



TIME-FREQUENCY ANALYSIS
OF
TERAHERTZ RADAR SIGNALS
FOR
RAPID HEART AND BREATH RATE DETECTION
THESIS

Melody L. Massar

AFIT/GAM/ENC/08-05

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT/GAM/ENC/08-05

TIME-FREQUENCY ANALYSIS
OF
TERAHERTZ RADAR SIGNALS
FOR
RAPID HEART AND BREATH RATE DETECTION

THESIS

Presented to the Faculty
Department of Mathematics and Statistics
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
In Partial Fulfillment of the Requirements for the
Degree of Master of Science

Melody L. Massar, BS

June 2008

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

TIME-FREQUENCY ANALYSIS
OF
TERAHERTZ RADAR SIGNALS
FOR
RAPID HEART AND BREATH RATE DETECTION

Melody L. Massar, BS

Approved:

/signed/

03 June 2008

Dr. Matthew Fickus (Chairman)

Date

/signed/

03 June 2008

Dr. Douglas T. Petkie (Member)

Date

/signed/

03 June 2008

Dr. Andrew J. Terzuoli, Jr. (Member)

Date

/signed/

03 June 2008

Dr. Mark E. Oxley (Member)

Date

Abstract

We develop new time-frequency analytic techniques which facilitate the rapid detection of a person's heart and breath rates from the Doppler shift the movement of their body induces in a terahertz radar signal. In particular, the Doppler shift in the continuous radar return is proportional to the velocity of the person's body. Thus, a time-frequency analysis of the radar return will yield a velocity signal. This signal, in turn, may undergo a second time-frequency analysis to yield any periodic components of the velocity signal, which are often related to the heart and breath rates of the individual. One straightforward means of doing such an analysis is to take the spectrogram of the ridgeline of the spectrogram of the radar signal. Instead of exactly following this approach, we consider an alternate method in which the ridgeline of the radar signal's spectrogram is replaced with a signal computed from spectral centroids. By using spectral centroids, rather than the ridgeline, we produce a smooth signal that avoids some traditional problems with ridgelines, such as jump discontinuities and overquantization. This new method for time-frequency analysis uses a Toeplitz matrix-based algorithm that has a fast Fourier transform-based implementation, and permits centroids of the vertical strips of the spectrogram of the radar signal to be computed without ever having to explicitly compute the spectrogram itself. This algorithm has a lower computational cost than the ridgeline method, and allows us to increase our frequency resolution. We conclude by testing these ideas on real-life data, successfully determining the heart and breath rates of a subject a distance of 10 meters from the radar aperture.

Acknowledgements

First and foremost, I would like to thank my friends and family for all of their encouragement and support. I would also like to thank Dr. Terzuoli for introducing me to all the opportunities at AFIT. Thank you for all of the valuable insight you provided me during my transition into graduate school. Also, thank you for allowing me to work under your supervision as a SOCHE student, research assistant. I would also like to express my appreciation to Dr. Petkie for providing real-life data upon which we tested our ideas. To Dr. Fickus, my thesis advisor, I am very thankful. The time and help you have provided on this thesis was invaluable. Working with you on this research has pushed me to do my best and work hard. I have learned so much in the last four quarters. I know that teaching me the new material took a great amount of time. You were so willing to help, no matter how long it took. I am extremely grateful for your patience and all of the time you have spent working with me.

Melody L. Massar

Table of Contents

	Page
Abstract	iv
Acknowledgements	v
List of Figures	viii
I. Introduction	1
II. Literature Review	6
2.1 Time-Frequency Analysis	6
2.2 Terahertz Radar	7
2.2.1 History of Terahertz Radar	7
2.2.2 Applications of Terahertz Radar	8
2.3 Heart and Breath Rate Detection	8
2.3.1 Contact Methods	8
2.3.2 Non-Contact Methods	10
III. Time-Frequency Analysis and Continuous Wave Radar	15
3.1 The Terahertz Radar System	15
3.2 Frequency	17
3.3 Instantaneous Frequency	21
3.3.1 The Doppler Effect and Instantaneous Frequency	21
3.3.2 The Windowed Fourier Transform	23
3.4 Time-Frequency Analysis of Digital Signals	25
3.5 Spectrograms	27
3.6 Time-Frequency Analysis of Periodic Digital Signals	35
3.6.1 The Fast Fourier Transform	35
3.6.2 Convolutions	37
3.6.3 Computing Discrete Spectrograms	38
IV. A New Method for Computing Spectral Centroids	41
4.1 Weighted Spectral Sums	41
4.2 Spectral Moments	48
4.2.1 Preprocessing by High-Pass Filtering	52

	Page
V. Experimental Results	54
5.1 Examples	54
5.1.1 Example 1	54
5.1.2 Example 2	57
5.1.3 Example 3	59
5.2 Spectral Centroids of the Velocity Centroid's Spectrogram	61
5.2.1 Example 4	62
5.2.2 Example 5	63
5.2.3 Example 6	64
5.2.4 Example 7	64
Bibliography	69

List of Figures

Figure		Page
1	Properties of terahertz radar signals.	2
2	Signal, its corresponding spectrogram, the ridgeline of the spectrogram, and the spectrogram of the ridgeline.	4
3	EKG and respiration belt signals and their corresponding spectrograms.	9
4	Simulated data.	18
5	A sinusoidal curve and a sinusoidal curve with noise.	19
6	The advantage of using Fourier transforms versus counting zero crossings.	21
7	Comparing the Fourier transform with the windowed Fourier transform.	22
8	Three-dimensional view of Figure 7(c).	28
9	Comparing the ridgeline and centroid of a spectrogram that was computed over frequencies with high resolution.	32
10	Comparing the ridgeline and centroid of a spectrogram that was computed over frequencies with coarse resolution.	33
11	Time-frequency analysis of the simulated data and its centroids . .	36
12	The discrete Fourier transforms of ψ_{zp} for $w(x) = 1, x$, and x^2 . . .	49
13	Spectral centroids (plus or minus their spectral deviations) of Figure 7(c).	52
14	Spectral deviations and time-frequency analysis of the spectral centroid for Example 1.	55
15	Truth data for heart and breath rate and their corresponding spectrograms for Example 1.	56
16	Spectral deviations and time-frequency analysis of the spectral centroid for Example 2.	57
17	Truth data for heart and breath rate and their corresponding spectrograms for Example 2.	58


Figure		Page
18	Spectral deviations and time-frequency analysis of the spectral centroid for Example 3.	59
19	Truth data for heart and breath rate and their corresponding spectrograms for Example 3.	60
20	Spectrogram showing heart and breath rate.	62
21	Heart and breath rate detection in Example 4.	65
22	Heart and breath rate detection in Example 5.	66
23	Heart and breath rate detection in Example 6.	67
24	Heart and breath rate detection in Example 7.	68

TIME-FREQUENCY ANALYSIS
OF
TERAHERTZ RADAR SIGNALS
FOR
RAPID HEART AND BREATH RATE DETECTION

I. Introduction

Time-frequency analysis is the study of how the frequency of a signal changes over time. The Fourier transform is typically used to extract information regarding the frequency of a signal. However, the Fourier transform is not appropriate when dealing with a signal whose frequency content changes over time. A time-frequency transform extracts both time and frequency information of a signal, that is, it indicates what frequencies are present in the signal at any given time.

The data used in this study of time-frequency analysis is provided by Dr. Douglas Petkie's terahertz radar system, located in the Department of Physics at Wright State University. This radar emits a continuous wave of electromagnetic radiation at either 120 or 240 GHz in frequency. Radiation at these frequencies has the remarkable ability to pass through clothing, but reflect off of skin. In particular, even minute motions of a person standing in the radar beam can be detected as a result of the Doppler shift this movement induces in the frequency of the reflected wave. In particular, biological features such as heart and breath rates can be measured. This property of the radar's radiation penetrating clothing can be seen in Figure 1(a), provided by Dr. Petkie, where at 120 and 240 GHz the fraction of radar energy that penetrates the clothing is very high. Another important property of these particular frequencies of radiation is the narrow beamwidth. In particular, as seen in Figure 1(b), even at



(a) The fraction of radar energy that penetrates clothing is very high, thus allowing the radiation to penetrate clothing, but reflect off of the chest wall. Therefore, we can measure the Doppler shift caused by the motion of the chest wall and use this to measure biological features such as heart and breath rates.

(b) Beam size as a function of distance with a 6 inch aperture. When the subject is at 50 meters, we can see that the beam width of a 240 GHz wave is approximately the width of an adult facing the transmitter.

Figure 1: Properties of terahertz radar signals.

distances of 50 meters, the beam is narrow enough so as to only illuminate a single individual, which helps mitigate the effects of clutter.


A basic outline of the pertinence of time-frequency analysis to the study of these radar signals may be seen in Figure 2. In particular, Figure 2(a) shows a small segment of 120 second's worth of a 240 GHz radar signal which has reflected off a test subject a distance of 10 meters from the transmitter/receiver. To be precise, Figure 2(a) shows the digital signal obtained by demodulating the received radar signal by 240 GHz, and then sampling the result at 10 kHz. As discussed in detail below in Chapter III, the instantaneous frequency of this signal at any given time is the Doppler shift of the radar signal which is proportional to the velocity of the object from which the beam is reflecting. As seen in Figure 2(b), this instantaneous frequency may be numerically estimated using a common tool of time-frequency analysis called the spectrogram. By taking the *ridgeline* of this spectrogram, that is, by determining the dominant frequency in the signal at any given time, we produce the curve seen in Figure 2(c). As this *ridgeline* provides a good numerical approximation of the velocity

of the subject at any given time, this signal must be further analyzed for the detection of breaths and heartbeats. In particular, as breathing and the beating of the heart cause periodic motions, the ridgeline in Figure 2(c) may be time-frequency analyzed using a spectrogram, as in Figure 2(d). This second time-frequency analysis serves to separate gross motion (the subject walking towards or from the radar) from periodic motion (heartbeats, breathing, and gait). In this particular example, the subject was at rest, and so only heart and breath motions were registered. In particular, as seen in Figures 2(b) and (d), the subject is breathing between 10 and 40 seconds and between 55 and 100 seconds. It is clear from Figure 2(d) that the subject's breath rate is about 0.3 Hz, meaning the subject takes almost one breath every 3 seconds. Similarly, as seen in Figures 2(b) and (d), the subject is holding his breath between 40 and 55 seconds. From Figure 2(d), the subject's heart rate is about 1 Hz, meaning the subject's heart beats about once every second.


The goal of this research is to improve the analysis presented in Figure 2. In particular, we develop a Toeplitz matrix-based, FFT implemented method for quickly computing a signal close to, but less noisy than, the ridgeline in Figure 2(c). This approach completely bypasses the need to compute the first spectrogram in Figure 2(b), and provides cleaner data to the second spectral analysis in Figure 2(d). These improvements make it possible to process radar data in real-time using standard hardware and, more significantly, determine the heart rate even when the subject is not holding his breath.

Real-life applications of this theory include the areas of security and ground combat. For example, a high heart and breath rate would alert security personnel to consider having the subject undergo additional screening. In terms of ground combat, being able to identify the living from the dead at large distances would better inform commanders in further mission planning.


We begin in Chapter II with a short review of time-frequency analysis, Terahertz radar, and heart and breath rate detection literature. In Chapter III, we take a closer




(a) A five second interval of the radar signal sampled at 10 kHz (real part plotted in blue and imaginary part plotted in red).



(b) The spectrogram of the radar signal shows us what frequencies are present at what times. These frequencies are proportional to the velocity of the target.



(c) The ridgeline of (b), which represents the dominant frequency at any time, is a quantity which estimates the Doppler shift in the radar signal at any given time. The ridgeline should be a close approximation of the velocity signal. The large sinusoidal motion, seen from 10 to 40 seconds and from 55 to 100 seconds, corresponds to the subject's breathing. The portion of the ridgeline between 40 and 55 seconds has a higher frequency and is a result of the subject holding his breath, showing only motion caused by the heart.



(d) The spectrogram of the velocity signal (c) should show the periodic components related to heart and breath motion. Here we can see that between 40 and 55 seconds, the subject is holding his breath. During this time, we can more clearly see the subject's heartbeat at about 1 Hz, meaning the subject's heart beats about one time every second. We can also see that the subject's breath rate is about 0.3 Hz, meaning the subject takes almost 1 breath every 3 seconds.

Figure 2: Signal, its corresponding spectrogram, the ridgeline of the spectrogram, and the spectrogram of the ridgeline.

look at the THz radar system that provides the data for our time-frequency analysis. We also review some of the major topics that are used throughout the thesis such as the Fourier transform, the Doppler effect, and the windowed Fourier transform. The advantages of using spectral centroids to extract heart and breath rate information are also discussed. In Chapter IV, we introduce a new Toeplitz matrix-based approach for the computation of spectral centroids, and also provide an algorithm that may be used to implement this new approach. We conclude in Chapter V by validating the algorithm on real-life data.

II. Literature Review

2.1 *Time-Frequency Analysis*

Consider obtaining the individual densities of height and weight of a certain type of animal. These densities tell us about the distribution of height and distribution of weight. However, they do not tell us how height and weight are related, as this would require joint density of height and weight. Likewise, a signal may not be fully understood from simply obtaining its time energy density and frequency energy density. For this reason, it became important to understand and devise a distribution that represents the energy of a signal in both time and frequency. There are several examples in everyday life that relate to this idea of time-varying spectra. For example, the frequency of light changes during sunset, the pitch in our voice changes as we speak, and the notes a musician plays change over time [5].

One of the most commonly used methods for signals of non-stationary frequency is the windowed Fourier transform. Suppose we were listening to an hour long flute recital that had piano accompaniment. If we Fourier analyze the entire hour, we will see frequency peaks in the spectrum corresponding to the flute and the piano. This only tells us that the flute and piano were being played and does not tell us when they were being played. Suppose we then break up the recital into five minute segments and then Fourier analyze each segment. We can then determine in which five minute segments the piano and/or flute were being played. Making the time intervals even smaller, we can localize even better and therefore can determine exactly when the flute and piano were being played [5]. To obtain this localized spectrum we multiply the signal $f(s)$ by the window $g(s)$ centered at time $s = t$ giving us:

$$f_w(t, s) = f(s)g(s - t). \quad (1)$$

Taking the Fourier transform of (1) gives us the short-time Fourier transform. As noted in [1], using the window $g(s)$ “allows localization of the spectrum in time, but

also smears the spectrum in frequency in accordance with the ‘uncertainty relationship’, leading to a trade-off between time resolution and frequency resolution.”

One tool of time-frequency analysis is the *spectral centroid*, which measures the spectral energy distribution of a signal. In [17], the *spectral centroid* is said to be calculated as “the sum of the frequencies weighted by the amplitudes, divided by the sum of the amplitudes, which is the first moment of spectrum with respect to frequency.” In most applications, it is used to measure the distribution of a tone. The *spectral centroid* can be used for musical instrument recognition, instrument sound description, and auditory scene recognition, as seen in [8], [16], and [17].

