Journal of Magnetic Resonance **144,** 305–315 (2000)
doi:10.1006/jmre.2000.2074, available online at http://www.idealibrary.com on IDEAL®

Detection of Nuclear Resonance Signals: Modification of the Receiver Operating Characteristics Using Feedback

A. J. Blauch,* J. L. Schiano,* and M. D. Ginsberg†

**Department of Electrical Engineering, Pennsylvania State University, 227D Electrical Engineering West, University Park, Pennsylvania 16802; and †United States Army Construction Engineering Research Laboratories, P.O. Box 9005, Champaign, Illinois 61826*

Received November 24, 1999; revised March 10, 2000

The performance of a nuclear resonance detection system can be quantified using binary detection theory. Within this framework, signal averaging increases the probability of a correct detection and decreases the probability of a false alarm by reducing the variance of the noise in the average signal. In conjunction with signal averaging, we propose another method based on feedback control concepts that further improves detection performance. By maximizing the nuclear resonance signal amplitude, feedback raises the probability of correct detection. Furthermore, information generated by the feedback algorithm can be used to reduce the probability of false alarm. We discuss the advantages afforded by feedback that cannot be obtained using signal averaging. As an example, we show how this method is applicable to the detection of explosives using nuclear quadrupole resonance.

Key Words: **binary signal detection; receiver operating characteristics; feedback control; nuclear quadrupole resonance; explosives detection.**

1. INTRODUCTION

Nuclear resonance provides a powerful tool for the noninvasive detection of materials such as underground water (*1*), explosives (*2*), and narcotics (*3*). Detection performance is limited by a poor signal-to-noise ratio (SNR). As a specific example, we consider a nuclear quadrupole resonance (NQR) system that detects explosives by revealing the presence of ^{14}N. Although ^{14}N is essentially 100% abundant, the small zero-field NQR splitting reduces the SNR causing missed detections and false alarms. In potential applications, such as humanitarian demining, the probability of correct detection must be over 90%, and preferably over 99.9%, while maintaining a reasonably small probability of false alarm (*4*).

Signal averaging and feedback optimization of pulse parameters are two methods for improving the SNR in detection experiments. Signal averaging increases the SNR by decreasing the amplitude of the noise component in the average signal (*5*). In the case of uncorrelated noise, the noise amplitude is reduced by the square root of the number of signals averaged. The SNR improvement is limited in practice by two factors. First, signal averaging is not as effective in reducing the amplitude of correlated noise (*6*). Second, the number of waveforms available for averaging is constrained by the time allotted for detection. Although signal averaging decreases the noise amplitude to some extent, it does not increase the amplitude of the nuclear resonance signal.

The amplitude of the nuclear resonance signal is determined by the pulse parameters. In most detection experiments there is insufficient a priori information to choose pulse parameter values that maximize the signal amplitude. For the case of landmine detection using NQR, the optimum values depend on unknown factors such as the temperature of the explosive and the location of the explosive with respect to the search coil (*7*). Under these circumstances one cannot obtain the maximum SNR using a fixed set of pulse parameters.

We previously demonstrated that feedback can be used to optimize a single pulse parameter, either pulse width or offset frequency, in the strong off-resonant comb (SORC) sequence (*8, 9*). By adjusting pulse parameters automatically, feedback provides a way of increasing the amplitude of the nuclear resonance signal.

This paper quantifies the effect of feedback on the performance of a nuclear resonance detection system using binary signal detection theory. The effectiveness of the detection system is measured in terms of the probability of correct detection and the probability of false alarm. These probabilities depend on the statistics of the noise source, the amplitude of the nuclear resonance signal, and a detection threshold used to decide whether a signal is present. By plotting the probability of correct detection versus the probability of a false alarm for varying threshold values, a receiver operating characteristic (ROC) curve is generated. This curve provides a graphical means of assessing the performance of the detection system.

We examine the effect of feedback on the ROC curve of an NQR detection system. Both theoretical predictions and experimental results reveal that the feedback algorithm raises the probability of correct detection by increasing the signal amplitude. Furthermore, the feedback algorithm generates information that can be used to reduce the probability of false alarm.

1090-7807/00 $35.00

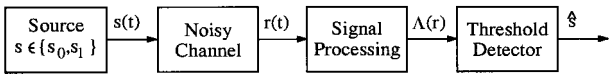

FIG. 1. Block diagram of a binary signal detection system.

Section 2 contains a brief review of binary signal detection theory and a discussion of how signal averaging and feedback affect the ROC curve. An NQR detection system and a feedback algorithm that maximizes signal power by automatically adjusting two pulse parameters is described in Section 3. Experimental results are presented in Section 4.

2. BINARY SIGNAL DETECTION

Binary signal detection theory (*10*) is a useful tool for studying the effect of signal averaging and feedback on the detection of nuclear resonance signals in noise. We first consider the simplified communication system in Fig. 1, where the source s is restricted to be in one of two states, either s_0 or s_1. In state s_i the source transmits a signal $s(t) = s_i(t)$, where $i = 0$ or 1. The received signal $r(t)$ is a combination of the source signal $s(t)$ and noise from the channel. By observing $r(t)$, the receiver must decide whether the source is in state s_0 or s_1. This is accomplished in two stages. First, a signal processing unit transforms the received signal $r(t)$ into a scalar metric $\Lambda(r)$. Second, a threshold detector compares $\Lambda(r)$ to a threshold Γ to determine the receiver output $\hat{s}$ using the decision rule

$$\hat{s} = \begin{cases} s_0, & \Lambda(r) < \Gamma \\ s_1, & \Lambda(r) \geq \Gamma. \end{cases} \quad [1]$$

The function $\Lambda(r)$ and the threshold value Γ are chosen by the designer.

The simplified communication system in Fig. 1 provides a model of the nuclear resonance detection system. In this context the state s_1 represents the presence of a material that produces a nuclear resonance signal, while s_0 represents the absence of the material. The channel accounts for the presence of any noise source that corrupts the measurement of $s(t)$. The receiver is the spectrometer, and the scalar metric may be the signal power, magnitude of a spectral peak, or some other function of the received signal. By comparing the metric to a threshold, the system decides whether a nuclear resonance signal is present.

