Document Title: Demonstration of a Concealed Weapons Detection System Using Electromagnetic Resonances, Final Report

Author(s): AKELA

Document No.: 190134

Date Received: September 10, 2001

Award Number: 1997-IJ-CX-K013

190134
C.1

Final Report - Demonstration of a Concealed Weapons Detection System Using Electromagnetic Resonances

January 26, 2001

Sponsored by
U.S. Department of Justice
Office of Justice Programs
Grant Number 97-IJ-CX-K013

AKELA
5276 Hollister Avenue, Suite 263, Santa Barbara, CA 93111

Final Report - Demonstration of a Concealed Weapons Detection System Using Electromagnetic Resonances

January 26, 2001

Sponsored by
U.S. Department of Justice
Office of Justice Programs
Grant Number 97-IJ-CX-K013

AKELA
5276 Hollister Avenue, Suite 263, Santa Barbara, CA 93111

Table of Contents

List of Figures

Introduction

Concealed weapons pose a significant threat to individuals charged with keeping the peace, whether they are our police forces here in the U.S. or our peacekeeping forces overseas. Between 1987 and 1996, 637 law enforcement officers were killed in the U.S. by firearms. In 1996 there were 2,758 assaults on law officers that involved the use of firearms or knives. The move toward relaxation of concealed weapons laws here in the U.S. and the uncontrolled environments associated with peacekeeping operations, provide a strong motivation for developing weapons detection technologies which are non-invasive and can function non-cooperatively. Detecting concealed weapons and safely resolving potential conflicts without violating the privacy of the individuals involved would be an important contribution to the public safety.

Currently available weapons detection systems are primarily oriented to detecting metal, have limited range, have a high false alarm rate, and are either large as in walk through detectors, or, for portable instruments, require close proximity to the person being searched. The walk through detector lends itself to circumvention. The handheld detector puts its user at risk.

These limitations have stimulated the development of a new generation of detectors designed to detect weapons from a distance. Most new approaches rely on imaging methods. An unfortunate side effect of imaging methods is that they tend to be very revealing and, at least in the U.S., there are still significant privacy issues that must be overcome before they see widespread use.

AKELA has developed a weapons detection concept that is non-imaging, portable, detects concealed weapons at a distance, and uses affordable technology. Our concept relies on the fact that weapons have "fingerprints" just like people do. In the case of a weapon, the "fingerprint" is a set of electromagnetic resonances that are determined by its size, shape, and physical composition. These resonances occur over a range of frequencies that extend from about 200 MHz to 2 GHz. Our detector uses a radar to sweep through this range of frequencies and record the resonant response. There is software in the detector to process the radar signals, extract the resonant response, and decide whether there is a weapon. By using the unique "fingerprints" associated with weapons, the detector is able to reduce the chance that objects commonly carried by individuals will cause the detector to alarm. Since there are more people that do not carry weapons than those who do, it is operationally important to be able to keep the false alarm rate to an acceptable level so that the detector has a good degree of predictability.

In addition to ignoring common nuisance objects, the advantages to using this type of radar approach include 1) operation at ranges of 20 feet, 2) detection of weapons concealed behind the back, 3) operation at power levels below the safe limit for human exposure, and 4) providing probable cause for a search without invading the privacy of the individual. It merely excites and detects a response unique to a weapon and sounds an alarm.

The feasibility of using this approach was demonstrated during a Phase I SBIR contract funded by the Defense Advanced Research Projects agency and monitored by the U.S. Air Force Rome Laboratory. However, the results were obtained under ideal conditions in a well calibrated and quiet radar range.

The National Institute of Justice awarded AKELA Grant 97-IJ-CX-K013 to build on the feasibility results of the SBIR program to investigate the effects of the operational law enforcement environment on detector performance. Our grant was divided into two phases. The focus of the first phase was to build a breadboard model using cost effective technology to show that a wide bandwidth radar sensor can detect a weapon being carried by an individual in a noisy operational environment. Major objectives were to determine the ability of the detector to tell the difference between a weapon and common nuisance objects such as cellular phones, and to understand how the various elements of the detector system contribute to overall detection performance.

The focus of the second phase was to improve the signal to noise ratio of the breadboard detector and to integrate the breadboard electronics into a brassboard package suitable for field testing. Major objectives were to determine the variability in detector system performance through limited field testing, and to get a better idea of the impact of operational constraints.

These objectives have been largely met. We have developed a brassboard detection system, as shown in Figure 1, that performs consistently in a field environment. It has been used to show that it is possible to detect concealed weapons on individuals at distances of 15 feet and that it is possible to detect weapons being concealed behind the back. While not hand held, it is portable, self contained, and can operate under battery power for 8 hours. Its power output is 100 mW which is well within the range for human safety, and it is likely that the manufacturing cost of a single unit would be between $500 and $1000.

However, we have also found that the false alarm rate with the existing signal classification algorithm is currently too high for operational use. The signals due to weapons are buried in a large "noise" component due to the person and there are large variations in the signature from person to person. This makes relying on matching patterns as a means of classification unreliable. It will be necessary to reduce the effect of the body signal so that of the weapon stands out in order to ensure better detection performance. While the detection problem has not yet been solved, we have encountered no fundamental limits to success. There are signal processing and additional hardware development techniques that can be used improve the detection performance. However, the grant level of effort was insufficient to pursue them. The remainder of this report summarizes the results of grant activities.

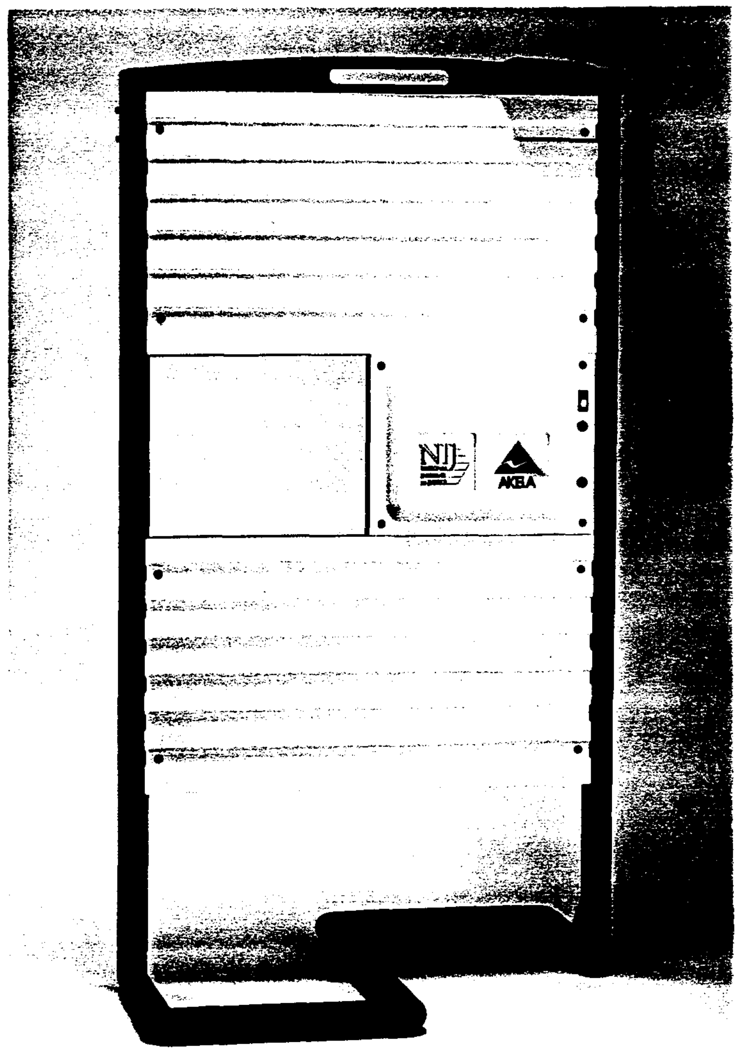

Figure 1 - Brassboard weapons detector.

Background and Theory of Operation

Characterizing the physical properties of the elements by resonance techniques, such as magnetic resonance, has been used by physicists and chemists for many years. Mechanical and electrical systems also exhibit resonance properties which are related to their physical size and composition. The resonant response of a system is related to the natural resonances of that system. These natural resonances are a key component in the solution of the mechanical or electrical equations of motion for a system.

The analog between an object illuminated with an incident impulsive radar signal and that of a vibrating string or drum head is very close. In both cases the objects are driven into oscillation. In the case of a drum, a mallet striking the head produces mechanical vibrations which in turn radiate a musical sound. For the radar, the incident impulsive electromagnetic signal induces resonant or "standing wave" currents on the object which produce the radar returns. Since these resonances are determined by the physical characteristics of the object, the radar return is its electromagnetic "fingerprint". Resonance based scattering exhibits some additional features which make it attractive for object identification schemes. They are:

- The scattered return is larger in the resonance region.
- The natural resonances seen in a scattered return are independent of the orientation of the object.
- An object's resonance patterns uniquely identify it within a class of objects.
- A few natural resonances characterize an object over a large frequency band.

However, in order to induce a resonant response in an object, it is necessary to illuminate it in the frequency band of the natural resonances. In general, the frequency dependence of an object's radar return can be categorized into three overall regions: the Rayleigh, resonance, and optical regions.

The scattering characteristics of the simple sphere are exemplary of the three scattering regions described above. Figure 2 shows the radar cross section (RCS) of a sphere as a function of its circumference measured in wavelengths, $2\pi a/\lambda = ka$, where a is the radius of the sphere and $k = 2\pi/\lambda$ is the wavenumber. The RCS is plotted normalized to the geometric cross section of the sphere, πa^2.

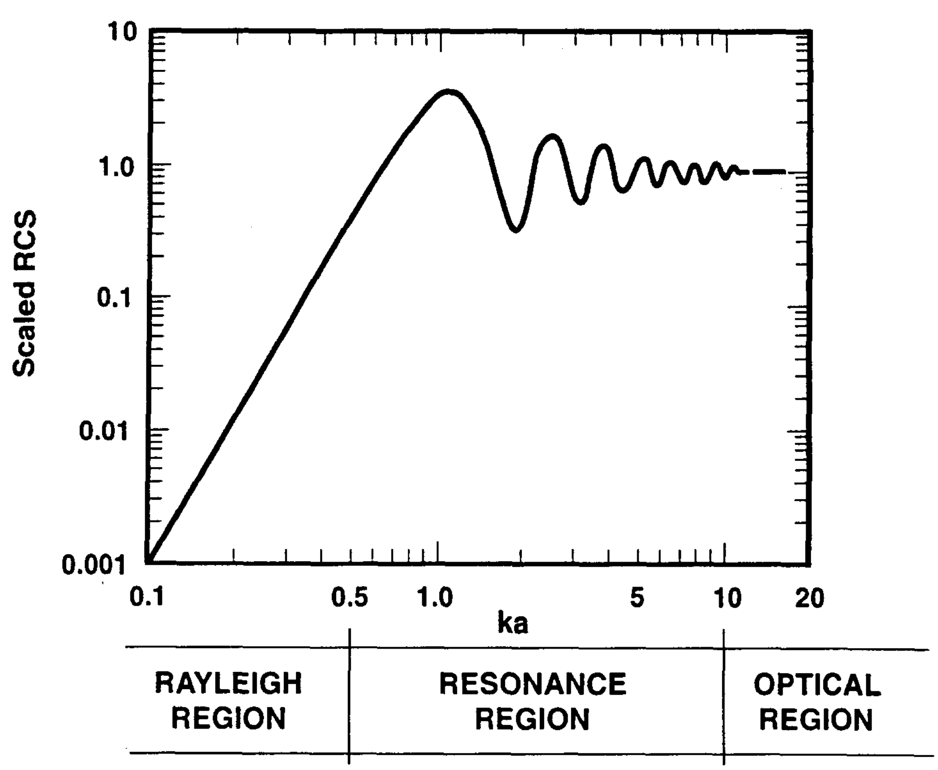

Figure 2 - Radar Cross Section is Enhanced in the Resonance Region

The resonance region return of the sphere has several peaks which correspond to the natural electromagnetic resonances of the sphere. They actually continue into the optical region but converge asymptotically to the physical cross section. The location of the resonances in frequency is directly related to the circumference of the sphere. The first resonant peak, and the highest RCS, corresponds to the point where the circumference equals exactly one wavelength. At this frequency, a wave traveling completely around the sphere's surface constructively adds to a wave scattered from the specular point at the sphere's front surface.

Figure 3 shows these three general scattering regimes in terms of the characteristic dimension of a target, in meters, and the wavelength and frequency of the radar. The shaded region on the figure shows the frequency band which represents the resonance region for concealed weapons. This is a frequency region where antenna sizes are small and electronic components are widely available commercially.

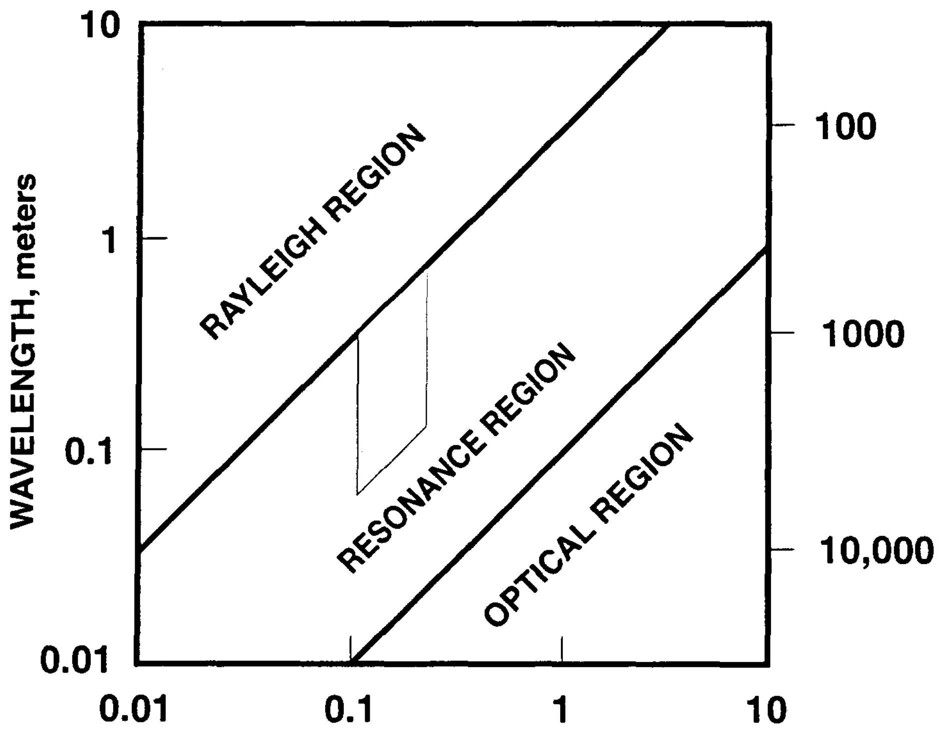

Figure 3 - The Resonance Region of Concealed Weapons

During our DARPA funded SBIR program we made measurements in a quiet electromagnetic environment on a variety of weapons and common nuisance objects. Experimental results confirmed the existence of a unique, aspect independent, radar signature that can be used as a "fingerprint".

The results also showed that the cross section of a weapon is relatively large, making the power requirements for detection modest. Additional tests were conducted with both weapons and nuisances being carried by a human. Results of these tests confirmed that it is possible to distinguish the armed from the unarmed case, even when the weapon is concealed behind the back.

With these results as a point of departure for our grant activity, we set out to develop an experimental detection system that would allow us to investigate the effects of the operational environment on detector system performance. The technical tasks of our grant were separated into the areas of hardware development, software development, and data analysis/classifier development.

Hardware Development

Hardware development activities were divided into three parts - development of a laboratory system designed to get early data and guide the selection of specific methods of implementing a breadboard system, development of a breadboard system that would help focus attention on hardware parameters having the most impact on operational performance, and development of a brassboard system that incorporated to the extent possible the lessons learned from the breadboard system.

Laboratory Test System

In order to induce a resonant response in an object, it is necessary to illuminate it in the frequency band of the natural resonances. As shown previously in Figure 3, for the weapons of interest this frequency band extends from about 200 MHz to 2 GHz. This is a very wide band of frequencies. Either frequency or time domain techniques can be used to generate this wide band of frequencies. In the frequency domain, this is accomplished by sweeping through the band one frequency at a time. In the time domain, this is accomplished by generating a short duration, impulsive waveform. We constructed a laboratory test system using each technique in order to determine the relative advantages of each.

Figure 4 shows a block diagram of the frequency domain system. As can be seen, this system was constructed from a set of Hewlett Packard laboratory test equipment, specially designed wide band antennas, and a computer interface. Using a vector network analyzer enabled us to obtain both amplitude and phase information from our test data. This allowed us to take advantage of various methods of preprocessing the data to take into account the effects of background noise, experimental system response, and antenna amplitude/phase response. With this setup we were able to sweep between frequencies of 10 MHz and 2 GHz. Figure 5 shows the major pieces of test equipment that constituted the

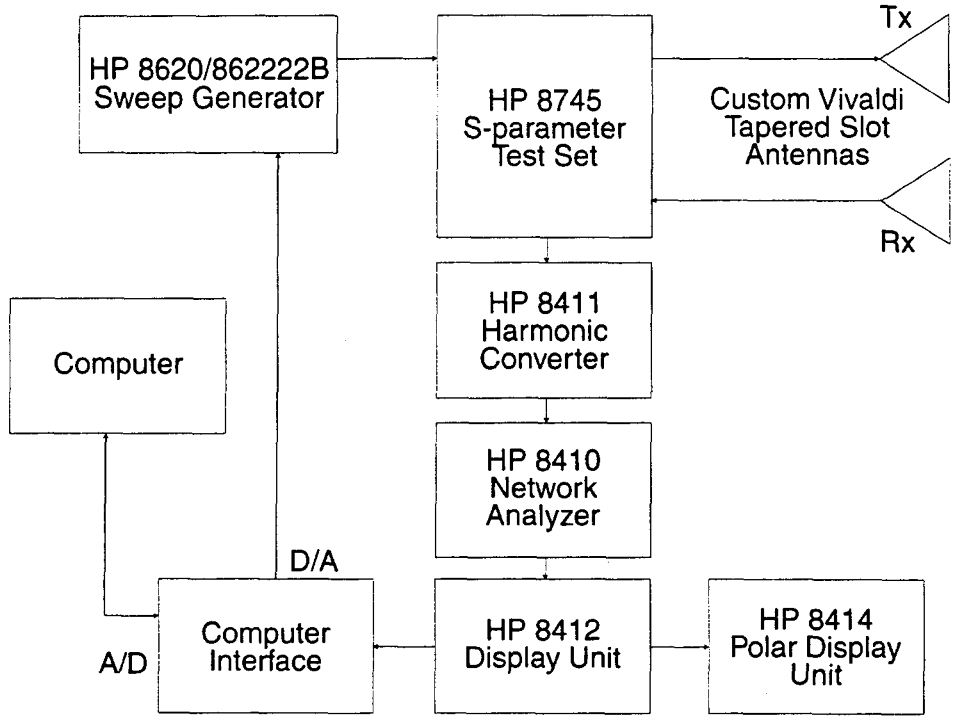

Figure 4 - Swept Frequency Test System Functional Diagram

swept frequency system. The system was used to take data over a wider variety of test conditions than was possible during our SBIR program so that we could refine our signal processing algorithms and characterize sources of system noise to provide guidance to our breadboard development activities.

