Third Annual Technical Report
Project A-3273

DEVELOPMENT OF AN EM-BASED LIFEFORM DETECTOR

By

Joseph Seals
Steven M. Sharpe
Scott Crowgey

Research Contract No. N00014-82-C-0390

Prepared for

NAVAL MEDICAL RESEARCH AND DEVELOPMENT COMMAND
National Naval Medical Center
Bethesda, Maryland 20014

March 1985

GEORGIA INSTITUTE OF TECHNOLOGY
A Unit of the University System of Georgia
Atlanta, Georgia 30332

GTRI
GEORGIA TECH RESEARCH INSTITUTE

DEVELOPMENT OF AN EM-BASED LIFEFORM DETECTOR

Third Annual Technical Report

Project A-3273

Prepared for

NAVAL MEDICAL RESEARCH AND DEVELOPMENT COMMAND
National Naval Medical Center
Bethesda, Maryland 20014

By

BIOMEDICAL RESEARCH DIVISION

DEVELOPMENT OF AN EM-BASED LIFEFORM DETECTOR

Third Annual Technical Report

Project A-3273

March 1985

Research Contract No. N00014-82-C-0390

By

Joseph Seals, Steven M. Sharpe and
Scott Crowgey

Prepared for

NAVAL MEDICAL RESEARCH AND DEVELOPMENT COMMAND
National Naval Medical Center
Bethesda, Maryland 20014

By

BIOMEDICAL RESEARCH DIVISION
Electronics and Computer Systems Laboratory
Georgia Tech Research Institute
Georgia Institute of Technology
Atlanta, Georgia 30332

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF FIGURES
Continued

LIST OF FIGURES
Continued

LIST OF FIGURES
Continued

LIST OF TABLES

FOREWORD

This annual report summarizes research efforts performed on this program from 1 May 1984 to 31 January 1985 by personnel of the Biomedical Research Division in the Electronics and Computer Systems Laboratory of the Georgia Tech Research Institute (formerly the Engineering Experiment Station) at the Georgia Institute of Technology, Atlanta, Georgia. The program is sponsored by the Naval Medical Research and Development Command contracting through the Office of Naval Research. Within the Navy, the program is identified as Contract No. N00014-82-C-0390 and it is designated by Georgia Tech as Project A-3273. Dr. Elliot Postow of the Naval Research and Development Command is serving as Program Monitor. Mr. Joseph Seals and Dr. Steven M. Sharpe of the Biomedical Research Division are serving as Project Director and Assistant Project Director, respectively. Staff members from the Biomedical Research Division that are assigned to this program include Dr. Scott R. Crowgey, Mr. Robert E. Voeks, and Mr. B. M. Jenkins. The help of Ms. Melinda Harp in preparing the figures in the report is gratefully acknowledged.

Respectfully submitted,

Joseph Seals

Joseph Seals,
Project Director

APPROVED:

J. C. Toler

J. C. Toler, Chief
Biomedical Research Division

SECTION I
INTRODUCTION

Georgia Tech is currently conducting a successful program to develop an electromagnetically-based lifeform detector capable of detecting vital signs-related information in battlefield casualties from extremely long ranges (10-100 meters). The original program goal was to develop a capability to detect respiratory activity in battlefield casualties from ranges up to 30 meters. However, because of the success of early program efforts, and the Navy's wishes to also detect cardiac activity, the program scope was expanded to include detection of respiratory and cardiac activity and the desired detection range was extended to 100 meters.

The lifeform detector (LFD) being developed on this program measures and analyzes scattered electromagnetic fields to detect respiratory and cardiac motions associated with the casualty being evaluated. Because antenna-based techniques are being used to perform the required scattered-field measurements, true remote-sensing is achieved. That is, a casualty being evaluated does not have to be wearing or carrying a biomedical transducer or other type of auxillary device. This fact greatly enhances the attractiveness and potential usefulness of the LFD.

A number of applications can be envisioned for a device which can function as an extremely sensitive, non-contact indicator of personnel-related motion. The instrument being developed on this program is primarily intended to be used for long-range detection of casualty status. However, it is potentially useful for purposes such as intruder detection and location of hidden personnel, as well as a variety of triage and diagnostic applications on the battlefield, as the following examples suggest.

- Rapid evaluation and prioritizing of casualties when they are dispersed over a broad battlefield area,
- Preliminary long-range evaluation of casualties isolated by distance, chemical agents, weapons fire, dangerous terrain, mine fields, or other threats,
- Evaluation of casualties wearing protective clothing or equipment that interferes with the operation of conventional medical instrumentation (e.g., chemical warfare suits), and

- Evaluation of casualties that have suffered burns or other trauma requiring that patient contact or movement be minimized.

Development of a LFD with the required range capability presents a significant challenge. The electromagnetic fields scattered by a casualty's body are extremely weak (especially when the casualty is lying on the ground) and measurements can easily be contaminated by the effects of numerous types of noise. In addition, since it is anticipated the LFD will not be operated by highly-trained personnel, the Navy has asked that the LFD include signal processing capabilities that will essentially make target detection an automated process. To meet this challenge, an intensive effort has been made by personnel at Georgia Tech to thoroughly analyze factors potentially capable of limiting the performance of the LFD, and to develop methods for effectively overcoming identified range-limiting problems.

To meet the required performance goals, features planned for the LFD include: (1) the use of a millimeter-wave operating frequency (35 GHz) to enhance motion sensitivity and to achieve high antenna directivity using compactly-sized antennas, (2) a sophisticated frequency modulation scheme that permits implementation of quadrature detection channels and improved rejection of both external noise (through range-gating) and internal system noise (through frequency translation), (3) a homodyne demodulation scheme that makes it possible to achieve excellent receiver performance using a compact system configuration, and (4) a microprocessor-based signal processing system that employs a conveniently implemented FFT-processor. Based on preliminary results, the combined performance improvements resulting from use of these features will eventually make it possible to achieve the long-range detection capabilities being sought on this research program.

The current version of the LFD includes a 35-GHz operating frequency and a sophisticated frequency modulation capability. The performance of this prototype system has been tested under a number of conditions. Test results have indicated that under light to moderate clutter conditions, the LFD is capable of detecting respiratory and cardiac activity from ranges up to 50 meters. The test results have also made it possible to determine system improvements needed to insure reliable detection performance and to extend the detection range beyond current limits.

The main performance-limiting problem with the current LFD has been found to be external noise (usually referred to as clutter) due to reflection

from objects other than the target of interest. The key to achieving greater discrimination against clutter is development of a refined range-gating capability that will enable narrower range-cell sizes to be achieved and make it possible to reduce the effects of range-sidelobes. An important step in improving the current range-gating capability was the recent purchase of a new 35 GHz Gunn oscillator. The new oscillator has a significantly wider tuning response than the Gunn oscillator in the current system which will enable narrower range-cell sizes. Analysis has also been performed to determine techniques for reducing the effects of range-sidelobes. It has been found that appropriately weighting the demodulated return signals from the LFD can significantly reduce range-sidelobe levels. Plans for implementing and testing an improved range-gating system are now underway.

Related to these studies, investigations are planned of methods for using narrow range cells without requiring precise range knowledge. Possibilities include an auto-ranging approach (a system that would actually measure the target range once the LFD was aimed at the desired target), a second approach that would search through adjacent range cells once the approximate target range was specified, and a third approach in which all range cells between specified minimum and maximum ranges are examined for the presence of target information. Each of these approaches is intriguing, but further feasibility studies are required.

Another key to improving the performance of the LFD is the use of improved signal processing techniques. At ranges less than 50 meters, it is often possible to observe distinct respiratory and cardiac waveforms in the raw signals outputted by the LFD. However, at extended ranges where the strength of target return signals will be reduced, and under high-clutter conditions, it is likely the raw data from the LFD will have to be carefully processed to extract useful casualty information.

Currently, spectral studies are being performed on samples of signals and noise outputted by the LFD, as well as on simultaneously recorded EKG's, to establish an information-base for designing a suitable signal processing system. The results of these studies have made it possible to obtain useful information about an important class of clutter and have pointed out the importance of obtaining information on other types of clutter. In addition, comparison of the spectra of subject returns and baseline returns (i.e., subject removed) have shown that periodic components associated with respiratory and cardiac motion can often be distinguished from the

approximately white, baseline spectrum. This suggests that the use of signal processing techniques based on spectral analysis and this possibility is being investigated.

The spectral studies that have been performed on this program indicate detection performance can be improved through signal processing. However, significantly improving detection performance through signal processing usually requires the use of long observation times. Ideally, the level of required signal processing and the corresponding observation period could be adjusted based on the difficulty of the detection problem. Otherwise, there may be times when excessive signal processing is used and times when inadequate signal procesing is used. A possible approach for a signal processing technique with this capability is a detector called the Sequential Observer [1]. This detector would be able to rapidly make a yes/no decision on casualty status when conditions were favorable (e.g., signal is present and very strong; signal is not present and noise is very low). However, when detection was difficult (e.g., the target information is weak or clutter is high), this detector would indicate that a yes/no decision was not possible then switch itself to a optional detection mode in which a longer observation period and a higher level of signal processing would be used. This detector appears well-suited for the LFD and investigations are needed to determine if this approach can be implemented in the microprocessor-based signal processing system planned for the LFD.

As indicated in the preceding description, efforts during the initial phases of this program have been extremely promising. However, it has also been noted that a number of system refinements are still needed to achieve the program objective of being able to reliably detect respiratory and cardiac activity from extremely long ranges. The next required step is continuation of the current program efforts. Key points of future efforts should include (1) improving the range-gating capability of the current LFD (narrower range cells, lower range sidelobes, and perhaps, some type of automatic ranging), (2) reducing system size, weight, and power drain to improve system portability, (3) additional spectral analysis of signals and noise to establish a meaningful information base for designing a suitable signal processing system, (4) implementation of a microprocessor-based signal processing system, and (5) extensive field testing. The basis for these needed improvements as well as specific details on the status of current program efforts, are presented in the technical report in the next section of this report.

SECTION II
REPORT OF TECHNICAL PROGRESS

A. Summary

The prototype lifeform detector (LFD) being developed on this program uses radiated electromagnetic fields to detect body motions associated with casualty respiratory and cardiac activity. Under the best conditions, development of such a system would be difficult because a major portion of the reflected energy detected by the LFD would be due to objects such as the ground, grass, bushes, and trees, and not to the target of interest. In many instances, the effects of these extraneous reflections is to strongly mask the desired casualty-related information. In addition to having to deal with the problem of detecting extremely weak signals in the presence of strong noise, the LFD must satisfy several other requirements that complicate its development.

One of these requirements is that the size, weight, and power consumption of the LFD must permit the system to operate in a portable mode. This requirement poses limitations on the type and amount of equipment that can be incorporated into the LFD, and in some cases, may make it necessary to use less than optimum techniques. A second requirement is that the LFD must provide an output that can be easily interpreted by personnel with minimum technical training. This increases the needed signal processing capabilities and effectively results in the requirement that the LFD perform automatic detection.

The effort to develop the LFD is being divided into two parts, (1) development of a motion sensitive, RF system and (2) development of a computer-based signal processing system. The general design being used for the LFD is shown in Figure 1. The RF system is comprised of an RF source, an RF network, an antenna, and a receiver. The purpose of the RF system is to maximize the target-related information and minimize the noise in the signal outputted to the signal processing system. The RF system performs this task by simultaneously achieving high motion sensitivity and minimizing the effects of various types of noise and interference.

The purpose of the signal processing system is to optimize the ability to detect the presence or absence of target-related information in the signal

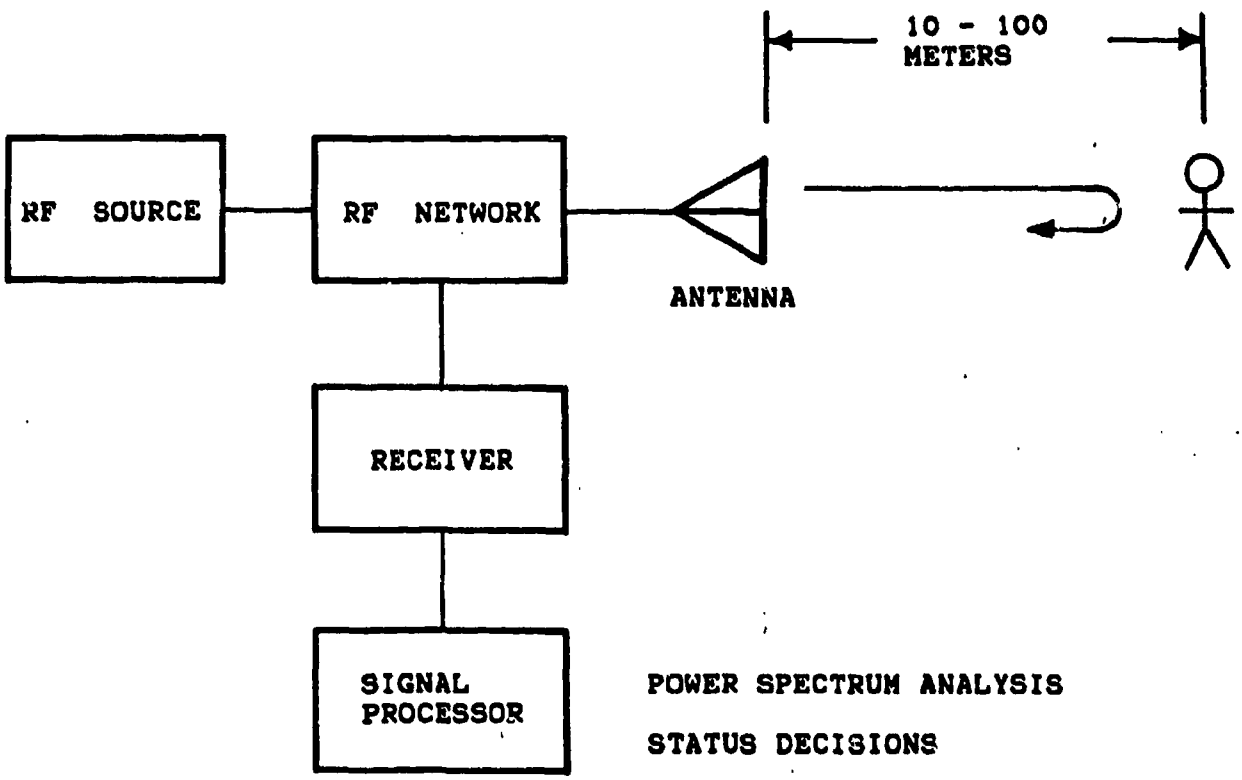

Figure 1. Conceptual block diagram of EM-based lifeform detector (LFD).

received from the RF system. A detection technique based on power spectrum analysis is currently being tested. When properly utilized, power spectrum analysis appears capable of providing a high level of detection performance [2-4]. This technique also has the added advantage that it can be implemented with a relatively simple microprocessor system interfaced with a special purpose FFT-processor.

Significant advances have been made in development of both the RF and signal processing systems. Efforts involving the RF system have been more concrete in that a prototype unit that includes a number of range-enhancement features has been implemented and tested. The performance of the prototype unit has been satisfactory for the ranges (50 meters or less) and conditions (low clutter) that have been tested. However, additional improvements will be needed to fully meet program goals. Improvements planned for the RF system include increasing the RF system's range-gating capabilities and reducing its size and weight.

Efforts involving the signal processing section have been more investigative in nature but useful experimental results have been obtained using in-house equipment. These results indicate the use of appropriate signal processing can improve the ability to observe target-related information in the output of the LFD's RF section. However, long observation periods are required if significant improvements in detection performance are to be achieved. A portable version of the in-house equipment being used to perform the power spectrum analysis must now be developed.

Further details on development of the LFD's RF and signal processing systems are provided in the remainder of this section. With the exception of the discussion on range-gating, much of the information pertaining to the RF system is summarized since it has been discussed in the previous two technical reports [5,6]. The signal processing discussion is presented in detail since it has received only cursory treatment in the previous reports. Both background information and experimental results obtained in recent months are presented in the signal processing discussion.

B. RF System Development

The initial step in implementing a suitable RF system was development of a model that satisfactorially described the problem of using EM-based techniques to detect cardiac and respiratory related events. This model was

needed to identify mechanisms on which operation of the LFD could be based and to identify factors capable of degrading performance of the LFD. The resulting model showed the LFD should be optimized to detect subtle body motions associated with cardiac and respiratory activity. It was found that the ability to detect these motions could be enhanced by using a high operating frequency. The primary problems predicted by this model were (1) the detrimental impact of reflections from objects other than the target of interest (generally referred to as clutter) and (2) sinusoidal variations in motion sensitivity that result in an effect being described as "range-deadspots".

1. Motion Detection

The phenomenon, as well as associated problems, of EM-based motion detection can be understood by analyzing the case in which a simple continuous wave (CW) system is used to detect a moving, reflective target where w represents angular frequency. This situation is shown in Figure 2. The transmitter produces a signal denoted x, of the form Acos(wt), which is initially transmitted at the target. If noise free conditions are assumed, a signal denoted y, of the form Bcos w(t-T) , is subsequently returned from the target, and detected in the LFD. The amplitude parameter, B, includes the effects of components in the LFD (e.g., couplers, antenna, circulator), attenuation during propagation to and from the target, and reflections at the target surface. The parameter T represents the time-delay due to propagation to and from the target, which is at a range denoted as R.

Detection is achieved by multiplying the signals x and y together in a nonlinear device such as a square-law detector or mixer. The exact output of the nonlinear device will depend on the device's design. However, the output will generally contain a cross-product term denoted z, of the form, ABcos(wT). The delay term T, is related to the target range and propagtion velocity v, as follows:

$$R = vT/2 \qquad (1)$$

The factor of two in this equation is due to the round-trip propagation to and from the target.

If the target is a casualty lying on the ground, body motions associated

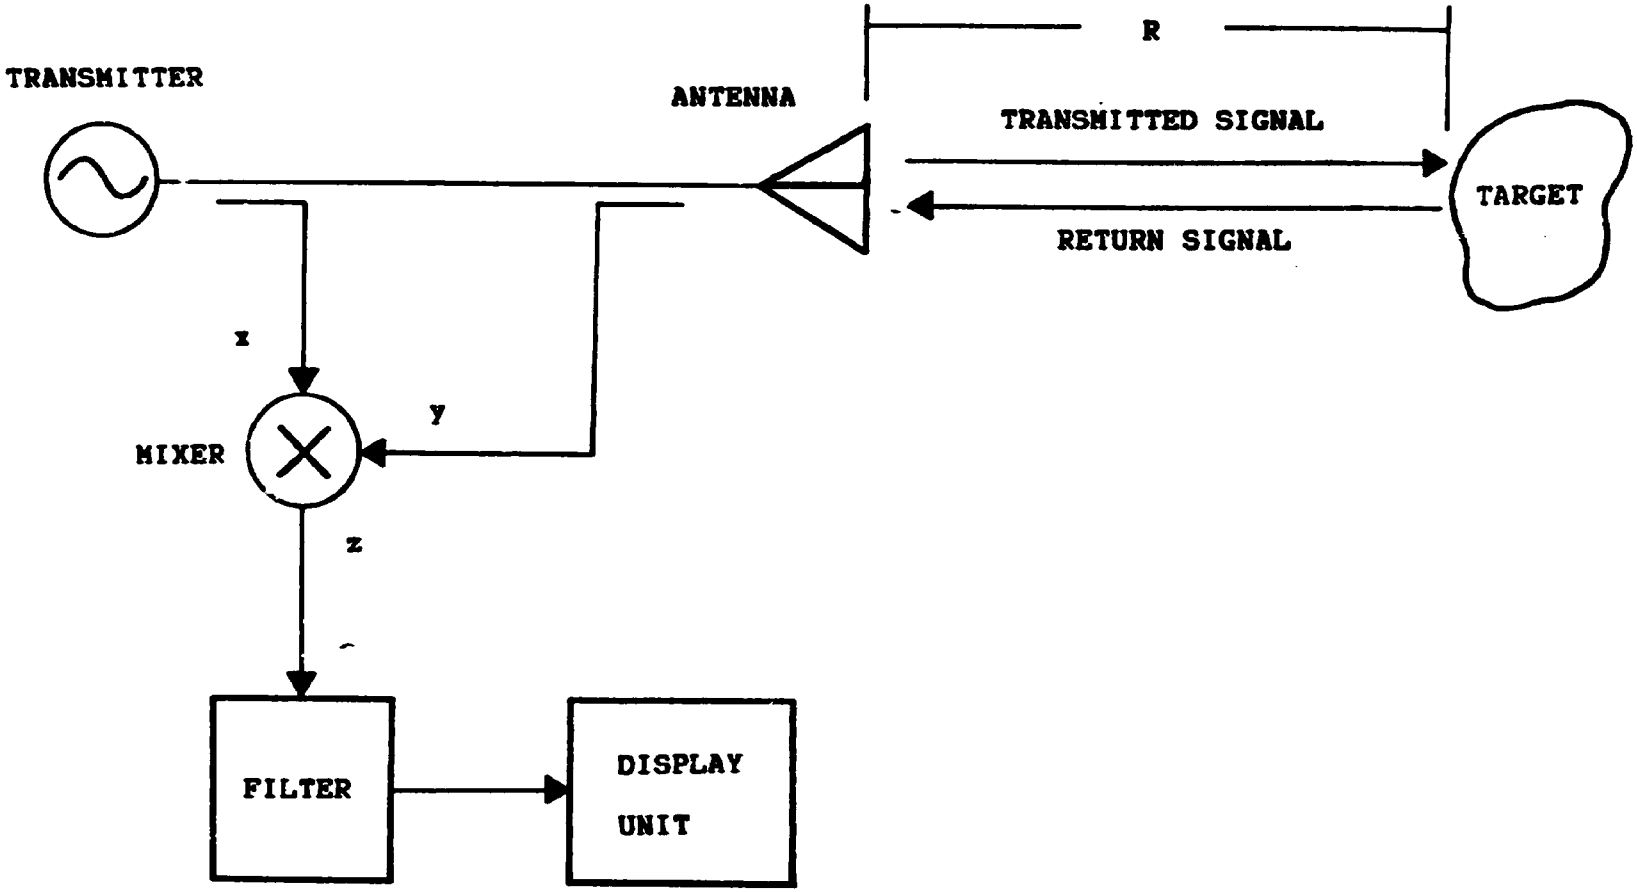

Figure 2. Example of simple continuous-wave (CW) motion detection system.

with the respiratory and cardiac activity result in minute range variations. The target range is therefore a time-varying quantity which can be expressed as:

$$R(t) = R + r(t), \tag{2}$$

where, R is now a constant representing the nominal distance to the target and r(t) represents small changes in range due to respiratory and cardiac related motion. If the results of equations 1 and 2 are substituted into the expression for z, the demodulation product becomes,

$$z = AB\cos 4\pi r(t)/\lambda + 4\pi R/\lambda \tag{3}$$

If it is assumed that the product AB remains relatively constant (a reasonable assumption since the target's minute motion should have little effect on its reflection cross section), several interesting conclusions can be drawn from equation 3. One important result is that the motion information, r(t), is contained in the argument of z. Therefore, phase sensitive techniques are required to extract the information of interest. A second result is that the amount of phase change produced by the target motion is inversely proportional to the operating wavelength. Thus, shorter wavelengths (or equivalently, higher operating frequencies) result in greater motion sensitivity.

Equation 3 can be rewritten in the following more compact form,

$$z = AB\cos(p(t) + P). \tag{4}$$

The parameter p(t) represents the motion-related information of interest and P is a phase offset term due to the nominal target range. Although phase detection is required for true extraction of p(t), under appropriate conditions, p(t) can be approximated by simple amplitude detection of z. The required conditions are that the offset term, P, equal an odd multiple of 90 degrees and that the excursions of p(t) be limited to less than $\pm$ 60 degrees. Under these conditions, z will be almost linearly proportional to p(t).

If the phase offset, P, is arbitrary (which would be the case for the LFD since the target range cannot be controlled), simple detection of z will

result in regions of poor sensitivity (or range-deadspots) at ranges where P is a multiple of 180 degrees. Re-examination of equation 3 shows range-deadspots will occur at range increments equal to one-fourth the operating wavelength. This phenomenom is due to the simple fact that the signals x and y are not truly coherent because of the phase offset, P. Although this effect is easy to understand, its significance should not be overlooked. In fact, any attempt to implement a lifeform detector without taking this effect into account should be viewed questionably.

The effects of range-deadspots was demonstrated by using an early version of the LFD to detect the motion of a metal-foil covered sphere that was moving sinusoidally with respect to the LFD. The peak-to-peak distance of the sphere's motion was set to produce a maximum phase change (i.e., p(t) in equation 4) of approximately 30 degrees. That is, the amount of target motion was small enough to permit p(t) to be approximated from detection of z, provided a suitable phase offset was attained.

The motion-related output of the LFD was viewed on a dual-trace oscilloscope along with the signal being used to control the sphere's motion. Initially, the target range (approximately 3 meters) was carefully adjusted so the phase offset, P, was an odd multiple of 90 degrees. Under this condition, it was expected the detected motion displayed on the oscilloscope would be similar in shape to the sinusoidal control voltage. Figure 3a shows that the expected effect did occur and that simple detection of z permitted extraction of the motion-related phase term, p(t).

The target range was then increased by one-eighth wavelength (approximately 2 millimeters at the test frequency of 35 GHz). This should have placed the sphere at one of the previously discussed range-deadspots and the ability to observe p(t) through detection of z should have been greatly diminished. The results obtained for this second set of conditions is shown in Figure 3b. As expected, the motion of the sphere appears to be substantially reduced even though the only difference in the two cases was a two millimeter difference in range. This result clearly demonstrates the range-deadspot effect predicted by the earlier analysis. As a further check, the target range was increased an additional one-eighth wavelength. This should have placed the sphere back at a range where it should have again been possible to observe the target's motion. As predicted, the motion of the sphere was again observable as shown in Figure 3c.

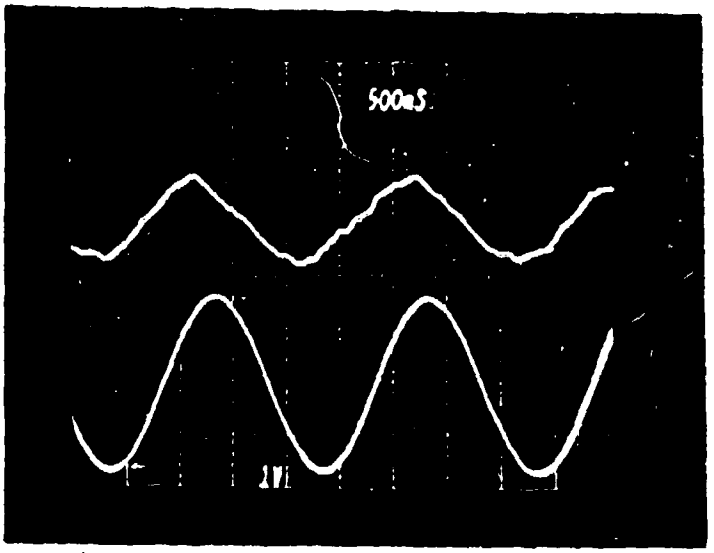

a. Sphere with small motion positioned for maximum phase sensitivity.

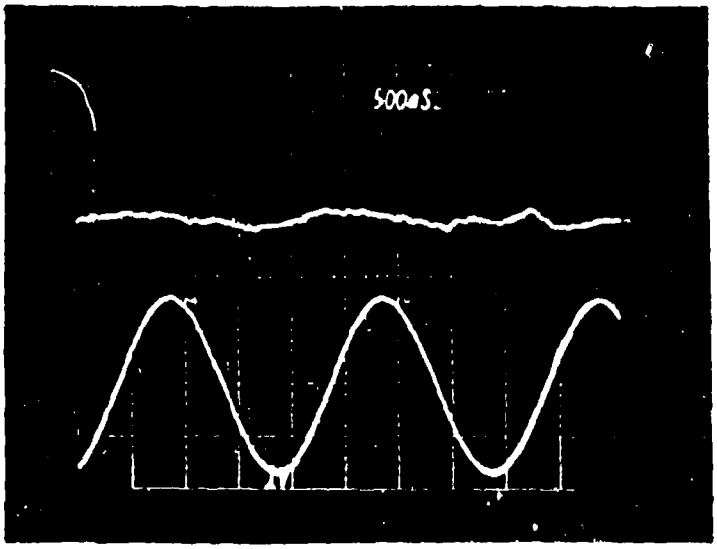

b. Sphere with small motion positioned for minimum phase sensitivity (range-deadspot).

