A SEQUENTIAL METHOD FOR PASSIVE DETECTION, CHARACTERIZATION, AND LOCALIZATION OF MULTIPLE LOW PROBABILITY OF INTERCEPT LFMCW SIGNALS

by

Brandon M. Hamschin

B.S.E.E, Penn State Erie, The Behrend College 2005

B.S. Mathematics, Penn State Erie, The Behrend College 2005

M.S.E.E, University of Pittsburgh, 2011

Submitted to the Graduate Faculty of

the Swanson School of Engineering in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2015

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Brandon M. Hamschin

It was defended on

July 22, 2015

and approved by

Patrick J. Loughlin, Ph.D., Professor

Amro A. El-Jaroudi, Ph.D., Associate Professor

Michael T. Grabbe, Ph.D., Senior Engineer, JHU/APL

Steven P. Jacobs, Ph.D., Assistant Professor

Zhi-Hong Mao, Ph.D., Associate Professor

Ervin J. Sejdic, Ph.D., Assistant Professor

Dissertation Director: Patrick J. Loughlin, Ph.D., Professor

A SEQUENTIAL METHOD FOR PASSIVE DETECTION, CHARACTERIZATION, AND LOCALIZATION OF MULTIPLE LOW PROBABILITY OF INTERCEPT LFMCW SIGNALS

Brandon M. Hamschin, PhD

University of Pittsburgh, 2015

A method for passive Detection, Characterization, and Localization (DCL) of multiple low power, Linear Frequency Modulated Continuous Wave (LFMCW) (*i.e.*, Low Probability of Intercept (LPI)) signals is proposed. We demonstrate, via simulation, laboratory, and outdoor experiments, that the method is able to detect and correctly characterize the parameters that define two simultaneous LFMCW signals with probability greater than 90% when the signal to noise ratio is -10 dB or greater. While this performance is compelling, it is far from the Cramer-Rao Lower Bound (CRLB), which we derive, and the performance of the Maximum Likelihood Estimator (MLE), whose performance we simulate. The loss in performance relative to the CRLB and the MLE is the price paid for computational tractability. The LFMCW signal is the focus of this work because of its common use in modern, low-cost radar systems.

In contrast to other detection and characterization approaches, such as the MLE and those based on the Wigner-Ville Transform (WVT) or the Wigner-Ville Hough Transform (WVHT), our approach does not begin with a parametric model of the received signal that is specified directly in terms of its LFMCW constituents. Rather, we analyze the signal over time intervals that are short, non-overlapping, and contiguous by modeling it within these intervals as a sum of a small number sinusoidal (*i.e.*, harmonic) components with unknown frequencies, deterministic but unknown amplitudes, unknown order (*i.e.*, number of harmonic components), and unknown noise autocorrelation function. It is this model of the

data that makes the solution computationally feasible, but also what leads to a degradation in performance since estimates are not based on the full time series. By modeling the signal in this way, we reliably detect the presence of *multiple* LFMCW signals in colored noise without the need for prewhitening, efficiently estimate (*i.e.*, characterize) their parameters, provide estimation error variances for a subset of these parameters, and produce Time-Difference-of-Arrival (TDOA) estimates that can be used to estimate the geographical location (*i.e.*, localize) of each LFMCW source. We demonstrate the performance of our method via simulation and real data collections, which are compared to the Cramer-Rao Lower Bound (CRLB).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

PREFACE

The main three chapters that comprise this dissertation are drawn almost entirely from three journal papers. At the time of the defense of this dissertation, the first [19] is in the final stages of preparation for submission to the IEEE Transactions on Aerospace and Electronic Systems and the other two [18, 20] were nearly complete. Chapter 1, which is drawn from [19], focuses on the theory of the method developed for Detecting, Characterizing, and Localizing (DCL) multiple low power Linear Frequency Modulated Continuous Wave (LFMCW) signals. Chapter 2 is drawn from [18] and develops the Cramer Rao Lower Bound (CRLB) associated with the LFMCW signal model. Finally, Chapter 3 is drawn from [20] and focuses on experimental results. We expand on some aspects of these chapters and provide preliminary results associated with areas of future work in the Appendices. Since the three chapters are so closely related to [18–20], each of which require introduction and background material, some content in the chapters may appear redundant.

ACKNOWLEDGMENTS

I would like to acknowledge and extend my sincere gratitude to several individuals and two financial channels that have made this work possible. First to the Office of the Assistant Secretary of Defense (ASD), Space and Sensor Systems for providing the initial funding that started this project. In addition to the financial support provided by the ASD, the sponsor, Dr. John Novak added invaluable technical insight and continued motivation to focus this work in a direction that would make it rapidly valuable to the men and women in the armed forces that his office is chartered with developing technology to support. Also, to Dr. John Clancy for including me in the initial concept that was eventually funded by the ASD and ultimately gave rise to the research contained in this dissertation. Next, to Lou Colangelo, Brian Geesaman, Jack Keane, and Glenn Mitzel at The Johns Hopkins University Applied Physics Laboratory (JHU/APL) [1] for their willingness to provide financial and managerial support to take the evaluation of these ideas from simulation to the real world by funding their testing on real data. It was this support that enabled much of what is in Chapter 3. Also to the JHU/APL Stuart S. Janney Publication Program that funded portions of the time spent developing this manuscript and those planned for submission to the IEEE Transactions on Aerospace and Electronic Systems. Next, to my supervisors, Andrew Newman and Mike Foust for not only giving me the freedom but also the support to pursue areas of work that are somewhat outside the scope of our group's focus. Next, to John Ferguson, without whom Chapter 2 would not have been possible. His combined knowledge of the theoretical aspects of signal processing and their practical application on real hardware devices was invaluable. A special thanks to Professor Pat Loughlin for his commitment to guiding me through the PhD process while I was both a full-time graduate student and a full-time employee at JHU/APL. Without his willingness to remain my advisor during both phases of my career, especially the latter, I would have never made it this far. Finally, to Mike Grabbe who, in addition to his technical contributions to this dissertation as a committee member,

[1] Distribution Statement A: Approved for Public Release; Distribution is Unlimited

has patiently mentored me during my development as a technical professional by providing countless hours of technical teaching and advice, without which I would be far less equipped to contribute to the field.

1.0 THEORY

This dissertation is devoted to developing and testing a practical method for Detecting, Characterizing, and Localizing (DCL) an unknown number of low-power Linear Frequency Modulated Continuous Wave (LFMCW) (*i.e.*., Low Probability of Intercept (LPI)) radar signals. In this chapter we focus on developing the method analytically and analyzing its performance via simulation, wherein we found that practicality comes at the expense of a loss in estimation accuracy when compared to the Maximum Likelihood Estimator (MLE). In contrast to the MLE, the practicality of the method is a consequence of it being simultaneously 1) computationally tractable, 2) capable of determining the number of LFMCW signals in the environment, 3) capable of estimating the parameters that define each LFMCW signal, 4) able to operate without knowledge of the background noise power or autocorrelation function and 5) capable of yielding estimates of Time Difference of Arrival (TDOA) between multiple receivers which enables localization of the radar source, without requiring high capacity data-links. Simulation results indicate that the parameters of two simultaneous LFMCW signals in complex additive white Gaussian noise can be estimated to within 10% of their true values with probability greater than 90% when input Signal to Noise Ratios (SNR) are -10 dB and above. In future chapters we evaluate the performance of the method on data collected in the laboratory using software defined radios and outdoors, in the presence of non-stationary noise with unknown autocorrelation function (*i.e.*, Chapter 3) and develop the Cramer-Rao Lower Bound (CRLB) of the parameters that define the multiple LFMCW signal model (*i.e.*, Chapter 2).

1.1 INTRODUCTION

The recent development of low cost radar systems that employ low power waveforms with long duration and large bandwidths were, in part, born out of the advent and proliferation of solid-state amplifier technology [37]. This technology enabled the production of small, lightweight radar systems that are capable of reliable detection (due to the signal's long duration), and precise ranging (due to the signal's wide bandwidth) at moderate standoff ranges. In many maritime environments, these radar systems are not only attractive in terms of size, cost, and performance, but also because of their potential to operate covertly in the presence of many current ELectronic INTelligence (ELINT) systems.

Since the majority of ELINT receivers are designed for high power, pulsed signals that were employed by most past threat radars [41–43], current ELINT systems suffer in the presence of low power, long duration, wide band signals since it was not until recently that these signals were feasible for practical radar applications [37, 56]. Consequently, these low power, wide band radars stand to replace high-power, pulsed radars of the past, which in turn prompts the need for modern ELINT techniques. The development of a set of techniques tailored to low power, Linear Frequency Modulate Continuous Wave (LFMCW) signals is the focus of this paper.[1]

[1] In the literature the terms LFMCW and Frequency Modulated Continuous Wave (FMCW) are sometimes used to refer to the same signal structure, while other times they refer to different signal structures. For example, sometimes FMCW refers to a signal whose instantaneous frequency has a periodic triangular structure. As we will see later when our LMFCW signal model is defined, our model assumes an instantaneous frequency that has a periodic *sawtooth* structure.

While numerous practical approaches exist for passively Detecting, Characterizing, and Localizing (DCL) high power, pulsed radar systems [56], relatively few exist for modern radar systems employing low-power LFMCW signals, in spite of their popularity [37, Ch. 2]. For this reason low power LFMCW signals are examples of radar waveforms that have become known as Low Probability of Intercept (LPI) signals. In Pace [37], LPI and Low Probability of Identification (LPID) are defined as:

Definition 1 *A low probability of intercept (LPI) radar is defined as a radar that uses a special emitted waveform intended to prevent a non-cooperative intercept receiver from intercepting and detecting its emission [37, pgs. 3,4].*

Definition 2 *Low probability of identification (LPID) radar is defined as a radar that uses a special emitted waveform intended to prevent a non-cooperative intercept receiver from intercepting and detecting its emission but if intercepted, makes identification of the emitted waveform modulation and its parameters difficult [37, pgs. 3,4].*

In Wiley [56], LPI radar is defined as:

Definition 3 *...a radar system whose signal is below the level of threshold of detection of opposing ELINT receivers while still being able to detect targets at useful ranges [56, pg. 211].*

The common thread between both Pace's and Wiley's definitions of LPI and LPID [2] radar is that the non-cooperative (*i.e.*, passive) intercept receiver in question is ill-suited to collect ELINT for the waveform employed. Given that the majority of commercial and military grade radars operate using high peak power, short duration, pulsed signals and, therefore, that most ELINT systems have been designed for those threats, a case can be made that modern signals are considered LPI because outdated ELINT receivers are being used to judge their intercept-ability and identifiability. In other words, attempts to use ELINT receivers designed for a different class of radar signals to detect and identify (*i.e.*, characterize) new threats yields results that have a *Low Probability* of success. Hence, new algorithms for detecting the presence and characterizing the structure of these signals are needed.

[2] For the remainder of this chapter we refer to the low-power, long-duration, LFMCW signals that we focus on as LPI, rather than, say, LPID.

All current methods for detecting, characterizing, and localizing LFMCW signals are afflicted by one or more of the following shortcomings:

1. computational requirements limit applicability
2. unable to handle case when more than one LFMCW is present
3. unable to automatically and efficiently determine number of LFMCW signals present
4. unable to set a meaningful detection threshold without knowledge of the statistics of the additive noise process corrupting the LFMCW signals
5. unable to operate at input Signal to Noise Ratio (SNR) levels characteristic of LPI radar signals observed at long ranges
6. unable to provide parameter estimates to enable localization, such as Time Difference of Arrival (TDOA), without access to high speed data links

The main contribution of this work is the development of a sequential method that, while suboptimal relative to the MLE, addresses each of these shortcomings. To summarize, the most salient features of our work are:

1. no calibration is necessary, either on-line or off-line in the laboratory, to determine detection threshold η settings as a function of Probability of False Alarm, P_{FA}
2. one does not need to assume the system noise present in the received time-series is white or has a known autocorrelation function
3. environments containing multiple LFMCW signals can be treated without adding a stifling computational burden
4. automatic recognition of the number of LFMCW signals present in the environment is achieved
5. in addition to estimates of the parameters that define the instantaneous frequency of each LFMCW signal, our approach produces estimates of the variances of a subset of these parameters directly from the data;
6. estimates of the TDOA associated with each LFMCW signal are obtained without the need for a high speed data-link to transmit raw signal samples to a central node.

In the next section we provide background on current methods to justify our statements regarding their shortcomings. In Section 1.1.3 we overview the sequential method developed in this chapter that is designed to exhibit the above features.

1.1.1 Background: Detection and Characterization

Whether they are designed for high power, pulsed signals or modern wide-band Continuous Wave (CW) signals, passive detection and characterization approaches can be broken into three categories: Noncoherent, Coherent, and Sequential. In the following three sections we describe each of these approaches. Section 1.1.1.1 provides background on noncoherent approaches, which requires our focus to be on legacy high power, pulsed signals, since they were the signals for which noncoherent methods were designed. Next, in Sections 1.1.1.2 and 1.1.1.3 our background discussion shifts to modern signals that utilize phase modulation to increase their bandwidths, with a particular emphasis on the LFMCW variety. In each of these three sections we assume that the signal in question is the complex envelope representation [21, pgs. 493-496] of the real-valued signal present at the output of the receiver's antenna. The reason for using this model of the data is that many Radio Frequency (RF) receivers designed for radar and communications applications utilize quadrature receiver architectures [52, pgs. 566-572], providing the complex envelope representation of the "real" RF signal. Doing so often leads to simplified analysis and processing, particularly for demodulation. We develop the complex envelope signal representation in Appendix A and relate it to the analytic signal representation [4, Ch. 2].

1.1.1.1 Noncoherent Approaches Noncoherent approaches remove signal phase information during processing, usually by considering only the magnitude of the signal, thus simplifying the resulting approaches and making them well suited for high power, pulsed threats. Some examples of noncoherent detection are integrated energy detectors and M of N detectors, [56, pg. 212-226], [44, Ch. 6], [22]. In practice, both of these approaches require one to assume that the noise is white and its power level is known and constant [44, pg. 347] over the collection duration in order to set a threshold that relates to the P_{FA}. Another approach, known as Constant False Alarm Rate (CFAR) detection [44, pgs. 347-382], does not require that the noise power is known or constant, but does require that it is white and Gaussian for a P_{FA} based threshold selection.

A vice and virtue of all noncoherent approaches is that they operate using detection statistics that are functions of energy alone, which makes them independent of any potential phase modulation[3]. This is a virtue because one does not need to assume anything about the structure of the transmitted signals, making the resulting approaches very general and well suited for the high power pulsed waveforms that typically contain no intentional phase modulation. The downside to this, which is especially pronounced when attempting to collect ELINT on LFMCW signals, is that in the absence of further processing one cannot ascertain information that characterizes the signal through its instantaneous frequency, since this information is typically embedded in the phase. Hence, the determination of the number of signals and a characterization of their full structure is impossible with noncoherent methods.

1.1.1.2 Coherent Approaches Coherent, in contrast to noncoherent approaches, use signal phase during processing. Since most systems that employ signals with sophisticated modulation schemes to make themselves LPI are based on Linear Frequency Modulation (LFM) [37, Ch. 1,2, and 4] [56, pg. 226], the most popular of the coherent approaches for LPI ELINT are based on the Wigner-Ville Transform (WVT). From the viewpoint of time-frequency analysis [4], the prominence of the WVT is owed to the fact that for a single LFM chirp the Wigner-Ville Distribution (WVD) is completely concentrated along the signal's instantaneous frequency [4, pg. 121], which is defined by a single straight line in the time-

[3]This is true since we are assuming a complex envelope representation of the signal.

frequency plane, and characterized by its slope (a/k/a chirp-rate) β and frequency intercept f^{int}. Since these are precisely the parameters that serve to characterize LFM signals, the WVT has become the most common among the many Time-Frequency Distributions (TFD's) that have been proposed to analyze LFM and LFMCW signals. In particular, the WVT of an LFM or LFMCW signal is commonly used as input to image processing techniques that are designed to detect edges or lines, such as the Radon or Hough Transforms. Unfortunately, these approaches tend to be computationally demanding and lack meaningful guidance for threshold selection.

