S 1B0

2. SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

Unclassified

3. TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title.)

The detection performance of a long-term spectral integrator using normalized and quantized data

4. AUTHORS (Last name, first name, middle initial)

MARANDA, Brian H.

5. DATE OF PUBLICATION (month and year of publication of document)

February, 1993

6a. NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)

30

6b. NO. OF REFS (total cited in document)

23

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include the address.)

9a. PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)

DRDM-08

9b. CONTRACT NO. (if appropriate, the applicable number under which the document was written)

10a. ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)

DREP TM 93-11

10b. OTHER DOCUMENT NOS. (Any other numbers which may be assigned this document either by the originator or by the sponsor)

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Distribution limited to defence departments and defence contractors; further distribution only as approved
( ) Distribution limited to defence departments and Canadian defence contractors; further distribution only as approved
( ) Distribution limited to government departments and agencies; further distribution only as approved
( ) Distribution limited to defence departments; further distribution only as approved
( ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)

Unlimited

13 ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both offical languages unless the text is bilingual).

The effect of normalization and quantization on the detection performance of a long-term spectral integrator is examined. The performance is described by the receiver operating characteristic, which is computed numerically from the probability density function of the detection statistic. It is found that the loss in performance due to normalization is small, less than 0.2 dB for an example case. The loss due to quantization depends on the number of bits used. When the quantizer uses only one bit (two levels), the loss is approximately 1 dB. This loss can be decreased to less than 0.1 dB by the use of three bits (8 levels) in the quantizer.

(See report for French version of the abstract.)

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Detection
Fast Fourier transform
FFT
Lofargram
Narrowband
Noncoherent integration
Normalization
Quantization
Spectrogram