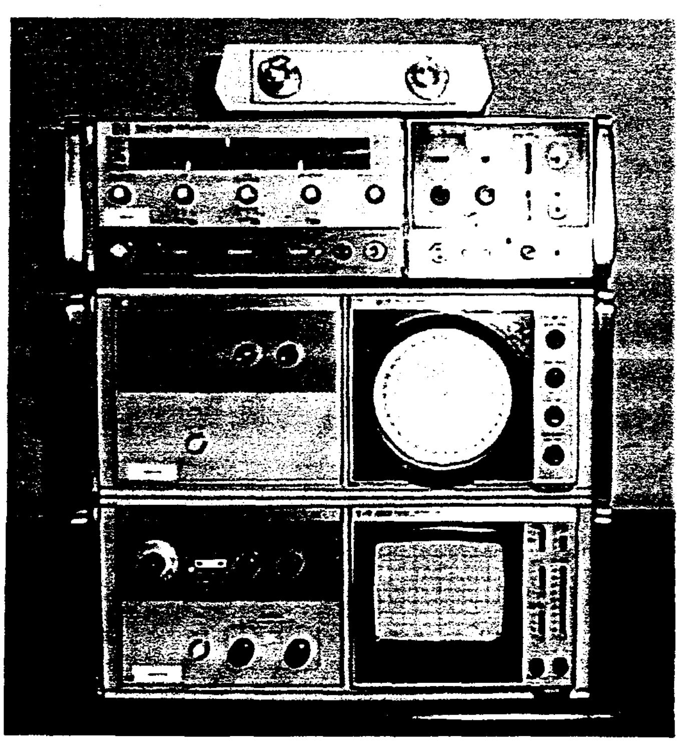

Figure 5 - Swept Frequency System Components

Initial testing of the system with a set of wide band discone antennas uncovered signal to noise ratio problems. The major sources of noise were from adjacent interfering commercial broadcast band (FM, TV, UHF) signals, and antenna coupling. The former was eliminated by a combination of power amplification at the transmitter and high pass filtering at the receiver end. An attempt to eliminate the latter by the use of a set of high isolation spiral antennas did not work as anticipated. While isolation was improved, the nonlinear phase response of the spiral antennas made the use of windowing in the time domain to isolate the desired test signal unrealizable.

This result focused attention on the importance of the antenna to performance of the detector system. As a consequence, we developed a set of custom antennas that provided us with both increased directivity and linear phase response. These antennas are tapered slot antennas of a type commonly referred to as Vivaldi antennas. Figure 6 shows two different implementations of this type of antenna. The antenna on the left is 24" long and 12" tall and was the first implementation we used. It was capable of radiating between 250 MHz and 2 GHz. The antenna on the right at 18" long and 12" tall was an attempt to reduce the size of the antenna by trading off its low frequency response characteristics. It was capable of radiating between 400 MHz and 2 GHz. While there is still antenna coupling between antennas of this type, the directional characteristics of these antennas make separating transmit and receive antennas by approximately 3 feet a solution sufficient to reduce the impact of antenna noise on system performance.

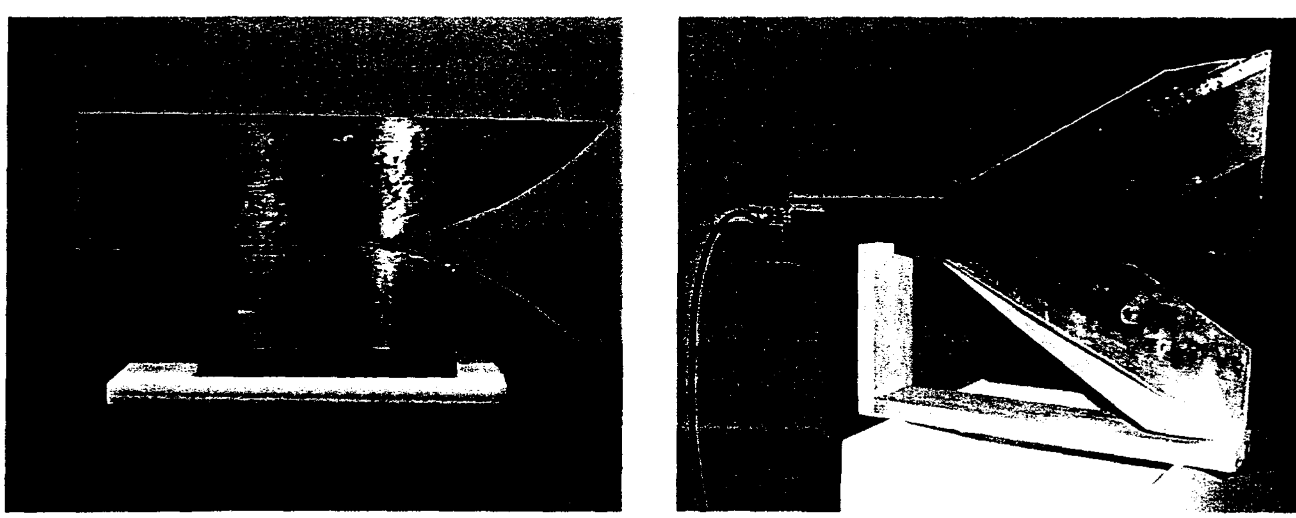

Figure 6 - Custom Vivaldi Antennas for Experimental Test System

The frequency domain system was used to record 160 tests on varying configurations in an outdoor environment over a frequency range of 250 MHz to 2 GHz. Data was taken at distances of 6, 8, 10, and 15 feet for both armed and unarmed individuals. Tests included 5 different individuals, two of which were female, standing with either their front, side or back to the detector antennas. In half of the

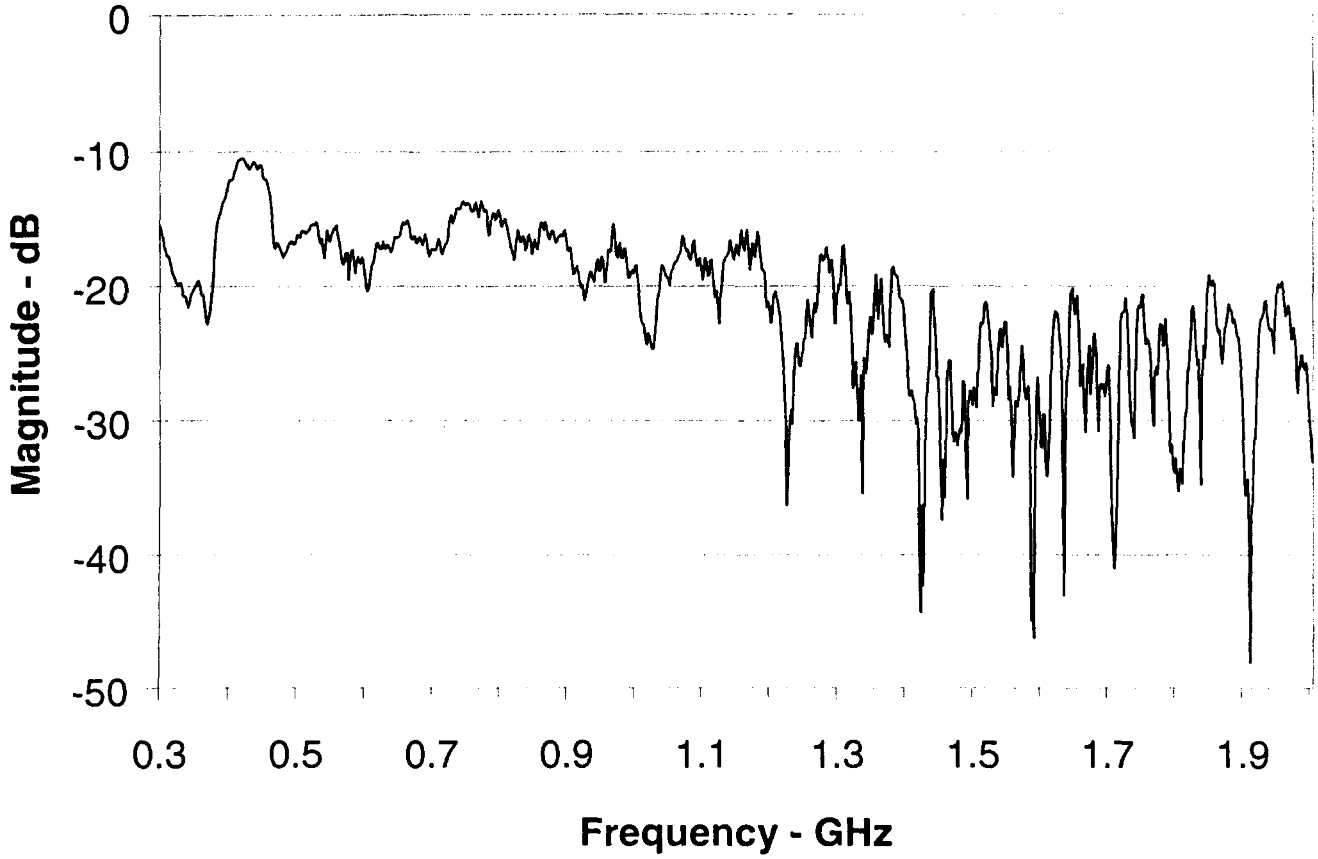

Figure 7 - Representative Data From Swept Frequency Test System

weapon tests the gun was on the opposite of the body from the detector antennas. Figure 7 shows a representative data trace. The signal to noise ratio achieved with the test system was between 10 dB and 15 dB. Data taken below 300 MHz was not usable because the noise from antenna coupling below that frequency was sufficient to saturate the network analyzer.

Analysis of this data showed that different individuals have different wideband radar cross sections with the dominant low frequency peak moving around both in amplitude and frequency. However, the cross sections are consistent for a single individual. Also, the weapon causes a distinct difference in the cross section and can be seen when carried behind the person. This data confirmed the general trend in the data that was obtained from the radar range at Pt. Mugu.

Figure 8 shows a block diagram for the time domain laboratory test system. As was the case for the frequency domain test system, we used off the shelf test equipment to assemble the test systems. Functionally this system was much simpler than the frequency domain system consisting of a pulse generator, a digitizing oscilloscope, control computer, and antennas. The output of the pulse generator is a 40 volt, 250 picosecond risetime pulse. The bandwidth of the HP digitizing oscilloscope is 12 GHz. Figure 9 shows the major pieces of test equipment that constituted the time domain system. Since subtracting background response is easier in the time domain, we used that system to more closely

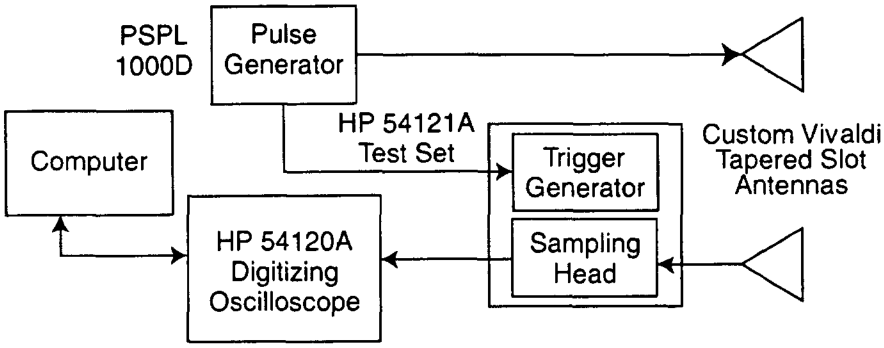

Figure 8 - Time Domain Test System Functional Diagram

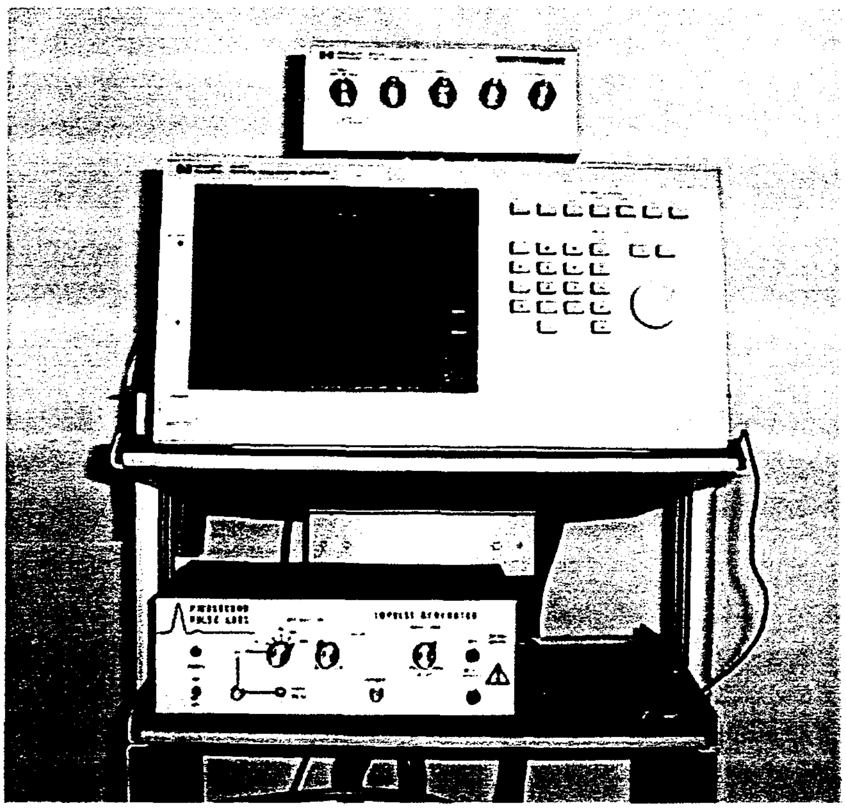

Figure 9 - Time Domain System Components

examine the nature of the weapons detector response to individuals of different builds, standing in different orientations with respect to the detector.

Initial testing with the time domain system focused on characterizing the frequency response of the Vivaldi antennas and the pulser. Figure 10 shows both the time and frequency response of the pulser. The risetime of the pulser limits the amount of energy present at higher frequencies. As can be seen from the figure, the power output decreases monotonically as frequency increases with there being a

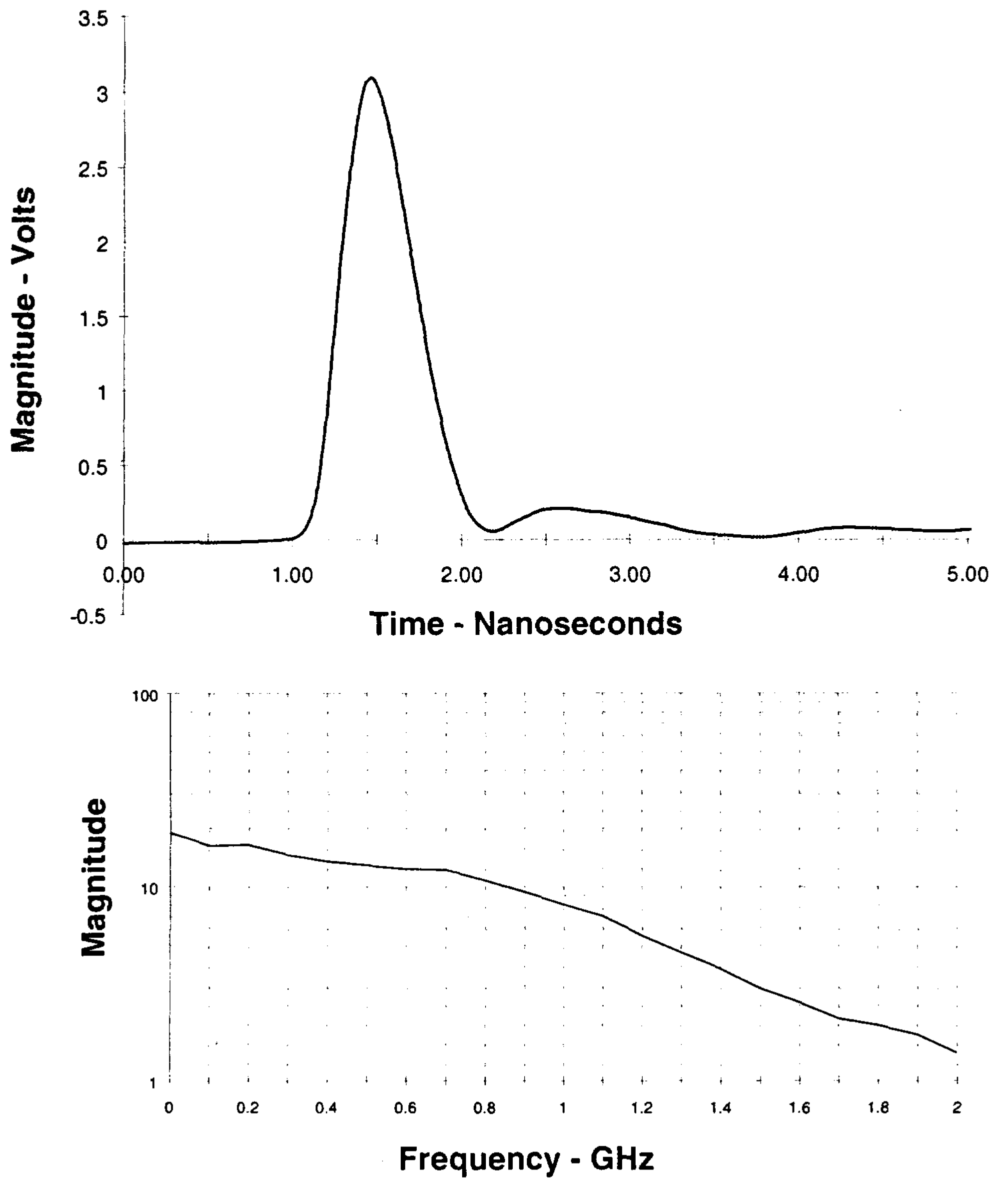

Figure 10 - Time Response and Power Spectrum of Pulser

10 dB decrease between 500 MHz and 2 GHz. Using this information it was possible to determine the frequency response of the antennas themselves. Figure 11 shows both the time and frequency response of the Vivaldi antennas measured by the time domain system. Note in the power spectrum the roll off of energy at both the low and high ends of the spectrum. This spectrum represents the combined response of both the pulser and the antenna with the low frequency roll off being due to the antenna and the high frequency roll off due to the pulser. Removing the response due to the pulser gives an accurate picture of the frequency response of the antennas themselves. Figure 12 shows the frequency response of the antenna alone. Note that it is quite flat over the range of 400 MHz to 2 GHz making this type of antenna quite well suited to meet the requirements for our wideband detection system.

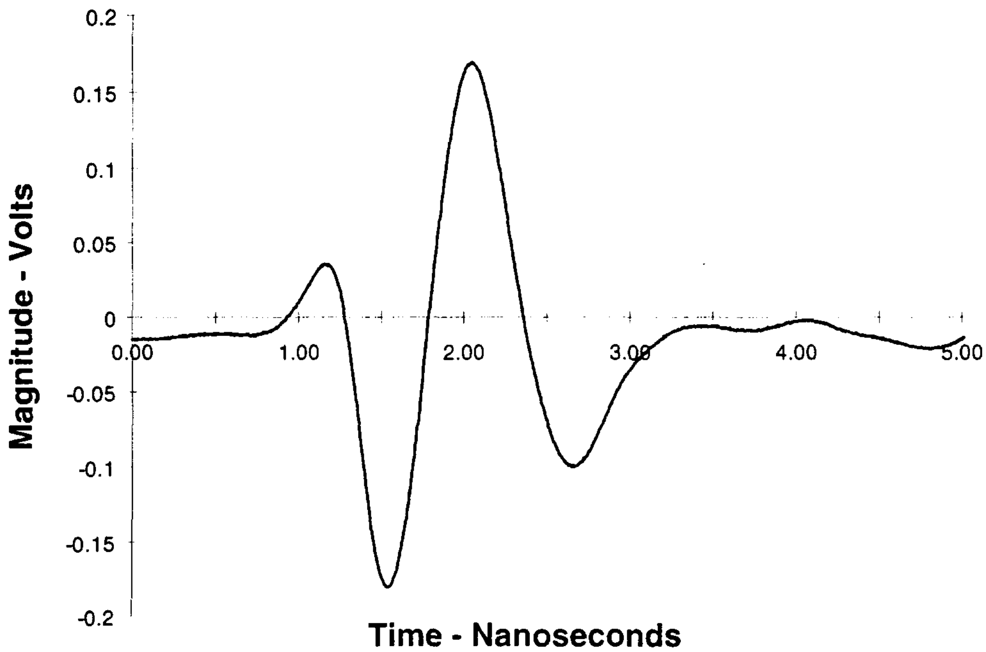

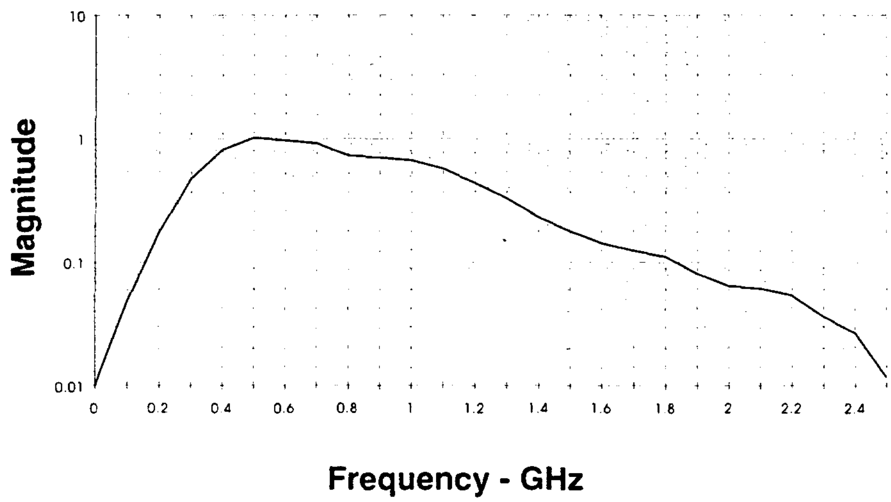

Figure 11 - Time and Power Spectrum of Antenna and Pulser

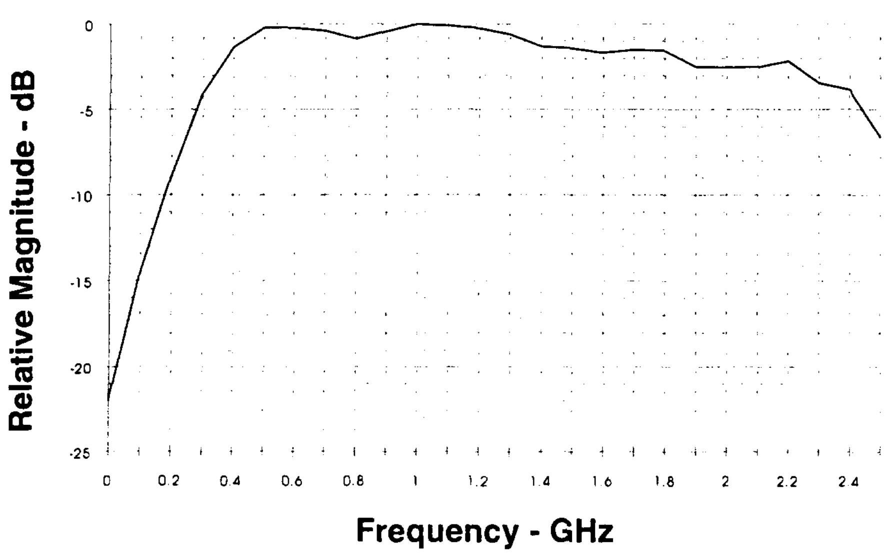

Figure 12 - Frequency Response of Vivaldi Antennas

The time domain system was used to collect data on 8 different individuals, standing at distances varying between 6 and 15 feet in an indoor environment. Figure 13 shows a picture of the indoor test setup. A total of 212 tests were recorded representing varying combinations of distance, individual, orientation, and the presence or absence of a weapon. About half of the tests included a weapon, and of those, in about a third the weapon was on the opposite side of the body from the detector.

