Using Ultrasound in Concealed Weapons Detection

John Edward DuChateau
Prof. Mark Hinders
NDE Lab, Department of Applied Science, College of William and Mary

April 2005

Table of Contents

1. Table of Contents
2. Abstract
3. Introduction
4. Physics Background
    I. Ultrasound
    II. Passage of sound through media
    III. Transducers
5. Description of Methods
    I. Equipment used
    II. Construction of Target
        A. Finding the best transducer
        B. Fabric
        C. Tissues
    III. Proof-of-Concept with Weapons
    IV. Using Wavelet Transforms for Data Interpretation
        A. Experimental Setup
        B. Computer programs and settings used
6. Data and Analysis of the Wavelet Transform graphs
7. Results
8. Conclusions

Acknowledgements
Bibliography

Appendix
    A. The maxcurve and clothmax programs
    B. Reflection coefficients between different materials and air
    C. Corbin's data on ballistic gel
    D. Key wavelet graphs

Abstract

We have explored using ultrasound, that is, sound that occurs above 20 kHz, for the purpose of detecting weapons that may be concealed on a human body beneath clothing. We first set up an apparatus with which to conduct experiments, finding which frequency suits our purposes and building a suitably realistic target to scan. From these initial experiments, we found that ultrasound is feasible as a weapons detection method, at least at short distances. We then continued our work by concentrating our efforts on using frequency analysis techniques to analyze our data. As a result, we found patterns within the data using said techniques which may be useful in future weapons detection algorithms.

Introduction

Since the unfortunate events of September 11, 2001, the scientific and security communities have focused on finding different and accurate ways to detect weapons and explosives concealed on a person's body. Current technology in that field includes x-rays and metal detectors. In our experiments, we explored using ultrasound as an alternative to these two detection methods. The advantage of ultrasound over magnetic metal detection is that ultrasound does not need the object being examined to be ferromagnetic; it can detect materials regardless of their magnetic properties. It also offers a cost-competitive alternative to x-rays and does not involve ionizing radiation. However, in order to use ultrasound for such purposes, personnel using ultrasound technology or a machine equipped to send and receive ultrasound signals must know what to look for in an ultrasound signal that is backscattered off a target. Thus, the main focus of our experiments was to figure out what part of a return signal contains the data concerning whether or not the target is armed and how that data can be read and interpreted best.

Physics Background

I. Ultrasound

Ultrasound is sound that occurs beyond the upper limit of human hearing, 20,000 Hertz (or 20 kilohertz). Ultrasound can reflect off of objects and return to its source, a property that is used by bats to detect prey. Bats can also use ultrasound to determine characteristics of their target, such as the size and shape of the target, and how far away it is. Bats use ultrasound in the range of 30 to 200 kHz [1], which is the same range in which we started our investigation.

II. Passage of sound through media

If we define a coefficient called the acoustic impedance of a material as the product of its density and the speed of sound through the material

$$Z = \rho \times c \tag{1}$$

then we may define another coefficient called the reflection coefficient of the material relative to air (R), which can be expressed as

$$R = \left| \frac{Z_{mat} - Z_{air}}{Z_{mat} + Z_{air}} \right|^2 \tag{2}$$

In Eq. (2), Z_{mat} is the acoustic impedance of the reflecting material and Z_{air} is the acoustic impedance of the conducting material, the air (0.0004 megarayleighs). The reflection coefficient of a material describes how much energy is reflected off an object when a wave encounters it as an interface (*i.e.* the wave encounters it as the wave travels

through another material, which was air in this case). As one can see in Eqs. (1) and (2), the higher the density of the reflecting object, the more energy will be reflected back to its source-in our case, the transducer.

III. Transducers

To produce ultrasound beams, we used a transducer which converts electrical energy to sound energy, and vice-versa [2]. It uses a piezoelectric crystal which changes shape when an electric potential is applied across it; conversely, when its shape is changed, it generates an electric potential. Thus, the piezoelectric crystal is capable of generating an ultrasonic pulse from an electric signal. When the pulse reflects off an object and returns an echo to the transducer, the echo hits the crystal and distorts it, which generates an electric signal that can be converted to digital data that is stored in a computer file. Transducers include other necessary components, such as a damping material behind the crystal to shorten the length of the emitted pulse, and an impedance-matching material to control the flow of energy into and out of the transducer. Fig. 1 is an illustration of a cut-away view of a transducer. Fig. 2 shows the four transducers that we used in our experiments.

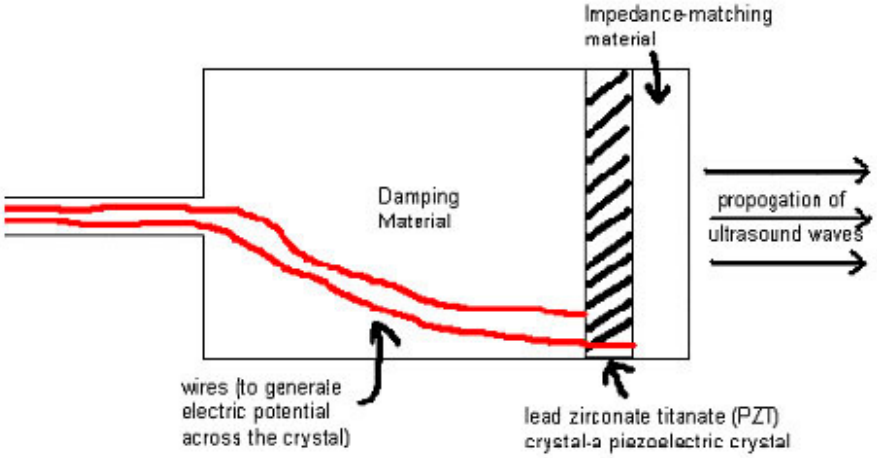

Figure 1: The basic structure of a transducer, as seen from the side. The important parts of the transducer are labeled in the diagram.

(a)

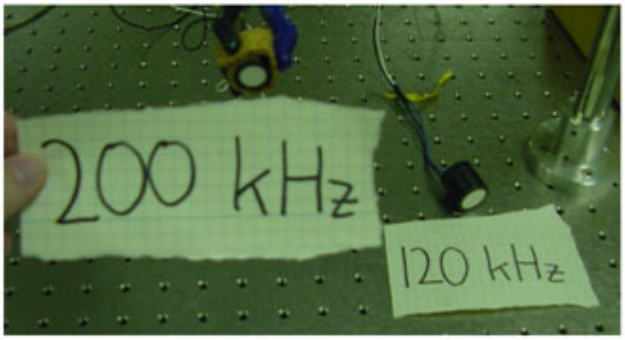

(b)

Figure 2: Photographs of the four transducers that we used. In (a), we see the 41 and 50 kHz transducers. In (b), we see the 200 and 120 kHz transducers.

Description of Methods

I. Equipment and Technology

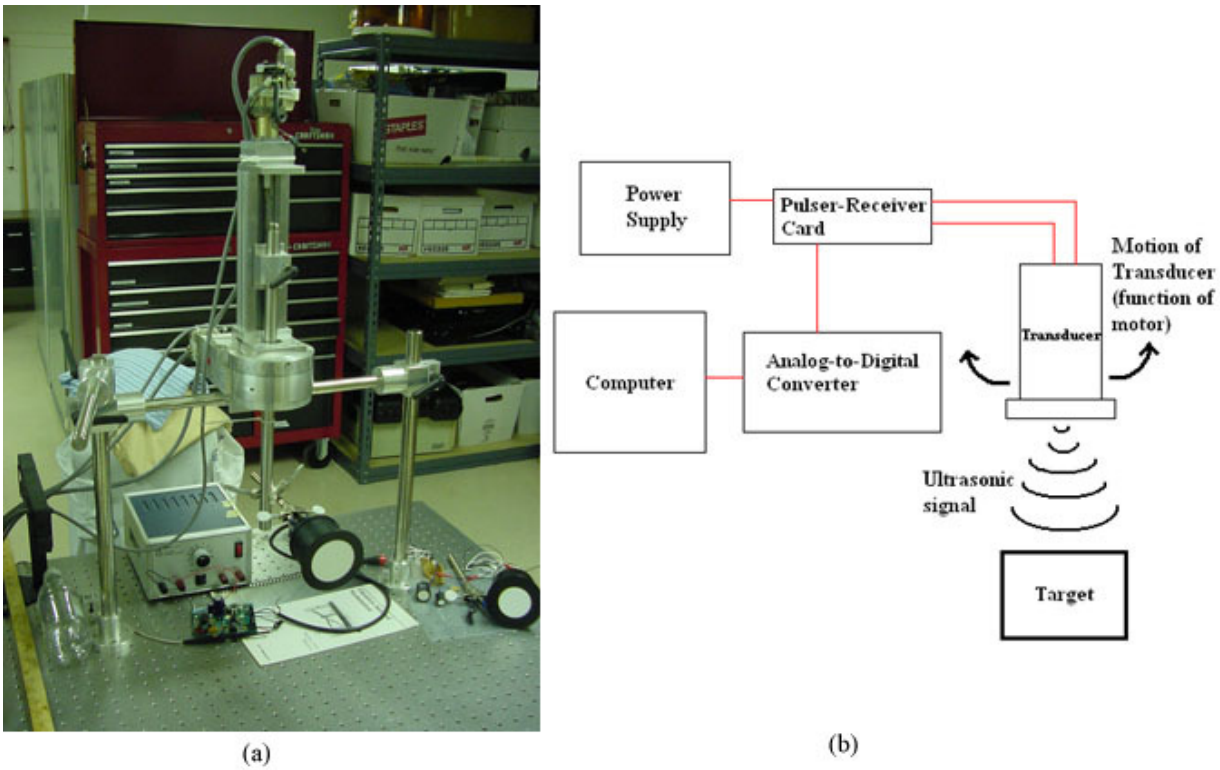

(a) (b)

Figure 3: (a) A photograph of our scanning setup. Included in the photo are the power supply, the pulser-receiver card, the transducer, and the motor. In (b), we see a diagram of the entire experimental setup, except for the motor (what the motor does, however, is clearly labeled).

Fig. 3 shows a photograph of our scanning setup and a diagram of the setup. Our scanning setup included the power supply to generate the electric potential, a pulser-receiver card to deliver potential to the transducer and send electrical energy generated by the return signal to the analog-to-digital converter, and the analog-to-digital converter to convert the return signal to digital data, which is sent to the computer. There is also a motor which sweeps the transducer across whatever target is in front of it. LabVIEW is used to control the motor and the operation of the transducer, as well as collect data. At

the start of our experiments, LabVIEW was set to tell the transducer to take 20 scans of whatever is in front of it, moving through 10 degrees. The transducer swept 10 degrees to the left from looking straight ahead (0°), taking scans at every angle from 0° to 9°, then returned to its starting position and repeated the process. From these scans, LabVIEW collected sample points, which are data points that can represent either distance or time. In our experiments, we took sample points to represent distance. 200,000 sample points were initially collected by LabVIEW for every scan, which was filtered into 2000 sample points per scan and stored in a text document as a matrix of 20x2000 elements. The 20 rows are the angles and the 2000 columns are the sample points. We later reset LabVIEW to take 400,000 sample points per scan, which were filtered down into 4000 sample points per scan and stored in a text document as a matrix of 20x4000 elements (the number of scans can also be changed, which affects the number of rows in the matrix, but we will discuss that later).

This data file is easily analyzed and graphed by a mathematical analysis program called MATLAB, which can produce a graph of the data called an A-scan. A-scans are graphs of the signal data that display distance from the transducer on the horizontal axis versus signal amplitude on the vertical axis (examples will be shown in Fig. 5). Every A-scan shows what we call "initial ringing" in the first few hundred sample points caused by the transducer and the machine themselves. We found that the smaller the transducer and the higher the frequency, the fewer sample points the ringing covered.

II. Construction of the Target

A. Finding the best transducer

In the first phase of our experiments, we searched for the transducer with the best balance between attenuation (covering an appreciable distance without much beam energy loss and penetrating clothing) and resolution (being able to clearly make out details of concealed weapons). In other words, we wanted a balance between distance and detail. We used four transducers of four different frequencies: 41 kHz, 50 kHz, 120 kHz, and 200 kHz. We attached a linear retroreflector (pictured in Fig. 4) to a trolley to use as a reflecting surface for the ultrasonic waves emitted from the transducers.

