ORNL/CP-96389

INVESTIGATION OF WAVELET-BASED ENHANCEMENTS TO NUCLEAR QUADRUPOLE RESONANCE EXPLOSIVES DETECTORS*

CONF-980412--

Stephen W. Kercel and William B. Dress
Oak Ridge National Laboratory
P.O. Box 2008
Oak Ridge, Tennessee 37831-6011
(423) 574-5278

Andrew D. Hibbs and Geoffrey A. Barrall
Quantum Magnetics, Inc.
7740 Kenamar Court
San Diego, CA 92121
(619) 566-9200

To be presented at the
SPIE Aerosense Symposium
Orlando, Florida
April 16, 1998

*This research was performed at OAK RIDGE NATIONAL LABORATORY, managed by LOCKHEED MARTIN ENERGY RESEARCH CORP. for the U.S. DEPARTMENT OF ENERGY under contract DE-AC05-96OR22464.

DISCLAIMER

Investigation of Wavelet-Based Enhancements to Nuclear Quadrupole Resonance Explosives Detectors

Stephen W. Kercel and William B. Dress

Oak Ridge National Laboratory, PO Box 2008, MS 6011, Oak Ridge TN 37831-6011
Phone: (423) 574-5278 FAX: ((423) 574-6663 e-mail: kzo@ornl.gov

Andrew D. Hibbs and Geoffrey A. Barrall

Quantum Magnetics, Inc., 7740 Kenamar Court, San Diego CA 92121
Phone: (619) 566-9200 FAX: (619) 566-9388 e-mail: andy.hibbs@qm.com

ABSTRACT

Nuclear Quadrupole Resonance (NQR) is effective for the detection and identification of certain types of explosives such as RDX, PETN and TNT. In explosive detection, the NQR response of certain 14N nuclei present in the crystalline material is probed. The 14N nuclei possess a nuclear quadrupole moment which in the presence of an electric field gradient produces an energy level splitting which may be excited by radio-frequency magnetic fields. Pulsing on the sample with a radio signal of the appropriate frequency produces a transient NQR response which may then be detected. Since the resonant frequency is dependent upon both the quadrupole moment of the 14N nucleus and the nature of the local electric field gradients, it is very compound specific.

Under DARPA sponsorship, the authors are using multiresolution methods to investigate the enhancement of operation of NQR explosives detectors used for land mine detection. For this application, NQR processing time must be reduced to less than one second. False alarm responses due to acoustic and piezoelectric ringing must be suppressed. Also, as TNT is the most prevalent explosive found in land mines NQR detection of TNT must be made practical despite unfavorable relaxation times. All three issues require improvement in signal-to-noise ratio, and all would benefit from improved feature extraction. This paper reports some of the insights provided by multiresolution methods that can be used to obtain these improvements. It includes results of multiresolution analysis of experimentally observed NQR signatures for RDX responses and various false alarm signatures in the absence of explosive compounds.

Keywords: "quadrupole resonance," "explosives detection," "wavelet feature extraction," "multiresolution analysis," "RDX signatures"

1. INTRODUCTION

Despite the profusion of sensor technologies being applied to the problem, the fact remains that the technology to detect deployed landmines has simply not progressed to a satisfactory level. The essence of the present state of countermine sensing is captured in a quote from Colonel Robert Greenwalt, "Today, highly trained, scared soldiers use all their senses, augmented with a coin detector and a stick." [1] The Defense Advanced Research Projects Agency (DARPA) offers the following advice to the developers of mine detectors. "When a soldier believes that a mine may be present - what does he reach for, a stick or your system?" [2]

A major issue in mine detection is, "What should one actually try to detect?" There are many hundreds of mine types, and many interchangeable detonation schemes. Together, these lead to many thousands of combinations of mine/detonator configurations, almost all of which have been deployed somewhere in the world. There are features that are detectable in some configurations, but not others. Some mines have significant metal content, and some do not. Some mine casing materials have distinctive chemical signatures and some do not. A practical mine detector must search for a common attribute.

The one attribute common to *all* configurations is that they contain a non-trivial quantity of chemical high explosive. Furthermore, there are comparatively few suitable explosive compounds. There is an inherent conflict in the requirements for a military high

explosive. When they explode, they must do so with great force, but in order to be safe to handle, it must be possible for them to explode only under precisely controlled conditions. Only about a half dozen chemical compounds meet this oxymoronic requirement of being a "safe explosive." Of these, TNT is the most commonly used in landmines. RDX and PETN are also used in some landmines.

If the explosive is the common attribute of all mines, then does it not make sense to detect mines by directly detecting the explosive? There are several methods of explosive-based mine detection. These include chemical vapor methods, neutron methods, X-ray backscatter, and nuclear quadrupole resonance (NQR).

Each method has its drawbacks. Chemical vapor methods suffer from the fact that mines emit extremely low vapor pressures (on the order of parts per trillion). [2] Neutron methods require an ionizing radioactive source. X-ray backscatter requires too much equipment to be practical on the battlefield. [3] NQR is not yet practical for battlefield detection of TNT. In fact, one of the goals of the present research is to move NQR-TNT detection from the laboratory to the battlefield.

In NQR explosive detection, the response of certain 14N nuclei present in the crystalline explosive material is probed. The 14N nuclei possess a nuclear quadrupole moment which in the presence of an electric field gradient produces an energy level splitting which may be excited by radio-frequency magnetic fields. Pulsing on the sample with a radio signal of the appropriate frequency produces a transient NQR response, which may then be detected. Since the resonant frequency is dependent upon both the quadrupole moment of the 14N nucleus and the nature of the local electric field gradients, it is very compound specific.