From a detection and estimation theory viewpoint, the use of the WVD was first motivated by [25], wherein the Generalized Likelihood Ratio Test (GLRT) [22, pgs. 187-189] was derived to decide between the following hypotheses

$$\mathcal{H}_0: \quad r(t) = w(t) \tag{1.1}$$

$$\mathcal{H}_1: \quad r(t) = b \cdot s\left(t;\bar{\theta}\right) + w(t) \tag{1.2}$$

where $s\left(t;\bar{\theta}\right)$ is a complex, unit magnitude (*i.e.*, $\left|s\left(t;\bar{\theta}\right)\right| = 1$) deterministic signal with known structure but unknown parameters $\bar{\theta}$, $r(t)$ is the noisy received signal, and $b = b_R + jb_I$ is a zero mean complex Gaussian random variable with variance σ_b^2 that models the signal amplitude. Additionally, $w(t)$ is band-limited, zero mean, complex, white Gaussian noise with *known* variance. Under these assumptions, the GLRT statistic was found to be the maximum of the magnitude-squared of the 1-D correlation between the received signal $r(t)$ and all possible templates of the transmitted signal $s\left(t;\bar{\theta}\right)$. More specifically, the GLRT was shown to decide $\mathcal{H}_1$ if

$$\mathcal{L} = \arg\max_{\bar{\theta}} \left| \int_0^T r(t)\, s^*\left(t;\bar{\theta}\right) dt \right|^2 > \eta \tag{1.3}$$

In accordance with the theory associated with the GLRT, the Maximum Likelihood Estimate (MLE) of $\bar{\theta}$, denoted by $\hat{\bar{\theta}}$, is the value of $\bar{\theta}$ that maximizes (1.3) when the threshold η is exceeded. Unfortunately, while the detection and characterization scheme is optimal in the GLRT sense, no guidance is given for selecting η, which implies no control over P_{FA}. In other words, when the signal is actually present the detector will likely do a very good job of indicating such, but when the signal is not present false detections are beyond the

designer's control. The reason for the lack of guidance in selecting η is because there are no general results for the Probability Density Function (PDF) of $\mathcal{L}$ under $\mathcal{H}_0$ [22, pg. 187]. The main mathematical challenge is in deriving a PDF of a random variable that is defined through a max operation. However, even if an informative threshold could be set, when the dimension of $\bar{\theta}$ is large and the number of samples of $r(t)$ necessary for the full signal structure to be observable is high, a computationally practical implementation of (1.3) is typically impossible [9, 34]. More will be said on this topic in Section 1.5.

The WVD was brought into the solution found in [25] by invoking Moyal's formula [4, pgs. 128-129] to show that 1-D correlation in (1.3) is equivalent to a 2-D correlation between the WVD of the received signal $\mathcal{W}_r(t,\omega)$ and the WVD of $s\left(t;\bar{\theta}\right)$, denoted by $\mathcal{W}_s\left(t,\omega;\bar{\theta}\right)$. Specifically,

$$\left|\int_0^T r(t)\, s^*\left(t;\bar{\theta}\right) dt\right|^2 = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathcal{W}_r(t,\omega)\, \mathcal{W}_s\left(t,\omega;\bar{\theta}\right) dt \frac{d\omega}{2\pi} \tag{1.4}$$

where

$$\mathcal{W}_x(t,\omega) \triangleq \frac{1}{2\pi}\int x^*\left(t-\tfrac{1}{2}\tau\right) x\left(t+\tfrac{1}{2}\tau\right) e^{-j\tau\omega} d\tau \tag{1.5}$$

If $s\left(t;\bar{\theta}\right)$ is assumed to be a single LFM chirp, defined by

$$s\left(t;\bar{\theta}\right) = e^{j\left(2\pi f^{\text{int}} t + \pi\beta t^2\right)} \tag{1.6}$$

where $\bar{\theta} = \begin{bmatrix} f^{\text{int}} & \beta \end{bmatrix}^T$, then $\mathcal{W}_s\left(t,\omega;\bar{\theta}\right)$ reduces to

$$\mathcal{W}_s\left(t,\omega;\bar{\theta}\right) = 2\pi\delta\left(\omega - 2\pi f^{\text{int}} - 2\pi\beta t\right) \tag{1.7}$$

where $\delta(\cdot)$ is the Dirac delta function. Substituting (1.4) and (1.7) into (1.3) yields

$$\mathcal{L} = \underset{f^{\text{int}},\beta}{\arg\max} \int_{-\infty}^{\infty} \mathcal{W}_r\left(t, 2\pi f^{\text{int}} + 2\pi\beta t\right) dt > \eta \tag{1.8}$$

Hence the GLRT for a single LFM chirp is equivalent to path integration of the WVD along paths determined by different chirp-rates and frequency intercepts. If the value of any such path integral, which is equivalent to a point in the Hough or Radon Transform domain of the WVD, exceeds the detection threshold η, then a single LFM signal is detected with the MLE of its parameters being the $\hat{f}^{\text{int}}$ and $\hat{\beta}$ that defined the maximum path integral exceeding

the threshold. Unfortunately, even for this special case, no direction on how to select η was given, which suggests one's only recourse is to either assume a PDF for $r(t)$ under $\mathcal{H}_0$, estimate it on-line with built-in testing equipment, or estimate it off-line with laboratory experiments. Each of these options is potentially costly and error prone but we speculate that it must be common-practice given the absence of any alternatives in the literature.

In [9] the authors show that if one assumes that $r\left(t;\bar{\theta}\right)$ is the sum of multiple LFMCW signals with deterministic but unknown amplitudes, the resulting GLRT is equivalent to that derived in [25] and given above in (1.3). They also demonstrate the relationship between the GLRT and a form of the WVD. Their model for the received signal under $\mathcal{H}_0$ is the same as in [25] but $\mathcal{H}_1$ becomes

$$\mathcal{H}_1: \quad r(t) = s\left(t;\bar{\theta}\right) + w(t) \tag{1.9}$$

where

$$s\left(t;\bar{\theta}\right) = \sum_{m=1}^{M} s_m\left(t;\bar{\theta}_m\right) \tag{1.10}$$

$$s_m\left(t;\bar{\theta}_m\right) = b_m e^{j\Phi_m(t)} \tag{1.11}$$

$$\Phi_m(t) = \varphi_m + 2\pi f_m^{\text{start}} t + \pi\beta_m \bmod \left(t+\Delta_m, T_{\text{swp},m}\right)^2 \tag{1.12}$$

$$\bar{\theta} = \left[\bar{\theta}_1^T \quad \ldots \quad \bar{\theta}_M^T\right]^T \tag{1.13}$$

$$\bar{\theta}_m = \left[f_m^{\text{start}} \quad \beta_m \quad \Delta_m \quad T_{\text{swp},m}\right]^T \tag{1.14}$$

where each of the parameters in this model is defined in Table 2. In order to reduce the dimensionality of their estimation problem they exclude the amplitude b_m and phase φ_m from (1.14) and circumvent the need to estimate them by treating them as nuisance parameters. Figure 1 shows the spectrogram of a received signal containing two LFMCW components, each observed with high input Signal to Noise Ratio (SNR), and of the form used in [9] and summarized in Equations (1.9)-(1.14). For the purposes of this chapter we define the input SNR of the m^{th} LFMCW signal as

$$\text{SNR}_m = \frac{b_m^2}{\sigma_w^2} \tag{1.15}$$

This example signal, with slight modifications to reflect the more general signal model we consider, is used throughout the paper to highlight various aspects of the sequential detection and characterization process.

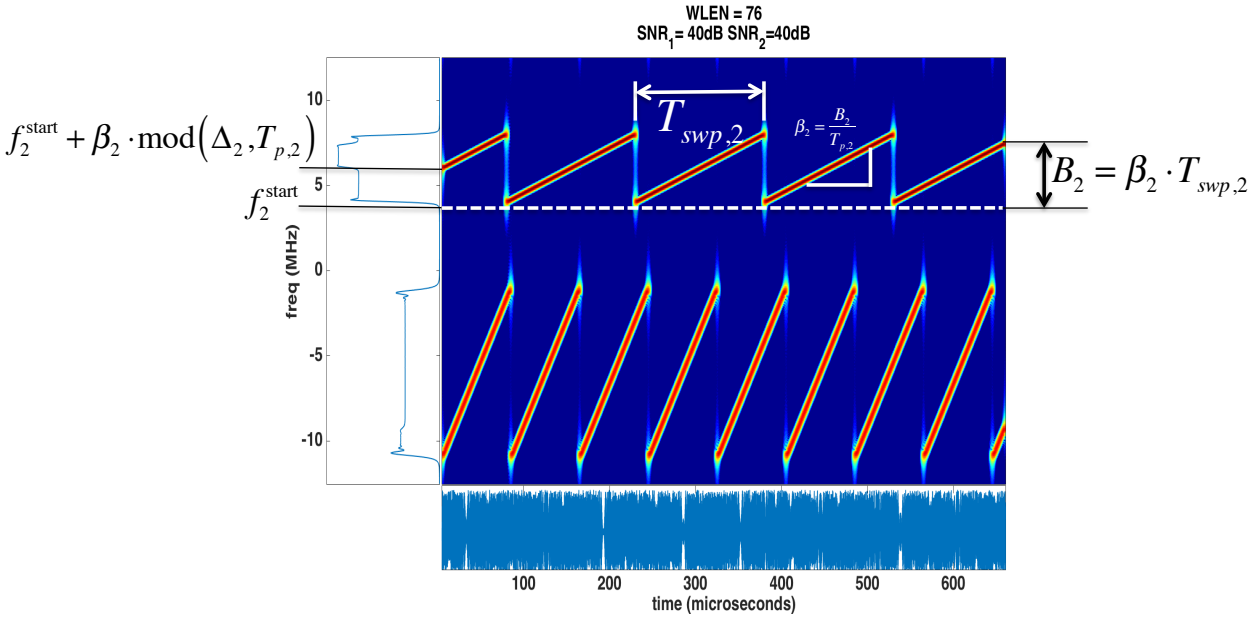

Figure 1: (main panel) Spectrogram (20 dB dynamic range) of two LFMCW signals signals based on signal model given in (1.11) by way of [9] (left panel) Frequency marginal of spectrogram (bottom panel) Real part of time-series $r(t)$

While the characteristically high SNR at the output of a GLRT detector is desirable since it yields high Probability of Detection (P_D), its practicality is limited since the detection statistic admits no systematic method for determining η. Furthermore, the authors of [9] make it clear that for a practical (*i.e.*, computationally feasible) implementation of their detector for the simplest case when $M = 1$, *a priori* knowledge about *each* dimension of $\bar{\theta}$ is required to define a search space. Additionally, limits on the number of grid points used to define each dimension of the search space must be imposed if computational tractability is to be achieved, which will ultimately limit the accuracy of the estimator. This, however, is an issue that afflicts all grid based techniques. Furthermore, if $M > 1$, or even worse, if M is unknown, the resulting computational burden makes their optimal estimator impractical. Finally, careful consideration of (1.11) reveals that their signal model forces the entire ob-

Table 1: Signal Parameters for Figure 1

Variable Description	Variable Name	Values (m=1, m=2)	Units
Amplitude	b_m	1,1	$\sqrt{J/s}$
Phase	φ_m	0,0	rad
Starting Frequency	f_m^{start}	-11,4	MHz
Chirp-Rate	β_m	0.125,0.0267	MHz/μs
Time-Offset	Δ_m	0,75	μs
Sweep Time	$T_{\text{swp,m}}$	80,150	μs
No. of LFMCW Signals	M	2	—
Sampling Rate	F_s	25	MHz
Observation Duration	T	655.36	μs

servation interval to be occupied by LFMCW energy, which is a limitation for two reasons. First, a passive receiver cannot control when the emitter beings to transmit relative to when the receiver begins to record. A way around this might be to limit $r(t)$ so that it adheres to (1.11), but such a procedure runs the risk of losing valuable signal information, such as Time of Arrival (TOA). Second, without a parameter that models delay relative to $t = 0$, the TOA of each signal may be unobservable. Without the capacity to estimate TOA, multi-platform localization, which we discuss in Section 1.1.2, based on Time Difference of Arrival (TDOA) is impossible.

1.1.1.3 Sequential Approaches A third category of methods, which can be thought of as "hybrid" approaches, break the process of detecting and characterizing LFMCW signals into a sequence of subproblems, each leading to a tailored solution. The first example of such a solution that is tailored to estimating the parameters of (1.12) is found in [34], where the author develops a method to sequentially estimate each of the parameters in (1.14) for

the case when $M = 1$. The method is shown in simulation to perform well, by comparison to the MLE, for SNR≥ 0 with a markedly lower computational expense. However, because of the tools used in the sequential method proposed, the approach is limited to the single signal case. Furthermore, the author focuses completely on the parameter estimation problem (*i.e.*, characterization) and therefore gives no guidance on how to select a threshold that indicates the presence of an LFMCW signal. In other words, the detection problem is not considered.

The approach we develop in this chapter, which is based upon our previous work [16,17], most appropriately falls in the category of *Sequential Approaches*. Our method differs from [34] by the subproblems chosen, the sequence in which they are addressed, and the tools that are applied to them. These differences account for the ability of the method to handle a problem formulation that allows M to be estimated, but at an increased computational expense compared to [34].

1.1.2 Background: Localization

In addition to detecting and characterizing LPI signals, ELINT systems are often tasked with estimating the location of the sources that transmit them. In passive localization applications [11], the three primary measurement types used for estimating source location are Direction of Arrival (DOA) [13], Frequency Difference of Arrival (FDOA), and Time Difference of Arrival (TDOA).

The primary advantage of DOA based localization systems is that a single moving ELINT receiver can estimate source locations [13]. However, this method often requires the use of an antenna array and accompanying receivers that are not only costly, but also necessitate time-consuming calibration in order to produce accurate DOA estimates. In contrast to DOA based methods, TDOA and FDOA methods require less sophisticated receiver systems and as few as one antenna element per receiver. Some level of calibration is required of the TDOA/FDOA measurement generation process, but it is typically much less involved and, therefore, less time-consuming and costly than that required for DOA antenna arrays. However, TDOA/FDOA systems require a network of at least two spatially separated receivers, synchronized in time, to operate collaboratively to produce TDOA/FDOA estimates. On

its own, this multi-node requirement limits the applicability of TDOA/FDOA solutions in many scenarios. Further compounding the practicality of classical TDOA/FDOA methods is the need to transmit signal samples to a central node so that methods based on cross-correlation [48] can be implemented to produce estimates of TDOA and FDOA. In scenarios where wired links between nodes are unavailable, this requirement becomes even more limiting since wireless data-links often have insufficient bandwidth to handle timely transmission of the quantity of data required to form accurate TDOA/FDOA estimates [48].

1.1.3 Overview of This Work

This section is devoted to summarizing the method developed in this chapter by putting it into the context of some closely related past work.

While the detection problem framed in (1.1) and (1.2) and the time series defined by (1.9) - (1.14) are necessary starting points to pose the detection and characterization problems, the complexity of the signal model (*i.e.*, the large number of parameters necessary to define it) and the number of samples required for a LFMCW signal to show its full structure in a received time series makes the optimal correlation-based procedures, like that given in (1.3), doomed to computational limitations. To depart from this course we adopted an approach rooted in the field of *spectral estimation* [40, 49] that attempts to extend methods developed for stationary processes to processes with time-varying spectral content. In particular, our approach breaks the signal down by analyzing the received time-series over short, non-overlapping, contiguous time intervals. Within each interval the received signal is modeled as a sum of harmonic components with unknown order and, within this interval, zero or more frequency components are detected and estimated. Section 1.2 develops this simplified signal model and the associated detection and estimation procedure. Over time, the frequency estimates are collected and used to produce estimates of the parameters that define a more general model of the received signal than that given in (1.9)-(1.14). The methods used to process these frequency samples are described in Section 1.3.