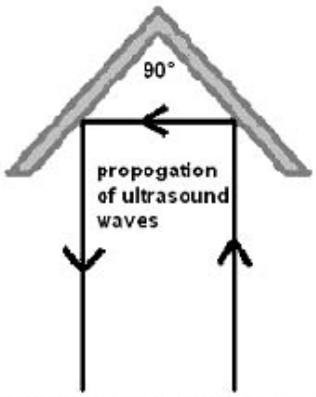

Figure 4: The linear retroreflector, as seen from the top.

For each transducer, we measured the signal returning to the transducer for the following distances between the retroreflector and the transducer: 0.3 meters, 0.7 meters, 1.2 meters, 1.8 meters, 2.5 meters, 3.3 meters, and 4.2 meters. We selected 0.3 meters as our first point because it was the closest we could get the trolley to the transducer. We wanted appreciable distances between each point, so we chose 0.7 meters as the next point, a 0.4 meter distance between the two points. We decided to make the next point 0.5 meters away from 0.7 meters (1.2 meters), and the next point after that 0.6 meters (1.8 meters), and so on. Later, we realized that we needed more points to graph the data, so we chose 0.5 meters, 0.9 meters, 1.1 meters, and 1.3 meters.

After we collected the data from all the different frequencies and distances, we plotted the angle of the transducer versus the maximum signal strength detected by the transducer (excluding ringing). This was accomplished using the maxcurve program, a custom program developed in MATLAB (detailed in Appendix A). Its parameters are the frequency of the transducer and the file name under which the data is stored. The files were stored in the following format:

cornerscan_<x>m_<y>khz

Where <x> is the distance from the retroreflector to the transducer (1.2 meters is written as 1_2m), <y> is the frequency of the transducer (120khz), and 2000 sample points were collected. In other words, the data collected by the transducer when it emits waves at a frequency of 120 kHz and the retroreflector is 1.2 meters away would be stored in the file cornerscan_1_2m_120khz. If the file is created with 4000 sample points, then the format is:

corner4000sp_<x>m_<y>khz

Displayed in Fig. 5 are A-scans taken when the retroreflector is 1.2 meters from the 50 kHz transducer. The left A-scan was created from a file based off of a 2000 sample points per scan system, while the right A-scan was created from a 4000 sample points per scan system.

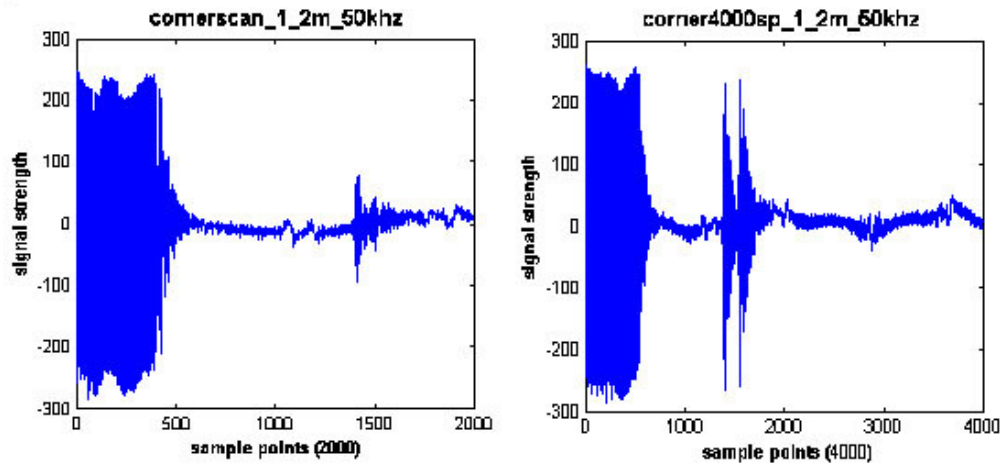

Figure 5: A-scans created from data collected by the 50 kHz transducer when the retroreflector (pictured in Fig. 4) was 1.2 meters from the transducer. The only difference between the scans is the distance viewed by the transducer, as is seen in terms of sample points.

From many A-scans like these, we calculated the sample-point to distance conversion to be about 1100 sample points per meter. Notice that the large signals from the retroreflector occur a little before sample point 1500, which is (1/1100)*1500=1.36 meters from the transducer, close to the observed 1.2 meters. The A-scans in Fig. 5 are those taken when the transducer was pointing directly at the retroreflector; there are actually 20 of these which, taken in sequence, show a gradually decreasing return signal, as is seen in Fig. 6.

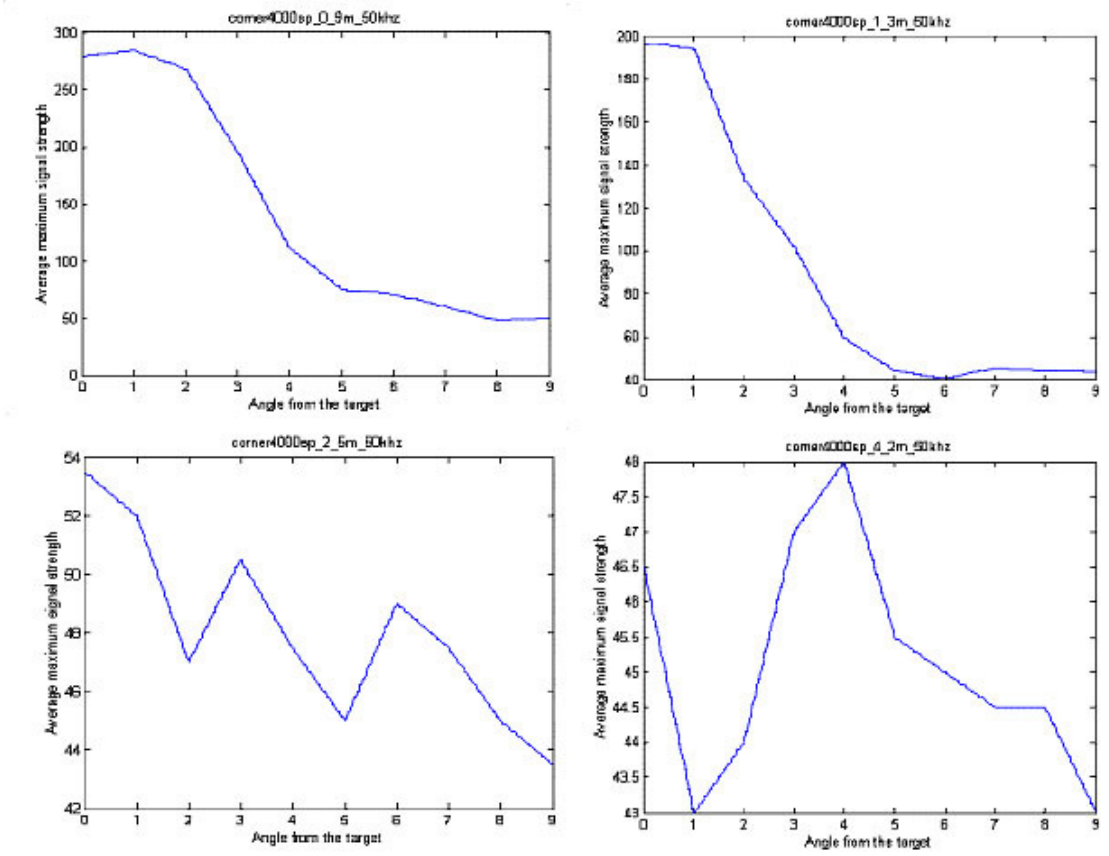

Figure 6: The angle versus signal strength graphs. The maxcurve program was used to construct these graphs. To construct them, the strongest non-noise signal from each of the 10 angles were averaged and plotted on the vertical axis, with their corresponding angles on the horizontal axis.

The angle versus signal strength data graphed as a half-bell curve in Fig. 6, with the strongest signals being at 0° and the weakest being at 9°. As the retroreflector is placed farther from the transducer, the bell curve shape breaks down and the noise takes over as the highest signal. From calculations of the conversion of sample points to meters (1100 sample points per meter, as mentioned above), 2000 sample points only cover a real-world range of about 1.8 meters. This can be improved with 4000 sample points, which cover 3.6 meters, though the signal may be drowned out by noise at such distances.

After the angle plots were collected, we plotted distance versus reflected signal strength for each transducer. Since the transducer points directly at the retroreflector at

0°, and should thus have the strongest return signal from the retroreflector, we plotted the 0° points' corresponding signal strengths on these new graphs. We plotted the four graphs on one page for easy viewing. We also fit a curve to the four graphs, which can be seen in Fig. 7.

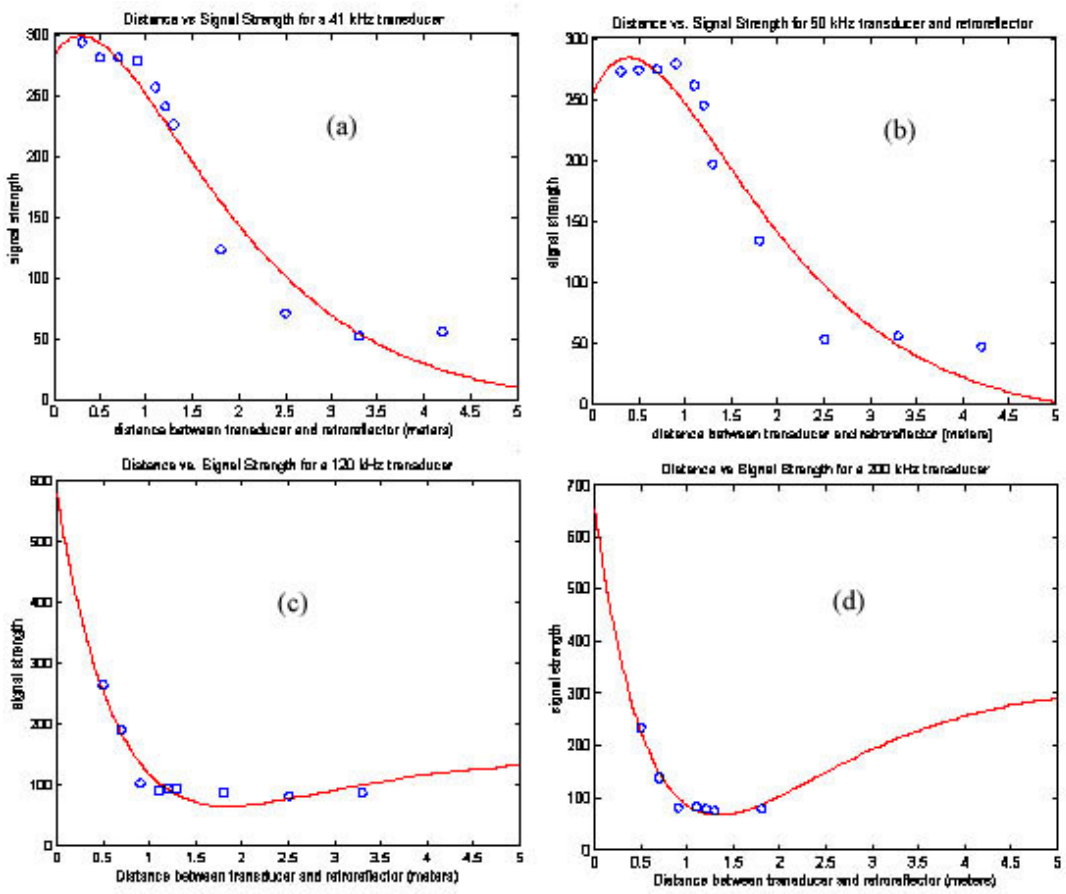

Figure 7: The distance-versus signal strength graphs for (a) 41 kHz, (b) 50 kHz, (c) 120 kHz, and (d) 200 kHz. In each of these graphs, the average of the strongest non-noise return signals at 0 degrees were plotted on the vertical axis versus their respective distances on the horizontal axis. From these graphs, we see that (b) 50 kHz has the best balance between resolution and attenuation.

From these graphs, we see that 200 kHz is essentially useless-it starts out with a good strong reflected signal at short distances, but quickly drops into noise even before the distance gets to 1 meter. We get slightly better signals with the 120 kHz transducer, but not by much; the signal drops to noise level at around 1 meter. These transducers have

great resolution-they can "see" the object with more detail-but they can only see for a short distance. The 41 kHz and 50 kHz transducers, on the other hand, cover more distance (i.e. smaller attenuation), but have less resolution. We determined the best balance between resolution and attenuation between the four transducers used in this experiment was the 50 kHz transducer.