2. THE TECHNICAL PROBLEM

The practical problems with NQR stem from the fact that it is a weak physical effect. The random noise inherent in the electronics of the detection equipment can completely mask the NQR signature unless great care, or excessively long data collecting intervals, are used. A related, but different, issue is that non-random artifacts inherent in the electronics of the detection equipment must be removed, without removing the desired signature. Furthermore, although no other compounds are known to have NQR signatures that could be confused with RDX or TNT, there are effects that can produce false alarms. These effects are attributed to sand and/or plated metal objects in the field of view of the instrument.

From the perspective of signal processing, the problem is as follows. Given a signal that is known to contain random noise and unknown but non-random experimental artifacts, and given the possible simultaneous occurrence of obscuring signals originating from sand and plated metal objects, is the NQR signature present in the signal or not? The solution requires that the physical effects and their distinguishing features be reasonably well characterized, that the experimental procedure does not discard the distinguishing features, and that a proper feature extraction method be used.

A number of different advanced signal processing methods are currently being investigated to solve this problem. This paper is primarily concerned with the insights that have been revealed thus far by multiresolution methods. (Note: multiresolution analysis is a general set of principles out of which discrete wavelets arise as a special case. This research has used both discrete wavelet methods, and other multiresolution techniques.) Two particular insights have been gleaned from the research thus far. First, a signal's evolution in time is a distinguishing feature. Second, multiresolution methods might lead to an innovative hardware approach to NQR-TNT detection.

3. EVOLUTION OF THE SIGNAL IN TIME IS IMPORTANT

Commercial NQR explosive detectors have been available for several years. The reason that they are not more widely used than they are at present is that they are vulnerable to false alarms due to environmental effects. Although these effects are not well characterized physically, the effect arising from sand is conventionally called piezo-electric ringing (PER), and the effect arising from plated metal objects is conventionally called magneto-acoustic ringing (MAR). Since these effects occur at or very near the frequency of the NQR signature, they are nearly impossible to distinguish from NQR, when the decision is based only on frequency data.

Can a simple wavelet processing scheme provide real advantages over frequency-based methods on real-world NQR data? The short answer is yes, but under very carefully contrived conditions. In the following experiments, a simple wavelet method can distinguish whether the signal contains a comparatively strong PER or NQR signal, but the condition of simultaneous NQR, PER, and noise leads to unreliable results.

The experimental samples were collected from approximately 320 trials for which sand, but no explosive was in the field of view of the instrument, and from approximately 320 trials for which both sand and RDX were simultaneously in the field of view of the instrument. For the following comparisons, a two-dimensional vector was computed for each sample. Each sample is a 100-point list of complex time-domain data resulting from an experimental trial. The first element of the vector is the strength of the peak of the discrete Fourier transform (DFT) of the sample. The second element measures the evolution of the sample over time.

To compute the second element, the sample was analyzed with a five-level Haar wavelet transform. All output levels except the coarsest were assumed to be noise and zeroed out. The resulting "denoised" signal was inverted to the time domain, and the log of the absolute value of each datum was computed. A straight line was fitted to the resulting data. The slope of the line indicates the decay rate of the sample.

Classification of the vectors was based on a linear decision surface. The two-dimensional decision space has a decay-axis and, normal to that, a frequency-axis. When the decision surface has a rotation of 0 degrees compared to the decay-axis, the decision is based strictly on the frequency information. When the decision surface has a rotation of 90 degrees compared to the decay-axis, then the decision is based strictly on the decay information. For other rotations, the decision is based on contributions from both effects.

The decision as to whether or not the signature contains RDX depends on which side of the decision surface the vector appears. For repeated trials, the two classes "RDX" and "no RDX" are not linearly separable. The two classes of samples each occupy a region of the decision space, but there is some overlap, and any linear decision surface will lead to some misclassifications. There is a tradeoff between false alarms (PER or noise being mistaken for the RDX signature when no RDX is present), and missed detections (no signature detected when the explosive is actually present).

The tradeoff can be controlled by the placement of the decision surface. To assure no false alarms. the decision surface can be placed such that all samples are on the "no RDX" side; this would eliminate false alarms, but would also guarantee that all genuine detections would be missed. At the other extreme, the decision surface can be placed such that all samples are on the " RDX" side; this would guarantee that no genuine detections would be missed, but would always give false alarms when no explosive is present. Other locations of the decision surface could lead to other tradeoff values in the error rates.

A receiver operator characteristic (ROC) curve is a graphic representation of the tradeoff in error rates as some parameter in a decision process is varied. If the slope of a linear decision surface is fixed, then the offset is a single parameter that can be varied to give different tradeoffs in error rates, and to generate a ROC curve. Comparative ROCs can reveal at a glance the relative merits of alternative detection schemes.

Figure 1 shows ROC curves for an experiment using approximately 320 trials in which RDX and sand are present and approximately 320 trials for which sand but no RDX is present. The dashed ROC is based on a decision surface with a rotation of 85 degrees compared to the decay-axis; decay rate dominates the decision, and frequency makes a tiny contribution. The solid ROC is based on a decision curve with a rotation of 0 degrees; decay rate is ignored, and the entire decision is based on frequency information. (Note: many rotations were tried, but 85 degrees gave the most favorable ROC for these data.)

Figure 2 shows ROC curves for an experiment using approximately 320 trials in which RDX, but no sand, is present and approximately 320 trails for which neither sand nor RDX is present. The dashed ROC is based on a decision curve with a rotation of 5 degrees compared to the decay-axis; frequency dominates the decision, and decay rate makes a tiny contribution. The solid ROC is based on a decision curve with a rotation of 0 degrees; decay rate is ignored, and the entire decision is based on frequency information. (Note: many rotations were tried, but 5 degrees gave the most favorable ROC for these data.)

It is important to keep in mind what these results reveal, and what they do not. They reveal that knowledge of the evolution of the signal over time should improve the reliability of the decision, as opposed to using strength in the frequency domain as the only feature. These results do not provide a simple means for overcoming the problems of ringing effects produced by sand. The most favorable ROC when no ringing effect was present used a decision curve rotation of 5 degrees (decision principally based on frequency), and the ROC for PER data used a decision curve rotation of 85 degrees (decision principally based on decay rate). Knowing how much to rotate the decision curve was based on *a priori* knowledge of what contaminants were in the instrument's field of view when the data were collected . This *a priori* knowledge would not be available in field deployment.