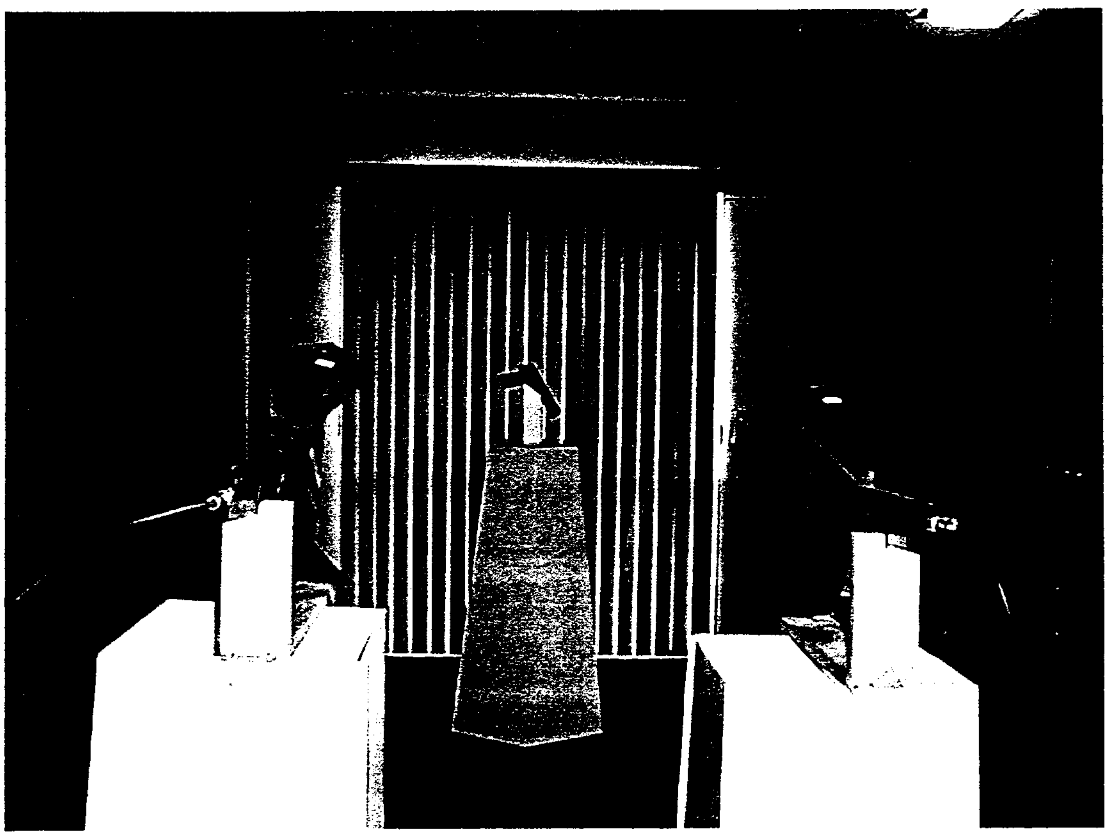

Figure 13 - Time Domain System Experimental Setup

The first tests performed were to more closely examine the nature of the weapons detector response to individuals of different builds, standing in different orientations with respect to the detector. Figure 14 shows how the signature of two different individuals varies as a function of orientation. Note that for each individual the left and right side signatures match closely. There is a distinct difference between the signature of an individual facing toward the detector and facing away. These results confirmed the observations made from the data collected with the frequency domain test system.

Because background subtraction was much easier with the time domain system, a second set of tests focused on showing the presence of a weapon in the detected signature was performed. In these tests the individual stood facing the detector in all cases, but the location of the weapon was changed for each test. For the cases where the gun was carried either in front or behind the person, it was oriented so that it's flat side faced the detector. For the cases where the gun was carried at the side of the person

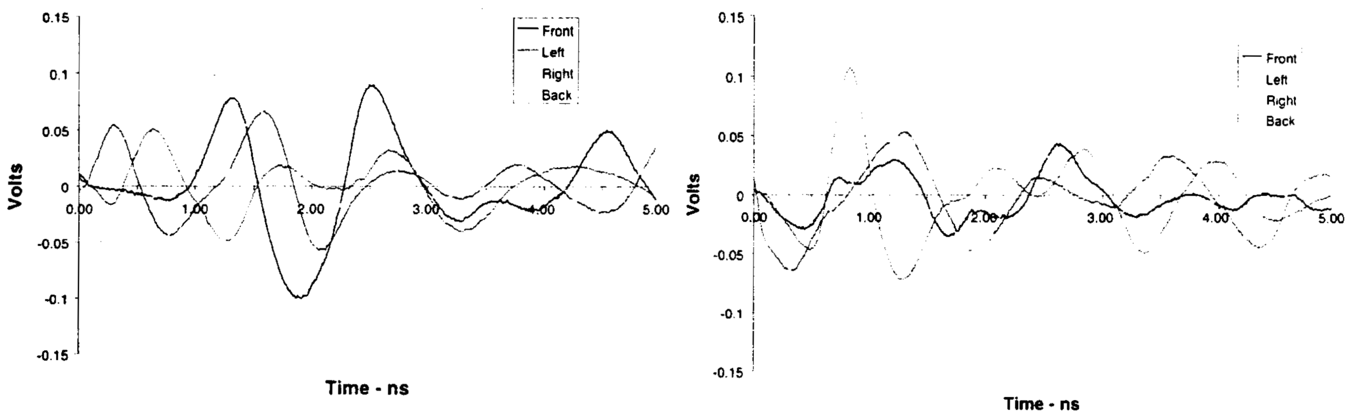

Figure 14 - Variation of Signature With Body Orientation

(carried in the pants pocket), it was oriented so that it's narrow side faced the detector. Because we were using the time domain system we were able to subtract the response of tests of the person and background alone from the response of tests which included the weapon. This allowed us to see the response of the gun alone. Figure 15 shows the results of these tests. As can be seen from the figure, the signature of the gun shows up in all cases, even the case where the gun is blocked from the view of the detector by the human body. These data showed that the gun signature is always in the radar return from the individual and are significant because they show that we can see weapons concealed from view by an individual's body.

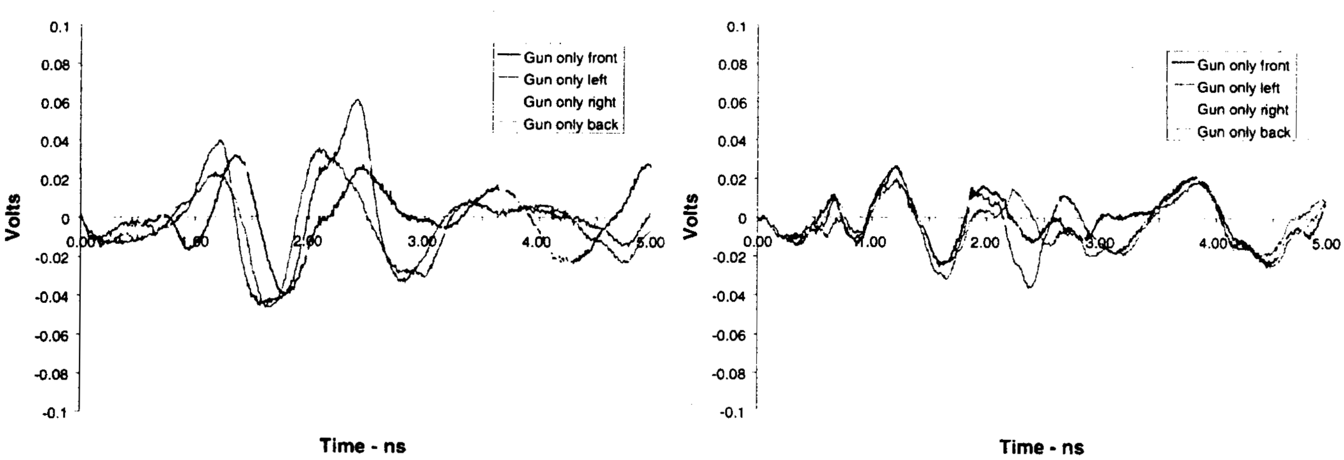

Figure15 - Signature of Gun in Various Locations

Based on the performance of the laboratory systems and an analysis of the ease and cost of implementation of each method, it was decided to construct the breadboard system using the swept frequency method of exciting the resonances.

Breadboard System

We began the process of shrinking the size of the frequency domain system by dividing it into three phases so that we could more easily characterize the changes in system performance that resulted from the redesign process. Figure 16 shows a functional block diagram of the breadboard design and the relationship each part has to the laboratory test system.

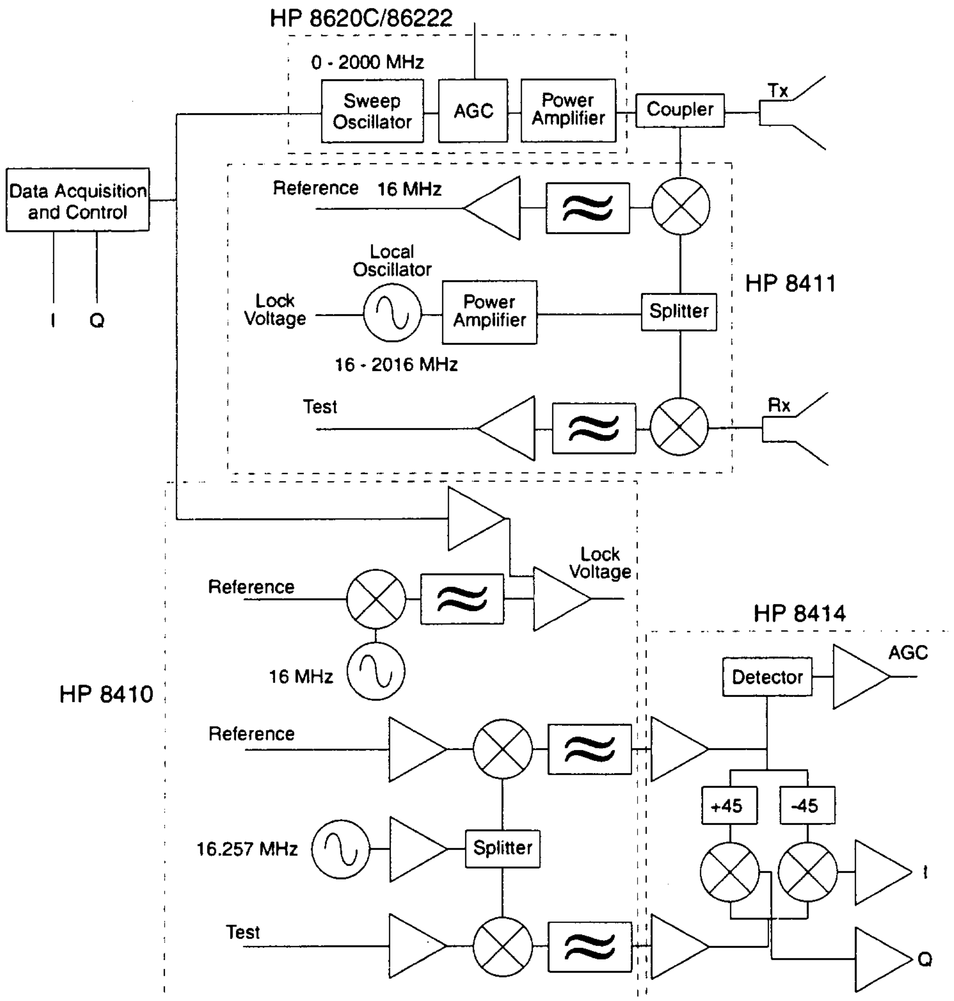

Figure 16 - Breadboard Detector Functional Design

The first phase consisted of replacing the front end electronics which were made up of the power amplifier, signal reference, mixer, and intermediate frequency amplifier. A design using building block microwave components was developed, a printed circuit board made, and the circuit was assembled. This circuit board was approximately 4" by 6" and was designed to replace the separate power amplifier and HP 8411 Harmonic Converter.

The second phase of design was the replacement of the intermediate frequency processing and detection electronics. This portion of the breadboard was performed by the HP 8410/8414 network analyzer and polar display. It is what detects the return signal and separates it into its in phase and quadrature components providing us with both amplitude and phase information. We developed a design using building block microwave components. This board presented an unexpected challenge. The quadrature detection portion of the board required adjustment because of the poor match between the electrical characteristics of the parts from which it was assembled. Debugging the performance of this board was more difficult than anticipated because of the need to ensure that the quadrature detector produced signals with the correct phase differences. In order to reduce the overall foot print of the breadboard electronics, a second layout of the IF board was made to add a microprocessor controller, A/D converter, and power conditioning circuitry to the board. Software was written for the microcontroller to enable it to control the frequency sweep and digitize and format the I and Q test data.

The third phase of design was the replacement of the sweep frequency generator. Our initial design incorporated two fixed and one variable YIG oscillators in order to produce the sweep signal, the offset sweep signal, and the reference signal. Unfortunately, as the testing of this board proceeded, it was determined that this design was susceptible to unwanted, inband, harmonics. As a consequence, a new design was developed consisting of four narrower frequency spread variable oscillators controlled by a frequency synthesizer which would allow us to step through each band in a well controlled manner. This eliminated the harmonic generation problems of the previous approach but resulted in a higher parts count sweep generator. Fortunately, the parts were less expensive and the reliability of the overall sweep generator was improved. As this board was built and tested it was found that there were coupling problems between the two oscillator circuits. To solve this problem it was necessary to add an additional frequency synthesizer integrated circuit to the board. A quick patch to the breadboard was made and it was verified that the problem had been eliminated. Characterization tests of the board showed that the reference and sweep oscillator circuits were well matched over the entire 500 to 2000 MHz band.

Once the individual circuit boards were tested and debugged, they were integrated into an enclosure measuring 5" wide, 6" tall, and 11" deep. Figure 17 shows a photograph of the packaged breadboard electronics. The package contained a power supply and operated off of 110 VAC. The breadboard was controlled by a laptop computer, could operate off of a large 12 volt battery, had a radiated power of 100 mW, and took 17 seconds to perform a single 1024 point sweep between 450 MHz and 2 GHz. Figure 18 shows a picture of the assembled breadboard detection system components.

The reduced size of the breadboard made it easy to take to the field to gather data. This system was used to take over 900 tests at distances varying from 10 to 20 feet, with three different individuals in

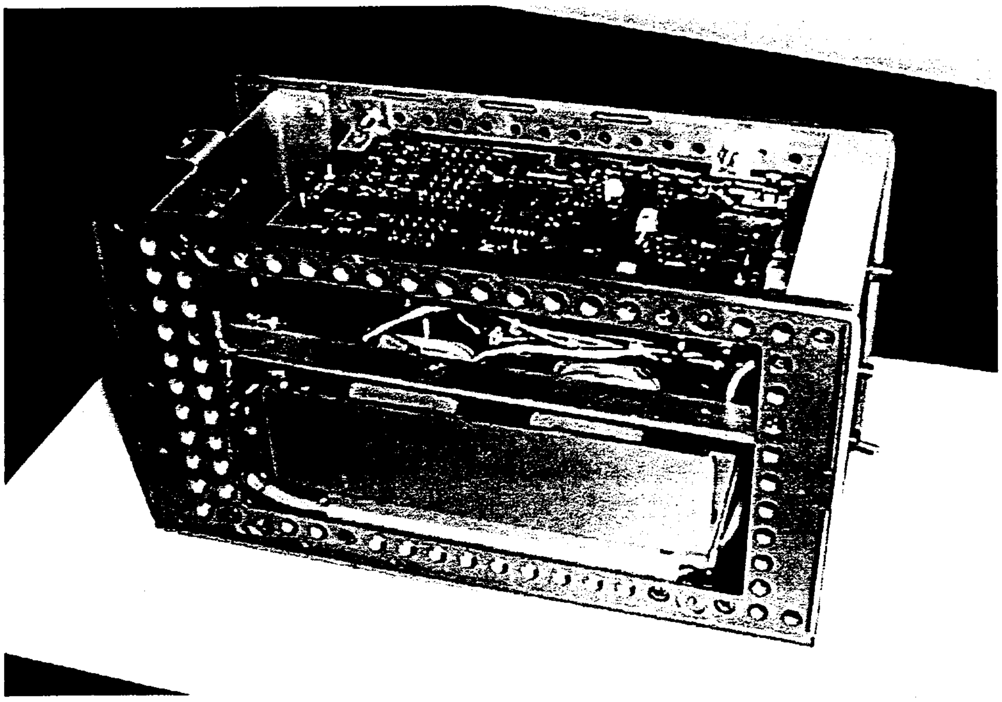

Figure 17 - Weapons Detector Breadboard Electronics

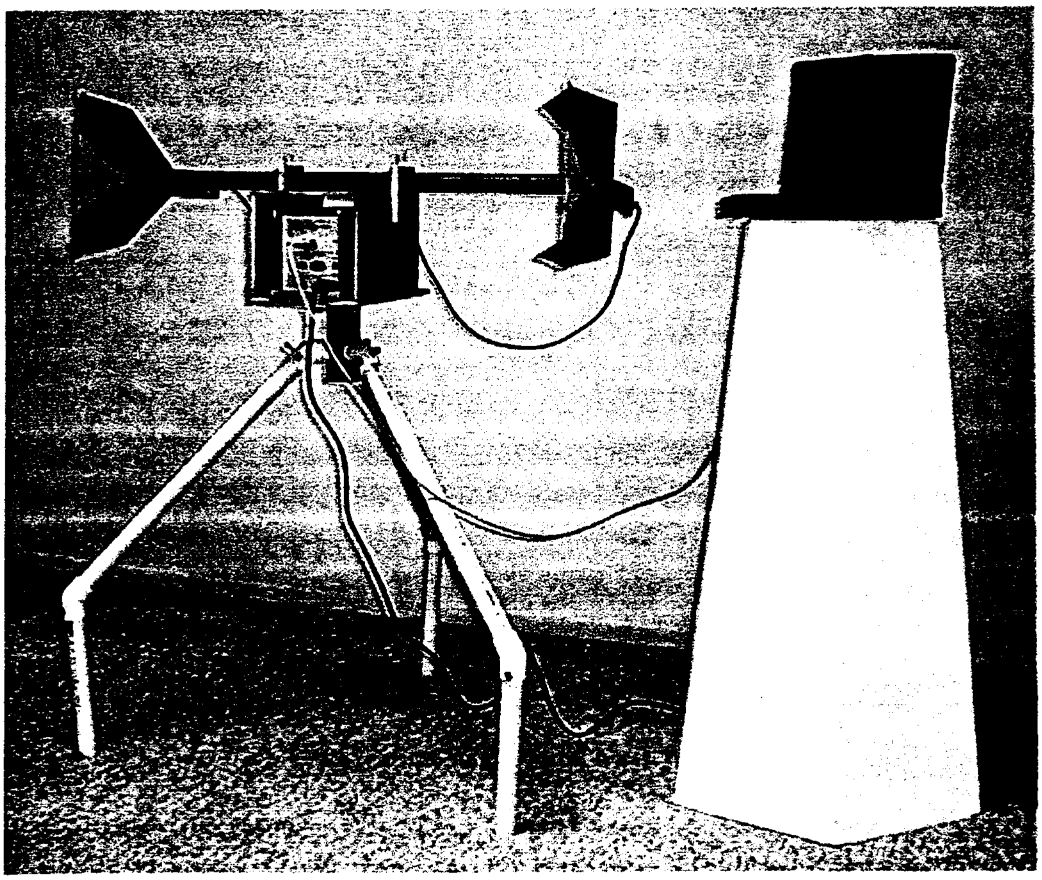

Figure 18 - Breadboard System Components

a variety of orientations both armed and unarmed. Figure 19 shows an example of the frequency and time data for a test at a distance of 15 feet with a person facing the detector and concealing a handgun behind his back. The peak in the time domain plot at 30 ns is the person with the gun. The frequency response data is comparable to that taken with the swept frequency laboratory system shown previously in Figure 7.