Figure 3ab. Detected motion of metal foil-covered sphere to demonstrate effects of range-deadspots: (a) sphere positioned for maximum phase sensitivity, (b) sphere positioned for minimum phase sensitivity (range-deadspot).

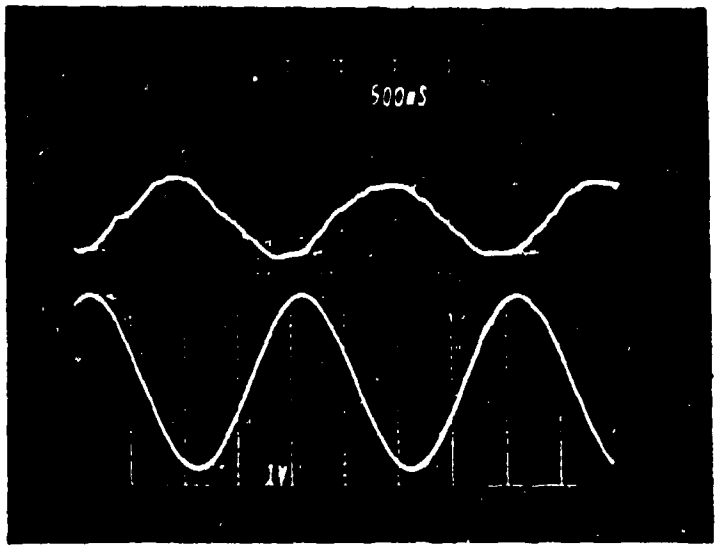

DETECTED MOTION
(50 μv / div)

CONTROL VOLTAGE
(1 v / div)

c. Sphere with small motion re-positioned for maximum phase sensitivity.

Figure 3c. Detected motion of metal foil-covered sphere to demonstrate effects of range-deadspots: (c) sphere repositioned for maximum phase sensitivity.

As mentioned previously, the problem of range-deadspots is due to the phase offset between the signals x and y used in the discussion of simple CW systems. A possible approach for dealing with this problem would be to slightly adjust the frequency of the transmitter to set the phase offset to a suitable value. This could be performed manually or through some type of automated scheme. However, a better approach, referred to as synchronous or quadrature detection, is being used in the LFD [7,8]. In this approach, the return signal (i.e., y) is initially split in half. One of these halves is multiplied with the transmitter reference signal, x. The remaining half is multiplied with a version of x that has been delayed in phase by 90 degrees. When this is done, the LFD will provide output signals of the form,

$$I = AB\cos(p(t) + P). \tag{5a}$$

$$Q = AB\sin(p(t) + P). \tag{5b}$$

The advantage of having quadrature channels is that if the target range results in poor sensitivity in one channel, the sensitivity of the remaining channel will remain high. Alternately, since the signals I and Q are in phase quadrature, they could be inputted to a processor which could actually compute the value of the motion-related quantity, p(t).

2. Noise Problems

Various types of noise can compromise the performance of an EM-based motion detection system. Noise sources that affect the performance of the LFD can be divided into two categories, internal system noise and external noise (or clutter). Internal system noise is generated by active and passive components within the LFD while external noise results from return signals detected from objects other than the target of interest. Noise, regardless of its source, is capable of masking casualty-related information contained in the detected return signal. This masking effect degrades system sensitivity and produces a higher minimum detectable signal. As shown by the following expression for the radar-range equation, a higher minimum detectable signal results in a poorer detection range [9].

$$R = \left[\frac{P A^2 \sigma}{4 \pi S \lambda^2} \right]^{1/4} \qquad (6)$$

where R = estimated maximum range, meters
P = transmitted power, watts
A = effective antenna aperture, meters^2
σ = target cross section, meters^2
s = minimum detectable signal, watts, and
λ = operating wavelength, meters.

The radar-range equation indicates that under ideal conditions, detection range is inversely proportional to the fourth root of the minimum detectable signal. Based on this assumption, a 7 dB improvement in the signal-to-noise ratio (SNR) will increase the detection range by 50 percent, a 12 dB improvement in SNR will double the detection range, and a 19 dB improvement in SNR will double the detection range, and a 19 dB improvement in SNR will triple the detection range. These figures should be noted for reference during subsequent discussions in this report.

It is convenient to first discuss internal system noise. There are three main types of internal system noise. Thermal noise results from random electron motion and is constant with frequency. Its impact depends on temperature and the system bandwidth, and the thermal noise floor generally represents the lowest level of practically achievable noise reduction. Shot noise results from the discrete nature of DC current flow and is thought to be constant to approximately 100 MHz. Its impact depends on the particular system design, and in most applications, precautions can be taken to limit shot noise to a negligible level.

The final, and most troublesome type of internal system noise in the LFD is flicker noise (sometimes referred to as semiconductor noise or 1/f noise). Flicker noise is due to various semiconductor effects and is typically inversely proportional to frequency (hence, the description, 1/f noise) [10-12]. In the LFD, primary sources of flicker noise are the transmitter [13] (noise sidebands in the transmitter's output that are translated to baseband during demodulation) and detector-related noise [14]. An example of the flicker noise measured for a mixer and oscillator similar to those used in the LFD is shown in Figure 4. This example demonstrates the manner in which flicker noise decays as a function of frequency.

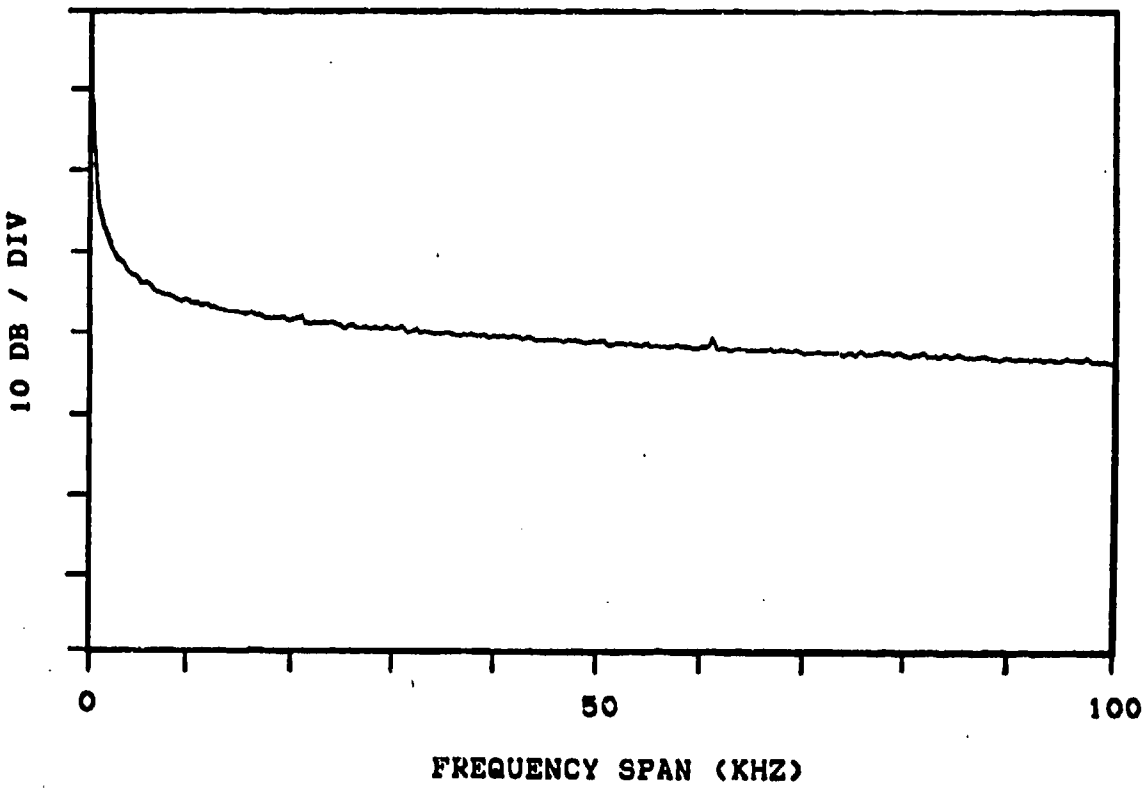

Figure 4. Measured noise spectrum of a single-ended mixer and a Gunn oscillator similar to units currently employed in LFD.

Flicker noise is a problem because of its large low-frequency content. In the simple CW system described in the previous part of this report, the respiratory and cardiac information in the received return signal was directly translated to baseband following demodulation. Since the strongest component of flicker nose exists at low-frequencies, the sensitivity and detection range of a simple CW system would be greatly compromised by flicker noise.

There are several techniques that can be used to reduce flicker noise effects. It is important to note that amplification of the casualty information following baseband translation is not one of these techniques since both the noise and the information would be amplified. An obvious approach to reducing flicker noise is to use low-noise components and a Gunn oscillator and silicon-Shottky diode mixer are being used in the LFD because of their inherent, low noise characteristics. However, even these low-noise devices exhibit higher levels of flicker noise than can be tolerated in the LFD.

Another approach for reducing flicker noise is to amplify the received return signal prior to demodulation. This permits the information of interest to be amplified independently of the troublesome flicker noise. Problems with this approach include the power consumption of the amplifier and the possibility of amplifier saturation from strong return signals due to antenna mismatch, ground reflections, or local oscillator/transmitter leakage. The saturation problem could be treated with feed-through nulling techniques but the cost would be increased system complexity. Because of these potential problems, this approach was not considered for use in the LFD.

Flicker noise can also be reduced by using a balanced detection scheme. In balanced detection, the demodulation process is split in half with complimentary local oscillator signals being used for each half. When the two demodulated outputs are added, noise components due to the local oscillators will cancel because of their complimentary nature. When carefully designed, balanced schemes can be very effective in cancelling the effects of local oscillator-related noise. However, noise from leakage due to antenna mismatch or ground reflections will not be cancelled and nulling techniques may again be necessary. The single-ended mixer being used in the LFD currently appears adequate. However, balanced schemes are still under consideration and may be adopted in the final prototype.

A third technique for reducing flicker noise effects is currently being used in the LFD. This technique involves translating the received return signal to an intermediate frequency (IF) prior to detection. The IF must be high enough in frequency to permit the information to be amplified independently of the flicker noise produced during the demodulation process, but must also be low enough that signal amplification is convenient and economical. With adequate IF amplification and filtering, the information will be significantly stronger than any existing flicker noise when the final translation to baseband is made and good receiver sensitivity will be achieved. Nulling still may be required with this technique. However, instead of performing nulling at a high RF (e.g., 35 GHz in the current LFD), it would be performed at the significantly more convenient IF (125 kHz is the maximum IF in the current LFD).

Frequency translation can be achieved with heterodyne or homodyne techniques. Heterodying requires a transmitter and local oscillator having offset frequencies. Following demodulation, the information of interest is translated to an IF equal to the difference between the transmitter and local oscillator frequencies. Although widely used, heterodyne techniques can be cumbersome because of their greater power consumption and the required coherence between the transmitter and local oscillator (phaselocking or single-sideband techniques are normally used to achieve the required coherence). Therefore, a homodyne approach that is better suited for portable applications is being used in the LFD.

Homodyne techniques have the advantage of being able to employ a single oscillator to provide both a transmitter signal and a reference signal. The reference signal is analogous to the local oscillator signal used in heterodyne techniques and enables demodulation of the received return signal. The reference signal also enables true mixer operation to be used. This latter property enables conversion efficiencies comparable to that of high performance heterodyne systems. The required frequency translation is achieved by appropriately modulating the oscillator and performing demodulation in a manner that results in the information of interest being translated to an appropriate IF. An IF of 125 kHz is currently being used in the LFD. Figure 4 shows that an IF greater than approximately 10 kHz provides a high degree of flicker noise reduction with the inherently low-noise components being used in the LFD.

The preceding discussion indicates that internal system noise, particularly flicker noise, must be minimized to achieve good receiver performance (i.e., a low minimum detectable signal) and long detection ranges. It was also shown that the LFD includes features such as homodyne demodulation that insure good receiver performance. In fact, based on a consideration of internal system noise, the current LFD should be capable of detecting casualty motions from distances greater than 100 meters. However, experience has shown it is external noise, and not internal system noise, that ultimately limits the range of the LFD. Therefore, the LFD must also include provisions for reducing the effects of external noise.

External noise (or clutter) is that portion of the received return signal that is not related to the target of interest. Clutter can result from reflections from stationary objects (the ground, buildings) as well as from moving objects (vehicles, personnel, trees, bushes, grass, etc.). When clutter is weak, it distorts the amplitude and phase of the target information. Thus, the ability to extract diagnostically useful information is impaired. However, good detection is still possible. When clutter is strong, it is capable of masking the target information and detection becomes difficult. This latter condition describes the problem with using techniques such as the LFD to detect casualty motion from long-ranges.

Since the purpose of the LFD is detection of low-frequency motion associated with respiratory and cardiac activity (0.05 to 20 Hz appears to be the approximate information bandwidth), the effects of clutter from stationary or fast moving objects can be filtered, provided they don't saturate the receiver. Therefore, it is clutter from slowly moving objects such as foliage around the target casualty, and personnel motion in proximity to the LFD, that is the biggest problem. From these observations, it is clear that an effective method of minimizing the effects of clutter is to reduce the volume of space interrogated by the LFD. That is, use spatial discrimination to separate target and clutter returns. This can be achieved by using appropriately designed antennas and by range-gating.

To aid clutter reduction, the antenna used in the LFD should have a narrow beamwidth (i.e., high directivity) to minimize the area scanned by the antenna. In addition, narrowing the antenna beamwidth increases the strength of the return signal received from the target as was seen from the radar-range equation. The antenna beamwidth required to achieve satisfactory clutter

rejection is a function of the target range since the area scanned will become larger as the range is inceased (because of beam spreading), and of the specific clutter conditions (different requirements would exist in barren terrain than in high-foilage terrain). For illustrative purposes, it is useful to assume good clutter rejection requires the diameter of the scanned area (as defined by the antenna's 3-dB beamwidth) to be equal twice the height of a 6-foot (1.8 meter) tall man. The resultant range that could be achieved for this condition with antenna beamwidths of 2, 4, and 10 degrees, is shown in Figure 5. These results indicate that extremely narrow antenna beamwidths are required at extended ranges. The beamwidth of the antenna currently in the LFD is approximately 4 degrees. For the example case in Figure 5, a 4 degree beamwidth would be adequate for ranges in the vicinity of 100 meters. Narrower beamwidths would be required to achieve clutter reduction at longer ranges if the clutter is particularly strong.

There are two factors associated with achieving narrow antenna beamwidths that should be noted. One factor is that antenna beamwidth is proportional to the ratio between antenna aperture size and operating wavelength. That is, high operating frequencies should be used to achieve narrow beamwidths and minimum antenna size. Once the operating frequency is fixed, the aperture size must be increased if a narrower beamwidth is required. Thus, one of the costs that must be paid to achieve improved clutter rejection through narrower antenna beamwidths is larger antenna sizes.

A second antenna-related factor is that as the width of the main antenna beam is narrowed, there can be an increase in the sidelobe levels radiated by the antenna. Although the sidelobe-radiation may be well below that of the main beam, sidelobe-radiation can be troublesome because it can result in reflections from objects in close proximity to the LFD. Since reflections from close-in objects will not experience significant range loss, they are capable of masking the information of interest. An example of the effects of sidelobe-radiation is operator-related motion that shows up in the output of the LFD. Reduction of sidelobe-radiation can be achieved by controlling the aperture illumination of the antenna. One of the reasons a lens antenna is being used in the LFD is that its illumination pattern can be controlled in a manner that permits effective sidelobe reduction without significant sacrifice in antenna beamwidth [15].

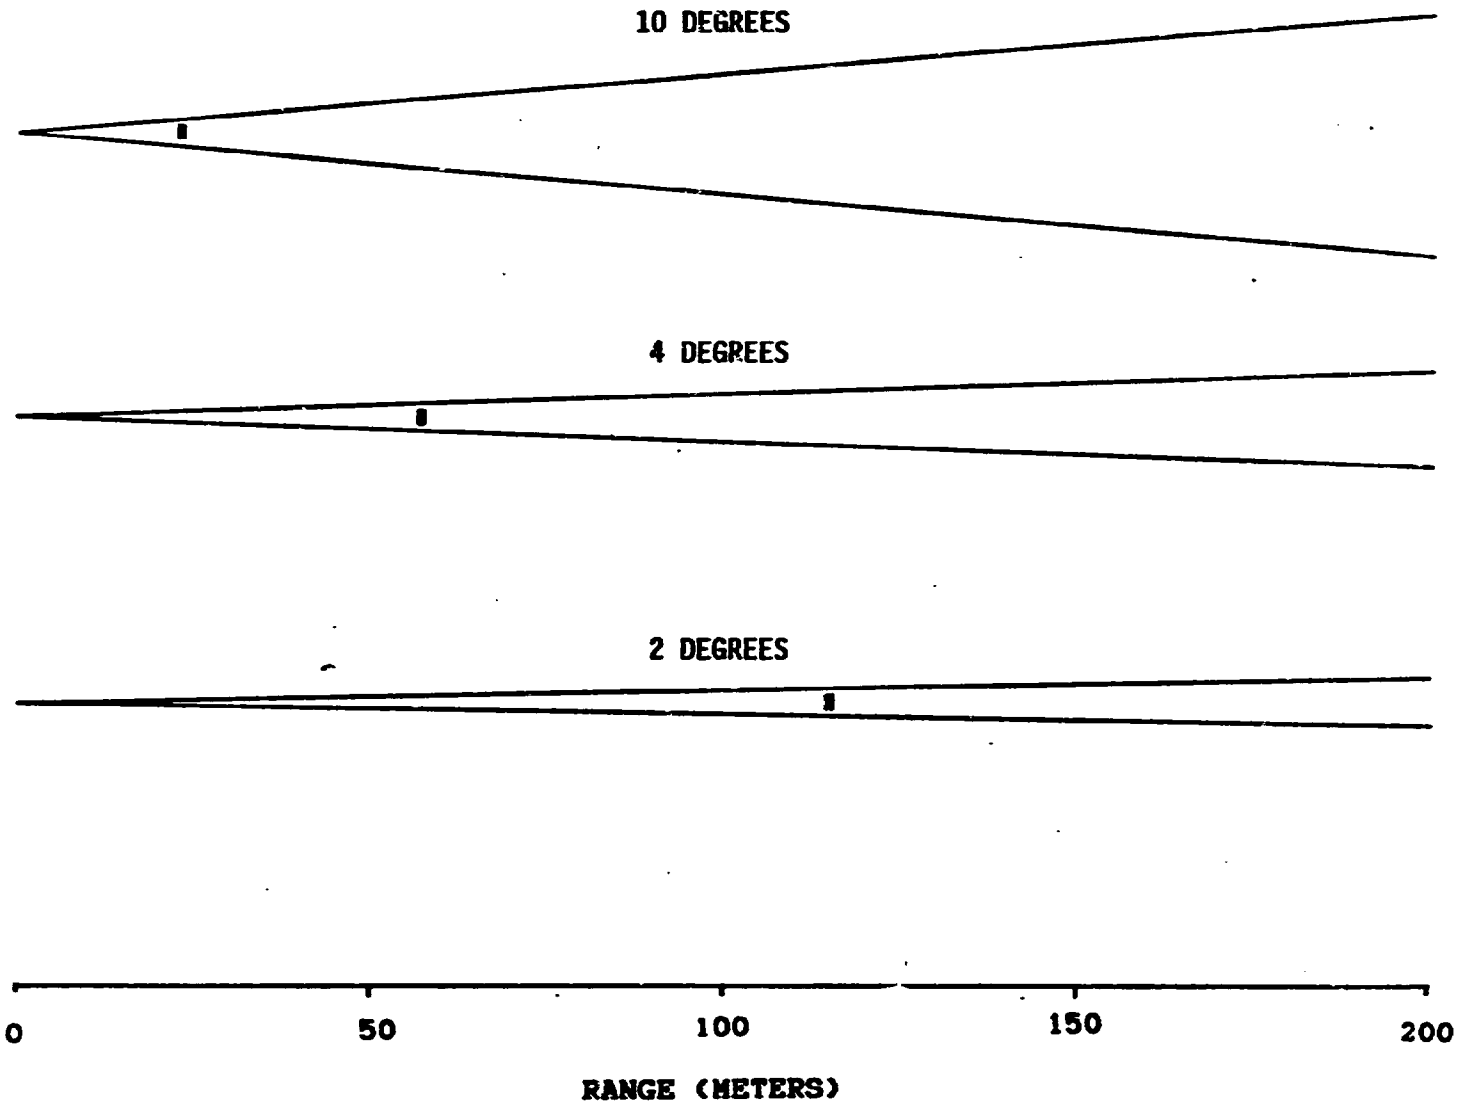

Figure 5. Relationship between antenna-beamwidth, target size (black rectangles denote a 1.8 meter-tall man), and target range, for 3-dB beamwidths of 2, 4, and 10 degrees.

3. Range-gating

In addition to using appropriately-designed antennas, the effects of clutter can be reduced by using range-gating techniques. Range-gating (or ranging) aids clutter rejection by making it possible to selectively examine return signals from specific ranges. The limited ranging capability of the earliest versions of the LFD was useful for eliminating the effects of operator motion and other close-in clutter but was not sophisticated enough to adequately reduce clutter in proximity to the target. Several improvements in the current version of the LFD permit a useful level of ranging to be achieved. However, problems still arise because of an inability to adequately negate the effects of large or extended clutter sources such as trees and grassy fields. Because of this inadequacy, and the fact that clutter is the predominate range-limiting problem of the LFD, a great deal of program effort has been expended on investigating and developing techniques for improving the ranging capabilities of the LFD.

Ranging is achieved by placing a timing mark on the signal transmitted at the target. The timing marks on received return signals can then be used as reference points for selecting or rejecting information from specific ranges. Frequency modulation is being used in the LFD. The choice between pulse and frequency modulated CW implementation depends on a number of technical tradeoffs including, peak to average power capability of the source, burnout susceptibility of the receiver and various noise mechanisms. Whether pulse or frequency modulation is used, the bandwidth of the transmitted signal detrermines the range selectivity (range-cell size) that can be achieved. In the case of the LFD, these tradeoffs favor frequency modulated CW over pulse modulated techniques. When pulse modulation techniques are used, range-cell size is proportional to the pulse width. Thus, narrow pulse widths are required to achieve narrow range cells. If frequency modulation is used, a wide frequency bandwidth is required to achieve narrow range cells. Examples of the pulse widths and equivalent frequency bandwidths required to achieve different range-cell sizes are shown in Table 1.

An ideal range discrimination function is shown in Figure 6a. In this example, only returns from within a selected range cell centered at the desired range (100 feet in this example) are accepted; all others are completely suppressed. As might be expected, the performance of a practical system falls short of the ideal case, although good performance can be

TABLE 1

BANDWIDTH REQUIREMENTS FOR RANGE-GATING

PULSE WIDTH (nanoseconds)	FREQUENCY BANDWIDTH (MHz)	RANGE CELL SIZE (Feet)
1	1000	0.49
5	200	2.5
10	100	4.9
25	40	12
50	20	25
100	10	49
200	5	98
1000	1	490

(a) Ideal case.

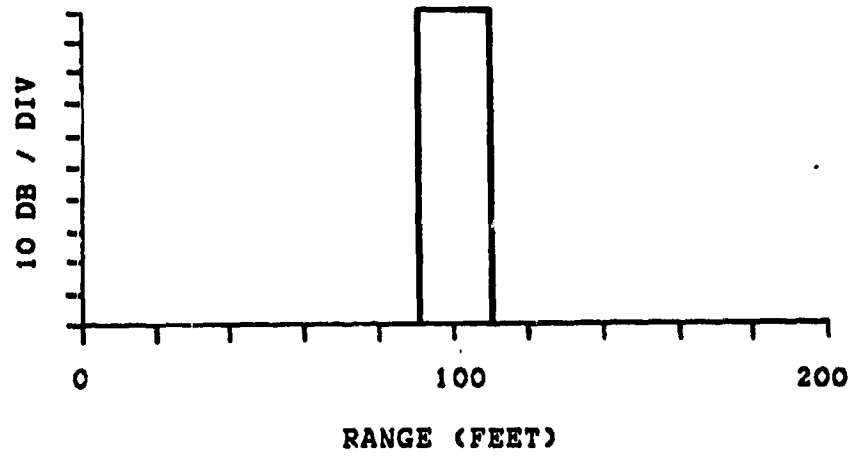

(b) Current (non-ideal) system.

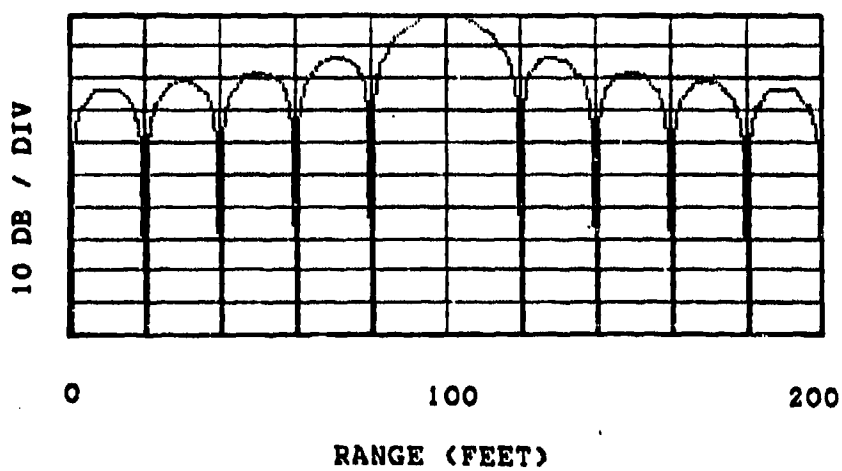

Figure 6ab. Graphs of range-discrimination functions for several different conditions: (a) ideal range-discrimination function, (b) non-ideal range-discrimination function of current LFD.

achieved with careful design.

The computed range-discrimination function of the current LFD is shown in Figure 6b. The target range and range-cell size in this example have been set for 100 feet and 20 feet, respectively (the range-cell size is defined as the distance between the 3-dB points of the range-discrimination function). The most serious departure from the ideal case is the incomplete suppression of return signals from ranges outside the main range cell as evident from the multiple secondary peaks in Figure 6b. These secondary peaks are referred to as range-sidelobes. Some suppression is clearly evident. For example, returns from 150 feet are suppressed by approximately 18 dB. However, this degree of suppression is not sufficient to adequately attenuate returns from large clutter sources such as a stand of trees.

With the frequency modulation-based ranging scheme being used in the LFD, the demodulation output of the LFD will contain sinusoidal bursts due to returns from the target being evaluated. The frequency of these bursts is a function of the target range which enables ranging to be performed. If the sinusoidal burst contains an integral number of cycles, a range-discrimination function similar to the ideal case shown in Figure 6a could be achieved. However, the number of cycles in the sinusoidal burst will vary as a function of target range. Thus, while complete suppression is achieved at some ranges as evident from the nulls in the range-discrimination function shown in Figure 6b, incomplete suppression is achieved for in-between ranges.

Range-sidelobes can be reduced by applying a weighting function 16-18 (or window) that tapers the ends of the sinusoidal bursts to decrease the contributions of any incomplete cycles within the outputted bursts. Several possible weighting functions have been evaluated during the third year. Results computed for one such weighting function is shown in Figure 6c. In this case, returns from 150 feet are now suppressed by 23 dB, which is an improvement over the result obtained with no weighting. Thus, better clutter discrimination could be achieved if weighting were used to reduce range-sidelobes.

An effect of using weighting functions is an increase in the width of the main range cell. The original width can be restored by increasing the amount of frequency deviation that is employed as shown in Figure 6d. In this example, returns from 150 feet are reduced to an inconsequential level. It may be noted that in addition to decreasing the level of the range sidelobes,

(c) Using weighting without increasing deviation.

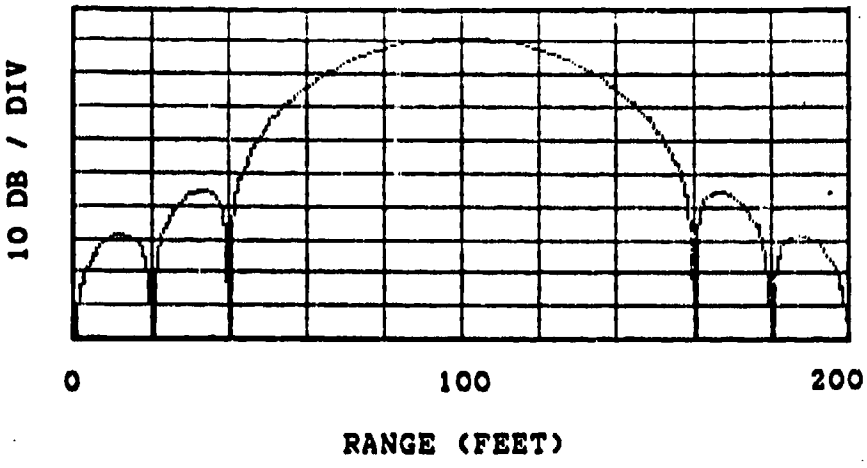

(d) Using weighting and increased frequency deviation.