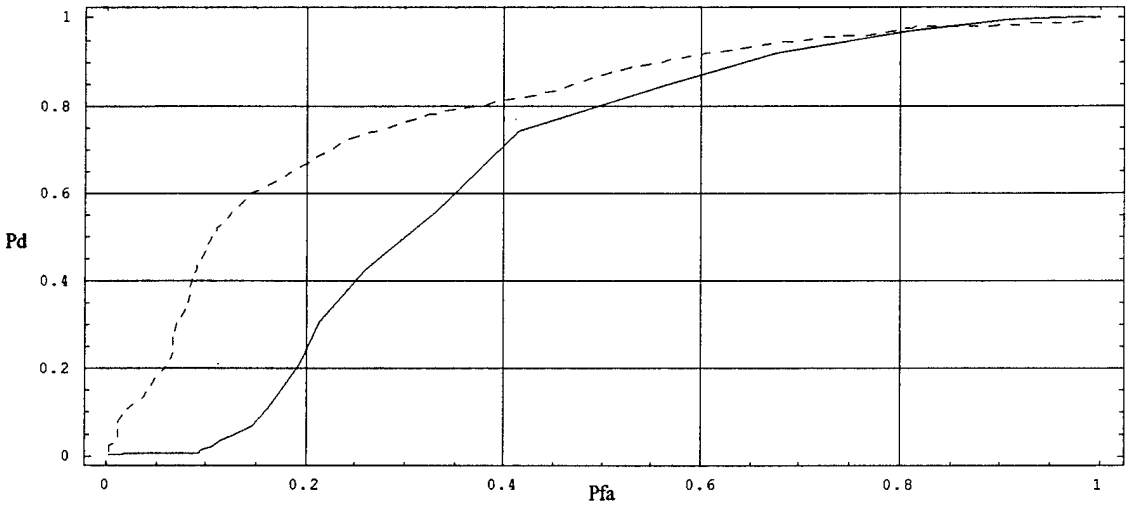

Figure 1. ROC for RDX Contaminated by PER (Decision is Dominated by Decay Rate)

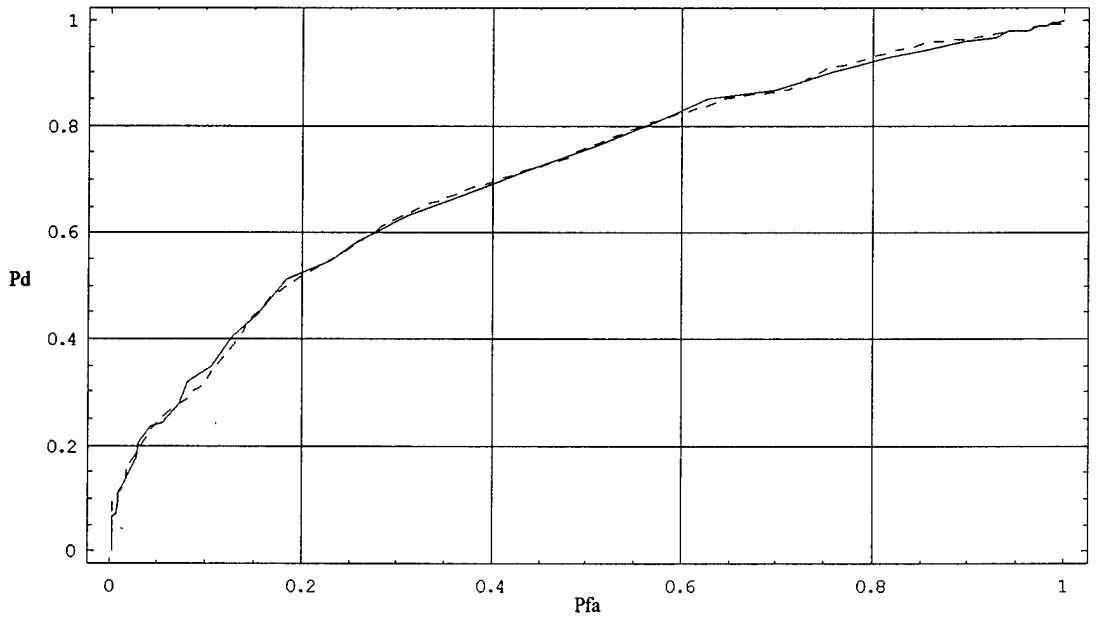

Figure 2. ROC for RDX Contaminated by Instrument Artifacts (Decision is Dominated by Frequency)

However, these results do bear out a crucial point. What distinguishes the obscuring signal from a genuine NQR response is the time-domain response. To reliably determine whether or not the NQR signal is present when a "ring" signal is present, a method of detecting time domain features *must* be used. The crude method described above indicates that it should be practical to do so.

The fact that time information matters indicates that the signatures of NQR and the more common interferents are not mathematically stationary. It is not valid to characterize the signature of NQR or its interferents as "narrowband." From the experimental data, it is clear that a reasonable approximate description of the NQR response is a damped sinusoid. The dominant frequency is temperature dependent, but known to be in the vicinity of 3410 kHz for RDX. The damping rates in the observed data covered the range from 750 to 1250 μs. Thus, a sinusoid with a frequency of 3410 kHz and a damping rate of 1 millisecond is a representative NQR signature.

The magnitude of the Fourier spectrum of this representative signal is shown in the solid line in Figure 3. The information in the signal that distinguishes it from a pure sinusoid is in sidebands 20 to 60 dB down from the dominant component, and covers a bandwidth of about 100 kHz. Thus, to preserve the waveshape of the signal (as a function of time) and to correctly extract the features that characterize it as an NQR signature, it is necessary to use hardware that preserves a comparatively wide bandwidth and dynamic range.