We apply these ideas to a radar signal (electromagnetic wave) rather than to a sound signal (pressure wave). More specifically, we apply these ideas to the signal of a continuous-wave terahertz radar. Here, the changing frequency is caused by the Doppler effect induced by the motion of the target. However, unlike many classical radar applications, the target here is not an aircraft. Here, the target is a human being, and the Doppler shift is caused by the motion of their body.

2.2 *Terahertz Radar*


2.2.1 History of Terahertz Radar. The term *terahertz* is applied to submillimeter wave energy that fills the wavelength range between 1 mm (300 GHz) and 100 μm (3 THz). The wavelength range between 1 cm (30 GHz) and 1 mm (300 GHz) crosses into the *millimeter-wave bands* [18] and [20]. Throughout the last century, the millimeter and submillimeter wave generation started to become a topic of interest to scientists [6]. By 1984, the frequency range from 30 to 100 GHz was in a state of advanced development, but above 100 GHz was still very exploratory. Wiltse said, “Perhaps what is most needed is the invention of new sources that would provide reasonable coherent power with simplicity and good efficiency at frequencies from 100 GHz into the submillimeter” [20]. By 2002, the terahertz frequency range was one of the least explored regions of the electromagnetic spectrum [18].

2.2.2 Applications of Terahertz Radar. “The universe is bathed in terahertz energy; most of it going unnoticed” [18]. From about 1990 to 2004, some of the major applications of terahertz technology included applications in the areas of space sciences, molecular spectroscopy, and plasma diagnostics. Around the beginning of the 21st century, terahertz technology started expanding its applications to areas such as biology and medicine. Some of these applications include disease diagnostics, tumor recognition, and performing label-free DNA sequencing [19]. In 2002, Siegel said, “All of these exciting applications and countless undiscovered ones remain in wait while terahertz technology enters adulthood—it has been a long time coming and there is still much work to be done” [18].


2.3 Heart and Breath Rate Detection

Whether it is simply monitoring an individual’s health, diagnosing chronic health conditions, or even searching for humans in earthquakes or avalanches, being able to accurately detect heart and breath rates has a wide range of security applications. There are many different methods to heart and breath rate detection. Some of these methods are contact based and some are non-contact based. Almost all of the papers discussing measuring an individual’s heart and breath rates require some type of frequency estimate. As technology becomes more advanced and the number of biological and biomedical applications of radar are increasing, much of the recent focus has been on non-contact methods.


2.3.1 Contact Methods. A person’s heart rate may be accurately measured using an electrocardiogram (EKG). The muscles in one’s heart contract via electricity. The electrocardiogram records this electrical activity over time. An example of the EKG signal can be seen in Figure 3(a). Electrodes, or electrical contacts, are attached to the skin and measure the electrical waves that pass through the body. Meanwhile, a person’s breath rate may be measured using a *respiration belt*. The belt is strapped around the abdomen or chest and then inflated. The changes in air pressure caused




(a) The EKG signal which records the electrical activity that causes the heart to contract.



(b) The spectrogram of (a). This spectrogram shows us what frequencies are present in the EKG signal (a) at a given time. We see that the heart rate is around 1 Hz, meaning that the subject's heart beats about 1 time every second. The harmonic at 2 Hz is a result of the EKG signal (a) having sharp peaks. This harmonic is not present in (d) because the respiration signal (c) is sinusoidal.



(c) The respiration belt signal which is obtained by strapping the respiration belt around the abdomen or chest, inflating it, and then measuring the changes in pressure caused by inhalation and exhalation.



(d) The spectrogram of (c). This spectrogram shows us what frequencies are present in the respiration belt signal (c) at a given time. From this, we can see that the respiration rate is around 0.25 Hz, meaning that the subject takes about 1 breath every 4 seconds.

Figure 3: EKG and respiration belt signals and their corresponding spectrograms.

by inhalation and exhalation are then recorded. An example of the respiration belt signature can be seen in Figure 3(c). As seen in Figures 3(b) and (d), standard time-frequency analysis of the EKG and pressure belt signals will reveal heart and breath rates, respectively. From the spectrogram in Figure 3(b), we can see a dominant frequency at slightly more than 1 Hz, indicating the subject’s heart beats about 1 time per second. Also, we can see from the spectrogram in Figure 3(d) that the dominant frequency occurs at about 0.25 Hz, indicating the subject takes approximately one breath every four seconds.

Though invaluable medical tools, pressure belts and EKGs require sensors to be attached to the subject, and as such, are inappropriate for the security checkpoint and battlefield environments. Instead, for these situations, what is needed is a non-contact method for measuring heart and breath rates. Indeed, as noted in [12]: “A more promising way of detecting heart rate, are contact-free measurements through radar: Radar waves are emitted and frequency analysis of the reflected signal from human body reveals heart rate based on Doppler phenomena.” We now review the existing technology to this end in detail.

2.3.2 *Non-Contact Methods.*

2.3.2.1 Micro-Impulse Radar (MIR). Micro-impulse radar (MIR) uses short radar impulses which last for only a few nanoseconds. The goal is to shape a “wearable heart rate sensing solution” which will operate in close proximity to the user. “The radar unit is worn by the user without the need of placing any additional sensor onto the body” [13]. The apparatus itself is detailed in [13]:

A MIR exists of several units. A pulse generator defines when the transmitter should emit a pulse over the antenna. Simultaneously, the pulse-generator activates a so-called delay line. This delay line is used for controlling the sampling of the received echoes at the receiver: the receiver is only activated at very short time intervals triggered by the delay-line. Thus, the length of the delay-line ensures that only pulses back-scattered from a certain distance are received. In the context of heart rate detection

the goal is to adjust the delay-line such that the receiver is activated only if echoes from the heart wall can be expected.

The algorithm that uses MIR to extract heart rate has four steps: filtering, local maxima detection, evaluation of distance between maxima, and division. Filtering the radar signal according to:

$$\tilde{x}(t_{i+1}) = Cx(t_{i+1}) + (1 - C)x(t_i)$$

smoothes the data to allow better processing. In a signal sampled at about 80 Hz, they found $C = 0.15$ to be an appropriate weight. The algorithm then finds local maxima in the signal, and then calculates the distances (in time) between them. These distances are then analyzed for regularly occurring patterns: “occurrences of a distance four times or higher has proven as a reliable indication of a regularly occurring maximum” [13]. Maxima with this distance are typically related to the heart rate. All the remaining non-contact methods that we discuss are based on Doppler principles, as is Dr. Petkie’s system.

2.3.2.2 Radar Vital Signs Monitor (RVSM). The Radar Vital Signs Monitor, developed by researchers at Georgia Tech Research Institute, has been used to measure human heart and breath rates at distances exceeding 10 meters. This radar system operates in the continuous wave mode at 24.1 GHz. A mixer diode is used to detect energy reflected from the target. When subject movement occurs, a Doppler shift between the transmitted and received signals is present, and there is a change in the phase between the transmitted and received signals. Part of this movement is the result of the heart beats and breathing of the subject. The phase difference resulting from this movement is then amplified and filtered. Several filters are applied to the signal to separate the heart and breath rate signals: “the filter bandpass is specified such that it can attenuate the respiration signature while allowing the fine resolution heart signature to pass without inducing filter effects” [9]. The filter is designed so that plasma modulation caused by fluorescent lights at 60 Hz is eliminated. Finally,

the signal is then amplified and sent to a low impedance output stage. The respiration channel is formed using a similar approach.

2.3.2.3 Microwave Techniques. Microwave radiation has been used in many medical applications, such as for the diagnosis and monitoring of pulmonary edema and other pathological cardiopulmonary conditions [15]. It has also been used to record apexcardiograms, which indicate precordial movements [11]. Microwave radiation can also be used to sense vital signs such as heart and breath rates. Coupled with its ability to penetrate rubble, microwave-based devices can be used to locate avalanche and earthquake victims.

One microwave technique [10] for measuring respiration is based on scattering of continuous wave radiation, and operates at a frequency of 10 GHz. Once the signal reaches the subject, it is modulated in amplitude and phase. This is a result of the movement of the chest wall due to breathing. The scattered energy modulated by this movement is detected by a crystal detector that is located on the receiving horn. The ratio meter compares the amplitude with part of the forward signal that is detected by a crystal located on the directional coupler. The instantaneous ratio between the reference and scattered signals is measured by the ratio meter. The ratio meter then outputs a voltage that has frequency corresponding to the breath rate.

Another technique [7] uses a microwave radio for sensing vital signs by Doppler radar. To remove the DC component, the signal is filtered using a bandpass filter between 0.03 Hz and 10 Hz. This filter also minimizes aliasing error and out-of-band noise. After this filtering, the breath rate is clearly visible. Applying a low-pass filter better resolves the respiration signal. To isolate the heart rate signal, a band-pass filter between 1 Hz to 3 Hz with a 6 dB roll-off is applied to the signal. This method can detect heart and breath rates at distances as large as one meter.

The microwave life-detection system provides applications such as locating earthquake and avalanche victims that are trapped under rubble. In [2], an “X-band microwave life-detection system” is discussed. This system was able to remotely detect

the heart and breath rate signals of subjects at distances of 100 feet or located behind a barrier wall. To detect the heart and breath rate signals, the system illuminates the subject with a low-intensity microwave beam. As with the other Doppler schemes, the small amplitude body-vibrations, caused by the breathing and heartbeat of the subject, modulate the backscattered microwave signal. The heart and breath rate signals are then extracted from the backscattered signal by phase-detection in the microwave receiving system. In this system, clutter from the rubble or surface of the ground was cancelled using a very slow manual-adjustment process. This would not be practical for real-world applications in emergency rescues where fast processing is essential [4]. To handle this problem, a new system with an automatic clutter-cancellation subsystem was developed. The microwave life-detection system, discussed in [4], operates at 2 or 10 GHz and remotely senses movements such as breathing and heartbeat. The automatic clutter-cancellation subsystem uses a microprocessor-based control unit to scan the attenuator and phase-shifter. This minimizes the input signal to the microwave amplifier and therefore cancels the clutter component of the signal. Experiments show that the 2 GHz system penetrated more rubble than the 10 GHz system, concluding that lower frequencies will be more capable of detecting a human's vital signs through thick barriers. The 2 GHz system penetrated rubble of up to 3 feet of thickness.

As the previously mentioned life-detection systems were still not sufficient to locate humans that could be buried deep in earthquake rubble or hidden behind very thick barriers, a “new sensitive microwave life-detection system” was constructed. This new system operated at 1150 or 450 MHz and was used to locate humans buried under earthquake rubble or hidden behind very thick barriers of approximately 10 foot thickness. These different frequencies had different advantages and disadvantages:

An EM wave of 1150 MHz can penetrate a rubble with layers of reinforced concrete slabs with metallic wire mesh easier than that of 450 MHz. However, an EM wave of 450 MHz may penetrate deeper into a rubble without metallic wire mesh than that of 1150 MHz. [3]

The idea behind extracting the heart and breath rates is quite similar to the other systems. After penetrating the rubble to reach the subject, the microwave beam illuminates the human subject. The movement of the human body induces a Doppler shift in the reflected wave. If the reflected wave from the background is cancelled and the reflected wave from the human subject is properly demodulated, the heart and breath rate signals can be extracted.

2.3.2.4 Heart Beat Monitoring by Laser Doppler Interferometry. Another non-contact method for heartbeat monitoring uses laser Doppler interferometry to optically record the movements of the chest wall. This method allows measurements for distances at tens of meters, and is based upon the fact that:

During ventricular contraction, the heart undergoes changes in volume as well as variations in position. The resulting combination of motions is transmitted to the surface of the skin and could be picked up by a laser displacement sensor pointing at a point on the thorax, near the heart. [14]

A laser beam passes through a beam splitter, creating a measuring beam which focuses on the vibrating object, and a second reference beam. Comparing the frequency of the reflected beam to the reference beam, one computes a Doppler shift which, in a manner similar to the technique above, may be further processed to determine heart rate.

III. Time-Frequency Analysis and Continuous Wave Radar

In this chapter, we introduce some of the major topics of time-frequency analysis. We begin by taking a closer look at the actual radar system from which our real-world data is obtained. We then focus on the areas of frequency and instantaneous frequency including topics such as the Fourier transform, the Doppler effect, and the windowed Fourier transform. In particular, we take a closer look at how the processing of analog signals may be approximated by similar processing of their digital samples.

3.1 The Terahertz Radar System

Dr. Petkie’s radar system operates on principles similar to those of Michelson interferometers and operates at either 120 or 240 GHz. Using “binary frequency shifting,” the real-valued radar signal is preprocessed to produce a complex-valued digital representation of the received radar signal.

Formally speaking, let us assume that the transmitter and receiver are located at the origin. Let the target be defined at some point $x(t) \in \mathbb{R}$. The signal transmitted at time t can be defined as

$$S_T(t) := e^{2\pi i \alpha t}, \quad (2)$$

where the symbol “ $:=$ ” denotes “defined as”. Here, α is the frequency of the transmitted signal. In practice, α will be either 120 or 240 GHz. The signal received at time t is assumed to be

$$S_R(t) := e^{2\pi i \alpha (t - 2x(t)/c)}, \quad (3)$$

or simply the signal transmitted at time $t - 2x(t)/c$, where c is the speed of light. This assumption is known as the *start-stop approximation*. The idea behind the assumption is as follows: since the target is located at some point $x(t)$, we expect the signal to travel a distance of $2x(t)$, the distance from the transmitter to the target and then back to the receiver. This signal is travelling at the speed of light and so the time it takes to travel this distance is $2x(t)/c$. So, a part of the signal received at a time t was transmitted at a time $t - 2x(t)/c$. The reason this analysis is somewhat flawed is

that the target itself is moving as the chest wall moves very slightly as a result of the subject's heart and breath motion. In particular, the wave does not actually travel $2x(t)$ units of distance, but rather $2x(\tau)$ units of distance, where $\tau < t$ is the time at which the wave bounced off the subject. However, as the speed of light is much larger than the speed of the subject's body, the assumption (3) is correct to a high degree of accuracy. Note that in (3) we have also ignored the possibly time-varying amplitude of the received signal, which in real-life applications is a result of changes in radar cross-section.