There are four possible outcomes of a binary detection experiment as shown in Fig. 2. The performance of the detection system is typically characterized by determining the probability of correct detection P_d and probability of false alarm P_f. These probabilities can be either predicted when the statistics of the signal metric $\Lambda(r)$ are known or computed from experimental data. The probabilities P_d and P_f are a function of the threshold Γ. The ROC curve in Fig. 3 is a two-dimensional plot

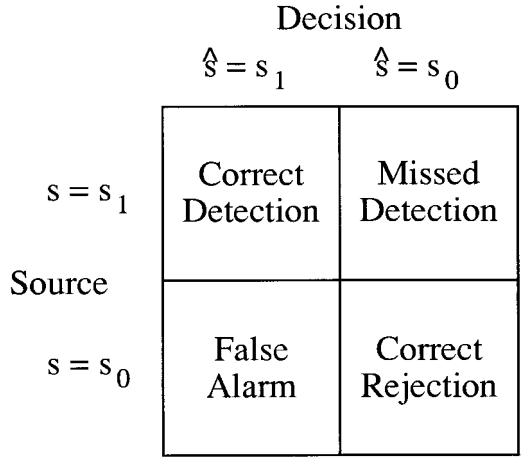

FIG. 2. Matrix of binary decision results.

of these probabilities, where each point on the curve corresponds to a fixed threshold Γ. Ideally, we would like P_f and P_d to be as close as possible to 0 and 1, respectively.

To better understand the information displayed by an ROC curve, we assume that the signal metric $\Lambda(r)$ has a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with mean μ and variance σ^2. The mean is determined by the state of the source; we let $\mu = 0$ when $s = s_0$ and $\mu = d$ when $s = s_1$. In addition, we assume that the variance σ^2 is the same for both states of the source. Under these assumptions the probability density function of the signal metric is

$$p(\Lambda(r)) = \begin{cases} \mathcal{N}(0, \sigma^2), & s = s_0 \\ \mathcal{N}(d, \sigma^2), & s = s_1. \end{cases} \quad [2]$$

Figure 4 shows an example of the probability density func-

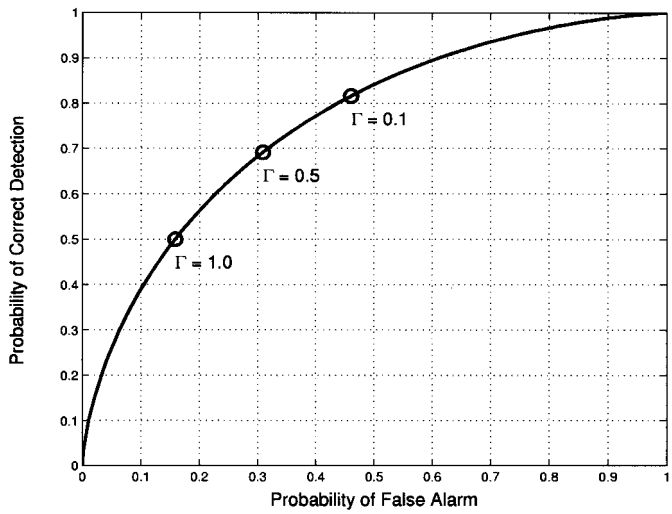

FIG. 3. Receiver operating characteristic curve.

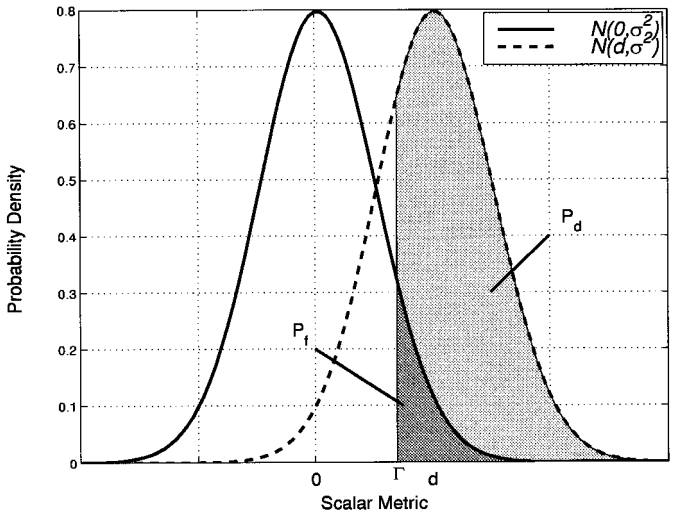

FIG. 4. Probability density functions of $\Lambda(r)$ for $s = s_0$ and $s = s_1$.

tion $p(\Lambda(r))$ for the two different states of the source. Given these distributions the probabilities of a correct detection and a false alarm can be calculated. The probability of a correct detection P_d is the light gray area under the density function $\mathcal{N}(d, \sigma^2)$ for $\Lambda(r) \geq \Gamma$. Similarly, the probability of a false alarm P_f is the dark gray area under the density function $\mathcal{N}(0, \sigma^2)$ for $\Lambda(r) \geq \Gamma$. These probabilities can be expressed as

$$P_d = \begin{cases} 1 - Q((d - \Gamma)/\sigma) & \Gamma < d \\ Q((\Gamma - d)/\sigma) & \Gamma \geq d \end{cases}$$

$$P_f = \begin{cases} 1 - Q(-\Gamma/\sigma) & \Gamma < 0 \\ Q(\Gamma/\sigma) & \Gamma \geq 0, \end{cases} \qquad [3]$$

where $Q(x)$ is a decreasing function equal to the integral of the tail of a normalized Gaussian distribution

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty \exp\left(-\frac{y^2}{2}\right) dy, \quad x \geq 0. \qquad [4]$$

It is evident from [3] and Fig. 4 that the probabilities P_d and P_f are a function of Γ, σ, and d.

The threshold Γ is chosen so that the performance of the receiver is optimum in some sense. There are several different definitions of optimality that result in different values of Γ. Three common approaches for choosing Γ are the Bayes, minimax, and Neyman–Pearson methods. In the Bayes method the overall cost of making a decision is obtained by assigning risks to the four possible outcomes of the binary detection experiment. The value of Γ is chosen to minimize the overall cost. This approach requires a priori knowledge of the probability of the events $s = s_0$ and $s = s_1$. In situations where these a priori probabilities are unknown, the other two methods are used. The minimax approach chooses Γ to minimize the maximum risk associated with choosing either s_0 or s_1. In the Neyman–Pearson method, the threshold Γ is chosen to place a bound on the maximum probability of a false alarm. The Neyman–Pearson method is used in detection applications where the a priori probabilities are unknown and it is unacceptable to have a large false alarm rate.

It is intuitive that the SNR of the signal metric should have an effect on the ROC curve. This relationship becomes evident if the SNR of $\Lambda(r)$ is expressed as (*10*)

$$\text{SNR} = 20 \log\left(\frac{d}{\sigma}\right). \qquad [5]$$

Observe that the SNR of the signal metric can be increased by either reducing the variance σ^2 and/or increasing the mean d. Each of these approaches has a different effect on the ROC curve.

For a fixed threshold Γ, decreasing the variance σ^2 increases the probability of correct detection and decreases the probability of false alarm. In contrast, increasing the mean d raises the probability of correct detection, but has no effect on the probability of false alarm. These effects can be inferred from Fig. 4 and are illustrated by the ROC curves in Fig. 5 for an operating point (P_d, P_f) defined by $\Gamma = 0.5$. Figure 5a shows the ROC curves for a fixed mean $d = 1$ and three different variances, while Fig. 5b shows the ROC curves for a fixed variance $\sigma^2 = 1$ and three different means.