Although the obscuring signals are not well understood, experience suggests that they are broader band than the NQR signature. For an obscuring signal consisting of a damped sinusoid at 3410 kHz with a decay rate of 100 μs, the spectrum is shown in the dashed line. Again, a bandwidth of at least 100 kHz and a substantial dynamic range are needed to capture the distinguishing features.

Suppose a trial contains the sum of the desired signal an obscurant. Suppose the desired signal is at 3410 kHz, with a damping rate of 1 ms. Suppose the obscurant is at exactly the same frequency and phase, and its only distinction is that it has a damping rate of 100 μs. Further suppose that the whole signal is contaminated by strong Gaussian noise (signal-to-noise ratio of less than 0 dB). Such a signal would completely confuse both Fourier and wavelet methods. However, it should be relatively straightforward for sophisticated methods such as Bayesian analysis to accurately estimate the parameters of both the desired signal and the interferent, provided the information actually describing signals is preserved by the instrument used to collect the data.

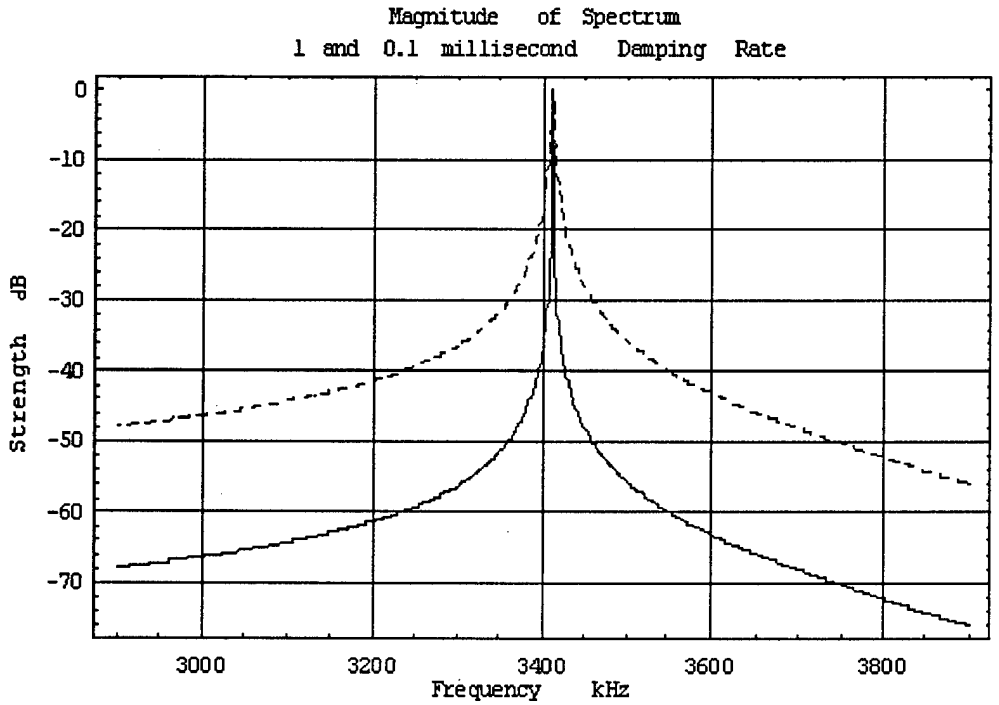

Figure 3. Magnitude Fourier Spectrum of Typical RDX-NQR Signature

4. HOW MUCH IMPROVEMENT MIGHT TIME-DOMAIN INFORMATION PROVIDE?

While a crude wavelet analysis will not provide the decision engine for this problem, it does indicate that the time-evolution contains valuable information. However, exploiting that information carries a cost. The physical hardware must operate at relatively large bandwidth and dynamic range, and a real-time parameter estimation scheme would probably require a dedicated DSP chip.

Are the benefits of preserving and exploiting the time-domain information really worth the trouble? Indeed they are. An improvement of even one or two dB could be significant. The following wavelet-based simulation indicates that improvement in performance which can be obtained by exploiting the time-evolution is similar to the improvement that might be obtained by increasing the signal-to-noise ratio by five dB.

Actually, it is more convenient to consider the converse problem, to estimate the effect of the loss of distinguishing information resulting from the assumption that the NQR signature is stationary. For example, suppose that a noisy data set contain a signal taken from one of two possible classes of damped sinusoids. One class has a carrier frequency of 50 kHz and a decay rate taken from a uniform distribution from 0.2 to 0.4 milliseconds (call this the fast class). The other class has a carrier frequency of 50 kHz and a decay rate taken from a uniform distribution from 0.6 to 0.8 milliseconds (call this the slow class). All members of both classes are damped cosines with a value of 1 at time = 0. Added to the signal is Gaussian noise at some selectable root mean square (rms) value.

The signal processing problem in this example is to guess whether the signal in the noisy data resulting from a given trial is a member of the fast class or the slow class. If the process completely separates the two classes for many repeated trials at a given noise level, (manifested as an ideally rectangular ROC) we say that the distinguishing features are strongly present. If the process at a given noise level is no more reliable than a random coin toss, (manifested as a straight line ROC) we say that the distinguishing features are completely lost. At some range of noise levels in between, the ROC will fall between these extreme shapes and we would say that the distinguishing features are present in the data but degraded by the noise.

This suggests a strategy for comparing two processes. The process whose ROC better approximates a rectangle is the superior process. To quantify how much better it is, consider the following. Suppose that for no noise or low noise the superior process completely separates the two classes, but the inferior process does not. By gradually increasing the noise level by the same amount for both processes, a threshold noise level will eventually be found where the superior process just barely fails to perfectly separate the two classes (producing a not quite perfectly rectangular ROC). At the threshold noise level, the ROC of the inferior process will be noticeably degraded from the ROC taken in the absence of noise.