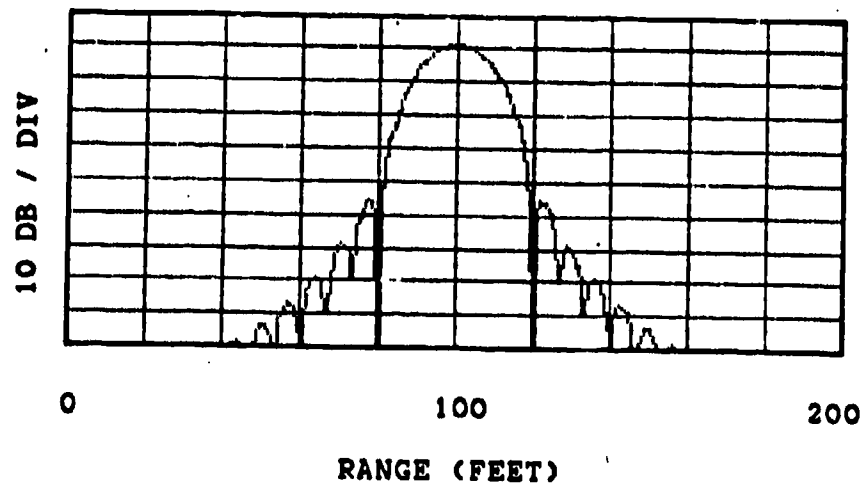

Figure 6cd. Graphs of range-discrimination functions for several different conditions: (c) computed range-discrimination funciton using weighting but no increase in modulation deviation, (d) computed range-discrimination function using weighting and increased modulation deviation.

the weighting function also reduces the peak amplitude of the range-discrimination function. However, since the noise is reduced by almost as much as the signal, there is a very small net loss. Implementation of this improved range-gating system appears practical for the LFD since it requires only a modest increase in system complexity. Currently, plans are being made to obtain the additional hardware needed to implement an improved range-gating capability.

4. Description of Current Prototype System

The preceding discussion indicated a number of performance features that should be incorporated into the lifeform detector to achieve good performance. These features included the use of (1) a high operating frequency to enhance motion sensitivity, (2) a directive, low-sidelobe antenna and range-gating to aid in clutter rejection, (3) synchronous (quadrature) detection to avoid range-deadspots that result in poor motion sensitivity, and (4) reduction of internal system noise, particularly flicker noise, to insure high receiver sensitivity. At least preliminary versions of each of these features has been incorporated into the current version of the LFD.

A block diagram of the current LFD is shown in Figure 7. The first point to note about this system is its high operating frequency of 35 GHz. The operating wavelength at this frequency is only 8.57 millimeters. Preliminary results indicates this is short enough to permit motions of less than 0.05 millimeters to be detected. In general, motions associated with respiratory and cardiac activity appear to range from a few millimeters to less than a millimeter, depending on the subject's condition and position. Greater motion sensitivity could have been obtained by using an even higher operating frequency. However, 35 GHz was selected to avoid the higher cost and poorer reliability of components currently available at higher frequencies.

The short wavelength at 35 GHz also makes it possible to achieve high antenna directivity using reasonably-sized antennas. Two different-sized antennas are currently available for use in the LFD. A six-inch diameter antenna that provides a 3-dB beamwidth of approximately four degrees has been used in the majority of studies performed to date. A nine-inch diameter antenna that provides a 3-dB beamwidth of approximately 2.5 degrees is also available and may be employed when the LFD is operated at longer ranges.

A scalar feedhorn-dielectric lens arrangement is being used for the

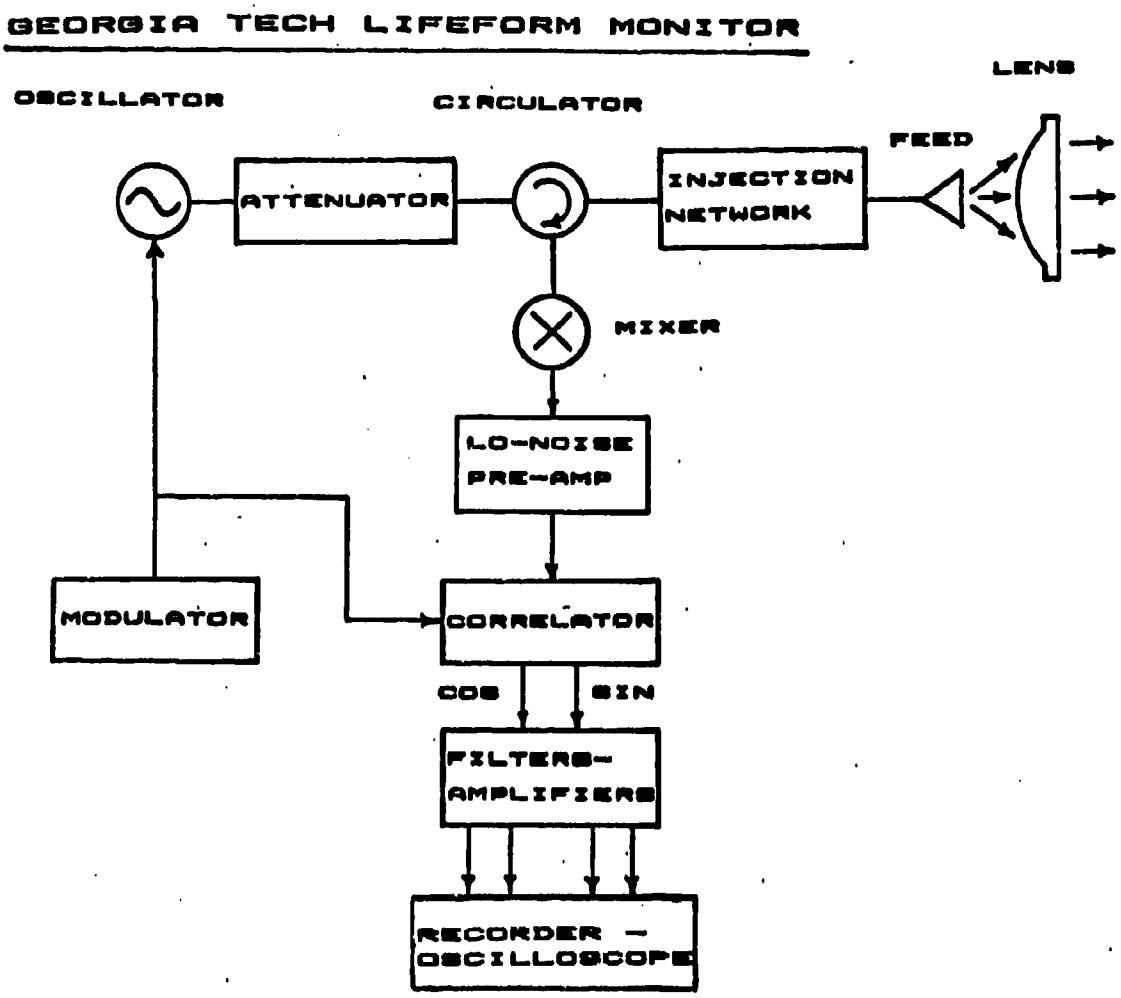

Figure 7. Block diagram of 35 GHz prototype EM-based LFD being developed at Georgia Tech.

antennas as shown in Figure 8. A single feedhorn can be used with either a six-inch or a nine-inch diameter lens. The feedhorn and dielectric lenses were all custom-fabricated for use on this program. The feedhorn has a relatively wide 3-dB beamwidth of approximately 20 degrees because of its small, one-inch aperture. However, the diverging beam from the feedhorn is collimated or "corrected" by refraction produced by the dielectric lens to produce the desired narrow beamwidth.

A scalar feedhorn design was specified instead of a more commonly available horn antenna because the inherently-low sidelobe radiation of the scalar horn helps minimize the overall sidelobe-radiation exhibited by the feedhorn-dielectric lens combination. In addition, the size of the feedhorn was selected so the lens illumination pattern produced when the feedhorn is positioned at the lens focal point is approximately 20 dB lower at the outside edge of the lens than at the lens center. By tapering the illumination in this manner, effects of the discontinuity at the lens edge can be minimized and overall sidelobe radiation can be reduced. However, a slightly wider beamwidth must be accepted for employing this type of illumination pattern.

The lens antenna is currently the largest single component in the LFD and it does not appear that it will be practical to use a lens larger than nine inches in diameter. If future testing indicates narrower antenna beamwidths are needed because of clutter-related problems, it may be desirable to consider using a higher operating frequency. For example, at a frequency of 94 GHz, a six-inch dielectric lens and suitable feed horn could be used to achieve a 3-dB beamwidth of approximately one degree. However, as noted earlier, problems could be encountered with component cost and reliability if a higher operating frequency were used. An alternative to changing the antenna beamwidth would be to improve clutter rejection by using a more refined range-gating scheme.

The LFD shown in Figure 7 is based on a homodyne design. As explained in a preceding part of this report, homodyne techniques make it possible to achieve efficient, low-noise detection of the received return signals without requiring the complexity and higher power drain of more commonly used heterodyne approaches. A single transmitter is required when a homodyne receiver is employed. The transmitter being used in the LFD is a 35 GHz voltage tuned oscillator (VTO). The VTO can be modulated by applying a suitable modulating waveform to the VTO's varactor tuning point. The output

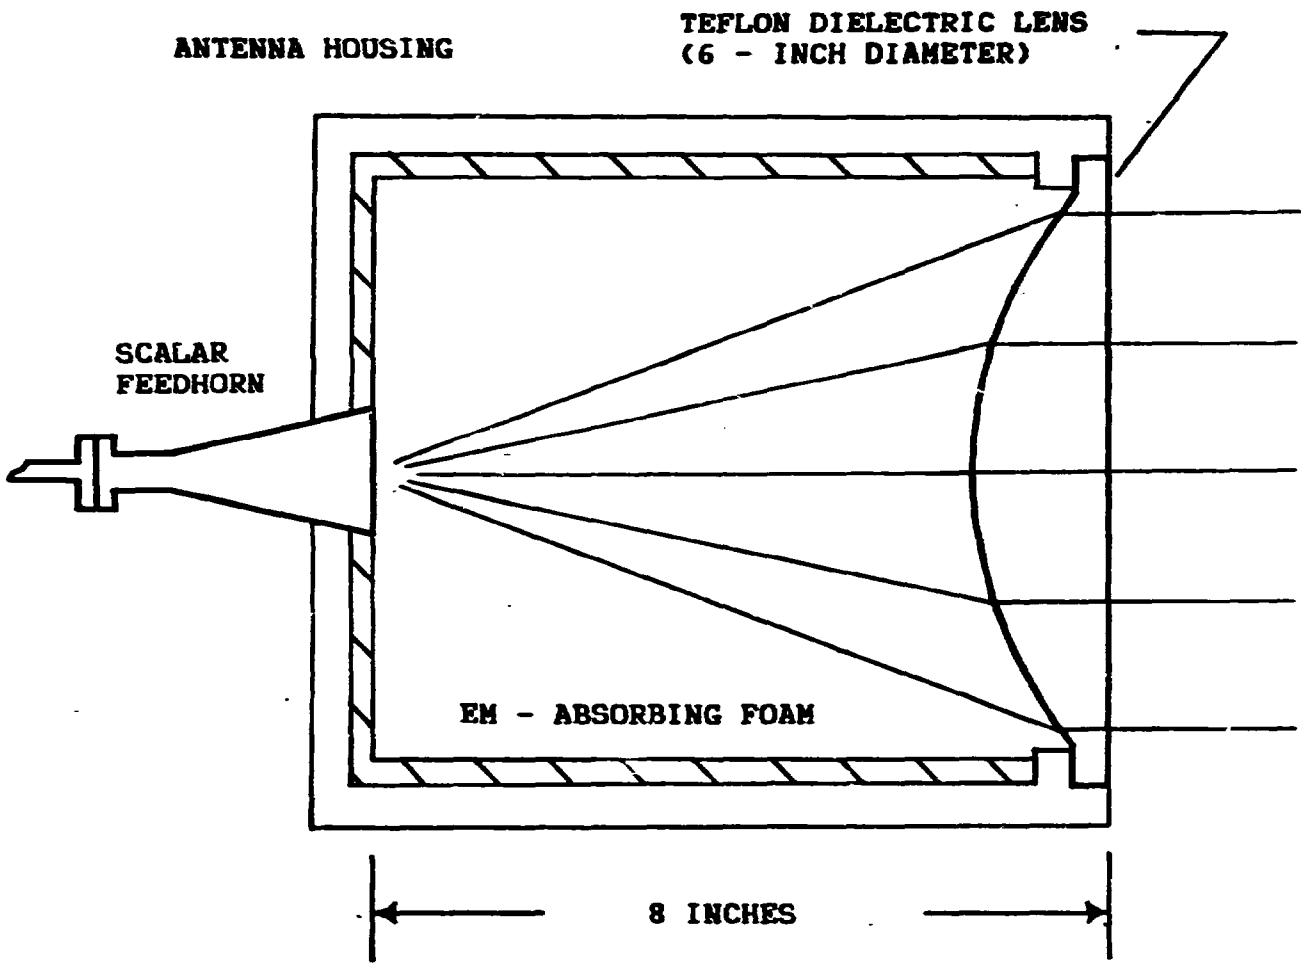

Figure 8. Scalar feedhorn and dielectric lens configuration comprising one of the antennas developed for use in the LFD.

level of the VTO is approximately 50 milliwatts, however, the attenuator shown in Figure 7 reduces the actual radiated power level. The VTO dissipates approximately 4.5 watts of power which is the main power drain in the LFD. A recently purchased replacement oscillator uses only three watts of input power.

Demodulation in the homodyne LFD is achieved by mixing the received return signal and a reference signal in a single-ended mixer. A silicon, Shottky-barrier diode mixer is being used because of its inherent low noise. The injection network shown in Figure 7 is used to supply the needed reference signal. By controlling the reference signal level and applying a suitable DC bias current to the mixer, excellent conversion efficiency and high receiver sensitivity are achieved. Numerous mixer tests performed during this program indicate that since receiver sensitivity is a function of both conversion efficiency and noise, the reference signal level should be relatively low (0.1-0.5 milliwatts as compared to the 5-20 milliwatts used in normal mixer operation). The ability to use a low reference signal level is advantageous from a viewpoint of power consumption and personnel safety.

The key feature of the current lifeform detector is its use of frequency modulation. Frequency modulation is achieved by applying a periodic ramp waveform to the varactor port of the VTO. As the ramp voltage rises from its starting to its peak voltage, the frequency of the VTO is linearly swept. Modulating frequencies as high as 125 kHz can be used in the current system. The amplitude of the modulating waveform controls the frequency deviation of the VTO, and hence, the range-cell size. To achieve ranging, the amplitude of the modulating waveform is adjusted so that information from the desired target range occurs at a specified frequency (in this case, a selected harmonic of the modulating frequency). Thus, a fixed frequency receiver can be used at the output of the mixer.

With this modulation/demodulation approach, the frequency deviation and range-cell size vary as a function of target range. The largest frequency deviation is used at the closest range which makes the range-cell size proportional to the target range. That is, the range-cell size increases as the target range is increased. This is undesirable since it is at long ranges that the greatest clutter-rejection is needed. This problem can be overcome by decreasing the modulating frequency as the target range is increased. This permits a large frequency deviation and narrow range-cell size to always be

used. During measurements with the current LFD, premeasured charts of the LFD's response are used to set the modulating frequency and deviation for the desired target range.

Several experiments were conducted to determine if the range-gating system of the current LFD behaved similarly to that predicted in Figure 6b (since no weighting was being used). These experiments also provided information for charts that would be needed to set the frequency deviation and modulating frequency when the LFD was field tested. In the first set of experiments, a target consisting of a sinusoidally moving corner reflector was positioned a known distance from the LFD. With the modulating frequency fixed at 62.5 kHz, the level of the return signal from the corner reflector was measured as the frequency deviation was adjusted from its minimum value (0 percent) to its maximum value (100 percent). The resulting data was used to plot the strength of the target return as a function of the frequency deviation.

Data obtained for a target range of 100 feet was used to produce the graph shown in Figure 9. The information in this graph shows that for best performance given a target range of 100 feet and a modulating frequency of 62.5 kHz, the frequency deviation should be set to approximately 40 percent. The effects of range-sidelobes is apparent in Figure 9 since the target could be seen to some extent at every deviation setting. The fact that the best frequency deviation setting was only 40 percent is an indication the modulating was too high for this target range. Recall that for best performance, the modulating frequency should be selected to permit full frequency deviation (i.e., 100 percent) to be used so that the minimum range-cell size can be achieved.

This initial experiment was repeated for target ranges from 75 to 175 feet in 25 foot increments, and for modulating frequencies from 12.5 to 125 kHz in 12.5 kHz increments (N will be used to denote the specific modulating frequency with N = 1 corresponding to the highest modulating frequency and N = 10 corresponding to the lowest modulating frequency). The resulting data was used to prepare graphs that illustrate how the modulating frequency and frequency deviation (0 to 100 percent) should be set as function of the target range. A graph showing results for N from 1-4 is presented in Figure 10. It should be apparent that a great deal of effort was required to obtain the information in this figure. To avoid this situation with future range-gating systems, efforts are underway to compare the results in Figure 10 to theoretically predicted results. If this comparison proves favorable, it

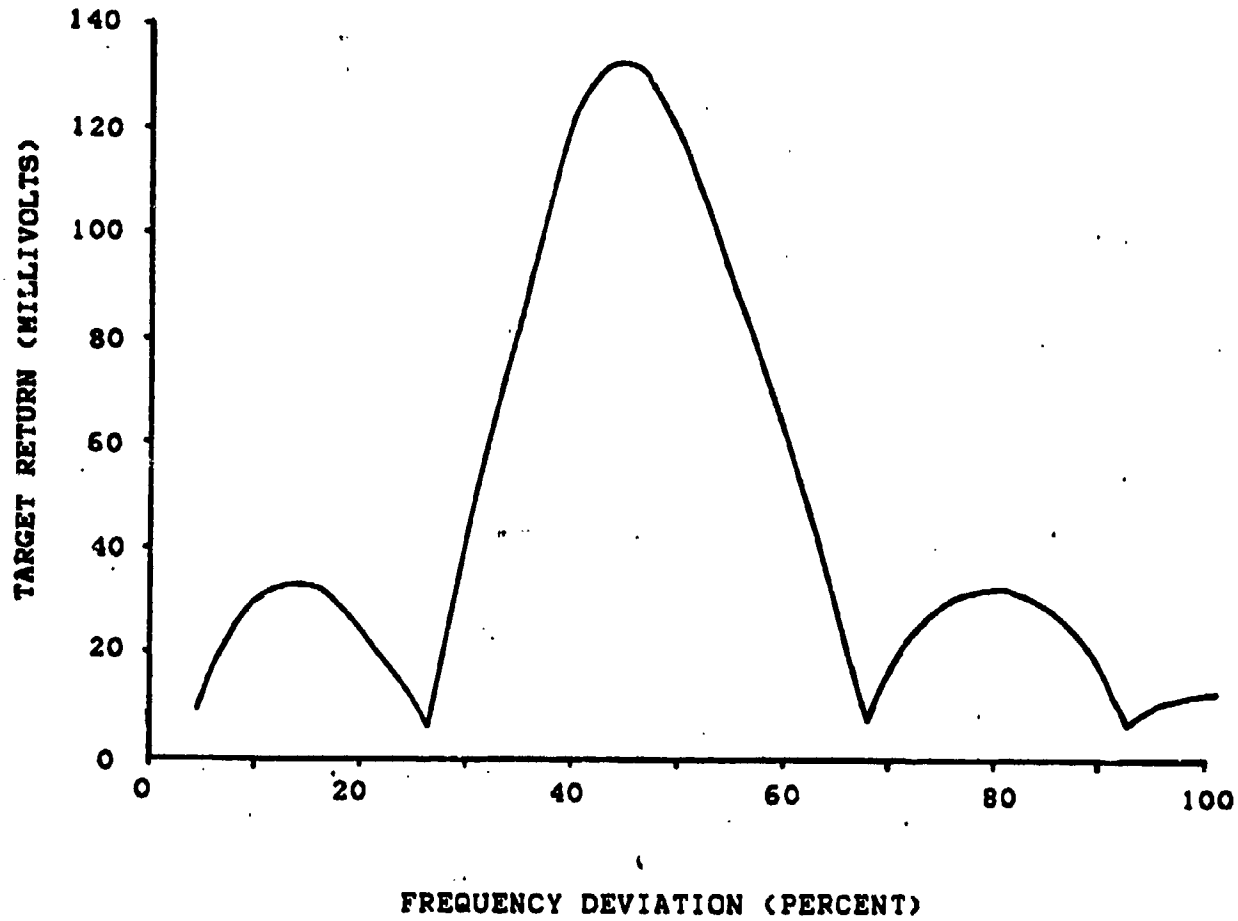

Figure 9. Graph of return signal strength as a function of range adjustment (deviation setting) for a target at a range of 30 meters.

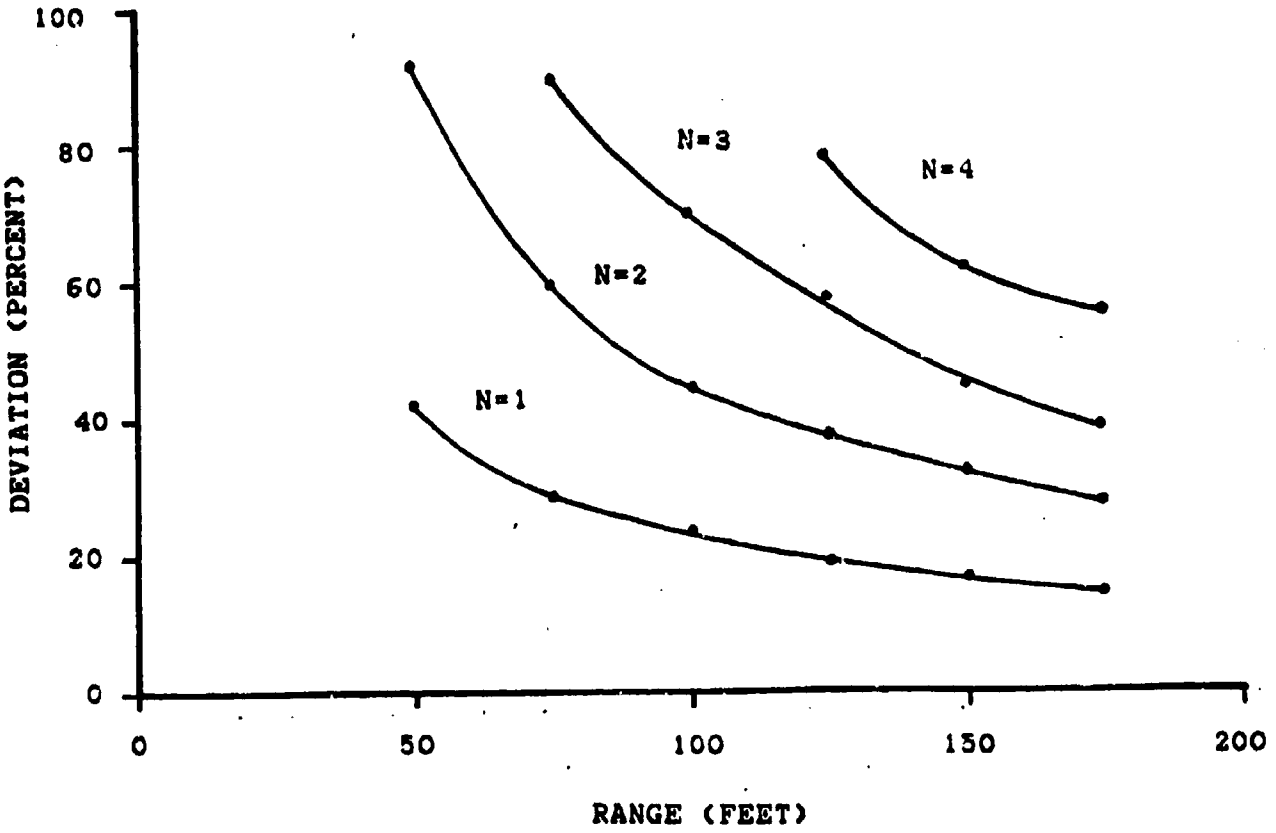

Figure 10. Graph of range adjustments (deviation settings) as a function of range for several different modulating frequencies (N=1 to 4).

will be possible to generate the required ranging data once the tuning response of the main oscillator is known.

In another informative experiment, the LFD detector was aimed at a stand of trees at a range of approximately 200 to 250 feet. With the modulating frequency set at 125 kHz (N = 1), the frequency deviation was set for a range of 175 feet. Because of the high modulating frequency (or low N), the required deviation was small (14 percent) and the resulting range-cell width was large (approximately 150 feet). Because the range-cell size was so large, the distant trees were partially within t.. main range cell and clutter effects were pronounced. The modulating frequency was then decreased in 12.5 kHz steps to 17.86 kHz (N = 7). At each new modulating frequency, the frequency deviation was reset for a range of 175 feet and the clutter from the trees remeasured. Since the range-cell size was decreasing as the modulating frequency was decreased, the effects of the clutter should have decreased at each new modulating frequency.

The results of this simple experiment are shown in Figure 11. As expected, the clutter was highest when the highest modulating frequency (lowest N and widest range-cell size) was used. As the value of N was changed from 1 to 2, the effects of the clutter were found to decrease because the trees were no longer in the main range cell. However, as the value of N was increased above 2, there was little improvement in the detected clutter level even though the range-cell sizes should have been small enough to discriminate against the trees. The fact that clutter from the trees could still be seen for high values of N, is evidence of the effects of range-sidelobes and demonstrates the need for refinement of the LFD's range-gating capability.

5. Experimental Data (Time-Waveforms)

Currently, a dual-channel strip-chart recorder and a digital storage scope are being used to monitor the output of the LFD. This makes it possible to directly observe respiratory and cardiac motion when the signal-to-noise ratio is sufficiently high. Examples of respiratory motion detected with an earlier version of the LFD for ranges of 30, 40, and 50 meters are now reviewed.