From (3), the signal received at time t is equal to $e^{2\pi i(\alpha t - 2\lambda^{-1}x(t))}$, where $\lambda = c/\alpha$ is the wavelength of the radar wave; for $\alpha = 300$ GHz, $\lambda = 1$ mm. Once the radar signal bounces off the subject, its energy is collected at the receiver. This analog signal is then demodulated by mixing it with a reference signal or a local oscillator, shifting the frequency α to zero. That is, the modulated received signal is

$$S_{MR}(t) := e^{-2\pi i \alpha t} S_R(t) = e^{-2\pi i(2\lambda^{-1}x(t))}. \quad (4)$$

The real-world data we receive is essentially a digitally-sampled version of (4). In the next section, we begin to discuss how a spectral analysis of this data may be used to gain information about the subject. In particular, the velocity of the subject's body may be determined from (4) as a result of the Doppler effect. It turns out that the target's velocity $\dot{x}(t)$ is proportional to the *instantaneous frequency* of $S_{MR}(t)$. We conclude this section by discussing how (4) may also be used to generate simulated data for the evaluation of heart and breath rate detection algorithms. In particular, for a person at rest whose heart and breath rates are constant, $x(t)$ may be approximated as the sum of the target distance, a breath rate term, and a heart rate term. In particular, we model the displacement $x(t)$ of the person from the radar receiver as:

$$x(t) = \beta + \gamma_1 \sin(2\pi\omega_1 t) + \gamma_2 \delta(\omega_2 t - \tau), \quad (5)$$

where β is the overall fixed distance from the person to the radar, γ_1 is their breathing displacement, ω_1 is their breath rate, γ_2 is their heartbeat displacement, ω_2 is their heart rate, τ is some heart rate shift, and δ is defined as:

$$\delta(t) := \frac{\gamma_2}{1-2a} \left(\left| t - \frac{1}{2} - \lfloor t \rfloor \right| - a + \left| \left| t - \frac{1}{2} - \lfloor t \rfloor \right| - a \right| \right)$$

for $a = 1/2 - r\omega_2$, where r is their *heartbeat radius*, a horizontal scale parameter which determines the width of the peaks in Figure 4(b).


Using the parameters in Table 1, we can model the displacement $x(t)$ by (5) and then model the modulated received signal $S_{MR}(t)$ by (4). The first 6.25 seconds of the respiration signal $\gamma_1 \sin(2\pi\omega_1 t)$ can be seen in Figure 4(a). Figure 4(b) shows the first 6.25 seconds of the heartbeat signal $\gamma_2 \delta(\omega_2 t - \tau)$. Figures 4(b) and (c) show the first 6.25 seconds of $x(t)$ and $S_{MR}(t)$.

Table 1: Parameters Used in Simulating Data in Figure 4


Name of Parameter	Symbol	Value
radar frequency	α	240 gigahertz
maximum time		2^5 seconds
sampling rate		2 kilohertz
target distance	β	8 meters
breathing displacement	γ_1	4.5 millimeters
breath rate	ω_1	0.23 breaths per second
heartbeat displacement	γ_2	0.7 millimeters
heart rate	ω_2	1 beat per second
heartbeat radius	r	0.25 meters
heart rate shift	τ	0 seconds

3.2 Frequency


Frequency can be defined as the number of instances a certain event occurs over a given amount of time. As the frequency of an event may itself change over time, time-frequency analysis, that is, the study of how the frequency of a signal changes over time, has been developed. There are many ways to calculate the frequency of a signal. For a purely sinusoidal signal, frequency is inversely related to the wavelength




(a) The simulated respiration signal from 0 to 6.25 seconds. The vertical axis is on a scale of 1 mm.



(b) The simulated heartbeat signal from 0 to 6.25 seconds. The vertical axis is on a scale of 0.1 mm.



(c) The sum of the target distance, the simulated heartbeat signal (a), and the simulated respiration signal (b) from 0 to 6.25 seconds. This sum is denoted $x(t)$. The vertical axis is on a scale of 1 m.



(d) The real part of the simulated data from 0 to 6.25 seconds.

Figure 4: Simulated data.




Figure 5: A sinusoidal curve and a sinusoidal curve with noise.

for a traveling wave. Another way to calculate frequency is to measure the time it takes for an event to repeat (the period). The frequency of the event is then simply the reciprocal of this time.

For example, one may use the distance between zero crossings of a real-valued signal to estimate its frequency. To be precise, a *zero crossing* of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is a point at which the horizontal axis is crossed, the sign changes from positive to negative or vice versa. For sinusoidal signals, a single period will contain two zero crossings, this is illustrated in Figure 5(a). Even with a small amount of noise present, it is not a good idea to use this method to calculate frequency, as the noise causes many zero crossings to occur near a true zero crossing of the clean signal as seen in Figures 5(b) and (c). In such a noisy environment, the frequency of a signal is better estimated using a technique that sums together data of a relatively long period of time, hopefully allowing some of this noise to cancel itself out. The canonical example of such a technique is the *Fourier transform*.

To be precise, let us consider the function space:

$$L^2(\mathbb{R}) := \left\{ f : \mathbb{R} \rightarrow \mathbb{C} \mid f \text{ is Lebesgue measurable, } \int_{-\infty}^{\infty} |f(t)|^2 dt < \infty \right\},$$

which is a Hilbert space under the inner product:

$$\langle f, g \rangle := \int_{-\infty}^{\infty} f(t) \overline{g(t)} dt.$$

We use the space $L^2(\mathbb{R})$ as a model of the continuous radar signal. The Fourier transform on $L^2(\mathbb{R})$ is the operator $\mathcal{F} : L^2(\mathbb{R}) \rightarrow L^2(\mathbb{R})$, defined as:


$$(\mathcal{F}f)(x) := \int_{-\infty}^{\infty} f(t) e^{-2\pi i x t} dt$$

for all $f \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. For any given $x \in \mathbb{R}$, the quantity $(\mathcal{F}f)(x)$ represents the degree to which the signal f contains a component whose frequency is x . Indeed, in light of the Fourier inversion formula:


$$f(t) = \int_{-\infty}^{\infty} (\mathcal{F}f)(x) e^{2\pi i x t} dx,$$

the values of $\mathcal{F}f$ tell one how to decompose any $f \in L^2(\mathbb{R})$ as a sum of sinusoids.


Figure 6 illustrates just how superior the Fourier transform is compared to counting zero crossings when it comes to estimating frequency. In particular, in Figure 6(a), we see a typical signal of Gaussian amplitude and near constant frequency. We note the frequency of the clean signal may be either determined by counting zero crossings or by computing its Fourier transform, as shown in Figure 6(b). However, when the original signal is buried in noise, as in Figure 6(c), the zero crossing method will fail to indicate the signal's frequency, whereas this information is still discernible in the Fourier domain (Figure 6(d)). This is not to say the Fourier transform is perfect. Indeed, by integrating over the entire real line, the Fourier transform may have the unfortunate effect of blending all of the signal's frequencies together. As we now discuss, the solution to this problem is to instead consider a windowed Fourier




(a) A modulated Gaussian.



(b) The magnitude of the Fourier transform of (a). The frequencies of the signal are clearly visible.



(c) A plot of (a) with random noise added. It does not look as if we could recover the signal (a) from this noisy signal by counting zero crossings.




(d) The magnitude of the Fourier transform of (c). Even though there is noise present in the signal, we are nevertheless able to determine the dominant frequencies in this domain.

Figure 6: The advantage of using Fourier transforms versus counting zero crossings.

transform, which better permits one to estimate the signal's frequency at any given time.

3.3 *Instantaneous Frequency*

3.3.1 The Doppler Effect and Instantaneous Frequency. Consider analyzing the frequencies of a signal whose frequencies are changing over time. Figure 7(a) shows a portion of a simulated signal (4) that represents a signal that has been reflected off of a person that was breathing at a constant rate. As the Fourier transform of this



(a) The simulated data without heartbeat over the time interval 34.6 to 35 seconds.

(b) The magnitude of the Fourier transform of (a) which shows us that many frequencies are present in the signal, but does not give us any indication of when these frequencies occur.

(c) The spectrogram of (a) shows us what frequencies are present at what time. We can see that this gives us more beneficial information, especially when our signal has frequency content that changes over time.

Figure 7: Comparing the Fourier transform with the windowed Fourier transform.

signal (Figure 7(b)) combines information from all times, it is difficult to tell what the frequency of the signal is at any given moment. To remedy this problem, one may compute the windowed Fourier transform of the signal, defined formally below, and produce the spectrogram given in Figure 7(c). Indeed, from Figure 7(c), we can see that the frequency of our simulated signal is changing over time. To be more precise, however, we must first define what is meant by *instantaneous frequency*.

In particular, consider a particle travelling along the unit circle whose position at any time t is $e^{2\pi i\gamma(t)}$. For any two distinct times t and $t + h$, the distance of the point $e^{2\pi i\gamma(t+h)}$ from the point $e^{2\pi i\gamma(t)}$ along the arc of the circle is $2\pi[\gamma(t+h) - \gamma(t)]$. Since the total circumference of the unit circle is 2π , the net number of cycles made by the point between the time t and time $t + h$ is

$$\frac{2\pi[\gamma(t+h) - \gamma(t)]}{2\pi} = \gamma(t+h) - \gamma(t).$$

Dividing by the change in time, the average frequency, that is, the average number of cycles made by this particle per unit of time is

$$\frac{\gamma(t+h) - \gamma(t)}{h}.$$

The instantaneous frequency of $e^{2\pi i\gamma(t)}$ at t is then defined as the limit of these average frequencies as the time interval h over which these averages are taken tends to zero. That is, the instantaneous frequency of $e^{2\pi i\gamma(t)}$ at time t is defined to be the derivative of $\gamma(t)$ at t :

$$\dot{\gamma}(t) = \lim_{h \rightarrow 0} \frac{\gamma(t+h) - \gamma(t)}{h}.$$

In the particular case of our continuous wave radar signal $S_R(t) = e^{2\pi i\alpha(t-2x(t)/c)}$, the instantaneous frequency is

$$\frac{d}{dt}\alpha\left(t - \frac{2x(t)}{c}\right) = \alpha\left(1 - \frac{2\dot{x}(t)}{c}\right). \quad (6)$$

The Doppler shift is calculated by subtracting the constant frequency α of the transmitted signal $S_T(t)$ from the instantaneous frequency of $S_R(t)$, that is:

$$\begin{aligned} \text{Doppler shift} &= \alpha\left(1 - \frac{2}{c} \frac{d}{dt}x(t)\right) - \alpha \\ &= -2\frac{\alpha}{c} \frac{d}{dt}x(t) \\ &= -\frac{2\dot{x}(t)}{\lambda}, \end{aligned}$$

where λ is the wavelength of the transmitted signal. Thus, the Doppler shift in the received radar signal is proportional to the target's velocity. With respect to our application, this means that we can determine the velocity of the subject's body, provided we can accurately measure the instantaneous frequency of our data at any given time. However, as we noted above in Figure 7(b), the Fourier transform is not always the best way to measure frequency as it blends together all frequency information over all times. Instead, what is needed is a transform that computes frequencies over shorter periods of time.

3.3.2 The Windowed Fourier Transform. Consider listening to an orchestra which is composed of many musicians playing many notes simultaneously. An individual hears the sum of all these notes being played. The purpose of time-frequency

analysis is to undo this sum, that is, to determine what notes are being played at which times.

The windowed Fourier transform is the standard tool used for such an analysis. To be precise, for any $t, x \in \mathbb{R}$, the *modulation* and *translation* operations on the space $L^2(\mathbb{R})$ are the operators $\mathcal{M}^x, \mathcal{T}^t : L^2(\mathbb{R}) \rightarrow L^2(\mathbb{R})$, that are defined as:

$$(\mathcal{M}^x f)(s) := f(s)e^{2\pi i x s}, \quad (7)$$

$$(\mathcal{T}^t f)(s) := f(s - t). \quad (8)$$

For any fixed windowing function $g \in L^2(\mathbb{R})$, the windowed Fourier transform of $f \in L^2(\mathbb{R})$ with respect to g is

$$\begin{aligned} (\mathcal{W}_g f)(t, x) &:= \langle f, \mathcal{T}^t \mathcal{M}^x g \rangle \\ &= \langle \mathcal{T}^{-t} f, \mathcal{M}^x g \rangle \\ &= \int_{-\infty}^{\infty} f(s + t) \overline{g(s)} e^{-2\pi i x s} \, ds. \end{aligned} \quad (9)$$

That is, for any time t , the windowed Fourier transform translates f by a factor of $-t$, then multiplies this function by $\overline{g}$, and then takes the Fourier transform of the result. When g is a Gaussian of standard deviation σ , the values $(\mathcal{W}_g f)(t, x)$ are essentially the Fourier transform of the portion of f over $(t - 3\sigma, t + 3\sigma)$. As both t and x vary, the windowed Fourier transform is an overcomplete representation of f ; a function of one variable has been transformed as a function of two variables, the magnitude of which may either be viewed as a surface (Figure 8) or, more conventionally, as a color coded image (Figure 7(c)). In either case, by the Cauchy-Schwarz inequality, one expects $|(\mathcal{W}_g f)(t, x)|^2 = |\langle f, \mathcal{T}^t \mathcal{M}^x g \rangle|^2$ to be large when f is a scalar multiple of g , and be small otherwise. As $\mathcal{T}^t \mathcal{M}^x g$ is essentially a bump function centered at t and of frequency x , one therefore expects $|(\mathcal{W}_g f)(t, x)|^2$ to be large when f contains a component of frequency x around time t , and to be zero otherwise. Indeed, as seen in

the red sinusoidal curve in Figure 7(c), we see that the simulated data in Figure 7(a) essentially has one dominant frequency at any given time, and that this frequency itself changes sinusoidally over time.

We conclude by noting that even though the radar signal is analog, it is passed through an A/D converter. Therefore, all we have are digital samples of the original signal. For this reason, we need to do time-frequency analysis on a space of discrete signals.

3.4 *Time-Frequency Analysis of Digital Signals*

We now turn to a formal analysis of this problem. Let f, g be continuous signals, and let their digital samples $f_\sigma, g_\sigma : \mathbb{Z} \rightarrow \mathbb{C}$ be defined:

$$\begin{aligned} f_\sigma(n) &= f(\sigma n), \\ g_\sigma(n) &= g(\sigma n), \end{aligned}$$

where σ is equivalent to the inverse sampling rate. That is, $f(\sigma n)$ is the n th sample of f and, similarly, $g(\sigma n)$ is the n th sample of g . The windowed version of the analog signal is $f(u)g(u - t)$. When the sampling rate σ^{-1} grows large, (9) may be approximated by a sum, that is,

$$(\mathcal{W}_g f)(t, x) \approx \sigma \sum_{n=-\infty}^{\infty} f(\sigma n + t) \overline{g(\sigma n)} e^{-2\pi i \sigma x n},$$

provided f and g are sufficiently smooth and have sufficient decay. In particular, for t of the form $t = \sigma m$ for some $m \in \mathbb{Z}$,

$$\begin{aligned} (\mathcal{W}_g f)(\sigma m, x) &\approx \sigma \sum_{n=-\infty}^{\infty} f(\sigma n + \sigma m) \overline{g(\sigma n)} e^{-2\pi i \sigma x n} \\ &= \sigma \sum_{n=-\infty}^{\infty} f_\sigma(n + m) \overline{g_\sigma(n)} e^{-2\pi i \sigma x n}. \end{aligned} \tag{10}$$

Equation (10) leads to a definition of a windowed Fourier transform on a space of discrete signals, in particular, the space $\ell^2(\mathbb{Z})$.