The increased generality comes by way of two modifications to the signal model. First, we include two temporal parameters that allow for the specification of both start-time, t_m^{start},

and stop-time, t_m^{stop}. Second, we no longer restrict $w(t)$ to be white, stationary, or have completely known autocorrelation properties. Mathematically, this means that the method can accommodate $w(t)$ such that

$$E\left\{w(t_1)\, w^*(t_2)\right\} = R(t_1, t_2) \tag{1.16}$$

without assuming $R(t_1, t_2) = R(t_1 - t_2) = R(\tau)$ or precise knowledge of the autocorrelation function $R(t_1, t_2)$. One does need to assume that $w(t)$ is Gaussian, locally wide-sense stationary over short time intervals [8], and that the Power Spectral Density (PSD) associated with each short time interval does not change abruptly in frequency intervals of width $2W$. Both the width of a short time interval and the resolution bandwidth $2W$ are defined in Section 1.2. For the purposes of this dissertation, noise processes that do not change abruptly on over frequency intervals of $2W$ are referred to as "slowly-varying" [40, pgs. 370, 498]. However, if any or all of these assumptions do not hold for some short time intervals, we believe the overall method can recover[4]. The modifications to (1.9)-(1.14) are summarized as follows, where again

$$\mathcal{H}_1: \quad r(t) = s\left(t; \bar{\theta}\right) + w(t) \tag{1.17}$$

and

$$s\left(t, \bar{\theta}\right) = \sum_{m=1}^{M} s_m\left(t; \bar{\theta}_m\right) \tag{1.18}$$

$$\bar{\theta} = \begin{bmatrix} \bar{\theta}_1^T & \dots & \bar{\theta}_M^T \end{bmatrix}^T \tag{1.19}$$

but now

$$s_m\left(t; \bar{\theta}_m\right) = \Lambda_m\left(t; \bar{\theta}_m\right) \cdot b_m e^{j\Phi\left(t; \bar{\theta}_m\right)} \tag{1.20}$$

where

$$\Lambda_m\left(t; \bar{\theta}_m\right) = u\left(t - t_m^{\text{start}}\right) - u\left(t - t_m^{\text{stop}}\right) \tag{1.21}$$

$$\Phi\left(t; \bar{\theta}_m\right) = \varphi_m + 2\pi \int_0^t f_m(\nu)\, d\nu \tag{1.22}$$

$$f_m(t) = f_m^{\text{start}} + \beta_m \bmod \left(t + \Delta_m, T_{\text{swp},m}\right) \tag{1.23}$$

[4]We believe that the assumptions for $w(t)$ to have a Gaussian distribution and its power spectrum to be slowly-varying can be broken over short, infrequent periods of time. We describe the theoretical basis for this belief in Section 1.2 and test it in Chapter 3.

and $u(t)$ is the unit-step function. In contrast to the definition of the instantaneous phase of the m^{th} signal given in (1.12), we define $\Phi\left(t, \bar{\theta}_m\right)$ in (1.22) as a function of the integrated instantaneous frequency defined by (1.23). The reason for this is that our detection and characterization procedures are more naturally suited to a definition of instantaneous frequency than a definition of the instantaneous phase. Consequently, we avoid the mathematical difficulties of expressing the instantaneous frequency as the time derivative of the instantaneous phase, which is undefined since the $\text{mod}(\cdot,\cdot)$ operation used in (1.23) is not differentiable. Finally, the unknown parameter vector is modified to include b_m, φ_m, t_m^{start}, and t_m^{stop}

$$\bar{\theta}_m = \begin{bmatrix} b_m \\ \varphi_m \\ t_m^{\text{start}} \\ t_m^{\text{stop}} \\ f_m^{\text{start}} \\ \beta_m \\ \Delta_m \\ T_{\text{swp},m} \end{bmatrix} \tag{1.24}$$

Since the signal model given in (1.17)-(1.24) is a generalization of that given in (1.9)-(1.14), Figure 1 and the parameters given in Table 1 serve as a concrete example of our signal model with $t_m^{\text{start}} = 0$ and $t_m^{\text{stop}} = T$.

The remainder of this chapter is devoted to developing our approach theoretically and analyzing its performance via simulation. To do so it is organized as follows. Sections 1.2-1.4 cover the algorithmic details associated with each part of the DCL process. In particular, the detection and characterization algorithms are detailed in Sections 1.2 and 1.3. To describe the method we carry an example through these sections that is based on the signal shown in Figure 1 and Table 1. The noisy signal used in these examples is at high SNR (*i.e.*, 0 dB) in order to highlight the algorithm's operation. The localization concept is described in Section 1.4 along with the relevant TDOA measurement model. In Section 1.5 we provide mean and Mean Squared Error (MSE) estimates derived from Monte Carlo experiments that compare our sequential method to the MLE for $M = 1$. For the $M = 2$ case, we show similar

performance metrics but omit a comparison to the MLE because implementing it in a Monte Carlo simulation was too time consuming. In contrast to the example signal used to develop the theory in Sections 1.2 and 1.3, we reduce the SNR until the algorithm breaks down to determine its limitations. We found that the breakdown point occurred at -10 dB SNR. Finally, we summarize the strengths and weaknesses of our approach in Section 1.6 and give suggestions for future work in Section 1.7.

1.2 DETECTION

In this section we describe the *detection* step of the algorithm. The term 'detection' is used to to refer to this step since it is where we formulate the first hypothesis test and where control over P_{FA} is obtained through our choice of a detection threshold. To be clear, in our problem formulation a false alarm is not the incorrect declaration of the presence of one or more LFMCW signals, as it would be if we were deciding between (1.1) and (1.2). Instead it is the incorrect declaration of the presence of one or more complex exponentials in a short-time interval. A more precise definition of what is meant by a short-time interval is given below.

Consider the following hypothesis test [5]

$$\mathcal{H}_0 : \; r_{n,k} = w_{n,k} \tag{1.25}$$

$$\mathcal{H}_1 : \; r_{n,k} = s_{n,k} + w_{n,k} \tag{1.26}$$

where $r_{n,k}$ is the n^{th} sample, for $n = 1, \ldots, N$, of $r(t)$ in the k^{th} time interval. The k^{th} time interval is given by

$$\delta_k = \{t : (k-1) \cdot NT_s \leq t \leq k \cdot NT_s)\} \tag{1.27}$$

where $k = 1, \ldots, K = \lfloor \frac{T}{N \cdot T_s} \rfloor$. That is,

$$r_{n,k} = r\left((k-1) \cdot NT_s + nT_s\right) \tag{1.28}$$

Within this short time interval we model the received time-series as a sum of a small number harmonic components, which gives rise to the following Short Time Harmonic Model (STHM) of the data

$$s_{n,k} \approx s_{n,k}^{\text{STHM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi f_{c,k} n T_s} \tag{1.29}$$

where $L_k \leq M$ is the number of constant frequencies $f_{c,k}$ present in δ_k. The complex amplitude is given by $C_{c,k} = b_{c,k} e^{j\phi_{c,k}}$ where $b_{c,k}$ is as defined in (1.20) and $\phi_{c,k}$ is the phase of the c^{th} harmonic component in the k^{th} short-time interval. We emphasize that L_k and each $b_{c,k}$, $f_{c,k}$, and $\phi_{c,k}$ are assumed to be unknown, but deterministic. From this point forward

[5] Samples of a continuous time signal $x(t)$ are denoted by $x_n \equiv x(t)\,|_{t=nT_s}$

we take the approximation in (1.29) to be an equality. Furthermore, in δ_k we assume that $w_{n,k}$ is a wide sense stationary, band-limited, zero mean, complex Gaussian random process with unknown PSD $P_{w_{n,k}}(f)$ that is not necessarily constant (*i.e.*, $\{w_{n,k}\}$ can be colored noise), but is slowly-varying. [6] The ability to handle an unknown colored noise process is what justifies Item 2 from the features list given in Section 1.1, and is a key feature for a practical implementation since the PSD of the noise is rarely known *a-priori*. This feature is a byproduct of Thomson's method for harmonic analysis, which we apply to (1.28) to decide between (1.25) and (1.26), and results because no further assumptions on the noise autocorrelation function are required in the method's development. Making few assumptions on the structure of the autocorrelation function is common in the field of spectral analysis, from which this method is drawn.

The STHM of the received multicomponent LFMCW signal in δ_k is

$$r_{n,k} = s_{n,k}^{\text{STHM}} + w_{n,k} \tag{1.30}$$

So, within δ_k selection between $\mathcal{H}_0$ and $\mathcal{H}_1$ can be achieved using Thomson's multi-taper method for harmonic analysis, which was first proposed in [51] and later summarized in [40, pgs. 331-374, 496-514]. In what follows we overview the method by following [40, pgs. 331-374, 496-514] closely and highlighting its main features in the context of our LPI DCL problem.

[6] The assumption that the PSD is slowly-varying is a mathematical necessity for the development of the test statistic outline in Section 1.2.1. However, in practice we only need the PSD to not exhibit dramatic variations (*i.e.*, be slowly-varying) in most of the short-time intervals since a small percentage of missed detections or erroneous detections that are a result of instability in the detector can be tolerated by the later steps devoted to characterization outlined in Section 1.3.

1.2.1 Steps for Detecting and Estimating STHM Parameters via Thomson's Method

Under $\mathcal{H}_1$, $\{r_{n,k}\}$ is a set of samples with a sinusoidally varying mean. Accordingly, the goal of the decision procedure designed to decide between (1.25) and (1.26) is to determine whether the mean is statistically different from zero and consistent with the harmonic model of the mean given in (1.29). The key idea behind making this determination is to express (1.30) in the frequency domain and, upon application of Thomson's multi-taper method, recognize the regression problem that results. The decision procedure then amounts to determining the statistical significance of the computed regression coefficients, which leads to estimates of $b_{c,k}$, $f_{c,k}$ and, $\phi_{c,k}$. For the sake of brevity the detailed development of Thomson's method is omitted from this chapter, but the reader can consult [40, pgs. 331-374, 496-514] for an excellent description of the method for harmonic analysis of real-valued signals or Appendix B for the derivation for the complex envelope data model adopted in this chapter. The following steps amount to an extension of the approach given in [40, pgs. 331-374, 496-514] to apply to the complex envelope model adopted in this chapter.

We now give a summary of the steps necessary to test whether the samples in δ_k contain significant frequencies, determine estimates of their values, and compute estimates of L_k, $b_{c,k}$, $f_{c,k}$ and $\phi_{c,k}$.

Step 0 – For $q \in \{0 \ldots Q-1\}$, with $Q = 2NW - 1$, compute the associated length N DPSS data taper[7] $\{h_{q,n}\}$ [40, Ch. 8] where N is the number of samples taken in δ_k at the sample rate F_s. The resolution bandwidth W, for W in the normalized frequency range $0 \leq W \leq 0.5$, determines the minimum separation between frequencies in the same δ_k that can be resolved. In other words, distinct frequencies $f_{1,k}$ and $f_{2,k}$ will be indistinguishable if $|f_{1,k} - f_{2,k}| / F_s < 2W$. Assuming N and W are fixed for all δ_k, this step only needs to be performed once and the tapers applied to each δ_k.

Step 1 – Let $J_{q,k}(f)$ be the Fourier Transform of $\{\sqrt{T_s} \cdot h_{q,n} \cdot r_{n,k}\}$. Then, for $p = -\frac{N_{FFT}}{2}, \ldots \frac{N_{FFT}}{2} - 1$, compute samples of $J_{q,k}(f)$ using an N_{FFT} point Fast Fourier

[7] The MATLAB© Signal Processing Toolbox [30] provides a routine that computes the required Q tapers. In particular, `h = dpss(N,NW,Q)` is a matrix of size NxQ whose q^{th} column is $\{h_{q,n}\}$.

Transform (FFT). Specifically, let $J_{q,k}(p \cdot \Delta f)$, where $\Delta f \equiv F_s/N_{FFT}$, be the p^{th} point of the FFT of $\{\sqrt{T_s} \cdot h_{q,n} \cdot r_{n,k}\}$.

Step 2 – For $q \in \{0, \ldots, Q-1\}$, compute $H_q(0) = T_s \sum_{n=1}^{N} h_{q,n}$

Step 3 – For $p = -\frac{N_{FFT}}{2}, \ldots, \frac{N_{FFT}}{2} - 1$, compute $C(p \cdot \Delta f) \equiv \sqrt{T_s} \cdot \frac{\sum_{q=0}^{Q-1} J_{q,k}(p \cdot \Delta f) H_q(0)}{\sum_{q=0}^{Q-1} H_q^2(0)}$

Step 4 – For each $q \in \{0, \ldots, Q-1\}$ and each $p = -\frac{N_{FFT}}{2}, \ldots, \frac{N_{FFT}}{2} - 1$, compute $\hat{J}_{q,k}(p \cdot \Delta f) \equiv C(p \cdot \Delta f) \frac{H_q(0)}{\sqrt{T_s}}$

Step 5 – For $p = -\frac{N_{FFT}}{2}, \ldots, \frac{N_{FFT}}{2} - 1$ compute $D_k(p \cdot \Delta f) = \frac{(Q-1)|C(p \cdot \Delta f)|^2 \sum_{q=0}^{Q-1} H_q^2(0)}{T_s \sum_{q=0}^{Q-1} \left| J_q(p \cdot \Delta f) - \hat{J}_q(p \cdot \Delta f) \right|^2}$

Step 6 – For a given P_{FA}, compute the detection threshold η by $\eta = \frac{(Q-1)\left(1 - {P_{FA}}^{1/(Q-1)}\right)}{P_{FA}^{1/(Q-1)}}$, since under $\mathcal{H}_0$, $D_k(p \cdot \Delta f) \sim F_{2,2Q-2}$ [40, pg. 501], where $F_{2,2Q-2}$ is an F-distribution with 2 and $2Q-2$ degrees-of-freedom.

Step 7 – Select the indices, denoted by p_c, corresponding to the largest values of $\{D_{p,k}\}$ such that $D_k(p \cdot \Delta f) \geq \eta$ and $|p_i - p_j| > \left\lceil \frac{2W \cdot F_S}{\Delta f} \right\rceil$. The number of peaks meeting these criterion is the estimate of L_k, with associated frequency estimates given by $\hat{f}_{c,k} = p_c \Delta f$. Amplitude and phase estimates are given by $\hat{b}_{c,k} = \left| \hat{C}_{c,k} \right|$ and $\hat{\phi}_{c,k} = 2 \cdot \tan^{-1}\left(\frac{\mathrm{Im}\{\hat{C}_{c,k}\}}{\mathrm{Re}\{\hat{C}_{c,k}\}} \right)$, respectively.

Following Step 6 we have a (potentially empty) set of ordered pairs

$$\mathcal{S}_k = \left\{ \left(t_k, \hat{f}_{1,k} \right), \ldots, \left(t_k, \hat{f}_{L_k,k} \right) \right\} \tag{1.31}$$

corresponding to frequency estimates made at time t_k, where t_k is defined to be the sample time closest to the center of δ_k. Hence, the P_{FA} specified controls the probability of detecting the presence of a harmonic component when one is not actually present. This way of specifying P_{FA} leads to a different interpretation than if P_{FA} were specified in the context of (1.1) and (1.2), but doing so still provides the designer control over incorrect decisions related to the detection of one or more LFMCW signals. Implementing this sequence of steps for all k produces the set $\mathcal{S} = \{\mathcal{S}_1, \ldots, \mathcal{S}_K\}$ containing all frequency estimates obtained in the time interval $(0, T)$. The set $\mathcal{S}$ is the main input to the Characterization step, which is described in Section 1.3.