The respiratory data in Figure 12a was taken with the subject lying supine and perpendicularly to the antenna beamwidth with his left side toward the LFD. Examination of this data reveals respiratory information is present

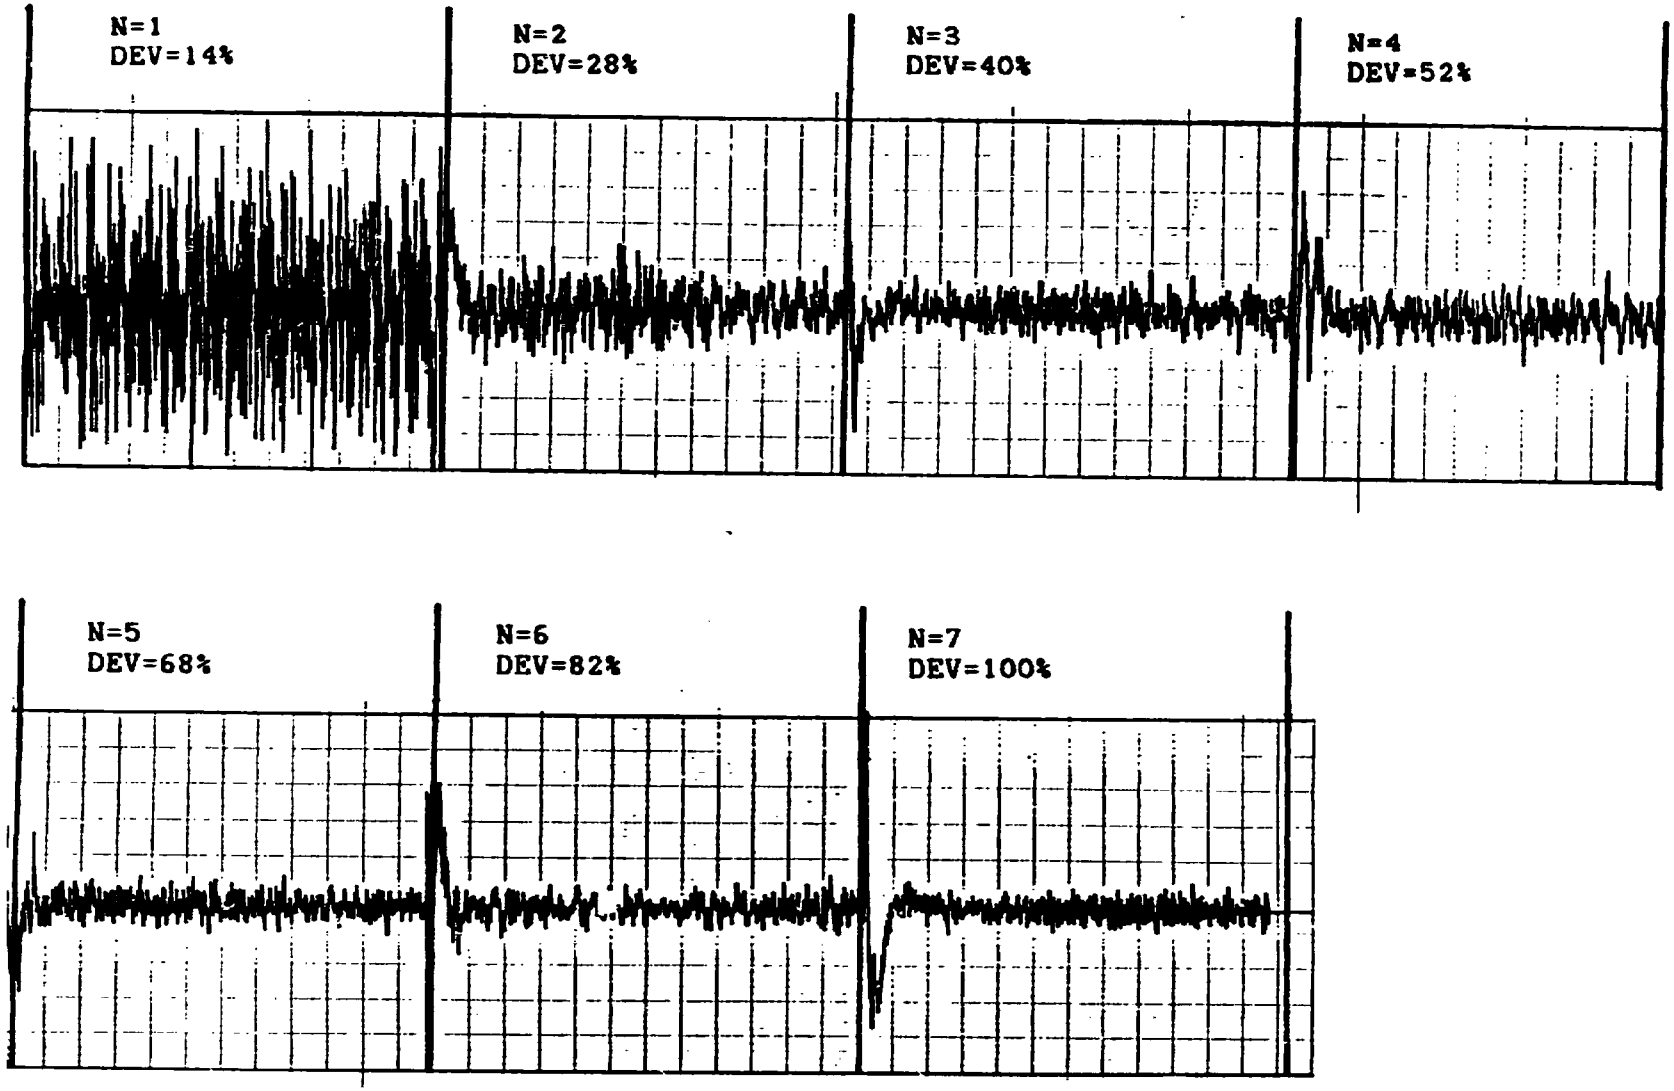

Figure 11. Example of effects of range-cell size and range-sidelobes on discrimination of clutter from distant stand of trees. LFD set for range of 175 feet. Trees at a range of approximately 200-250 feet.

CODE:	221-4	POWER:	0.05 mW
RANGE:	30 M	SUBJECT:	JS, SUPINE, SIDE
FREQ:	35 GHZ	FILTERS: (HZ)	HP 0.07 LP 2.0

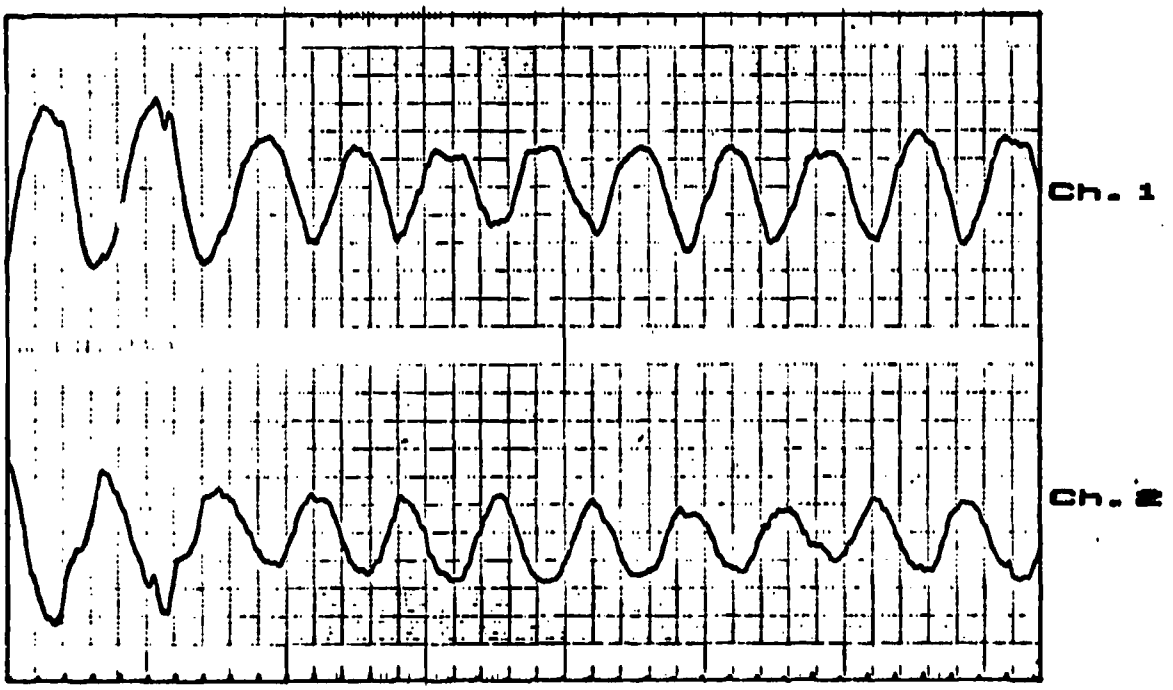

Figure 12a. Respiratory motion detected from 30 meters for subject lying supine with left side to beam from LFD.

in both Channel-1 and Channel-2 (these are the quadrature channels from the LFD). The fact that neither channel appears distorted indicates the detected motion was approximately one-eighth wavelength (one millimeter). The subject's chest was certainly moving more than a millimeter. The actual respiratory motion was probably more than a centimeter. However, most of this motion was in a vertical direction while the LFD was mainly sensitive to horizontally-direction motion. Thus, body position may prove to be an important factor.

To obtain some insight into the effect of body position, a second set of data was taken at 30 meters using a subject whose positon was rotated 90 degrees so that his head was toward the LFD. In this case, respiratory motion was again observed in both channels as shown in Figure 12b. However, the information in Channel-2 was distorted. This distortion was a result of the range-deadspots described previously. In this case, the target motion was large enough that the information in Channel-2 was not entirely suppressed. However, it was severely distorted and the respiratory frequency actually appears to be doubled for the first 4 to 5 breathing cycles. This result demonstrates one type of problem that can occur with systems not possessing quadrature-channel capability.

Examples of respiratory motion detected from a longer range of 40 meters are shown in Figure 13. Respiratory data is present in both channels in the data shown in Figure 13a. However, the data does appear to contain more noise than the data obtained from 30 meters. This was a result of weaker target returns from the greater target range and increased clutter from a large stand of trees that was behind the target at a range of 60 to 70 meters from the LFD. Respiratory data was present in only one channel in the data shown in Figure 13b. This demonstrates the loss in motion sensitivity that can occur because of the effects of range-deadspots.

Respiratory data measured from a range of 50 meters are shown in Figures 14. The results are similar to those obtained at closer ranges. In one case (Figure 14a), respiratory data can be clearly observed in both channels, while it is present in only one channel for the case shown in Figure 14b. A greater noise level can be noted in the data obtained from 50 meters. This was again a result of weaker return signals and increased clutter from the trees which were only 20 to 30 meters behind the target.

In examining the data in Figures 12-14, little evidence of cardiac

CODE:	224-12	POWER:	0.31 mW
RANGE:	30 M	SUBJECT:	MS, SUPINE, HEAD
FREQ:	35 GHZ	FILTERS: (HZ)	HP 0.07 LP 16
COM'ts:	RESP, POLZ. HORZ.		

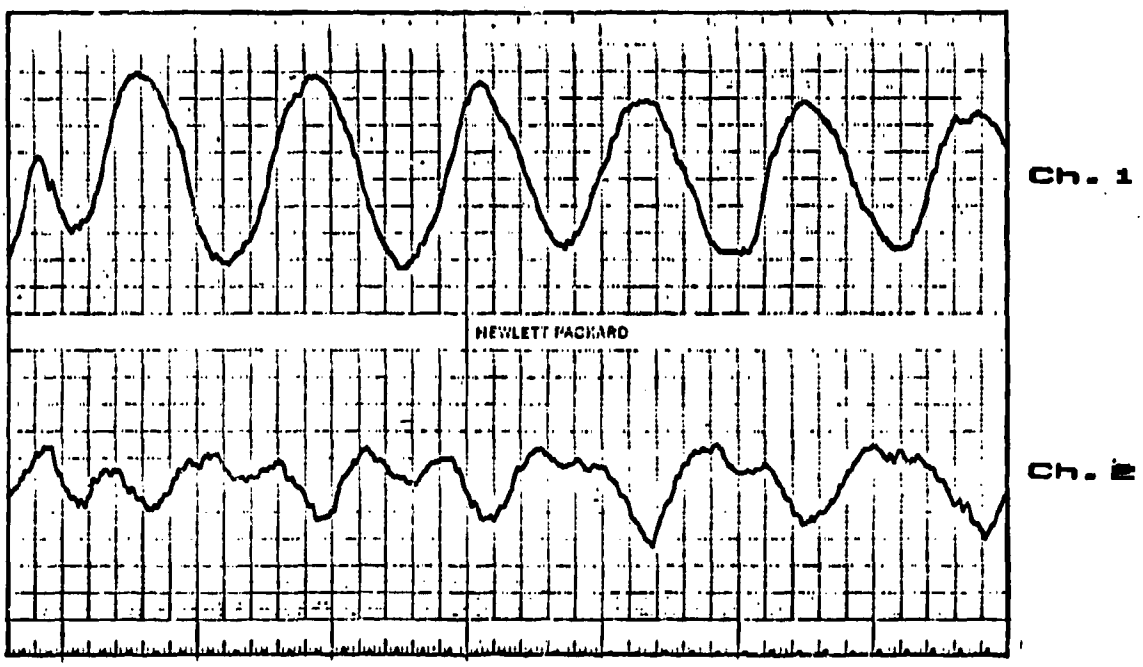

Figure 12b. Respiratory motion detected from 30 meters for subject lying supine with head to beam from LFD.

CODE: 224-3 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 0.07 (HZ) LP 16

COM' =: RESP.

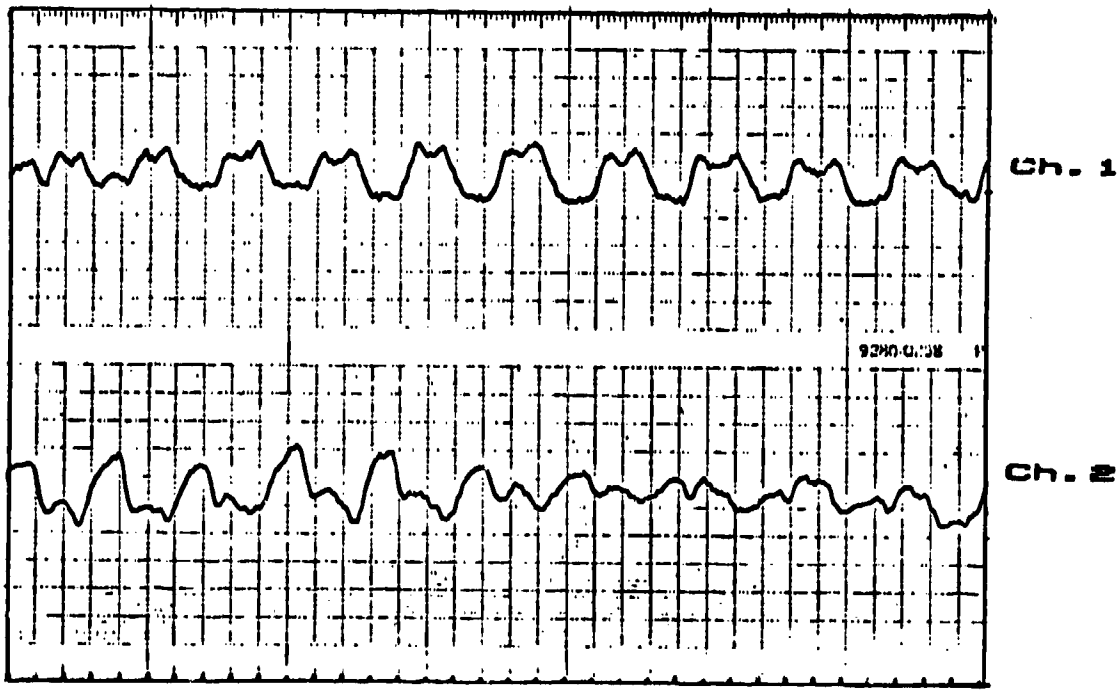

1 SECOND / TIMING MARK

Figure 13a. Respiratory motion detected from 40 meters for subject lying supine with left side to beam from LFD.

CODE: 224-4 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 0.07 (HZ) LP 16

COMMENTS: RESP.

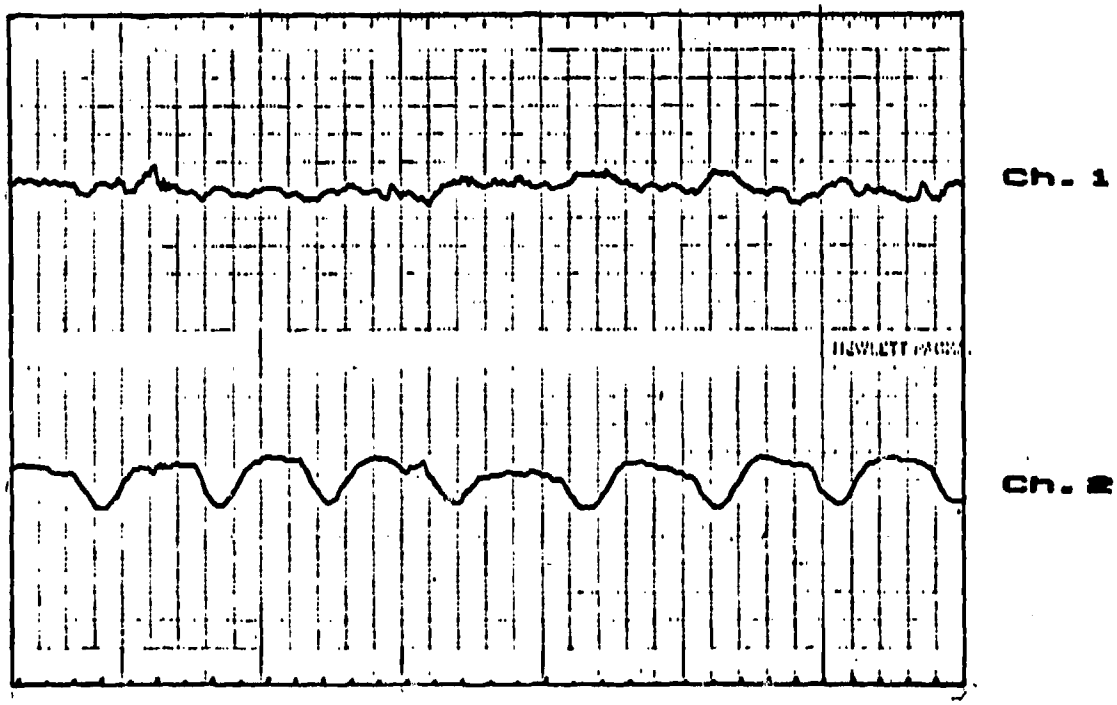

Figure 13b. Respiratory motion detected from 40 meters for subject lying supine with left side to beam from LFD showing effects of range-deadspots in Channel 1.

CODE: 224-8 POWER: 0.31 mW

RANGE: 50 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 1.0 (HZ) LP 16

COMMENTS: RESP, POLZ. HORZ.

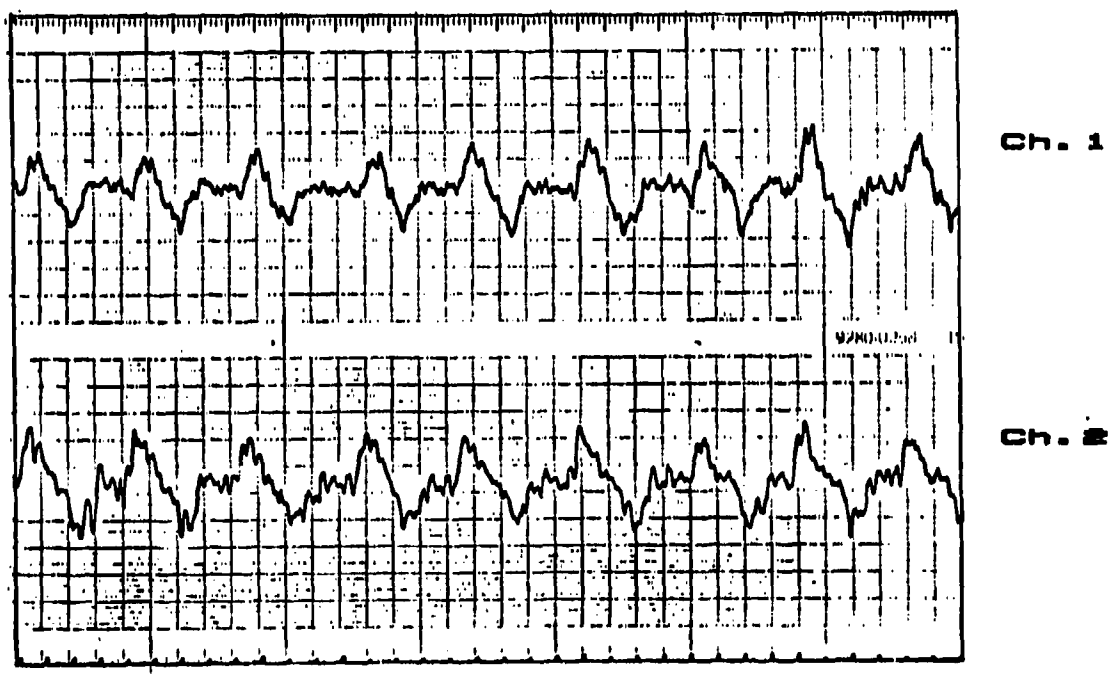

1 SECOND / TIMING MARK

Figure 14a. Respiratory motion detected from 50 meters for subject lying supine with left side to beam from LFD.

CODE:	224-9	POWER:	0.31 mW
RANGE:	50 M	SUBJECT:	SS, SUPINE, SIDE
FREQ:	35 GHZ	FILTERS: (HZ)	HP 1.0 LP 16
COM'TS:	RESP, POLZ. HORZ.		

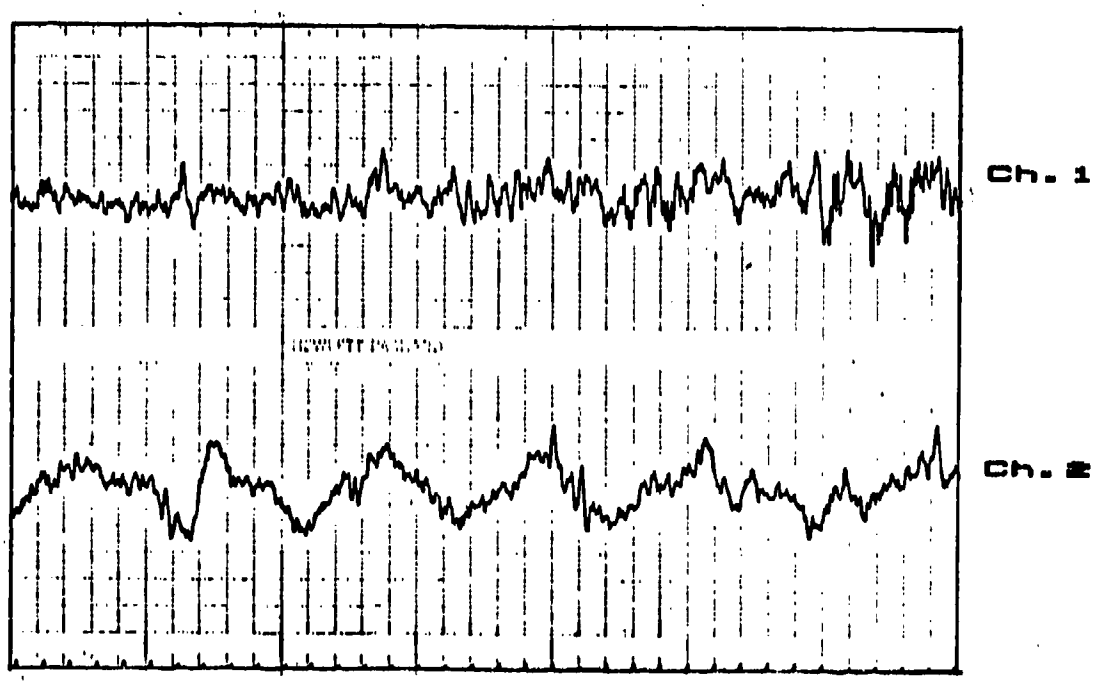

Figure 14b. Respiratory motion detected from 50 meters for subject lying supine with left side to beam from LFD showing effect of range-deadspots in Channel 1.

activity can be observed. Experience indicates that cardiac information is present, but is so weak that it cannot be seen on the recorder sensitivities required to observe respiration. To offset this problem, subjects were asked to stop breathing for short periods (20 to 30 seconds) so that the recorder sensitivities could be increased enough to permit cardiac activity to be observed. This approach proved successful. In general, it was found that the recorder sensitivity had to be increased by a factor of ten to enable cardiac information to be observed.

An example of cardiac motion detected from a range of 30 meters is shown in Figure 15. Although relatively weak, comparison of the suspected cardiac data in Figure 15a to the corresponding baseline data (recorded with no target present) in Figure 15b clearly shows the presence of a distinct heartbeat signal. Similar results were obtained from a range of 40 meters as shown in Figures 16a and 16b. In both cases, the cardiac information is masked by a strong noise component. Since this noise is not evident in the baseline data, it is probably not due to clutter. Instead, it appears to be due to the combined effects of clothing and slight body motions such as twitches, head motions, chest wall motions as subjects strained to hold their breath, and possibly gastrointestinal motions. It was found that much of the noise in the cardiac signal could be removed with a high-pass filter. When the filter cutoff-frequency was changed from 0.07 Hz to 1.0 Hz, the results in Figures 16c and 16d were obtained.

It was generally difficult to detect cardiac signals from a 50 meter range because of the large amount of background noise. The problem was clutter from the trees beyond the target location since a high level of noise was also observed for the baseline case. There were periods, however, when cardiac information could be observed from 50 meters. Results of one set of successful measurements are shown in Figures 17a and 17b (note that the recorder's horizontal and vertical sensitivities were increased for these measurements to make it easier to observe the weak cardiac signals). The cardiac information in Figure 17a is partially masked by the background noise but a distinct cardiac signature can be observed in Channel-1. It may be noted that this cardiac signal is very similar in shape to the signal seen in Channel-2 of Figure 16c.

The preceding examples show that the ability to observe detected respiratory and cardiac motion provides a convenient test of the performance

CODE:	221-5	POWER:	0.05 mW
RANGE:	30 M	SUBJECT:	JS, SUPINE, SIDE
FREQ:	35 GHZ	FILTERS: (HZ)	HP 0.07 LP 2.0
COMMENTS:	HB		

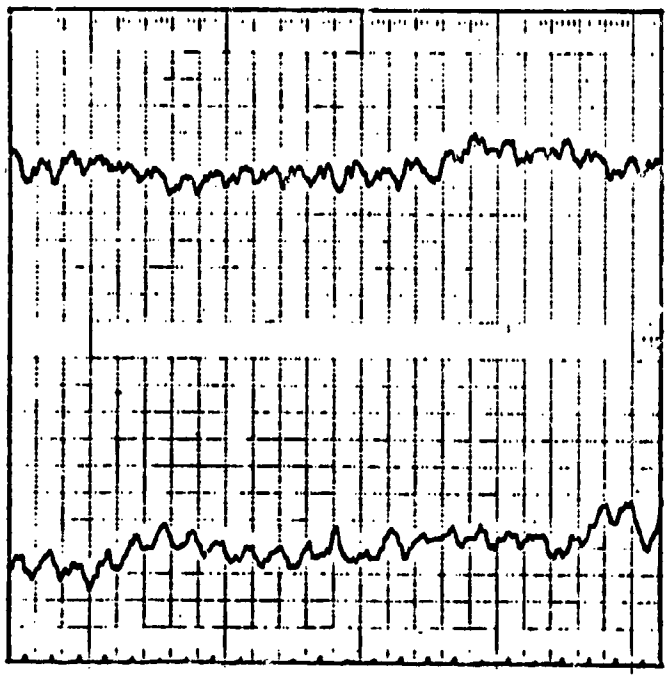

Figure 15a. Cardiac motion detected from 30 meters for subject lying supine and holding breath.

CODE: 221-5 POWER: 0.05 mW

RANGE: 30 M SUBJECT: JB, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 0.07 (HZ) LP 2.0

COMMENTS: BASELINE

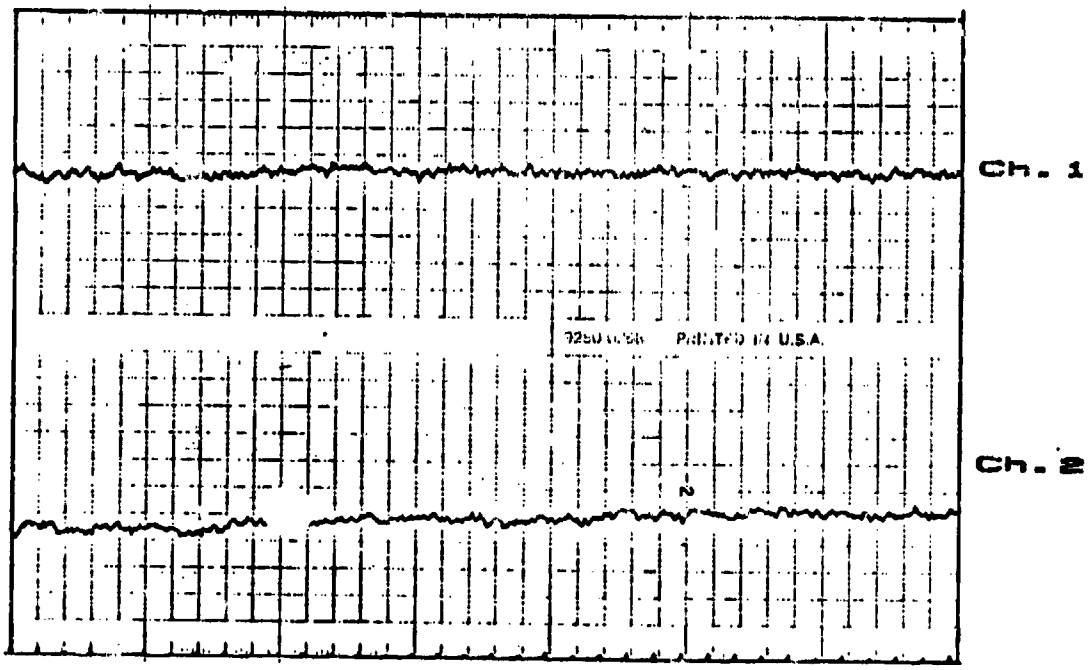

1 SECOND / TIMING MARK

Figure 15b. Baseline data for heartbeat data in Figure 15a.