Further experiments in this phase involved changing the settings on the pulser-receiver card. One thing we did to the card was move the jumpers, which drives the signal based on the desired frequency (in this case, the desired frequency was the transducer frequency). Initially, the card had been set to drive a 120 kHz signal, but this was changed to drive the signal of each transducer. In addition, we rewired the card to amplify the return signal by only 100 times instead of its initial 1000 times and we minimized the pulse width by adjusting it with a small knob on the card. This eliminated the “double image” effect that we saw in the first A-scans, and that is illustrated by the A-scan of corner4000sp_1_2m_50khz presented in Fig. 5. These changes to the card had the effect of cleaning up the signal and allowing it to “see” a little farther, but overall, we came to the same conclusion: that 50 kHz has the best balance between attenuation and resolution (the graphs in Fig. 7 result from the jumper changes). From then on, we used 50 kHz in all of our experiments, except in explicitly stated cases where we used 41 kHz.

B. Fabric

In the real world, no criminal or terrorist that is trying to conceal a weapon is going to go without clothes. So the next step in our experiments was to scan the linear retroreflector with and without fabric in front of it and compare the two scans.

To perform these experiments, we used a total of six fabrics. The first was a stretchy, blue striped polyester. The second was a dull yellow cotton. The third was a smoother, floral print cotton. The fourth was a stretchy dull-gray polyester-cotton mix. The fifth was a blue denim. The sixth "fabric" was saran wrap (a thin plastic film which was suggested by our colleagues).

After scanning the retroreflector with and without the fabrics in front of it, we used a different form of the maxcurve program, called clothmax, to analyze the data (also detailed in Appendix A). This program plots the angle of the transducer from dead-on the retroreflector versus the maximum signal strength at each angle for the retroreflector with the fabric ("clothed") and without the fabric ("naked") on the same graph. From these graphs, we did not learn much. Below one meter, the graphs were chaotic because of near-field effects. Above one meter, the expected half-bell curve was observed. The cloth always seemed to drown out the signal from the retroreflector. A few of these graphs can be seen in Fig. 8.

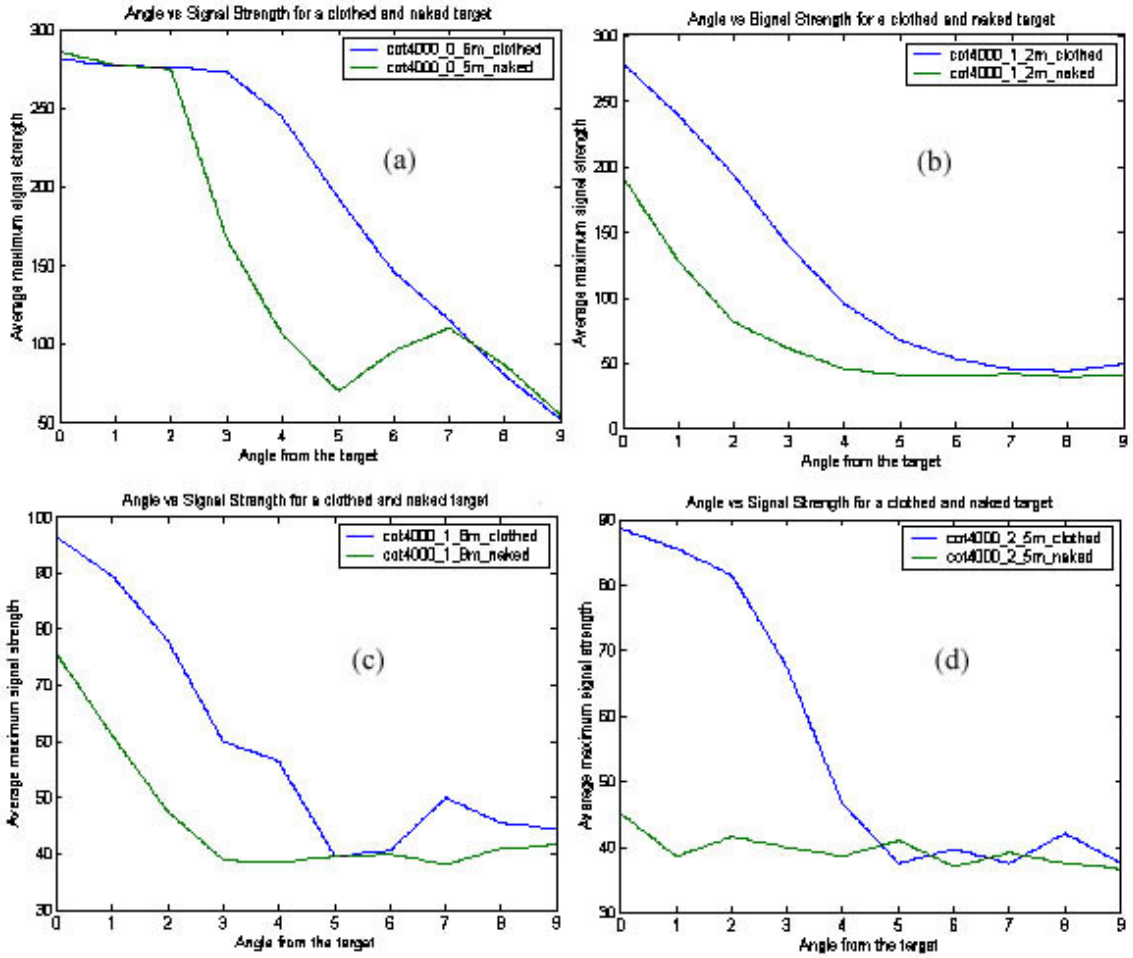

Figure 8: The angle versus signal strength graphs for a fabric-covered target and non-fabric-covered target, colored in blue and green, respectively. The program used to create these graphs was clothmax. The fabric used was cotton, and the distances depicted are (a) 0.5 meters, (b) 1.2 meters, (c) 1.8 meters, and (d) 2.5 meters from the transducer. Otherwise, the graphs are just like those in Fig. 6.

The blue lines are the retroreflector with the cloth (in this case, cotton) in front of it, while the green lines are the retroreflector by itself. There does not seem in these graphs to be any way to tell if a person has hidden a weapon beneath his or her clothing or not.

In order to get a clearer picture, we took a step backwards and looked at the A-scans. Fig. 9 shows the retroreflector with and without cotton at 1.2 meters from the transducer.

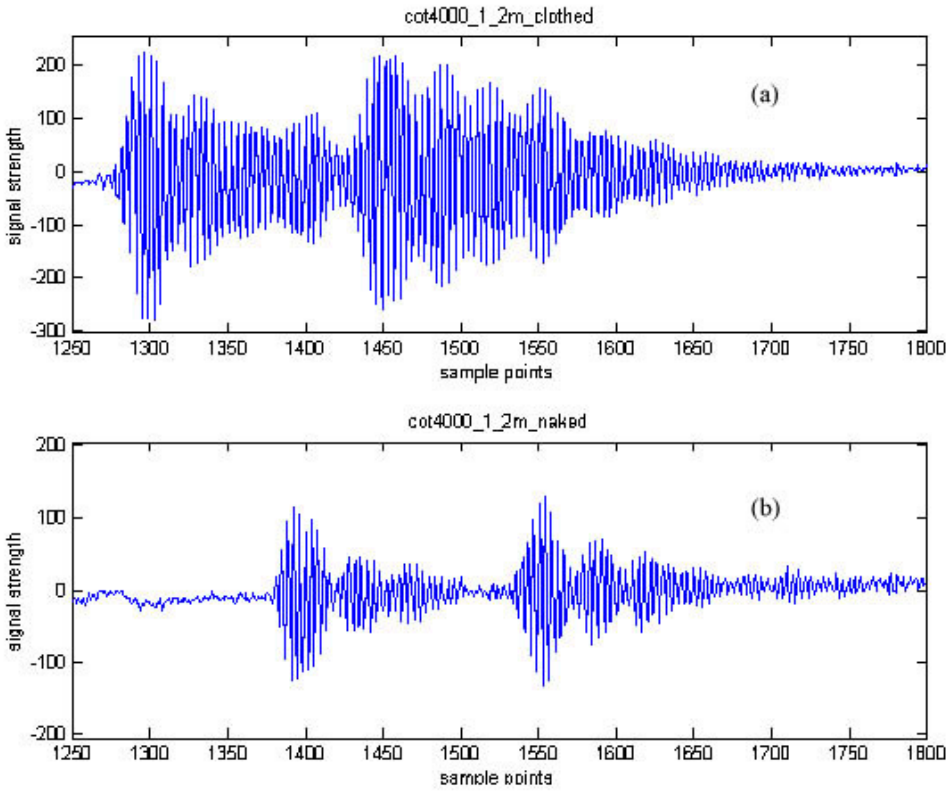

Figure 9: Two A-scans of the linear retroreflector at 1.2 meters from the transducer, taken (a) with cotton draped in front of it, and (b) without any fabric in front of it. These were taken between sample points 1250 and 1800, which cover a real-world range from 1.14 meters to 1.64 meters. Notice the double signals in each graph, an error that was corrected, and the bumps from the retroreflector at sample points 1400 and 1550.

We made these A-scans before we adjusted the pulser-receiver card to get rid of the "double image" effect. Nonetheless, it is possible to see in the upper A-scan around sample point 1400 a "bump" in the image. It also occurs around sample point 1550. Looking at the bottom A-scan, we see that these bumps correspond to the images of the retroreflector. We would expect the fabric by itself to bounce off some of the wave energy and send it back to the transducer, without any irregularities in the signal-just a spike in the A-scan which gradually decreases back to noise. But this is not the case here. What is happening is when the ultrasound wave hits the fabric, some of the wave bounces off the fabric, while some of the wave passes through the fabric and bounces off

the retroreflector beneath. After adjusting the pulser-receiver card, we see this more clearly in Fig. 10 with the polyester-cotton mix at 1.1 meters.

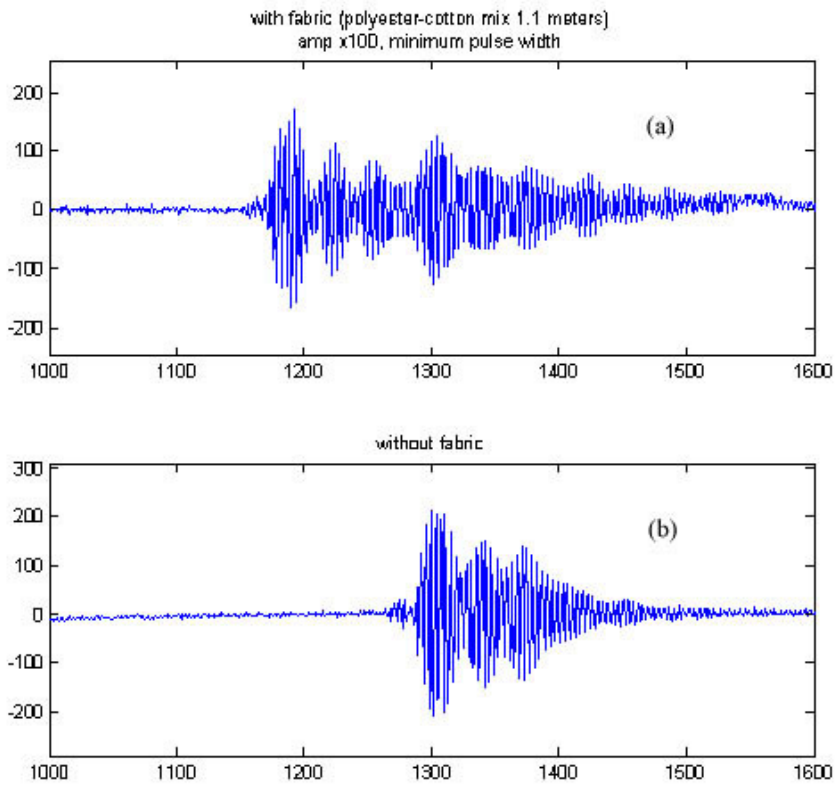

Figure 10: Two A-scans of the linear retroreflector at 1.1 meters from the transducer, (a) with, and (b) without a swatch of polyester-cotton mix fabric draped in front of it. The real-world distance ranges go from 0.91 meters to 1.45 meters. Notice in (a) how the fabric shows up as a high-amplitude signal at sample point 1190, while there are bumps corresponding to the retroreflector at sample point 1300 in both (a) and (b).