If no significant frequency components exist in any of the time intervals δ_k then the DCL process is terminated and declares no LFMCW signals are present in the environment.

1.2.2 Selecting N and W

In this section we provide guidance for selecting N and W. To do so it is instructive to summarize the trade-offs one faces as N and W are varied.

1. **Pro**: As N grows detection performance increases as long as the STHM model given in (1.30) holds [22, pgs. 484,485]
2. **Con**: As N grows the less the samples of (1.17) will adhere to (1.30) since the *true* instantaneous frequency of the signal (1.23) is increasing or decreasing linearly with time
3. **Pro**: As W increases the number of usable tapers $Q = 2NW - 1$ increases, which yields an estimate of $C_{c,k}$ with lower variance [40, pg. 499]
4. **Con**: As W increases the minimum separation between frequencies in same δ_k that can be resolved increases [40, pg. 335], as does the computational effort required to implement **Steps 0-7**

Based on items 1 and 2 above, it becomes evident that we would like to find the largest possible N such that the STHM holds. To find such an N we analyze the spectrogram of a single LFM chirp. The spectrogram, which we have already utilized in Figure 1 to visualize the time-frequency content of two simultaneous LFMCW signals, is a general analysis tool that can be used to uncover properties of signals whose frequency content varies with time. The spectrogram accomplishes this goal by breaking the signal into small time segments, applying a window that is centered on each segment and falls off rapidly outside it, then applies a Fourier Transform. The time ordered collection of each of these spectra is called the Short-Time Fourier Transform, and its magnitude-squared is referred to as the spectrogram [4, Ch. 7]. Given its nature, it is well suited to draw conclusions about our STHM. In particular, to determine the relationship between N and β, such that N is as large as possible without deviating too much from a constant frequency, we model a single segment

of an LFMCW signal as

$$s(t) = \left(\frac{1}{2\sigma_s^2\pi}\right)^{1/4} e^{-t^2/\left(4\sigma_s^2\right)+j\beta t^2/2+j\omega_0 t} \tag{1.32}$$

which is an LFM chirp with a Gaussian amplitude. Next, suppose that the window used to construct the spectrogram is also Gaussian such that

$$h(t) = \left(\frac{1}{2\sigma_h^2\pi}\right)^{1/4} e^{-\frac{t^2}{2\sigma_h^2}} \tag{1.33}$$

where σ_s and σ_h are related to the temporal duration of $s(t)$ and $h(t)$, respectively. Then, it can be shown [4, pgs. 104-105] that the spectral variance of $s(t)$ conditioned on t, computed with a spectrogram having window defined by $h(t)$ is

$$\sigma_{\omega|t}^2 = \frac{1}{2}\left(\frac{\sigma_s^2+\sigma_h^2}{2\sigma_s^2\sigma_h^2} + \frac{2\beta^2\sigma_s^2\sigma_h^2}{\sigma_s^2+\sigma_h^2}\right) \tag{1.34}$$

This expression quantifies the amount of spectral deviation that occurs at a given instant in time for a fixed β as a function of the temporal duration of the chirp and the window. Hence, it can be used to determine how to control spectral spread based on the window chosen [5]. So, if we assume β and σ_s^2 are fixed, minimizing $\sigma_{\omega|t}^2$ with respect to σ_h^2 gives us an indication of the maximum amount of time that can elapse while keeping the deviation of frequency about the conditional mean small. The necessary condition $\frac{\partial\sigma_{\omega|t}^2}{\partial\sigma_h^2} = 0$ yields [8]

$$\frac{\sigma_h^2\sigma_s^2}{\sigma_h^2+\sigma_s^2} = \frac{1}{2|\beta|} \Rightarrow \sigma_h \approx \frac{1}{\sqrt{2|\beta|}}, \text{ for } \quad \sigma_h^2 \ll \sigma_s^2 \tag{1.35}$$

Requiring that $\sigma_h^2 \ll \sigma_s^2$ is equivalent to requiring that the observation duration of an entire chirp segment is much greater than the duration of a short-time interval δ_k. This is true by definition δ_k. So, a reasonable starting point for selecting an upper bound on N is to let

$$N_{max} = \lceil 2\sigma_h F_s \rceil = \left\lceil \frac{2F_s}{\sqrt{2\,|\beta_{max}|}} \right\rceil \tag{1.36}$$

where $|\beta_{max}|$ is the maximum possible absolute chirp-rate expected and $\lceil\cdot\rceil$ is the ceiling operator.

[8] It can also be shown that the second derivative of $\sigma_{\omega|t}^2$ with respect to σ_h evaluated at $\sigma_h = \frac{1}{\sqrt{2|\beta|}}$ is greater than zero, which guarantees that our solution is a local minimum as desired.

With a value for N chosen as given above, selecting NW requires one choose W, which fixes Q. From Items 3 and 4 above, the main trade-off faced by the selection of W is a common one faced in spectral analysis: *spectral estimator variance vs. resolution*. In general spectral estimation, where one applies a spectral estimator in order to uncover the structure of the signal's power distribution over frequency, often little can (and should) be assumed about that structure *a-priori*. However, our application of Thomson's Method as a spectral estimator is tailored for a specific model of the data given through (1.29). Furthermore, in practice it is possible that multiple LFMCW signals are present in the same operating band, like those given in the example shown in Figure 1, and possibly even have instantaneous frequencies that cross in the time-frequency plane. However, it is unlikely that multiple signals overlap for a significant percentage of time, such that their instantaneous frequencies are within $\pm WF_S$ of one another. To illustrate this, consider the example of the signal given by (1.18) and defined in Table 1. If we assume that the maximum possible chirp-rate is $1.8\beta_1$ then $N_{max} = 76$ and a selection of $NW = 15$, which is very large compared to more typical selections [40, pg. 335] of $NW = 2, 3, \text{or } 4$, yields a resolution half-bandwidth of only $W \cdot F_s = 0.023$ MHz. It is reasonable to assume that no other LFMCW signals will exist within ± 0.023 MHz of one another at a given instant of time, and even less likely to be the case over long periods of time. If this were not the case the radar systems employing these waveforms would likely interfere with one another and, if possible, at lease one of the radars would modify its signal properties to avoid the interference. While NW is not limited in the same way it is when multitaper methods are applied to general spectra, NW cannot be chosen arbitrarily large though since the more tapers that are used the more computation that is required, as noted in Item 4 at the beginning of this subsection.

1.2.3 Detection Example

In this section we give an example of detection based on the STHM. The signal used is very close to that given in Figure 1 and Table 1 except that noise is added with $\sigma_\omega^2 = 1$ so that $\text{SNR}_1 = \text{SNR}_2 = 0$ dB. Additionally, t_m^{start} and t_m^{stop} parameters are included. The specific parameters of the example signal used throughout the rest of this chapter are given in Table 2. The particular parameters required of Steps 0-7 are summarized in Table 3.

Table 2: Signal Parameters for Detection and Characterization Examples

Variable Description	Variable Name	Values (m=1, m=2)	Units
Amplitude	b_m	1,1	$\sqrt{J/s}$
Phase	φ_m	0,0	rad
Start Time	t_m^{start}	10,20	μs
Stop Time	t_m^{start}	650,620	μs
Starting Frequency	f_m^{start}	-11,4	MHz
Chirp-Rate	β_m	0.125,0.0267	MHz/μs
Time-Offset	Δ_m	10,20	μs
Sweep Time	$T_{\text{swp,m}}$	80,150	μs
No. of LFMCW Signals	M	2	—
Sampling Rate	F_s	25	MHz
Observation Duration	T	655	μs

The top panel of Figure 2(a) shows the whole time-series $\{r_n\}$ in blue and $\{r_{n,15}\}$ overlaid in red. A closer look at $\{r_{n,15}\}$, which we assume adheres to the STHM given in (1.29), is shown in the lower panel of Figure 2(a). Figure 2(b) shows the detection statistic that results from Steps 0-5, plotted as a function of frequency, and the associated detection threshold computed based on Step 6. Following execution of Step 7, we arrive at the following set of

frequency estimates for δ_{15} and $t_{15} = 44.08\ \mu_s$

$$\mathcal{S}_{15} = \{\,(44.08\ \mu_s, -6.753\ \text{MHz})\,, \\ (44.08\ \mu_s, 4.639\ \text{MHz})\} \tag{1.37}$$

which are plotted, along with truth at t_{15} on the spectrogram given in the main panel of Figure 2(c). For comparison, the true instantaneous frequency at t_{15} for $s_1\left(t;\bar{\theta}_1\right)$ is -6.740 MHz and 4.642 MHz for $s_2\left(t;\bar{\theta}_2\right)$.

Table 3: Implementation Parameters for Detection Example

Variable Description	Variable Name	Value	Units
FFT Length	N_{FFT}	2^{14}	—
SHTM Window Length	N	76	—
Resolution Bandwidth	$W \cdot F_s$	0.023	MHz
Number of Tapers	Q	29	—
False Alarm Probability Per Short Time Interval	P_{FA}	10^{-4}	—

Performing Steps 1-7 for each k yields the set $\mathcal{S}$, which is plotted in Figure 2(d), along with truth. A close look at Figure 2(d) reveals that the estimates associated with the instantaneous frequency of each LFMCW signal appear to be very close to their true instantaneous frequencies. While visual confirmation is satisfying, any firm conclusion on the estimation accuracy is deferred to Section 1.5, wherein we determine how well these frequency estimates enable the estimation of each element of each $\bar{\theta}_m$. The set $\mathcal{S}$ is the main input to the Characterization step, which is described next.

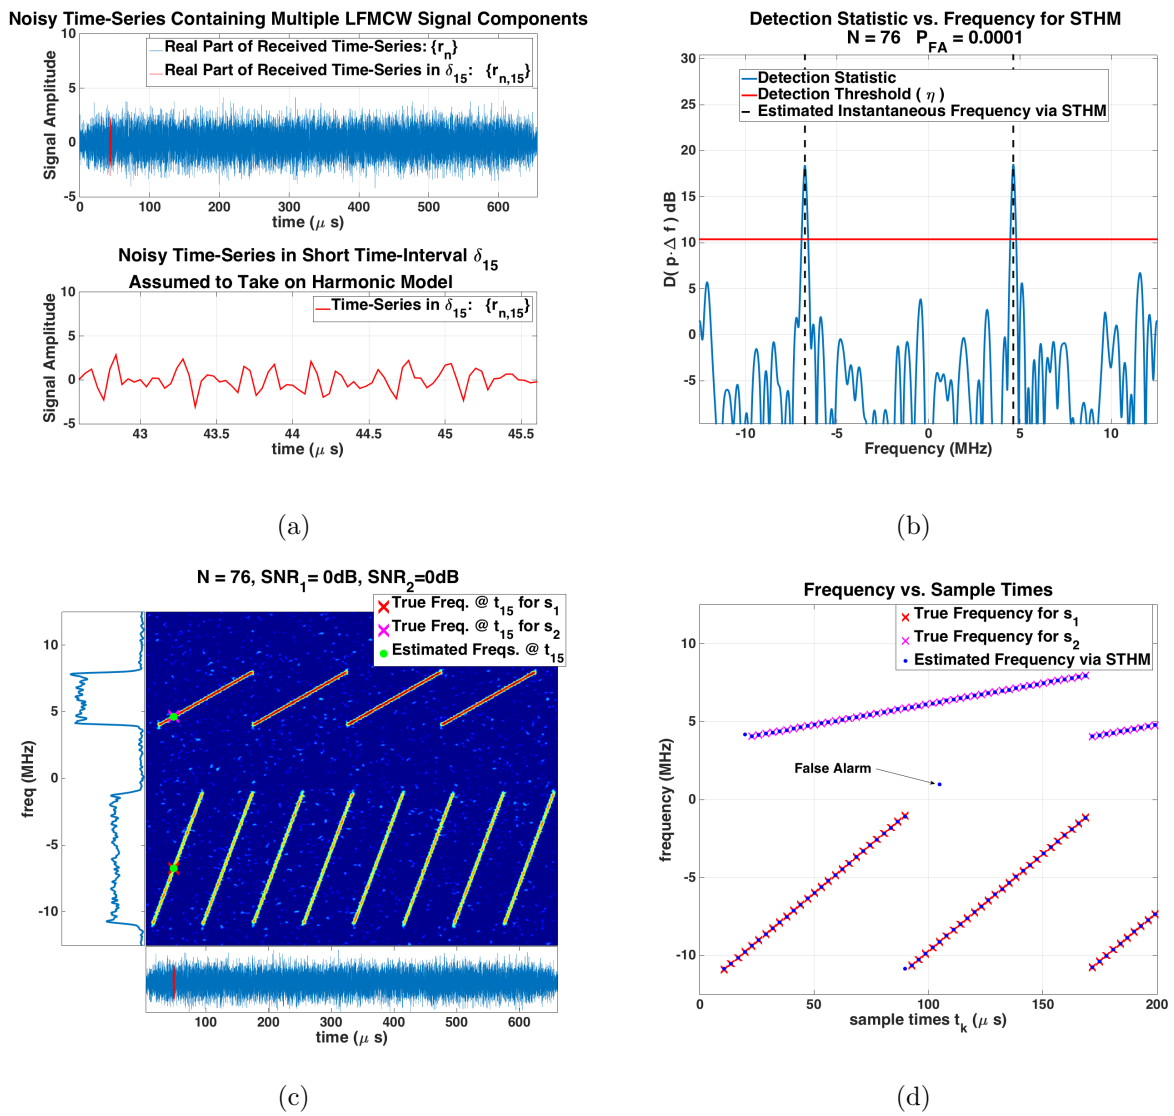

Figure 2: Detection Example: (a) Received Time-Series $\{r_n\}$ (Blue) and Received Time-Series in δ_{15} (Red), (b) Detection Test Statistic Computed based Steps 0-6 using $\{r_{n,15}\}$, (c) (main panel) Spectrogram (20 dB dynamic range) of $\{r_n\}$ with Frequency Estimates via STHM in δ_{15} (left panel) Frequency Marginal of the Spectrogram (lower panel) Real part of time-series $\{r_n\}$, (d) Frequency Estimates vs. Sample Time t_k

1.3 CHARACTERIZATION

The characterization stage picks up where the detection stage leaves off by operating on $\mathcal{S}$. The purpose of this step is to produce estimates of $\bar{\theta}_m$ given by the signal model that accompanies (1.24). The process by which these estimates are obtained is broken into 3 steps:

1. *Clustering*, wherein the frequency samples in $\mathcal{S}$ are clustered by chirp segment;
2. *Association*, wherein the chirp-segments, and their associated frequency estimates, are associated to individual LFMCW signals;
3. *Estimation*, wherein the estimates of all unknowns specified in (1.20) and (1.24) are computed for each LFMCW signal.

1.3.1 Clustering

The core ideas associated with this stage are drawn from the clustering method proposed in [10, 15], wherein the authors seek to estimate the locations of an unknown number of Radio Frequency (RF) sources using Line-of-Bearing (LOB) measurements by clustering the measurements into groups using a Mahalanobis distance criterion. Among the potentially numerous candidate clusters that arise due to one RF source, the optimal cluster is chosen to be the one that maximizes a likelihood function. Once the LOB measurements associated with the Maximum Likelihood (ML) cluster are determined, they are removed from the overall list of LOB measurements and the process is repeated until no feasible clusters are generated. For the purposes of our work the frequency estimates in $\mathcal{S}$ are analogous to the LOB measurements and the chirp segments are analogous to RF source locations.