To be precise, we note that since we only have discrete data, we will drop the σ in our notation and from this point forward let f and g denote discrete functions, and let

$$\ell^2(\mathbb{Z}) := \left\{ f : \mathbb{Z} \rightarrow \mathbb{C} \left| \sum_{n=-\infty}^{\infty} |f(n)|^2 < \infty \right. \right\},$$

which is a Hilbert space under the inner product

$$\langle f, g \rangle := \sum_{n=-\infty}^{\infty} f(n) \overline{g(n)}.$$

In a manner similar to the space $L^2(\mathbb{R})$, one may define a windowed Fourier transform on $\ell^2(\mathbb{Z})$ provided one first defines translation and modulation operators on these spaces. In particular, we define the translation operation, $T^n : \ell^2(\mathbb{Z}) \rightarrow \ell^2(\mathbb{Z})$, as:

$$(T^n f)(m) := f(m - n),$$

where $m, n \in \mathbb{Z}$. Also, we can define the modulation operation, $M^x : \ell^2(\mathbb{Z}) \rightarrow \ell^2(\mathbb{Z})$, as

$$(M^x f)(m) := f(m) e^{2\pi i x m},$$

where $m \in \mathbb{Z}$ and $x \in \mathbb{R}$. Note, however, that with respect to modulation, we have:

$$(M^{x+n} f)(m) := f(m) e^{2\pi i x m} (e^{2\pi i m})^n = f(m) e^{2\pi i x m} = (M^x f)(m),$$

that is $M^{x+n} = M^x$ for all $n \in \mathbb{Z}$. Thus, the modulation parameter x is most appropriately regarded as an element of the circle group $\mathbb{T} := \mathbb{R}/\mathbb{Z}$. The Fourier transform on $\ell^2(\mathbb{Z})$ is defined as $F : \ell^2(\mathbb{Z}) \rightarrow L^2(\mathbb{T})$,

$$(Ff)(x) = \sum_{n=-\infty}^{\infty} f(n) e^{-2\pi i n x}$$

for all $f \in \ell^1(\mathbb{Z})$. Inspired by the definition of the continuous windowed Fourier transform (9), and the relation (10), for any fixed $g \in \ell^2(\mathbb{Z})$, we define the windowed Fourier transform of $f \in \ell^2(\mathbb{Z})$ with respect to g as:

$$\begin{aligned} (W_g f)(n, x) &:= \langle f, T^n M^x g \rangle \\ &= \langle T^{-n} f, M^x g \rangle \\ &= \sum_{m=-\infty}^{\infty} f(m+n) \overline{g(m)} e^{-2\pi i x m}. \end{aligned} \tag{11}$$

In terms of this notation, (10) becomes: $(\mathcal{W}_g f)(\sigma m, x) \approx \sigma(W_g f)(m, \sigma x)$. Having discussed how the analog windowed Fourier transform may be discretized, we now turn to the fundamental tool of time-frequency analysis: the spectrogram.

3.5 Spectrograms

The *spectrogram* of $f \in L^2(\mathbb{R})$ with respect to some window $g \in L^2(\mathbb{R})$ is the square of the modulus of the windowed Fourier transform of f with respect to g , that is $|(\mathcal{W}_g f)(t, x)|^2$. The spectrogram of a signal shows how the frequency content of that signal varies over time. Being a function of two variables, the spectrogram may be viewed as a surface, as in Figure 8. More commonly, the spectrogram is depicted as an image, where the magnitude of $|(\mathcal{W}_g f)(t, x)|^2$ is indicated by the color of the pixel, as in Figure 7(c).

When viewing the spectrogram as a surface, one sees peaks and valleys. Peaks correspond to large values of $|(\mathcal{W}_g f)(t, x)|^2$, that is, the time t at which f has an especially large component which is oscillating at frequency x . Indeed, for any given time t , one may find the particular x 's for which the value of $|(\mathcal{W}_g f)(t, x)|^2$ is maximized, that is, the frequency that is most present in f at time t . Writing these dominant frequencies as a function of t , one produces a curve known as the *ridgeline* of the spectrogram, as it corresponds to the crests of the hills of the surface $|(\mathcal{W}_g f)(t, x)|^2$. To be precise, the *ridgeline* of the spectrogram $|(\mathcal{W}_g f)(t, x)|^2$ is the function $\mathcal{R}_g f : \mathbb{R} \rightarrow \mathbb{R}$,




Figure 8: Three-dimensional view of Figure 7(c).

where

$$(\mathcal{R}_g f)(t) := \operatorname{argmax}_{x \in \mathbb{R}} |(\mathcal{W}_g f)(t, x)|^2. \quad (12)$$

Though a standard tool of time-frequency analysis, the ridgeline (12) of a spectrogram still leaves much to be desired. For instance, for signals which contain multiple frequency components at any given time, the ridgeline may be discontinuous. In particular, the ridgeline will have a jump discontinuity whenever one frequency peak rises above another. Such jump discontinuities are especially troubling for our heart and breath rate application, in which the ridgeline (12) serves as an estimate of the velocity of the subject's body. In particular, to detect the periodic motions of the heart and lungs, one needs to subject the velocity signal itself to a time-frequency analysis. Here, any jump discontinuities will manifest themselves as noise which could possibly drown out the relatively weak heartbeat component. In short, when using the ridgeline as our velocity signal, experimentation has revealed the heartbeat component of the signal to have a dangerously small signal-to-noise ratio. Because of this, we aban-

doned the use of the ridgeline as a velocity estimate, and instead began to use *spectral centroids*. In particular, for any fixed time $t \in \mathbb{R}$, we may regard the corresponding vertical strip of the spectrogram $|(\mathcal{W}_g f)(t, x)|^2$ as a distribution over the real line. Indeed, considering the definition (9) of the windowed Fourier transform, the values $\{|(\mathcal{W}_g f)(t, x)|^2\}_{x \in \mathbb{R}}$ give the spectral energy density of the function $f(s+t)g(s)$. The centroids (centers of mass) of these distributions may themselves be viewed as functions of time. To be precise, the *centroid* of the spectrogram $|(\mathcal{W}_g f)(t, x)|^2$ is the function $\mathcal{C}_g f : \mathbb{R} \rightarrow \mathbb{R}$, where

$$(\mathcal{C}_g f)(t) := \frac{\int_{-\infty}^{\infty} x |(\mathcal{W}_g f)(t, x)|^2 dx}{\int_{-\infty}^{\infty} |(\mathcal{W}_g f)(t, x)|^2 dx}. \quad (13)$$

As noted in the previous section, the actual data we receive is discrete-time, that is, digital samples of the analog radar signal. This poses no serious problem for the ridgeline (12) and centroid (13) definitions introduced above for analog signals. Indeed, as the windowed Fourier transform for discrete signals was generalized to the discrete case in (11), we define the *discrete spectrogram* of $f \in \ell^2(\mathbb{Z})$ as $|(W_g f)(n, x_j)|^2$, where $n \in \mathbb{Z}$, and $\{x_j\}_{j=1}^J$ is some fixed finite set of real numbers. In this setting, the *ridgeline* is then

$$(\mathcal{R}_g f)(n) := \operatorname{argmax}_{j=1, \dots, J} |(W_g f)(n, x_j)|^2 \quad (14)$$

and the *centroid* is then

$$(\mathcal{C}_g f)(n) := \frac{\sum_{j=1}^J x_j |(W_g f)(n, x_j)|^2}{\sum_{j=1}^J |(W_g f)(n, x_j)|^2}. \quad (15)$$

In particular, using the fact that for any differentiable function $z : \mathbb{R} \rightarrow \mathbb{C}$, we have:

$$\frac{d}{dt}|z(t)|^2 = z(t)\overline{z'(t)} + z'(t)\overline{z(t)} = 2\operatorname{Re}[z(t)\overline{z'(t)}],$$

we can evaluate the derivative of the numerator of (13) with respect to t as

$$\begin{aligned} \frac{d}{dt} \int_{-\infty}^{\infty} x |(\mathcal{W}_g f)(t, x)|^2 dx &= \int_{-\infty}^{\infty} x \left[\frac{d}{dt} |(\mathcal{W}_g f)(t, x)|^2 \right] dx \\ &= \int_{-\infty}^{\infty} 2x \operatorname{Re} \left[(\mathcal{W}_g f)(t, x) \overline{\frac{d}{dt} (\mathcal{W}_g f)(t, x)} \right] dx, \end{aligned} \quad (16)$$

where the remaining derivative of $(\mathcal{W}_g f)(t, x)$ in (16), is given by

$$\begin{aligned} \frac{d}{dt} (\mathcal{W}_g f)(t, x) &= \frac{d}{dt} \int_{-\infty}^{\infty} f(s+t)g(s)e^{-2\pi i x s} ds \\ &= \frac{d}{dt} \int_{-\infty}^{\infty} f(s)g(s-t)e^{-2\pi i x (s-t)} ds \\ &= \int_{-\infty}^{\infty} f(s)e^{-2\pi i x s} \left[\frac{d}{dt} g(s-t)e^{2\pi i x t} \right] ds. \end{aligned}$$

Thus, we expect the numerator of (13) to be differentiable as long as $g(s-t)e^{2\pi i x t}$ is, and as $e^{2\pi i x t}$ is infinitely differentiable, the smoothness of $(\mathcal{W}_g f)(t, x)$ is truly dependent upon the smoothness of g .


A similar argument shows that the denominator of (13) is also differentiable, and so (13) itself will be differentiable whenever its denominator is nonzero, which, according to the Parseval-Plancherel identity will happen whenever $f(s+t)\overline{g(s)}$ is not identically zero. We note that being only an informal derivation, we shall not formally justify the interchanging of the derivative and the integral above in (16).

To summarize, though the ridgeline (12) is not continuous, in general, one should expect the centroid (13) to be differentiable, provided g satisfies mild smoothness as-


sumptions. We note that even though the above analysis only applies to the analog ridgeline (12) and centroid (13), we nevertheless expect a similar result to hold in the discrete case for the ridgeline (14) and centroid (15), provided one first appropriately defines what is meant by the smoothness of a discrete function. Thus “discrete smoothness” seems to be highly dependent on the vertical resolution of the spectrogram, that is, at how many frequencies $\{x_j\}_{j=1}^J$ the spectrogram is computed.

Indeed, consider the spectrogram in Figure 9(a), which was the result of using a windowed Fourier transform on a 60 second sample of 120 GHz CW radar data. Both the ridgeline and the spectral centroid, seen in Figures 9(b) and (d), respectively, were easily computed from Figure 9(a). As the radar signal had only one dominant frequency at any given time, these two curves are nearly identical, a fact which may be further confirmed by comparing their spectrograms, as seen in Figures 9(c) and (e). Meanwhile, when the vertical resolution of the spectrogram is decreased, as in Figure 10(a), the differences in smoothness between the ridgeline and the centroid become very apparent. Indeed, as the discrete ridgeline (14) produces an index $j = 1, \dots, J$, it may appear heavily overquantized when J is small, as seen in Figure 10(b). Meanwhile, the centroid curve derived from the same data, is still relatively smooth, as seen in Figure 10(d). This is because the centroid is a weighted average of the x_j ’s, rather than just picking whichever one is largest. As noted above, with respect to our particular application of heart and breath rate detection, this difference in smoothness is critical. In particular, the spectrogram of the centroid curve (Figure 10(e)) exhibits well-defined horizontal strips around 0.3 Hz and 1 Hz, indicating breath and heart rates, respectively. Meanwhile, the spectrogram of the ridgeline signal (Figure 10(c)) is much more noisy, making it difficult to accurately determine the heart and breath rates.


To summarize, the discrete ridgelines (14) and discrete centroids (15) seem comparable when the frequency resolution is high, that is, when the x_j ’s are closely spaced. Meanwhile, the centroid seems to dramatically outperform the ridgeline when the x_j ’s are coarse. Importantly, as the length of time needed to compute the discrete spectro-




(a) The spectrogram of a 120 GHz radar signal that was sampled at 1 kHz. This spectrogram was computed over frequencies with high resolution.




(b) The ridgeline of (a) over a 20 second interval, which measures the Doppler shift at a given time. Since (a) was computed over high resolution frequencies, we get a smooth ridgeline.



(c) The spectrogram of the ridgeline. Since (a) was computed over frequencies with high resolution, we get a smooth signal in (b) that gives a spectrogram from which we can extract heart and breath rate information. The horizontal strip around 0.3 Hz represents the breath rate of the subject while the faint horizontal strip around 1 Hz represents the heart rate.




(d) The spectral centroids of (a) over a 20 second interval, which gives us another way to measure the Doppler shift at a given time. We see that since the frequency resolution is high in (a), this signal is similar to that of the ridgeline (b). We get a smooth, well-behaved signal by taking the centroids of the vertical strips of (a).




(e) The spectrogram of (d), gives us heart and breath rate information. This spectrogram is similar to (c) since the spectrogram of the radar signal was computed over high resolution frequencies, giving us a smooth ridgeline signal (b) that is very similar to the centroid signal (d). Again, the horizontal strip around 0.3 Hz represents the breath rate of the subject while the faint horizontal strip around 1 Hz represents the heart rate.


Figure 9: Comparing the ridgeline and centroid of a spectrogram that was computed over frequencies with high resolution.




(a) The spectrogram of a 120 GHz radar signal that was sampled at 1 kHz. This spectrogram was computed over frequencies with coarse resolution.




(b) The ridgeline of (a) over a 20 second interval, which measures the Doppler shift at a given time. Since (a) was computed over coarse resolutions for frequency, our ridgeline is overquantized.



(c) The spectrogram of the ridgeline, which should provide information about the heart and breath rates of the subject. However, since (b) is overquantized, we get an ambiguous result.



(d) The spectral centroids of (a) over a 20 second interval, which gives us another way to measure the Doppler shift at a given time. Even though (a) is very coarse, we get a smooth, well-behaved signal by taking the centroids of the vertical strips of (a).



(e) The spectrogram of (d), which is much clearer than (c), gives us heart and breath rate information. The horizontal strip around 0.3 Hz represents the breath rate of the subject and the horizontal strip around 1 Hz represents the heart rate.

Figure 10: Comparing the ridgeline and centroid of a spectrogram that was computed over frequencies with coarse resolution.

gram $|(W_g f)(n, x_j)|^2$ grows linearly with J , this smoothness advantage translates into a speed advantage. In the main results of this research, presented in the next chapter, we show that this speed advantage may be made even more pronounced. In particular, whereas the computation of the ridgeline necessitates that one first compute the spectrogram $|(W_g f)(n, x_j)|^2$ for all $n \in \mathbb{Z}$ and $j = 1, \dots, J$, in the next section, we show how the discrete centroid (15) may be computed without ever needing to compute $|(W_g f)(n, x_j)|^2$ explicitly. That is, though both Figures 10(b) and (d) may be computed from Figure 10(a), it turns out that Figure 10(b) must be computed this way, whereas a computational trick will allow us to compute Figure 10(d) directly using FFT's. In short, for the purposes of heart and breath rate detection, the centroid is much cleaner than the ridgeline, and, as shown in the next chapter, may actually be computed faster than the ridgeline, that is, we show that a better signal can be obtained at a lower cost.