We can exploit the fact that further increasing the noise level degrades the ROCs for both processes. Eventually, we will find a noise level (let's call it the degradation noise level), for which the ROC of the superior process approximately overlays the ROC for the inferior process operating at the threshold noise level. We can argue that the degradation in performance at the threshold noise level due to going from the superior to the inferior process is the same as the degradation in performance of the superior process due to going from the threshold noise level to the degradation noise level.

In this simulation, the two processes that are compared are as follows. The wideband process takes the output of the dominant scale of a 3-level wavelet packet transform, and has an effective bandwidth of 125 kHz. For the wideband process, the input data stream is discretized as it would be for 16-bit integer representation (84 dB usable dynamic range). The narrowband process takes the output of the dominant scale of a 10-level wavelet packet transform, and has an effective bandwidth of 1 kHz. For the narrowband process, the input data stream is discretized as it would be for 12-bit integer logic (60 dB usable dynamic range). In both processes, the underlying filter is a Daubechies 8-coefficient wavelet filter. In both processes, the base 10 log of the absolute value of each output point is computed, and the resulting list is fitted to a straight line. The distinguishing feature is the slope of the fitted straight line; the slope is correlated with the decay rate.

One might wonder, if the simulation produces such good results, why we don't use the "wavelet packet-log-slope trick" in the actual mine detector, and be done with it. Although the "wavelet packet-log-slope trick" is excellent for distinguishing whether a fast or slow decaying signal is present in noisy data, it becomes totally confused when a fast and slow decaying signal are present simultaneously, as would frequently occur in mine detection. The reason for using the "wavelet packet-log-slope trick" is that it is much simpler than Bayesian parameter estimation for determining the degradation of the features that distinguish the slow class from the fast class as bandwidth and dynamic range go from wide to narrow.

The first data set was generated with no noise added to the input signal. 100 trials were run for each combination of slow and fast signals into narrowband and wideband processes. For each trial the decay rate is chosen at random from a uniform distribution; the fast and slow distributions have already been described. For these data, the broadband process completely separates the slow decaying class from the fast decaying class. The narrowband process almost completely separates the two classes; only one point from slow decaying class overlaps into the fast decaying class. There is no need to plot the ROCs. One is rectangular, and the other is rectangular except at one point.

The ROCs for the data set including noise at the threshold level is shown in Figure 4. For each combination of signal class and bandwidth 150 trials were run. For the wideband process, an rms noise level of 0.424 leads to a very slight overlap of features in the slow class with features of the fast class. This is seen as a very small knee in the leftmost ROC curve. At an rms noise level of 0.424 the narrowband process is clearly inferior to the wideband process at distinguishing the slow class from the fast class. This is evident in the rightmost ROC curve.

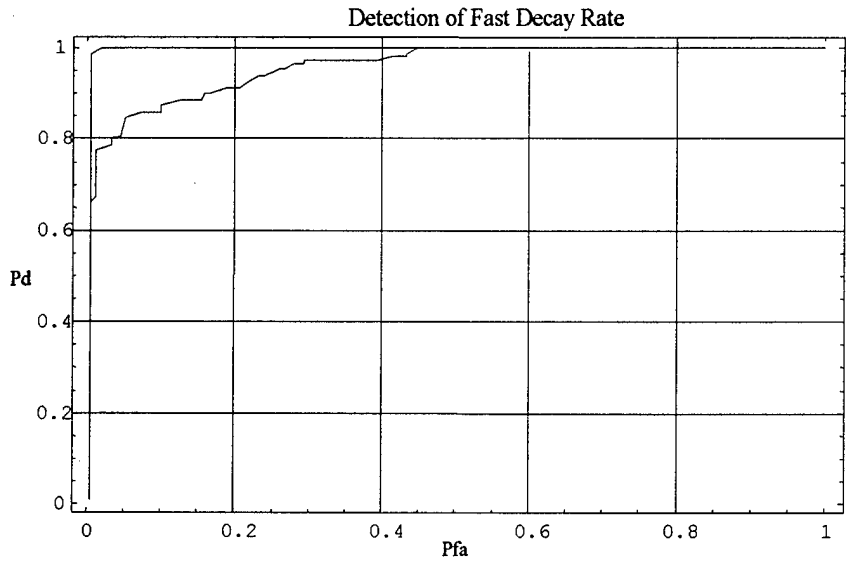

Figure 4. ROCs for Noise at Threshold Level

(Note: The interpretation of the ROCs is as follows. P_d is the probability of correctly detecting a member of the fast class when a member of the fast class is actually contained in the noisy data. P_{fa} is the probability of mistakenly identifying a trial as a member of the fast class when a member of the slow class is actually contained in the noisy data.)

The data set including noise at the degradation level is shown next. For each combination of signal class and bandwidth 150 trials were run. For the wideband process, an rms noise level of 0.754 leads to the leftmost ROC curve in Figure 5. For the narrowband process, an rms noise level of 0.754 leads to the rightmost ROC curve in Figure 5.