Here, the "double image" has been eliminated, as is obvious from the bottom A-scan of the retroreflector by itself. But in the upper A-scan of the polyester-cotton mix in front of the retroreflector, we see a bump in the image at sample point 1300, corresponding to the image of the retroreflector by itself. Furthermore, taking into consideration that the reflection from the fabric occurs around sample point 1190 and the reflection from the retroreflector occurs around 1300, we calculate the retroreflector to be (1/1100)*(1300-1190)=0.1 meters, or 10 centimeters from the fabric. A quick measurement of the trolley used to hold up both the retroreflector and the fabric shows

that that is a reasonable distance; when the target apparatus was set up, the retroreflector was about that distance from the fabric. Fig. 11 is a photograph we took measuring the distance from the retroreflector (at the 90 cm point) to where the fabric hangs down (denim at the 100 cm point).

Figure 11: A photograph showing a measurement of the distance between the fabric (right) and the retroreflector (middle). From the meter stick, we see that there is a distance of 100 cm - 90 cm = 10 cm, or 0.1 meters, between the fabric and retroreflector, as was predicted from the data in Fig. 10(a).

C. Tissues

We started our experiments in using ultrasound for concealed weapons detection by investigating the air-weapon interface and the air-fabric-weapon interface. The data that were collected from these experiments were incomplete. In a real-world situation, an aggressor's body would consist of muscle, skin, fat, and other tissues, as well as the weapon and clothing over it. So the next phase of our experiment was to investigate the interface between air, fabric, weapon, and human tissues.

The problem with this, though, is that we could not use actual human beings in these experiments. Using any one or two people to model for the scans would not produce adequate data because people come in all sorts of body types. Using a model that is either in perfect health or obese would only give us data of what we could expect to see in the scans of a portion of the population. Also, it would be difficult to position and keep still a model just right so that he or she would reflect ultrasound waves back to the transducer. To solve these problems, we decided to use a dressmaker's dummy we named "No-head." No-head can be positioned anywhere in front of the transducer without the hassle of shuffling a person around to the correct spot.

The problem with using No-head is that it does not have the tissues of a living human being. In order more closely simulate a human being using the dummy, we had to find some kind of substance that could simulate human tissues. We concluded that rubber would be best for the job, since it is cheap and flexible, and can be stretched around the dummy's torso. We thus had to ask ourselves what kind of rubber would be best suited for simulating human tissues. Since we were working with sound, it is only logical that a rubber that would be well suited to simulating human tissues would have an acoustic impedance equal or close to that of human tissues.

Our first step was to find the acoustic impedance of living human tissues. We found that skin has an acoustic impedance of 1.6 MRayls (megarayleighs), fat has 1.35 MRayls, and muscle 1.7 MRayls. Using this data, we went to the Onda Corporation website [3], where we found acoustic properties tables for several different kinds of substances, including rubber. From this table, we found that there were three kinds of rubber that had acoustic impedances similar to those of human tissues: ecothane,

polyurethane, and RTV (Room Temperature Vulcanizing) rubber. We did not want to use three kinds of rubber, so we investigated which rubber has the best balance of simulating the three kinds of tissues.

To find the best rubber of the three, we plotted the six materials (skin, fat, muscle, ecothane, polyurethane, and RTV) versus their reflection coefficients relative to air using MS Excel 97. Figs. 12 and 13 are plots of the reflection coefficients of the six materials-the three tissues and three rubbers-relative to air. Fig. 14 is a plot of the reflection coefficients of various non-tissues, for comparison.

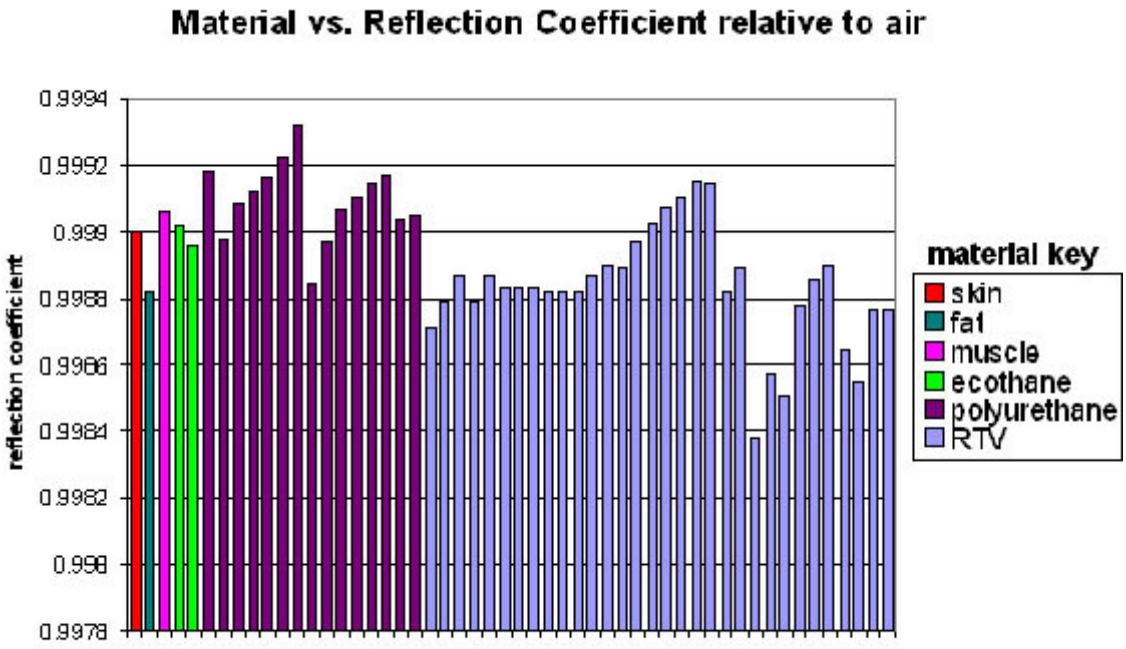

Figure 12: A plot of the reflection coefficients of six materials-skin, fat, muscle, ecothane, polyurethane, and RTV rubber-relative to air. The graph is color-coded by material.

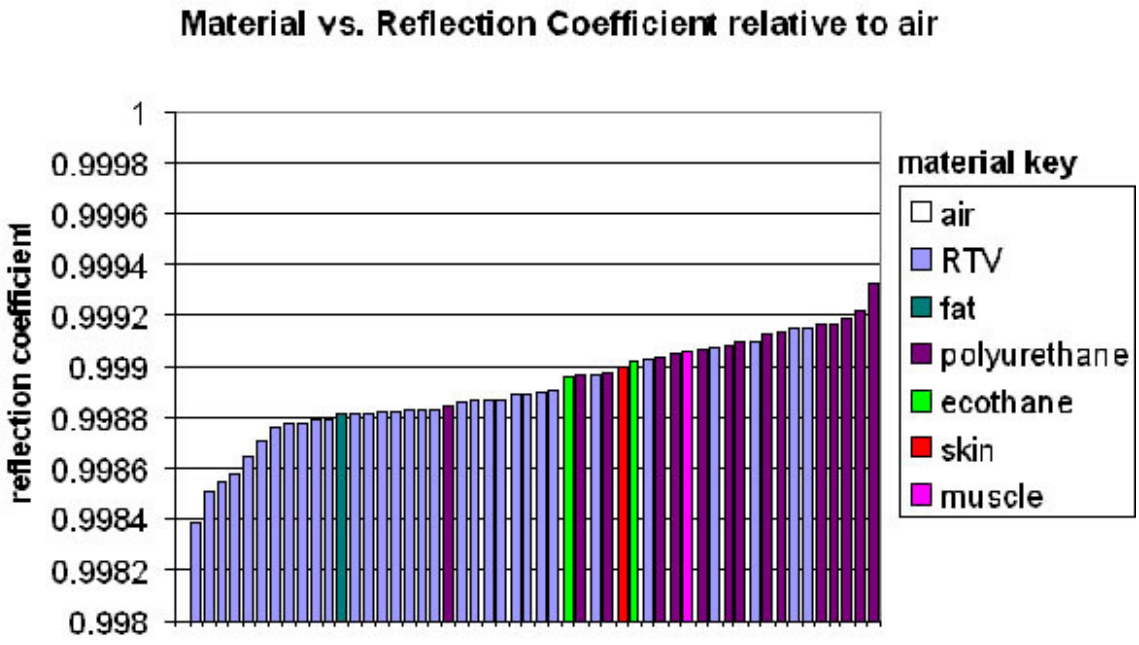

Figure 13: The same as the plot shown in Fig. 12, only with the reflection coefficients in ascending order, for ease of comparison between the three tissues and the three rubbers.

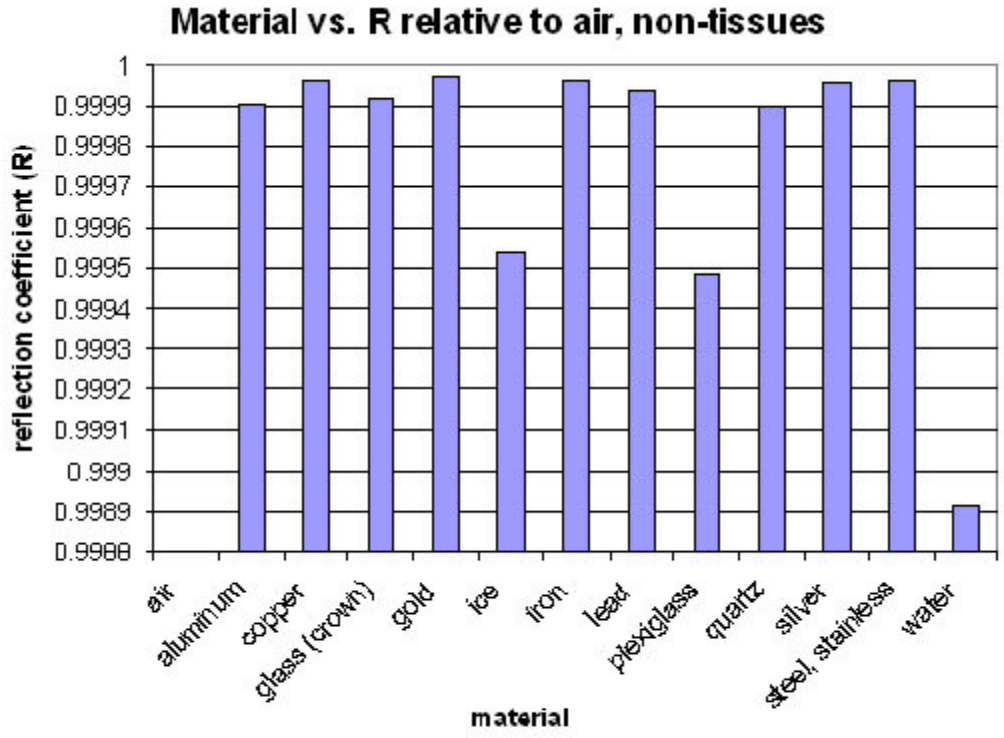

Figure 14: A plot of the reflection coefficients of various non-living materials relative to air. This was created simply for comparison with Fig. 12 and 13.

There are multiple bars for each rubber because each rubber has multiple varieties with different acoustic impedances (see Appendix B for details). What we see from these graphs is that the reflection coefficient of fat is closest to that of the RTVs. Skin is closest to ecothane, as well as some polyurethanes and RTVs. Muscle is closest to the polyurethanes, but not so much the RTVs or ecothane.

We can see with even more detail how close the rubbers come to simulating the tissues with graphs of material vs. the reflection coefficient percent difference of the rubbers from the tissues. In this case, the formula used was

$$Percent\ Difference = \left| \frac{R_{mat} - R_{air}}{R_{air}} \right| X\ 100 \qquad (3)$$

where R_{mat} and R_{air} are the reflection coefficients of the rubber and the air, respectively. Figs. 15, 16, and 17 show these graphs.