We conclude this section by performing a time-frequency analysis on our simulated data generated using the data model discussed in Section 3.1. In so doing, we get some preliminary idea as to what our data should look like. Figure 11(a) shows $x(t)$ and Figure 11(b) shows the simulated data $S_{MR}(t)$ for a 32 second sample. The spectrogram of the simulated data is shown in Figure 11(c). This spectrogram shows us what frequencies are present at what time and is proportional to the velocity of the target. Figure 11(d) shows the ridgeline of Figure 11(c), which represents the dominant frequency at any time. The ridgeline is a quantity which estimates the Doppler shift in the radar signal at any given time and should be a close approximation of the velocity signal. Figure 11(e) shows the centroids of the spectrogram of the radar signal. The centroids give us another way to measure the Doppler shift at a given time and are smooth and well-behaved. These centroids can be computed from Figure 11(c) or directly from Figure 11(b) using the ideas presented in the next chapter. Figure 11(f) shows the spectrogram of Figure 11(e), which indicates a breath rate of 0.23 Hz, and a heart rate of 1 Hz, where the harmonic at 2 Hz is a result of the peaks

in the heartbeat signal. That is, we obtain the numbers ω_1 and ω_2 in Table 1 that were used to generate the data in the first place.

We conclude this chapter with a discussion of topics of discrete Fourier analysis, such as the FFT and circular convolutions. Though not directly applicable to the above ideas, these concepts shall become important in the algorithmic implementation of our main results in the following chapter.

3.6 Time-Frequency Analysis of Periodic Digital Signals

For any positive integer N , the set of integer-indexed, N -periodic, complex-valued sequences is

$$\ell(\mathbb{Z}_N) := \{f : \mathbb{Z} \rightarrow \mathbb{C} \mid f(n + N) = f(n) \ \forall n \in \mathbb{Z}\},$$

which is a Hilbert space under the inner product


$$\langle f, g \rangle = \sum_{n \in \mathbb{Z}_N} f(n) \overline{g(n)}.$$

Here, the summation over $\mathbb{Z}_N$ denotes a sum of the coset representatives of $\mathbb{Z}$ over $N\mathbb{Z}$, that is, any set of indices congruent to $\{0, \dots, N-1\}$ modulo N . The Fourier transform is $F : \ell(\mathbb{Z}_N) \rightarrow \ell(\mathbb{Z}_N)$:


$$(Ff)(m) = \sum_{n \in \mathbb{Z}_N} f(n) e^{-2\pi i m n / N}. \quad (17)$$

This discrete Fourier transform is often referred to as the Fast Fourier Transform (FFT), as it may be implemented in a suprisingly fast manner. In particular, whereas a direct implementation of (17) requires $\mathcal{O}(N^2)$ operations, the FFT algorithm, as we now discuss, permits (17) to be evaluated in $\mathcal{O}(N \log N)$ operations.


3.6.1 The Fast Fourier Transform. Any integer can be written as a product of its primes. In our computations, we shall always use integers of the form $N = 2^k$




(a) The distance $x(t)$ (in meters) of the target from the transmitter. The motion of the target is a result of heart and breath motion.




(b) The real part of the simulated data $S_{MR}(t)$.




(c) The spectrogram of (b), which shows us what frequencies are present at what times. These frequencies are proportional to the velocity of the target.



(d) The ridgeline of (c), which represents the dominant frequency at any time, is a quantity which estimates the Doppler shift in the radar signal at any given time. The ridgeline should be a close approximation of the velocity signal.



(e) The spectral centroids of (b), which gives us another way to measure the Doppler shift at a given time. We get a smooth, well-behaved signal by taking the centroids of the vertical strips of (c). In fact, using a computational trick developed in Chapter IV, we can compute the spectral centroids without computing (c) explicitly.



(f) The spectrogram of (e), gives us heart and breath rate information. The horizontal strip around 0.25 Hz represents the breath rate of the subject and the horizontal strip around 1 Hz represents the heart rate. From our data simulation, we expect the breath rate to be 0.23 breaths per second and the heart rate to be 1 beat per second.

Figure 11: Time-frequency analysis of the simulated data and its centroids

for some $k \in \mathbb{N}$. Letting $f \in \ell(\mathbb{Z}_N)$, the Fourier transform of f is

$$\begin{aligned}
(Ff)(m) &= \sum_{n \in \mathbb{Z}_N} f(n) e^{-\pi i m n / N} \\
&= \sum_{n=0}^{N-1} f(n) e^{-2\pi i m n / N} \\
&= \sum_{p=0}^{\frac{N}{2}-1} f(2p) [e^{-2\pi i m 2p / N} + f(2p+1) e^{-2\pi i m (2p+1) / N}] \\
&= \sum_{p=0}^{\frac{N}{2}-1} f(2p) e^{-2\pi i m p / (N/2)} + e^{-2\pi i m / N} \sum_{p=0}^{\frac{N}{2}-1} f(2p+1) e^{-2\pi i m p / (N/2)} \\
&= (Ff_0)(m) + e^{-2\pi i m / N} (Ff_1)(m),
\end{aligned} \tag{18}$$

where $f_0, f_1 \in \ell(\mathbb{Z}_{N/2})$ are the even and odd components of f , respectively, defined as $f_0(p) := f(2p)$, $f_1(p) := f(2p+1)$. We see that a discrete Fourier transform of size N may be written as two Fourier transforms of size $N/2$ along with N additional multiplications. Letting $\text{op}(N)$ denote the number of operations needed to evaluate a discrete Fourier transform of size N , (18) gives that $\text{op}(N) \leq N + 2\text{op}(N/2)$. This then implies that $\text{op}(2^k) \leq k2^k$. Indeed, by induction, we know this is true for $k = 1$, and assuming it is true for a given k , then

$$\text{op}(2^{k+1}) \leq 2^{k+1} + 2\text{op}(2^k) \leq 2^{k+1} + 2k2^k = (k+1)2^{k+1}.$$

Put another way, for $N = 2^k$, we have $\text{op}(N) \leq N \log_2 N$. We now turn to the main application of the FFT that we shall employ, namely its ability to permit a very fast evaluation of circular convolutions.

3.6.2 Convolutions. The *circular convolution* of $f, g \in \ell(\mathbb{Z}_N)$ is $f * g \in \ell(\mathbb{Z}_N)$:

$$(f * g)(m) := \sum_{n \in \mathbb{Z}_N} f(m-n)g(n). \tag{19}$$

Under the action of the discrete Fourier transform, a convolution becomes a pointwise product:

$$\begin{aligned}
[\mathcal{F}(f * g)](p) &= \sum_{m \in \mathbb{Z}_N} (f * g)(m) e^{-2\pi i m p / N} \\
&= \sum_{m \in \mathbb{Z}_N} \sum_{n \in \mathbb{Z}_N} f(m - n) g(n) e^{-2\pi i m p / N} \\
&= \sum_{m \in \mathbb{Z}_N} f(m - n) e^{-2\pi i (m - n) p / N} \sum_{n \in \mathbb{Z}_N} g(n) e^{-2\pi i n p / N} \\
&= (\mathcal{F}f)(p) (\mathcal{F}g)(p),
\end{aligned}$$

that is,

$$(f * g) = \mathcal{F}^{-1}[\mathcal{F}(f)\mathcal{F}(g)]. \quad (20)$$

Calculating the convolution directly using (19) takes N^2 operations, whereas calculating the convolution using (20) takes $N \log N$ operations. That is, by calculating our convolutions using FFT's, we can save time, which is especially important when N is large.

3.6.3 Computing Discrete Spectrograms. In the previous sections, we have defined windowed Fourier transforms on both $L^2(\mathbb{R})$ and $\ell^2(\mathbb{Z})$. In fact, one may also perform a similar analysis on $\ell(\mathbb{Z}_N)$. In particular, defining translation and modulation operators as $T^n, M^m : \ell(\mathbb{Z}_N) \rightarrow \ell(\mathbb{Z}_N)$:

$$\begin{aligned}
(T^n f)(m) &:= f(m - n), \\
(M^m f)(n) &:= f(n) e^{2\pi i m n / N},
\end{aligned}$$

we define the discrete windowed Fourier transform as

$$\begin{aligned}
(W_g f)(n, m) &:= \langle f, T^n M^m g \rangle \\
&= \sum_{l \in \mathbb{Z}_N} f(l + n) \overline{g(l)} e^{-2\pi i l m / N}.
\end{aligned}$$

We now show how using FFT's we can compute the spectrogram in two different ways. In the first method, we can compute the spectrogram for all frequencies and some given times. In particular, letting $f, g \in \ell(\mathbb{Z}_N)$ and $m, n \in \mathbb{Z}_N$, we have:

$$\begin{aligned}
(W_g f)(m, n) &= \langle f, T^m M^n g \rangle \\
&= \langle T^{-m} f, M^n g \rangle \\
&= \sum_{p \in \mathbb{Z}_N} (T^{-m} f)(p) \overline{(M^n g)(p)} \\
&= \sum_{p \in \mathbb{Z}_N} f(p+m) e^{-2\pi i n p / N} \overline{g(p)} \\
&= \sum_{p \in \mathbb{Z}_N} [f(p+m) \overline{g(p)}] e^{-2\pi i n p / N} \\
&= [F((T^{-m} f) \bar{g})](n).
\end{aligned}$$

That is, for any fixed m , we may compute $(W_g f)(m, n)$ for all $n \in \mathbb{Z}_N$ as the FFT of $(T^{-m} f) \bar{g}$. This approach works well if we wanted to compute the spectrogram over all frequencies for just a few times. However, as our data tends to exist over long periods of time, but is only nonzero for a few frequencies, this is not our best choice. Instead, we may use a second method:

$$\begin{aligned}
(W_g f)(m, n) &= \langle f, T^m M^n g \rangle \\
&= \sum_{p \in \mathbb{Z}_N} f(p) \overline{(T^m M^n g)(p)} \\
&= \sum_{p \in \mathbb{Z}_N} f(p) \overline{(M^n g)(p-m)} \\
&= \sum_{p \in \mathbb{Z}_N} f(p) \widetilde{(M^n g)(m-p)} \\
&= [f * \widetilde{(M^n g)}](m), \tag{21}
\end{aligned}$$

where $\tilde{f}(n) := \overline{f(-n)}$ is the *involution* of f . Unfortunately, (21) is itself not a great method, as it involves knowing f at all times. In particular, (21) is not a real-time

method. Thus, both of these “immediate” applications of the FFT to our time-frequency analysis problem are inappropriate.

In Chapter IV, we introduce a new application of the FFT to time-frequency analysis, one which suits our heart and breath rate application well. In particular, we show how the spectral centroids, discussed above in Section 3.5, may be computed directly using Toeplitz matrices. Multiplication by these matrices, in turn, may be expressed in terms of circular convolutions, which, as noted above, may be quickly implemented using FFT’s.

IV. A New Method for Computing Spectral Centroids

In the previous chapter, we discussed the spectrogram, which is a fundamental tool of time-frequency analysis. We also discussed that, for the purpose of heart and breath rate detection, taking spectral centroids of the radar signal's spectrogram yields a smoother estimate of the velocity signal than that provided by the spectrogram's ridgeline. In this section, we prove that spectral centroids have an additional important advantage over ridgelines: speed. In particular, we will show that spectral centroids of a spectrogram can be computed without ever needing to explicitly compute the spectrogram itself. In fact, we will show that spectral centroids may be directly computed from a signal using an FFT-implemented, Toeplitz matrix-based algorithm.

4.1 Weighted Spectral Sums

Theorems 4.1.1 and 4.1.3 are the most fundamental original results in this thesis. Theorem 4.1.1 shows how weighted sums of the vertical strips of a discrete spectrogram may be computed as a Toeplitz matrix multiplication. A classical result, summarized in Lemma 4.1.2, shows how such a multiplication may be written as a circular convolution, which may then be evaluated using FFT's. The implications of these two results are then explicitly determined in Theorem 4.1.3, which provides the foundations for a fast algorithm for the computation of spectral centroids, as discussed in Section 4.2.

Theorem 4.1.1. *For any weighting function $w : \mathbb{R} \rightarrow \mathbb{R}$, any fixed frequencies $\{x_j\}_{j=1}^J \subset \mathbb{R}$, any window $g \in \ell^2(\mathbb{Z})$ with $g(n) = 0$ whenever $|n| > N$, and any $f \in \ell^2(\mathbb{Z})$, we have:*

$$\sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) = z^* A z, \quad (22)$$

where $z \in \mathbb{C}^{2N+1}$ has entries:

$$z_k = f(k + n - N - 1) \overline{g(k - N - 1)},$$

and A is the $(2N+1) \times (2N+1)$ Toeplitz matrix whose (k, k') th entry is:

$$A_{k,k'} = \sum_{j=1}^J w(x_j) e^{2\pi i x_j (k-k')}.$$

Proof. By the definition (11) of the discrete windowed Fourier transform,

$$\begin{aligned} |(W_g f)(n, x_j)|^2 &= \sum_{m=-\infty}^{\infty} f(m+n) \overline{g(m)} e^{-2\pi i x_j m} \overline{\sum_{m'=-\infty}^{\infty} f(m'+n) \overline{g(m')} e^{-2\pi i x_j m'}} \\ &= \sum_{m=-\infty}^{\infty} \sum_{m'=-\infty}^{\infty} f(m+n) \overline{g(m)} e^{2\pi i x_j (m'-m)} \overline{f(m'+n) g(m')}. \end{aligned}$$

Since $g(n) = 0$ for all $n \in \mathbb{Z}$ such that $|n| > N$, then

$$|(W_g f)(n, x_j)|^2 = \sum_{m=-N}^N \sum_{m'=-N}^N f(m+n) \overline{g(m)} e^{2\pi i x_j (m'-m)} \overline{f(m'+n) g(m')}. \quad (23)$$

Using (23), we have:

$$\begin{aligned} \sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) & \quad (24) \\ &= \sum_{m=-N}^N \sum_{m'=-N}^N f(m+n) \overline{g(m)} \left(\sum_{j=1}^J w(x_j) e^{2\pi i x_j (m'-m)} \right) \overline{f(m'+n) g(m')} \\ &= \sum_{m'=-N}^N \overline{f(m'+n) g(m')} \sum_{m=-N}^N \left(\sum_{j=1}^J w(x_j) e^{2\pi i x_j (m'-m)} \right) f(m+n) \overline{g(m)}. \quad (25) \end{aligned}$$

Letting $z \in \mathbb{C}^{2N+1}$ be taken as it is in the statement of the result, note that we have

$$z_{m+N+1} = f(m+N+1+n-N-1) \overline{g(m+N+1-N-1)} = f(m+n) \overline{g(m)}$$

whenever $m + N + 1 = 1, \dots, 2N + 1$, that is, whenever $m = -N, \dots, N$. Thus, (25) may be written in terms of z as:

$$\sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) = \sum_{m'=-N}^N \overline{z_{m'+N+1}} \sum_{m=-N}^N \left(\sum_{j=1}^J w(x_j) e^{2\pi i x_j (m'-m)} \right) z_{m+N+1}. \quad (26)$$

Making the changes of variables $k = m + N + 1$, $k' = m' + N + 1$ for $m, m' = -N, \dots, N$, and recalling the definition of A given in the statement of the result, (26) becomes

$$\begin{aligned} \sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) &= \sum_{k'=1}^{2N+1} \overline{z_{k'}} \left[\sum_{k=1}^{2N+1} \left(\sum_{j=1}^J w(x_j) e^{2\pi i x_j ((k'-N-1)-(k-N-1))} \right) z_k \right] \\ &= \sum_{k'=1}^{2N+1} \overline{z_{k'}} \left[\sum_{k=1}^{2N+1} \left(\sum_{j=1}^J w(x_j) e^{2\pi i x_j (k'-k)} \right) z_k \right] \\ &= \sum_{k'=1}^{2N+1} \overline{z_{k'}} \left(\sum_{k=1}^{2N+1} A_{k',k} z_k \right) \\ &= \sum_{k'=1}^{2N+1} \overline{z_{k'}} (Az)_{k'} \\ &= z^* Az. \end{aligned}$$

To show the matrix A is Toeplitz, we must show that A is constant along diagonals, that is, $A_{k+1,k'+1} = A_{k,k'}$ for all $k, k' = 1, \dots, 2N$. From (25),

$$\begin{aligned} A_{k+1,k'+1} &= \sum_{j=1}^J w(x_j) e^{2\pi i x_j (k+1-(k'+1))} \\ &= \sum_{j=1}^J w(x_j) e^{2\pi i x_j (k-k')} \\ &= A_{k,k'}. \end{aligned} \quad \square$$

Theorem 4.1.1 shows how weighted sums of vertical strips of a discrete spectrogram may be written as a quadratic form involving a Toeplitz matrix. We now present a classical result which shows how Toeplitz matrix multiplication may be written in

terms of circular convolutions, which, as described in Section 3.6, have an FFT-based implementation.