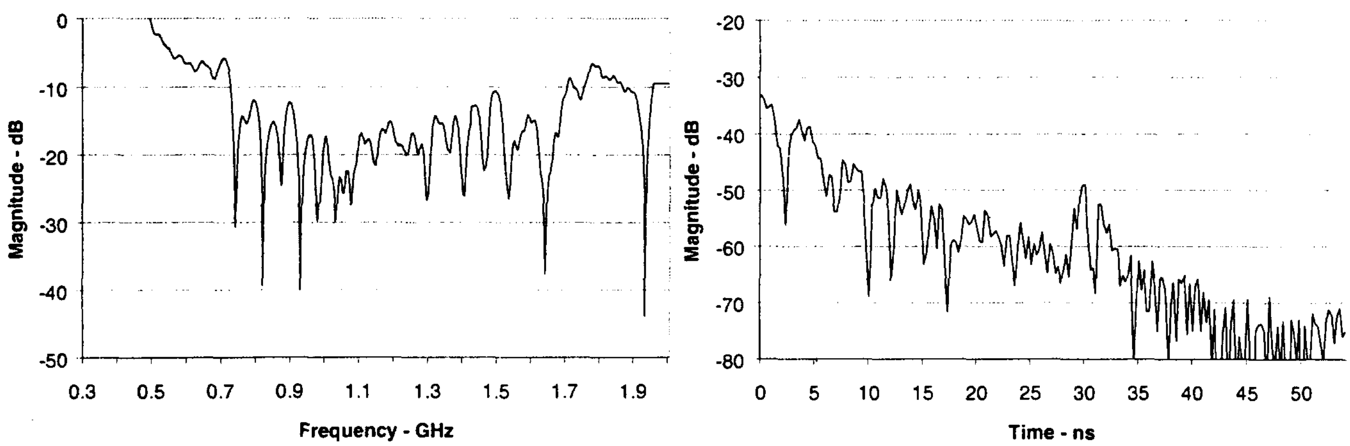

Figure 19 - Breadboard System Test Data

Processing of the data taken with the breadboard system showed that signal classification performance was limited by antenna coupling noise. At this point the grant was modified to allow us to explore antenna, packaging and processing alternatives to improve system performance. As a result we embarked on the third part of our hardware development activities that resulted in a brassboard design.

Brassboard System

Development of electronics for the brassboard system proceeded in two steps. In the first, we made quick fixes to the breadboard electronics to improve some operational performance parameters. This allowed us to actually see the impact of the changes in the operational performance of the breadboard system before committing to them in a redesign. In the second, we incorporated the modifications made to the breadboard and some additional performance improvement measures into a redesign of the electronics.

The first modifications made to the breadboard were to reduce a single sweep to 2 seconds from 17 seconds and to change the sweep range so that it began at 400 MHz. We did this to avoid the antenna cross talk present below that frequency and to improve the frequency resolution over the range of interest. At the same time, the breadboard was modified so that it would run from a single 12 volt power supply. This allowed us to package the detector in a box 2" tall, 6" wide, and 9" long. With these modifications we were able to support NIJ's attendance at the DoD's Force Protection Equipment Demonstration where we performed demonstrations with the breadboard and showed attendees the capability of the detector system to see a weapon concealed behind the back. Figure 20 shows the repackaged breadboard.

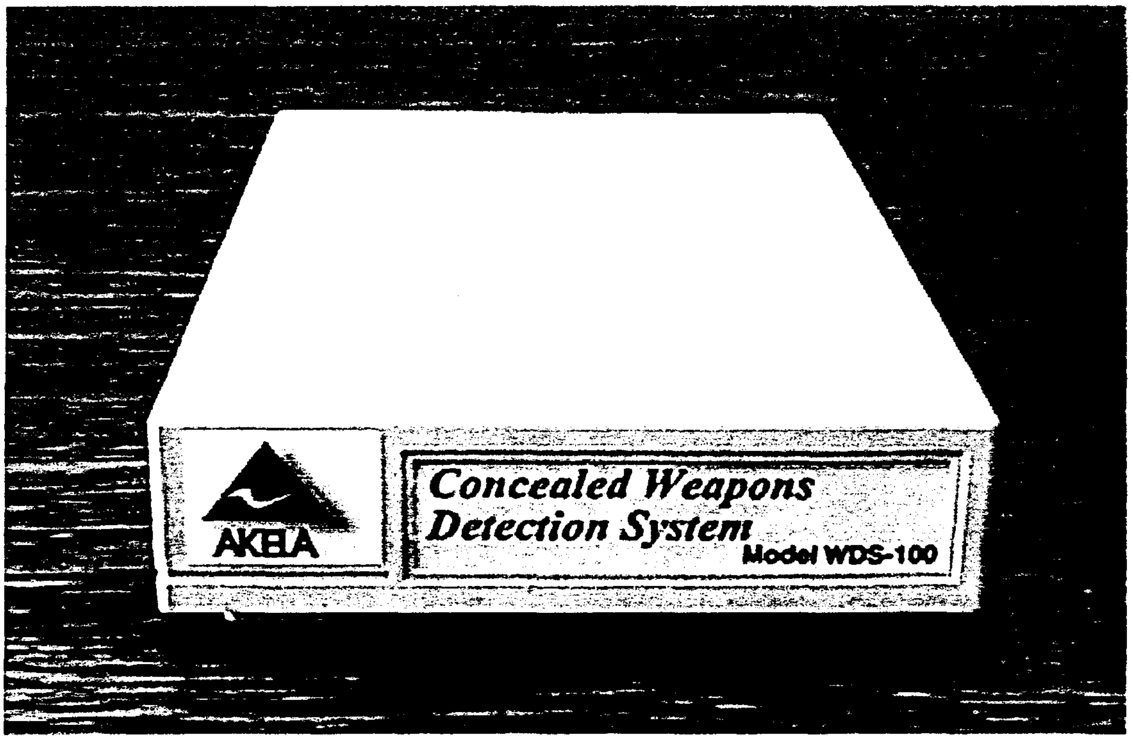

Figure 20 - Repackaged and Modified Breadboard Electronics

Further modifications were made to the breadboard to improve its dynamic range, its isolation between transmit and receive sections, and its power consumption. These improvements allowed us to quickly improve the breadboard further to give us a system that we could use to focus on data collection and processing activities while brassboard design efforts proceeded.

This system was used for conducting two days of sensor fusion data collection at Chang Industries at the request of NIJ. Over 650 tests were performed with 7 different handguns, 2 plastic replica handguns, 3 different people, and 2 distances. Data were taken with weapons in a variety of configurations and people in different orientations. AKELA's modified hardware operated without any problems during the entire time. This effort resulted in the creation of a CDROM data disk with 40 Mb of data which was delivered to AFRL. Figure 21 shows the configuration that was used for this testing and an example of the quality of the data taken.

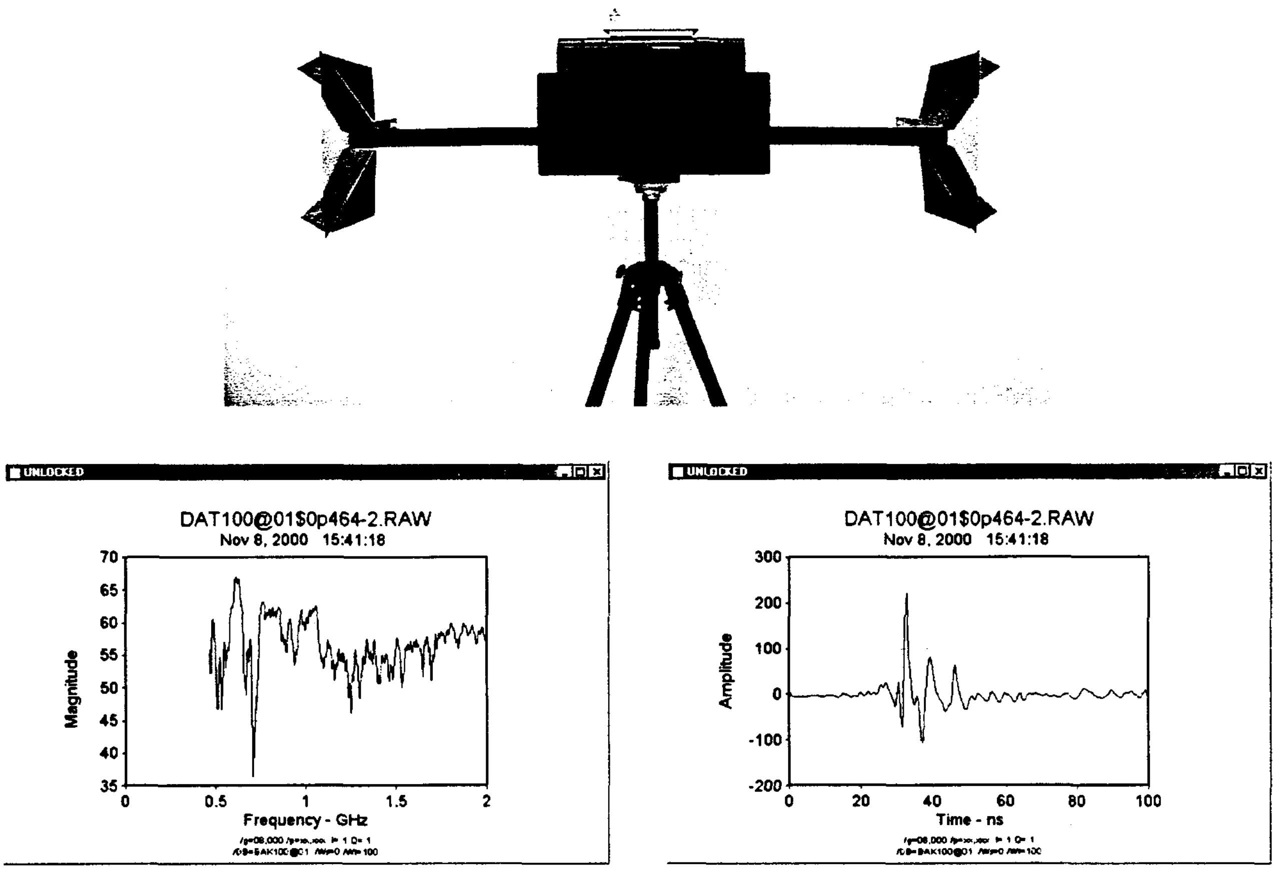

Figure 21 - Early Brassboard Test System and Example Data Quality

At this point our brassboard design activity began with the goal of reducing the size of the electronics, extending the usable frequency range, and reducing the power consumption yet further. The RF portion of the board was modified to reduce power consumption and to improve the isolation between transmit and receive sections. The IF, processing, and power supply board was extensively modified. These modifications included improving the dynamic range by installing a 14 bit ADC and reworking the power supply so that the electronics would run off of a single 12 volt supply. Circuit boards were manufactured, parts ordered, and four sets of boards were fabricated. During checkout of these boards we found and fixed some design problems which improved the overall performance of the brassboard electronics. One set of boards was reserved for evaluating some new electronic components - mixers and oscillators - that would allow us to extend the frequency range of the system. The brassboard electronics run from a 12 volt VCR battery that is 7" x 2.5" x 1" and weighs a little more than a pound. It would operate the weapons detector continuously for about 2 hours. Figure 22 shows the completed brassboard electronics hardware.

Because our goal was to reduce the size of the antenna so that our detector could be handheld, our antenna work began by looking at more compact antenna forms such as the spiral. Unfortunately,

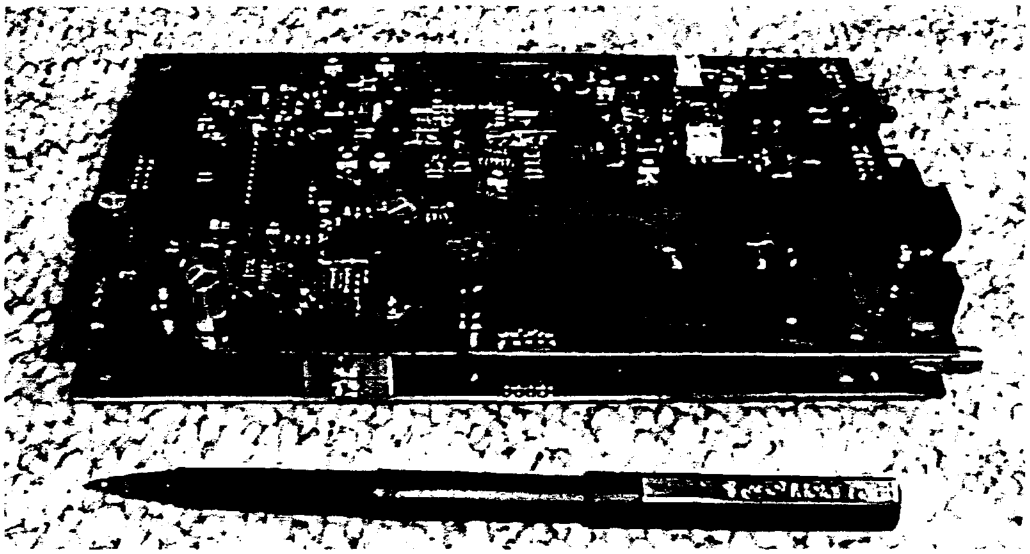

Figure 22 - Brassboard Electronics

many of these antennas have a poor phase response. Since our brassboard electronics had the ability to control frequency and measure amplitude and phase response very precisely, we anticipated being able to develop algorithms to correct for this nonlinear response. This would allow us to reduce the total size of the system without losing any performance.

We made phase response measurements on a series of antennas to get data that we used to develop correction algorithms to improve antenna performance. A spiral antenna with a diameter of 9" was tested over the frequency range of interest with a large aluminum disk used as a calibration target. Since the phase dispersion of a spiral is large, we were unable to see the calibration target in the signal. (The effect of phase dispersion produces the same type of signal you would see as that of an antenna physically being moved back and forth during the sweep interval.) After making corrections to the data using the algorithms we developed, the target showed up clearly. This important result gave us confidence to further investigate non-ideal but small antennas for our system.

Two different 8.25" diameter spiral antennas, one an Archimedes spiral and the other a logarithmic spiral, and a printed circuit balun were developed and fabricated by a printed circuit techniques. The balun is a device that matches the antenna to the brassboard electronics and allows efficient transmission and reception of the detector signals. Measurements on the balun showed a need for improvement in order to give an acceptable return loss over the band of frequencies of the brass-board and subsequent modifications resulted in acceptable performance. Measurements on each of the spiral antennas showed that they radiated in the designed band. Since these designs are planar antennas, they radiate bidirectionally and required an enclosure and absorbing material to decrease the back radiation. Spiral antennas are circularly polarized. Measurements made with the spiral antennas showed that there was a 30 dB difference between a received signal that is of the same polarization at the antenna and that which is polarized in the opposite direction. A series of measurements were made to determine the amount of isolation we were able to achieve between the transmit and receive antennas as a function of their spacing. We found that we were able to place the antennas side by side touching each other and still achieve the same isolation as with the Vivaldi antennas spaced 3 feet apart.

We also designed and fabricated a planar log periodic antenna to cover the same bandwidth as the spiral antennas. Unlike the spiral antenna, the log periodic antenna is linearly polarized. This antenna was tested and measurements made of its bandwidth, return loss, and isolation. The antenna had a relatively flat response over the design range with the exception of the low end. At 450 MHz there was unexpectedly poor behavior which was due to the enclosure that housed the antenna. The diameter of the enclosure was very close to the diameter of the low frequency radiating elements of the antenna and interfered with the low frequency transmission response. The return loss measurements showed excellent performance being between -10 dB and -20 dB over the frequency band of interest. Figure 23 shows the planar logperiodic and Archimedes spiral antennas that were developed.

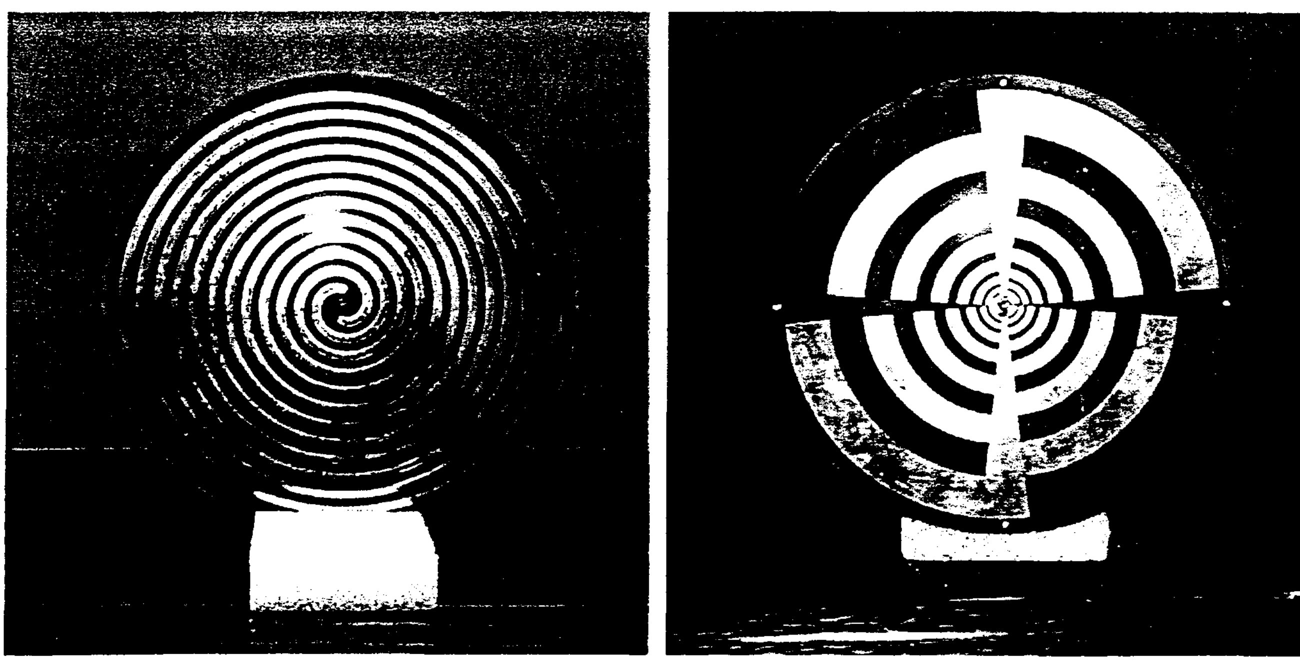

Figure 23 - Archimedes Spiral and Planar Log Periodic Antennas

Because of the good laboratory measurements of the planar antennas, we performed a series of field tests with the brassboard electronics to see whether we had achieved any performance improvement. We were surprised at the results since the signal from the detector was weaker than expected. Subsequent testing was performed to determine whether the reduced signal was due to poor radiation efficiency of the antenna or due to pattern gain. Results indicated that the radiation efficiency was commensurate with commercial planar antennas but that the pattern efficiency was not as good as the Vivaldi antennas.