CODE: 224-5 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JB, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 0.07 (HZ) LP 16

COMMENTS: HB

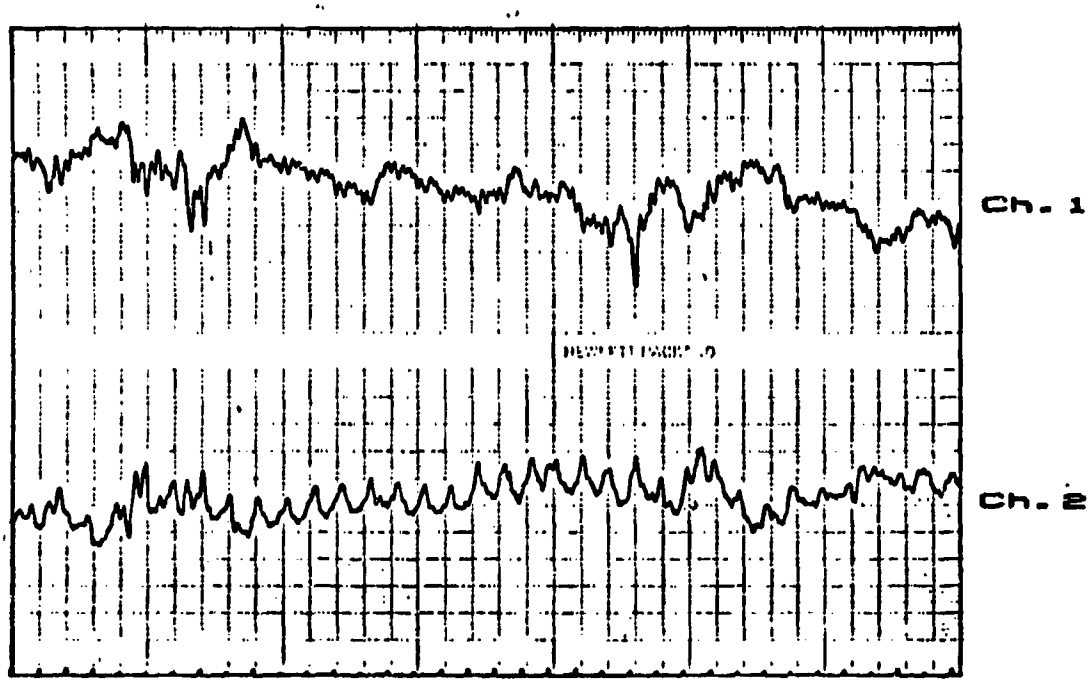

1 SECOND / TIMING MARK

Figure 16a. Cardiac motion detected from 40 meters for subject lying supine and holding breath.

CODE: 224-5 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 0.07 (HZ) LP 16

COMM's: BASELINE

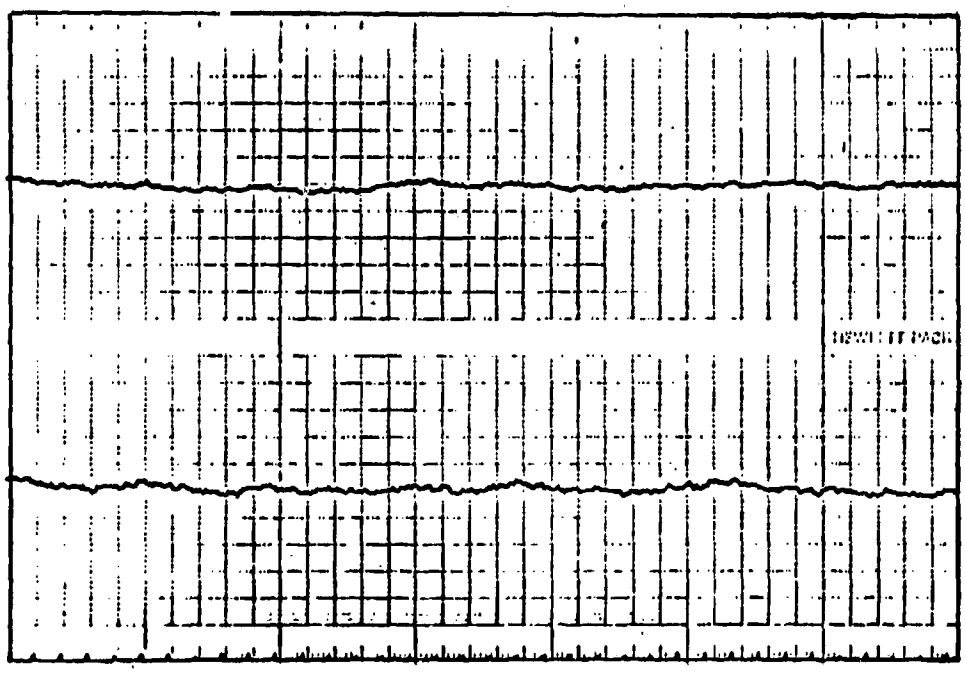

1 SECOND / TIMING MARK

Figure 16b. Baseline data for heartbeat data in Figure 16a.

CODE: 224-7 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 1.0 (HZ) LP 16

COMMENTS: HB

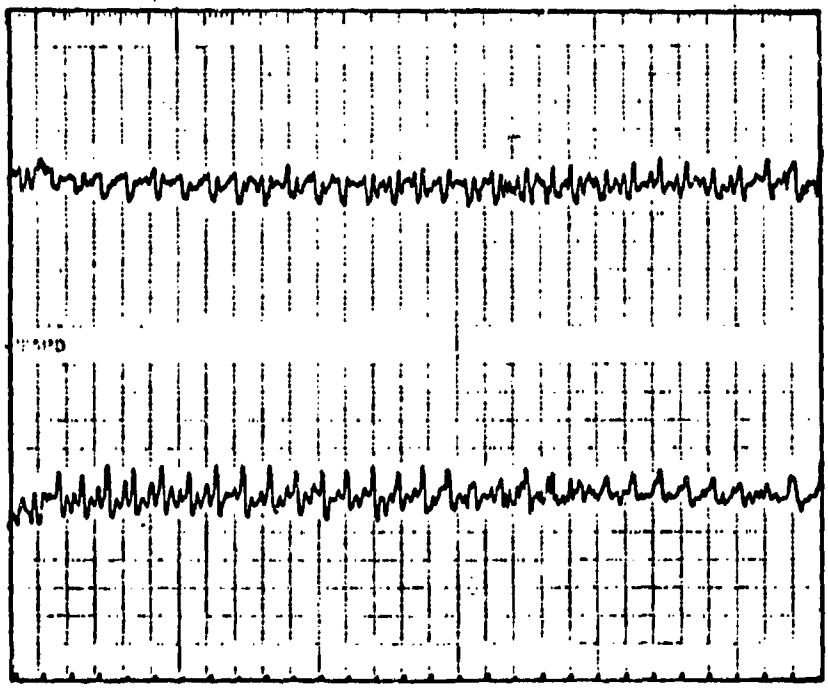

1 SECOND / TIMING MARK

Figure 16c. Cardiac motion detected from 40 meters for subject lying supine and holding breath. One Hz highpass filter used to remove low frequency noise.

CODE: 224-7 POWER: 0.31 mW

RANGE: 40 M SUBJECT: JB, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 1.0 (HZ) LP 16

COMMENT: BASELINE

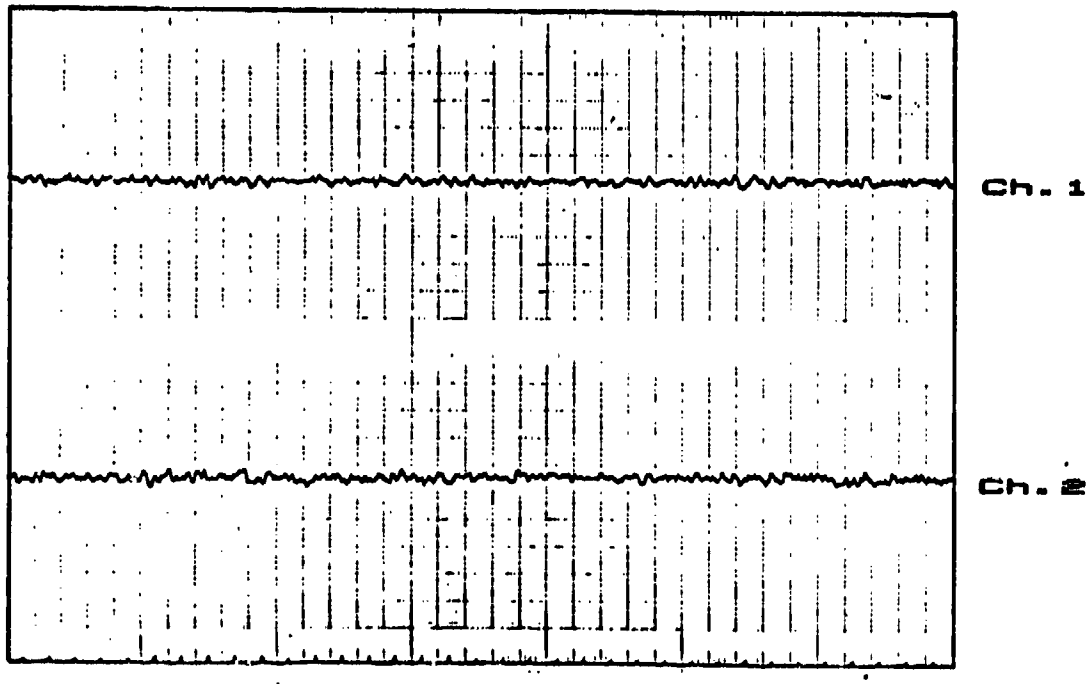

1 SECOND / TIMING MARK

Figure 16d. Baseline data for heartbeat data in Figure 16c.

CODE:	224-8	POWER:	0.31 mW
RANGE:	50 M	SUBJECT:	JS, SUPINE, SIDE
FREQ:	35 GHZ	FILTERS: (HZ)	HP 1.0 LP 16
COM'N:	HB, POLZ. HORZ.		

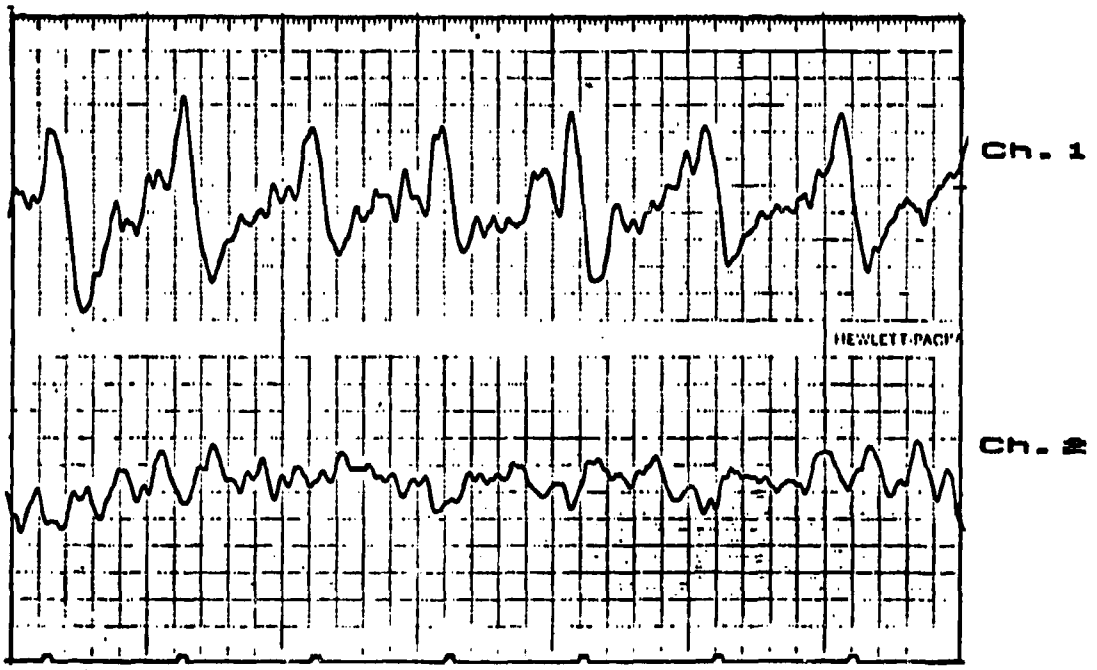

1 SECOND / TIMING MARK

Figure 17a. Cardiac motion detected from 50 meters for subject lying supine and holding breath. Recorder's horizontal and vertical sensitivities both increased.

CODE: 224-8 POWER: 0.31 mW

RANGE: 50 M SUBJECT: JS, SUPINE, SIDE

FREQ: 35 GHZ FILTERS: HP 1.0 (HZ) LP 16

COM'TS: BASELINE, POLZ. HORZ.

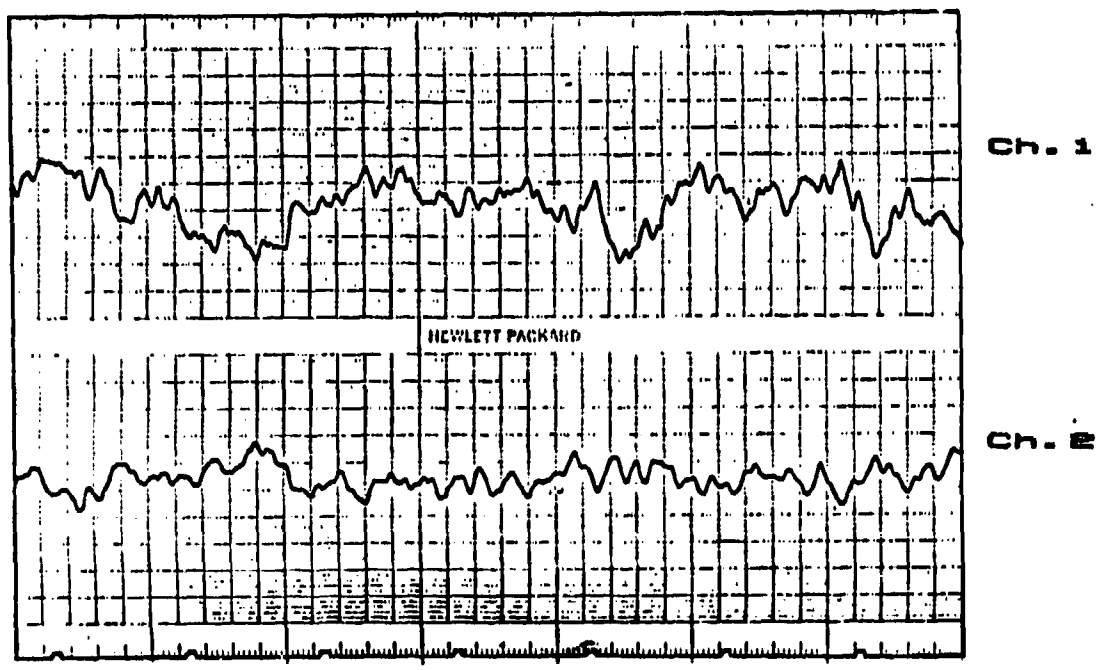

Figure 17b. Baseline data for heartbeat data in Figure 17a.

of the LFD. However, it does not appear to represent a fair test of performance. This assumption is based on the fact that a very high signal-to-noise ratio (SNR) is required to clearly observe respiratory and cardiac waveforms in the output of the LFD. With appropriate signal processing, it should be possible to achieve reliable detection with a lower SNR. Therefore, improving the current signal processing capability is a major requirement.

The signal processing capability of the LFD currently consists of a set of bandpass filters with cutoff frequencies based on the range of respiratory and cardiac rates that are to be detected. Since allowances must be made for the wide variations that can occur in these rates, the filters have relatively wide bandwidths. Typical bandwidths are 0.05-1 Hz for the respiratory filters and 0.5-10 Hz for the cardiac filters (although an upper cutoff of 20 Hz may be a better choice). Since these filter bandwidths are wide relative to the information bandwidth typically observed during a single measurement, it is not possible to achieve very good noise reduction. This is not a serious problem when the SNR ratio is high, but is a problem when the signal-to-noise ratio is low such as at long ranges or under high clutter conditions.

A possible improvement would be to replace the current wideband filters with a narrow-band filter that could be tuned over the frequency band of interest. A narrower filter bandwidth would provide greater noise reduction, but the added inconvenience and time required for tuning the filter over the entire possible information band would have to be suffered.

A bank of narrow-band, bandpass filters that cover the required information band could be substituted for the single tunable filter. An appropriate detector could be used to check the output of each filter in the bank to determine the presence or absence of target information. It may be noted that if the individual filter bandwidths were selected properly, a sinusoidal signal would be outputted by one or more of the filters in the bank whenever any type of periodic information was processed through the filters. By using a suitable detection scheme, it appears this would make it possible to discriminate against strong, but irregular noise components. However, it would be necessary to observe the output of each filter for a period sufficiently long to permit the effects of random noise fluctuations to average out.

To achieve best performance with the filter-bank approach, the filter bandwidths should approximately match the bandwidth of components in the

spectrum of respiratory and cardiac information. Experience indicates that individual filter bandwidths as narrow as 0.05 Hz might be necessary. Thus, if the filter bank were to cover the entire 0.05-10 Hz band covered by the current respiratory and cardiac filters, a suitable filter bank would have to contain 199 individual filters. Such a network would of course be troublesome to implement. A smaller number of filters could be used, but there would be a corresponding loss in detection performance.

If the output power of the proposed filter bank was detected and plotted as function of each filter's center frequency, the resulting graph would represent an estimate of the power spectrum of the information passed through the filter bank [19-22]. Therefore, an alternative to using a filter bank would be to process the output of the LFD through a system capable of using efficient algorithms to compute the power spectrum of the inputted signal. Individual components in the computed power spectrum would be analogous to the outputs of individual filters in the filter bank. With this approach, the required detection could be performed by comparing the target power spectrum to a noise-only spectrum. In this case, knowledge of the characteristics of the noise becomes an important part of the overall detection process.

The following part of this report discusses the general subject of signal processing. Background information on target detection as well as information on obtaining good power spectrum estimates is included. Some of the advantages as well as some of the costs of performing signal processing are also identified through results of several signal processing experiments performed during the third year.

C. Signal Processing

1. Basic Information

The term signal processing can be used to describe anything done to enhance the ability of the LFD to detect the presence or absence of target information. Therefore, in the strictest sense, signal processing includes capabilities such as range-gating, directive antennas, and low-noise receivers, that are being used to enhance the performance of LFD. However, since these topics have already been reviewed, the current discussion will focus on the remaining aspect of signal procesing which consists of operations performed to reliably and efficiently detect the presence of target information in the noise-contaminated output of the LFD.

This discussion will eventually lead to the idea of achieving target

detection based on power spectrum analysis of samples of the LFD's output. Under appropriate conditions, the literature [2,3], as well as results of preliminary experiments, indicate power spectrum estimates can be used to achieve very-good, possibly near-optimal, detection performance. In general, the price that must be paid for improving detection performance through signal processing is observation time (observation time is the period over which target information is collected and should not be confused with computation time which can be relatively short with a suitable processor). When power spectrum analysis is used, it will be shown that increased observation periods are needed to achieve narrower resolution bandwidths and to permit an effective level of averaging to be performed.

The term target-detection has been used frequently during this report to describe the process of determining the presence or absence of target information (where target information is considered to be respiratory and cardiac activity). For the LFD, target-detection can be defined as the process of deciding between two hypotheses [23-28]. Hypothesis-1 is that the target is not present while hypothesis-2 is that the target is present. A decision between the two possible hypotheses could be made by passing the information from the LFD through a signal processing system that would provide an output that could be compared to a suitable threshold level. If the output of the processor exceeds the threshold, a decision of target presence is made. Conversely, if the processor output does not exceed the threshold, the alternate decision of target absence is made.

Selecting a detection threshold level requires making a compromise between the two possible types of detection error. One type of error occurs when system noise exceeds the established threshold level and a target presence is falsely indicated. This type of false-alarm error is statistically described as Type-1 error and can be characterized by a corresponding probability of false-alarm. A second type of error occurs when a target signal does not exceed the threshold and a target's presence is not detected. This type of missed-detection error is statistically described as a Type-2 error and can be characterized by a corresponding probability of missed detection (or more commonly by the related probability of detection).

For good detection performance, the threshold level is set to achieve a high probability of detection and a low probability of false-alarms. A common procedure in radar applications is to fix the false-alarm probability at an

acceptable value, then adjust system parameters to optimize the detection probability [1]. System parameters that can be varied to adjust the detection and false-alarm probabilities are the signal-to-noise ratio (SNR) and the observation period.

As might be expected, high SNR's and long observation periods generally result in the best detection performance. If a short observation time is required, the SNR must be maximized to achieve good detection performance. Conversely, if the SNR is low, a long observation time is necessary to achieve good detection performance. Since short observation times are desired for the LFD, the SNR must be made as high as is practical. This is why improving the clutter-rejection capabilities of the LFD is considered to be so important. When the target range and clutter level become so great that respiratory and cardiac information are heavily corrupted by noise (i.e., low SNR), signal processing can be used to offset the loss in receiver performance, but at a price of longer observation times.

2. Types of Detectors

When signal processing is used to improve detection performance, precautions must be taken to insure that the selected detector is suitable for the particular detection problem. Detection problems vary as a function of the specific signal and noise characteristics. They range from the signal-known-exactly (or SKE) case in which the amplitude, phase, and frequency of the signal being detected are known exactly, to the case where so little is known about the signal being detected that it must be treated as a random process. The more that is known about a particular detection problem, the better the detector that can be derived for that problem. An optimum detector can be derived in most cases. In each case, the definition of optimum depends on the measure of goodness selected to evaluate the detector's performance. A common detector choice for detection problems such as the LFD is the matched filter, which is a special case of an optimum linear filter that has been designed to maximize the output SNR [1].

The SKE case is one detection problem that has been heavily analyzed. An example of this type of problem is detection of a sinusoidal signal of known amplitude, phase, and frequency that is corrupted by additive white noise. The optimum detector for this case is the matched filter detector. This detector can be implemented using a correlator to multiply the input signal-

plus-noise with a replica of the signal, then integrating the product from the correlator. At the end of the specified observation period, the integrator output can be compared to a suitable threshold level to make the required detection decision.

The SKE case is one of the few detection problems for which relatively simple mathematical results are available. Analysis of the SKE case is informative because the optimum detector for this case defines the limits of the improvements that can be achieved with signal processing. Analysis of this case also provides a convenient way to evaluate the relationship between detection probability, false-alarm probability, signal-to-noise ratio, and observation period.

A graph of detection performance for the SKE-case is shown in Figure 18 [29]. This graph shows detection probability (PD) and false-alarm probability (PF) as functions of a parameter called the detectability index (d) which is simply the square-root of the product of SNR (1 Hz bandwidth) and the observation period (T). There are several ways to interpret and utilize the graph in Figure 18. One informative use of this graph is to weigh the cost of receiver performance (i.e., SNR) versus the cost of signal processing (i.e., the observation period).

For example, suppose the SNR and T are set to produce a specific detectability index, d. Given this detectability index, a desired false alarm probability can be specified. The resultant detection probability can then be read from the graph. If a higher detection probability is to be achieved without changing the detectability index, a higher false alarm probability must be accepted. That is, improving the detection probability results in a worse false-alarm probability. If the detection probability is to be increased while maintaining the same false-alarm probability, the detectability index must be improved. Recall that improving the detectability index requires increasing the SNR and/or the length of the observation period.

As a second example, suppose the LFD could be treated as a SKE case. Because of the LFD's intended use, an extremely high detection probability is required, perhaps 99.9 percent. Although not as detrimental as missed detections, false-alarms are also undesired. An acceptable false-alarm probability for the LFD might be one percent. For these probability values, Figure 18 shows that a detectability index of approximately 5.5 would be

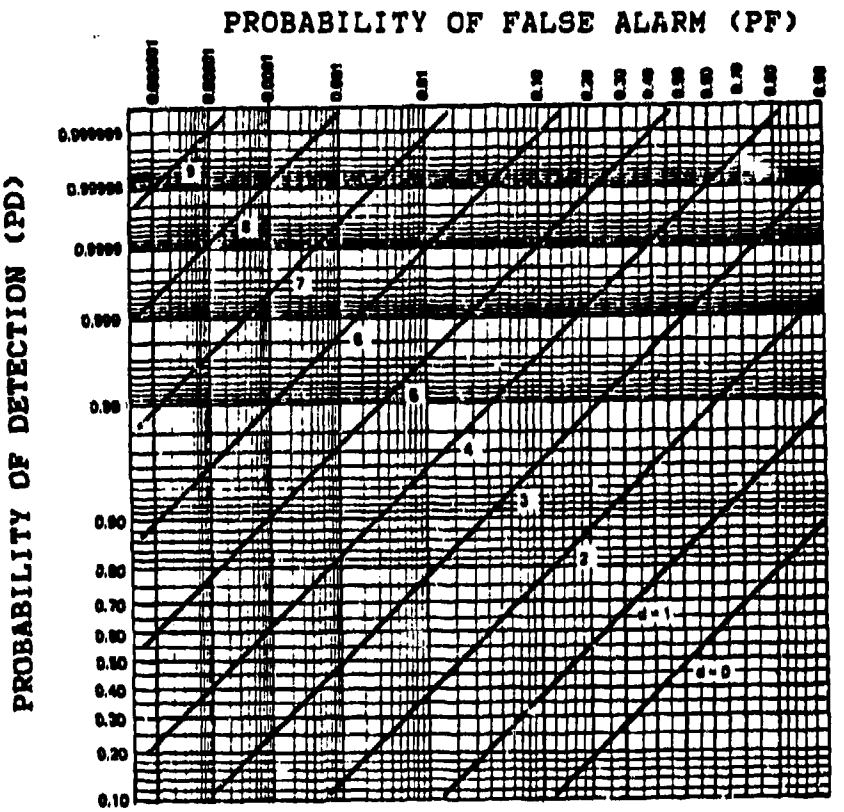

Figure 18. Detection probability and false-alarm probability as a function of the detectability index, d, for the signal-known-exactly case [29].

required. If the SNR in a one Hz bandwidth equaled 0.5 (or 3 dB), the observation period required to achieve the desired false-alarm and detection probabilities can be computed as 60.5 seconds.

Detection performance is critically dependent on the detectability index, particularly in the region of good detection performance (i.e., high PD and low PF). If the SNR were to drop (e.g., because of an increase in range or clutter), the observation time would have to be increased if the detection and false-alarm probabilities were to be maintained. If the SNR drops, but the observation period is not increased, a worse detection probability and/or false alarm probability must be accepted. For illustrative purposes consider that if the target range was doubled, the SNR would decrease by at least 12 dB due to a loss in return signal strength. To maintain the detection and false-alarm probabilities in the preceding example at 99.9 percent and one percent, respectively, the observation time would have to be increased from 60.5 seconds to 968 seconds. It should be clear that achieving good detection performance is a combination of using a good receiver and adequate signal processing. Signal processing can definitely be beneficial, but can not be used to compensate for poor receiver performance unless one is willing to pay the required price of extended observation periods.

On the other extreme from the SKE-case is a case where so little is known about the signal being detected that it must be treated as a random process. In this case, detection could be achieved using a technique usually referred to as an energy detector which requires little knowledge except for the approximate frequency limits of the signal being detected [30-32]. Detection consists of determining if any energy in the suspected frequency band exceeds an established threshold. This type of approach could be used with the LFD if the outputs of the existing respiratory and cardiac filters were integrated and then compared to suitable threshold levels.

The performance of the simple energy detector is poorer than that of the matched filter used in the SKE-case. However, the simple energy detector has the advantage that its performance will not be seriously degraded by variations in signal characteristics. For example, if frequency variations occurred in a sinusoidal signal being detected with a matched filter detector, the performance of the SKE detector could be seriously degraded. Conversely, frequency variations would have little impact on the performance of the simple energy detector.

The LFD is not perfectly described by either of the two extreme cases discussed in the preceding examples. Casualty respiratory and cardiac information as well as clutter-related noise will vary widely. Therefore, it would be difficult to treat the LFD as a SKE problem (e.g., see Figures 12-17). However, it is known that the detected waveforms from the LFD are often periodic in nature and reasonably accurate frequency limits can be defined for the detected cardiac and respiratory information. Therefore, it should be possible to derive a detector that provides better detection performance than the simple energy detector.