As in [10,15], we assume the number of chirp segments present in $\mathcal{S}$ is unknown *a-priori*. It is the ability to make this assumption in conjunction with the assumption that L_k from (1.29) is unknown that allows our sequential method to handle situations where M is not only greater than unity, but also unknown and possibly zero. Additionally, we assume the frequency estimates can be modeled as $\hat{f}_{c,k} = f_{c,k}(\mathbf{x}_{c,k}) + \zeta_{f_{c,k}}$, where $f_{c,k}(\mathbf{x}_{c,k})$ is the true instantaneous frequency and $\mathbf{x}_{c,k} = \begin{bmatrix} f_{c,k}^{\text{int}} & \beta_{c,k} \end{bmatrix}^T$ is the vector of unknown parameters of the c^{th} chirp-segment in time interval δ_k corresponding to measurement time t_k. Hence $\hat{f}_{c,k}$ can now be interpreted as a measurement of the instantaneous frequency of the c^{th} chirp-segment at time t_k. Since the instantaneous frequency of each chirp-segment is, by definition, a linear function of time, $f_{c,k}(\mathbf{x}_{c,k}) = f_{c,k}^{\text{int}} + \beta_{c,k} t_k$ we have that

$$\hat{f}_{c,k} = f_{c,k}^{\text{int}} + \beta_{c,k} t_k + \zeta_{f_{c,k}} \tag{1.38}$$

$$= \mathbf{H}_k \mathbf{x}_{c,k} + \zeta_{f_{c,k}} \tag{1.39}$$

where $\mathbf{H}_k = \begin{bmatrix} 1 & t_k \end{bmatrix}$ is referred to as the model matrix. To account for the error induced on $\hat{f}_{c,k}$ by the random noise $\{w_{n,k}\}$ and the modeling error resulting from the assumed STHM found in (1.29), we include an Additive White Gaussian Noise (AWGN) error term $\zeta_{f_{c,k}}$ that we assume is zero mean with variance $\sigma^2_{\hat{f}_{c,k}}$. The fact that $\zeta_{f_{c,k}}$ is assumed to be zero mean is equivalent to assuming that $\hat{f}_{c,k}$ is an unbiased estimator for $f_{c,k}$. This is a reasonable assumption since we ensure, by Step 7 in Section 1.2, that multiple frequency components within the same δ_k are well-separated[9] and, therefore, avoid influence, or bias, due to spectral leakage from neighboring frequencies. Whiteness of $\zeta_{c,k}$ is justified if $\{w_n\}$ is white and the error terms are from different δ_k's, but perhaps only approximately true otherwise. The assumption that $P_w(f)$ in each δ_k is slowly varying is, however a reasonable justification for why this assumption will hold approximately when $\{w_n\}$ is colored. Assuming that the noise has a Gaussian distribution is a convenient mechanism for developing the maximum likelihood solution. Finally, the estimation error variance $\sigma^2_{\hat{f}_{c,k}}$ can be approximated by [40, Eq. 477b, pg. 477]

$$\sigma^2_{\hat{f}_{c,k}} \approx \frac{3}{N^3 \hat{R}_{c,k} (\pi T_s)^2} \tag{1.40}$$

[9] With reference to **Step 7** from Section 1.2.1, "well-separated" means adjacent peaks are at least $\lceil \frac{2W \cdot F_s}{\Delta f} \rceil$ Hz apart.

where $\hat{R}_{c,k} = \frac{\hat{A}^2_{c,k}}{2 \cdot \hat{\sigma}^2_{w_{n,k}}}$ is an estimate of the SNR associated with the c^{th} sinusoid in the k^{th} time interval, and $\hat{\sigma}^2_{w_{n,k}}$ is the sample variance of the residuals $\hat{w}_{n,k} = r_{n,k} - \hat{s}_{n,k}$ where $\hat{s}_{n,k} = \sum_{c=1}^{\hat{L}_k} \hat{C}_{c,k} \cdot e^{2\pi \hat{f}_{c,k} t_k}$. Then, $\left(\frac{\hat{f}_{c,k} - \mathbf{H}_k \mathbf{x}_{c,k}}{\sigma_{\hat{f}_{c,k}}} \right)^2$ is approximately Chi-Square distributed with one degree-of-freedom, an approximation we take to be a fact in what follows. So,

$$\left(\frac{\hat{f}_{c,k} - \mathbf{H}_k \mathbf{x}_{c,k}}{\sigma_{\hat{f}_{c,k}}} \right)^2 \sim \chi^2_1 \tag{1.41}$$

where χ^2_1 denotes a chi-squared distribution with one degree-of-freedom.

With all the preliminaries in place we are now positioned to describe the 4 steps of the clustering algorithm. Before we do so we point out a simplification in notation that we will adopt throughout the remainder of this Section. In particular, we index the sample times and associated frequency estimates in $\mathcal{S}$ with a single index, rather than the $(\cdot)_{c,k}$ used up to this point. As a result, when a time-interval δ_k gives rise to more than one frequency estimate there will be indices i and j, where $i \neq j$, such that $t_i = t_j$ and $\hat{f}_i \neq \hat{f}_j$. While this notation is generally more convenient, we will occasionally need to refer to the time interval δ_k that gave rise to a particular frequency estimate $\hat{f}_i$. To do so we define the mapping $\mathcal{I}(i) = k$, which is not one-to-one. This mapping amounts to book keeping in a computer implementation.

Step 1 – *Compute Candidate Chirp-Segments*: The first step is to compute all possible chirp-segment parameters from pairs of frequency estimates $\left(\hat{f}_i, \hat{f}_j \right)$, such that $|\mathcal{I}(i) - \mathcal{I}(j)| > g$, where $g > 0$ is a user-defined integer. This integer specifies the minimum number of time intervals that must separate frequency samples used to compute candidate chirp-segment parameters. A trade-off exists in the selection of g. Choosing g to be too small will result in an unnecessarily large computational burden since a very large number of candidates will result. Furthermore, solutions resulting from sample times too close together will be ill-conditioned. On the other hand, choosing g to be too large will limit the number of candidates available from which to form clusters. For the time-scales and sample-rates considered in this example, $g = 5$ works well. However, if

the conditions on N given in Section 1.2.2 are met then $g = 5$ will likely work well in general. Hence, the MLE of $\mathbf{x}_{ij}$ [21, pg. 186] is

$$\hat{\mathbf{x}}_{ij} = \left(\mathbf{H}_{ij}^T \mathbf{R}_{ij}^{-1} \mathbf{H}_{ij}\right)^{-1} \mathbf{H}_{ij}^T \mathbf{R}_{ij}^{-1} \hat{\mathbf{f}}_{ij} \tag{1.42}$$

$$= \begin{bmatrix} \hat{\beta}_{ij} \\ \hat{f}_{ij}^{\text{int}} \end{bmatrix} \tag{1.43}$$

where

$$\mathbf{H}_{ij} = \begin{bmatrix} \mathbf{H}_i \\ \mathbf{H}_j \end{bmatrix} \tag{1.44}$$

$$\hat{\mathbf{f}}_{ij} = \begin{bmatrix} \hat{f}_i \\ \hat{f}_j \end{bmatrix} \tag{1.45}$$

$$\mathbf{R}_{ij} = \begin{bmatrix} \sigma_{\hat{f}_i}^2 & 0 \\ 0 & \sigma_{\hat{f}_j}^2 \end{bmatrix} \tag{1.46}$$

The associated estimation error covariance [21, pg. 186] is

$$\mathbf{P}_{ij} = E\left[\left(\mathbf{x} - \hat{\mathbf{x}}_{ij}\right)\left(\mathbf{x} - \hat{\mathbf{x}}_{ij}\right)^T\right] \tag{1.47}$$

$$= \left(\mathbf{H}_{ij}^T \mathbf{R}_{ij}^{-1} \mathbf{H}_{ij}\right)^{-1} \tag{1.48}$$

and the associated estimation error covariance is [21, pg. 186]

$$\mathbf{P}_{ij} = E\left[\left(\mathbf{x} - \hat{\mathbf{x}}_{ij}\right)\left(\mathbf{x} - \hat{\mathbf{x}}_{ij}\right)^T\right] = \left(\mathbf{H}_{ij}^T \mathbf{R}_{ij}^{-1} \mathbf{H}_{ij}\right)^{-1} \tag{1.49}$$

Each computed $\hat{\mathbf{x}}_{ij}$ and associated $\mathbf{P}_{ij}$ are only considered chirp-segment candidates if the resulting $\hat{\beta}_{ij}$ is within a specified range of chirp-rates. In practice, chirp-rate can be bounded since the physical limitations of hardware will prohibit arbitrarily large or small chirp-rates.

Step 2 – *Determine Frequency Estimates that are Statistically Similar to each Candidate:* In this step we use the statistical distance metric known as the Mahalanobis distance to compare each frequency estimate from $\mathcal{S}$ to each candidate chirp-segment computed in the previous step. Specifically, for frequency estimate $\hat{f}_i$, for $i = 1, \ldots, N_{\text{est}}$, and chirp-segment candidate $\mathbf{x}_j$, where $j = 1, \ldots, N_{\text{cand}}$, the Mahalanobis distance is $\mathcal{M}_{ij} = \left(\frac{\hat{f}_i - \mathbf{H}_j \hat{\mathbf{x}}_j}{\sigma_{\hat{f}_i}} \right)^2$, which, like in (1.41), we assume is approximately χ^2_1 distributed when $\hat{\mathbf{x}}_j$ is the actual chirp-segment associated with $\hat{f}_i$. We use this approximation to formulate a hypothesis test whose null hypothesis is that the frequency estimate $\hat{f}_i$ falls on the candidate chirp-segment defined by $\hat{\mathbf{x}}_j$. Let α^{C} be the probability of Type I error (*i.e.*, the probability that the null hypothesis is rejected when it is true) and let γ be the critical value from the χ^2_1 distribution such that

$$\Pr\left(\chi^2_1 \leq \gamma\right) = \int_0^{\gamma} \chi^2_1(\xi)\, d\xi = 1 - \alpha^{\text{C}} \tag{1.50}$$

which can be determined numerically [10] or using standard tables [7]. So, we let $\hat{f}_i$ be a member of cluster $\mathcal{C}_j$ if $\mathcal{M}_{ij} \leq \gamma$. In other words, $\mathcal{C}_j$ is the set of indices such that $\mathcal{C}_j = \{i : \mathcal{M}_{ij} \leq \gamma\}$. Since the two indices associated with the frequency estimates that are used to compute $\hat{\mathbf{x}}_j$ from Equation (1.43) will always appear in $\mathcal{C}_j$, we require that the number of elements in a feasible cluster, denoted $|\mathcal{C}_j|$, be at least 3. Finally, we note that elements of $\mathcal{C}_j$ are indexed by $j_1, \ldots, j_{|\mathcal{C}_j|}$.

Step 3 – *Find ML Cluster Among Feasible Clusters*: The third step is to determine the ML cluster. This is accomplished by first using all frequency estimates in each cluster to compute a combined chirp-segment parameter estimate, $\hat{\mathbf{x}}_{\mathcal{C}_j}$. Based on this chirp-segment parameter estimate and the associated frequency estimates in $\mathcal{C}_j$ we compute the likelihood value, $\mathcal{L}_j$.

[10] With the Statistics Toolbox in MATLAB®, the Chi-Squared critical value can be computed as $\gamma = \text{chi2inv}\left(1 - \alpha^{\text{C}}, \nu\right)$, where ν is the number of degrees of freedom and α is the probability of Type I error.

First, we note that similar to Equation (1.43) the MLE of $\mathbf{x}_{\mathcal{C}_j}$ is

$$\hat{\mathbf{x}}_{\mathcal{C}_j} = \left(\mathbf{H}_{\mathcal{C}_j}^T \mathbf{R}_{\mathcal{C}_j}^{-1} \mathbf{H}_{\mathcal{C}_j}\right)^{-1} \mathbf{H}_{\mathcal{C}_j}^T \mathbf{R}_{\mathcal{C}_j}^{-1} \hat{\mathbf{f}}_{\mathcal{C}_j} \tag{1.51}$$

where

$$\mathbf{H}_{C_j} = \begin{bmatrix} \mathbf{H}_{j_1}^T & \cdots & \mathbf{H}_{j_{|C_j|}}^T \end{bmatrix}^T \tag{1.52}$$

$$\widehat{\mathbf{f}}_{\mathcal{C}_j} = \begin{bmatrix} \hat{f}_{j_1} & \cdots & \hat{f}_{j_{|C_j|}} \end{bmatrix}^T \tag{1.53}$$

$$\mathbf{R}_{\mathcal{C}_j} = \begin{bmatrix} \sigma^2_{\hat{f}_{j_1}} & \cdots & \mathbf{0} \\ \vdots & \ddots & \vdots \\ \mathbf{0} & \cdots & \sigma^2_{\hat{f}_{j_{|\mathcal{C}_j|}}} \end{bmatrix} \tag{1.54}$$

and the associated estimation error covariance is

$$\mathbf{P}_{\mathcal{C}_j} = \left(\mathbf{H}_{\mathcal{C}_j}^T \mathbf{R}_{\mathcal{C}_j}^{-1} \mathbf{H}_{\mathcal{C}_j}\right)^{-1} \tag{1.55}$$

Hence, the value of the likelihood function is

$$\mathcal{L}_j = \prod_{l=1}^{|\mathcal{C}_j|} \frac{1}{\sigma_{\hat{f}_{j_l}} \sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\hat{f}_{j_l} - \mathbf{H}_{j_l} \hat{\mathbf{x}}_{j_l}}{\sigma_{\hat{f}_{j_l}}}\right)^2} \tag{1.56}$$

which is equivalent to the PDF of $\hat{\mathbf{f}}_{\mathcal{C}_j}$ evaluated at $\hat{\mathbf{f}}_{\mathcal{C}_j}$ and $\hat{\mathbf{x}}_{\mathcal{C}_j}$. Hence, the cluster giving the largest likelihood value is

$$j_{\max} = \arg\max_j \mathcal{L}_j \tag{1.57}$$

Step 4 – *Retain ML Cluster Estimates and Remove Associated Frequency Estimates from* $\mathcal{S}$: The fourth step is to retain the ML chirp segment estimate $\hat{\mathbf{x}}_{j_{\max}}$ and the associated estimation error covariance matrix $\mathbf{P}_{j_{\max}}$. Finally, we remove the elements of $\mathcal{C}_{j_{\max}}$ from $\mathcal{S}$ and repeat Steps 1-3 until no candidate clusters are formed.

1.3.1.1 Clustering Example In the remainder of this section we apply the clustering approach just developed to the results from the example given in Section 1.2.3. In particular, we use the frequency estimates $\hat{f}_i$ vs. sample time $\hat{t}_i$ shown in Figure 2(d) as the elements of $\mathcal{S}$ and attempt to find one cluster for each chirp-segment present in the data. In other words, if the clustering approach functions properly then we will produce twelve clusters $\mathcal{C}_j$ whose elements correspond to the indices associated with frequencies falling on the correct chirp-segment.