We now consider how signal averaging and feedback optimization can affect the ROC curve in nuclear resonance detection systems. Signal averaging is frequently used to improve the SNR of nuclear resonance signals (*5*). In reference to the model binary detection system in Fig. 1, the received signal $r(t)$ is replaced by an average signal $\bar{r}(t)$. For uncorrelated noise, the variance of the noise in the average signal decreases by a factor N_A, where N_A is the number of signals averaged. Depending on the signal metric, a similar reduction is obtained in the variance of $\Lambda(r)$. As a result, signal averaging increases the probability of correct detection and decreases the probability of a false alarm.

In the general case of a communications system, the signal $s_1(t)$ is typically fixed, resulting in a signal metric with a constant mean d. However, for nuclear resonance detection systems, the signal $s_1(t)$ can be altered by adjusting the pulse parameters. The signal mean $d = d(\Theta, \Phi)$ is a function of controlled Θ and exogenous Φ inputs. The controlled inputs are adjustable and include parameters such as pulse width and pulse separation. Exogenous inputs cannot be altered by the detection system. If the exogenous inputs are unknown, it is not possible to choose, prior to the experiment, controlled inputs that maximize the signal metric mean. As an example, nuclear magnetic resonance provides a means to detect and determine

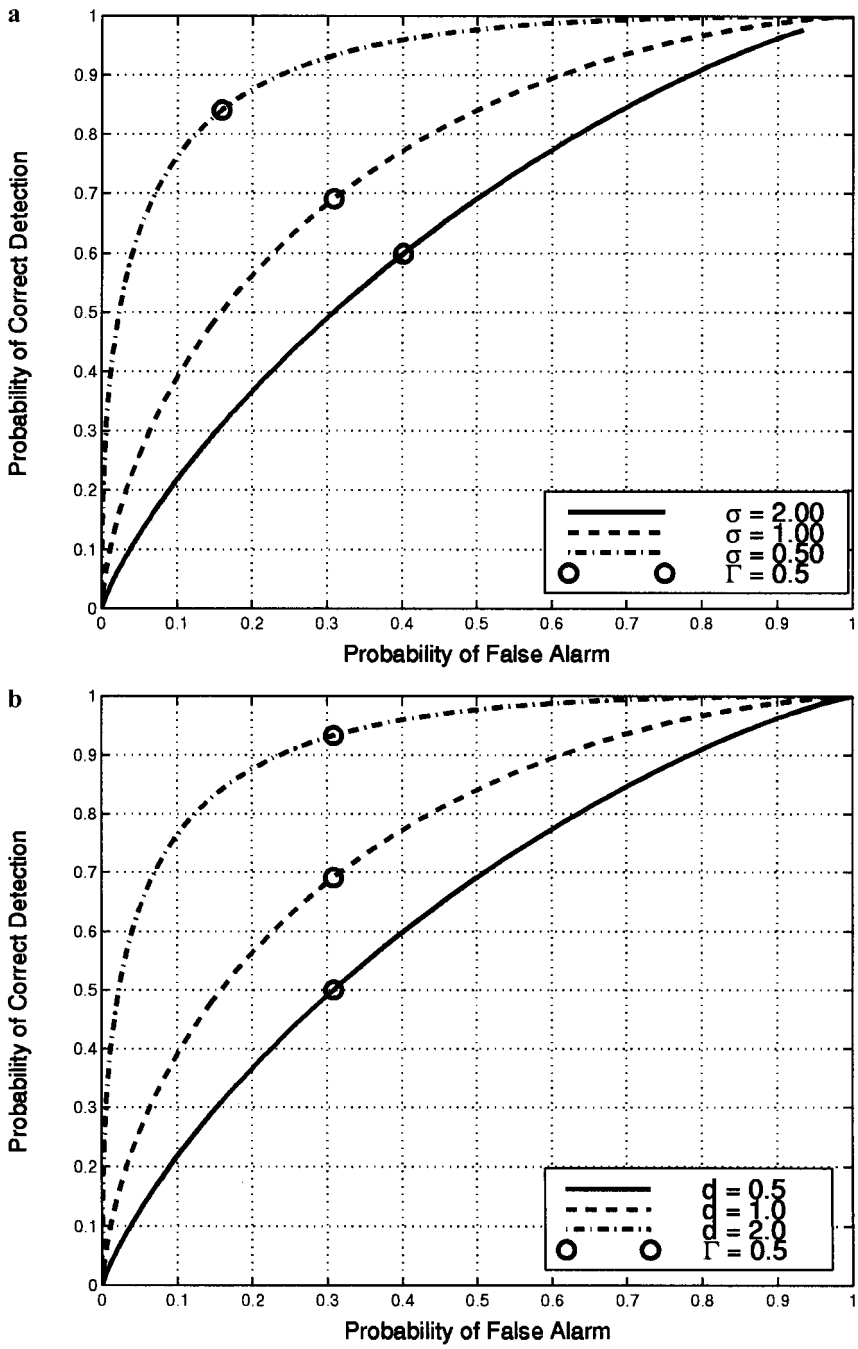

FIG. 5. ROC curves for (a) fixed mean $d = 1$ and three different variances and (b) fixed variance $\sigma^2 = 1$ and three different means.

the depth of groundwater (*1*). In this case, water depth is an exogenous input that is unknown, and so it is not possible to initially choose pulse parameters that yield the largest nuclear resonance signal.

In a series of earlier papers we showed that even when the exogenous inputs are unknown, it is possible to maximize the mean of the signal metric by automatically adjusting the controlled inputs (*8, 9*). As illustrated in Fig. 5b, maximizing the signal mean d increases the probability of correct detection. If only a single signal metric $\Lambda(r)$ is used, then the probability of false alarm remains unchanged. It would be advantageous if this probability could be lowered using data provided by the feedback algorithm.

We now show that the feedback system does provide information that reduces the probability of a false alarm. The algorithm presented in Section 3 iteratively tunes two separate

pulse parameters to maximize the mean d. The kth iteration generates a signal metric $\Lambda(r(k))$, and the algorithm terminates after completing N iterations. The detection rule in [1] uses the single metric $\Lambda(r(N))$ to make a decision. We consider a modified detection rule based on majority voting, where the receiver output is set to state s_1 only if the majority of the signal metrics $\Lambda(r(1)), \ldots, \Lambda(r(N))$ exceed the threshold value (*11*).

The majority-voting rule produces a probability of correct detection P_d^m and probability of false alarm P_f^m. These probabilities are calculated in terms of the number of iterations N and the probabilities of correct detection $P_d(k)$ and false alarm $P_f(k)$ based on a single value of the signal metric $\Lambda(r(k))$. We assume that the probability of false alarm is the same for each iteration because the control algorithm does change the variance σ^2 and the threshold Γ is fixed. For now, we also assume that the pulse parameters are fixed, so that the probability of correct detection is the same for each iteration.