The performance of any detector developed for the LFD will depend on the level of information known about the signals and noise outputted by the LFD. The less knowledge that can be assumed about the problem, the greater the SNR and observation period must be to achieve a desired level of detection performance. Therefore, a large part of the third-year efforts on this program have focused on spectral analysis of the signals and noise. Spectral studies were performed primarily to characterize the signals and noise associated with the LFD in order to provide data for deriving a suitable signal processing system. However, in addition to this basic data, the spectral studies provided ideas for performing signal detection based on power spectrum estimates and and also provided information relevant to selecting spectral parameters to obtain useful power spectrum estimates.

3. Power Spectrum Analysis

Power spectrum analysis provides a means of evaluating the spectral behavior of random processes by providing an indication of the average power as a function of frequency [19-22]. Power spectrum analysis is useful for a number of detection applications including detection of periodic signals in noise [2-4]. Examples of the time-waveforms and corresponding power spectrum estimates for a noise-free sinusoidal signal and the same sinusoidal signal corrupted by noise are shown in Figure 19. For the noise-free case, the sinusoid can be clearly observed in both the time and frequency domains. For the noise corrupted case, the sinusoid is too heavily masked to be observed in the time domain. However, examination of the power spectrum of the noisy signal reveals the presence of a spectral component at the proper frequency. By comparing this spectrum to a suitable threshold, detection of the noise-corrupted sinusoid would be possible.

(a) Low-noise.

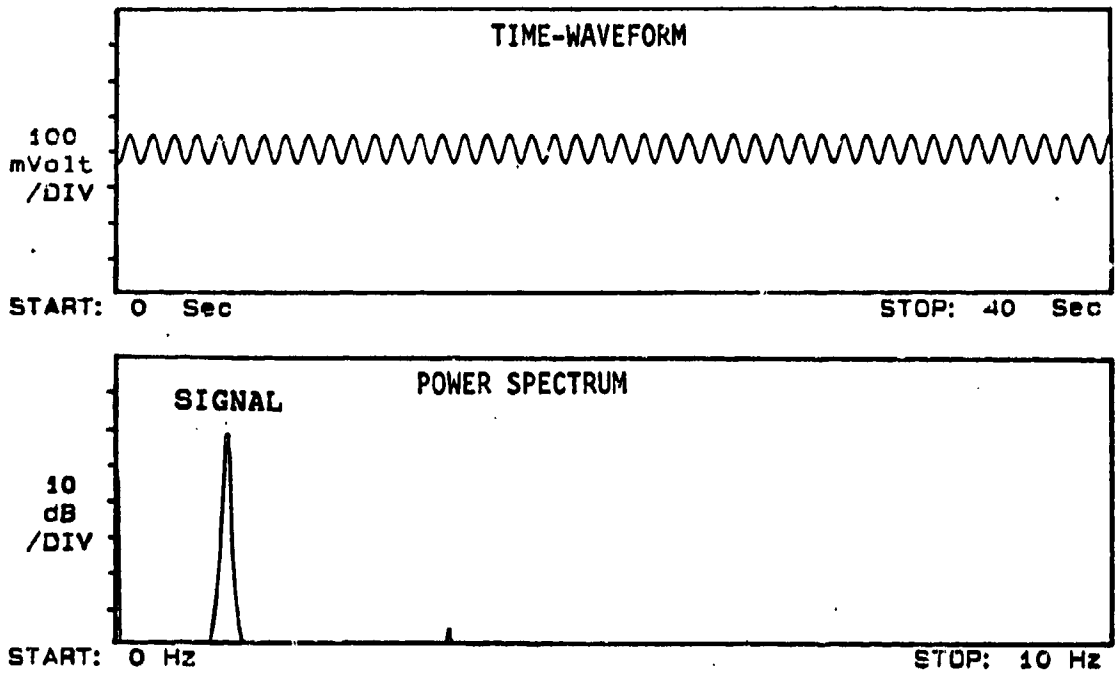

(b) High-noise.

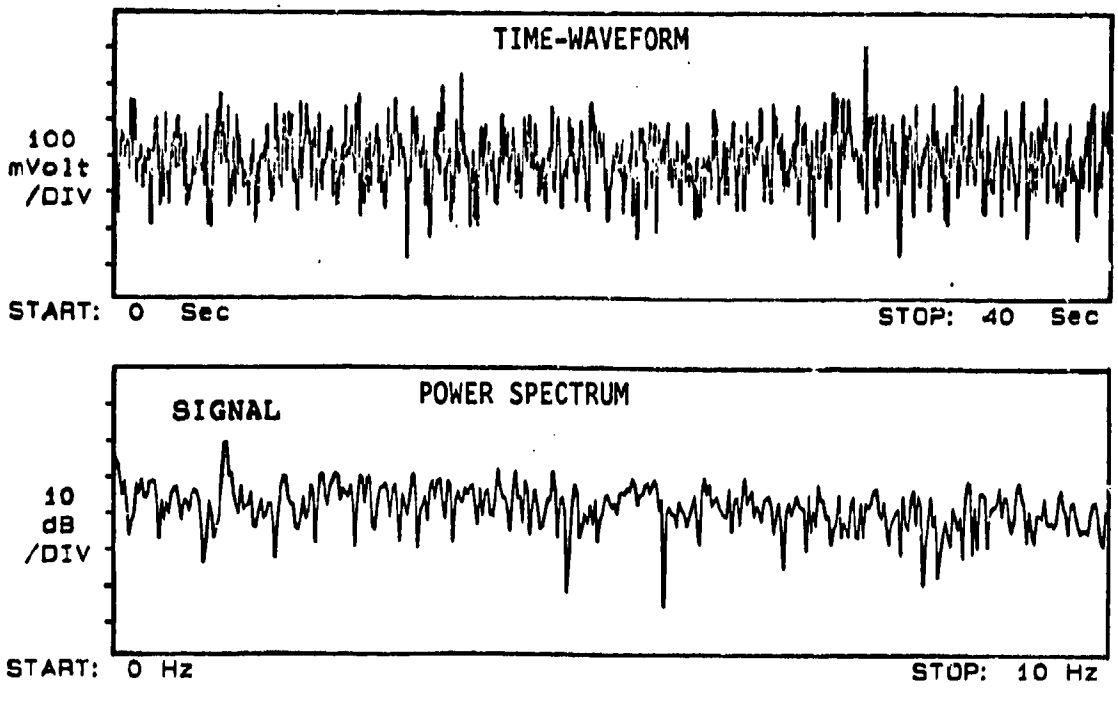

Figure 19. Examples of time-waveforms and power spectrum estimates for sinusoidal signal with (a) low noise and (b) high noise.

To achieve good power spectrum estimates when dealing with narrowband signals, it is necessary to use a narrow resolution bandwidth (BW) to minimize the effects of noise. Recalling the earlier analogy between the power spectrum and a filter bank, a narrow resolution bandwidth is equivalent to a filter bank comprised of a large number of narrowband filters. It should be noted that the BW is inversely proportional to the length of the observation period. Thus, narrowing the BW to achieve lower noise requires increasing the observation period.

The significance of the BW can be observed by performing power spectrum analysis on a sinusoidal signal plus noise using three different BW's. It is convenient to describe the three cases as the low, medium, and high resolution cases, where the term "high" refers to the narrowest resolution. The power spectrum estimates for these three cases are shown in Figure 20. A narrow peak corresponding to the sinusoidal signal can be observed in all three cases. To permit the relationship between the BW and noise to be observed, a line corresponding to a imaginary threshold has been drawn at an identical level on all three estimates.

For the medium-resolution case shown in the top graph, the selected threshold is slightly above the noise but well below the signal peak. This should produce a high detection probability and a low false-alarm probability (where the terms high and low are used in a relative sense). The middle graph in Figure 20 shows results obtained when a low BW was used. In this case, the signal peak is still well above the threshold level and a high detection probability would be maintained. However, the higher noise level causes the threshold to be exceeded at many frequencies where it is known no information is present. Thus, the false-alarm probability would be significantly worse. If the false-alarm probability had to be kept low, the threshold level would have to be raised. This would of course also decrease the detection probability.

Results obtained when the high BW was used are shown in the bottom graph in Figure 20. The signal peak is well above the threshold so a high detection probability is maintained. In addition, the decreased noise resulting from the higher resolution is now well below the established threshold level. This results in a lower false-alarm probability since there is less chance of the lower noise-level exceeding the threshold. If the false-alarm probability was to remain unchanged from that for the original medium-resolution case, the threshold for the high-resolution case could be lowered. This would have the

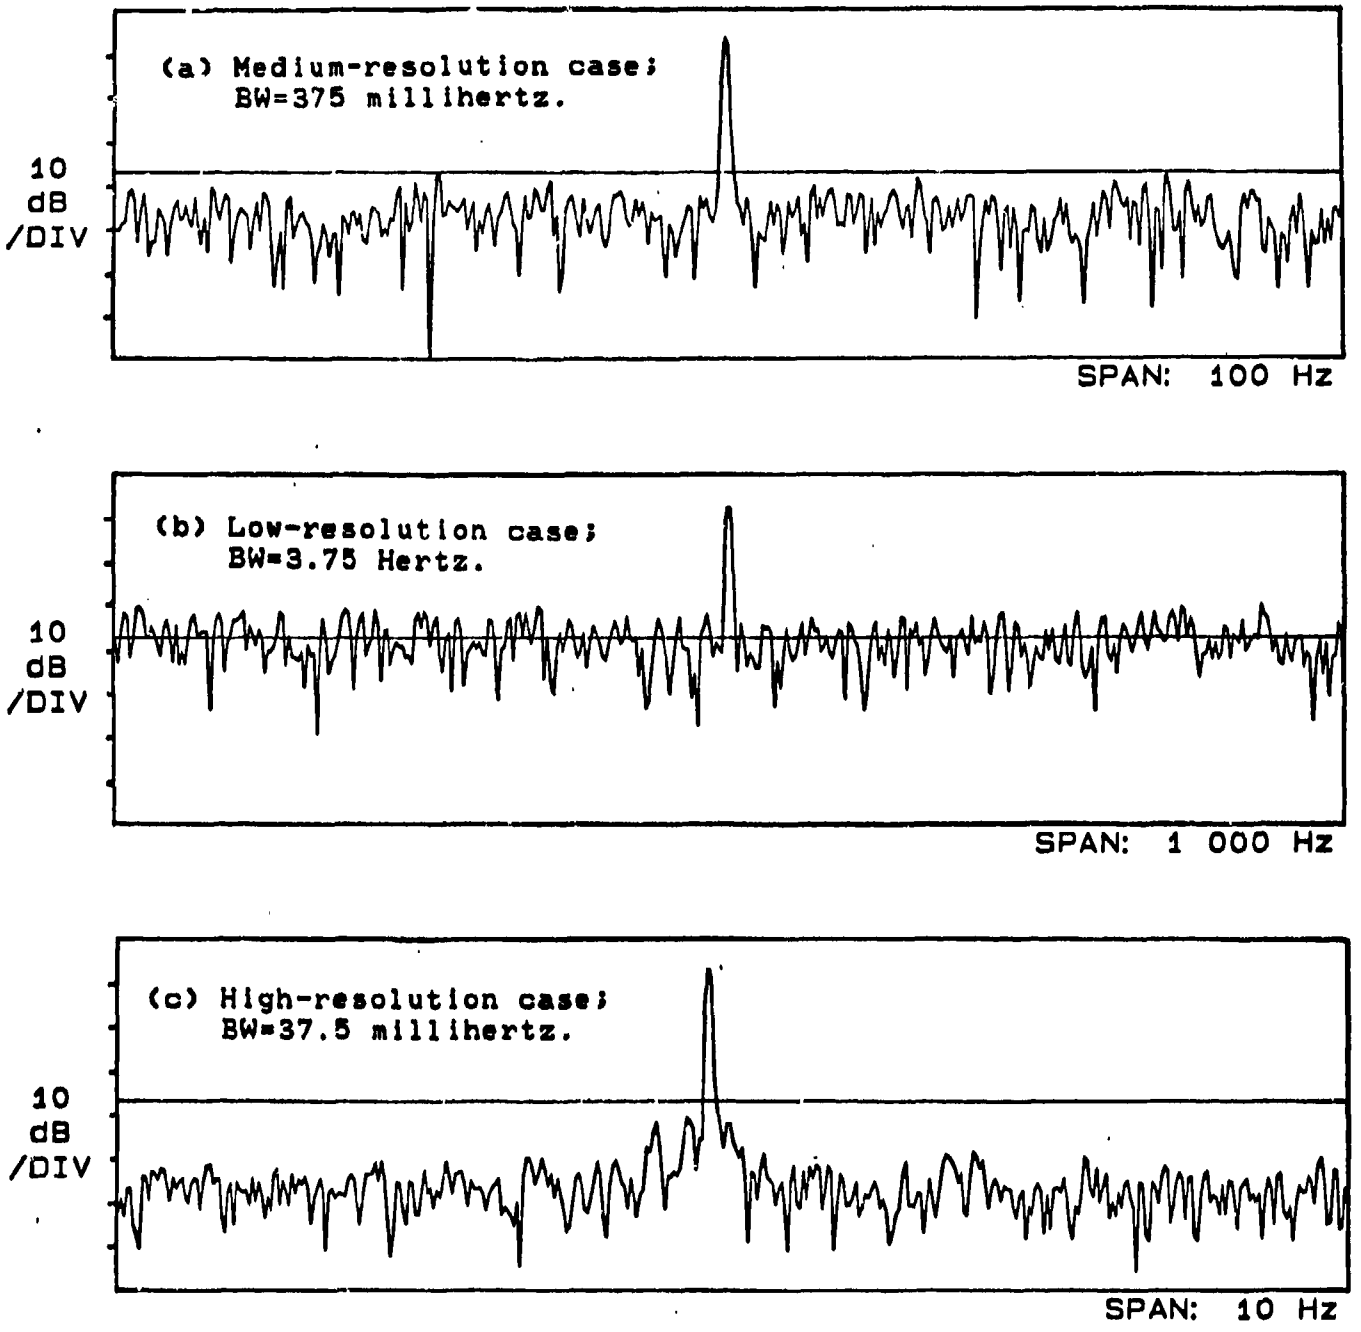

Figure 20. Power spectrum estimates of noisy sinusoid computer for (a) medium-resolution bandwidth, (b) low-resolution bandwidth, and (c) high-resolution bandwidth. Horizontal line in each graph represents identical threshold levels.

added benefit of increasing the detection probability. This demonstrates the benefits of using high resolution when using power spectrum analysis for making detection decisions.

Re-examination of the power spectrum estimates in Figure 20 shows that more reliable detection could be achieved if the variations in the noise portion of the spectrum could be reduced (or smoothed). This would permit the use of lower threshold levels that would result in improved detection probabilities without sacrificing the false-alarm probability. Smoothing can be achieved by averaging the results of power spectrum estimates of multiple samples of data. Since the noise in the individual spectra is uncorrelated, variations in the averaged spectrum will be reduced while the signal information will be minimally affected.

The benefits of averaging can be demonstrated by comparing the unaveraged power spectrum of a weak sinusoid in noise to two other cases where 10 averages and 100 averages have been used. These results are shown in Figure 21. It can be seen that the predicted smoothing significantly improves the ability to detect the weak signal peak. However, it should be noted that if averaging is to be employed without changing the BW, the required observation time must be multiplied by the number of averages used. If the observation time cannot be increased, averaging can be performed by dividing the original sample sequence into a number of shorter sample sequences. Results of power spectrum analysis performed on the shorter sample sequences could be computed and then averaged to obtain the desired smoothed spectrum. Of course, resolution must be sacrificed if smoothing is to be achieved in this manner.

The possible tradeoffs between BW and averaging present an interesting dilemma. Given a maximum, acceptable observation period, how should the BW and number of averages be determined? The answer depends on the signal being detected. If the signal has an extremely narrow bandwidth, the highest resolution should be selected (i.e., no averaging). If the information bandwidth is finite, the resolution should be selected to match the information bandwidth. If the required resolution can be achieved using less than the maximum permissable observation period, additional power spectrum estimates should be computed and averaging should be performed.

To demonstrate this point, the power spectrum of a weak, noise-modulated sinusoid with a 3-dB bandwidth of approximately 400 millihertz was computed using a BW of 375 millihertz. Although this was an appropriate BW for this

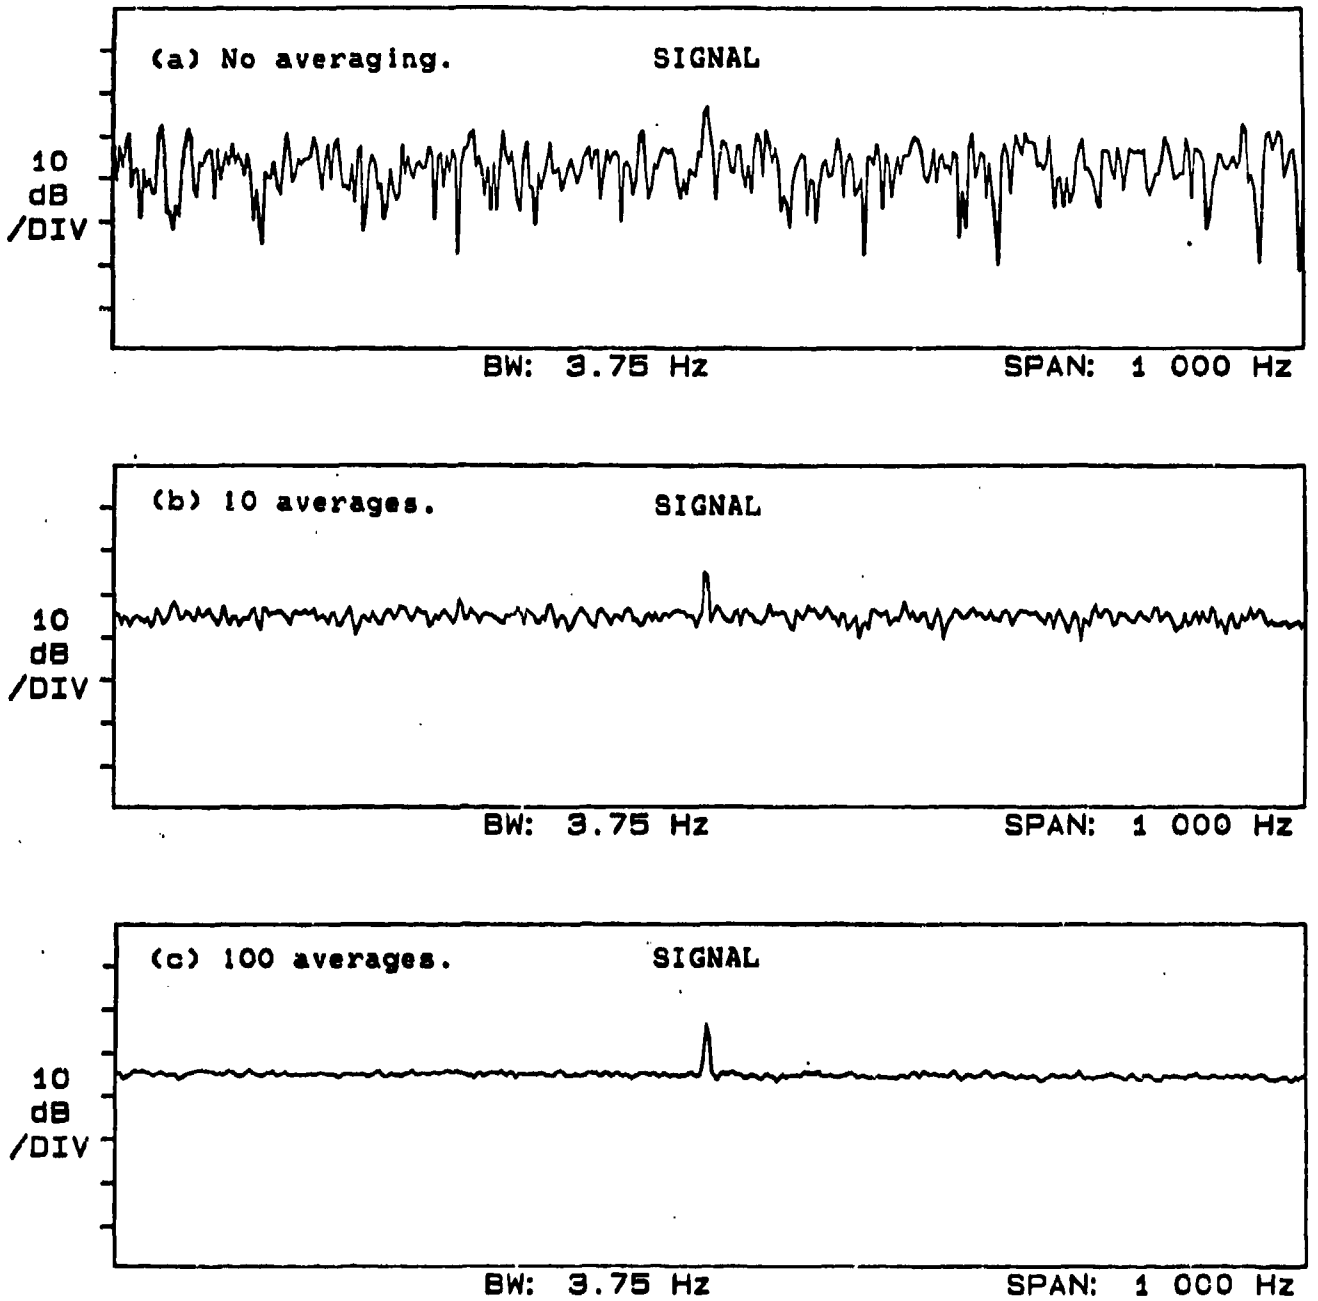

Figure 21. Examples of effects of averaging on noise variations in power spectrum estimates of a weak sinusoidal signal in noise, (a) no averaging, (b) 10 averages, and (c) 100 averages.

case, the signal was difficult to detect because of noise. The spectrum was then recomputed in two different ways, each requiring the observation period to be increased by a factor of ten. In one case, the BW was increased by a factor of ten to 37.5 millihertz. In the second case, the BW was kept at 375 millihertz and the additional observation period was used to compute additional power spectrum estimates which were then averaged to obtain a smoothed power spectrum.

Results of these tests are shown in Figure 22. As predicted, the middle graph in Figure 22 shows that using a resolution substantially narrower than the information bandwidth did not significantly improve the ability to detect the signal peak. However, using the original 375 millihertz resolution plus averaging did significantly improve detection performance, as shown in the bottom graph. This example demonstrates that accurate information about the signal being detected is needed to determine the BW needed to achieve the best detection performance using a minimum observation period.

4. Results of Power Spectrum Analysis of LFD Signals and Noise

During the third-year of this program, a Hewlett-Packard Model 3561A Dynamic Signal Analyzer was used to perform power spectrum analysis on signals and noise outputted by the LFD. The data analyzed had previously been recorded on a multichannel FM-tape recorder. This data was recorded on the roof of the building housing our laboratories in the manner described for the data shown in Figures 12-17. The recorded data used in the analysis was collected for three different conditions. For the first set of conditions (referred to as the test case), the LFD was aimed at a subject lying 30 meters from the LFD. For the second set of conditions (referred to as the baseline case), the subject was removed and the LFD was aimed at the previously occupied target location to provide a measure of the clutter level. The LFD was then aimed at the empty sky to obtain samples of clutter-free data for the third set of conditons (referred to as the sky-noise case). In all cases, data was passed through 0.05-10 Hz bandpass filters prior to recording (the 10 Hz upper cutoff frequency will be increased in future tests to at least 20 Hz).

In the initial experiment, the power spectrum of recorded baseline data was compared the power spectrum of recorded sky-noise data. The purpose of

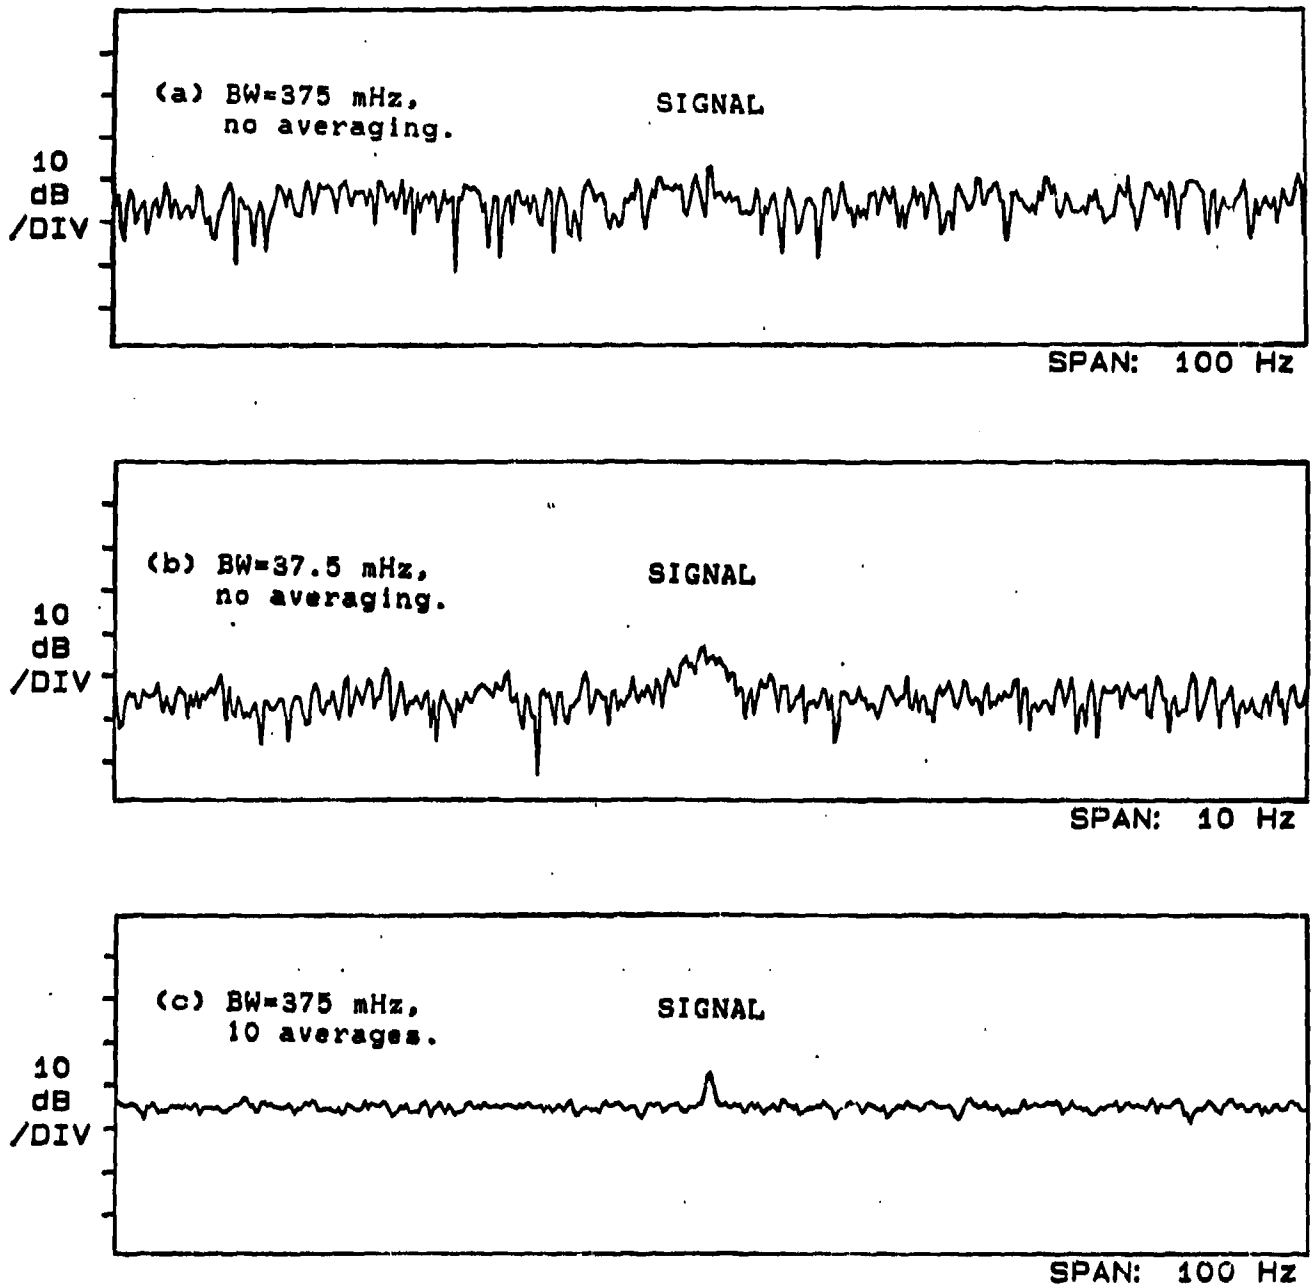

Figure 22. Graphs demonstrating tradeoffs between using an excessively-high resolution bandwidth and using the proper resolution bandwidth and averaging.

this experiment was to verify the existence of clutter-related noise in the baseline data. It was suspected that because of range-sidelobes associated with the current LFD, clutter from a large stand of trees approximately 30-40 meters beyond the target location would be evident. Comparison of the baseline power spectrum to the clutter-free, sky-noise spectrum verified that clutter was contaminating the baseline data as shown in Figure 23.