In order to address the pattern gain issue, we examined variations on the planar geometry. Two additional designs were performed and implemented with copper tape in order to get quick results. One of the designs was folded. Folding the antenna makes it radiate unidirectionally giving a 3 dB advantage over a conventional planar design. Measurements made of the antennas indicated that additional parameters in the design would have to be changed in order to improve the pattern gain. Since the level of effort to perform additional design activities on the planar type of antenna was greater than we had allotted for this portion of the grant, we concluded that making modifications to the baseline Vivaldi design was the most cost effective course of action.

Accordingly, several variations of the Vivaldi antenna were made to investigate the effects of antenna length and feed mechanism. One was a replica of the baseline antenna but with an end feed with slotline excitation. The other was a shortened version - 18" as opposed to 24" - and stripline fed from the bottom of the antenna. Pattern and gain measurements were made for all the variations. We found that the antenna pattern remains relatively constant with the different variations, however, gain drops off for the shorter configurations. A corresponding set of isolation measurements was also made with the antennas in both horizontal and vertical configurations. These measurements indicated that the antennas could be as close as 6" when spaced vertically and still avoid saturation of the system due to antenna coupling. As a result of these measurements we concluded that the baseline configuration for the weapons detector would be the original Vivaldi antenna design but with the antennas vertically separated.

We engaged the services of an industrial design firm to develop a prototype package for the brassboard system. The purpose of this packaging was to 1) enable us to have a simple method of taking experimental data in the field, and 2) to give us an idea of what a potential first generation weapons detector would look like. A quick set of conceptual designs was developed and we selected one for fabrication. The system is controlled by a Palm Pilot™ personal organizer and is capable of taking data. Figure 24 shows the relative size of the brassboard electronics and the control computer.

Initial experiments with the packaged detector showed its performance matched that of the system before packaging. The self contained nature of the detector made it very easy to do experimental work in the operational environment. Most of our testing was performed in the parking lot adjacent to AKELA's offices. One of the first things we noticed was that we had to be careful how long we tested with the unit since performance degraded as the battery voltage dropped below acceptable limits. That gave us about 1.5 hours worth of test time before having to recharge the battery. We

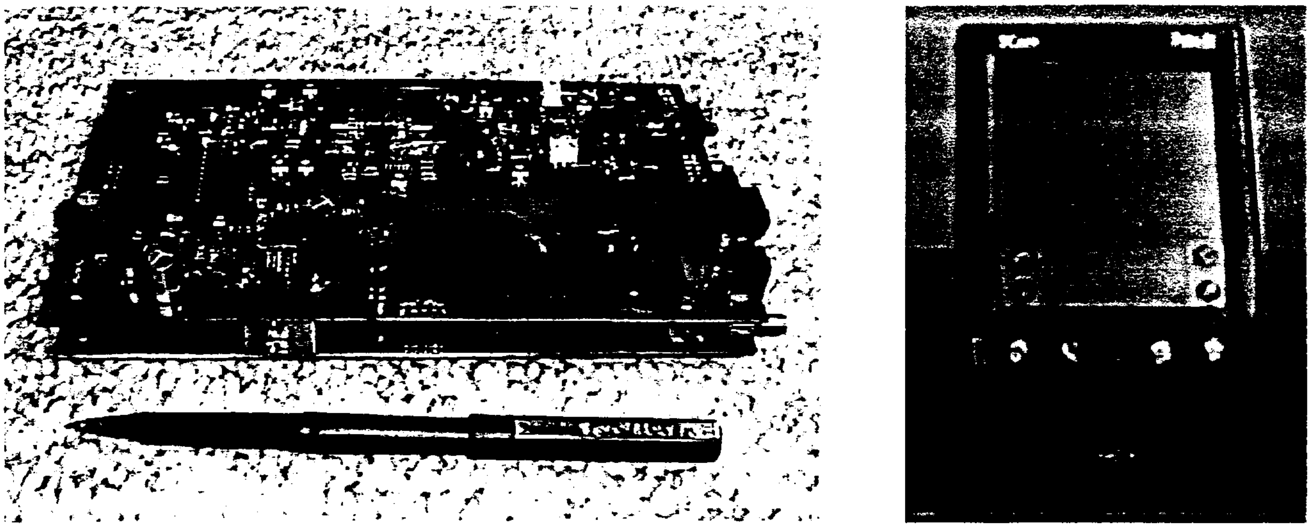

Figure 24 - Relative Size of Brassboard Electronics and Control Computer

subsequently modified the brassboard electronics so that the RF section is only turned on during data collection. That effectively extended the amount of time that could be spent in the field to about 8 hours before the battery needed to be recharged.

The performance of the unit was very repeatable. Because of this it was possible to perform a simple background subtraction to eliminate the effects of antenna coupling. Our data collection included testing to determine the effects of different people, orientations, and weapons. Analysis of this data indicated that the signatures from some people show more variation than those from others, that the signatures of different people look quite different as do the signatures of a single person from the front, back, and side, and that the signatures of a single person vary from day to day. These results confirmed observations made from earlier sets of tests. Figure 25 shows examples of the quality of data collected with the packaged brassboard.

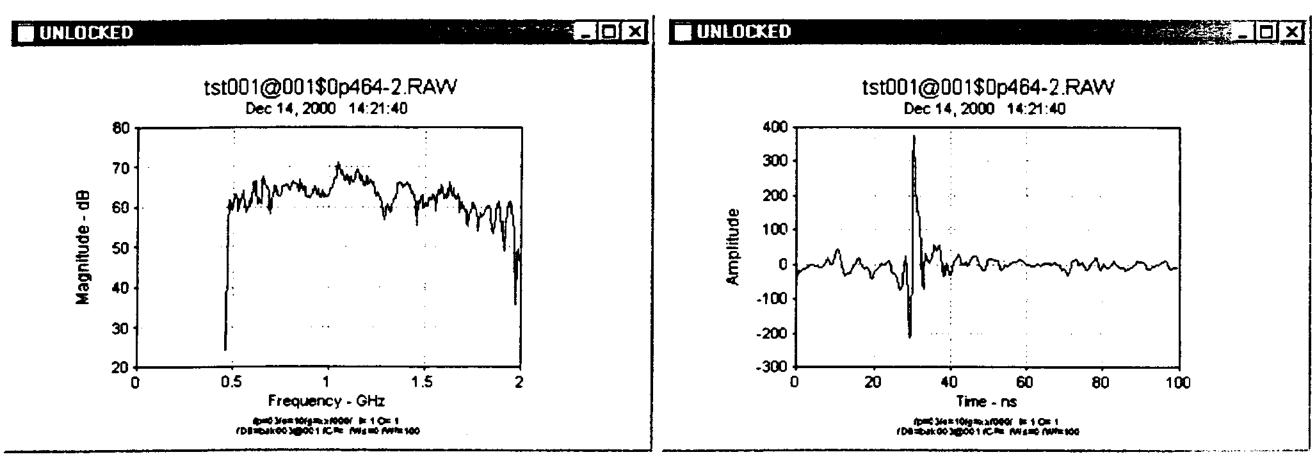

Figure 25 - Data Collected With Brassboard Detector

The control algorithms for the brassboard were modified so that it was possible to both take data for archive and training purposes, and to make classification judgements with the Palm Pilot™ control unit. Extensive testing of this capability was performed and has shown that the signal processing task is more difficult than originally anticipated. Signatures are very dependent on movement of the test subject, and the level of modification of the signature by a weapon is small. Techniques to improve the signal to "noise" ratio of the weapon against the background of the person need to be the focus of further signal processing activities. As a result we modified the brassboard control software to allow independently selecting the limits of the frequency sweep, the number of points recorded per sweep, and the number of sweeps to average for a single data point.

Software Development

The software which was developed for the grant evolved to support program needs. We developed a set of software programs integrated into an environment that made it easy to both collect and process data with the breadboard and brassboard test systems. These programs were accessed through a windowed interface that allowed us to both control hardware data collection and view processed results interactively. This interactivity allowed us to change test parameters, view results, and make judgements about how to change test parameters to investigate phenomena observed during the testing process. Since it wasn't clear at the beginning of the test program what types of data analysis would be necessary, the structure of the software had to be flexible enough to allow easy addition of different analysis methods suggested by the test data.

The software program that emerged from this effort combined the low level control provided by C++ with the high level flexibility of scripting provided by Scheme (a dialect of Lisp). All of the hardware interface functions and special signal processing algorithms were provided by the former while the ability to change hardware operating parameters and the ways the data and signal processing operated on the data were provided through the scripting language. With this architecture we were able to control laboratory hardware through a GPIB interface at the outset of grant activities, migrate to control thorough a standard RS232 serial interface as breadboard hardware was developed, and ultimately provide a stand-alone controller using a 3COM Palm Pilot™ for the brassboard detection system.

Figure 26 shows the scan control window from the control system software whose functions evolved as grant activities proceeded. Examination of the figure shows functions associated with specifying data file names and directories, the frequency range of interest, number of points to scan, number of scans to record for a single test point, the type of scan being made (included as an identifier so that later analysis could make selections more intelligently), and the mode of scanning. Single stepping through a scan made it easier to troubleshoot the hardware system and delayed scanning made it possible for a single individual to setup a test and then participate as the subject being scanned. Upon the completion of a scan, the software automatically plotted the results for immediate visual analysis.

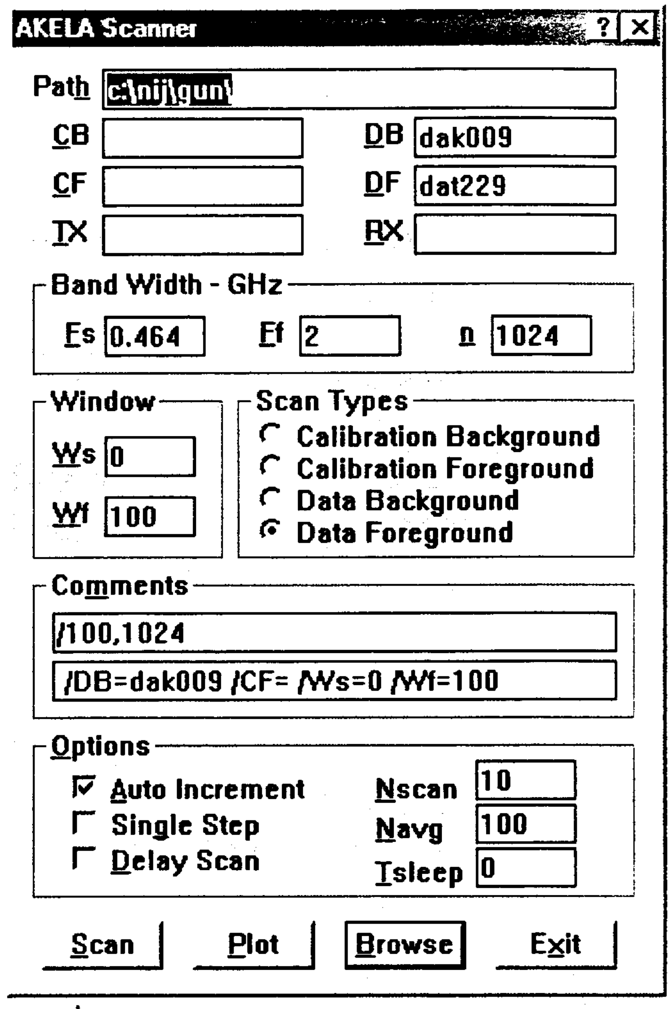

Figure 26 - Scan control window from software control system interface.

Figure 27 shows an example of the plot output of the software control program. By seeing results immediately, we were able to make decisions about adjusting test parameters to either improve system performance or to take additional data.

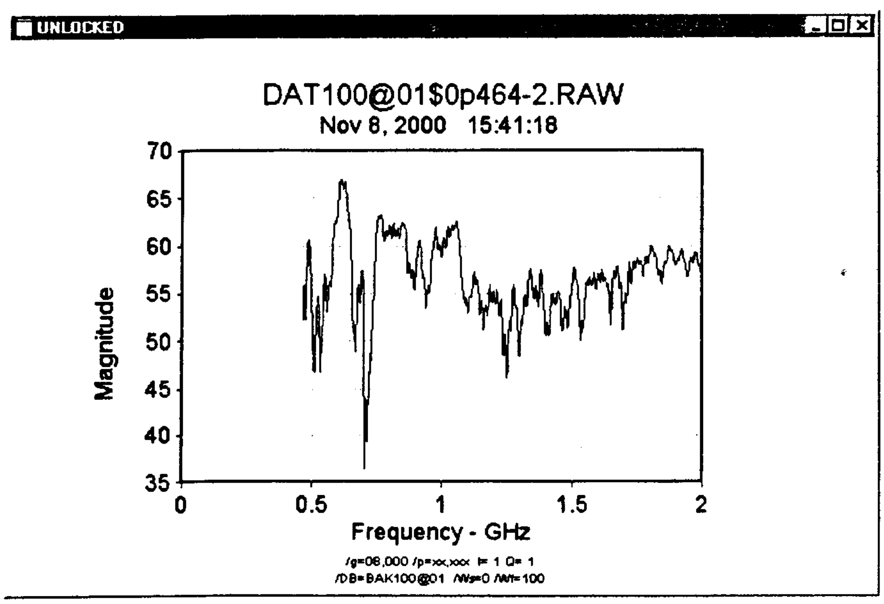

Figure 27 - Example frequency plot from test system control software.

A companion capability to this real time output was the ability to use the software to assist with analysis of the data. Figure 28 shows a browse control window that allowed us to select specific test data files, determine the type of plot to display, and specify a level of data analysis. As shown in the figure there were three analytical types that could be plotted either in the frequency or time domain. A raw plot showed the data just as it was collected from the test system. A cor(rected) plot was one that had a back-ground file subtracted from it. A nor(malized) plot was both background subtracted and normalized to the response of a calibration target to remove the effects of system parameters. Filter functions made it easy to search the information embedded in each file to select files with the desired characteristics from among a large amount of data.

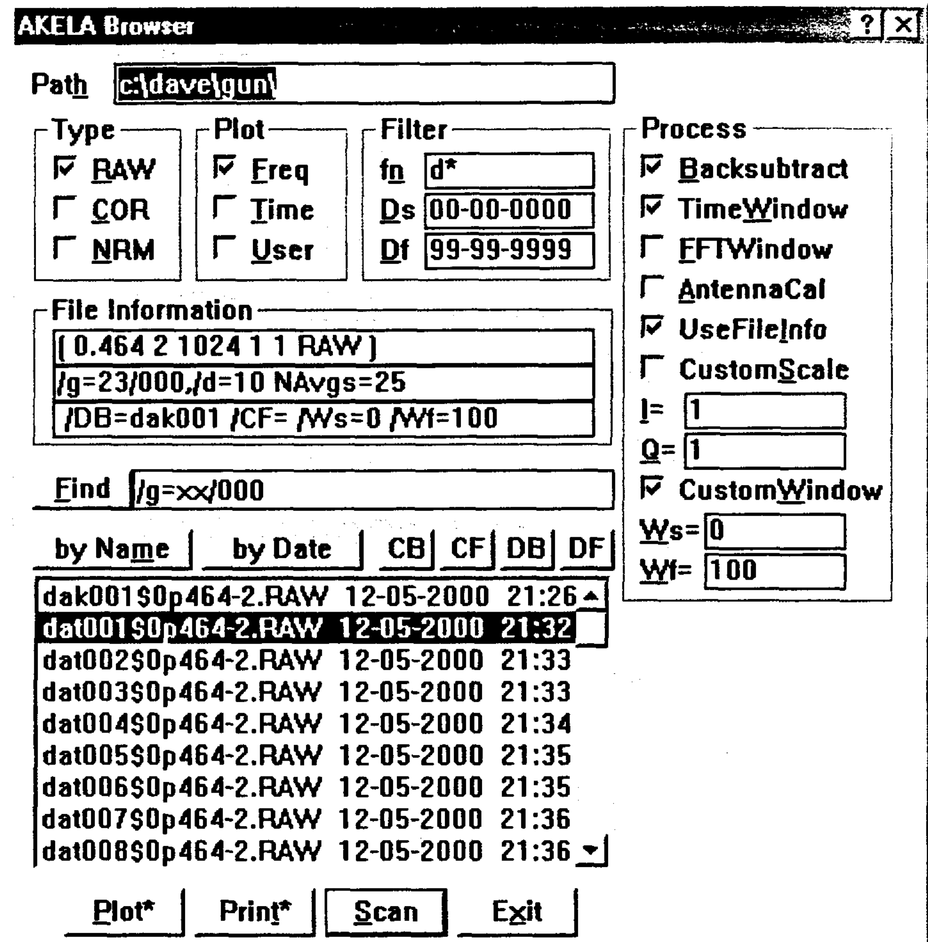

Figure 28 - Browse control window from software control system interface.

An additional set of processing and visualization functions was available directly through the plot window. Figure 29 shows the range of additional control that was possible. Low level functions such as changing the plot from a line plot to a scatter plot, enlarging a selected area of a plot, cross plotting one set of data with another, and automatically making the scales of two different plots the same enabled us to quickly discover both differences and similarities in the data that we collected. These functions could be performed while viewing the data in a variety of ways. Higher level signal processing functions let us view the data in the frequency, time, or phase domain, allowed us to see the effects of simple processing techniques such as data smoothing or averaging, and let us experiment with different methods of extracting features from the data.

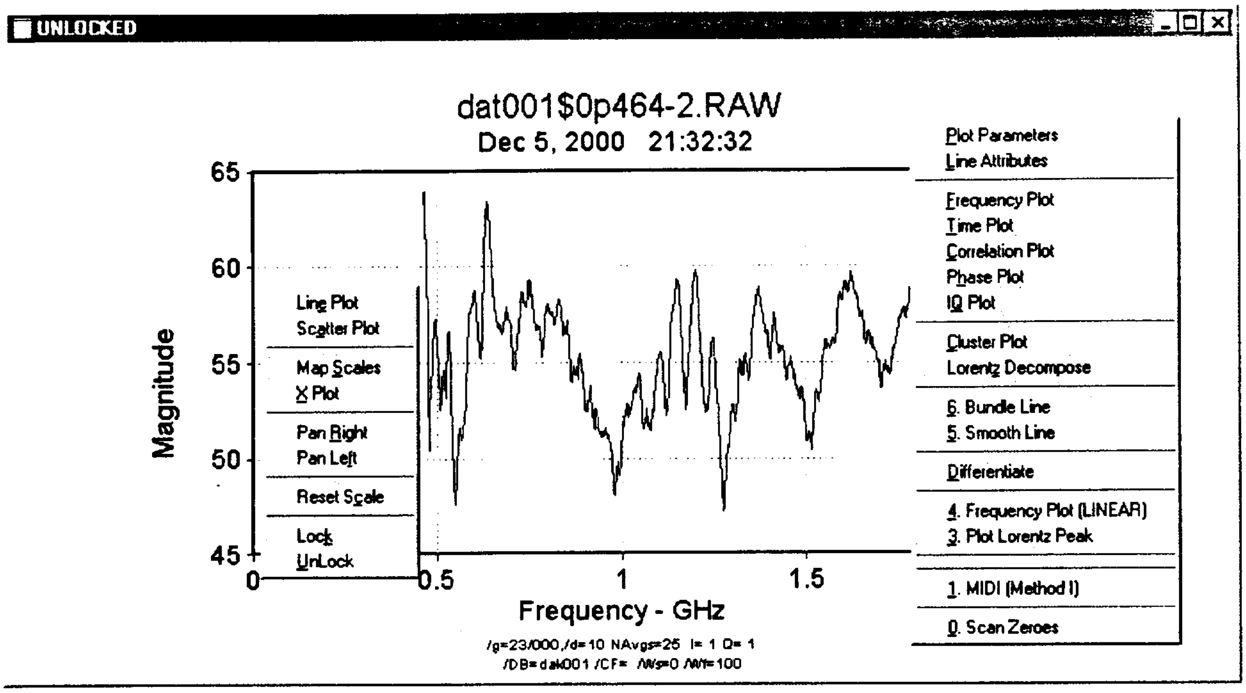

Figure 29 - Additional control options for test system control software plots.