Let P_i^m represent the probability of a correct detection ($i = d$) or false alarm ($i = f$) using the majority-voting rule for N iterations. The probability of correct detection or false alarm for a single iteration is denoted as P_i. In N iterations the number of correct detections or false alarms is a random variable with a binomial distribution. It follows that the probability of obtaining at least $\lceil N/2 \rceil$ correct decisions or false alarms is

$$P_i^m = \sum_{k=\lceil N/2 \rceil}^{N} \binom{N}{k} (P_i)^k (1 - P_i)^{N-k}, \qquad [6]$$

where $\lceil N/2 \rceil$ is the smallest integer greater than or equal to $N/2$. Figure 6 shows P_i^m as a function of P_i for three different values of N. The solid curve shows that when $N = 1$, the probabilities obtained using majority voting and a single vote are identical as expected. For N greater than 1, the majority-voting probability P_i^m is less than the individual probability P_i when $P_i < \frac{1}{2}$. On the other hand, when $P_i > \frac{1}{2}$, the majority-voting probability is greater than the single probability.

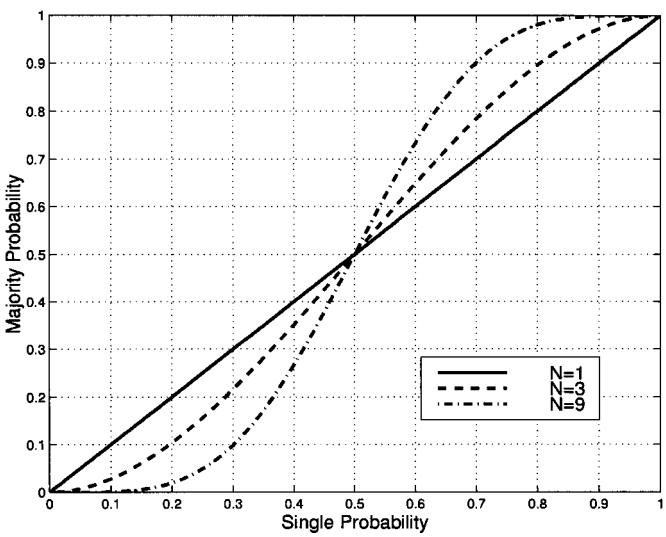

FIG. 6. Relationship between the majority probability P_i^m and the probability P_i of a single decision.

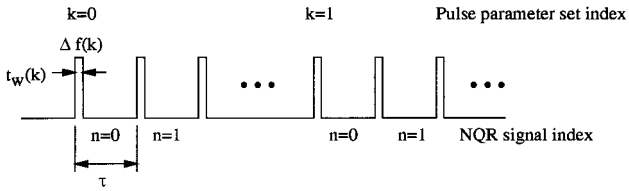

FIG. 7. SORC pulse sequence notation.

Figure 6 shows that the majority-voting rule decreases the probability of a false alarm when the probability of false alarm for each iteration is less than one-half. Likewise, the majority-voting rule increases the probability of correct detection as long as the probability of correct detection for a single iteration is greater than one-half.

Applying the majority-voting rule is not helpful when the pulse parameters result in a signal metric whose mean is less than the threshold. In this case, because the probability $P_d(k)$ of correct detection is less than one-half, a reduction in the probability P_f^m of false alarm is accompanied by a reduction in the probability P_d^m of correct detection. On the other hand, if a feedback algorithm raises $P_d(k)$ above one-half, then the majority-voting rule lowers P_f^m and raises P_d^m. Experimental results presented in Section 4 support this claim.

3. NQR DETECTION SYSTEM

This section describes an NQR detection system that utilizes feedback control to optimize multiple pulse parameters. NQR signals are generated using the SORC sequence, a periodic series of identical RF pulses (*12*). As the sequence progresses, the NQR signal observed between the RF pulses reaches a steady-state waveform. The amplitude of the SORC waveform depends on the pulse separation τ, pulse width t_w, and offset frequency Δf that define the controlled inputs Θ. The SORC signal intensity is also dependent upon temperature (*13*), distance between the sample and the coil (*2*), and composition of the substance to be detected (*14*). This collection of variables defines the exogenous inputs Φ.

The SORC pulse sequence is shown in Fig. 7. The control algorithm automatically adjusts the pulse width $t_w(k)$ and offset frequency $\Delta f(k)$ for a constant pulse separation τ. The controlled inputs are held constant over a fixed number of RF pulses so that the SORC signal reaches a steady-state waveform before N_A waveforms are acquired for signal averaging.

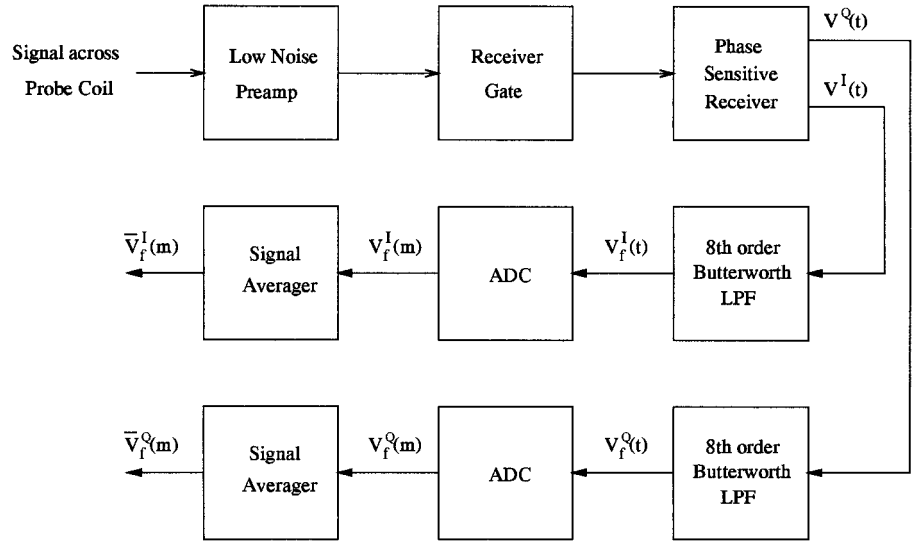

FIG. 8. Block diagram of the receiver and data acquisition system.

The index k references the pulse parameter set, while individual NQR waveforms for each parameter set are indexed by n.

Figure 8 shows the block diagram of the receiver and data acquisition system. The receiver is gated on 240 μs after the RF pulse to avoid saturation due to ringing in the probe coil. The phase-sensitive receiver mixes the gated signal induced in the probe coil with the excitation frequency $\nu_- + \Delta f$. The in-phase $V^I(t; k)$ and quadrature $V^Q(t; k)$ components of the receiver output are passed through identical eighth-order Butterworth lowpass filters with a cutoff frequency of 10 kHz to produce $V_f^I(t; k)$ and $V_f^Q(t; k)$, respectively. The signals $V_f^I(m; k)$ and $V_f^Q(m; k)$ are obtained by sampling the filtered signals at 100 kHz, where $m = 0, 1, \ldots, M - 1$ and M is the number of sample points per waveform. The sampled signals are added point by point to form the averages $\bar{V}_f^I(m; k)$ and $\bar{V}_f^Q(m; k)$ for N_A SORC signals.