The results in Figure 23 show that even with range-gating, clutter from the distant trees is degrading the SNR by 10-20 db in the frequency band of interest. This loss in SNR severely decreases the possible detection range and again points out the need for enhancing the current range-gating capability of the LFD. The clutter is seen to have a large low-frequency component. This component was possibly due to the relatively slow, swaying motion of large tree branches and whole trees. At frequencies between 1-10 Hz, the baseline spectrum is approximately white (i.e., flat), although clutter is still present since the baseline spectrum is above the sky-noise spectrum. These higher frequency effects are probably due to the more rapid motion of tree leaves. The fact that noise due to clutter is approximately white in the 1-10 Hz band represents useful information since a large part of the cardiac information detected by the LFD also appears to be concentrated in this band. The results in Figure 23 represent a small fraction of the studies that need to be performed on clutter-related noise. Additional data is needed for a variety of clutter sources under a variety of conditions.

In the next experiment, the power spectrum of recorded baseline data was compared to the power spectrum of recorded test data. This comparison was made to identify characteristics of the test data that could provide a basis for target detection based on power spectrum analysis. A test case with a high SNR (i.e., respiratory and cardiac information observable in the time-waveforms) was selected to insure that respiratory and cardiac components were present in the computed spectrum. Identical resolution bandwidths and averaging were used for both the baseline and test data. Results of this experiment are shown in Figure 24. The following observations can be made about these results.

In general, the spectrum of the test data is well above the baseline spectrum. This was expected since the SNR was known to be high. Therefore, one detection possibility would be to compare the total energy in the test data spectrum to a suitable threshold level. This would be equivalent to the

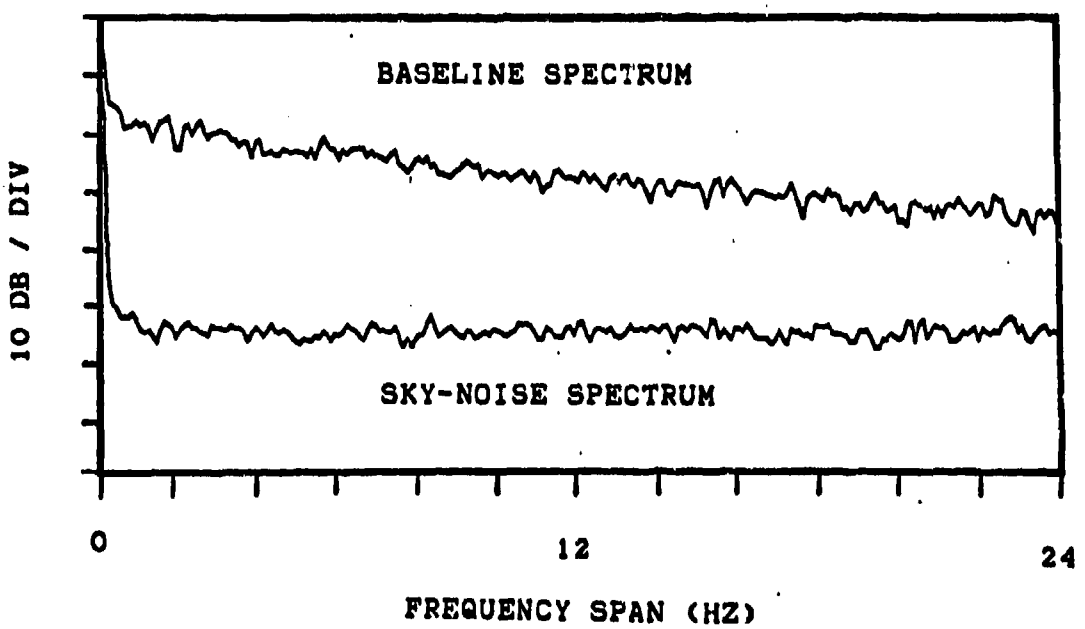

Figure 23. Comparison of the power spectrum estimates of recorded baseline data containing clutter-related effects and recorded sky-noise which is essentially clutter-free.

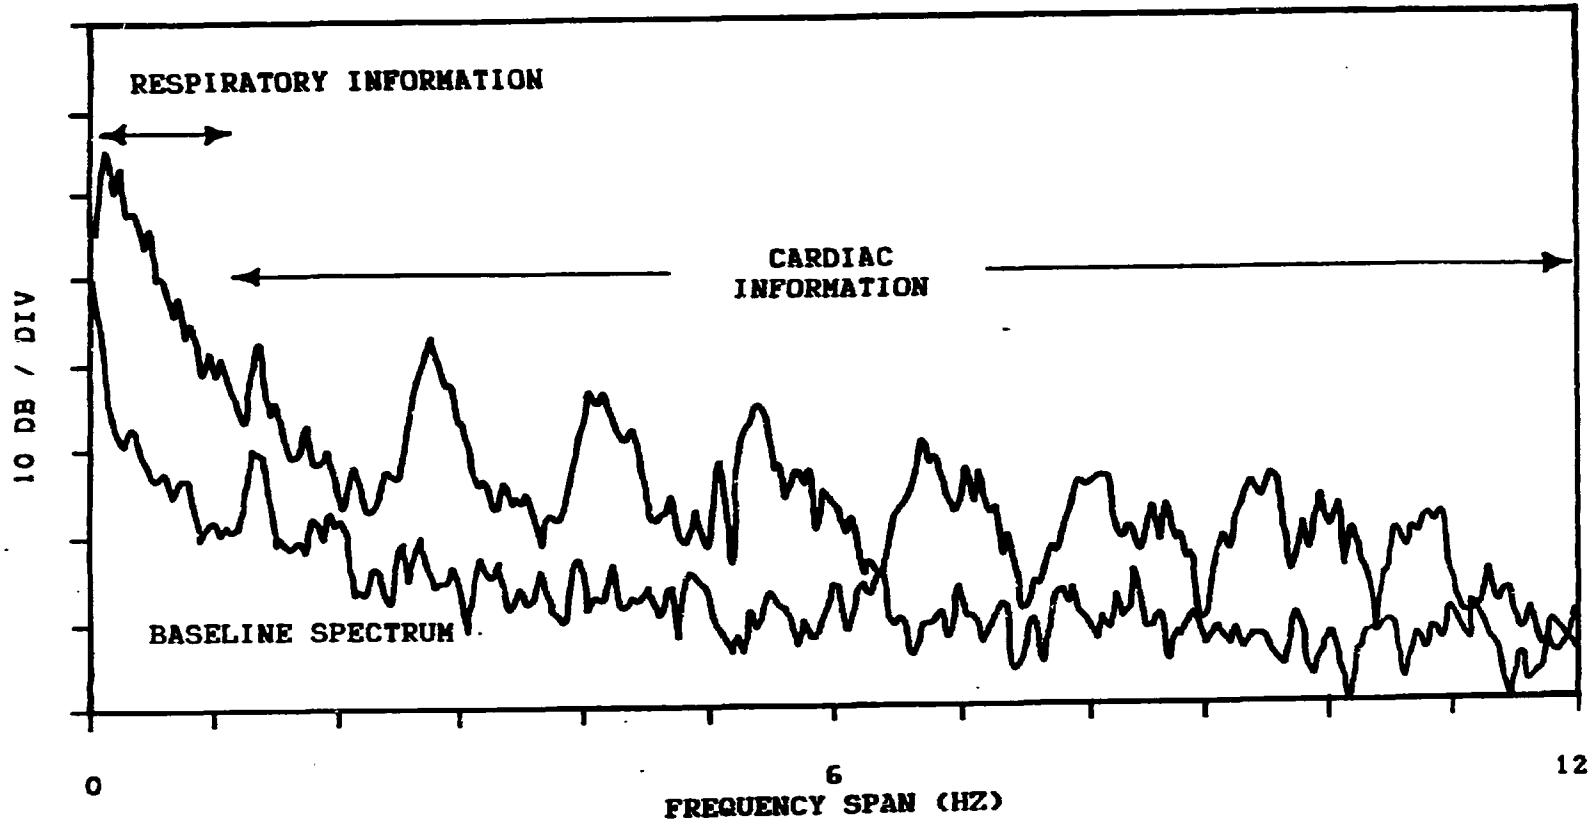

Figure 24. Comparison of results of power spectrum analysis on recorded baseline data and recorded test data for subject at target range of 30 meters.

simple energy detector approach described earlier in this report.

It also may be noted that there are several periodic components present in the test data. Components below one Hz are mainly due to respiratory motion while the higher frequency components can be attributed to cardiac motion. The detected respiratory motion does not have significant harmonic structure. This coincides with the sinusoid-like motion normally associated with respiratory activity. The cardiac information has a pronounced harmonic structure which is indicative of the more impulse-like motion associated with cardiac activity. Even with its more pronounced harmonic structure, the cardiac information in Figure 24 is significantly weaker than the respiratory information. The strongest cardiac component is approximately 20 dB (or a factor of 100) lower than the main respiratory component. Its lower strength and wider frequency distribution explains why it is difficult to observe cardiac information in the time-waveforms outputted by the LFD.

Because of the relative weakness of the cardiac information, it was judged that power spectrum analysis should be optimized for detecting cardiac information. This might result in some loss in the ability to detect respiratory motion (because of the use of a non-ideal BW), but any loss should be offset by the greater strength of the respiratory information. To determine an appropriate BW to use for detecting cardiac information, power spectrum analysis was performed on standard EKG recordings from five different test subject. The 3-dB bandwidths of the fundamental lobe in the power spectrum of the different EKG's was found to range from approximately 0.02-0.11 Hz (these results are shown in Figure 25). Variations of the higher-order harmonics can be obtained by multiplying by the appropriate harmonic number. Because of the wide range in the observed information bandwidths, some compromise was necessary in selecting an appropriate BW. Experimentation with different resolutions indicated the ideal BW can be increased or decreased by several factors without severely degrading performance. Therefore, a compromise BW of 0.150 Hz was selected. This BW has a corresponding observation period of 10 seconds (the BW and observation period are not exactly inversely related because of windowing employed by the spectrum analyzer). This BW appears to be adequate. However, the suitability of even lower BW's resolutions will be investigated to determine if the required observation times can be reduced. Reduced observation times would be especially important in situations where a large number of averages must be used to achieve reliable detection performance.

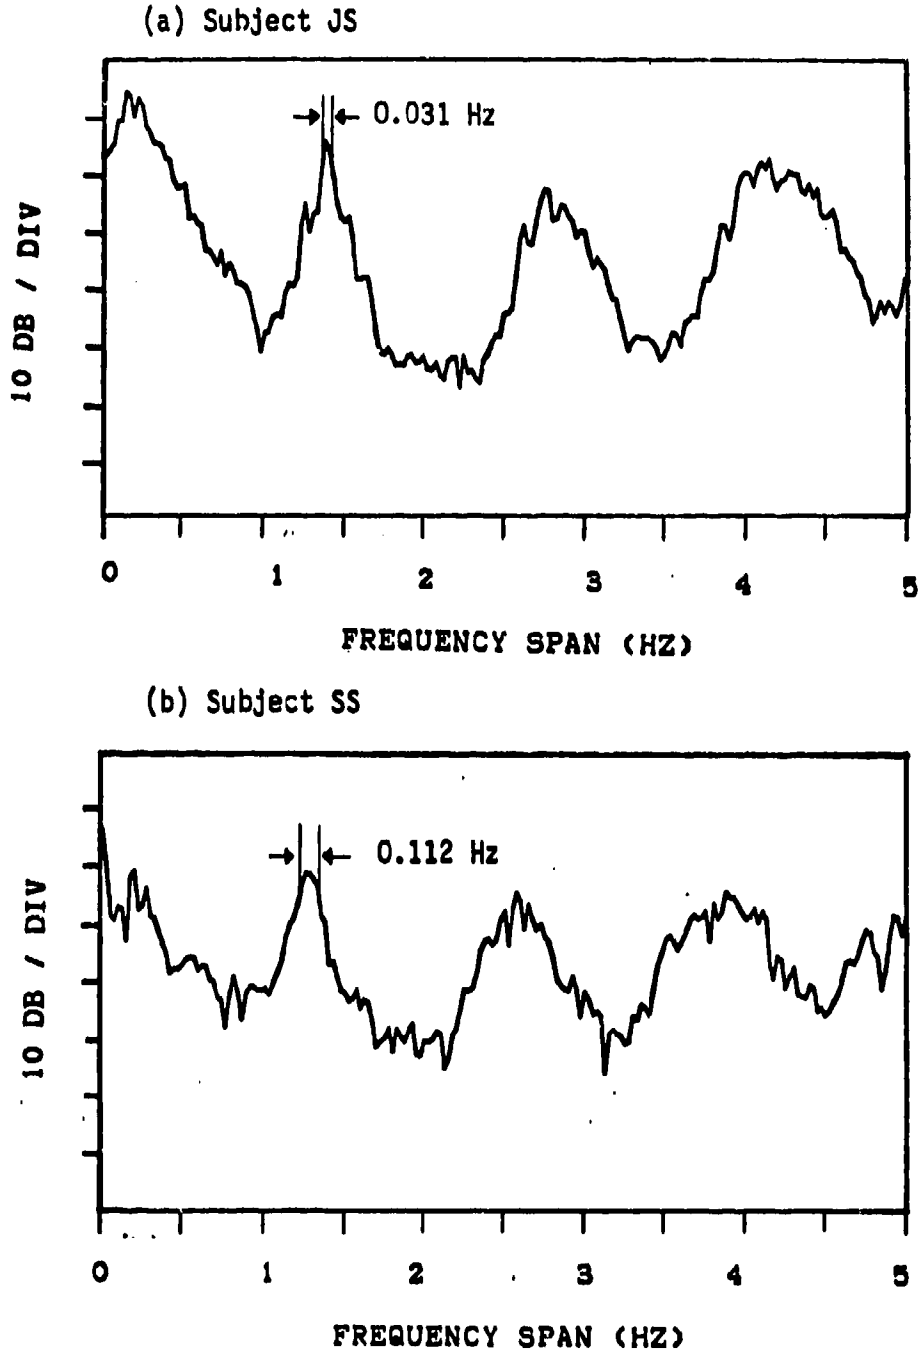

Figure 25ab. Results of power spectrum analysis of EKG data from five different test subjects showing approximate information bandwidths (as defined by 3-dB points) in fundamental cardiac component: (a) subject JS, (b) subject SS.

(c) Subject RV

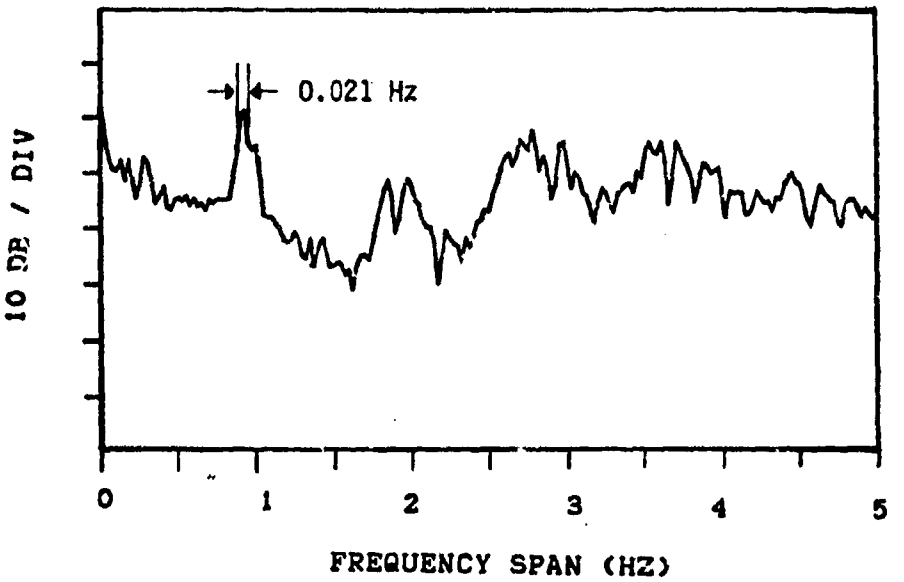

(d) Subject DC

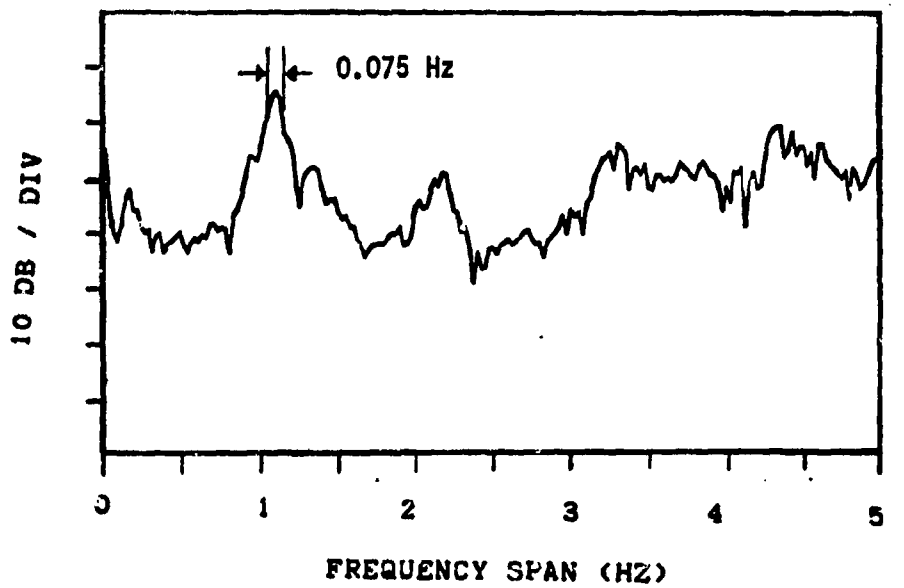

Figure 25cd. Results of power spectrum analysis of EKG data from different test subjects showing approximate information bandwidth (as defined by 3-dB points) in fundamental cardiac component: (c) subject RV, (d) subject DC.

(e) Subject SC

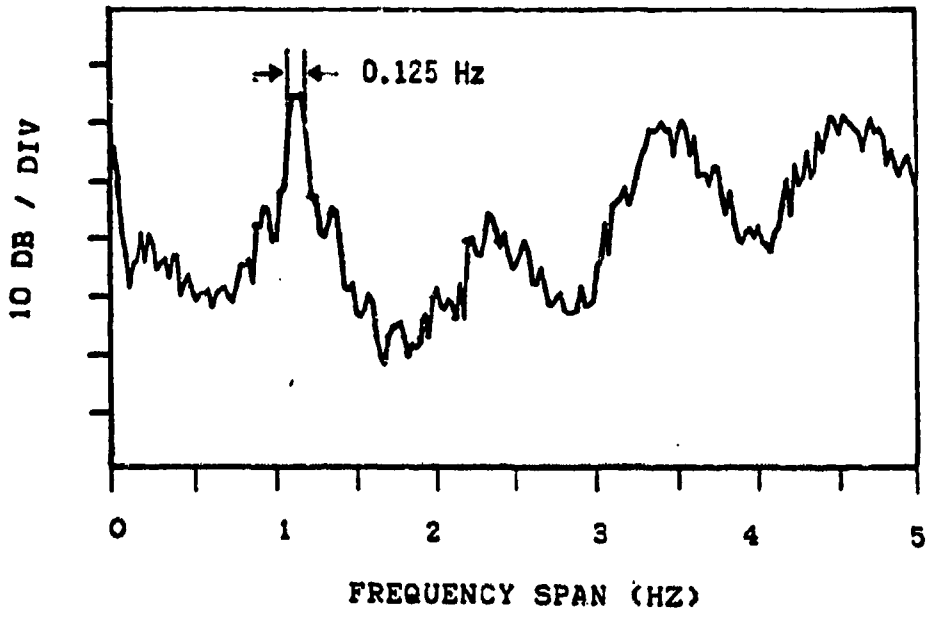

Figure 25e. Results of power spectrum analysis of EKG data from different test subjects showing approximate information bandwidth (as defind by 3-dB points) in fundamental cardiac component: (e) subject SC.

Once an appropriate BW had been selected, the spectrum analyzer was used to analyze test data recorded from the LFD for five different subjects. The target range was 30 meters. Three to four minutes prior to recording each set of test data, baseline data was recorded to establish a suitable noise reference. For each subject, the SNR was sufficient that respiratory information, but not weaker cardiac information, could be clearly observed in the time-waveforms. Identical BW's of 0.15 Hz and the same number of averages were used for analysis of the test data and baseline data. The observation period for each case was 40 seconds.

Results of the power spectrum analysis are shown in the five graphs in Figure 26. Comparison of the baseline spectrum to the test data spectrum show the existence of periodic components corresponding to respiratory and cardiac motions for all five test subject. As noted previously, detection could have been based on energy detection or on techniques that compared individual components in the test-data spectrum to suitable threshold levels. If the baseline spectrum in each of the five graphs is imagined to represent the detection threshold, reliable detection of casualty information would be possible for all five subjects.

The frequency of the respiratory component observed in each of the graphs was found to correspond to the respiratory rates observed in the corresponding time-waveforms. Verification of the periodic components attributed to cardiac motion could not be done from the time-waveforms since cardiac motion was too weak to be observed. To insure that the periodic components seen in the spectrum of each set of test data was indeed due to cardiac motion, the power spectrum of one of the sets of test data was compared to the spectrum of the simultaneous EKG record. The test data spectrum and the EKG spectrum were not expected to have the same overall shape (since one is due to mechanical events while the other is due to electrical events). However, the information in the two spectra was expected to appear at identical frequencies. The results shown in Figure 27 verify that periodic components observed from 1-10 Hz in the previous test data spectra were, in fact, related to cardiac events and not simply some type of periodic clutter.

A portion of the results in Figure 27 can be used to demonstrate another possible advantage of power spectrum analysis. The lower two graphs in this figure represent results of power spectrum analysis performed on the individual quadrature channels of the LFD. The differences apparent in these two graphs is a result of the range-deadspot problem discussed previously. Good detection will require testing both channels. Experimentation indicates

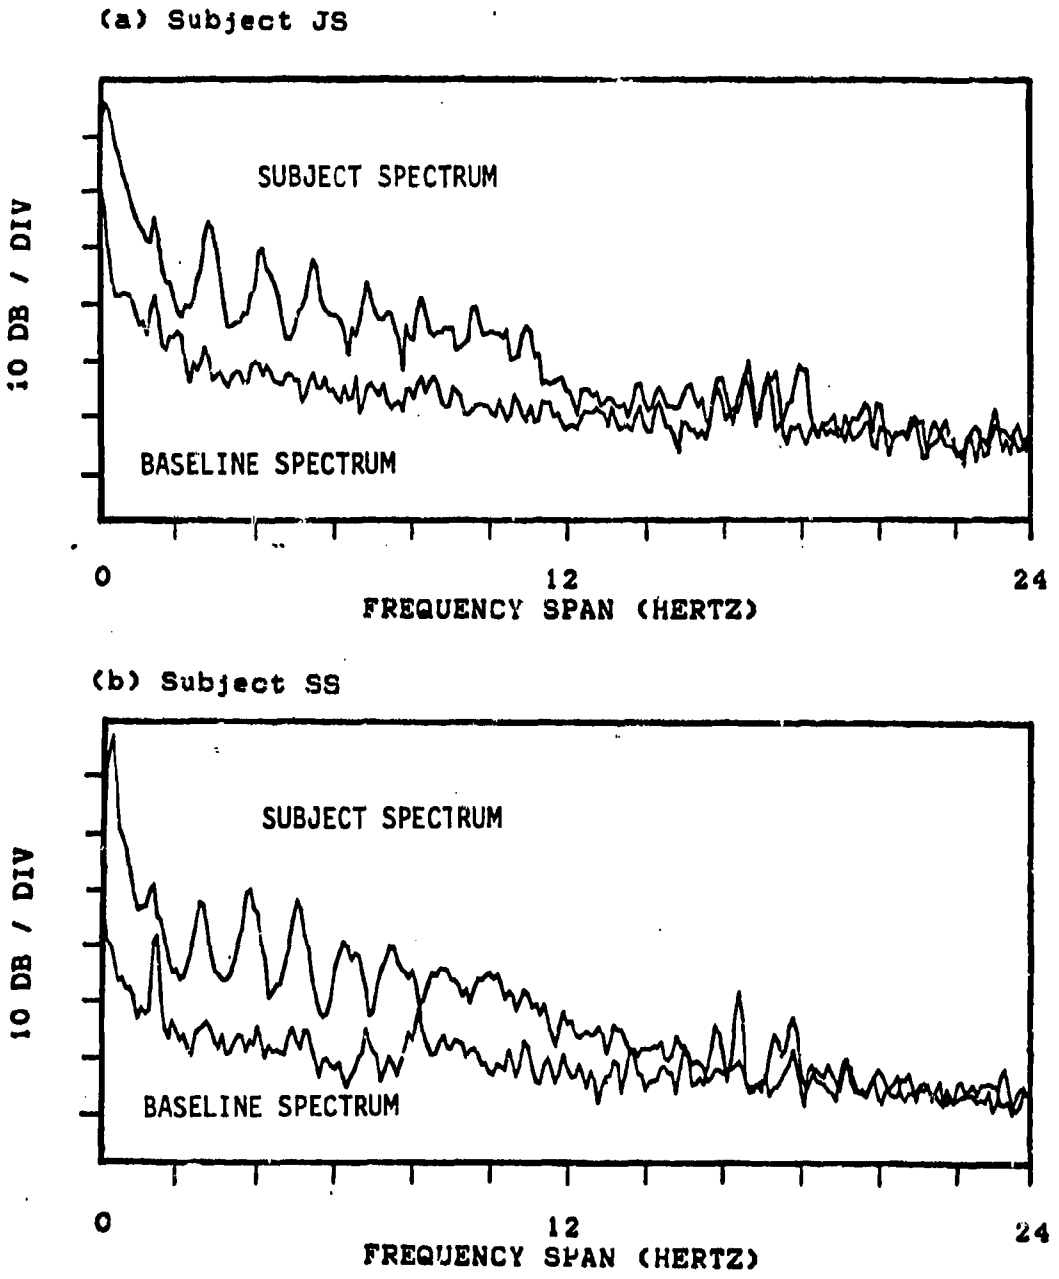

Figure 26ab. Comparison of power spectrum estimates of recorded test data to baseline data for different test subjects: subject JS, (b) subject SS. Target range was 30 meters and observation period was 40 seconds.