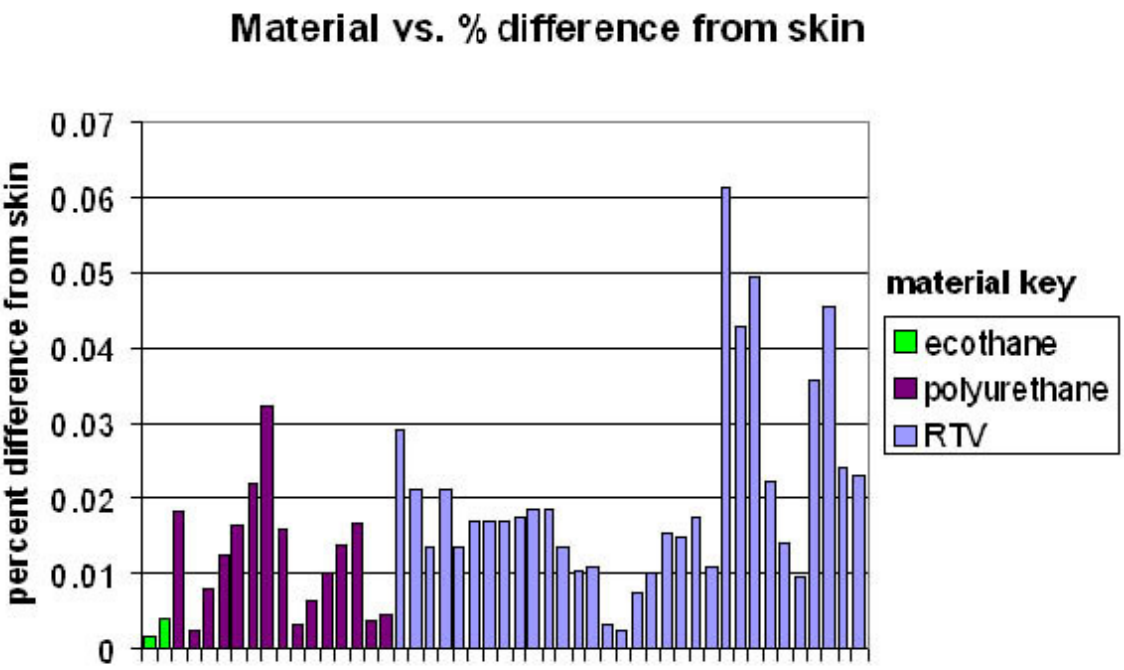

Figure 15: A plot of the percentage difference between the reflection coefficients of each of the three rubbers considered-ecothane, polyurethane, and RTV rubber-and the reflection coefficient of skin relative to air. There are multiple entries for each rubber because each rubber has multiple varieties. The closer a bar is to the horizontal axis, the more like skin the bar's respective rubber is.

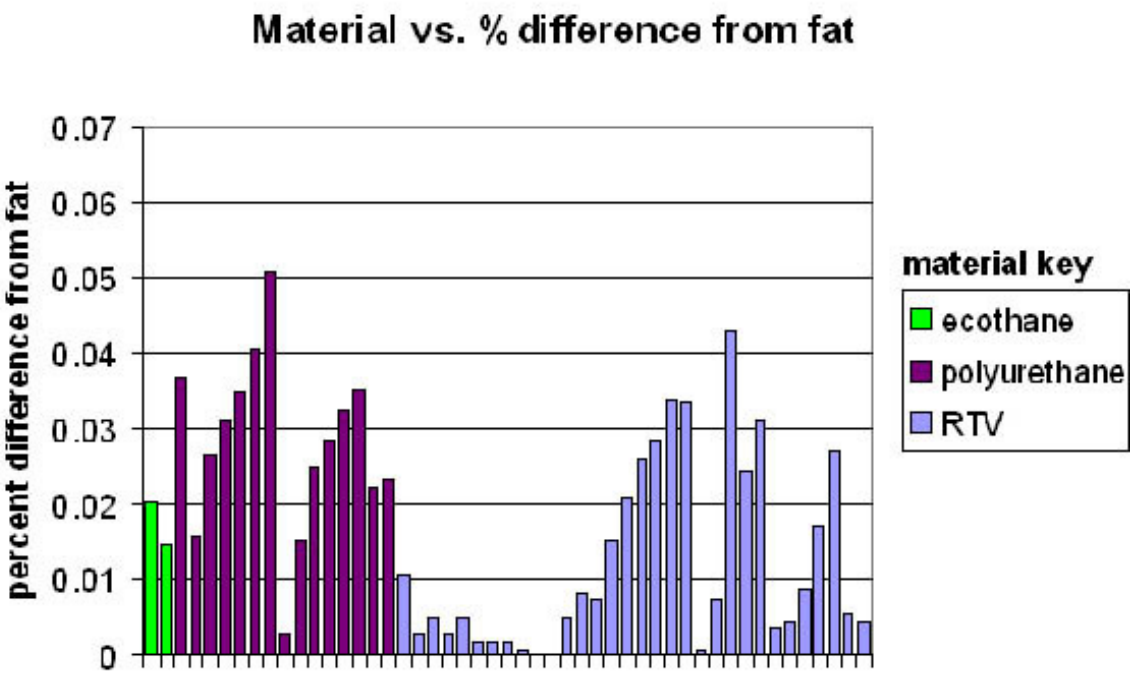

Figure 16: A plot of the percentage difference between the reflection coefficients of each of the three rubbers considered and the reflection coefficient of fat relative to air. Otherwise, just like Fig. 15.

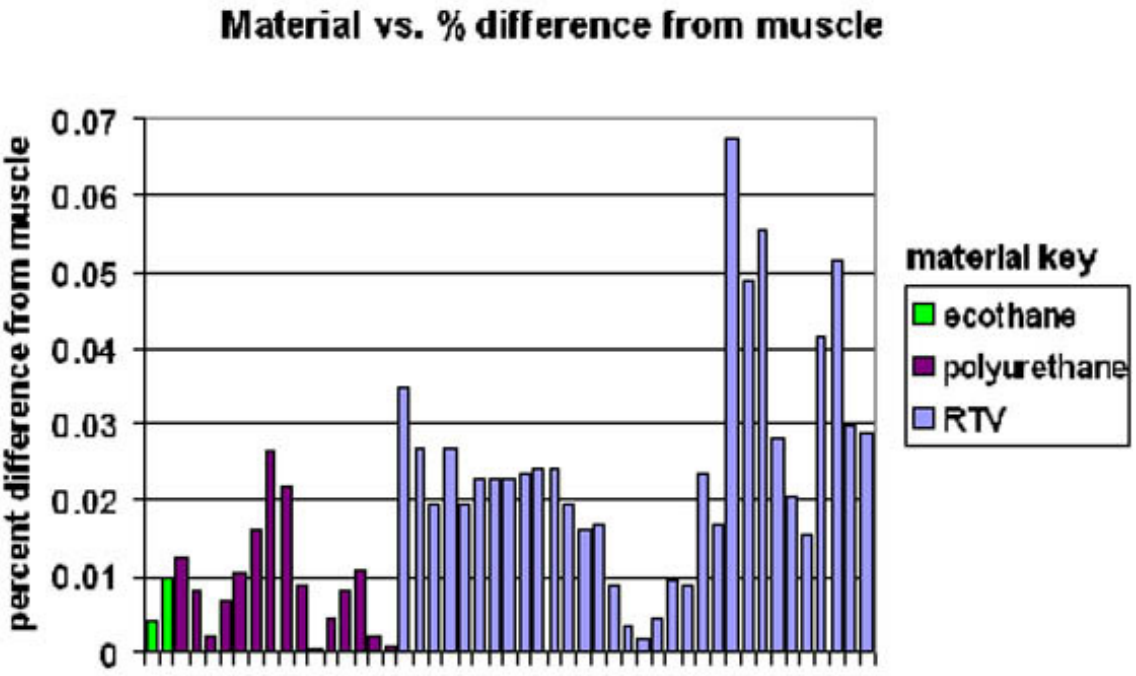

Figure 17: A plot of the percentage difference between the reflection coefficients of each of the three rubbers considered and the reflection coefficient of muscle relative to air. Otherwise, just like Fig. 15 and 16.

On these graphs, the closer the bars are to the horizontal axis, the closer is the material's acoustic impedance to that of the tissue. We see that RTVs go from simulating each tissue pretty well to not simulating it closely enough. Overall, it simulates fat the best, followed by skin. Ecothane simulates skin the best, followed by muscle. Polyurethane simulates muscle the best, followed by skin.

So, which rubber simulates human tissue the best? We like to think that a rubber can be chosen based on the body type of a target person one would want to scan. Ecothane would be best if one wanted to look at the skin without bothering too much with muscle or fat. In this case, it would look like the transducer is scanning a skinny person without much muscle or fat. The RTV, with fat being the closest to its reflection coefficient, could be used to simulate an overweight person. Polyurethane could be used to simulate a strong, muscular person without too much body fat.

If we had to choose one rubber out of the three, however, we would choose RTV rubber. It has a wide enough range of reflection coefficients and acoustic impedances to have a better chance at simulating all three tissues. Plus, it is very readily available online, and is sold as a modeling agent. Ecothanes, on the other hand, are not as readily available as RTVs-a Google search of the word "ecothane" only turns up 11 matches ("ecothane rubber" turned up 0 matches), while a search of the words "RTV rubber" turns up 1,570 matches ("RTV" turns up 1,170,000 matches; result as of July 16, 2004). Polyurethanes can have serious health risks associated with them, as they release hazardous chemicals called isocyanites into the air. A search of the OSHA (Occupational Safety and Health Administration, a subsection of the Department of Labor) website revealed no health risks associated with RTV rubber [4].

We did not use any RTV rubber in our experiments. Instead, we used another tissue-simulating substance called ballistic gel. Ballistic gel is used to test bullet performance. It is a very stretchy, rubber-like substance with a smooth, somewhat sticky texture, a yellow color, and a pleasant vanilla odor. It melts at 140°F and is 100% soluble in water. After melting it and letting it re-solidify five times, we found that its elasticity remains fairly consistent. We left some pieces in a tub of cool water overnight. The next morning, they had turned mostly white and had lost their vanilla odor. They also were easily pinched apart, unlike the ballistic gel which requires much more force to tear. A few days after melting and re-solidifying the gel that had been soaked, the gel became harder and almost inelastic. Overall, ballistic gel is very stable, and the company that makes it [5] claims that it simulates muscle tissue very well (see Appendix C). We also used a sheet of latex gum for preliminary scans. Pictured in Fig. 18 and 19 is a photograph and B-scans of No-head wrapped in the sheet of latex gum and bound with the denim cloth, respectively. There are 15 scans in each case and the transducer makes a more complete sweep of the target, which are changes to the scanning apparatus that we will discuss in the next section.

Figure 18: A photograph of No-head wrapped in a sheet of latex gum rubber and bound with denim cloth. We used this model in our first scans of No-head with rubber, but we later used ballistic gel on No-head.