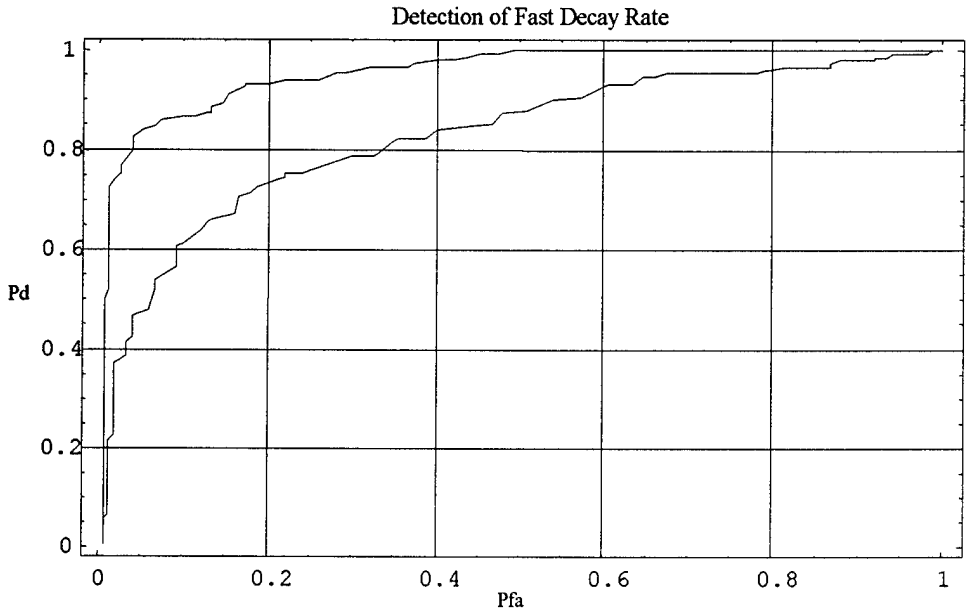

Figure 5. ROCs for Noise at Degradation Level

If the ROCs for the threshold and degradation noise levels for the wideband and narrowband processes are overlaid, it can be seen that increasing the rms level of the noise by 5 dB, from 0.424 (the threshold level) to 0.754 (the degradation level) produces the same degradation in performance as discarding the time evolution information. The ROC for the wideband wide dynamic range data at the higher noise level is practically identical to the ROC of the narrowband low dynamic range data at the lower noise level.

5. A POSSIBLE STRATEGY FOR TNT

It is well known that TNT-NQR signatures are much more difficult to detect than RDX-NQR signatures. Nevertheless, the majority of the existing landmines use TNT, and for practical NQR mine detection, the technology must be extended to TNT. To do so requires that the "front end" process impose as few new artifacts on the signal as possible.

Several insights from multiresolution analysis show how this might be done. Sampling rate change in the digital domain corresponds to heterodyne conversion in the analog domain. From the "noble identities" of multiresolution theory we can predict the effect of a change of sampling rate on a filter response, and thus contrive a filter that avoids the "image overlap" problem associated with downsampling. [4] From the foundations of multiresolution filter theory it is recalled that the binomial filter can be implemented without any multiply operations. [5] Thus, the binomial filter is particularly well suited to be implemented on a dedicated chip for real-time processing at radio frequencies (RF).

Direct digital conversion of the RF signal should impose fewer artifacts on the data than conventional spectroscopic methods. The signal of interest is in the vicinity of 890 kHz, and it can be conveniently digitized with a 2 MHz, 16-bit commercial analog to digital converter (ADC) chip, followed by a 71-stage binomial filter (easily realizable on a single field programmable gate array (FPGA) chip), subsampled by a factor of 10, and followed by a parallel interface to the computer. This would result in an effective data rate of 200-k samples per second that could be managed by a sophisticated feature extraction process on a dedicated DSP chip.

Can a digital filter be found that can be undersampled by a factor of 10 with no overlap of images in the passband? The following analysis uses a sampling rate of 2,094,118 samples/second. The resulting Nyquist limit is 1047.059 kHz. The digital frequency corresponding to 890 kHz is 2.67 radians. The 71-stage binomial filter whose peak is closest to 2.71 has a transfer function of $(1-1/z)^{67}(1+1/z)^{4}$. Digital implementation is shown in Figure 6, and the insertion loss is plotted in Figure 7.