While this interface worked well with both desktop and laptop computers, the requirement to make the brassboard detector portable led us to investigate alternative platforms to use for control and analysis of the detector signals. Initial examination of digital signal processor implementations showed serious limitations with respect to programming and display options. The introduction of palm top computers led us to investigate their use for controlling the brassboard system. We selected the 3COM Palm Pilot™ after an evaluation of the options because of its processing capability, display, availability of a suitable programming environment, and low cost.

The software that was developed for the Palm Pilot™ mirrored the capabilities of the PC control software. Its interface allowed us to take data with the detector, store and view the results in the Palm Pilot™, and later download the data to a computer. All of the processing elements in the Palm Pilot™ had to be implemented with routines that used integer operations. This allowed us to perform all of the processing and classification tasks in near real time. The resulting software gave us the ability to both take data for archive and training purposes and to make classification judgements with the Palm Pilot™ control unit.

Figure 30 shows the scanner interface developed for the Palm Pilot™. As shown in the figure, the user has the ability to completely configure the parameters of the detector by specifying the type of scan; distance to the test object; the identification and orientation of any people, weapons, or nuisances in the scan; the number of points in the scan; and the start and stop frequencies of the scan. The number of scans to be averaged is also selectable. This latter capability was added to allow the

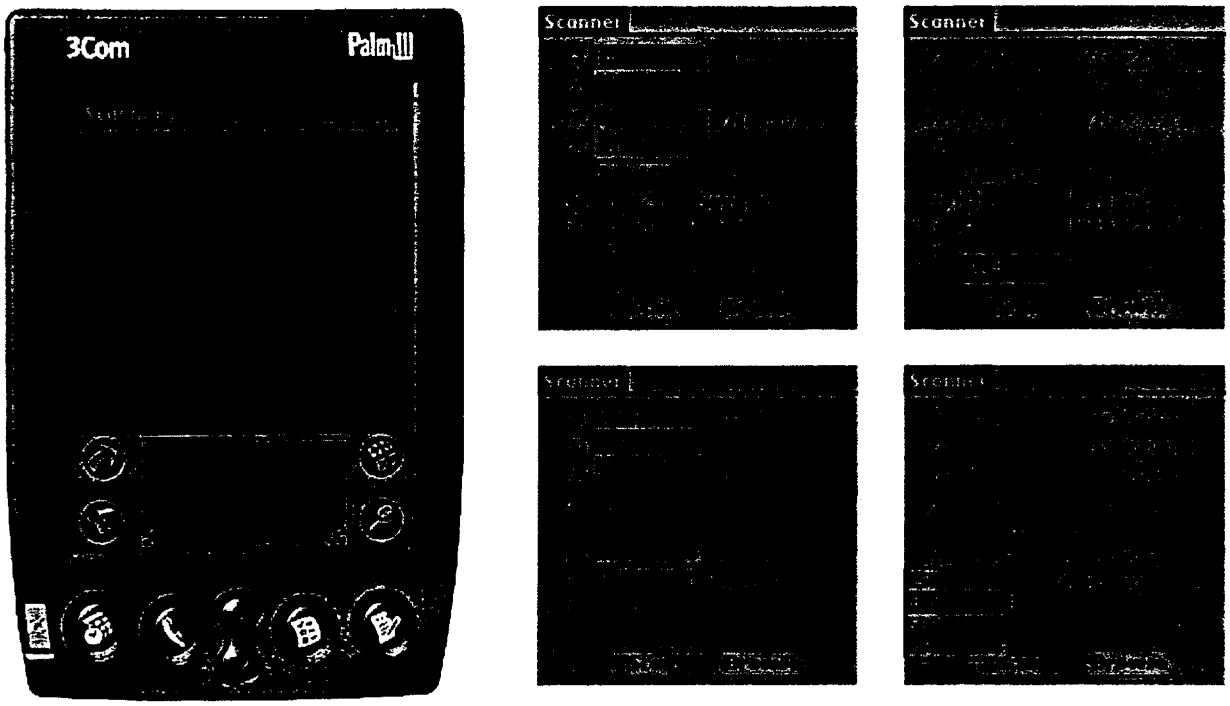

Figure 30 - Brassboard detector scanner interface and example menu selections.

user to take higher signal to noise ratio data. Each of the selections is made from a flyout menu. The available selections have been predetermined in order to ensure that the user selects combinations that are supported by the weapons detector hardware. This flexibility allows the detector to be used in the field as a data collection system. In addition, the Classifier option allows the user to put the weapons detector in a mode where as data is taken, a classification decision is made. These classifiers are developed on a desktop PC and then uploaded to the Palm Pilot™.

Figure 31 shows the browser interface developed for the Palm Pilot™. The browser lists all of the data that has been taken and stored by the detector and allows it to be reviewed. Selecting a data file and then pressing the plot button plots the raw frequency spectrum data. At the top of the plotter screen is a set of options that allows the user to toggle between frequency and time plots, select a window to plot in more detail, and pan right or left through the plot to examine sections in more detail. The Info button brings up a summary of information stored with the raw data that describes the individual scanning parameters for that scan.

All of the data that is collected with the brassboard detector is stored in the Palm Pilot™. The interface software includes a program that has been written that allows all of the data collected to be saved to a desktop PC during a standard Palm Pilot™ HotSync™ operation. The amount of memory in the Palm Pilot™ is sufficient to allow data collection to be performed over several days before data has to be downloaded to the PC desktop.

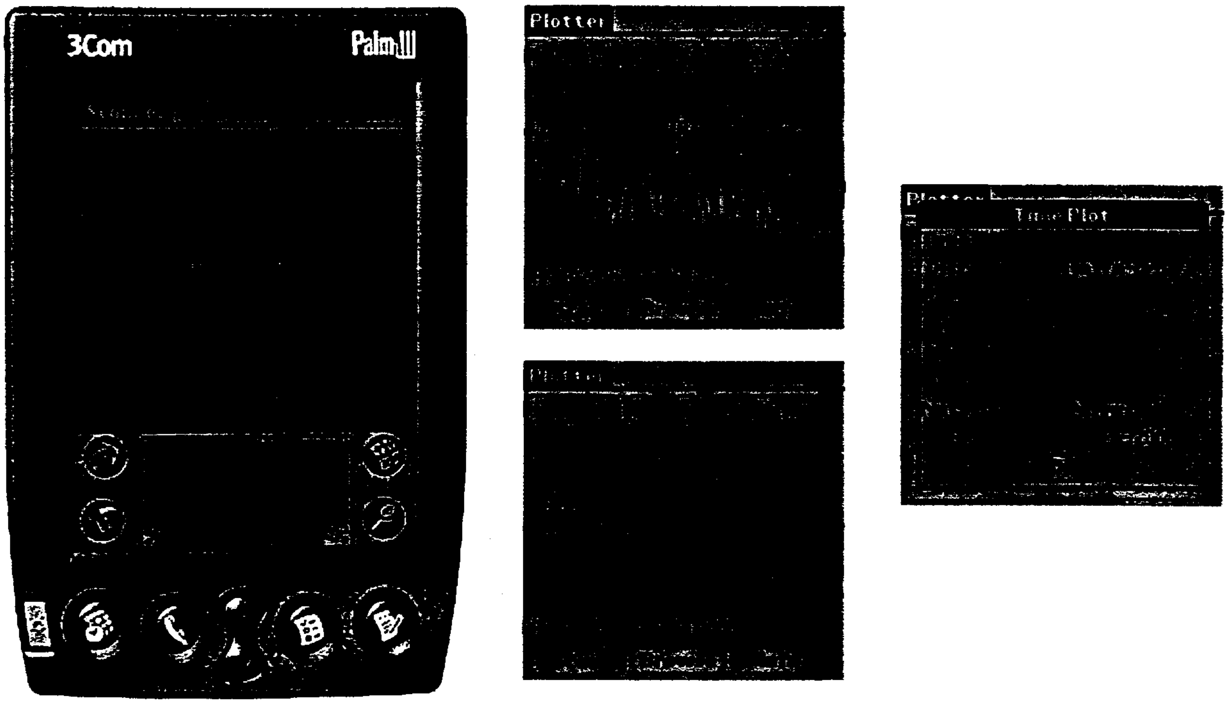

Figure 31 - Brassboard detector browser interface and example plot screens.

Data Analysis and Classifier Development

Data collection and analysis was performed continuously during the grant period but can be divided into three major phases. The first phase began as soon as both the time and frequency domain laboratory test systems were assembled, the second began with the completion of assembly of the first breadboard hardware, and the last commenced with the development of the first brassboard hardware.

Laboratory Test System

There were two areas of emphasis that guided data processing development during this phase of analysis. They were, 1) determining the effects of hardware and the operational environment on classifier performance, and 2) determining the adequacy of the signal preprocessing and classification techniques we planned for the detector system. While the classification results that we achieved during our SBIR program were quite good, we had found that they were sensitive to the number of features selected for training. Since we anticipated that the data we would be collecting with our laboratory system would pose additional challenges, we felt it was important to be able to determine in a more optimal way how many and what kinds of features were needed, how fine grained the features should be, and how all the possible combinations of features could be explored.

We selected a genetic algorithm technique to pursue and began by performing a reanalysis of our SBIR radar range data marrying our neural network software with the genetic algorithm. A primary reason for using this type of algorithm was to allow us to search a very large space of possible solutions in an efficient way. The genetic algorithm that we developed attempted to mimic the evolutionary process. We created a population of "creatures" each of which had a "gene" consisting of a random selection of the signal features from a data set. The gene (feature set) of each creature was used as input to the neural network training algorithm. We trained a classifier for each creature and then ranked all of the classifiers for how well they performed. The top performer represented the creature with the best initial gene. The process then repeated itself with the best creature mating with the others by exchanging parts of its gene. We also allowed parts of the genes of some of the creatures to be randomly changed to simulate a mutation. The next generation of classifiers was then trained, ranked, etc. and the process continued until we converged on a solution.

Twenty different classifiers were trained in this way. Our results showed that we were able to drive the false alarm rate close to zero without sacrificing any detection of weapons. Figure 32 shows a table comparing the results we obtained using this hybrid approach. In all cases using the genetic algorithm improved classifiers produced by our neural network software.

	ALN alone			ALN with GA		
Polarization	HH	45	VV	HH	45	VV
# Features	85	85	85	85	85	85
# Used	85	85	85	15	15	15
Goodness of fit	.4187	.4457	.4133	.3913	.3986	.3419
# Weapons detected	52	55	57	52	55	57
# Nuisances	15	16	16	15	16	16
# False Alarms	5	10	5	1	2	1

Figure 32 - Comparison of classifiers generated using neural networks and genetic algorithms.

The features we used represented the power in a narrow frequency band of the radar signature. Since the neural network generally only selected 6 to 10 features from the entire set which was used to train it, we hoped to use this to our advantage in designing a system that ignored the significant bands of interference that occur in the operational environment. Figure 33 shows a comparison of the frequency bands selected by the neural network and radar cross section data. What was unknown was the extent to which these frequency bands were operationally noisy.

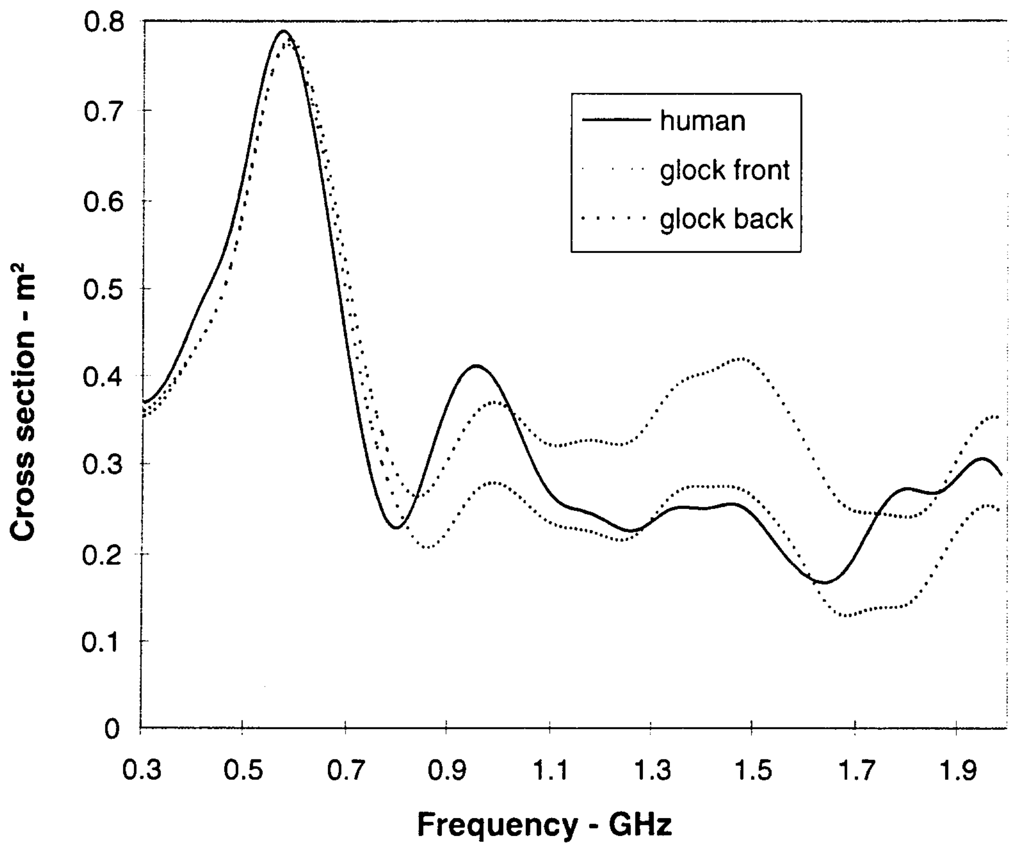

Figure 33 - Comparison of neural network frequency selection and radar cross section data.

We next began to examine the origin of noise sources which could degrade the performance of the classifier algorithms. These sources were caused both by the hardware and the operational environment and included undesired signals (e.g. radio and TV), undesired reflections, antenna feedthrough, and system noise. Since a big uncertainty with our wideband approach was the extent to which other signals in the electromagnetic environment would interfere, we made a series of measurements to characterize the extent of the problem. Figure 34 shows a measurement of the ambient signal spectrum overlaid with the frequency bands that the classifiers used. As can be seen from the figure some frequency bands that contain signature information also contain significant noise sources. Since both these undesired signals and system noise were relatively constant, we used their levels to set the output power level of the detector system.

Both antenna feedthrough and undesired reflections scale with power so different methods of dealing with them had to be considered. Geometry and antenna type were selected to deal with the former, while using a range gate was selected to deal with the latter. Reflections that come from close by the desired signal, however, had to be separated from the desired signal by the neural network classifier.

In order to get a better idea of how sensitive the neural network technique was to this type of interfering signal, we performed a series of processing experiments using the data from our SBIR pro-

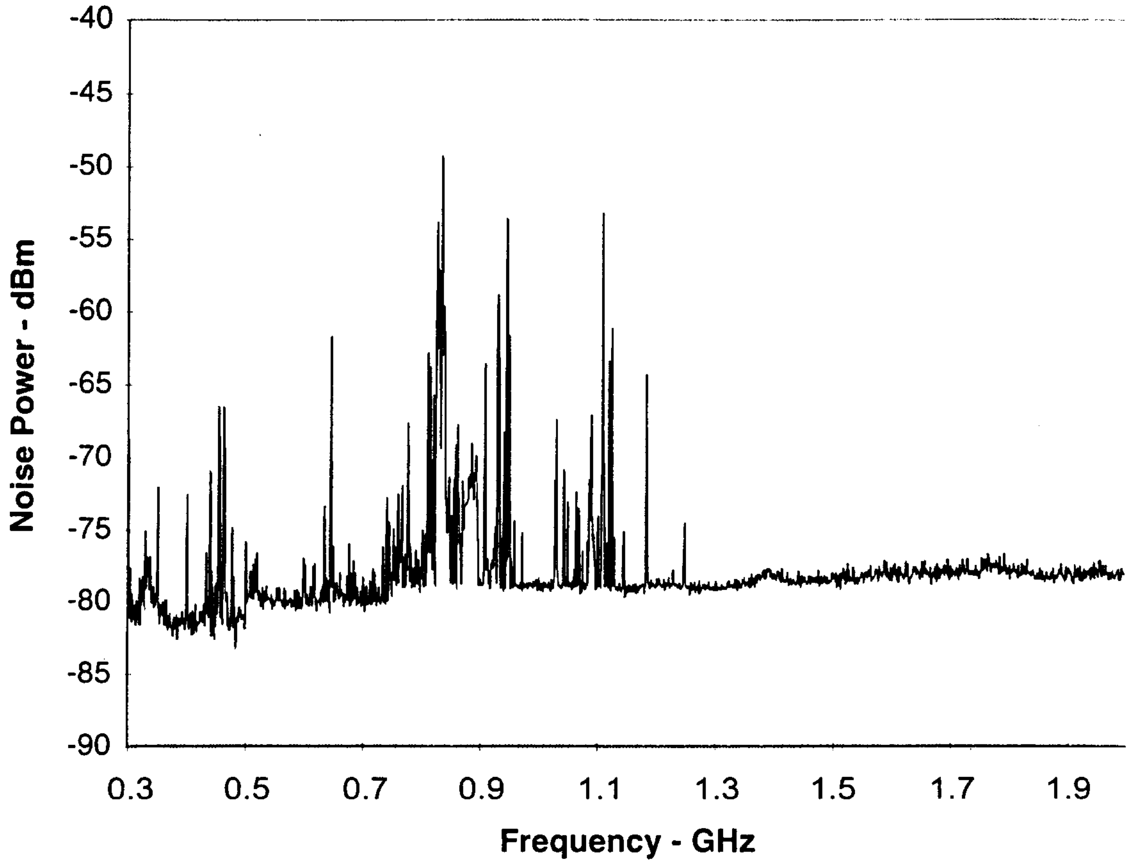

Figure 34 - Comparison of ambient noise spectrum and classifier frequency bands.

gram. We introduced various types of distortion into the data to simulate what we might see in operational data. Distortion types considered were a baseline offset, an offset ramp both increasing and decreasing with frequency, a sine wave, and random noise. The magnitude of the distortion was varied from 3% to 100%. Feature extraction and classifier development were performed for each combination. Figure 35 shows a table of the results of this processing experiment. As is evident from the table, a classification approach relying on frequency features alone looked to be very sensitive to signal distortion. As a result, our breadboard development activities focused on ensuring these types of distortion were not introduced by the hardware, and our signal processing activities expanded to consider other classification features.

Breadboard System

With the completion of breadboard hardware we were able to begin taking data in a more operationally relevant environment. Over 900 separate data points were collected in a variety of configurations that included a calibration target at various distances; two different individuals at distances of 10, 15, and 20 feet, in three different orientations (front, back, side); weapons at various distances; and individuals carrying weapons.

	Magnitude of Distortion						
	100%	50%	25%	12%	6%	3%	
No Distortion	100% 7%						
Offset	98	98	96	96	94	94	P_d
	100	60	47	47	47	47	P_{fa}
Ramp - right	84	94	98	96	96	94	P_d
	100	87	53	47	47	47	P_{fa}
Ramp - left	50	62	83	92	94	94	P_d
	13	20	27	27	47	40	P_{fa}
Sine wave	No valid result		60	90	96	98	P_d
			93	100	100	87	P_{fa}
Noise	63	98	96	94	98	96	P_d
	73	87	73	93	87	87	P_{fa}

Figure 35 - Classifier sensitivity to signal distortion.

This data was analyzed extensively in an attempt to determine the limits of performance of the breadboard hardware and the signal processing and classification methods. Because of the results of our earlier processing activities we began to use an additional feature extraction technique called Prony's method. Prony's method is a technique for directly extracting resonance information from a time domain signature. It is built around the assumption that the time signature is a sum of damped sinusoids and its application results in a set of features that describe the resonant characteristics of the signature. Each resonance is described by its frequency, amplitude, phase, and damping coefficient. This more complex set of time domain features was added to the frequency domain features and used to train neural network classifiers.