A signal metric is needed for both the control algorithm and the threshold detector. A wide variety of signal metrics is available, including peak-to-peak voltage, midsignal amplitude, signal energy, and peak FFT magnitude (*9*). We use the SORC signal power defined as

$$\mathcal{P}(k) = \frac{1}{M}\sum_{m=0}^{M-1} (\bar{V}_f^I(m; k) - V_{DC})^2, \qquad [7]$$

where V_{DC} is the DC offset measured at the filter output in the absence of an NQR signal.

The statistics of the signal metric $\mathcal{P}(k)$ are required to investigate the effect of noise on the dual tuning algorithm. In the absence of noise, the signal metric $\mathcal{P}_s(k)$ is defined as

$$\mathcal{P}_s(k) = \frac{1}{M}\sum_{m=0}^{M-1} s^2(m; k), \qquad [8]$$

where s is the steady-state SORC waveform obtained using the controlled input $\Theta(k)$. In the presence of noise the signal metric is expressed as

$$\mathcal{P}_{sn}(k) = \frac{1}{M}\sum_{m=0}^{M-1} [s(m; k) + n(m)]^2, \qquad [9]$$

where $n(m)$ is a random variable that does not depend on the controlled input $\Theta(k)$. We assume that the $n(m)$ are independent and identically distributed normal random variables with zero mean and variance σ_n^2. The signal metric $\mathcal{P}_{sn}(k)$ is also a random variable with a noncentral χ^2 density of degree M. Because M is on the order of 100, this density can be approximated by a normal distribution with mean $\mu_{sn}(k)$ and variance $\sigma_{sn}^2(k)$ (*10*),

$$\mu_{sn}(k) = \mathcal{P}_s(k) + \sigma_n^2$$
$$\sigma_{sn}^2(k) = \frac{4\sigma_n^2}{M}\mathcal{P}_s(k) + \frac{2\sigma_n^4}{M}. \qquad [10]$$

The dual tuning algorithm is an extension of methods developed earlier for optimizing the pulse width (*8*) and the offset frequency (*9*). A gradient ascent algorithm iteratively maximizes the signal metric $\mathcal{P}(k)$ using the update rules

$$t_w(k+1) = t_w(k) + \lambda_{t_w}(k)\nabla_{t_w}\mathcal{P}(k-1) \quad k = 2, 4, 6, \cdots \tag{11}$$

$$\Delta f(k+1) = \Delta f(k) + \lambda_{\Delta f}\nabla_{\Delta f}\mathcal{P}(k-1) \quad k = 3, 5, 7, \cdots \tag{12}$$

to automatically adjust the pulse width $t_w(k)$ and offset frequency $\Delta f(k)$. The initial pulse parameter values $t_w(0)$, $t_w(1)$, $\Delta f(0) = \Delta f(1)$, and $\Delta f(2)$ are chosen based on estimated values of the exogenous inputs. The gradients $\nabla_{t_w}\mathcal{P}(k-1)$ and $\nabla_{\Delta f}\mathcal{P}(k-1)$ are approximated as (*8, 9*)

$$\nabla_{t_w}\mathcal{P}(k-1) = G(\mathcal{P}(k-1) - \mathcal{P}(k-2)) \times G(t_w(k-1) - t_w(k-2)) \tag{13}$$

$$\nabla_{\Delta f}\mathcal{P}(k-1) = G(\mathcal{P}(k-1) - \mathcal{P}(k-2)) \times G(\Delta f(k-1) - \Delta f(k-2)), \tag{14}$$

where

$$G(x) = \begin{cases} 1 & x \geq 0 \\ -1 & x < 0. \end{cases} \tag{15}$$

The gradients $\nabla_{t_w}\mathcal{P}(k)$ and $\nabla_{\Delta f}\mathcal{P}(k)$ only take on values of ± 1. The sign determines the direction in which the parameter should be incremented in order to increase the signal metric. The learning factors $\lambda_{t_w}(k)$ and $\lambda_{\Delta f}$ determine the magnitude of the change in pulse parameters.

Because the gradients in [13] and [14] are always nonzero, the gradient algorithm converges to a limit cycle instead of a single point. If the metric maxima are sharp with respect to parameter variations, the limit cycle will occur about the optimal value. This is the case for offset frequency. In contrast, the metric maxima are not as sharp for the pulse width. To reduce the radius of the limit cycle about the optimal pulse width the learning factor $\lambda_{t_w}(k)$ is decremented as (*8*)

$$\lambda_{t_w}(k) = \begin{cases} 20\ \mu s & k = 0, 2, \ldots, 8 \\ 15\ \mu s & k = 10, 12, \ldots, 18 \\ 10\ \mu s & k > 20. \end{cases} \tag{16}$$

A fixed learning factor $\lambda_{\Delta f} = 50$ Hz is used for tuning the offset frequency (*9*). The control algorithm tunes the controlled inputs for a fixed number of iterations. The number of iterations is chosen sufficiently large to allow convergence when an NQR signal is present.

It is important to consider the effect of noise on the gradient tuning algorithm. We derive an expression for the probability that the algorithm moves the controlled input in a direction that maximizes the signal metric on any given iteration. Consider the simplified update algorithm for a single parameter θ,

$$\theta(k+1) = \theta(k) + \lambda\nabla_\theta\mathcal{P}(k), \tag{17}$$

where the gradient is approximated as

$$\nabla_\theta\mathcal{P}(k) = G(\mathcal{P}(k) - \mathcal{P}(k-1))G(\theta(k) - \theta(k-1)). \tag{18}$$

The following definitions are needed:

$$\Delta\mathcal{P}_s(k) = \mathcal{P}_s(k) - \mathcal{P}_s(k-1) \tag{19}$$

$$\Delta\mathcal{P}_{sn}(k) = \mathcal{P}_{sn}(k) - \mathcal{P}_{sn}(k-1). \tag{20}$$

From [10], it follows that $\Delta\mathcal{P}_{sn}(k)$ is a normal random variable with mean $\mu_{\Delta sn}(k)$ and variance $\sigma^2_{\Delta sn}(k)$,

$$\mu_{\Delta sn}(k) = \Delta\mathcal{P}_s(k)$$
$$\sigma^2_{\Delta sn}(k) = \frac{4\sigma_n^2}{M}(\mathcal{P}_s(k) + \mathcal{P}_s(k-1)) + \frac{4\sigma_n^4}{M}. \tag{21}$$