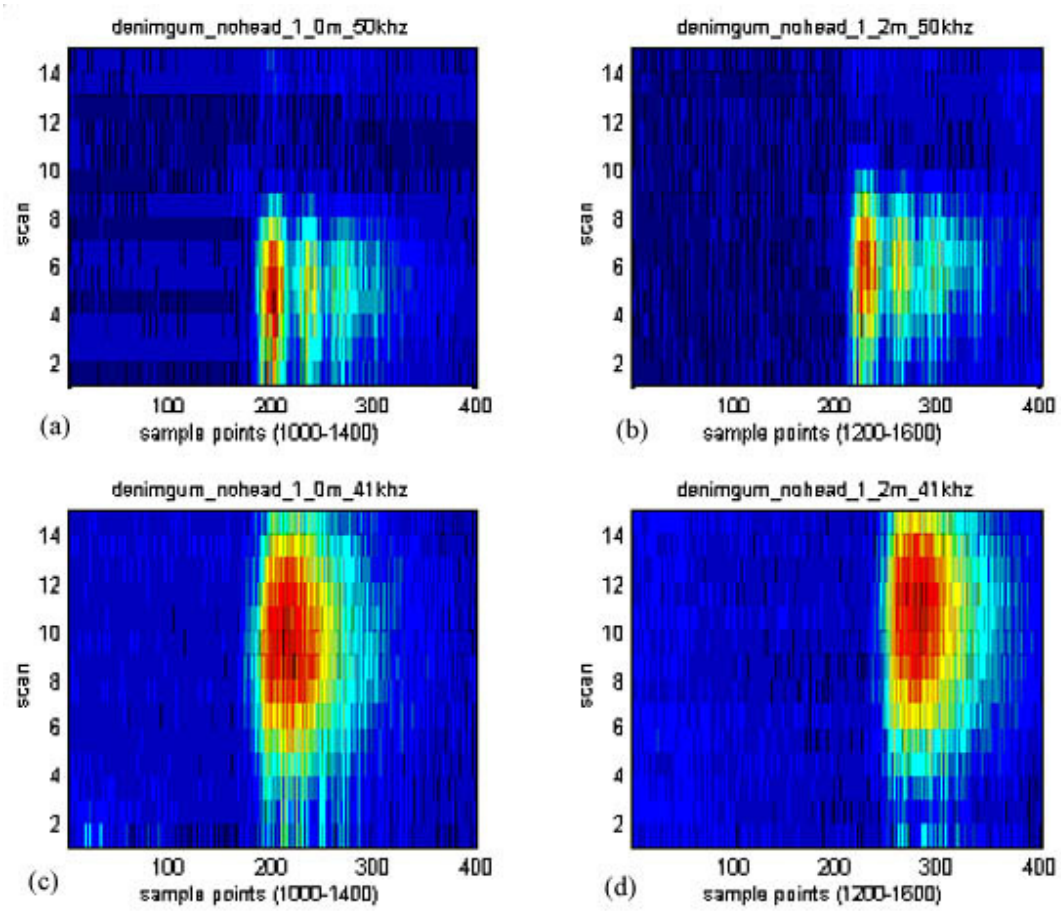

Figure 19: B-scans of the latex gum and denim-covered No-head, taken with (a) the target 1 meter from the 50 kHz transducer, (b) 1.2 meters from the 50 kHz transducer, (c) 1 meter from the 41 kHz transducer, and (d) 1.2 meters from the 41 kHz transducer. A B-scan is several A-scans placed together to form a three-dimensional image-that is what it means by "scan" on the vertical axes of these graphs. We are looking down the z-axis, or signal strength axis, of these graphs-the strongest signals are in red, and the weakest in blue.

The scans in Fig. 19 are B-scans, which are several A-scans placed adjacent to one another to form a three-dimensional image (imagine placing sheets of paper that show A-scans into a folder in the order in which they were made, and looking into the folder at just the lines of the A-scans). The B-scans above are arranged looking down the z-axis (the signal strength axis) at sample points on the x-axis (horizontal) and the scan number on the y-axis (vertical). In these B-scans of No-head, the blue areas represent the lowest signal strengths, while the red areas represent the highest signal strengths. Notice

how the highest peaks shift from one scan to another at different distances. This happened because the orientation of No-head was changed slightly when it was shifted around to different distances. Still, these scans show that it is possible to "see" the dummy with the fabric and the rubber, even though they are too close together to differentiate between the two in the scans.

In summary, we searched for the best rubber to simulate human tissues for the next round of experiments we have planned. We wanted a rubber, since it is flexible and cheap. Out of three rubbers-ecothane, polyurethane, and RTV rubber-we found that RTV rubber is the best at simulating human tissues, as it has the widest range of acoustic impedances, and thus the best chance at simulating skin, fat, and muscle. For our next series of experiments, however, we used ballistic gel.

III. Proof-of-Concept with Weapons

In this phase of the experiment, using No-head wrapped in ballistic gel and four "weapons" found around our laboratory as targets, we set out to prove that ultrasound is feasible for detecting concealed weapons.

We created the target by first melting down ballistic gel in an incubator. The incubator was very slow in melting the gel, but we managed to melt down three large chunks. We dumped the liquid gel into a large rectangular metal tub, where it was allowed to cool down and solidify. The gel solidified into three rectangular pieces, about 16.5 by 10 inches each, which we wrapped around the front of No-head. One piece went around the upper torso, or chest, one went around the middle torso, or midriff, and the third went around the lower torso, or below the waistline. A smaller fourth piece was

constructed to go around the neck. No pieces of gel were wrapped around the backside, as it was not necessary. Three photographs of No-head wrapped in the ballistic gel are shown in Fig. 20.

Figure 20: Three photographs of our dressmaker's dummy No-head wrapped in ballistic gel. It was constructed from three large rectangular pieces of ballistic gel melted down in an incubator and one small rectangular piece of ballistic gel for the neck.

Small flaws developed in the ballistic gel as they were melted down and shaped. First of all, small bubbles formed in the liquid gel. Second, the gel was removed from the incubator before it was fully melted, and a thin skin of gel was still left on top for the middle and lower pieces. When the liquid gel was poured into the tub to make those two

pieces, the skins went in the tub with the liquid and subsequently were encased within the gel pieces, creating uneven patches in the pieces. We do not think that this affected our scans, though, as the flaws were made of the same material as the surrounding body pieces.

After No-head was wrapped in the gel, we looked for weapons to place on No-head. We found four "weapons" around the laboratory that we believed would be effective in simulating actual weapons (in their defense, some of them really could be used as actual weapons). They were a 6 by 3 inch metal plate (or a "metal shard"), a ball-peen hammer, a laser pointer (or "pen laser"), and a boxcutter ("exactoknife"). The four weapons are shown in Fig. 21.

Figure 21: The four weapons used in our final proof-of-concept studies. The weapons pictured are, from left to right, a metal shard, a ball-peen hammer, a pen laser, and a boxcutter ("exactoknife").

The orientation and data collection parameters of the transducer apparatus were changed so it would collect more data and take a more complete scan of the target. First,

the transducer was moved so it would start scanning to the right of the target and move left, rather than starting directly on the target then moving left. We had it start scanning the target at around 7° to the right of the target, then move left to 0° and left of the target by about 8°. Second, we changed the LabVIEW program so it would take 15 scans in a sweep across the target, move back to its starting position, then take another 15 scans, rather than the 10 scans per sweep taken in previous experiments. With these changes and 4000 sample points per scan, we were able to collect more complete data from the target with the transducer. This explains the B-scans from the latex gum and denim-wrapped No-head.

Each of the four weapons was duct-taped to No-head and scanned. No-head without any weapons was also scanned. This was done at three different distances from the transducer: 0.8 meters, 1.0 meters, and 1.2 meters. No fabrics were used in this part of the experiment, although we intend to use them in future scans. Scans like that of the B-scan of the latex gum-denim covered No-head proved to be the most helpful in analyzing whether the weapon was "visible" to the transducer or not (keep in mind that the weapons were duct-taped to No-head and thus right against it, and that the signal from the weapon could blend in with the signal from the gel on an A- or B-scan). Four of these B-scans are shown in Fig. 22. They are scans of No-head covered in ballistic gel without any weapons and with the four weapons.

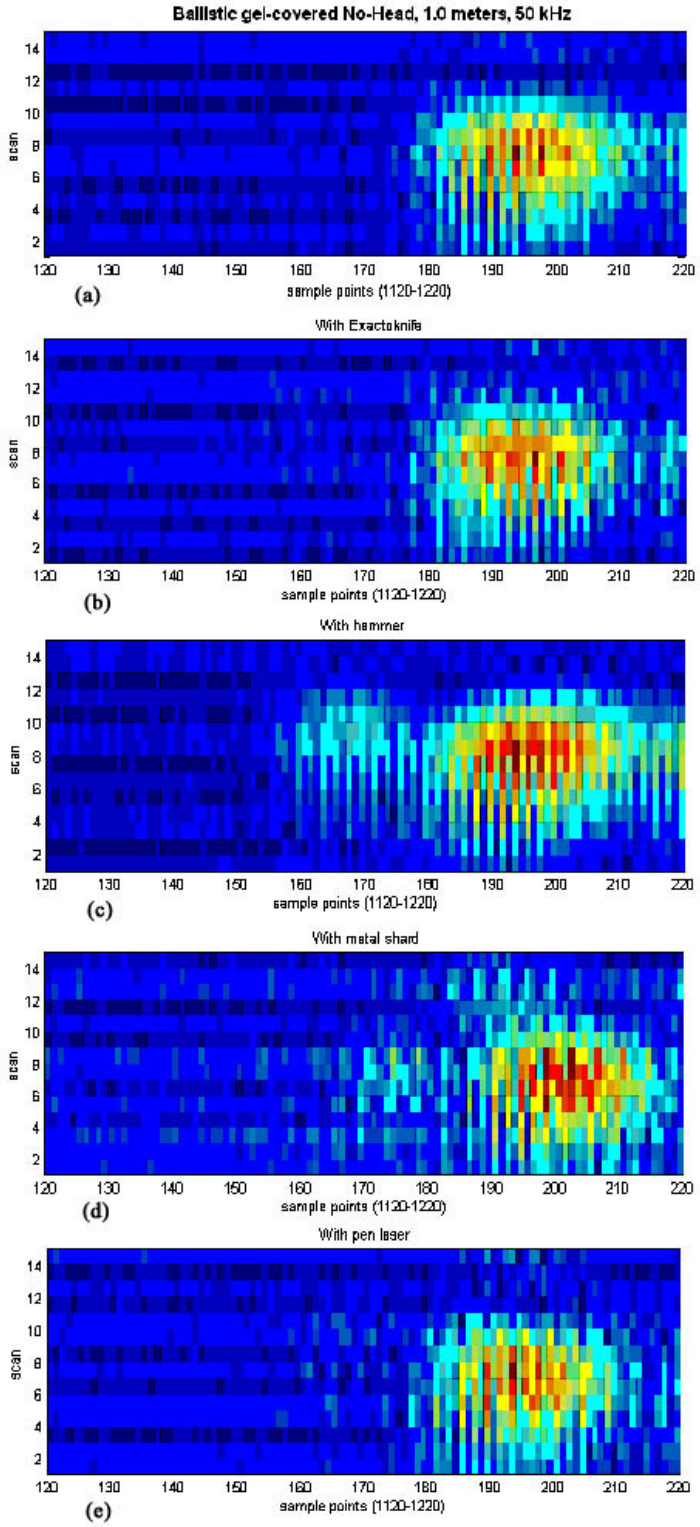

Figure 22: B-scans of No-head covered with ballistic gel and with the four weapons from Fig. 21 duct-taped to the front of it. The weapons in these graphs are (a) no weapon at all, (b) the boxcutter, (c) the hammer, (d) the metal shard, and (e) the pen laser. The large signals between sample points 180 and 220 are No-head. Every bump before that can be taken as a weapon. Notice how we get large return signals from (c) and (d), the densest weapons of the four.

The first thing we see in the B-scans in Fig. 22 is what No-head looks like without any weapons. The signal from No-head here is ellipsoid-shaped with the strongest signal in the middle of the ellipse. We then move onto the four weapons. We do not see much of a signal returned from the boxcutter or the pen laser. This makes sense, as both of those objects are small and difficult to see with the relatively poor resolution of the 50 kHz transducer. We move onto the metal shard, which can be seen in the small but visible signal (the light blue spots on the second B-scan from the bottom) to the left of No-head's strong signal. The hammer has the strongest signal of the four weapons. A signal is returned from the hammer that is not quite as strong as the signal from No-head, but is definitely visible as a solid light-blue ellipse on the B-scan. This also makes sense, as the high density of the hammer gives it a very strong interface with the air (*i.e.* the acoustic impedance of the hammer is very different from that of air, giving a high reflection coefficient).

We conclude from these scans that at least at short distances around 1 meter, a 50 kHz transducer can differentiate an unarmed person from a person armed with a blunt instrument or some other high-density object that can be used as a weapon. We reached this conclusion using four "weapons" found around our lab and a dressmaker's dummy covered in ballistic gel.

IV. Using Wavelet Transforms for Data Interpretation

Although we were able to see large, dense objects by the end of our proof-of-concept studies using ultrasound, like hammers and flat pieces of metal, we were unable to see smaller objects, like pen lasers and boxcutters. The reason we were able to see

dense objects was because they have a higher reflection coefficient relative to air. However, just because an object is small and less dense than, say, a hammer, does not mean it cannot be just as dangerous. We would like it if ultrasound could thus be able to pick out any weapon on a person's body regardless of the weapon's density. Fig. 22, though, does not offer us much hope in that respect.