Lemma 4.1.2. *If $A \in \mathbb{C}^{K \times K}$ is Toeplitz, then for any non negative integer P ,*

$$(Az)_k = (\psi_{zp} * \varphi_{zp})(k-1)$$

for all $z \in \mathbb{C}^K$ and all $k = 1, \dots, K$, where $\psi_{zp}, \varphi_{zp} \in \ell(\mathbb{Z}_{2K-1+P})$,

$$\begin{aligned} \psi_{zp}(k) &= \begin{cases} \psi(k), & k = -K+1, \dots, K-1, \\ 0, & \text{else,} \end{cases} \\ \varphi_{zp}(k) &= \begin{cases} \varphi(k), & k = 0, \dots, K-1, \\ 0, & \text{else,} \end{cases} \end{aligned}$$

where ψ and φ are given by:

$$A_{k,k'} = \psi(k-k'), \quad z_k = \varphi(k-1),$$

for all $k, k' = 1, \dots, K$.

Proof. For any $k = 1, \dots, K$,

$$(\psi_{zp} * \varphi_{zp})(k-1) = \sum_{k' \in \mathbb{Z}_{2K-1+P}} \psi_{zp}(k-1-k')\varphi_{zp}(k') = \sum_{k'=0}^{K-1} \psi_{zp}(k-1-k')\varphi(k'). \quad (27)$$

Since $k = 1, \dots, K$ then for any $k' = 0, \dots, K-1$ we have:

$$-K+1 \leq k-1-k' \leq K-1,$$

and so we may continue to simplify (27) as:

$$\begin{aligned}
(\psi_{\text{zp}} * \varphi_{\text{zp}})(k-1) &= \sum_{k'=0}^{K-1} \psi(k-1-k')\varphi(k') \\
&= \sum_{k'=1}^K \psi(k-k')\varphi(k'-1) \\
&= \sum_{k'=1}^K A_{k,k'} z_{k'} \\
&= (Az)_k.
\end{aligned}
\tag*{$\square$}$$

Having demonstrated the straightforwardness of this classical result, we now combine it with our original result (Theorem 4.1.1) to produce a new result which forms the foundation of our spectral centroid-computing algorithm given below in Section 4.2:

Theorem 4.1.3. *For any weighting function $w : \mathbb{R} \rightarrow \mathbb{R}$, any fixed frequencies $\{x_j\}_{j=1}^J \subset \mathbb{R}$, any window $g \in \ell^2(\mathbb{Z})$ with $g(n) = 0$ whenever $|n| > N$, and any $f \in \ell^2(\mathbb{Z})$, we have:*

$$\sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) = \frac{1}{4N+4} \sum_{m \in \mathbb{Z}_{4N+4}} (\text{F}\psi_{\text{zp}})(m) |(\text{F}\varphi_{\text{zp}})(m)|^2,$$

where the discrete Fourier transform of $\varphi_{\text{zp}} \in \ell(\mathbb{Z}_{4N+4})$,

$$\varphi_{\text{zp}}(k) := \begin{cases} f(k+n-N) \overline{g(k-N)}, & k = 0, \dots, 2N, \\ 0, & \text{else,} \end{cases}$$

may be computed numerically using FFTs while the discrete Fourier transform of $\psi_{\text{zp}} \in \ell(\mathbb{Z}_{4N+4})$ may be computed symbolically as:

$$(\text{F}\psi_{\text{zp}})(m) = \sum_{j=1}^J w(x_j) D_{2N} \left(2\pi \left(x_j - \frac{m}{4N+4} \right) \right),$$

where $D_{2N} : \mathbb{R} \rightarrow \mathbb{R}$ is the Dirichlet kernel of order $2N$, namely:

$$D_{2N}(x) := \sum_{k=-2N}^{2N} e^{ikx} = \frac{\sin((2N + \frac{1}{2})x)}{\sin(x/2)}. \quad (28)$$

Proof. Letting A and z be defined as in the statement of Theorem 4.1.1, we apply Lemma 4.1.2 to evaluate the matrix-vector product Az . Specifically, letting $K = 2N + 1$ and fixing $P = 3$, we have $\varphi_{zp} \in \ell(\mathbb{Z}_{4N+4})$ is:

$$\begin{aligned} \varphi_{zp}(k) &= \begin{cases} \varphi(k), & k = 0, \dots, (2N + 1) - 1, \\ 0, & \text{else,} \end{cases} \\ &= \begin{cases} z_{k+1}, & k = 0, \dots, 2N, \\ 0, & \text{else,} \end{cases} \\ &= \begin{cases} f(k + n - N) \overline{g(k - N)}, & k = 0, \dots, 2N, \\ 0, & \text{else,} \end{cases} \end{aligned}$$

as claimed. Meanwhile, as the ψ in Lemma 4.1.2 is defined by the relation $A_{k,k'} = \psi(k - k')$, the definition of A in the statement of Theorem 4.1.1 implies that:

$$\begin{aligned} \psi_{zp}(k) &= \begin{cases} \psi(k), & k = -(2N + 1) + 1, \dots, (2N + 1) - 1, \\ 0, & \text{else,} \end{cases} \\ &= \begin{cases} \sum_{j=1}^J w(x_j) e^{2\pi i x_j k}, & k = -2N, \dots, 2N, \\ 0, & \text{else.} \end{cases} \end{aligned}$$

Under these particular definitions of ψ_{zp} and φ_{zp} , Lemma 4.1.2 now tells us that the matrix-vector product Az of Theorem 4.1.1 may be written as a circular convolution.

In particular, we have:

$$\begin{aligned}
\sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) &= z^* A z \\
&= \sum_{k=1}^{2N+1} \overline{z_k} (A z)_k \\
&= \sum_{k=1}^{2N+1} \overline{\varphi(k-1)} (\psi_{\text{zp}} * \varphi_{\text{zp}})(k-1) \\
&= \sum_{k=0}^{2N} \overline{\varphi(k)} (\psi_{\text{zp}} * \varphi_{\text{zp}})(k) \\
&= \sum_{k \in \mathbb{Z}_{4N+4}} \overline{\varphi_{\text{zp}}(k)} (\psi_{\text{zp}} * \varphi_{\text{zp}})(k) \\
&= \langle \psi_{\text{zp}} * \varphi_{\text{zp}}, \varphi_{\text{zp}} \rangle.
\end{aligned}$$

To continue, we use the Parseval-Plancherel Identity for the discrete Fourier transform, namely that $F : \ell(\mathbb{Z}_N) \rightarrow \ell(\mathbb{Z}_N)$ satisfies $F^* F = NI$, and so for any $f, g \in \ell(\mathbb{Z}_N)$ we have $\langle Ff, Fg \rangle = \langle F^* Ff, g \rangle = N \langle f, g \rangle$. In particular, applying this identity to $\psi_{\text{zp}} * \varphi_{\text{zp}}, \varphi_{\text{zp}} \in \ell(\mathbb{Z}_{4N+4})$, and using the fact that the discrete Fourier transform distributes multiplicatively over convolutions (20) gives:

$$\begin{aligned}
\sum_{j=1}^J |(W_g f)(n, x_j)|^2 w(x_j) &= \frac{1}{4N+4} \langle F(\psi_{\text{zp}} * \varphi_{\text{zp}}), F\varphi_{\text{zp}} \rangle \\
&= \frac{1}{4N+4} \langle (F\psi_{\text{zp}})(F\varphi_{\text{zp}}), F\varphi_{\text{zp}} \rangle \\
&= \frac{1}{4N+4} \sum_{m \in \mathbb{Z}_{4N+4}} (F\psi_{\text{zp}})(m) (F\varphi_{\text{zp}})(m) \overline{(F\varphi_{\text{zp}})(m)} \\
&= \frac{1}{4N+4} \sum_{m \in \mathbb{Z}_{4N+4}} (F\psi_{\text{zp}})(m) |(F\varphi_{\text{zp}})(m)|^2,
\end{aligned}$$

as claimed. To conclude, we need only note that as ψ_{zp} depends solely on the fixed frequencies $\{x_j\}_{j=1}^J$ and the weighting function w , its discrete Fourier transform may

be computed symbolically:

$$\begin{aligned}
(F\psi_{\text{zp}})(m) &= \sum_{k \in \mathbb{Z}_{4N+4}} \psi_{\text{zp}}(k) e^{-2\pi i k m / (4N+4)} \\
&= \sum_{k=-2N}^{2N} \sum_{j=1}^J w(x_j) e^{2\pi i x_j k} e^{-2\pi i k m / (4N+4)} \\
&= \sum_{j=1}^J w(x_j) \sum_{k=-2N}^{2N} e^{2\pi i k (x_j - m / (4N+4))} \\
&= \sum_{j=1}^J w(x_j) D_{2N} \left(2\pi \left(x_j - \frac{m}{4N+4} \right) \right). \quad \square
\end{aligned}$$

Having proven the main results of the thesis, we now show how they may be implemented to the application of detecting heart and breath rates from terahertz radar data. Specifically, we now consider how Theorem 4.1.3 may be used to produce an algorithm which efficiently computes the centroids of vertical strips of a signal's spectrogram.

4.2 Spectral Moments

As discussed above in Section 3.5, the time-frequency estimate provided by the centroid of the spectrogram, namely:

$$(C_g f)(n) := \frac{\sum_{j=1}^J x_j |(W_g f)(n, x_j)|^2}{\sum_{j=1}^J |(W_g f)(n, x_j)|^2}, \quad (29)$$

is, in general, smoother than that provided by the spectrogram's ridgeline (14). This difference in smoothness is particularly noticeable when the spectrogram is only computed over a coarse set of frequencies $\{x_j\}_{j=1}^J$. We now discuss, a second, crucial advantage of centroids over ridgelines, namely how by Theorem 4.1.3, one may compute the centroid curve (29) without ever needing to explicitly compute the values




Figure 12: The discrete Fourier transforms of ψ_{zp} for $w(x) = 1, x$, and x^2 .

$(W_g f)(n, x_j)$ themselves. As a consequence, we have experimentally found that by using our Theorem 4.1.3-based algorithm, we may actually compute the smooth centroid signal (29) in much less time than it takes to compute the noisy ridgeline signal (14). That is, using Theorem 4.1.3, we can get better results in less time. The key idea is to notice that both the denominator and numerator of (29) may be computed using Theorem 4.1.3 where $w(x)$ is taken to be the constant function 1 and the linear function x , respectively. Indeed, it turns out that even more information about the signal may be found by also applying Theorem 4.1.3 where $w(x)$ is taken to be x^2 .

To be precise, let us assume that the frequencies $\{x_j\}_{j=1}^J$ and the window g of Theorem 4.1.3 have been fixed. Our first step is to compute the discrete Fourier transforms (Figure 12) of the three distinct ψ_{zp} 's we obtain by letting $w(x) = 1, x$ and x^2 , respectively. To do this, we first form the vectors $\{X_j\}_{j=1}^J \in \mathbb{R}^{4N+4}$ obtained by evaluating the $2N$ th Dirichlet kernel over $\mathbb{Z}_{4N+4}$:

$$X_j := \begin{bmatrix} D_{2N}\left(2\pi\left(x_j - \frac{0}{4N+4}\right)\right) \\ D_{2N}\left(2\pi\left(x_j - \frac{1}{4N+4}\right)\right) \\ \vdots \\ D_{2N}\left(2\pi\left(x_j - \frac{4N+3}{4N+4}\right)\right) \end{bmatrix}.$$

We then form three vectors $Y_1, Y_x, Y_{x^2} \in \mathbb{R}^{4N+4}$ as linear combinations of the X_j 's:

$$Y_1 := \sum_{j=1}^J X_j, \quad Y_x := \sum_{j=1}^J x_j X_j, \quad Y_{x^2} := \sum_{j=1}^J x_j^2 X_j.$$

By Theorem 4.1.3, these three vectors are precisely the discrete Fourier transforms needed to produce:

$$\sum_{j=1}^J |(\mathcal{W}_g f)(n, x_j)|^2, \quad \sum_{j=1}^J x_j |(\mathcal{W}_g f)(n, x_j)|^2, \quad \sum_{j=1}^J x_j^2 |(\mathcal{W}_g f)(n, x_j)|^2, \quad (30)$$

respectively. We further note that Y_1, Y_x, Y_{x^2} need only be computed once: in our current MATLAB implementation, they are computed at the beginning of the algorithm; they may also be computed in advance, stored offline, and simply read into the code when needed. Having these vectors, we then compute the quantities (30) (in real time) as follows: for any given time index n , we form the vector $Z_n \in \mathbb{C}^{4N+4}$ obtained by zero-padding the values $\{f(n+m)\overline{g(m)}\}_{m=-N}^N$:

$$Z_n := \begin{bmatrix} f(n-N)\overline{g(-N)} \\ \vdots \\ f(n+N)\overline{g(N)} \\ 0 \\ \vdots \\ 0 \end{bmatrix}.$$

That is, Z_n is a vectorization of the function $\varphi_{zp} \in \ell(\mathbb{Z}_{4N+4})$ given in Theorem 4.1.3. Letting $|\hat{Z}_n|^2$ be the vector obtained by pointwise-squaring the modulus of the FFT

$\hat{Z}_n$ of Z_n , Theorem 4.1.3 then gives:

$$\sum_{j=1}^J |(W_g f)(n, x_j)|^2 = \frac{1}{4N+4} Y_1^T |\hat{Z}_n|^2, \quad (31)$$

$$\sum_{j=1}^J x_j |(W_g f)(n, x_j)|^2 = \frac{1}{4N+4} Y_x^T |\hat{Z}_n|^2, \quad (32)$$

$$\sum_{j=1}^J x_j^2 |(W_g f)(n, x_j)|^2 = \frac{1}{4N+4} Y_{x^2}^T |\hat{Z}_n|^2. \quad (33)$$

That is, after the one-time cost of computing Y_1 , Y_x and Y_{x^2} , the amount of computation needed to find the quantities in (30) at any given time n consists of 1) a single FFT of size $4N + 4$ (and as such, N should be chosen so that $4N + 4$ is a power of 2), 2) taking a pointwise modulus square of the result, and 3) taking three dot products of this vector with Y_1 , Y_x and Y_{x^2} . Once (31), (32) and (33) are computed, the value of the centroid curve (29) at n is simply $Y_x^T |\hat{Z}_n|^2 / Y_1^T |\hat{Z}_n|^2$.