Figure 3 shows the clustering results in terms of the elements of each $\mathcal{C}_j$ and the resulting $\hat{\mathbf{x}}_{\mathcal{C}_j}$, with implementation parameters given in Table 4. Figure 3(a), which resembles 2(d), shows the frequency estimates within each of the twelve clusters automatically formed by the algorithm described in Steps 1-4 of this Section. Each color/marker-style combination corresponds to a unique $\mathcal{C}_j$. The most interesting result associated with Figure 3(b) is that all frequency estimates in $\mathcal{S}$ were correctly associated with estimates of chirp-segment parameters that were very close to their true values. In other words, each $\mathcal{C}_j$ contained frequency estimates from $\mathcal{S}$ that resulted from the associated true chip-segment, $\mathbf{x}_j$. Using the same color/marker-style combinations as Figure 3(a), Figure 3(b) shows the chirp-segment parameter estimates computed via (1.51) along with truth, $\mathbf{x}_j$. This figure is another way of showing that each of the twelve chirp-segments present in the data gave rise to a unique cluster that resulted in a chirp-segment parameter estimate that was very close to truth. Figure 3(c) is a magnified view of the frequency estimates in $\mathcal{C}_2$ from Figure 3(a). The intent of this figure is to show that the frequency estimates via the STHM fall along a straight line. Figure 3(d) is a magnified view of $\hat{\mathbf{x}}_{\mathcal{C}_2}$ from Figure 3(b). In addition to truth and $\hat{\mathbf{x}}_{\mathcal{C}_2}$, Figure 3(d) shows the 95% containment region resulting from the estimation error covariance matrix given by (1.55). The resulting ellipse is equivalent to the level-set of a bivariate Gaussian distribution with mean $\hat{\mathbf{x}}_{\mathcal{C}_2}$ and covariance $\mathbf{P}_{\mathcal{C}_2}$ that contains 95% of the total probability mass. The main takeaway from Figure 3(d) is that for this example and this cluster, all of the modeling assumptions leading up to it were accurate enough to produce estimates of $\hat{\mathbf{x}}_{\mathcal{C}_2}$ and $\mathbf{P}_{\mathcal{C}_2}$ that are statistically consistent since the ellipse contains truth. This same characteristic is true of the remaining 11 clusters. The observation of statistical consistency at this stage is compelling evidence for the validity of our modeling assumptions.

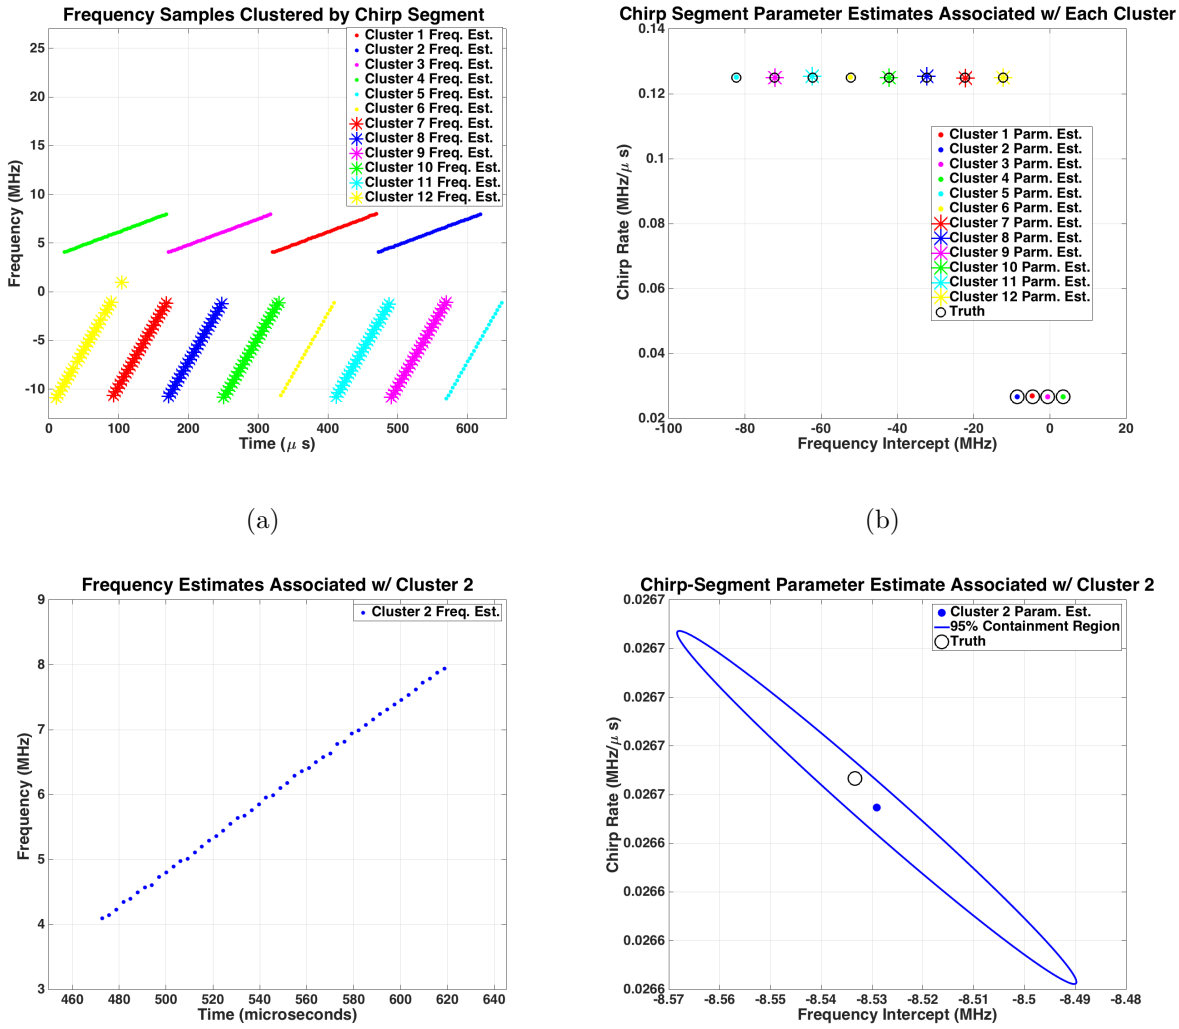

(a) (b)

(c) (d)

Figure 3: Clustering Example: (a) Chirp-Segments Clustered by Frequency, (b) Chirp-Segment Parameter Estimates Associated with Each Frequency Cluster, (c) Magnified View of Frequency Estimates Contained in Cluster 1, (d) Magnified View of Chirp-Segment Parameter Estimates Resulting from Cluster 1

Table 4: Implementation Parameters for Clustering Example

Variable Description	Variable Name	Value	Units	Relevant Clustering Step
Minimum # of time intervals δ_k that must separate frequency samples to compute candidate chirp segment	g	5	—	1
Maximum Feasible Chirp Rate	βmax	0.225	MHz/μs	1
Minimum Feasible Chirp Rate	βmin	-0.225	MHz/μs	1
Probability of Failing to Associate a Frequency Sample to Correct Chirp Segment	α^C	10^{-4}	—	2
Minimum Number of Elements in Cluster	$N_{\min}$	6	—	2

By contrast to all other detection and characterization steps of or sequential method, this step is the most time consuming and would benefit most of improvements related to increasing executing speed.

1.3.2 Association

The goal of the association stage is to associate $\mathcal{C}_j$'s corresponding to the same LFMCW signal. With reference to Figure 3(a), if the association stage performs properly the four higher frequency chirp-segments will associate with one LFMCW signal and the remaining eight lower frequency chirp-segments with associate with the other LFMCW signal. To accomplish this association within a statistical framework we will perform two hypothesis tests.

In the first hypothesis test we treat the samples within each $\mathcal{C}_j$ as realizations from an unknown PDF. This view of the data motivates a hypothesis test where the null hypothesis is that frequency estimates contained in $\mathcal{C}_i$ and $\mathcal{C}_j$ are realizations from the same underlying PDF. Rejecting this null hypothesis, and accepting the alternative hypothesis, amounts to concluding the data are from different PDF's. The most common approach for deciding between these hypotheses is the so-called Kolmogorov-Smirnov (KS) test [57, pgs. 620, 623-626]. The KS test statistic is

$$D_{ij}^{\mathrm{KS}} = \underset{-\infty<f<\infty}{\arg\max} \left| S_{\mathcal{C}_i}(f) - S_{\mathcal{C}_j}(f) \right| \tag{1.58}$$

where, for $|\mathcal{C}_j|$ frequency estimates $\hat{f}_{j_1}, \ldots, \hat{f}_{j_{|\mathcal{C}_j|}}$, $S_{\mathcal{C}_j}(f)$ is the function that gives the fraction of estimates less than a given value f (*i.e.*, an estimate of the Cumulative Distribution Function (CDF)). Under the null hypothesis, the CDF of D_{ij}^{KS} can be closely approximated by a computable function. So, the null hypothesis is rejected if

$$\alpha^{\mathrm{KS}} \leq \mathrm{CDF}_{D^{KS}}\left(D_{ij}^{\mathrm{KS}}\right) = \Pr\left(D^{\mathrm{KS}} < D_{ij}^{\mathrm{KS}}\right) \tag{1.59}$$

where α^{KS} is the significance level of the test ($\alpha^{\mathrm{KS}} = 0.05$ is a common value) and

$$\Pr\left(D^{\mathrm{KS}} < D_{ij}^{\mathrm{KS}}\right) \approx Q_{\mathrm{KS}}\left(\left[\sqrt{N_e} + 0.12 + 0.11/\sqrt{N_e}\right] \cdot D_{ij}^{\mathrm{KS}}\right) \tag{1.60}$$

is a useful approximation to the CDF of D^{KS} [53] where

$$Q_{\text{KS}}(\lambda) = 2\sum_{l=1}^{\infty}(-1)^{l-1}\, e^{-2j^2\lambda^2} \tag{1.61}$$

$$N_e = \frac{|\mathcal{C}_i| \cdot |\mathcal{C}_j|}{|\mathcal{C}_i| + |\mathcal{C}_j|} \tag{1.62}$$

The second hypothesis test uses the estimation error variance associated with chirp-rate $\sigma^2_{\hat{\beta}_{\mathcal{C}_j}}$, contained in $\mathbf{P}_{\mathcal{C}_j}$, to formulate a chi-squared test that, under the null hypothesis, assumes $\hat{\beta}_{\mathcal{C}_i} = \hat{\beta}_{\mathcal{C}_j}$. The test statistic is

$$D^{\beta}_{ij} = \left(\frac{\hat{\beta}_{\mathcal{C}_i} - \hat{\beta}_{\mathcal{C}_j}}{\sigma_{\hat{\beta}_{\mathcal{C}_i}}}\right)^2 \tag{1.63}$$

which has a χ^2_1 distribution when $\hat{\beta}_{\mathcal{C}_j}$ is the true chirp-rate of the i^{th} chirp-segment. Hence, the critical value γ for a probability of Type I error α^{β} is computed as in Equation (1.50). The following steps implement the association approach.

Step 1 – *Compute Test Statistics*: For $i, j = 1, \ldots, N_{\text{chirps}}$ compute D^{KS}_{ij} and D^{β}_{ij}.

Step 2 – *Find ML Association Based on Chirp-Rate*: For a fixed j and all i such that

$$\alpha^{\text{KS}} > Q_{\text{KS}}\left(\left[\sqrt{N_e} + 0.12 + 0.11/\sqrt{N_e}\right] \cdot D^{\text{KS}}_{ij}\right) \tag{1.64}$$

and $D^{\beta}_{ij} < \gamma$ (*i.e.*, $\mathcal{H}_0$ is accepted) compute the combined estimation error variance as $\sigma^2_{\hat{\beta}_j} = \frac{1}{\sum_i \frac{1}{\sigma^2_{\hat{\beta}_i}}}$ and the combined chirp-rate estimate as $\hat{\beta}_j = \sigma^2_{\hat{\beta}_j} \cdot \sum_i \frac{\hat{\beta}_i}{\sigma^2_{\hat{\beta}_i}}$. Based on these combined estimates, determine which set of associated chirp-segments maximizes its chirp-rate likelihood function

$$\mathcal{L}_j = \prod_{i=1} \frac{1}{\sqrt{2\pi} \cdot \sigma_{\hat{\beta}_j}} \cdot e^{-\frac{1}{2}\left(\frac{\hat{\beta}_i - \hat{\beta}_j}{\sigma_{\hat{\beta}_j}}\right)^2} \tag{1.65}$$

$$j_{\max} = \arg\max_j \mathcal{L}_j \tag{1.66}$$

Let $\mathcal{W}_m$ be the set containing $\hat{\beta}_{j_{\max}}$, $\sigma^2_{\hat{\beta}_{j_{\max}}}$, and all frequency estimates associated with $\mathcal{C}_i$. This set of parameter estimates and corresponding frequency estimates from each $\mathcal{C}_i$ serves to represent the m^{th} LFMCW signal component.

Step 3 – *Remove Contents of $\mathcal{W}_m$ from Consideration and Recompute:* Remove all $\hat{\beta}_i$, $\sigma^2_{\hat{\beta}_i}$, and associated $\mathcal{C}_i$ corresponding to $\mathcal{W}_m$ from consideration. If, after removal, only one $\mathcal{C}_i$ remains, assign it and the associated $\hat{\beta}_{\mathcal{C}_i}$ and $\sigma^2_{\hat{\beta}_{\mathcal{C}_i}}$ to $\mathcal{W}_M$, then terminate. If, after removal, no $\mathcal{C}_i$ remain, then terminate. Otherwise, repeat Steps 1-3.

1.3.2.1 Association Example In the remainder of this section we apply the association approach developed above to the results obtained in Section 1.3.1.1, with implementation parameters given in Table 5.

Table 5: Implementation Parameters for Association Example

Variable Description	Variable Name	Value	Units	Relevant Association Step
Probability of failing to associate two chirp segments that belong to same LFMCW signal based on distribution of their frequency samples	α^{KS}	0.05	—	2
Probability of failing to associate two chirp segments that belong to same LFMCW signal based on their chirp-rate estimates	α^{β}	10^{-20}	—	2

Figure 4 shows that the twelve clusters formed from the steps outlined in Section 1.3.1 were correctly associated to two LFMCW signals. This is represented in the Figure by using red markers to denote all frequency samples associated to $s_1\left(t;\bar{\theta}_1\right)$ and blue markers to denote all frequency samples associated to $s_2\left(t;\bar{\theta}_2\right)$.