The probability that the algorithm moves the parameter $\theta(k)$ in the correct direction at the kth iteration is

$$P_c(k) = P\{G(\Delta\mathcal{P}_{sn}(k)) = G(\Delta\mathcal{P}_s(k))\}. \tag{22}$$

Using Bayes theorem and the definition of $G(x)$, the probability $P_c(k)$ can be expressed as

$$P_c(k) = P(\Delta\mathcal{P}_{sn}(k) \geq 0 | \Delta\mathcal{P}_s(k) \geq 0)\, P(\Delta\mathcal{P}_s(k) \geq 0) + P(\Delta\mathcal{P}_{sn}(k) < 0 | \Delta\mathcal{P}_s(k) < 0)\, P(\Delta\mathcal{P}_s(k) < 0) \tag{23}$$

$$= \left[1 - Q\left(\frac{\mu_{\Delta sn}}{\sigma_{\Delta sn}}\right)\right] P(\Delta\mathcal{P}_s(k) \geq 0) + \left[1 - Q\left(\frac{-\mu_{\Delta sn}}{\sigma_{\Delta sn}}\right)\right] P(\Delta\mathcal{P}_s(k) < 0) \tag{24}$$

$$= 1 - Q\left(\frac{|\mu_{\Delta sn}|}{\sigma_{\Delta sn}}\right). \tag{25}$$

Figure 9 shows the probability of a correct update as a function of $|\mu_{\Delta sn}|/\sigma_{\Delta sn}$. This result shows that $P_c(k)$ is always greater than one-half and approaches 1 when $|\mu_{\Delta sn}|/\sigma_{\Delta sn} > 1$. It follows that the performance of the feedback algorithm is improved when signal averaging is used to reduce $\sigma_{\Delta sn}$. In addition, the more sensitive the signal metric is to changes in the pulse parameter, that is, the larger the value of $\mu_{\Delta sn}$, the more likely the algorithm will update correctly.

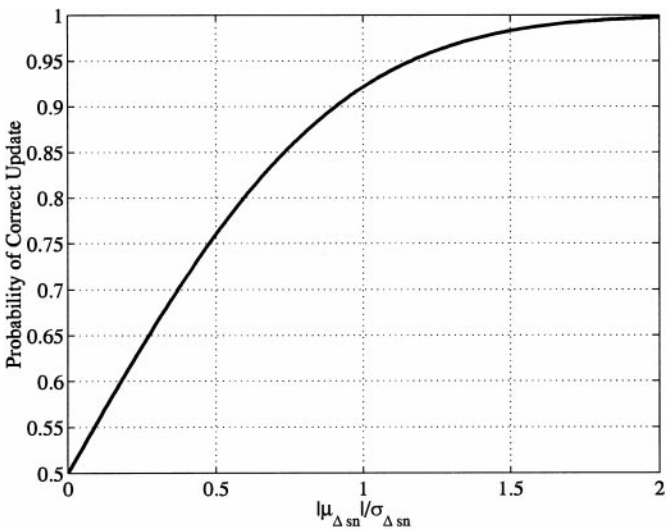

FIG. 9. Probability of gradient algorithm correctly updating parameter.

4. EXPERIMENTAL RESULTS

This section presents experimental data that demonstrate the ability of the dual tuning algorithm to optimize the pulse width and offset frequency. The results from a series of detection experiments show the effect of the dual tuning algorithm on the ROC curve. A 50-g sample of sodium nitrite at room temperature serves as a model compound for the explosive RDX. All experiments were conducted near the ν_- transition at 3.6 MHz. The ν_- transition is dominated by a single spin–lattice relaxation time $T_{1\ell}$ of 0.3 s. The spin–spin relaxation time T_2 and the T_2^* lineshape parameter are 6 and 1 ms, respectively.

A custom-made 1-kW pulsed spectrometer enables real-time adjustments of the controlled inputs during an experiment (*15*). The sodium nitrite is packed into a 1-inch diameter, 3-inch-long glass vial over which the probe coil is tightly wound. The quality factor of the probe coil is approximately 150. The statistical properties of the measurement noise are controlled by mounting the probe coil within a shielded enclosure along with a one-turn coil connected to an external lowpass Gaussian noise source with a cutoff frequency of 10 MHz. The amplitude of the noise source is adjusted so that the noise level at the receiver output matched that for an unshielded coil in our laboratory. The amplitude of the applied RF pulses is fixed so that the magnetic field amplitude at the coil center is approximately 5 G, which is comparable to the field strength used in a commercial NQR detection system (*7*). The corresponding output power of the transmitter is approximately 25 W. In all experiments the pulse separation is set at 1 ms.

The dependence of the SORC signal power on pulse width and offset frequency is experimentally measured. Holding the pulse width constant and with the external noise source turned off, the steady-state SORC signal was recorded for offset frequencies ranging from −4 to 4 kHz in 50-Hz steps. This experiment was repeated for pulse widths ranging from 10 to 300 μs in 10-μs steps. For each pulse width and offset frequency, $N_A = 1000$ steady-state SORC signals were filtered and averaged. The signal power $\mathcal{P}(t_w, \Delta f)$ for each pulse parameter set was calculated using [7].

The relationship between steady-state signal power and pulse parameters is shown in Fig. 10 and is consistent with earlier results (*12, 16*). As the offset frequency approaches zero, the NQR signal vanishes. The series of steep valleys arises from interference between the NQR signals immediately following and preceding consecutive RF pulses. The ratio of the maximum and minimum observed signal power is approximately 10, corresponding to a 20-dB variation in SNR across the pulse parameters used to generate Fig. 10.

The ability of the dual tuning algorithm to maximize the signal power and improve the ROC curve is demonstrated using a series of experiments with the external noise source connected. In each experiment the initial pulse widths and offset frequencies are chosen as $t_w(0) = 20$ μs, $t_w(1) = 40$ μs, $\Delta f(0) = \Delta f(1) = 0$ Hz, and $\Delta f(2) = 50$ Hz, respectively. These parameters are purposely chosen to produce a small signal power. The exact offset frequency is unknown due to temperature variations of the sample. In each experiment the dual tuning algorithm is terminated after 40 iterations, providing 20 updates in pulse width and 20 updates in offset frequency.