We used a complete set of data on two individuals at different distances and orientations, both armed and unarmed, to perform the classification processing. Each signature was transformed to the time domain and a fixed time window was applied around the peak in the data containing the person. Prony's method was run on the windowed data to extract the resonance information. This information was then used to create feature sets for the neural network to use to build classifiers. A matrix of different combinations of features (resonances alone, resonances with damping coefficient, etc.) was created and a family of neural network classifiers was trained.

Classifiers created using this technique combining data from all orientations gave mediocre results. Classifiers created using data characterized one orientation at a time performed much better. Training results indicated expected probability of detection $P_d = 1.0$ at a false alarm probability $P_{fa} = 0.1$. However, the performance of these classifiers on test data was "brittle". The results on test data sets were often worse than $P_d = 1.0$ at $P_{fa} = 0.5$.

This set of processing results highlighted several areas of improvement for the breadboard detection system. Perhaps the biggest was the importance of the antenna feedthrough signal. Even though time windowing was able to remove most of this noise component, FFT artifacts from the coupling at low frequencies tended to wash out the signal from the gun making classification difficult. In addition, the sensitivity of the time domain Prony's method technique to time window width and starting location, made it desirable to find an equivalent frequency domain technique for accomplishing the same result. Finally we learned that the neural network training algorithm seemed to work best with ordered feature sets. The four feature variables available from Prony's method application suggested that we seek algorithms that were better suited for pattern matching.

Brassboard System

In an effort to avoid the problems encountered with extracting resonance information using Prony's method in the time domain, a frequency domain technique was developed. To do this we used the fact that a damped sinusoid has the shape of a Lorentzian distribution in the frequency domain. We developed a curve fitting technique that successively found and subtracted Lorentzian curves directly from the frequency domain data until a lower limit magnitude was reached. Two features result from this method - frequency of the resonant peak and width of the resonance. Figure 36 shows an example of this curve fitting technique.

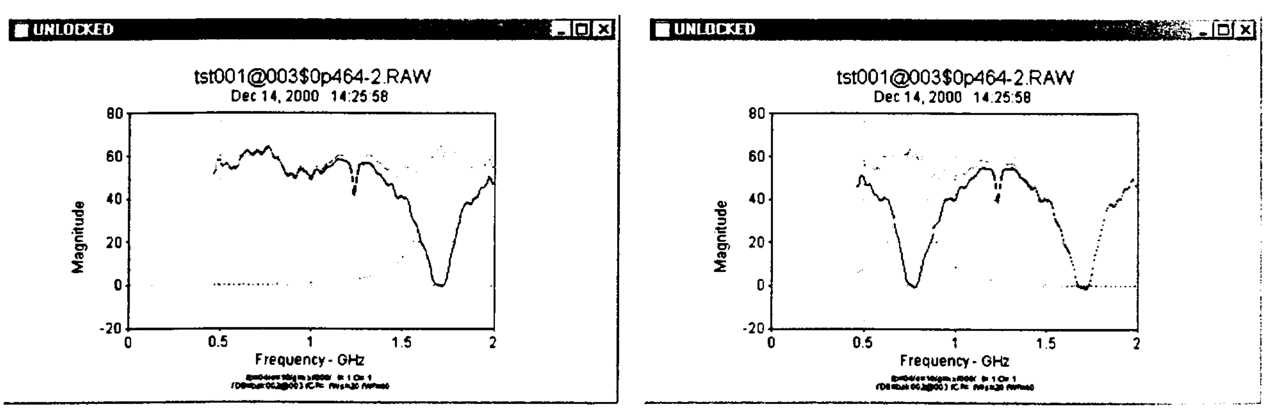

Figure 36 - Resonance extraction directly from the frequency domain.

This method was applied to an extensive set of multisensor fusion data that was taken with the brassboard system over a period of two days at Chang Industries. The test series was coordinated by the Air Force Research Laboratory and designed to provide a set of data that could be made available to other NIJ and Air Force contractors for performing signature processing. During the two days of testing over 650 data points were recorded on 3 different individuals and 9 different weapons. The weapons were supplied by the LA Sheriffs Department and consisted of 7 handguns and two plastic replica guns that had been confiscated by the Department.

Tests were performed at two different distances, in multiple orientations, and with the guns and people being tested by themselves and in combination. In addition, tests with different calibration objects were performed to provide baseline detector sensitivity data. Figure 37 shows some example data from this series of tests. This series confirmed the results of earlier testing that showed that a weapon changes the signature of a person, that weapons hidden behind the back are seen by the detector, and that body type and orientation change the signature of a person.

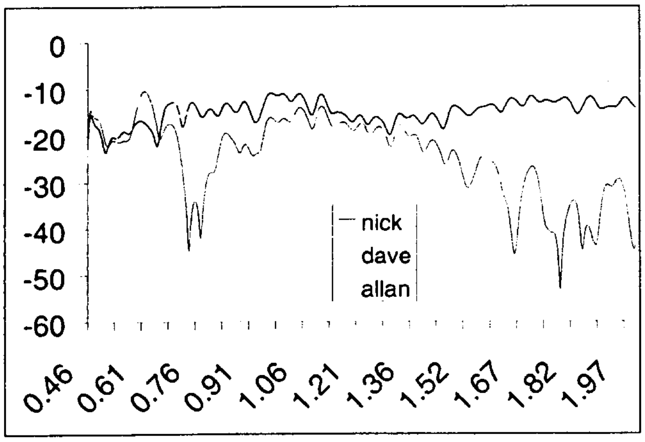

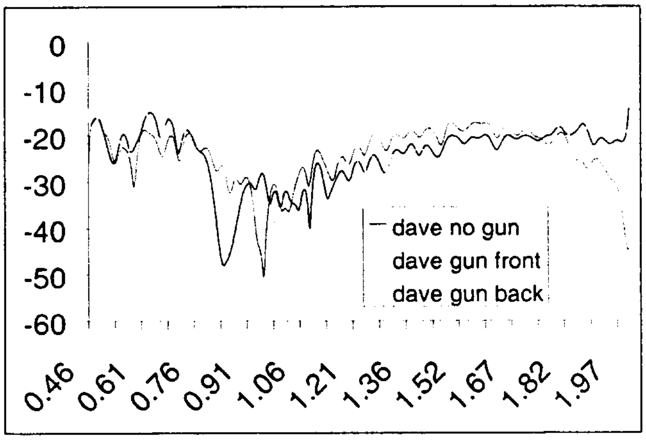

Figure 37 - Representative data from multisensor fusion test series.

All of this data was processed to extract the Lorentz features and a set of classifier algorithms was developed. Classifier training nearly always resulted in perfect separation between the weapon and no weapon case. However, when these classifiers were applied to the larger set of data the best performance achieved was a $P_d = 0.70$ at a $P_{fa} = 0.20$. The decrease in performance can be attributed to the small size of the signal contribution from the weapon in relationship to the size of the signal from the individual, and the variability in the signature of the individual. While not perfect detection, these results showed that there was a signal contribution from the weapon and that methods of enhancing it were needed.

An alternate classification method was also developed. Since the signatures associated with the presence or absence of a weapon create patterns, we decided that using a parametric technique for building a classifier would represent a good method for improving performance. The algorithm we chose was Singular Value Decomposition (SVD). SVD is a technique that is used for producing a best approximation to data in a least squares sense. It is particularly well suited for use on experimental data which is not always well behaved mathematically and results in a set of weights that describe the areas of the data that contain the most information for describing the signatures of interest. In effect, the SVD algorithm discovers the common features in the training data that can best differentiate the signatures of people carrying weapons from those that are not.

In order to develop a classifier using SVD, the algorithm was presented with a set of "training" data which consisted of representative detector signatures of individuals both with and without a gun. Separate training was performed on the gun and no gun cases. The output of the algorithm was a set of weights for each data set that represented the relative importance of each part of the signature. Figure 38 shows a comparison of the weights for both the gun and no gun case. The data used for training was taken by the brassboard unit set to scan the 450 MHz to 2 GHz region in 256 individual frequency steps. As can be seen in the figure the differences in the weights are small and confined to specific regions. In this case the differences show up in the 500 - 700 MHz, 1.25 - 1.4 GHz, and 1.6 - 2 GHz frequency bands.

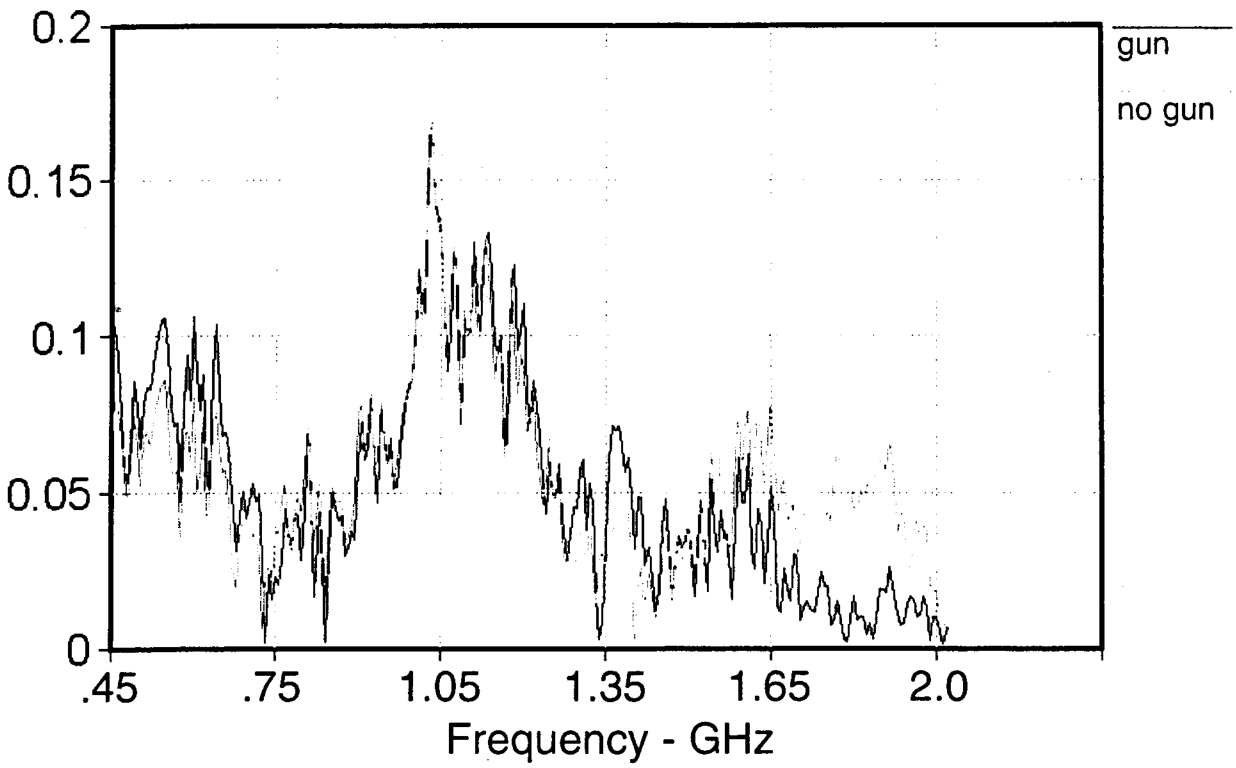

Figure 38 - Classifier weighting functions from SVD algorithm.

To determine the success of the SVD training, the weights are multiplied on a point by point basis with the training data. A summation across the resultant multiplications gives a number that represents the goodness of fit between the training data set and the individual signature. Figure 39 shows an example of the separation achieved using SVD training on 50 signatures equally divided between gun and no gun cases. Each of the circles and squares in the figure represents a separate signature. For this set of data the SVD classifier achieved a $P_d = 1.0$ at a $P_{fa} = 0.04$.

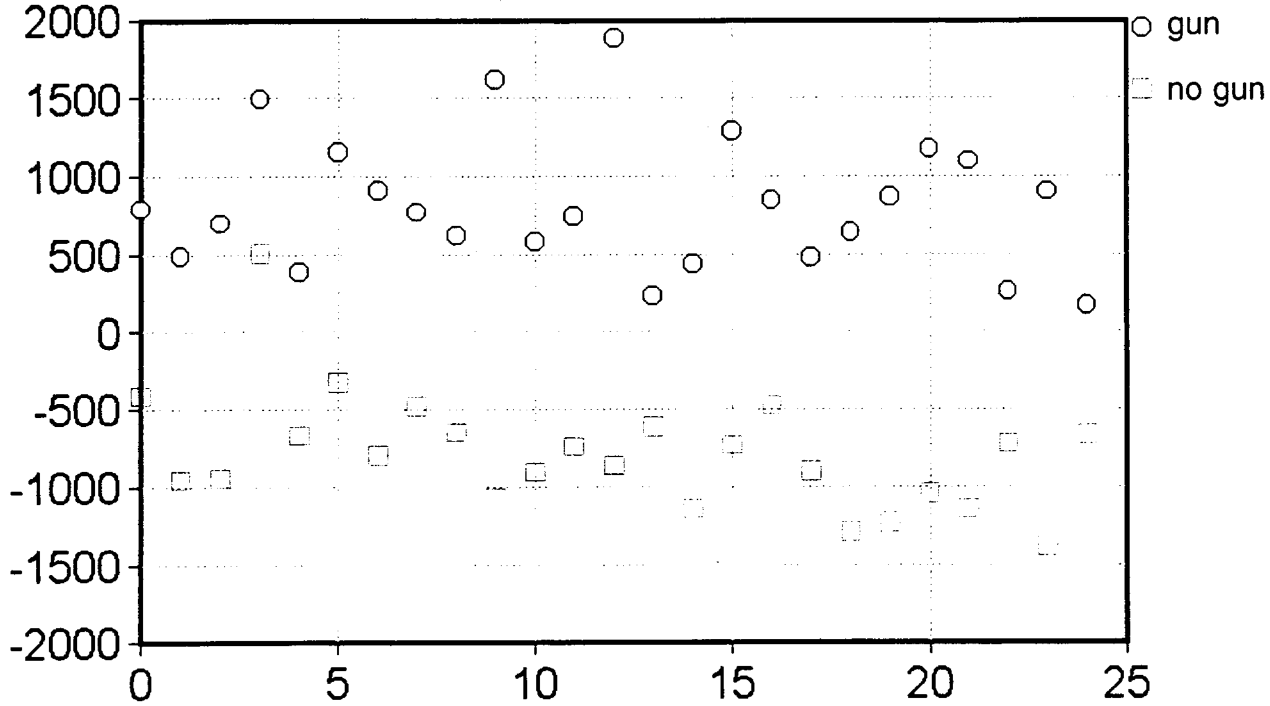

Figure 39 - Classifier separation using SVD algorithm.

The availability of a fully integrated, portable brassboard unit at this point in the program allowed us to collect an operationally realistic set of data that could be used to stress the classification processing algorithms. We collected over 450 data points on three different individuals. Two separate sets of data were taken, one to use exclusively as training data and the other to use as test data.

All of the training data was taken at a constant distance of 10 feet with the individual facing the weapons detector. We recorded 50 tests for each individual, 25 unarmed and 25 armed. In order to ensure that there would be sufficient variability in the data, after each signature was taken the test subject moved away from the test area and then back into place. The inability to get back in an identical position both in terms of distance and posture created differences in the signature for each test.

The test set data was taken in a similar manner, however, we included many test configurations that were not present in the training data set. This was an attempt to give the classification algorithm a very different type of data to identify and a means to determine how robust the classification method was. We took 105 separate data points for each individual. In this case we varied the distance between 8 and 12 feet, included orientations where the detector looked at the side and back of the individual, and included postures where the individual's arms were extended straight out from the side and straight up over the head.

We used these two sets of data with the Lorentz and SVD algorithms to perform 70 different classification experiments. The experiments were grouped into two different types. In the first, we used the Lorentz algorithm on a set of training data to try to determine the resonant regions characteristic of both the gun and individual. These regions were then used to limit the data that was given to the SVD algorithm for training and subsequent classification. This was an attempt to improve the signal to noise ratio of the data prior to SVD training and classification by limiting the signature pattern matching to only those areas where we would expect the most distinguishing information to be present. The second type used all of the data in each signature for the SVD training. In this case we let the SVD algorithm select those areas through its weighting function.

Figure 40 shows a table of all of the classification experiments that were performed. The individual columns in the figure describe the conditions of each experiment including the data that was used for testing the specific classifier, the two classifiers that were used to test the data against, whether the experiment used classifiers trained on segmented (Lorentz defined) signatures, the total number of data points being classified, and the classification results.

The first set of experiments used only the gun in various orientations and was designed to determine how well the classifier algorithms would train and classify on data of the same type. Classifiers were trained on odd numbered tests and tested against even numbered tests. With the exception of experiments 9 and 10 the classifiers had perfect detection capability. In experiments 9 and 10 we classified data with the gun in all orientations against a classifier trained on data from all orientations and one trained on data from only the 0 degree orientation. All of the missed detections were for data at the 0 degree orientation. This indicated that the classifier trained on 0 degree orientation data was better at classifying 0 degree data than the classifier that had been trained on all angles. This result is what should be expected from the experiment.

The final experiment in the first set was number 11. In this experiment we tested the gun data at all orientations against a gun classifier and a cell phone classifier. In this case the detection was perfect and the separation between the classifiers was significant indicating a very robust classification capability.

Case	Data being tested	Yes condition comparison	No condition comparison		Total Data Points	P_d	P_{fa}
1	gun 23 0 degrees odd	gun23 0 degrees even	gun23 90 degrees even	segmented	13	100	0
2	gun 23 0 degrees odd	gun23 0 degrees even	gun23 90 degrees even		13	100	0
3	gun 23 180 degrees odd	gun23 0 degrees even	gun23 90 degrees even	segmented	5	100	0
4	gun 23 180 degrees odd	gun23 0 degrees even	gun23 90 degrees even		5	100	0
5	gun 23 90 degrees odd	gun23 90 degrees even	gun23 0 degrees even	segmented	5	100	0
6	gun 23 90 degrees odd	gun23 90 degrees even	gun23 0 degrees even		5	100	0
7	gun 23 270 degrees odd	gun23 90 degrees even	gun23 0 degrees even	segmented	5	100	0
8	gun 23 270 degrees odd	gun23 90 degrees even	gun23 0 degrees even		5	100	0
9	gun 23 all angles odd	gun23 all angles even	gun23 0 degrees even	segmented	28	35	64
10	gun 23 all angles odd	gun23 all angles even	gun23 0 degrees even		28	53	46
11	gun 23 all angles odd	gun23 all angles even	cellphone 0-90 degrees even	segmented	28	100	0
12	Allan w/gun odd	Allan w/gun even	Allan w/o gun even	segmented	12	66	33
13	Allan w/gun odd	Allan w/gun even	Allan w/o gun even		12	75	25
14	Allan w/gun odd	Allan w/gun all	Allan w/o gun all		12	75	25
15	Doug w/gun odd	Doug w/gun even	Doug w/o gun even	segmented	12	75	25
16	Doug w/gun odd	Doug w/gun even	Doug w/o gun even		12	100	0
17	Doug w/gun odd	Doug w/gun all	Doug w/o gun all		12	100	0
18	Terry w/gun odd	Terry w/gun even	Terry w/o gun even	segmented	12	75	25
19	Terry w/gun odd	Terry w/gun even	Terry w/o gun even		12	83	16
20	Terry w/gun odd	Terry w/gun all	Terry w/o gun all		12	83	16
21	Allan test w/gun	Allan w/gun even	Allan w/o gun even	segmented	60	38	61
22	Allan test w/gun	Allan w/gun even	Allan w/o gun even		60	45	55
23	Allan test w/gun	Allan w/gun all	Allan w/o gun all		60	46	53
24	Allan test all	Allan w/gun even	Allan w/o gun even		105	62	38
25	Doug test w/gun	Doug w/gun even	Doug w/o gun even	segmented	60	56	43
26	Doug test w/gun	Doug w/gun even	Doug w/o gun even		60	65	35
27	Doug test w/gun	Doug w/gun all	Doug w/o gun all		60	61	38
28	Doug test all	Doug w/gun even	Doug w/o gun even		105	68	32
29	Terry test w/gun	Terry w/gun even	Terry w/o gun even	segmented	60	75	25
30	Terry test w/gun	Terry w/gun even	Terry w/o gun even		60	70	30
31	Terry test w/gun	Terry w/gun all	Terry w/o gun all		60	66	33
32	Terry test all	Terry w/gun even	Terry w/o gun even		105	73	27
33	Allan test w/gun	all people w/gun even	all people w/o gun even	segmented	60	38	61
34	Allan test w/gun	all people w/gun even	all people w/o gun even		60	45	55
35	Allan test w/gun	all people w/gun all	all people w/o gun all		60	46	53
36	Allan test all	all people w/gun even	all people w/o gun even		105	60	40
37	Doug test w/gun	all people w/gun even	all people w/o gun even	segmented	60	68	31
38	Doug test w/gun	all people w/gun even	all people w/o gun even		60	70	30
39	Doug test w/gun	all people w/gun all	all people w/o gun all		60	66	33
40	Doug test all	all people w/gun even	all people w/o gun even		105	70	30
41	Terry test w/gun	all people w/gun even	all people w/o gun even	segmented	60	61	38
42	Terry test w/gun	all people w/gun even	all people w/o gun even		60	75	25
43	Terry test w/gun	all people w/gun all	all people w/o gun all		60	75	25
44	Terry test all	all people w/gun even	all people w/o gun even		105	56	44
45	all people w/gun	all people w/gun even	all people w/o gun even	segmented	180	56	43
46	all people w/o gun	all people w/o gun even	all people w/gun even	segmented	135	48	51
47	all people w/gun	all people w/gun even	all people w/o gun even		180	63	36
48	all people w/o gun	all people w/o gun even	all people w/gun even		135	58	41
49	all people all data	all people w/gun even	all people w/o gun even		315	62	38
50	all people facing front w/gun	all people w/gun even	all people w/o gun even		150	65	34
51	all people facing front w/o gun	all people w/o gun even	all people w/gun even		105	60	39
52	all people facing front	all people w/gun even	all people w/o gun even		255	64	36
53	all people w/gun	Allan w/gun even	Allan w/o gun even		180	52	47
54	all people w/o gun	Allan w/o gun even	Allan w/gun even		135	59	40
55	all people all data	Allan w/gun even	Allan w/o gun even		315	56	44
56	all people facing front w/gun	Allan w/gun even	Allan w/o gun even		150	57	42
57	all people facing front w/o gun	Allan w/o gun even	Allan w/gun even		105	56	43
58	all people facing front	Allan w/gun even	Allan w/o gun even		255	57	43
59	all people w/gun	Doug w/gun even	Doug w/o gun even		180	58	41
60	all people w/o gun	Doug w/o gun even	Doug w/gun even		135	63	36
61	all people all data	Doug w/gun even	Doug w/o gun even		315	61	39
62	all people facing front w/gun	Doug w/gun even	Doug w/o gun even		150	58	41
63	all people facing front w/o gun	Doug w/o gun even	Doug w/gun even		105	68	31
64	all people facing front	Doug w/gun even	Doug w/o gun even		255	63	37
65	all people w/gun	Terry w/gun even	Terry w/o gun even		180	52	47
66	all people w/o gun	Terry w/o gun even	Terry w/gun even		135	66	33
67	all people all data	Terry w/gun even	Terry w/o gun even		315	58	42
68	all people facing front w/gun	Terry w/gun even	Terry w/o gun even		150	53	46
69	all people facing front w/o gun	Terry w/o gun even	Terry w/gun even		105	75	24
70	all people facing front	Terry w/gun even	Terry w/o gun even		255	62	38