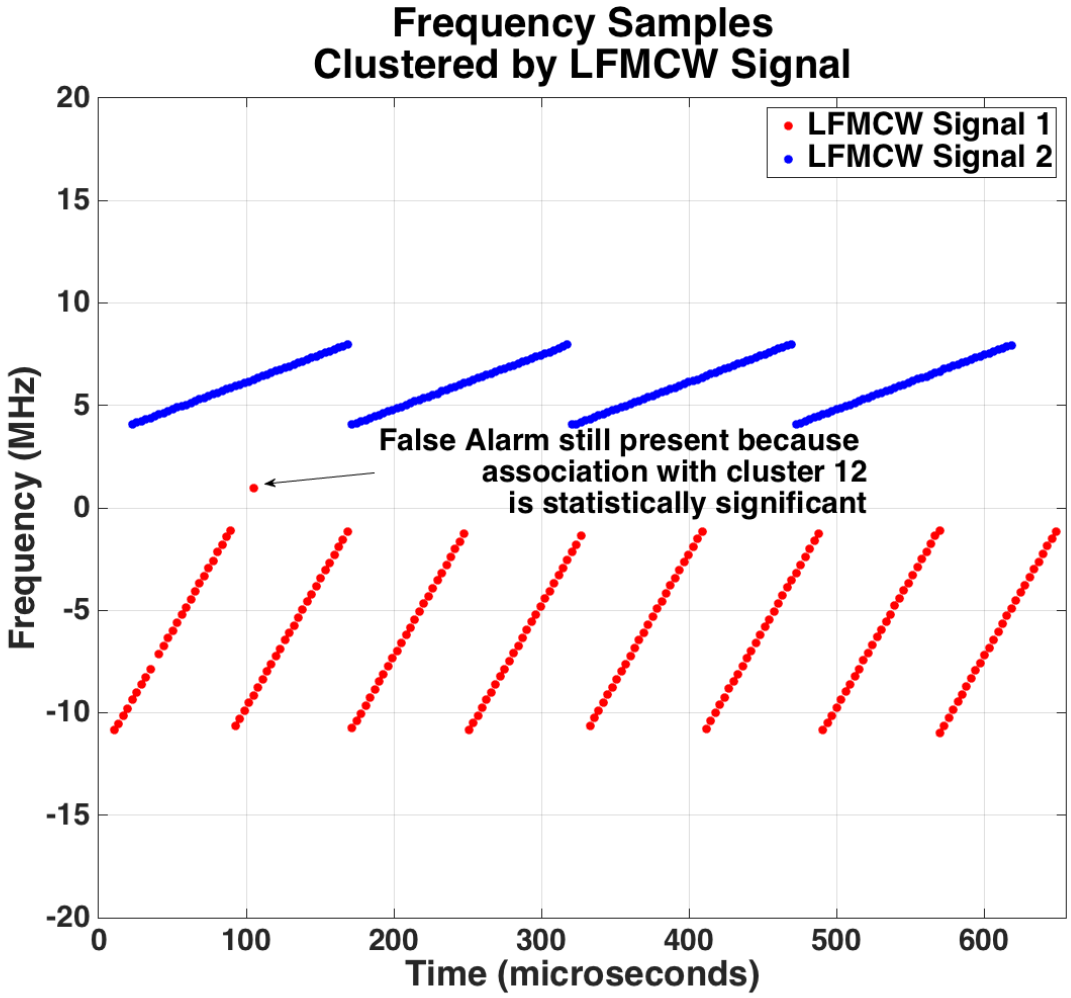

Figure 4: Association Example

1.3.3 Estimation

At this stage in the algorithm it is worthwhile to remind ourselves what our goal is and what we have at our disposal. The goal is to estimate the parameters of (1.20) given in (1.24). So far we have at our disposal

1. estimates of the mean and variance of the instantaneous frequency of each $s_m\left(t;\bar{\theta}_m\right)$
2. estimates of the mean and variance of the chirp rate associated with each $s_m\left(t;\bar{\theta}_m\right)$
3. estimates of the mean and variance of the frequency intercept of each chirp-segment, which can be thought of as conservative approximations to the mean and variance of the starting frequency of each LFMCW signal

In this section we propose a method, which we refer to as *initial* estimation, for estimating the elements of $\bar{\theta}_m$. We refer to this as initial estimation because one could use the output of this step as the initial conditions to a MLE procedure like that originally derived in [25]. This approach to implementing the MLE, in contrast to [9], may be computationally tractable for $M > 1$ since one has an estimate of the number of signals in the environment and estimates of the means and variances the parameters that define each. The author of [34] also proposed the initialization of the MLE at the conclusion of his sequential method, but was limited to the case for $M = 1$. We also note that while these estimates are referred to as initial, they may be sufficiently accurate for some applications. If very high accuracy is not required, then computational requirements of the MLE, which [34] clearly provides for $M = 1$, can be drastically reduced by avoiding the MLE search all together.

To obtain the initial estimates of $\bar{\theta}_m$, denoted $\hat{\theta}_m^0$, let the true instantaneous frequency of $s_m\left(t;\bar{\theta}_m\right)$ at the sample times in $\mathcal{W}_m$, denoted $t_{m,i}$, be given by

$$f_{m,i} = f_m^{\text{start}} + \beta_m \bmod \left(t_{m,i} + \Delta_m, T_{p,m}\right) \tag{1.67}$$

and let the associated frequency estimates in $\mathcal{W}_m$ be denoted by $\hat{f}_{m,i}$, for $i = 1, \ldots, N_m$, where N_m is the number of frequency estimates in $\mathcal{W}_m$. Then we can obtain initial estimates of f_m^{start}, β_m, Δ_m, and $T_{\text{swp,m}}$ by solving the following optimization problem

$$\underset{f_m^{\text{start}},\beta_m,\Delta_m,T_{\text{swp,m}}}{\arg\min} \left\{ \sum_{i=1}^{N_m} \left| \hat{f}_{m,i} - f_{m,i} \right| \right\} \tag{1.68}$$

Unfortunately, since the objective function in (1.68) is not only non-differentiable but also nonlinear, standard optimization techniques based on gradients of the objective function will likely fail. Therefore, to obtain initial estimates of f_m^{start}, β_m, Δ_m, and $T_{\text{swp,m}}$ we implement the Nelder-Mead simplex direct search algorithm [31], which is one example of an unconstrained, derivative-free optimization method. Finally, we can obtain initial estimates of t_m^{start} and t_m^{stop} by taking them to be

$$\hat{t}_m^{0,\text{start}} = \min_i t_{m,i} \tag{1.69}$$

$$\hat{t}_m^{0,\text{stop}} = \max_i t_{m,i} \tag{1.70}$$

In Section 1.5 we evaluate the performance of the method based on only initial estimates.

1.3.3.1 "Initial" Estimation Example In this example initial estimates of $\bar{\theta}_1$ and $\bar{\theta}_2$ are obtained by computing $\hat{t}_m^{0,\text{start}}$, $\hat{t}_m^{0,\text{stop}}$ as described above, and solving the optimization problem posed in (1.68) by using MATLAB's® implementation of the Nelder-Mead simplex direct search algorithm found in the Optimization Toolbox. The results are summarized in Table 6.

Table 6: Initial Estimation Results

Parameter	$\bar{\theta}_1$	$\hat{\bar{\theta}}_1^0$	$\bar{\theta}_2$	$\hat{\bar{\theta}}_2^0$	Units
$t_m^{0,\text{start}}$	10	10.640	20	22.800	μs
$t_m^{0,\text{stop}}$	650	649.04	620	618.64	μs
$f_m^{0,\text{start}}$	-11	-11.018	4	3.9890	MHz
β_m^0	0.125	0.12510	0.0267	0.02671	MHz
Δ_m^0	10	9.8556	20	19.9267	μs
$T_{p,m}^0$	80	79.998	150	149.95	μs

1.4 LOCALIZATION

While the use of Time Difference of Arrival (TDOA) in a passive sensor network is not a new *concept* for locating Electromagnetic (EM) signal sources, successful *implementation* is often complicated by any number of hardware or environmental limitations. One common approach for obtaining TDOA estimates is to cross-correlate two sampled versions of the same transmitted signal, taken at spatially separated receivers, and use the peak of the cross-correlation as the estimator for TDOA. This approach, which is explained in detail in [48] and [12], is attractive when hardware to support high bandwidth data links is available for the transmission of the data samples taken at each node to a central node that performs the cross correlation. This type of TDOA estimation is especially useful when very little is known, or needs to be known, about the detailed structure of the signal. However, this approach becomes less attractive when the number of signals residing in the operating bandwidth of the receiver is large or when the interference (*i.e.*, extraneous signals that are not of interest) environment becomes dense. Furthermore, if one has general knowledge of the structure of the signal, like that which we assume of the LFMCW signal modeled in (1.20), classical cross-correlation provides no means for exploiting this information to improve processing accuracy or to aid ambiguity resolution that inevitably arises in a multi-source environment.

If the structure of the transmitted signal is known, but the time the signal began to transmit is unknown, an alternative method for obtaining TDOA estimates is available that is better suited to multi-source environments. In particular, one can design a matched filter, which gives the MLE of TOA, to estimate the TOA of each source at each receiver in the network. Since the detection and characterization steps outlined above culminate in estimates of all of the parameters that define the signals, a matched filter based estimate of the TOA of each signal can be approximated by

$$\hat{\tau}_m = \arg\max_{\tau_m} \left| \int r(t)\, s^* \left(t - \tau_m; \hat{\theta}_m \right) dt \right| \tag{1.71}$$

where τ_m is the TOA of the m^{th} signal relative to a clock that is assumed to be synchronized between each node, and $\hat{\tau}_m$ is the associated estimate. Then, upon transmission of only the parameters that define the signal structure to a central node, in particular TOA, TOA

differences can be taken to eliminate the unknown transmit time and yield TDOA estimates, which are a function of only the source's position and the positions of receivers. Furthermore, since matched filtering is equivalent to correlating a known signal with an observed signal, additive interference components present in the observed signal that are uncorrelated with the signal defining the matched filter tend to have little effect on the output. In other words, interference that is uncorrelated with the EM source of interest is naturally suppressed. An additional benefit of our approach to TDOA generation and source localization is that since each TOA estimate can be unambiguously associated with a particular source, association of TOA estimates to the same source at a central node can be aided by the other parameters (*i.e.*, chirp-rate, sweep-time, starting frequency, etc.) that define the signal structure of the source. Associating observations to the same source is referred to as the *data association* problem, and is a major undertaking in all multi-source localization applications [2]. Our approach naturally provides additional information, often referred to as features in the multi-target tracking literature, that can improve data association performance.

1.4.1 Measurement Model

In this section we develop the measurement model that relates estimates of TOA $\hat{\tau}_m$ obtained via (1.71) to TDOA, and ultimately to the latitude and longitude of each LFMCW source. In order to do this we need at least three stationary or two moving receivers collaborating in the sensor network with synchronized clocks and an ability to resolve their own Earth Centered Earth Fixed (ECEF) [6, 35] position. We denote an estimate of the TOA of the m^{th} LFMCW source taken by the i^{th} receiver as $\hat{\tau}_{m,i}$, where $i = 1, \ldots, N_{\text{rx}}$ and N_{rx} is the number of receivers.[11]

[11] If any of the receivers are moving, each TOA estimate would need to be indexed by time as well. For simplicity, we do not add this extra index and assume the receivers are stationary.

Let $t = 0$ be defined as the common time at which all receivers begin to sample the signal environment and $t_{m,0}$ be defined as the time, relative to $t = 0$, at which the m^{th} source begins to transmit. The true TOA of the m^{th} LFMCW source, observed by the i^{th} receiver is then

$$\tau_{m,i} = t_{m,0} + r_{m,i}/c \tag{1.72}$$

where $r_{m,i} = \|\mathbf{p}_{s,m} - \mathbf{p}_{\text{rx},i}\|$ is the range from the m^{th} LFMCW source to the i^{th} receiver, $\|\cdot\|$ denotes the Euclidean vector norm, $\mathbf{p}_{s,m}$ is the ECEF position vector of the m^{th} source, $\mathbf{p}_{\text{rx},i}$ is the ECEF position vector of the i^{th} receiver, and c is the speed of light. TDOA can then be expressed as

$$\begin{aligned} \text{TDOA}_m^{ij} &= \tau_{m,i} - \tau_{m,j} && (1.73)\\ &= t_{m,0} + r_{m,i}/c - t_{m,0} - r_{m,j}/c && (1.74)\\ &= \left(r_{m,i} - r_{m,j}\right)/c && (1.75) \end{aligned}$$

where $i, j = 1, \ldots, N_{\text{rx}}$. We can relate the unknown position of the m^{th} source to TDOA_m^{ij} using our knowledge of the ECEF position of the i^{th} receiver $\mathbf{p}_{\text{rx},i}$. Hence, (1.75) can be rewritten as

$$\begin{aligned} \text{TDOA}_m^{ij} &= \left(r_{m,i} - r_{m,j}\right)/c && (1.76)\\ &= \left\|\mathbf{p}_{s,m} - \mathbf{p}_{\text{rx},i}\right\|/c - \left\|\mathbf{p}_{s,m} - \mathbf{p}_{\text{rx},j}\right\|/c && (1.77) \end{aligned}$$

As the name implies, the ECEF position of an object is defined relative to the earth's center and, therefore, is not constrained to the earth's surface. Since we are interested in estimating the location of stationary, ground-based LPI radar systems we can incorporate the shape of the earth as a constraint on the position by employing a model of the earth's surface. A particularly popular choice is the World Geodetic System 1984 (WGS84) earth model [6, 35], which models the earth's surface as an oblate spheroid (ellipsoid) and allows position on the earth's surface to be represented using the angles longitude and geodetic

latitude. In particular, the WGS84 earth model relates ECEF position to longitude Ψ and latitude Θ as

$$\mathbf{p}(\mathbf{q}) = \begin{bmatrix} (r_E + a)\cos(\Psi)\cos(\Theta) \\ (r_E + a)\sin(\Psi)\cos(\Theta) \\ (r_E(1 - ecc^2) + a)\sin(\Theta) \end{bmatrix} \tag{1.78}$$

where $\mathbf{q} = \begin{bmatrix} \Psi & \Theta \end{bmatrix}^{\mathbf{T}}$, a is the altitude above the WGS84 ellipsoid, and ecc is the earth's eccentricity. The term r_E is the earth's transverse radius of curvature defined by

$$r_E = \frac{r_{eq}}{\sqrt{1 - ecc^2 \sin^2(\Theta)}} \tag{1.79}$$

where r_{eq} is the earth's radius at the equator. Finally, we can rewrite (1.77) as

$$\begin{aligned} \text{TDOA}_m^{ij} &= \|\mathbf{p}_{s,m} - \mathbf{p}_{\text{rx},i}\|/c - \|\mathbf{p}_{s,m} - \mathbf{p}_{\text{rx},j}\|/c & (1.80)\\ &= \|\mathbf{p}(\mathbf{q}_{s,m}) - \mathbf{p}_{\text{rx},i}\|/c - \|\mathbf{p}(\mathbf{q}_{s,m}) - \mathbf{p}_{\text{rx},j}\|/c & (1.81)\\ &= h(\mathbf{q}_{s,m}; \mathbf{p}_{\text{rx},i}, \mathbf{p}_{\text{rx},j}) & (1.82) \end{aligned}$$

where $\mathbf{q}_{s,m}$ is the vector containing the longitude and latitude of the the m^{th} source.

In our localization approach we will assume that the altitude of each LFMCW source is either known or can be closely approximated by a constant value. This is typically a reasonable assumption in passive ELINT applications and will lead to algorithms for estimating $\mathbf{q}_{s,m}$ that are extremely stable and rapidly convergent [11, 13]. However, if the terrain is particularly mountainous, making the assumption of locally constant altitude invalid, one can incorporate terrain information, such as that found in Digital Terrain and Elevation Data (DTED), a data product published by the Department of Defense (DoD).

Equation (1.82) leads to the following model of the TDOA measurements associated with the m^{th} LFMCW source

$$\mathbf{z}_m = \mathbf{h}(\mathbf{q}_{s,m}; \mathbf{p}_{\text{rx}}) + \mathbf{v}_m \tag{1.83}$$

where $\mathbf{h}(\mathbf{q}_{s,m}; \mathbf{p}_{\text{rx}})$ is an $N_{meas} \times 1$ vector-valued function with elements defined by (1.82), with N_{meas} determined by the number of unique pairwise sensor node combinations. Additionally we define $\mathbf{p}_{\text{rx}}$ to be a $N_{\text{rx}} \times 3$ matrix containing all ECEF receiver positions, with N_{rx} being the number of receivers in the network. The term $\mathbf{v}_m$ is also $N_{meas} \times 1$ and models

the errors associated with the TDOA estimates, which we take to be independent Gaussian random variables with covariance $\mathbf{R}_m$.

One approach to estimating $\mathbf{q}_{s,m}$ is by solving the following optimization problem

$$\arg\min_{\mathbf{q}_{s,m}} \left(\mathbf{z}_m - \mathbf{h}\left(\mathbf{q}_{s,m}; \mathbf{p}_{\text{rx}}\right)\right)^T \mathbf{R}_m^{-1} \left(\mathbf{z}_m - \mathbf{h}\left(\mathbf{q}_{s,m}; \mathbf{p}_{\text{rx}}\right)\right) \tag{1.84}$$

which amounts to solving a set of nonlinear equations for $\mathbf{q}_{s,m}$. Since the argument in (1.84) is a differentiable function of $\mathbf{q}_{s,m}$, numerous numerical methods exist for determining the $\mathbf{q}_{s,m}$ that minimizes (1.84) [28]. The most popular of these is Newton's method, and a simplification of Newton's method known as Iterated Least-Squares (ILS). We develop the ILS method in detail in [17], but pursue it no further here in order to keep focus on the signal processing methods developed to yield $\hat{\theta}_m$. Further implementation issues, beyond those discussed here and in [17], are given in [14].