Results from a single tuning experiment without averaging ($N_A = 1$) are shown in Fig. 11. Figure 11a shows the normalized signal power at each iteration of the dual tuning algorithm, while Figs. 11b and 11c show the corresponding values of offset frequency and pulse width. The squares and circles show data obtained with and without a sample present, respectively,

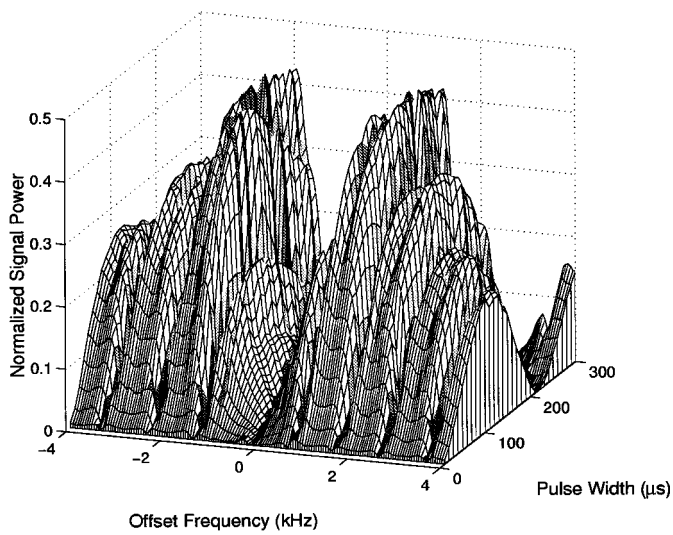

FIG. 10. SORC signal power as a function of pulse width and offset frequency.

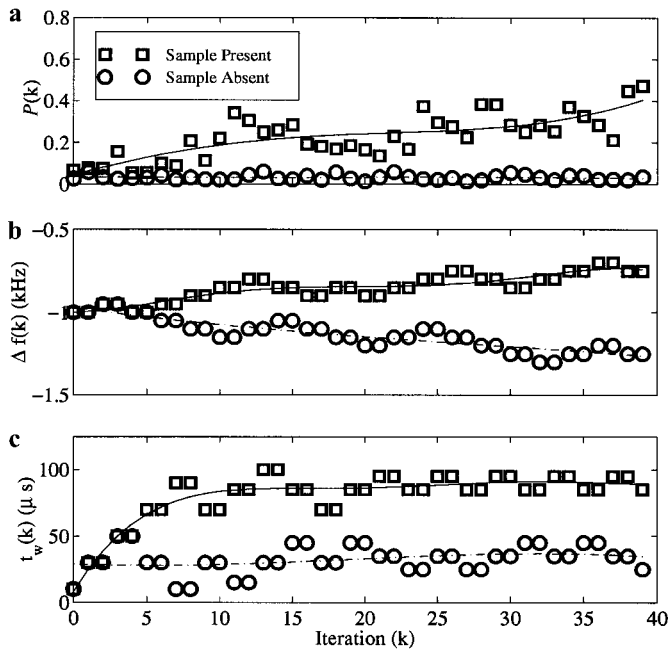

FIG. 11. Tuning history when the signal power is derived from a single ($N_A = 1$) SORC waveform.

whereas the solid curves are best-fit polynomials that show the trajectory of the signal power and tuned parameters. In the absence of a sample, only a small variation in the signal power is observed. In contrast, when the sample is present, the dual tuning algorithm significantly increases the signal power and hence the SNR of the NQR measurement.

The ROC curves in Fig. 12 demonstrate the effect of feedback on the probabilities of correct detection and false alarm. Data for the ROC curves are obtained in a series of 200 experiments, half with the sample present and half with the sample absent. For each experiment, the signal power is compared against the threshold Γ to decide whether a sample is present. Let $N_d(N_f)$ represent the number of experiments where the sample is present (absent) and the signal power is larger than Γ. The probabilities of correct detection and false alarm are $N_d/100$ and $N_f/100$, respectively. In order to generate the ROC curve the probabilities of correct detection and false alarm must be generated for a range of different threshold values. Taking into account the peak signal power in Fig. 10 and the variance of the external noise source, each ROC curve is generated using 100 values of Γ ranging from 0 to 1 in 0.01 steps.

The solid ROC curve in Fig. 12 is generated by fixing the pulse parameters at the initial values. The dashed and the dash–dotted ROC curves are generated using the signal power from the last iteration of the dual tuning algorithm for three and eight pulse parameter updates, respectively. The dual tuning algorithm increases the probability of a correct detection as indicated by the open circles that correspond to a fixed threshold Γ. Taking into account the statistical variation between experiments, there is little change in the probability of false alarm. Results obtained with the majority-voting rule for three interactions of the tuning algorithm are shown by the dotted curve. In comparison to the dashed ROC curve, the majority-voting rule increased the probability of correct detection and decreased the probability of false alarm.

5. DISCUSSION

Figure 12 shows that the dual feedback algorithm increases the probability of correct detection, but has limited effect on the probability of false alarm unless the majority-voting rule is used. Compare this result against the theoretical ROC curves shown in Fig. 5. A similar effect is achieved by increasing the signal metric mean. This comparison suggests two conclusions. First, the binary signal detection system shown in Fig. 1 can be used as a conceptual model for describing the performance of a nuclear resonance detection system. Second, adjusting the pulse parameters increases the signal metric mean, an improvement that is not attainable using signal averaging.

Signal averaging reduces the variance of the noise in the signal metric, but does not increase the mean of the signal metric. In situations where the noise is correlated or the available data are limited, the improvement afforded by signal averaging is constrained. Under these circumstances, adjusting the pulse parameters to increase the NQR signal power provides another tool for improving the ROC curve.

Having motivated the importance for adjusting the pulse parameters to improve the ROC curve, we now consider methods by which the adjustments can be made. One possibility is to systematically apply different pulse parameter sets and record the corresponding signal metrics. Depending upon the

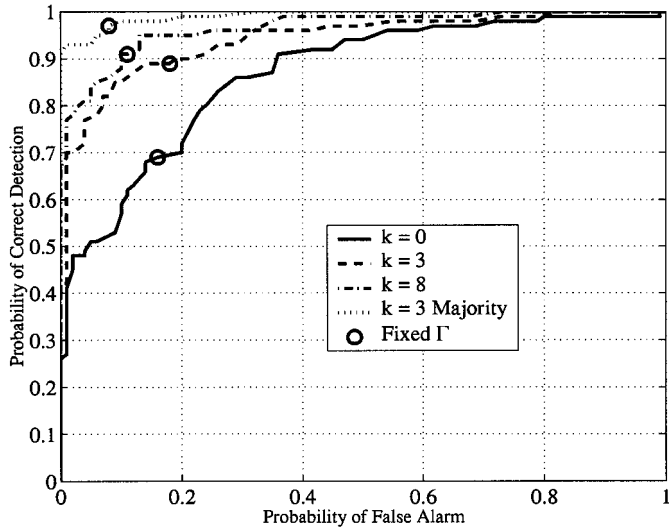

FIG. 12. ROC curves for tuned and untuned pulse parameters.