BINOMIAL DIGITAL FILTER

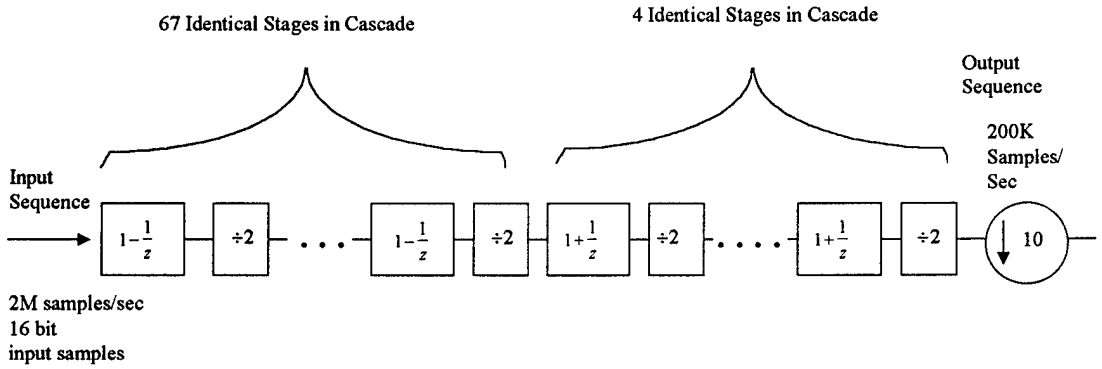

Figure 6. Implementation Binomial Anti-alias Filter

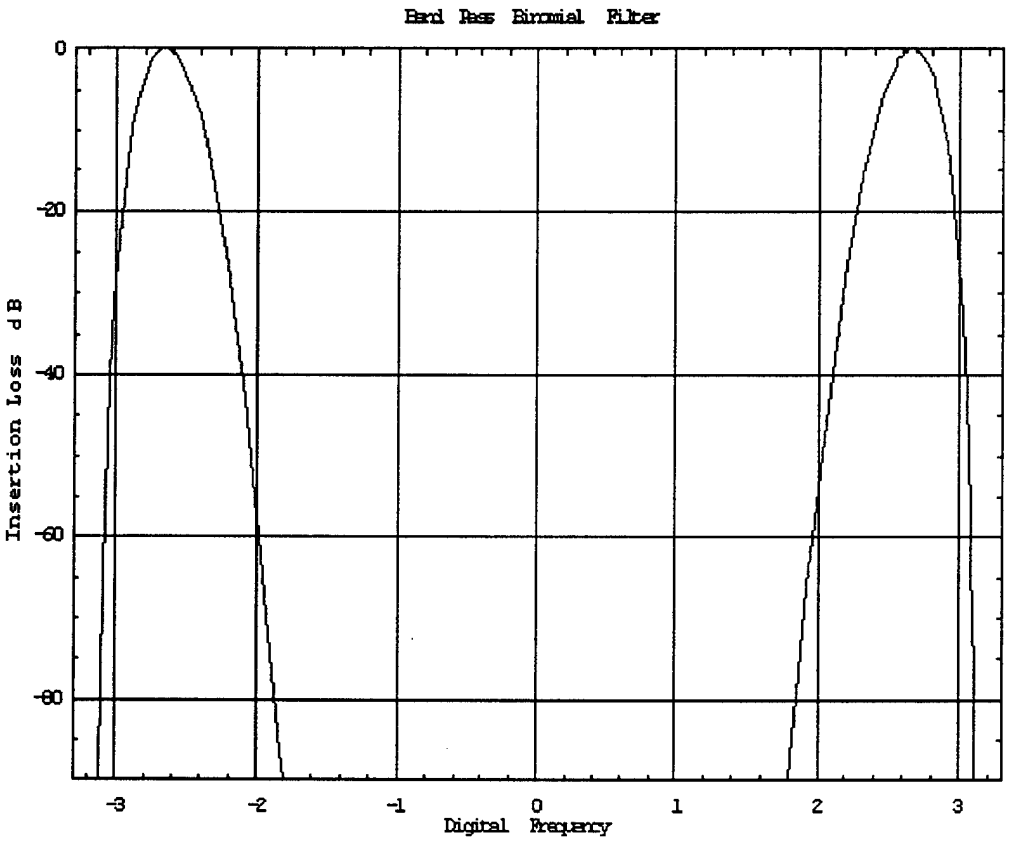

Figure 7. Frequency Response of Binomial Anti-Alias Filter

The geometric center frequency is 2.65 radians or 884 kHz. The half power bandwidth is 0.277 radians or 92.4 kHz. The filter response drops into the quantizing noise of 16-bit logic (84 dB down from peak) at a bandwidth of 1.277 radians or 425.7 kHz. The filter is linear phase, meaning that it does not change the shape of signals in the passband.

Assuming that the filter output is subsampled by a factor of ten, the following is the most straightforward way to interpret what happens. The process envisioned in two steps. The first step is to retain every tenth sample of the filter output and zero-out all the others. The second step is to discard the zeroed-out values. The first step will create nine images spaced at $\pi/5$ radian intervals around the unit circle in the z-plane.

Both the negative and positive frequency components of the original signal contribute to the images. However, in this case, the half-power passbands of the images do not overlap. The passband of the fourth image down from the positive original spans the digital frequencies of 0.004047 to 0.2815 radians (solid line in Figure 8). The passband of the third image down from the positive original spans the digital frequencies of 0.6324 to 0.9098 radians (dashed line in Figure 8). The passband of the fifth image up from the negative original spans the digital frequencies of 0.3469 to 0.6243 radians (dot-dashed line in Figure 8). The passband of the fourth image up from the negative original spans the digital frequencies of -0.004047 to -0.2815 radians (dotted line in Figure 8). The images interleave rather than overlap, and the binomial filter serves the role of a bandpass anti-aliasing filter.

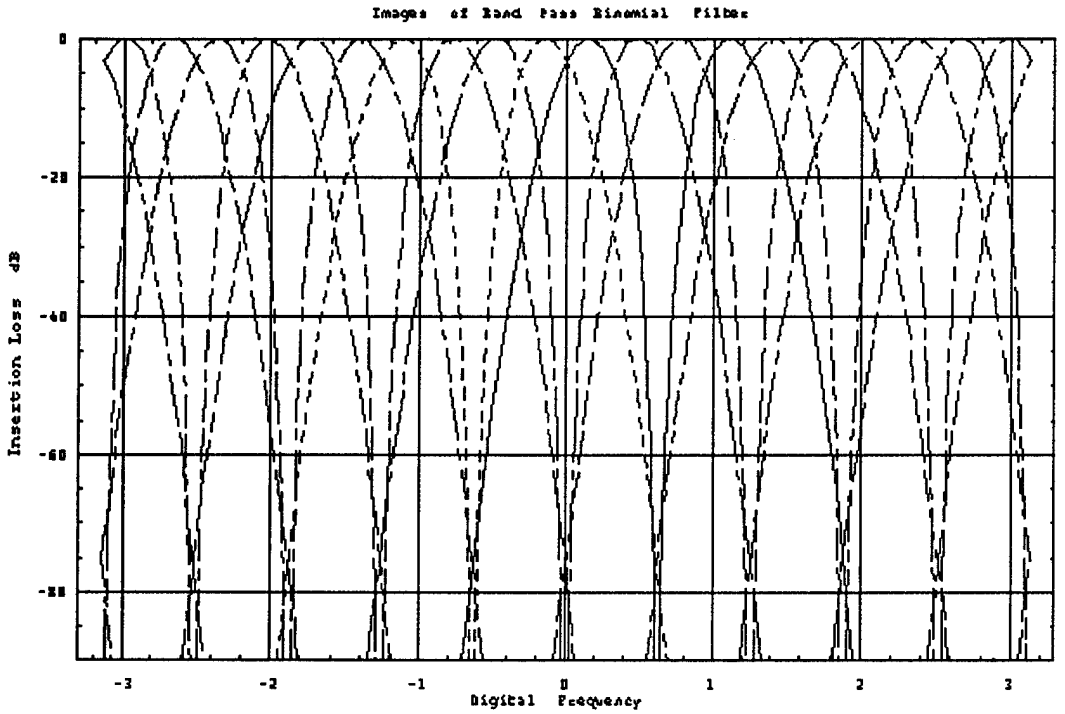

Figure 8. Interleaved Filter Images

The effect of discarding the zeroed-out samples is to magnify the passband by a factor of ten in the digital frequency dimension. Thus, the output of the 71-stage binomial filter decimated by ten has an effective half-power passband from 0.04047 to 2.815 radians, with a mirror image (reflected about the $\omega = 0$ axis) at the negative frequencies. These are shown in Figure 9.