Figure 40 - Classification experiments and results.

In experiments 12 through 20, we used data that had people and weapons in the same signature and focused on single individuals in an attempt to determine the sensitivity of the classification training technique to different body types. Classifiers were developed using both the segmented and unsegmented data and tested against the data from the training set that had not been included in classifier development. Examination of the results of this set of experiments shows several trends. First, for each of the three individuals, the classifier trained using segmented data points performed more poorly than that using all of the data points. While we expect that it will be necessary to ultimately limit the areas of the frequency spectrum where we look for information that will assist in helping us identify the difference between an armed and unarmed individual, the method we used to select the areas was not as good as letting the SVD algorithm do the selection itself.

Second, we were able to achieve perfect detection for only one of the three individuals and the detection results were different for each individual. This indicates that there is sensitivity to body type. In addition, the separation between the with and without gun cases was very small. This indicates that there is a substantial contribution to the classifier of the body component of the person and that some means of improving the signal to noise ratio of the gun is necessary to improve detection performance. Finally, using all of the data for training did not improve the classifiers. Operationally this is important should it become necessary to develop a wide range of classifiers to embed in the weapons detector, since the amount of data that would need to be collected to support classifier development would be reduced.

Experiments 21 through 32 were repeats of experiments 12 through 20 except that the classifiers were used to classify the test set data described earlier instead of a portion of the training set data. This was a more rigorous test of the classifier training technique and more representative of the performance that could be expected operationally. As in the previous experiments, the classifiers trained using only preselected segments of the training data performed more poorly than the classifiers trained using unsegmented data and those trained using all of the training data performed no better than those using only some of the training data. However, for the test data, classifier performance was not as good for each individual as it was for the training data. This matched our expectations since the test data included signatures where the body position had been changed significantly.

The results of experiments 24, 28, and 32 are important to highlight. In these experiments, all of the test data for each individual was presented to the classifiers for separation. In all cases the results showed that the classifiers had a good probability of detection. Specifically, for experiment 24 the results were $P_d = 0.62$ at $P_{fa} = 0.38$, for 28 they were $P_d = 0.68$ at $P_{fa} = 0.32$, and for 32 they were $P_d = 0.73$ at $P_{fa} = 0.27$. Although these classifiers did not achieve perfect detection, they were generally correct. These results show that there is unique information within the signatures taken by the weapons detector that can be used to distinguish armed from unarmed individuals, but that the signature from the body contributes significantly to the classifier. This reinforces the observation that finding a way to increase the signal to noise ratio of the weapon component of the signature is very important to improved detector performance.

This last set of experiments was repeated in experiments 33 through 44 with the exception that the test data for each individual was classified by a set of composite classifiers trained on data from each of the three individuals. The important experiment results to examine in this case are for experiments 36, 40, and 44 where the classifiers were tested on the entire set of test data for a single individual. For experiment 36 the results were $P_d = 0.60$ at $P_{fa} = 0.40$, for experiment 40 they were $P_d = 0.70$ at $P_{fa} = 0.30$, and for experiment 44 they were $P_d = 0.56$ at $P_{fa} = 0.44$. As expected, the results are not as good as those obtained using specific individual classifiers, but they still result in positive detection. Improving signal to noise ratio of the gun in the signatures would improve these results and potentially lead to a classifier that was not as body shape dependent.

The remainder of the classification experiments focused on using the test data from all of the individuals with a variety of single and multiple individual classifiers. Experiments 45 through 52 tested data from all of the individuals against classifiers built from data from all of the individuals. The results of experiments 49 and 52 are the most relevant from this set of experiments. In experiment 49 the results were a $P_d = 0.62$ at $P_{fa} = 0.32$. In experiment 52, only data where the individual was facing the detector was used and the results improved slightly to $P_d = 0.64$ at $P_{fa} = 0.36$. Since the classifiers were trained using data where the individual was facing forward, the small improvement in performance indicates that it is body type more than body orientation that is important at this level of signal to noise ratio for the gun.

Experiments 53 through 70 tested classifiers trained on individuals against data from all the individuals. The results of experiments 55, 61, and 67 should be compared to that of experiment 49. The best classification result for all of the test data classified by a single individual classifier was that of experiment 61 where the $P_d = 0.61$ at $P_{fa} = 0.39$ was slightly worse than experiment 49 where $P_d = 0.62$ at $P_{fa} = 0.32$. In experiments 58, 64, and 70 the test data was limited to the forward facing signatures. The best result in this case was for experiment where the $P_d = 0.63$ at $P_{fa} = 0.37$ was again just slightly worse than experiment 52 where $P_d = 0.64$ at $P_{fa} = 0.36$. As above, limiting the test cases to those where the individual faced forward only marginally improved the classifier performance.

Perhaps the most important conclusion to draw from the classification experiments is that there is information buried within the signatures taken with the weapons detector that can be used to detect weapons. However, at the current levels of performance, the false alarm rate in an operational situation would probably not be acceptable. Additional techniques need to be applied to reduce the effect of the body so that a more robust classifier can be trained.

Conclusion

The activities performed under this grant have been quite successful. We have developed a brass-board detection system, as shown earlier in Figure 1, that performs consistently in a field environment. It has been used to show that it is possible to detect concealed weapons on individuals at distances of 15 feet with a range resolution of 4 inches, and that it is possible to detect weapons being concealed behind the back while taking only one second to collect and classify a test. While not hand held, it is portable, self contained, and can operate under battery power for 8 hours.

We have shown the feasibility of reducing overall system size by using smaller planar antennas, however the gain penalty for this type of antenna is significant. Unfortunately the grant level of effort was insufficient to pursue a smaller, handheld design based on a single antenna of this type. Even so, the power output of the brassboard detector is 100 mW which is well within the range for human safety, and it is likely that the manufacturing cost of a single unit would be between $500 and $1000.

Antenna coupling limits the signal to noise response of the detector. Figure 41 shows a plot of a raw signature and the same one that has had the background subtracted from it. The gain of the weapons detector is set to avoid saturation by the antenna coupling that is predominant at low frequencies. As shown in the figure, this coupling component is about 20 dB. Reducing the coupling component can

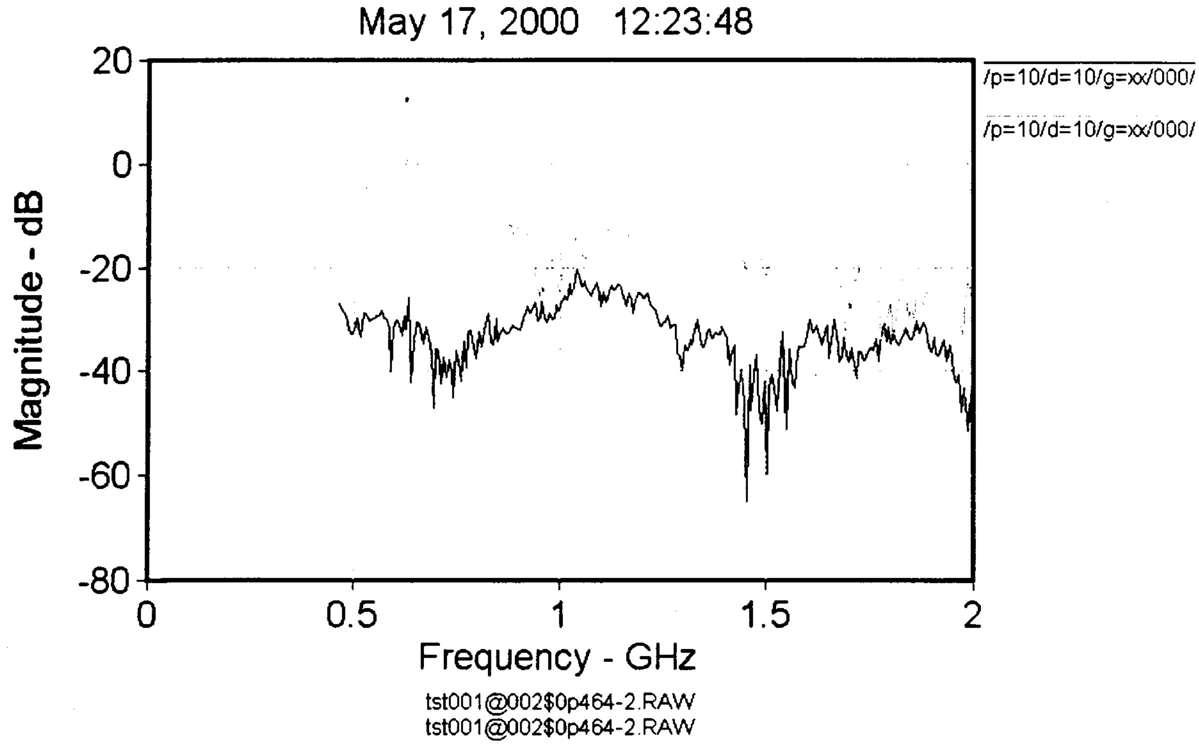

Figure 41 - Weapons detector signatures showing the effect of low frequency antenna coupling.

be accomplished either passively or actively and would allow system operation at greater ranges. For current antennas we have reached the passive limit.

The weapons detector operates reliably and consistently under field conditions. Figure 42 shows a comparison of 6 different background signatures taken with the brassboard detector over a three hour period. As can be seen in the figure, the data collected is very repeatable.

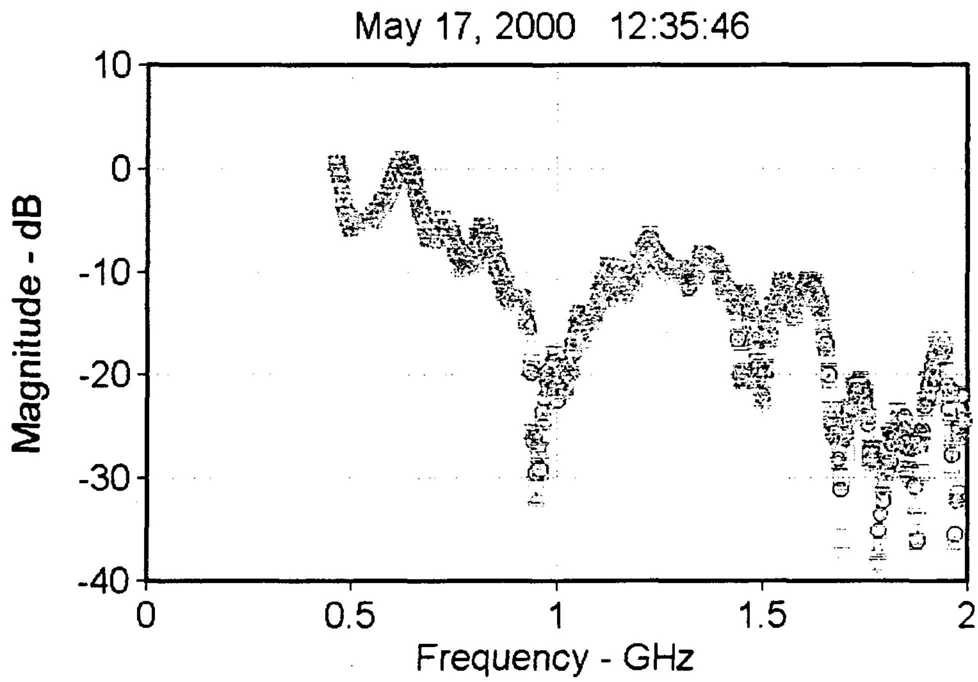

Figure 42 - Weapons detector data showing consistency of performance.

Because of this consistency we discovered that the signatures of different people vary widely and that the signature of a single individual varies from test to test. Figure 43 shows plots of three identical tests taken 30 seconds apart with each individual trace an average of 8 taken over a 2 second period. As can be seen in the time plot, the position of the signature shifts slightly in time indicating movement associated with tilting and swaying of the individual. These time shifts show up as differences in the frequency spectrum as well.

Similar results are seen when comparing the signatures of different individuals. Figure 44 shows a plot comparing the signatures of two different individuals of similar build. As seen in the plot the signatures are quite different. This variation may be due to changes in body water content in combination with movement. While portions of the signatures look similar, the large variations in the different signatures cause the processing algorithms problems. Additional testing is needed to better understand the origins of this signature variability.

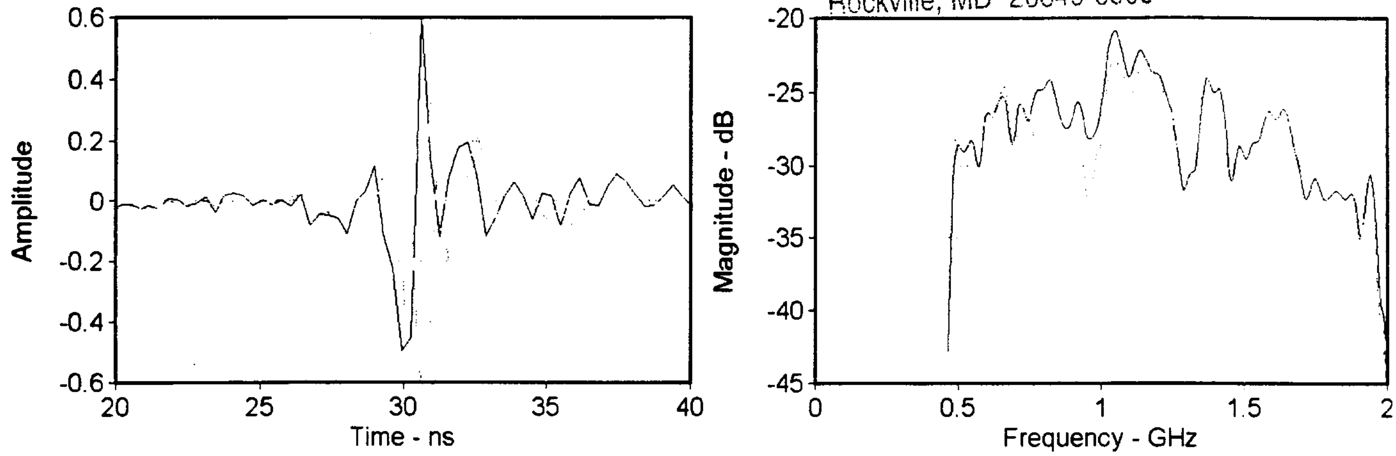

Figure 43 - Signatures of the same individual vary.

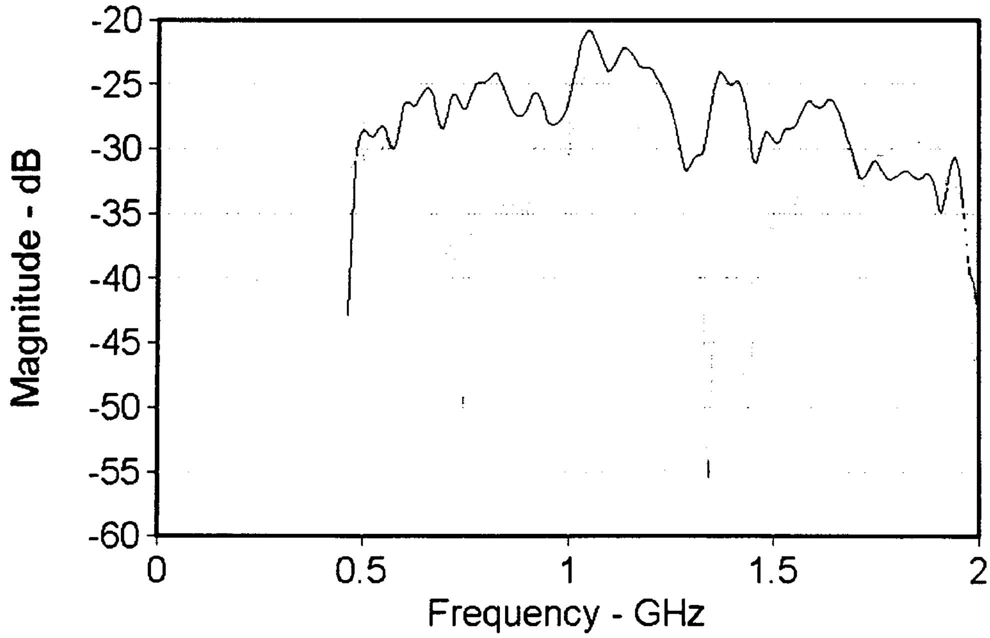

Figure 44 - Signatures of different individual of the same build vary.

While the signatures of individuals vary from one another, they also vary when a weapon is present. Unfortunately, the signature of an individual with a weapon is very similar to one without a weapon. Both of these effects had an impact on the success of our signal processing activities. The signals due to weapons are buried in the large and varying "noise" component due to the person. This makes relying on pattern matching as a means of classification unreliable.

Our signal processing and classification experiments have shown that there is information in the weapons detector signatures that can be used to detect weapons. The best results we were able to achieve were a $P_d = 0.64$ at $P_{fa} = 0.36$ for a classifier built to classify all individuals, and $P_d = 0.73$ at $P_{fa} = 0.27$ for a classifier built to classify a single individual. This level of false alarm rate is currently too high for operational use. It will be necessary to reduce the effect of the body signal so that of the weapon stands out in order to ensure better detection performance. While the detection problem has not yet been solved, we have encountered no fundamental limits to success. Additional effort needs to be applied to signal processing and hardware development techniques that can be used improve the detection performance.