We were not sure the smaller weapons went completely unseen, however. In our experiments, we used wavelet transforms on the signal returned to the transducer from the target to give us another perspective on the data contained within the signal. Instead of the time vs. amplitude graph of the signal that we usually see, the Continuous Wavelet Transform, expressed in Eq. (4), gives us a graph of time vs. frequency. The Continuous Wavelet Transform is expressed as [6]

$$CWT_x^{\psi}(\tau, s) = \Psi_x^{\psi}(\tau, s) = \frac{1}{\sqrt{|s|}} \int x(t)\psi^* \left(\frac{t-\tau}{s}\right) dt \qquad (4)$$

where x(t) is the original signal being transformed and ψ(t) is the transforming function (some predetermined wave that is fit to the data). The transformed signal Ψ is a function of translation (time, τ, also expressible as distance) and scale (1/frequency, s). Basically, Eq. 4 tells us where we can see certain frequencies in the space in front of the transducer. Using a contouring operation, we can make the graph appear as a series of fingerprints.

A. Experimental Setup

In our experiments, we once again used a dressmaker's dummy called "No-head" and sheets of ballistic gel to simulate the human body. We duct-taped a gun-shaped piece of metal to No-head's midriff and scanned it with and without one of six fabrics draped in

front of No-head and pinned in back of No-head. No-head was placed one meter from the transducer. The transducer swept across the front of No-head, collecting data in two one-dimensional sweeps, with 15 scans taken per sweep.

We had to account for how the orientation of the body with respect to the scanner and the orientation of the weapon on the body could affect the scan, so we included these factors in our scans as well. To factor in body orientation, we used clock as a kind of coordinate system. When the hour hand points directly at 12 ("straight forward"), it is 12:00; likewise, when No-head's stomach is oriented parallel to the face of the transducer ("dead on," or "straight forward"), we say that its orientation is at 12_00. An orientation of 9_00 or 3_00 would have No-head oriented at a 90° angle from the transducer (and 180° from each other), and an orientation of 10_30 or 1_30 would have No-head oriented at a 45° angle from the transducer (90° from each other). For body orientations, we used 10_30, 12_00, and 1_30 for the first few scans, then we later added the orientations 11_00 and 1_00 (30° from transducer each) for more detail. These orientations were chosen because they were less than 90°, at which we would not expect to receive any signal back from the target.

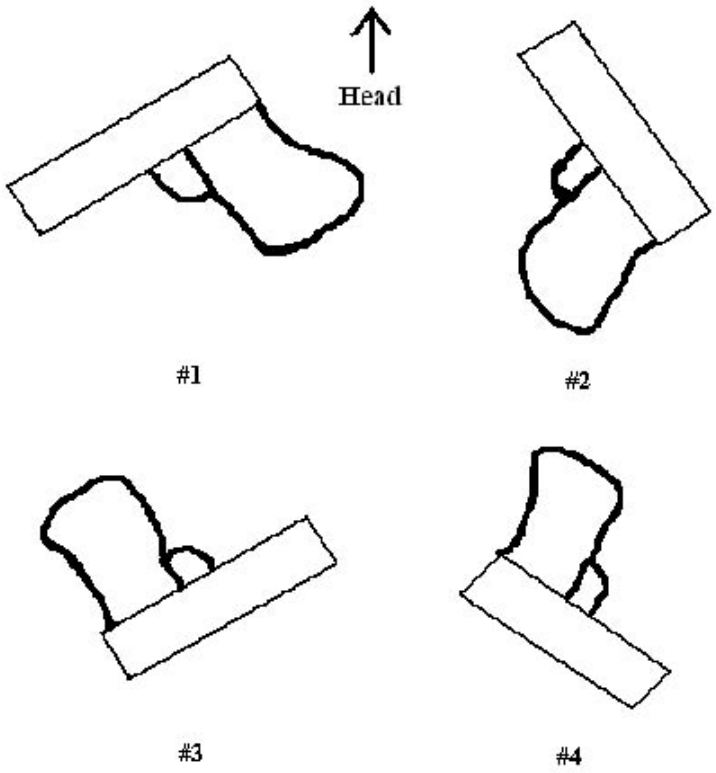

Figure 23: The four different gun orientations.

Fig. 23 shows the four different orientations of the gun on the body. If we were to once again imagine a clock, only this time on No-head's stomach with 12_00 pointing straight up, then gun orientation #1 would have the barrel of the gun pointed at 8:00, gun orientation #2 at around 11:00, #3 at 2:00, and #4 at around 4:30. These orientations were chosen randomly. The complete orientation of the body-gun system was labeled as "body orientation_gun orientation." As an example, 12_00_02 would be body orientation 12_00, gun orientation #2; 1_30_01 would be body orientation 1_30, gun orientation #1; 11_00_04 would be body orientation 11_00, gun orientation #4.

In addition to introducing the controlled inaccuracies of body and gun orientation, we also allowed for some small, uncontrolled inaccuracies. We wanted to see how much small changes in body orientation would affect our scans, so were not overly precise with

regards to body orientation. For example, the orientation 11_00, in some cases, might not have been exactly 30° to the left of the transducer, and 12_00 might not have been exactly dead-on the transducer. The gun orientations were not exactly precise, either-we sketched the initial orientations in a notebook. When we wanted to return the gun to that orientation, we looked in the notebook and moved the gun to most closely match the gun in the picture (which, of course, does not give us an entirely accurate representation of what the gun's orientation was before). We hoped that these small, uncontrolled errors would allow us to see how small changes in the orientation of a body and a weapon would affect how they are "seen" by the transducer. After all, in real life, people move a little even when standing still-they shift their weight on their feet, they stretch, and they perform other small motions. If a person is concealing a weapon, we want to make sure that our detection techniques are not prone to make errors simply because we could not get the person to stand completely still.

We also melted down a sheet of ballistic gel and used a roughly female-shaped plastic swimsuit dummy to mold it into a human torso shape. We backed up the gel with a piece of wood and drilled two holes in the shoulders so we could hang it from a movable trolley. We called this new dummy "No-back" and "torso plate." We scanned the metal gun on No-back at all four gun orientations as well, but only at the 12:00 body orientation.

B. Computer Programs and settings used

After we took the scans, we analyzed them using the Interactive Wavelet Thumbprint Tool, a computer program developed by Kevin Rudd. This program takes

the signal x(t) and transforms it using the continuous wavelet transform from Eq. (4). The Wavelet tool then arranges the transformed signal into a series of two-dimensional "thumbprints."

Fig. 24 shows us one example of a thumbprint diagram created using the Interactive Wavelet Thumbprint Tool. We would like to note here that the program has a 10:1 ratio of sample points plotted on the thumbprint graphs to sample points actually collected (sample points being a measure of distance or time). This exists because giving the program a more complete data set allows it to paint a more thorough picture of the thumbprints.

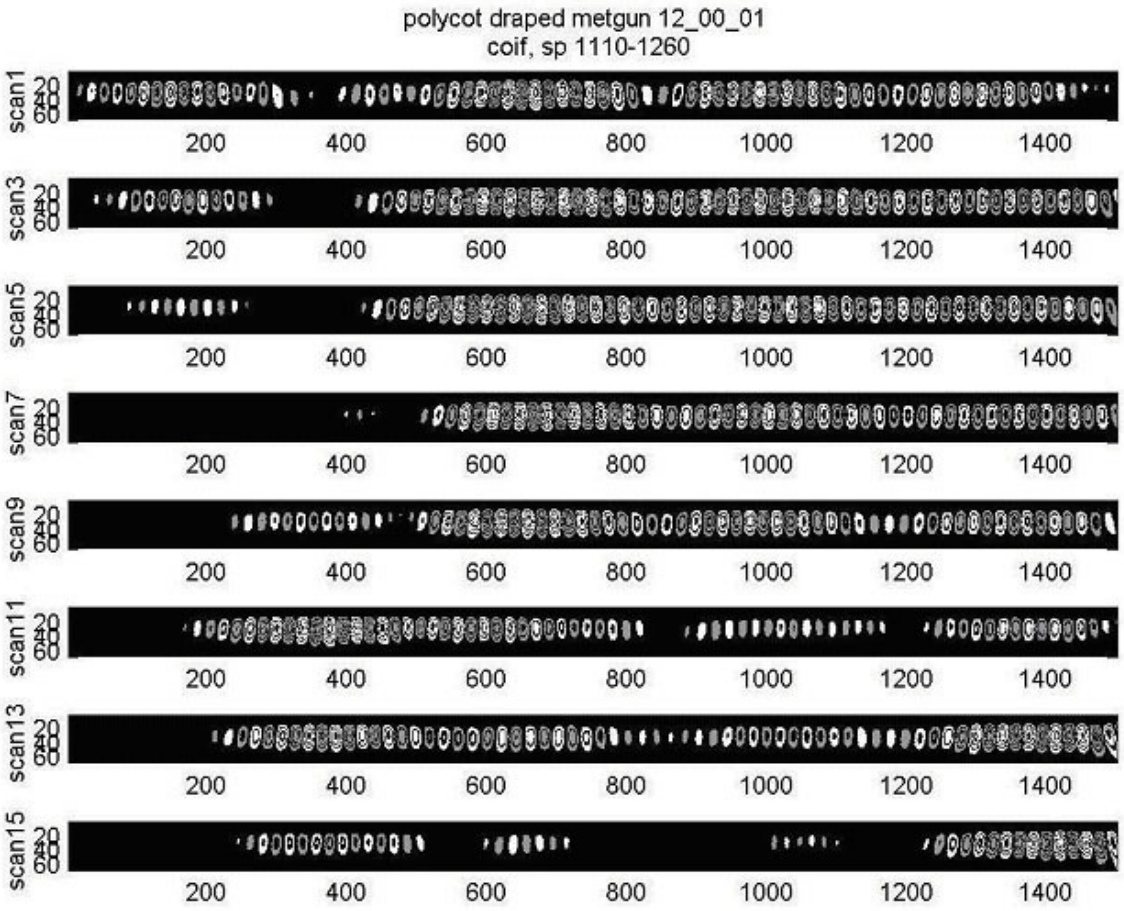

Figure 24: This is one of the thumbprint diagrams created using the Interactive Wavelet Thumbprint Tool. The above thumbprints were created from the first 15 odd-numbered scans of the polyester-cotton mix cloth-covered No-head armed with the metal gun. The body orientation here is at 12:00, and the gun orientation at #1 (see Fig. 23). The horizontal axes of these graphs are a measure of distance/time (with the leftmost edge being at about one meter from the transducer), and the vertical axes are an (indirect) measure of frequency. The word "coif" signifies that the thumbprint tool fit the signal with the Coiflets family of wavelets to collect the time-frequency data.

From initial observations, packets of high frequency appear and disappear frequently as the transducer moves across the target. Whether these packets are the weapons, the fabric, the body, or something else requires further analysis.

Data and Analysis of the Wavelet Transform graphs

After taking scans of all possible metal gun-equipped targets, we whittled down the data we looked at to a set of 108 graphs. Among the fabrics we looked at, we decided to look at the polyester-cotton mix, the denim, and the cotton. We only used the 11:00, 12:00, and 1:00 body orientations, and all four gun orientations. All three No-head dressing styles-naked, draped, and pinned-were used. We only used the Coiflets family of wavelets ("coif" on the graphs) to fit the original signal. The sample points are labeled with "sp" and also displayed on the graphs. The data on the graphs were usually taken between sample points 1050 and 1300 (with 1100 sample points per meter, that is between 0.95 meters and 1.18 meters from the transducer). We usually looked in a range between 120 and 150 sample points (for the graphs, multiply the sample points by 10 to get the measure that the wavelet transform program used). Among the various scans taken in each sweep of the target, we only looked at the first eight odd-numbered scans-that is, scans 1, 3, 5, 7, 9, 11, 13, and 15 (labeled scan1, scan3, etc.). Since we believe it is to the reader's advantage to know the researcher's perspective of the data, Fig. 25 displays how the 108 graphs were laid out.

	Body Orient.	11:00				12:00				1:00			
Type of Fabric	Fabric \ Gun	1	2	3	4	1	2	3	4	1	2	3	4
Denim	naked												
	draped												
	pinned												
Polyester-Cotton Mix	naked												
	draped					X							
	pinned												
Cotton	naked												
	draped												
	pinned												

Figure 25: The layout of the thumbprint graphs created from the Interactive Wavelet Thumbprint Tool. The position of the graph from Fig. 24 is indicated by a big black X.

One of the most immediately apparent things about these graphs is how the data tend to bunch together in ellipsoid shapes. We called these patterns "frequency bubbles." A simple example of what a frequency bubble looks like is shown in Fig. 26.