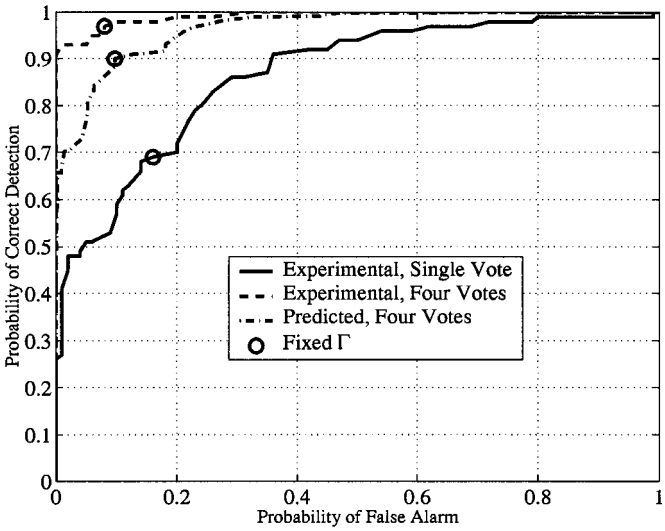

FIG. 13. Experimental and predicted ROC curves using the majority-voting rule.

degree of improvement sought, a large number of pulse parameters must be used, thereby increasing the detection time. We demonstrated another approach, based on feedback control concepts, that utilizes the measured signal to adjust pulse parameters automatically, thereby maximizing the NQR signal power. In comparison to the former approach, the feedback algorithm determines the optimal parameters using fewer iterations.

In addition to increasing the probability of correct detection, we also showed that the feedback algorithm generates a series of signal metric values that can be used to reduce the probability of false alarm. A majority-voting rule was implemented, and the results in Fig. 12 show a reduction in the probability of false alarm and an increase in the probability of correct detection.

Figure 13 shows that the observed reduction in the probability of false alarm is consistent with that predicted by [6]. The solid curve is the ROC curve obtained using the untuned parameters and a decision rule based on a single vote. These data, along with [6], were used to predict the ROC curve (dash–dotted) that would be obtained for the same fixed pulse parameters and a majority-voting rule based on four votes. In comparison, the experimentally generated ROC curve using the dual tuning algorithm with a majority-voting rule based on four votes is shown by the dashed curve. The improvement between the predicted and experimental ROC curves using majority voting is due to the increase in the probability of correct detection following each iteration of the dual tuning algorithm.

Two modifications are needed to further improve the performance of the dual tuning algorithm. First, the current algorithm may become trapped at a local maximum. Referring to Fig. 10, observe that the maxima on either side of exact resonance are smaller than the surrounding maxima. Second, the algorithm requires, as an input, the steady-state response of the magnetization to the pulse sequence. This is not a limitation when the spin–lattice relaxation time is short, as is the case for RDX. On the other hand, TNT has a much larger spin–lattice relaxation time than RDX and a short spin–spin relaxation time. As a result, for TNT, it is difficult to obtain a steady-state waveform. For these situations, the dual tuning algorithm must be modified to take into account the dynamic response of the magnetization signal to the pulse sequence. We are currently developing a dual tuning algorithm that overcomes both of these limitations by taking into account the dynamic behavior of the NQR response.

6. CONCLUSION

This paper investigated the effect of feedback on the ROC curve of an NQR detection system. We successfully demonstrated that a gradient tuning algorithm can automatically adjust two pulse parameters so that the nuclear resonance signal power is increased. As a result, the probability of correct detection is raised. We used a majority-voting rule that took advantage of the information generated at each iteration of the feedback algorithm. The majority-voting rule reduced the probability of false alarm and further increased the probability of correct detection. The use of feedback to improve the ROC curve in nuclear resonance detection systems is beneficial in situations where correlated noise is present and the pulse parameter values that yield the largest signal amplitude are initially unknown.

ACKNOWLEDGMENT

This work was funded in part by contracts from the United States Army Construction Engineering Research Laboratories.

REFERENCES

1. A. V. Legchenko and O. A. Shushakov, Inversion of surface NMR data, *Geophysics* **63**(1), 75–84 (1998).
2. T. Hirschfeld and S. M. Klainer, Short range remote NQR measurements, *J. Mol. Struct.* **58,** 63–77 (1980).
3. J. P. Yesinowski, M. L. Buess, A. N. Garroway, M. Ziegeweid, and A. Pines, Detection of ^{14}N and ^{35}Cl in cocaine base and hydrochloride using NQR, NMR, and SQUID, *Anal. Chem.* **67**(13), 2256–2263 (1995).
4. United States Department of State, Hidden killers: The global problem with uncleared landmines. United States Department of State Political–Military Affairs Bureau Office of International Security Operations (1993).
5. M. P. Klein and G. W. Barton, Enhancement of signal-to-noise ratio by continuous averaging: Application to magnetic resonance, *Rev. Sci. Instrum.* **34**(7), 754–759 (1963).
6. R. R. Ernst, Sensitivity enhancement in magnetic resonance. I.

Analysis of the method of time averaging, *Rev. Sci. Instrum.* **36**(12), 1689–1695 (1965).

7. A. D. Hibbs, G. A. Barrall, P. V. Czipott, D. K. Lathrop, Y. K. Lee, E. E. Magnuson, R. Matthews, and S. A. Vierkotter, Landmine detection by nuclear quadrupole resonance, *in* "Detection and Remediation Technologies for Mines and Minelike Targets III," (A. C. Dubey, J. F. Harvey, and T. Broach, Eds.), Vol. 3392, pp. 522–532 (1998).
8. J. L. Schiano, T. Routhier, A. J. Blauch, and M. D. Ginsberg, Feedback optimization of pulse width in the SORC sequence. *J. Magn. Reson.* **140,** 84–90 (1999).
9. A. J. Blauch, J. L. Schiano, and M. D. Ginsberg, Optimization of offset frequency in the SORC pulse sequence using feedback control, *J. Magn. Reson.* **139,** 139–151 (1999).
10. R. N. McDonough and A. D. Whalen, "Detection of Signals in Noise," 2nd ed., Academic Press, San Diego (1995).
11. P. J. Boland, Majority systems and the Condorcet jury theorem. *Statistician* **38,** 181–189 (1989).
12. S. S. Kim, J. R. P. Jayakody, and R. A. Marino, Experimental investigations of the strong off-resonant comb (SORC) pulse sequence in ^{14}N NQR. *Zeitschrift Naturforschung A J. Phys. Sci.* **47A,** 415–420 (1992).
13. T. P. Das and E. L. Hahn, Nuclear quadrupole resonance spectroscopy, *in* "Solid State Physics," Suppl. 1, Academic Press, New York (1958).
14. V. S. Grechishkin and N. Ya. Sinyavskii, New technologies: nuclear quadrupole resonance as an explosive and narcotic detection technique, *Physics-Uspekhi* **40**(4), 393–406 (1997).
15. J. L. Schiano and M. D. Ginsberg, A pulsed spectrometer designed for feedback NQR, *Zeitschrift Naturforschung A J. Phys. Sci.* **55A,** 61–66 (2000).
16. S. M. Klainer, T. B. Hirschfeld, and R. A. Marino, Fourier transform nuclear quadrupole resonance spectroscopy, *in* "Fourier, Hadamard, and Hilbert Transforms in Chemistry," pp. 147–182, Plenum Press, New York (1982).