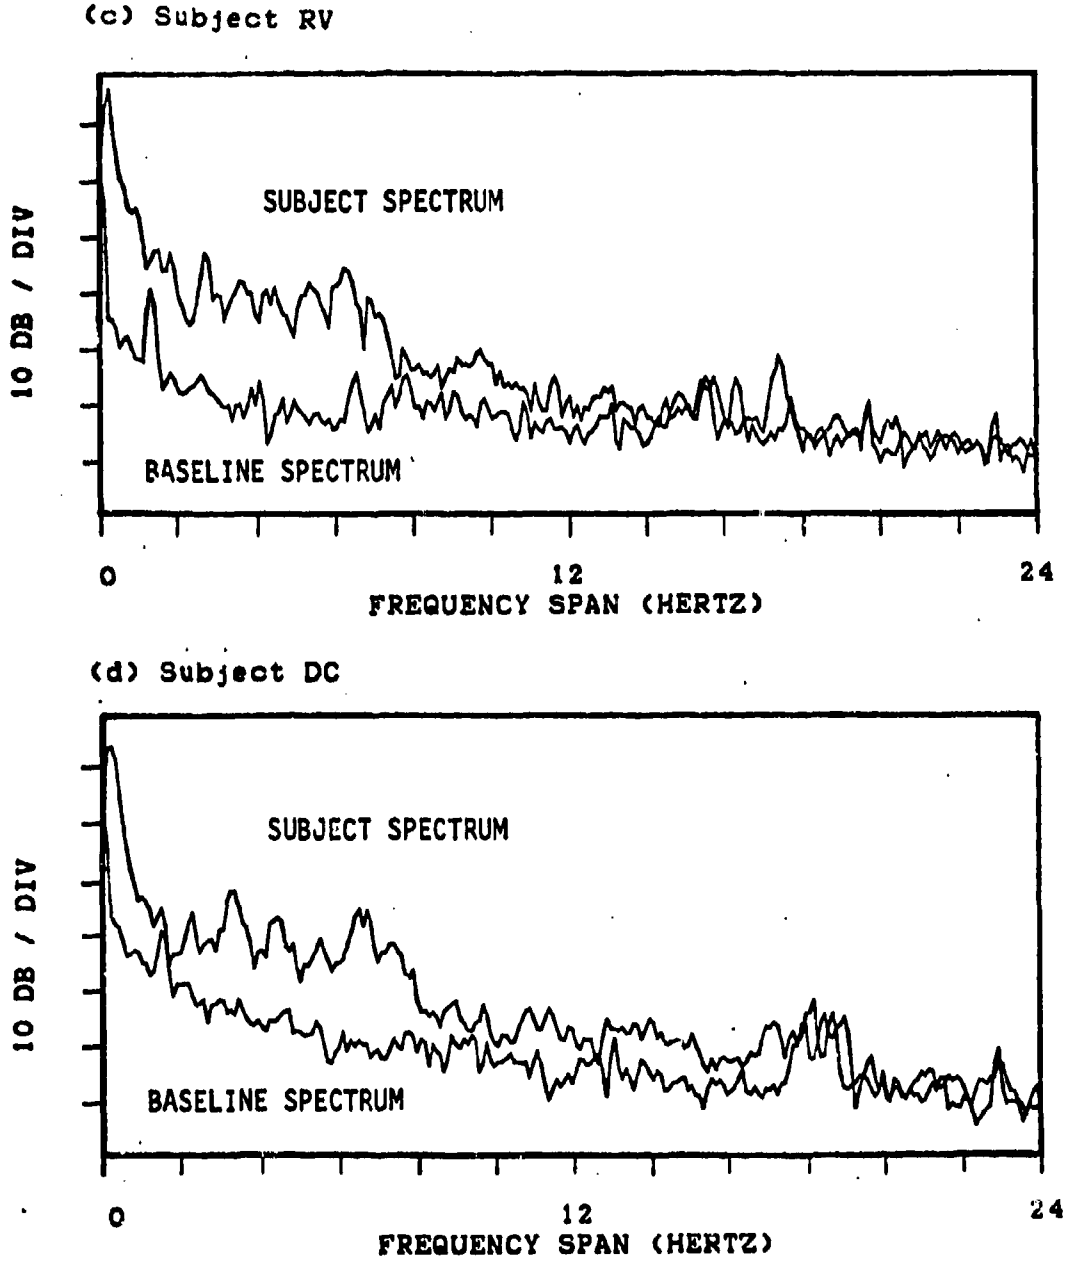

Figure 26cd. Comparison of power spectrum estimates of recorded test data to baseline data for different test subjects: subject RV, (d) subject DC. Target range was 30 meters and observation period was 40 seconds.

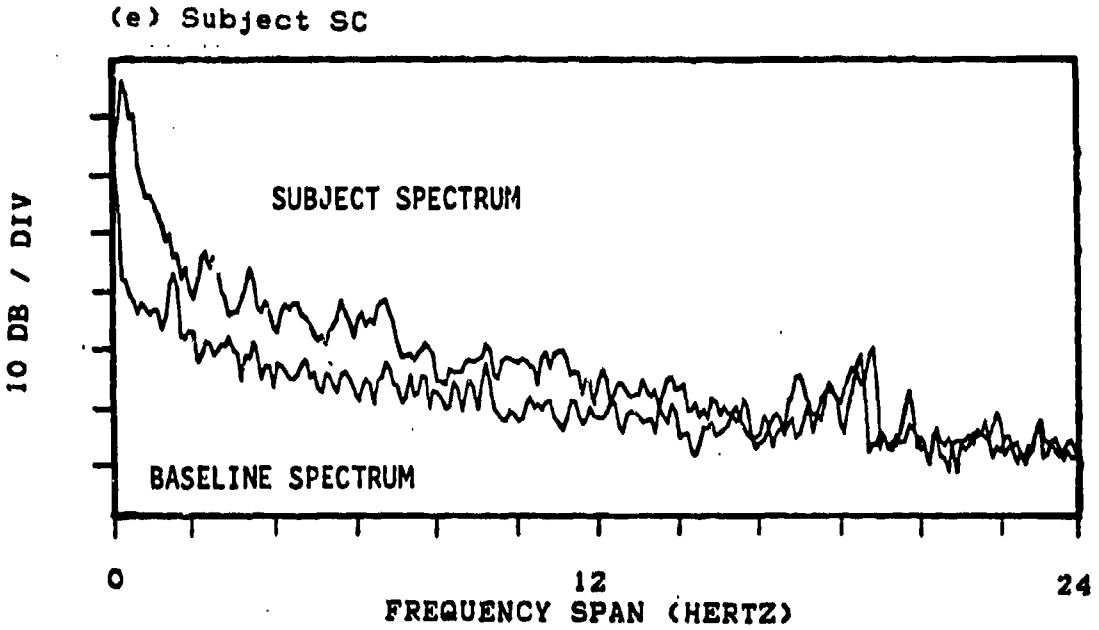

Figure 26e. Comparison of power spectrum estimates of recorded test data to baseline data for different test subjects: (e) subject SC. Target range was 30 meters and observation period was 40 seconds.

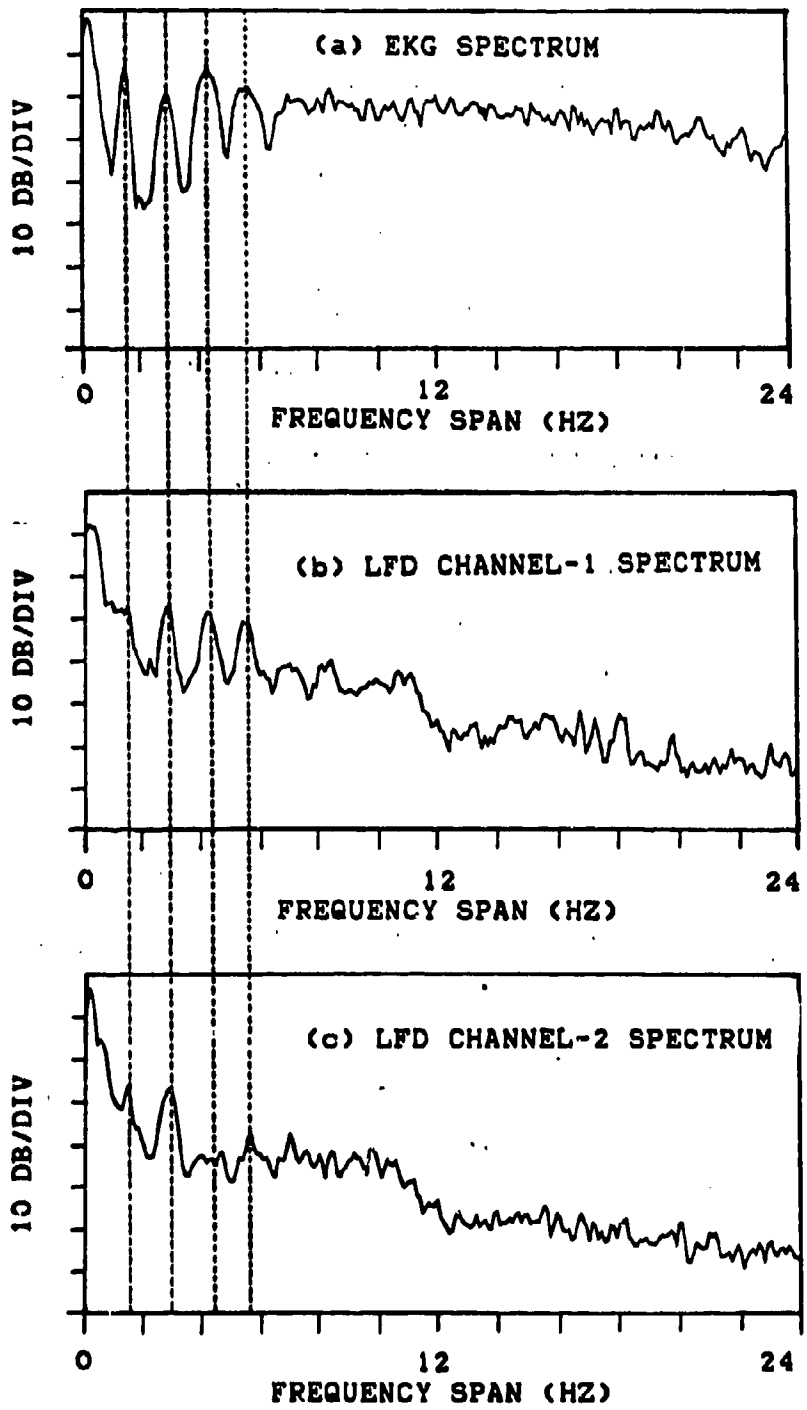

Figure 27. Comparison of power spectrum estimates of test data to corresponding EKG spectrum to verify existence of suspected cardiac information.

this is not an obstacle to using power spectrum analysis since results obtained for the individual channels may be averaged to produce a single spectrum that contains all the target information. When the two channels in this example case are combined, the result presented in Figure 26a was obtained. Comparison of this figure with the lower two graphs in Figure 27 indicates power spectrum analysis provides a convenient method for combining the information contained in the quadrature channels of the LFD.

The final example in this discussion involves power spectrum analysis of test data taken from a longer target range. Re-examination of the data in Figure 26 indicates the SNR was relatively high in each of the previous five test cases. Therefore, it is probable the target information could have been reliably detected from the time-waveforms without a significant amount of signal processing. To observe the impact of signal processing when the SNR was poorer, test data was recorded from a longer target range of 45 meters. Because of a weaker return signal and a higher clutter level, the ability to detect the presence of respiratory or cardiac information in the time-waveform was marginal. This is evident from the data in Figure 28a which represent the time-waveforms that were outputted by the quadrature channels of the LFD for this case.

Power spectrum analysis was subsequently performed on the test data and corresponding baseline data for this case. To minimize the effects of noise in these estimates, a BW of 0.0375 Hz was used. In addition, six averages were employed which resulted in a total observation period of 120 seconds. When the test data spectrum was compared to the spectrum of the corresponding baseline data, the presence of both respiratory and cardiac information was evident. This result is shown in Figure 28b.

It may be noted that only a small fraciton of the test-data spectrum in Figure 28b is significantly stronger than the baseline spectrum, unlike the results obtained in Figure 26 when the SNR was appreciably higher. Therefore, a detection approach such as the simple energy detector might not provide very reliable detection in this situation. Conversely, it appears that using a technique such as power spectrum analysis, which provides the resolution necessary to observe the presence of narrow-band components, would make it possible to reliably detect the presence of weak, but periodic casualty information. Of course, this ability is dependent on the casualty information remaining constant over an observation period sufficient to permit the required BW and an effective amount of averaging to be used.

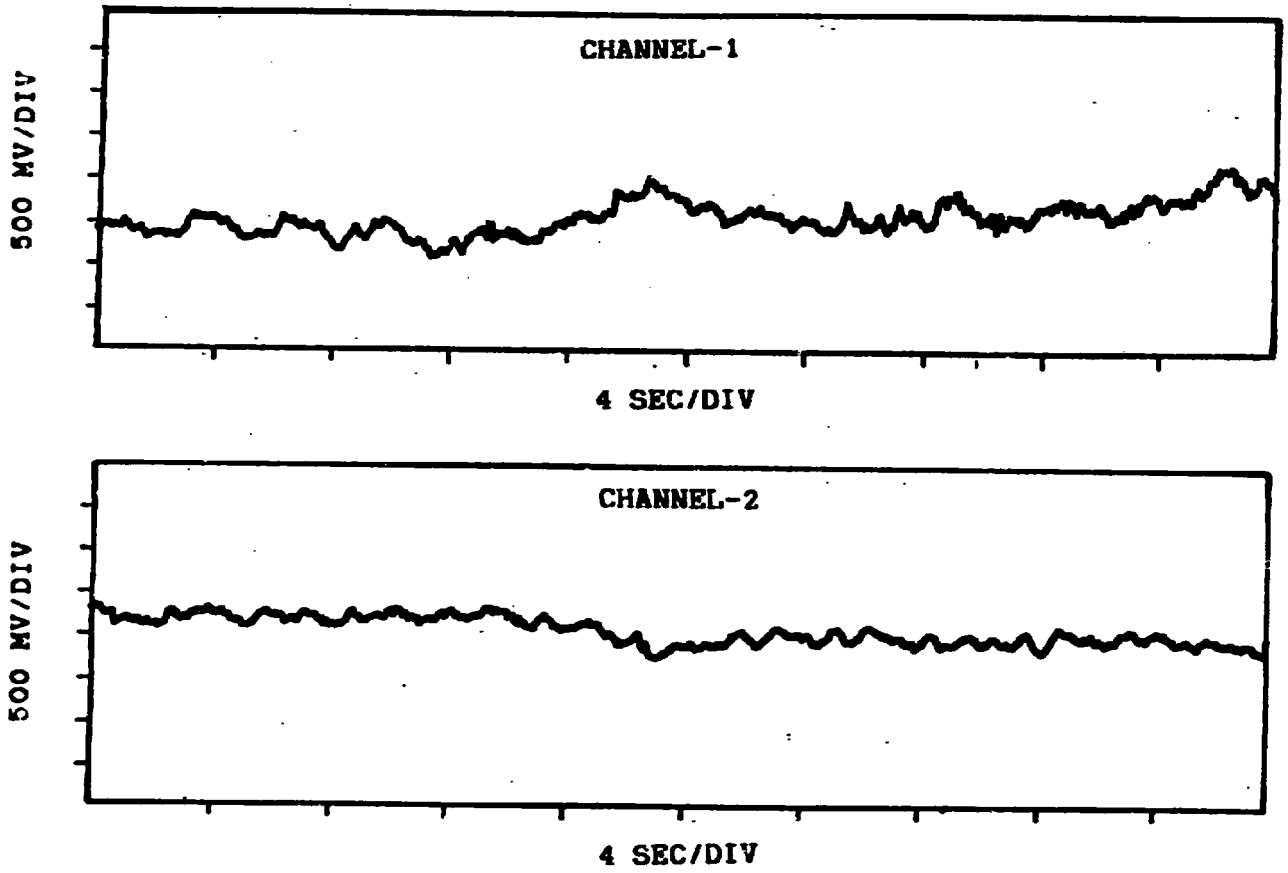

Figure 28a. Time-waveforms outputted by LFD for subject at target range of 45 meters.

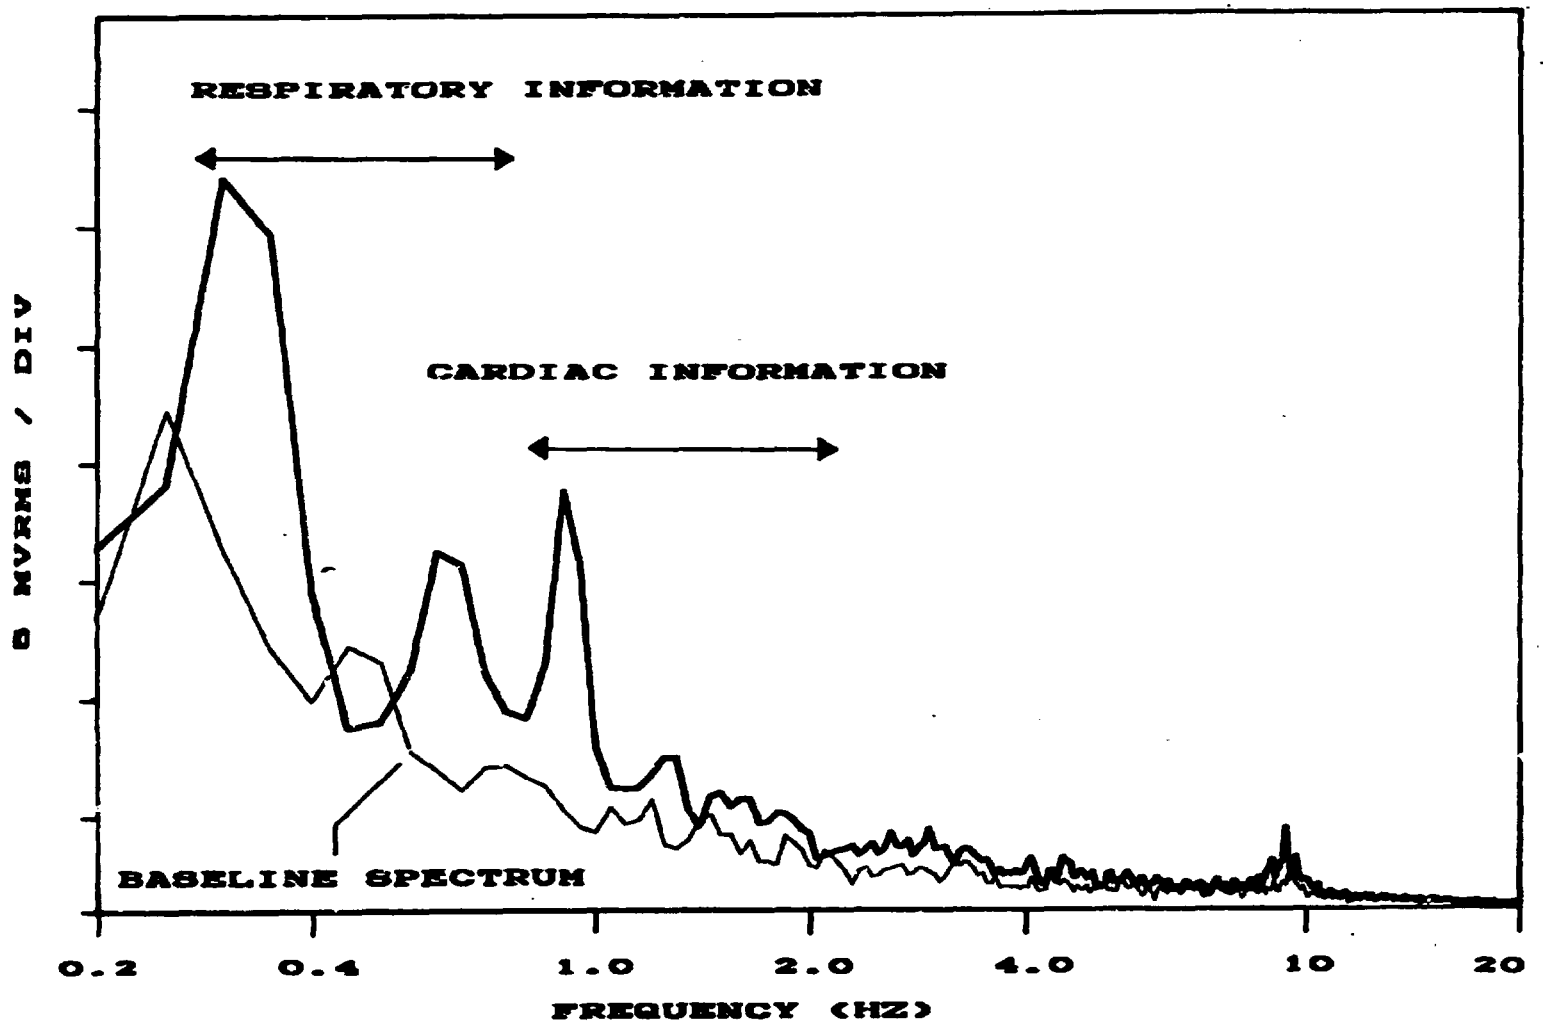

Figure 28b. Comparison of results of power spectrum analysis on time-waveforms in Figure 28a to corresponding baseline spectrum.

5. Sequential Detector Based on Power Spectrum Estimates

The preceding two parts of this report have shown that target detection can be achieved using power spectrum analysis. If the requied power spectrum estimates were computed with a microprocessor-based system, each point in the computed spectrum could be stored as an individual value in the microprocessor's memory. Detection could be achieved by comparing each of the stored values to a suitable threshold level. It has been seen that the best detection can be achieved by using narrow BW's and averaging multiple sets of data. Each of these operations requires that the observation period be increased.

The longer observation periods required to achieve high detection performance is acceptable provided a weak signal is being detected or the noise is excessively high. However, there will be times when high detector performance is not really needed to reliably make the correct detection decision. At these times, the use of narrow BW's and averaging would result in excessively long observation periods. One such instance would be when target information is present and is very strong in comparison to the noise. In this situation, optimum detection should not be required to reliably determine that target information is present. Another example would be a situation in which there is not any target information present and the noise is very low. In this situation, optimum detection should not be required to determine that target information is not present, since if it were present, it could easily be detected because of the low noise. In each of these example situations, the use of long observation periods would be a waste of valuable time.

It appears that a choice might be necessary between using long observation periods that enable high detection performance but will sometimes result in wasted observation time, or using short observation periods that result in rapid detection decisions but may not provide very reliable detection performance when the SNR is low or the clutter is high. However, an alternative approach exists. This alternate approach is the Sequential Observer (or sequential detector)[1].

A sequential detector would normally operate in a mode in which a moderate observation period would be used and decisions would be quickly made. Three decisions would be possible: (1) "yes", (2) "no", or (3) "maybe". The sequential detector would use two different threshold levels. The upper

threshold would be set high enough that it could be exceeded by strong target signals, but not by noise or weak target signals. If this threshold is ever exceeded, the sequential detector can quickly decide that a target is present. The lower threshold level would be set so that it would be exceeded by even the weakest target signal. If this threshold was not exceeded, the sequential detector would very quickly decide that a target was not present. Thus the sequential detector could very quickly make a yes/no decision for these two special cases.

It is apparent that a situation exists in which the lower threshold is exceeded (so a "no" decision cannot be made) but the upper threshold is not exceeded (so a "yes" decision cannot be made). In this situation, the sequential detector makes the third possible decision of "maybe" (or "don't know"). If a "maybe" decision was made, the sequential detector would switch from its normal operating mode to a second mode that permits higher detection performance to be achieved. This high-performance mode could be implemented in several ways. In each case the separation between the two thresholds could be decreased to reduce the detection uncertainty.

One possibility for switching the sequential detector from its normal mode to a high-performance mode would be to use a longer observation period and an increased BW. Another possibility would be to use the same BW but compute one or more additional power spectra that could be averaged with the original spectrum to produce a smoother spectrum that would enable a more reliable yes/no detection decision to be made. The sequential detector could be designed to continually improve its performance until it was finally able to produce a yes/no decision (or until reset by the operator), or it could be designed to stop once it had reached a maximum observation period, regardless of whether it was ever able to make a yes/no decision.

The sequential detection approach is not commonly used in traditional radar applications because it is not always practical to have a radar dwell on a given target or range-cell for an extended period of time. However, this would not be a problem in the LFD since the target is both stationary and at a known range. Thus, the sequential detection approach appears to be extremely well-suited for use in the LFD and will be investigated during the design of the target detection system for the LFD.

6. Establishing Suitable Detection Threshold Levels

Whether sequential detection or some other detection approach is employed in the LFD, it will be necessary to establish suitable threshold levels for basing detection decisions. If sufficient knowledge can be gained

about the signals and noise encountered by the LFD, it is plausible that the threshold level could be set once the target range and general clutter conditions were known. However, because of the wide range of possible signal and noise levels, it is likely the threshold level will have to be adjusted for each specific situation.

One method for establishing threshold levels would be to make use of information in the target power spectrum. One such possibility would be to use a portion of the spectrum that is known to be free of target information. For example, if future testing revealed that respiratory and cardiac information is limited to frequencies below 20 Hz, portions of the spectrum above 20 Hz could be used to establish the required threshold. This method could prove especially effective in frequency bands where the clutter spectrum is approximately white. Recall that the baseline spectrum was approximately white for frequencies above 1 Hz in several of the preceding examples (e.g., Figures 24 and 26).

Another possible approach would be to measure the clutter at a location near the target then use the target-free power spectrum for the detection threshold. However, this approach has several potential problems. One problem is that an additional observation period is required for computation of the power spectrum to be used for the threshold. This approach also assumes the clutter will be relatively constant as the aim of the LFD is changed from the target-free location to the target. Additional data is needed to determine if such an assumption is valid.

A final approach that is being considered for setting the threshold level is based on the improved range-gating capability planned for the LFD. In this final approach, the power spectrum of the clutter in individual range cells around the target would be computed, then averaged to establish the required threshold level. This approach would not require aiming the LFD at any positions other than the target and a single observation period could be used to collect all the data needed to compute the required power spectrum estimates. An added advantage of this approach is that since the individual power spectrum of each range-cell could be compared to the threshold once it was determined, the exact target range is not needed. Because of the emphasis planned for development of an improved range-gating capability, it is anticipated that this possible approach for establishing a threshold level will be seriously considered during future research efforts.

D. Conclusions

The preceding discussion of the RF and signal processing systems of the LFD has shown that good target detection requires good receiver performance as well as adequate signal processing. A major limitation of the current RF system is range-sidelobes that make it difficult to reject the effects of clutter from large or distributed clutter sources and can seriously degrade receiver performance. Plans have been developed to improve the current range-gating system by increasing the deviation used during modulation and by using an appropriate weighting scheme following demodulation of the target return signal. These improvements should result in a range-gating system capable of producing narrower range-cell sizes and reduced range-sidelobe levels.

Poor receiver performance as a result of high-clutter levels or weak return signals can be partially offset by using appropriate signal processing. However, improving detection performance through signal processing requires the use of increased observation times, and specific knowledge about the signals and noise is needed. Therefore, signal processing cannot be used to compensate for poor range-gating performance or other receiver inadequacies unless longer observation times are acceptable and sufficient information is available about the specific signals and noise to permit an appropriate detector to be derived.

One promising signal processing approach that has been investigated on this program is based on the use of power spectrum analysis. Power spectrum analysis makes it possible to detect the presence of narrowband signals such as the periodic components associated with respiratory and cardiac motion and also provides a convenient method for combining the information in the LFD's quadrature output channels. An added advantage of this approach is that power spectrum analysis can be performed with a conveniently implemented microprocessor-based system interfaced with a suitable FFT-processor. Achieving good power spectrum analysis requires using appropriate resolution bandwidths and averaging of multiple spectra. To insure optimum detection performance and minimum observation times with this approach, spectral information about the signals and noise is required.

Although intuitive arguments have been used to justify using a signal processing procedure based on power spectrum analysis, a rigorous detection theory analysis indicates that such an approach is optimum or near-optimum for cases similar to those encountered with the LFD. The exact detection

performance of such a system (with respect to a classical optimum detector) for the LFD has not yet been determined since the comparison depends on even more precise knowledge of the signal and noise than has been obtained. However, preliminary results indicate that in the case of the LFD, the spectral analysis detector is extremely close (within a very few decibels) to the optimum detector. Nevertheless, in the interest of achieving the best possible performance, alternative approaches are under consideration.

One such procedure developed by Kailath [33] uses a recent implementation of an optimum detector which is much more efficient numerically than traditional methods, although not as efficient as the (sub-optimum) spectral analysis method. The required calculations can be performed quickly with compact, economical digital processors; however power drain can become a serious problem in a portable instrument, particularly if many range cells must be examined. Thus, a larger and heavier power source (battery) would be required. This additional size and weight could instead be used to accomodate a larger antenna, which might improve overall performance. Such tradeoffs are currently under investigation.

Another interesting approach involves adaptive processors [34-37], which attempt to adjust themselves to the prevailing signal and noise characteristics as estimated from the incoming data. Under ideal conditions, such an approach might be expected to approach the SKE case, assuming that the learning phase is perfect. However, in practice, the learning phase is not perfect since it is based on noisy estimates obtained from limited data records. In addition, such processors also require some prior signal and noise assumption to ensure that they adapt to the desired signal and not to the noise. Nevertheless, good performance is claimed for such systems in certain cases [34,36,37]. A wide variety of adaptive algorithms of varying degrees of complexity and claimed performance exist. Some of the more promising algorithms (such as that of Widrow [35]) are under investigation for the LFD.

It is clear from the preceding discussion that for best overall performance, the design of the LFD requires a careful evaluation of the many tradeoffs involved, and, above all, a clear understanding of the signal and noise mechanisms.