The second moment (33) may be used to determine *spectral deviations* of vertical strips of the spectrogram, that is, the square root of the variance of the distribution $\{|(W_g f)(n, x_j)|^2\}_{j=1}^J$ about its mean. This number indicates how “spread out” the energy of the spectrogram is along the vertical axis, and is explicitly defined as:

$$\sqrt{\frac{\sum_{j=1}^J x_j^2 |(W_g f)(m, x_j)|^2}{\sum_{j=1}^J |(W_g f)(m, x_j)|^2} - \left(\frac{\sum_{j=1}^J x_j |(W_g f)(m, x_j)|^2}{\sum_{j=1}^J |(W_g f)(m, x_j)|^2} \right)^2} = \sqrt{\frac{Y_{x^2}^T |\hat{Z}_n|^2}{Y_1^T |\hat{Z}_n|^2} - \left(\frac{Y_x^T |\hat{Z}_n|^2}{Y_1^T |\hat{Z}_n|^2} \right)^2}.$$

Figure 13 shows the centroids of the first 30-seconds-worth of Figure 7(c) along with the centroids plus or minus the standard deviation. From Figure 13, we can see that the spectral deviations are very small implying that the spectral centroids are an accurate way of measuring the dominant frequency of the spectrogram.




Figure 13: Spectral centroids (plus or minus their spectral deviations) of Figure 7(c).

In the next chapter, we apply this algorithm to real-world terahertz radar data, successfully determining a subject's heart and breath rate in several instances. Before continuing however, we must take a moment to discuss an important preprocessing step we often applied to the radar data as a means for dealing with clutter.

4.2.1 Preprocessing by High-Pass Filtering. Our radar data often has a DC component that is the result of clutter. Specifically, while the beam created by Dr. Petkie's radar is often very narrow, and as such, most of the reflected energy has bounced off of a moving object, there is always some portion of the energy that reflects off of stationary objects. These clutter reflections cause no Doppler shift. In other words, the reflected wave has the same frequency as the transmitted wave, which, after demodulation, is manifested as a DC component in the radar data. As a strong DC component may significantly distort the spectral centroid, we often first preprocess the radar data to remove this component. Specifically, we first pass the data through a high-pass filter. Specifically, letting $e \in \ell^1(\mathbb{Z})$ be some low-pass filter, where

$$\sum_{n=-\infty}^{\infty} e(n) = 1,$$

we may first compute the local average of some signal f about some time m as:

$$\begin{aligned}
\check{x}(m) &= \operatorname{argmin}_{x \in \mathbb{R}} \sum_{n=-\infty}^{\infty} |f(m-n) - x|^2 e(n) \\
&= \operatorname{argmin}_{x \in \mathbb{R}} \sum_{n=-\infty}^{\infty} \{[f(m-n)]^2 - 2xf(m-n) + x^2\} e(n) \\
&= \operatorname{argmin}_{x \in \mathbb{R}} \sum_{n=-\infty}^{\infty} [f(m-n)]^2 e(n) - 2x \sum_{n=-\infty}^{\infty} f(m-n)e(n) + x^2 \sum_{n=-\infty}^{\infty} e(n) \\
&= \operatorname{argmin}_{x \in \mathbb{R}} \sum_{n=-\infty}^{\infty} [f(m-n)]^2 e(n) - 2x \sum_{n=-\infty}^{\infty} f(m-n)e(n) + x^2. \tag{34}
\end{aligned}$$

The explicit value of the local average $\check{x}(m)$ thus corresponds to the vertex of the quadratic in (34), namely:

$$\check{x}(m) = \sum_{n=-\infty}^{\infty} f(m-n)e(n) = (f * e)(m).$$

For significantly cluttered radar data f , this low-passed version $(f * e)$ of f contains much of the DC component. By subtracting this component off, that is, by computing

$$f(m) - (f * e)(m) = (f * \delta_0)(m) - (f * e)(m) = [f * (\delta_0 - e)](m), \tag{35}$$

where δ_0 is the Dirac- δ filter centered at the origin, we produce a new signal with little to no DC component. Note this new signal (35) is a high-passed version of f . Indeed, taking the Fourier transform of $\delta_0 - e$ is:

$$F(\delta_0 - e)(x) = \sum_{n=-\infty}^{\infty} (\delta_0 - e)(n)e^{-2\pi i n x} = 1 - \sum_{n=-\infty}^{\infty} e(n)e^{-2\pi i n x},$$

whose value at the origin is $F(\delta_0 - e)(0) = 1 - \sum_{n=-\infty}^{\infty} e(n) = 0$.


V. Experimental Results

5.1 Examples

The data for the following three examples was collected by Dr. Petkie’s 240 GHz radar system, where the transmitter/receiver was about 10 meters from the subject. The three 60 second samples have a sampling rate of 20 kHz.

5.1.1 Example 1. In our first example, the subject is known to be at rest. The subject alternates between fast and slow breathing rates, starting off breathing at a fast rate for about the first 15 seconds and then slowing down for the next 15 seconds and then repeating this behavior for the last 30 seconds. To see how “spread out” the energy of the spectrogram is, we plot the centroid along with the centroid plus or minus its standard deviation, shown in Figure 14(a). These values for the spectral centroids and spectral deviations can be computed directly from the IQ data using the Toeplitz matrix-based algorithm discussed above in Section 4.2. Figure 14(b) shows the spectrogram of the centroid, which should provide information about the heart and breath rates of the individual. Here, the red curve indicates a slowly varying breath rate around 0.5 Hz. In particular, at around 30 seconds, we can see the subject’s breath rate decrease, which agrees with the ground truth. Though we do not see the heart rate from Figure 14(b) directly, this information may be gleaned by summing across the rows of the spectrogram, as shown in Figure 14(c). Here, we can see a peak around 0.5 Hz that corresponds to the breath rate, while the smaller peaks in Figure 14(c) occurring at 1 and 1.5 Hz could correspond to the heart rate.

In particular, to determine how accurate these estimates of breath and heart rates are, we can compare them against ground truth data taken from the subject in the form of respiration belt and EKG signals. Figure 15(a) shows the EKG data which was collected at the same time as the radar data analyzed in Figure 14. The spectrogram of the EKG signal, shown in Figure 15(b), shows what frequencies are present in the EKG signal at a given time. These frequencies correspond to the actual heart rate of the subject. In particular, we can see that the heart rate of the subject is




(a) A 10 second interval of the centroid and its standard deviation. The centroid gives us a way to measure the Doppler shift at a given time and gives us a signal that should be a good approximation to the velocity signal.

(b) The spectrogram of the centroids should provide information corresponding to heart and breath motion. We can see the varying breath rate slightly below 0.5 Hz. More specifically, notice that at 30 seconds the breath rate decreases as a result of the subject alternating between fast and slow breathing approximately every 15 seconds.


(c) The sum of the rows of (b) shows us what frequencies occur the most over the 60 second sample. We can see a peak at about 0.5 Hz, which corresponds to the breath rate. The peaks at 1 and 1.5 Hz might correspond to the heart rate.

Figure 14: Spectral deviations and time-frequency analysis of the spectral centroid for Example 1.


about 1.5 Hz, showing that the peak in Figure 14(c) at 1.5 Hz does indeed correspond to the heart rate of the subject. Meanwhile, the respiration belt signal is shown in Figure 15(c). Its spectrogram, seen in Figure 15(d), shows that the breath rate of the subject is slightly less than 0.5 Hz. In short, for this particular data set, we were able to accurately determine both the heart and breath rates of the subject.




(a) The EKG signal for Example 1, which records the electrical activity that causes the heart to contract.



(b) The spectrogram of (a), which shows us what frequencies are present in the EKG signal (a) at a given time. We see that the heart rate is around 1.5 Hz, meaning that the subject's heart beats 3 times every 2 seconds. This heart rate was accurately determined from Figure 14(c)



(c) The respiration belt signal for Example 1, which is obtained by strapping the respiration belt around the abdomen or chest, inflating it, and then measuring the changes in pressure caused by inhalation and exhalation.



(d) The spectrogram of (c), which shows us what frequencies are present in the respiration belt signal (c) at a given time. From this spectrogram, we can see that the respiration rate is varying slightly below 0.5 Hz. This breath rate may be determined from the radar data. Again, we notice that at 30 seconds the breath rate decreases as a result of the subject alternating between fast and slow breathing approximately every 15 seconds.

Figure 15: Truth data for heart and breath rate and their corresponding spectrograms for Example 1.






Figure 16: Spectral deviations and time-frequency analysis of the spectral centroid for Example 2.


5.1.2 Example 2. In our second example, the subject is known to be at rest, breathing at a constant rate during the 60 second sample. Paralleling Figures 14 and 15 for Example 1, Figure 16(a) shows the centroid curve (plus or minus the standard deviation curve) obtained from the radar data using the algorithm given above in Section 4.2. Taking the spectrogram of this curve (Figure 16(b)) and summing the result across the rows (Figure 16(c)) we gain some insight into the frequencies of the velocity signal. In particular, experience tells us that the strong peak around 0.25 Hz corresponds to the subject's breath rate, a fact confirmed by taking the spectrogram (Figure 17(d)) of the ground truth respiration data (Figure 17(c)). However, we were unable to detect a meaningful second peak in either Figure 16(b) or Figure 16(c) which would indicate the 1.5 Hz heart rate we know the subject had (Figures 17(a) and (b)). In short, for this particular data set, we were able to accurately determine the breath rate from the radar data, but were unable to determine the heart rate.




(a) The EKG signal for Example 2.



(b) The spectrogram of (a), showing a heart rate of about 1.5 Hz, meaning that the subject's heart beats 3 times every 2 seconds. The harmonic at 3 Hz is a result of the EKG signal (a) having sharp peaks. This heart rate is undetectable in the corresponding radar data.




(c) The respiration belt signal for Example 2.



(d) The spectrogram of (c), showing a respiration rate of about 0.25 Hz, meaning the subject takes approximately 1 breath every 4 seconds. This breath rate may be determined from the radar data.

Figure 17: Truth data for heart and breath rate and their corresponding spectrograms for Example 2.




(a) The centroid and its standard deviation. We can see that between 25 and 40 seconds the subject is holding his breath, allowing the higher frequency of the heart rate to be seen.

(b) The spectrogram of the centroid. We can see that the subject is holding his breath between 25 and 40 seconds and, while the subject is breathing, we see that the subject's breath rate is about 0.25 Hz. The higher frequencies around 1.5 Hz could correspond to the heart rate of the subject.


(c) The sum of the rows of (b) again showing a peak at about 0.25 Hz, which corresponds to the breath rate of the subject. We also see a relatively strong second frequency peak right below 1.5 Hz that could correspond to the heart rate of the subject.

Figure 18: Spectral deviations and time-frequency analysis of the spectral centroid for Example 3.


5.1.3 Example 3. In our third example, the subject is known to be at rest, breathing normally for the first 20 seconds, then holding his breath for 20 seconds, and then breathing normally for the last 20 seconds. Figure 18(a) shows the centroid curve, in which we can see that between 25 and 40 seconds the subject is holding his breath, allowing the higher frequency of the heart rate to be seen. The spectrogram of this curve is shown in Figure 18(b), and the sum across its rows is seen in Figure 18(c). We can see in Figure 18(b) that the subject is holding his breath between 25 and 40 seconds, and otherwise has a breath rate of 0.25 Hz as confirmed by taking the spectrogram (Figure 19(d)) of the ground truth respiration data (Figure 19(c)). Meanwhile, the nontrivial second peak in Figure 18(c) indicates the subject's heart rate is approximately 1.5 Hz, a fact confirmed by the spectrogram (Figure 19(b)) of the EKG (Figure 19(a)). Thus, for this particular data set, we were able to accurately determine both the breath and heart rates from the radar data.




(a) The EKG signal for Example 3.



(b) The spectrogram of (a) showing a heart rate of about 1.5 Hz, meaning that the subjects heart beats 3 times every 2 seconds.



(c) The respiration belt signal for Example 3. We can see that between 25 and 40 seconds the subject is holding his breath.



(d) The spectrogram of (c) showing that the subject is holding his breath between 25 and 40 seconds. During the first 25 seconds, the subject is breathing, and we can see that the subject's breath rate is about 0.25 Hz, meaning the subject takes approximately 1 breath every 4 seconds. This breath rate may be determined from the radar data.

Figure 19: Truth data for heart and breath rate and their corresponding spectrograms for Example 3.

5.2 Spectral Centroids of the Velocity Centroid's Spectrogram

Up to this point, we have only used spectral centroids as a method to generate a velocity estimate, whose spectrogram is taken to reveal heart and breath rates. We now turn our focus to a repeated application of the spectral centroid idea. In particular, after generating the velocity signal as before using spectral centroids, we then take the spectral centroids of this velocity signal to get our heart and breath rate estimates. By taking spectral centroids of the velocity signal over a band of frequencies which typically contain only heart rate information, we form our heart rate estimate. Similarly, by taking spectral centroids of the velocity signal over a band of frequencies which typically only contain breath rate information, we form our breath rate estimate. To be more precise, we take a closer look at the spectrogram of a velocity signal obtained from a 240 GHz CW radar signal.

Consider, for example, the spectrogram shown in Figure 20, which was obtained by performing a time-frequency analysis on a velocity signal obtained from a 240 GHz CW radar signal. We can see a dominant frequency at 0.25 Hz, and a second, less intense frequency component at 1 Hz. From experience, we expect these two frequencies to be the breath and heart rates of the subject, respectively. Suppose we take the spectral centroids of the horizontal band of this spectrogram that lies between 0 and 0.4 Hz on the vertical axis. Doing this would give us a centroid curve that should do a decent job of indicating the breath rate of the subject at any given time. Similarly, we could compute the spectral centroids of the band of the spectrogram that lies between 0.75 and 1.5 Hz on the vertical axis, which would give us a centroid curve that should do a decent job of indicating the subject's heart rate. We note this method is far from perfect, as it makes some *a priori* assumptions on the range of a person's breath and heart rates. Nevertheless, the approach seemed to work very well in the next three examples, when compared with the ground truth data provided by the respiration belt and EKG signals.