If a properly bandlimited signal with a carrier frequency of 890 kHz and a sampling rate of 2,094,118 samples/second is applied to the 71-stage binomial filter, and the filter output is decimated by ten, then the following results ensue. The sampling rate is reduced to 209.412 kHz. The Nyquist frequency is 104.706 kHz. The 890 kHz carrier has been converted to 52.45 kHz. The band from 839 kHz to 931 kHz fills the digital spectrum. A digital conversion has been effected without resort to mixing in the analog domain or multiplication in the digital domain. The whole process is done with a pipeline of digital delay and add operations. The distinguishing features in the 100-kHz-wide band are preserved.

Perhaps most importantly, this multiresolution process introduces the bare minimum of artifacts that might obscure the weak TNT-NQR signal. Note that, at this point in the research, it has not been determined that this will ultimately be the strategy used for TNT processing. However, the preceding analysis shows that direct digital conversion by multiresolution methods is a promising alternative to conventional methods.

6. CONCLUSIONS AND FURTHER RESEARCH

This paper is an interim report of work still in progress. What has been learned through the use of multiresolution methods is the following. A signal's evolution in time is an important distinguishing feature, and exploiting that feature should lead to an improvement in performance that is similar to what would be obtained by improving the signal-to-noise ratio by 5 dB. Also, multiresolution methods may lead to a novel hardware approach to TNT detection.

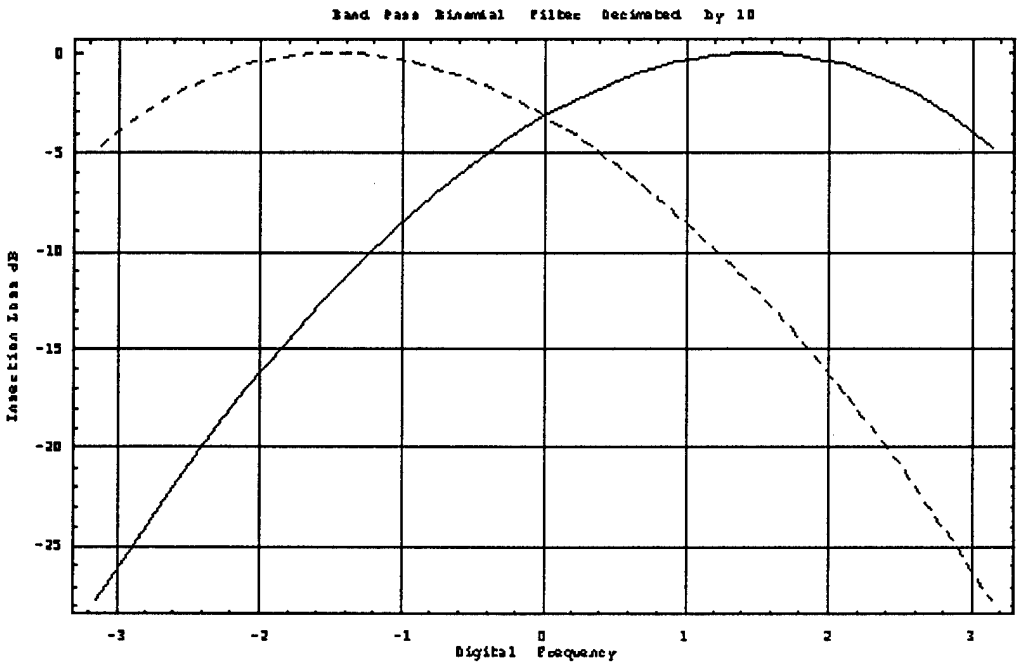

Figure 9. Response of Decimated Binomial Filter

7. ACKNOWLEDGMENTS

The authors wish to thank Dr. Regina Dugan for her support. This research was funded by the Defense Advanced Research Project Agency under Contract N00164-97-C-004.

8. REFERENCES

[1] Horowitz, P., *et. al.*, *New technological approaches to humanitarian demining*, Draft Report JSR-96-115, JASON, The MITRE Corporation, McLean VA, August 1996

[2] Dugan, R., "Detection of Landmines and Unexploded Ordinance by Exploitation of the Chemical Signature," presented at DARPA UXO Detection by the Chemical Signature Conference, Arlington VA, 3, June 1996.

[3] Shope, S., Lockwood, G., Bishop, L., Selph, M., Jojola, J., Wavrik, R., Turman, B., and Wehlburg, J., "Mobile, Scanning X-ray Source for Mine Detection Using Backscattered X-rays," in *Detection and Remediation Technologies for Mines and Minelike Targets II*, Abinash C. Dubey and Robert L. Barnard, Editors, Proceedings of SPIE Vol. 3079, pp. 400-407, 1997.

[4] Akansu, A.N., and Haddad, R.A., *Multiresolution Signal Decomposition*, Academic Press, San Diego, pp. 103-110, 1992.

[5] Akansu, A.N., and Haddad, R.A., *Multiresolution Signal Decomposition*, Academic Press, San Diego, pp. 47-53, 1992.

Report Number (14) ORNL/CP--96389
CONF-980412--

Publ. Date (11) 199801
Sponsor Code (18) DARPA, XF
JC Category (19) UC-000, DOE/ER

DOE