Figure 26: An MS Paint-drawn example of a frequency bubble. Many of these were observed in the thumbprint diagrams like the one shown in Fig. 24.

The ellipse around the thumbprints in Fig. 26 shows that the "bubble" has this general shape. The bubble may vary slightly in shape, like the ends may be pulled upward. It may also vary in size, which we took to signify differences in the sizes of parts of the target (the weapon is smaller than the body, for example, and thus we expect there to be a difference in how the A-scans and their corresponding wavelet transforms show them). An example of some real-life frequency bubbles is shown in Fig. 27.

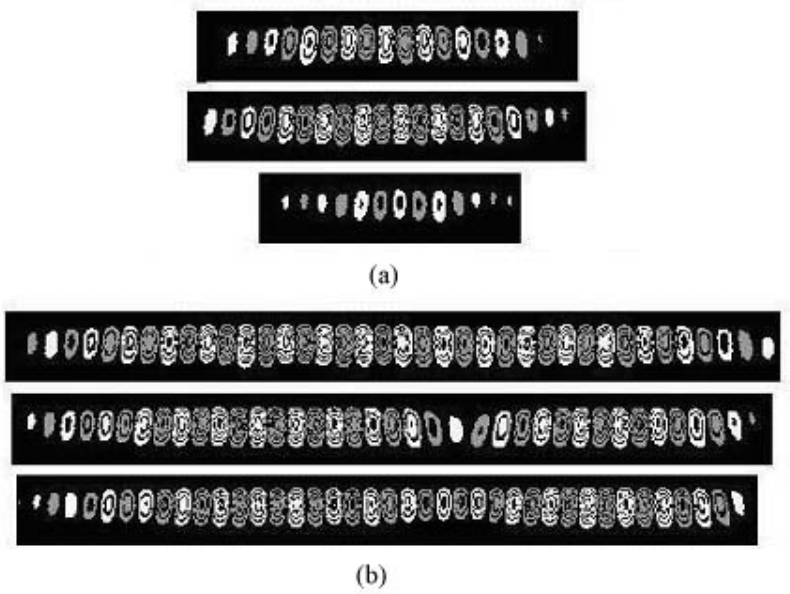

(a)

(b)

Figure 27: These are some examples of actual frequency bubbles taken from an actual thumbprint diagram. In (a), we see bubbles that occur all by themselves. We usually took these bubbles to be the front of something, like No-head, the weapon, or the fabric. Also notice how they tend to vary in size and shape. In (b), we see a common occurrence-frequency bubbles that group together in long strings. In these cases, the ellipsoids that constitute the frequency bubbles may overlap. Ususally we saw these patterns as the front of No-head, since No-head was the largest object in our target and thus the most likely to show up in our scans, followed by several smaller reflections.

Since we took the long strings of thumbprints to be the front of No-head and its reflections, we concluded that anything to the left of these strings was either one of three things:

1. Noise (which tends to show up as random strings of non-geometrically organized thumbprints; an example is shown in Fig. 28),

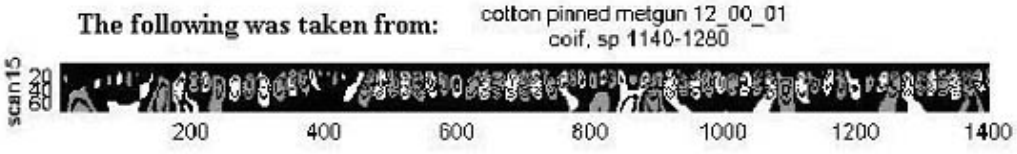

Figure 28: An example of noise in the thumbprint diagrams.

2. The return signal from the fabric, or

3. The return signal from the weapon.

We called this area on the graph the “area of interest” because it contained the data we were interested in.

Defining the area of interest was not easy. Since we placed No-head one meter from the transducer, we would expect to see the front of No-head at sample point 1100 (because there is a conversion of computerized units of measurement to real-world units of 1100 sample points per meter, as we mentioned previously). But, we see many of those long strings of thumbprints start around sample points 1140 to 1160, putting No-head around 1.04 to 1.05 meters from the transducer. It is possible that our measurements of how far No-head was from the transducer were somewhat inaccurate; there could have been an error of 4 or 5 centimeters that we did not notice. Even so, the consistency of the position of these strings of thumbprints suggests that something that changes very little is at that point in space, and that something can only be No-head. Therefore, the area of interest in our graphs tends to occur before sample point 1140.

Also, when viewing the graphs, one must take into consideration how the transducer was oriented relative to the target. The transducer swept across the target during the scans, moving over the surface of the target from left to right (from the perspective of the target), and then back again to make a second sweep. This means that the very first and very last scans were not taken with the transducer pointed directly at the target, and thus decreases the probability that the scans will show any return signal from the weapon. The scans that have the best probability of successfully receiving a signal from the weapon are scans 7 and 9 in our graphs, since those were taken when the

transducer was the closest to pointing directly at the target. So when we looked at the wavelet transform graphs, we focused the most on scan7 and scan9.

The following is a list of potential "hits" from our graphs, meaning these are places we spotted a possible signal from the weapon, fabric, or both on our graphs (the corresponding graphs, which include notes, are displayed in Appendix D):

-The signal from denim pinned metgun 11_00_02, scan9 is visible at sp 1130, but in its corresponding draped counterpart, the signal breaks up, with a small signal in front of a larger frequency bubble. That small signal takes the form of a frequency bubble itself, which suggests that it is a solid object instead of mere noise. So, we interpret the small signal at sp 1135 in the draped scan9 to be the weapon (with a trace of the fabric in front of it), and the large signal at sp 1155 to be No-head.

-In denim draped metgun 12_00_01, scan7 and scan9, we see signals between sp 1100-1140 that are small enough to be either the weapon or fabric or both. Later scans in the same graph show those signals to be bigger, suggesting that they are from No-head, but scan7 and scan9 already have larger signals around sp 1145 and 1140 (respectively) that account for the whereabouts of No-head in the scans.

-In denim pinned metgun 12_00_03, scan9, we see a signal before sp 1150 that "adds on" to the already existing signal from No-head, pulling No-head's frequency bubble out to the left. It is possible that that extra signal adding on to No-head's signal is either the weapon or the fabric or both.

-A small signal is seen in polycot pinned metgun 12_00_02, scan9, between the leftmost edge of the scan at sp 1130 and the front of No-head at sp 1150.

-In polycot pinned metgun 12_00_03, scan9, a frequency bubble suddenly appears in front of No-head's frequency bubble around sp 1160. This new frequency bubble did not appear in any of the scans in this graph, and hence can be interpreted as the frequency bubble of the weapon and/or fabric.

-In polycot naked metgun 12_00_04, scan7 and scan9, No-head's signal occurs around sp 1150, and there are very slight signals that occur just before it. Since they seem to be too small to be actual noise, we believe that they are a faint signal from the weapon (since there is no fabric in front of the weapon in these scans).

-If we may make assumptions from previous scans, signals occurring before sp 1160 in scan7 of polycot draped metgun 1_00_01 and _02 are the fabric and the weapon. At sp 1160 is supposedly No-head. But, the supposed signals from the fabric and the weapon seem to be too big to be signals from such objects (the polyester-cotton mix is a thin fabric, and the weapon has never before returned such a large signal). The signals, however, do get smaller in scan9 of those two graphs, while the "No-head" signal remains fairly constant. This assumption is further supported by the corresponding naked graphs, which show No-head to be around sp 1150.

From these graphs, we learned that the signals from the fabric and the weapon tended to be smaller than the signal from No-head. When we did see signals that suggested the presence of a weapon, the signal only occasionally took the shape of an ellipsoid prerequisite for a clear “frequency bubble.”

Results

From the analysis of our graphs and scans, we reach some tentative conclusions about how weapons can be detected using ultrasound. First, the scans of the target must be taken, with the individual or individuals taking the scans preferably standing directly facing the front of the target. The scanners may stand a little to the side of the target; however, if we notice how many signals we saw with the 12_00 body orientation, we see that facing the target dead-on would be best. There are ways to overcome this that we have not tested, such as using multiple transducers to send and receive signals, but that is being addressed by others. Second, the data received from the return signal must be processed by the wavelet transform function and turned from distance-amplitude data to distance-frequency data. Third, the distance from the scanner to the target needs to be assessed. This may be done with some approximation and guesswork by the individual taking the scan, or more precisely by the machine itself. In any case, finding such distance data would be necessary in order to assess where in the scan the target's body is located. The machine can fit a frequency bubble to the target bubble if it is possible, or it may request more information if there is too much noise. The other purpose of finding the distance from the scanner to the target is to narrow the data collected down to an area of interest that is only in the tens of centimeters instead of the hundreds of centimeters we see in the raw data. By doing so, we can make the area of interest more visible. Typically, the target's body will get the largest frequency bubble in our newly narrowed window of data. By looking to the left of that large frequency bubble (i.e. at points in

space closer to the scanner), we can determine if there are any other frequency bubbles that imply the presence of other objects on the target's body.

Then again, the data we are interested in may not come in bubble form. This will have to be determined in later experiments, where hopefully others will explore exactly how the fabric and weapon data accumulate in wavelet transform graphs. They will also need to ensure that the blips that we viewed in the scans are, indeed, from the weapon. There are still many questions to be asked regarding the wavelet transform graphs, but from our observations thus far, we are confident that viewing the data in this manner presents us with a good window with which to analyze whether a target is carrying a weapon or not.

Conclusions

In this project, we investigated the possibility of using ultrasound waves to detect weapons concealed on a person's body. We wanted to know if an ultrasound device could be used as an alternative to other nondestructive weapons detection methods, such as x-rays and metal detectors. Such a device could give much-needed assistance to security personnel and authorities, but only if ultrasound can detect weapons. If ultrasound can detect weapons, then we have to ask ourselves where to look and what to look for. So far, we have successfully shown that ultrasound waves can detect weapons beneath clothing and on a person's body. We understand how we need to organize our setup and the differences between an armed and unarmed person.

Further experimentation will require that researchers analyze the distance-frequency thumbprint diagrams and compare them with their corresponding distance-amplitude A-scans. Also, to facilitate further work, future researchers will have to look at how fingerprints are categorized and identified. These studies could further develop the algorithms of concealed weapons detection technology by shedding light on the thumbprint diagrams. Thus, we hope that our collected research will turn ultrasound into a reliable method of detecting concealed weapons.

Acknowledgements

I would like to thank Dr. Mark Hinders, Kevin Rudd, and Wen Gao for their help and support.

And last, but not least, I would like to thank my parents, Richard and Marie DuChateau, for their unwavering moral support.

Bibliography

[1] Griffin, Donald R. Listening in the Dark: The Acoustic Orientation of Bats and Men. Ithaca: Cornell University Press, 1986.

[2] Meire, Hilton B. and Farrant, Pat. Basic Ultrasound. Chichester: John Wiley & Sons, 1995.

[3] http://www.ondacorp.com/

[4] http://www.osha.gov/

[5] www.Corbins.com

[6] Polikar, Robi. "The Engineer's Ultimate Guide to Wavelet Analysis: The Wavelet Tutorial." http://users.rowan.edu/~polikar/WAVELETS/WTtutorial.html

http://isb.ri.ccf.org/biomch-1/archives/biomch-1-1997-12/00080.html

http://msds.ogden.disa.mil/msds/owa/web_msds.display?imsdsnr=192682

http://nyny.essortment.com/rtvrubberpolyu_rypr.htm

http://www.ndt-ed.org/EducationResources/CommunityCollege/Ultrasonics/Physics/acousticimpedance.htm

http://www.nottingham.ac.uk/physics/ugrad/courses/mod_home/f33ab5/notes/us/US001intro198.doc

http://www.semcosh.org/polyurethanes.pdf

http://www.stresstel.com/stresstelwebsite/velocity_acoustic.pdf

http://www.sonotech-inc.com/PDF/Medical_Technical_Paper_2003.pdf

http://www-ibt.etec.uni-karlsruhe.de/people/mag/frames/papers/EMC99-MD/node3.html

https://www.cstsales.com/Material_Data/hs_iii_rtv_silicone_mold_making.htm

Appendix A:
The maxcurve and clothmax programs

Appendix B:
Reflection coefficients between different materials and air

Appendix C:
Corbin's data on ballistic gel

Appendix D:
Key wavelet graphs

Appendix E:
Haaretz.com article-"The 100th suicide bomber" by Amos Harel