Figure 20: Spectrogram showing heart and breath rate.

The IQ data for Examples 4 through 6 was collected by Dr. Petkie’s 240 GHz radar system. For Examples 4 through 6, the data was sampled at a rate of 10 kHz and was collected for 80 seconds on a subject that was about 10 meters from the transmitter/reciever. For Example 7, the data was sampled at 600 Hz and was collected for 60 seconds on a subject that was about 30 meters from the transmitter/receiver. For Examples 4 through 6, the breath and heart rate centroids, computed over frequency ranges of 0 to 0.4 Hz and 0.75 to 1.5 Hz, respectively, are compared to pressure belt and EKG ground truth data. Though these breath and heart rate centroid curves are also computed for Example 7, no simultaneous ground truth data was available for comparison.

5.2.1 Example 4. In our fourth example, the subject is known to be at rest, breathing normally for the entire 80 second sample. Figure 21(a) shows a 10 second portion of the real part of the radar signal. The spectrogram of this signal, seen in Figure 21(b), shows us what frequencies are present in the signal at a given time. These frequencies are proportional to the velocity of the subject, which corresponds to the motion of the chest wall. In particular, the velocity of the subject’s body is well-approximated by the spectral centroids (Figure 21(c)) of the spectrogram. These centroids measure the Doppler shift at a given time and can be computed directly from the radar signal (Figure 21(a)) using the Toeplitz matrix-based algorithm discussed


above in Section 4.2. By doing a time-frequency analysis on Figure 21(c), we can extract the periodic components of the velocity signal which should correspond to heart and breath motion. The spectrogram of the centroids is shown in Figure 21(d), where we can see dominant frequencies at 1 and 0.25 Hz, corresponding to the heart and breath rates, respectively. By computing spectral centroids of the velocity signal's spectrogram (Figure 21(d)) over the frequencies between 0.75 and 1.5 Hz, we form our heart rate estimate, which is also plotted in Figure 21(d) as the upper black curve. By also plotting the same curve over the spectrogram of the EKG (Figure 21(e)), we see that our estimate is very close to the ground truth. Meanwhile, computing the spectral centroids of the velocity signal's spectrogram over the frequencies between 0 and 0.4 Hz yields our breath rate estimate. This estimate, plotted as the lower black curve in Figure 21(d), is a very good estimate of the subject's actual breath rate, as seen by comparing it to the spectrogram of the pressure belt signal (Figure 21(f)). In a real-life application of this theory, one would simply first compute the velocity centroids, and then compute both the heart and breath rate centroids. That is, one would have Figure 21(a), and compute Figure 21(c) and the two black curves in Figure 21(d) using three applications of the spectral centroid algorithm discussed in Section 4.2. That is, in a real-life application, the spectrograms depicted in Figures 21(b), (d), (e), and (f) would never be explicitly computed, as the whole purpose of our centroid computing algorithm is to bypass the large cost of computing Figures 21(b) and (d), while the ground truth data behind Figures 21(e) and (f) would be entirely unavailable. In short, both here and in the two examples to follow, the role of these six images is to convince the reader of the feasibility of accurately estimating heart and breath rates using only three computationally-fast applications of the main ideas of Chapter IV.

5.2.2 Example 5. In our fifth example, the subject is known to be at rest, and holds his breath in the middle of the 80 second sample. From Figure 22(c), we can see that between 15 and 55 seconds the subject holds his breath, allowing

the higher frequency of the heart rate to be seen. Because of this breath-holding, a breath-based component of the spectrogram (Figure 22(d)) of the velocity signal can only be seen after 55 seconds. Figure 22(e) shows that our heart rate estimate closely matches the ground truth throughout. However, Figure 22(f) shows that this method of estimating breath rate performs poorly when the subject is not breathing.

5.2.3 Example 6. In our sixth example, the subject is known to be at rest, and takes two yawns during the 80 second sample. From Figures 23(b) and (c), we can see that the subject yawns at 30 and 60 seconds, manifested in the radar signal as Doppler shifts of large magnitude. Here, both our heart and breath rate estimates are very accurate when compared to the ground truth. In particular, in Figure 23(f), one may see that our estimate even captures the slight change in breath rate caused by the yawns.


5.2.4 Example 7. In our seventh example, the subject is known to be at rest, breathing normally for the entire 60 second sample. Though no ground truth is available for this data set, we may nevertheless form our breath and heart rate estimates, as shown in Figure 24. These estimates seem consistent with data gathered from this particular subject in the past. The reason this data set is included in this study is that it was gathered at the relatively large distance of 30 meters. That is, Dr. Petkie’s radar shows promise in being able to facilitate heart and breath rate detection at distances not often seen before in the literature.



(a) A ten second portion of the real part of a high-pass filtered radar signal that was sampled at 10 kHz on a subject located about 10 meters from the transmitter/receiver.

(b) The spectrogram of the radar signal, which shows us what frequencies are present at what time. The frequencies are proportional to the velocity of the target, which corresponds to the motion of the chest wall.

(c) The spectral centroids of (b), which measure the Doppler shift at a given time. Using the algorithm of Section 4.2 these spectral centroids can be computed directly from the radar signal (a).




(d) The spectrogram of (c), which indicates any periodic component of the velocity signal (c) related to the heart and breath rates of the subject. Also plotted are the black centroid curves corresponding to the bands of this spectrogram from 0 to 0.4 Hz and from 0.75 to 1.5 Hz, which are our breath and heart rate estimates, respectively. Again, using the ideas of Chapter IV, these estimates can be computed directly from (c).


(e) The spectrogram of the EKG signal, which gives us a way to check our method for heart rate detection. The centroid curve corresponding to the band of the spectrogram (d) between 0.75 and 1.5 Hz is also plotted to see how well it matches the true heart rate. We can see that our estimated heart rate is very close to the actual heart rate of about 1 Hz. The harmonic at about 2 Hz is a result of a “peaky” EKG signal.

(f) The spectrogram of the respiration belt signal, which gives us a way to check our method for breath rate detection. The centroid curve corresponding to the band of the spectrogram (d) between 0 and 0.4 Hz is also plotted to see how well it matches the true breath rate. We can see that our estimated breath rate is very close to the actual breath rate of about 0.25 Hz.


Figure 21: Heart and breath rate detection in Example 4.




(a) The real part of the high-pass filtered radar signal that was sampled at 10 kHz on a subject located about 10 meters from the transmitter/receiver. We can see that between 15 and 55 seconds the subject is holding his breath.




(b) The spectrogram of the radar signal.




(c) The spectral centroids of (b). Again, we can see that between 15 and 55 seconds the subject holds his breath, allowing the higher frequency of the heart rate to be seen.



(d) The heart and breath rate estimates on top of the spectrogram of (c).




(e) The heart rate estimate on top of the spectrogram of the EKG signal. We can see that our heart rate estimate is very close to the actual heart rate of about 1 Hz.




(f) The breath rate estimate on top of the spectrogram of the respiration belt signal.


Figure 22: Heart and breath rate detection in Example 5.




(a) A ten second portion of the high-pass filtered radar signal that was sampled at 10 kHz on a subject located about 10 meters from the transmitter/receiver.




(b) The spectrogram of the radar signal. We can see that the subject yawns at 30 and 60 seconds.




(c) The spectral centroids of (b).



(d) The heart and breath rate estimates on top of the spectrogram of (c).




(e) The heart rate estimate on top of the spectrogram of the EKG signal. We can see that our heart rate estimate is very close to the actual heart rate of about 1 Hz.




(f) The breath rate estimate on top of the spectrogram of the respiration belt signal. We can see that our breath rate estimate is very close to the actual breath rate of about 0.25 Hz, with subtle variation in the rate caused at 30 and 60 seconds by the yawns.


Figure 23: Heart and breath rate detection in Example 6.




(a) The real part of the high-pass filtered radar signal that was sampled at 600 Hz on a subject located about 30 meters from the transmitter/receiver.



(b) The spectrogram of the radar signal, which shows us what frequencies are present at what time. The frequencies are proportional to the velocity of the target, which corresponds to the motion of the chest wall.



(c) The spectral centroids of (b), which measure the Doppler shift at a given time. Using a Toeplitz matrix-based, computational trick, the spectral centroids can be computed directly from the radar signal.



(d) The spectrogram of (c), which provides information corresponding to the periodic motion related to the heart and breath rates of the target. Also plotted are the centroid curves corresponding to the bands of this spectrogram from 0 to 0.4 Hz and from 0.75 to 1.5 Hz. We call these centroid curves the breath rate estimate and the heart rate estimate, respectively. Using a Toeplitz matrix-based, computational trick, the heart and breath rate estimates can be computed directly from (c). We see that our heart rate estimate is 1 Hz and our breath rate estimate is 0.25 Hz.

Figure 24: Heart and breath rate detection in Example 7.

Bibliography

1. Boashash, Boualem. *Time Frequency Signal Analysis and Processing*. UK: Elsevier Ltd., 2003.
2. Chen, Kun-Mu, Devendra Misra, Huei Wang, Huey-Ru Chuang, and Elliot Postow. "An X-Band Microwave Life-Detection System," *IEEE Transactions on Biomedical Engineering*, BME-33: 697–701 (July 1986).
3. Chen, Kun-Mu, Yong Huang, Jianping Zhang, and Adam Norman. "Microwave Life-Detection Systems for Searching Human Subjects Under Earthquake Rubble or Behind Barrier," *IEEE Transactions on Biomedical Engineering*, 27: 105–114 (January 2000).
4. Chuang, Huey-Ru, Y. F. Chen, and Kun-Mu Chen. "Automatic Clutter-Canceler for Microwave Life-Detection Systems," *IEEE Transactions on Instrumentation and Measurement*, 40: 747–750 (August 1991).
5. Cohen, Leon. *Time-Frequency Analysis*. New Jersey: Prentice-Hall PTR, 1995.
6. Coleman, P. D. and R. C. Becker. "Present State of the Millimeter Wave Generation and Technique Art–1958," *IEEE Transactions on Microwave Theory and Techniques*, 7: 42–61 (January 1959).
7. Droitcour, Amy, Victor Lubecke, Jenshan Lin, and Olga Boric-Lubecke. "A Microwave Radio for Doppler Radar Sensing of Vital Signs," *2001 IEEE MTT-S International Microwave Symposium Digest*, 1: 175–178 (May 2001).
8. Eronen, Antti and Anssi Klapuri. "Musical Instrument Recognition Using Cepstral Coefficients and Temporal Features," *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2: 753–756 (June 2000).
9. Greneker, E. F. "Radar Sensing of Heartbeat and Respiration at a Distance," *Radar*, 97: 150–154 (October 1997).
10. Lin, James C. "Noninvasive Microwave Measurement of Respiration," *Proceedings of the IEEE*, 96: 1530 (October 1975).
11. Lin, James C., Joseph Kiernicki, Martin Kiernicki, and Paul B. Wollschlaeger. "Microwave Apexcardiography," *IEEE Transactions on Microwave Theory and Techniques*, MTT-27: 618–620 (June 1979).
12. Michahelles, Florian, Peter Matter, Albrecht Schmidt, and Bernt Schiele. "Applying Wearable Sensors to Avalanche Rescue: First Experiences with a Novel Avalanche Beacon," *Computers and Graphics*, 27: 839–847 (2003).
13. Michahelles, Florian, Ramon Wicki, and Bernt Schiele. "Less Contact: Heart-Rate Detection Without Even Touching the User," *Proceedings of the Eighth International Symposium on Wearable Computers*, (2004).

14. Morbiducci, Umberto, Lorenzo Scalise, Mirko De Melis, and Muaro Grigioni. "Optical Vibrocardiography: A Novel Tool for the Optical Monitoring of Cardiac Activity," *Annals of Biomedical Engineering*, 35: 45–58 (January 2007).
15. Pedersen, P. C., C. C. Johnson, C. H. Durney, and D. G. Bragg. "An Investigation of the Use of Microwave Radiation for Pulmonary Diagnostics," *IEEE Transactions on Biomedical Engineering*, 410–412 (September 1976).
16. Peeters, Geoffroy, Stephen McAdams, and Perfecto Herrera. "Instrument Sound Description in the Context of MPEG-7," *Proceedings of International Computer Music Conference*, 1–4 (September 2000).
17. Peltonen, Vesa, Juha Tuomi, Anssi Klapure, Jyri Huopaniemi, and Time Sorsa. "Computational Auditory Scene Recognition," *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2: 1941–1944 (2002).
18. Siegel, Peter H. "Terahertz Technology," *IEEE Transactions on Microwave Theory and Techniques*, 50: 910–928 (March 2002).
19. Siegel, Peter H. "Terahertz Technology in Biology and Medicine," *IEEE Transactions on Microwave Theory and Techniques*, 52: 2438–2447 (October 2004).
20. Wiltse, James C. "History of Millimeter and Submillimeter Waves," *IEEE Transactions on Microwave Theory and Techniques*, MTT-32: 1118–1127 (September 1984).

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.						
1. REPORT DATE (DD-MM-YYYY) 04-06-2008		2. REPORT TYPE Master's thesis			3. DATES COVERED (From — To) Sep 2007–Jun 2008	
4. TITLE AND SUBTITLE TIME-FREQUENCY ANALYSIS OF TERAHERTZ RADAR SIGNALS FOR RAPID HEART AND BREATH RATE DETECTION				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Massar, Melody L.				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFIT/ENC 2950 Hobson Way Bldg 641, Rm 302 WPAFB, OH 45433					8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GAM/ENC/08-05	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT We develop new time-frequency analytic techniques which facilitate the detection of a person's heart and breath rates from the Doppler shift the movement of their body induces in a terahertz radar signal. One straightforward means of doing such an analysis is the following: take the spectrogram of the ridgeline of the spectrogram of the radar signal. Instead of following this approach exactly, we consider an alternate method in which the ridgeline of the radar signal's spectrogram is replaced with a signal computed from spectral centroids. By using spectral centroids, rather than the ridgeline, we produce a smooth signal that avoids traditional problems with ridgelines, such as jump discontinuities and overquantization. This new method for time-frequency analysis uses a Toeplitz matrix-based algorithm that has a fast Fourier transform-based implementation, and permits centroids of the vertical strips of the spectrogram of the radar signal to be computed without ever having to explicitly compute the spectrogram itself.						
15. SUBJECT TERMS terahertz, radar, heart, breath, time-frequency analysis, spectrogram, centroid						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES	
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU		80	
U	U	U	19a. NAME OF RESPONSIBLE PERSON Matthew C. Fickus			
						19b. TELEPHONE NUMBER (include area code) (937) 255-3636 x 4513