1.5 RESULTS

In this section we analyze the performance of the algorithms proposed in this chapter in terms of their ability to detect and characterize one (*i.e.*, $M = 1$) and two (*i.e.*, $M = 2$) LFMCW signals, with the focus being on evaluating the estimation accuracy of f^{start}, β, Δ, and T_{swp}. The true structure of the signals used in the simulations is defined in Table 2. The results for each of these cases divides this section into three main parts, wherein the first two parts quantify performance using estimates of the mean and variance of the parameters derived from Monte Carlo simulations. The third part concludes with a comparison of the estimated probability of correct characterization (P_C) between the $M = 1$ and $M = 2$ cases.

1.5.1 Comparison of Mean and Variance of Sequential Method and the MLE for $M = 1$

In this section we compare the performance of the methods developed in this chapter to the performance of the MLE given in (1.3) [25]. In general, the MLE of a parameter observed in

the presence of AWGN is known to be asymptotically unbiased and efficient [21, pg. 164]. In our context, asymptotic unbiasedness means that as the number of samples of (1.17) grows the expected value of the mean of the MLE is equal to (1.19). Similarly, asymptotic efficiency means that as the number of samples of (1.17) increases the covariance matrix associated with the MLE of (1.19) approaches the Cramer-Rao Lower Bound (CRLB) [21, pg. 27-62]. Hence, a comparison of the mean and variance of a suboptimal method, such as ours, to the MLE indicates the amount of performance that is lost.

The parameters that define the signal we used in the simulation are given in Table 2 for $m = 2$. Additionally, we set $t_2^{\text{start}} = 0$ and $t_2^{\text{stop}} = T$ and, for the sake of the MLE simulation, assumed they were known. To simulate the MLE we discretized (1.3) as

$$\mathcal{L}\left(f^{\text{start}}, \beta, \Delta, T_{\text{swp}}\right) = \left| \sum_{n=0}^{N_s - 1} r\left[n\right] s^*\left[n; f^{\text{start}}, \beta, \Delta, T_{\text{swp}}\right] \right| \tag{1.85}$$

where $N_s = \lfloor T/T_s \rfloor$ and $s\left[n; f^{\text{start}}, \beta, \Delta, T_{\text{swp}}\right]$ are discrete-time samples of (1.20) - (1.23). To obtained estimates of f^{start}, β, Δ, and T_{swp} based on the MLE we maximized (1.85) using a four-dimensional grid-search with the grid centered at the *true* parameter values and discretized to have 11 samples per dimension. The search range used for each dimension was

$$\begin{aligned}
4 - 1.54 \times 10^{-5} \leq f^{\text{start}} &\leq 4 + 1.54 \times 10^{-5} \\
0.0267 - 8.11 \times 10^{-7} \leq \beta &\leq 0.0267 + 8.11 \times 10^{-7} \\
0 - 2.68 \times 10^{-5} \leq \Delta &\leq 0 + 2.68 \times 10^{-5} \\
150 - 1.21 \times 10^{-4} \leq T_{\text{swp}} &\leq 150 + 1.21 \times 10^{-4}
\end{aligned}$$

The mean and variance for each parameter were estimated based on a Monte Carlo simulation with 50 trials per SNR level for the MLE estimator. To simulate the performance of the sequential method we used the implementation parameters described in Tables 3 through 5 and used 200 Monte Carlo trials per SNR level. The reason for the relatively few number of trials used in the MLE simulation, compared to that used to simulate the sequential method, was primarily due to the fact that the search ranges given above were so narrow, in order to make simulating the MLE computationally realizable, that few trials were necessary for estimates of the mean and variance of the parameter estimates to converge.

Figure 5 compares the mean of our proposed estimator to that computed with the MLE. Above -10 dB SNR we see that the proposed sequential estimator is reasonably unbiased in each of the parameters. Similarly, Figure 6 shows a comparison between our sequential method and the MLE based on the standard deviation estimated from the simulation for each SNR level. Again, above -10 dB SNR the MSE of each parameter achieved by the proposed sequential estimator improves more steadily to a lower standard deviation than when SNR is below -10 dB SNR. We see then that for the conditions we simulated, the

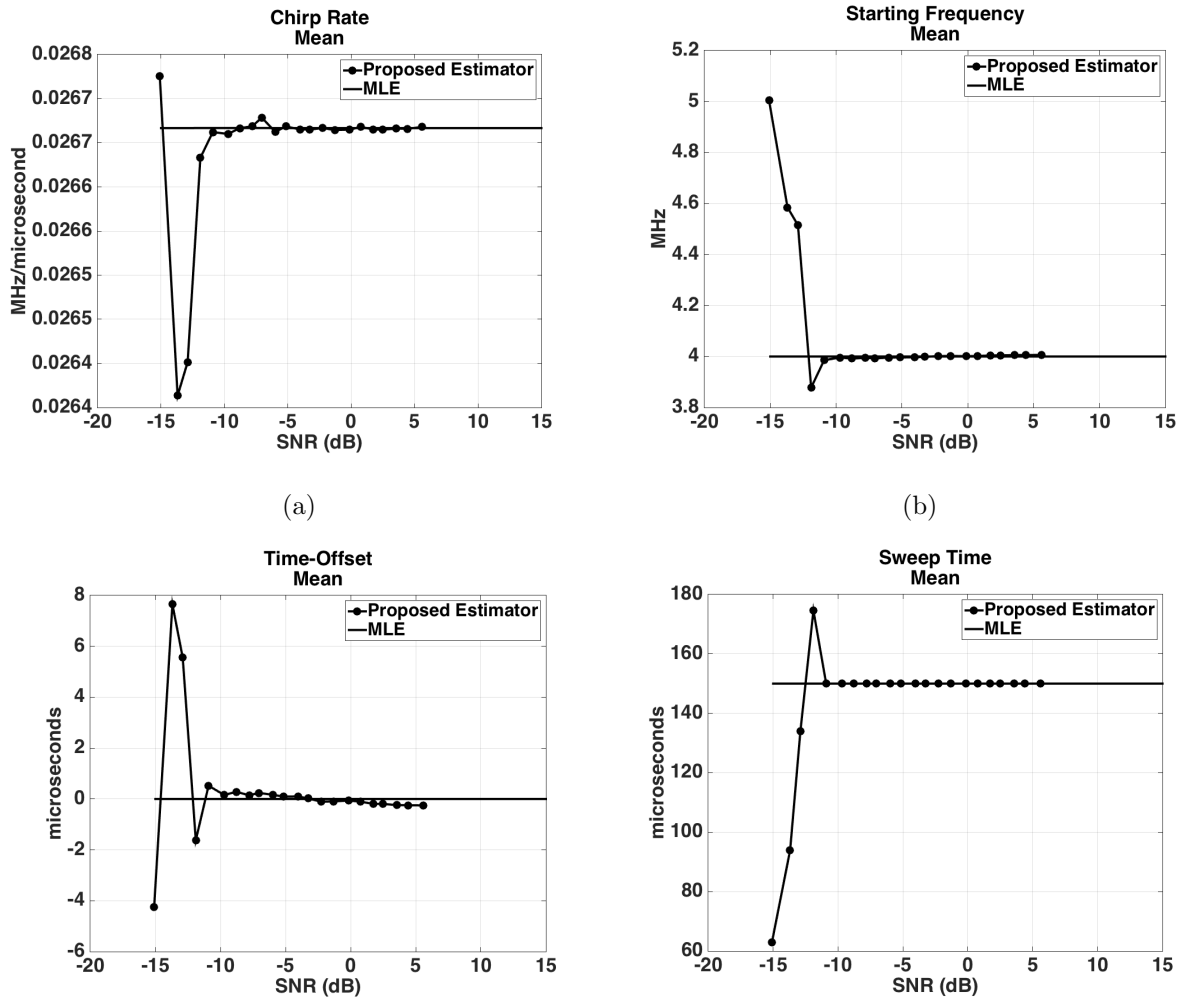

Figure 5: Monte Carlo Simulation Results for Mean of Parameter Values for $M = 1$ based on s_2

so-called *Threshold Region* of the sequential estimator is below -10 dB SNR. The threshold region of an estimator is typically defined as the level of SNR below which an estimators performance is dramatically worse than when operating just above this 'threshold.' Finding an acceptable balance between the trade-off driving sub-optimality (in our case the primary one is computational burden) and the size of the threshold region is a key consideration when designing suboptimal estimators. Figure 6 also makes it clear that each of the parameters estimated with the sequential approach yield performance gaps in their MSE that are several orders of magnitude; *this is the price we pay for a solution that is computationally tractable and practically useful when M is unknown.* The primary reason for the performance gap is that our method operates on short-time intervals of length $N \cdot T_s = 3.04\mu s$ while the MLE operates on the entire signal, which is of length $T = 655.36\mu s$.

As a point-of-reference for computational requirements, to obtain an estimate of each of the parameters in (1.18) using our method took approximately 2 seconds on a MacBook Pro with a 2.6 GHz Intel Core i7 processor and 16 GB of DDR3 RAM using MATLAB R2015a.

1.5.2 Mean and Variance of Sequential Method for $M = 2$

For simplicity, in this section we discuss results based on the performance of only the sequential estimator since a Monte Carlo simulation of the MLE where the grid is even restricted to a narrow regions about the true parameter values would have taken a prohibitively long time to complete for $M = 2$. The simulation uses all the same implementation parameters used to define the simulation described in Section 1.5.1. The main difference is that both signals from Table 2 are summed together with noise to get the simulated signal. Also, when SNRs are reported in this section the implication is that noise is added to (1.18) so that $\text{SNR}_1 = \text{SNR}_2$.

Figures 7 and 8 show that for both signals the sequential estimator is reasonably unbiased and has well-behaved and decreasing MSE for SNR levels, again, above -10 dB. We stress that, aside from the maximum and minimum chirp-rates given in Table 4, there is no *a-priori* knowledge required about the values of the parameters required, as in the MLE to define the search grid. However, the *a-priori* knowledge we do use on maximum and minimum

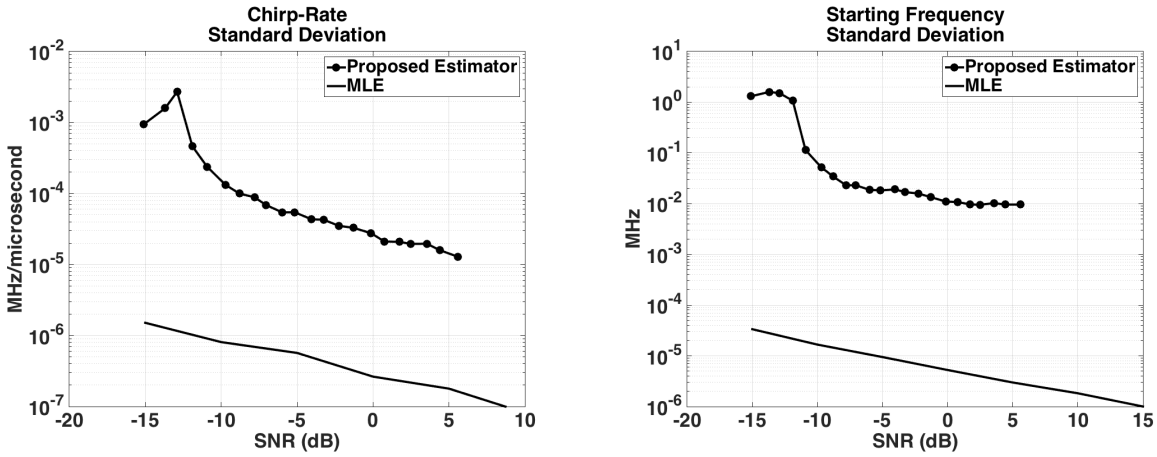

(a)

(b)

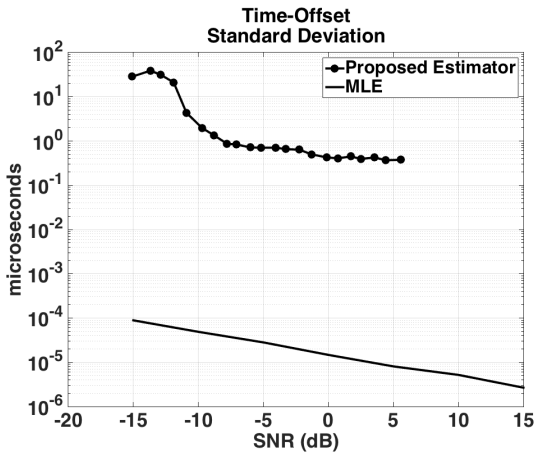

(c)

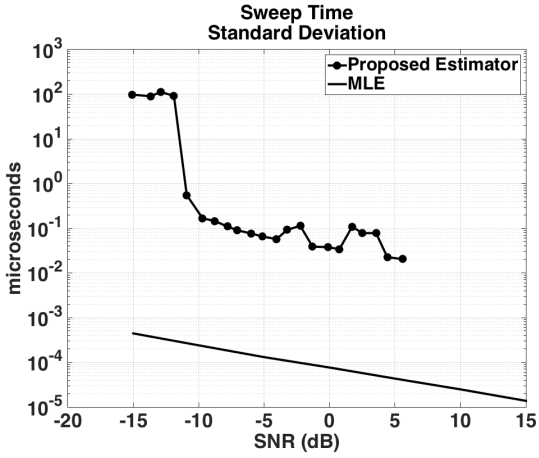

(d)

Figure 6: Monte Carlo Simulation Results for Standard Deviation of Parameter Values for $M = 1$ based on s_2

possible chirp-rates can correspond to a fairly large range of possibilities without imposing unmanageable additional computational burdens. We also emphasize that we did not inform the algorithm that two LFMCW signals were in the environment, rather the algorithm informed us. This is a feature that is absent from any other computationally feasible and practical approach. Finally, on the same MacBook Pro described in Section 1.5.1 one set of estimates of the parameters that defined both LFMCW signals took 30 seconds.

1.5.3 Probability of Correct Characterization for $M = 1$ and $M = 2$

In this section we investigate the performance of the algorithms by determining the probability of correct characterization P_C as a function of SNR and range form source to receiver. For our purposes we declare a correct characterization for signal m if f_m^{start}, β_m, and $T_{\text{swp,m}}$ are all within 10% of their true values. The results, which are given in Figures 9 and 10, are obtained from the same Monte-Carlo simulations described in Sections 1.5.1 and 1.5.2.

Figure 9 shows three P_C vs. SNR curves, which correspond to two separate simulation experiments. The curve that is the furthest to the left corresponds to the first simulation where only one signal is present (*i.e.*, $M = 1$) with specific parameters defined by $m = 2$ from Table 2. The second simulation, which was performed with both signals from Table 2 present, produced the middle curve (corresponding to signal $m = 2$) and the curve the furthest to the right (corresponding to signal $m = 1$). There are a few main takeaways that can be drawn from the comparison of these two experiments. First, the fewer the signals present in the environment, the better the algorithm will perform. This is evident from Figure 9 since the single signal case yields a P_C vs. SNR curve that has the greatest P_C for the lowest values of SNR (*i.e.*, it is the furthest to the left). Second, signals with lower chirp-rates yield better performance. For this experiment, this is the case because the signal with the greater chirp rate (*i.e.*, $m = 1$) also had the smaller sweep time. Hence for a fixed number of samples N per short time interval signal $m = 1$ yielded fewer frequency samples per chip segment, which degraded the algorithms performance relative to signal $m = 2$. To summarize Figure 9 we note that among both experiments, P_C exceeds 90% for SNR≥ -10dB. This was the claim made at the outset of this chapter. Finally, the value of SNR at which $P_C = 90\%$ for

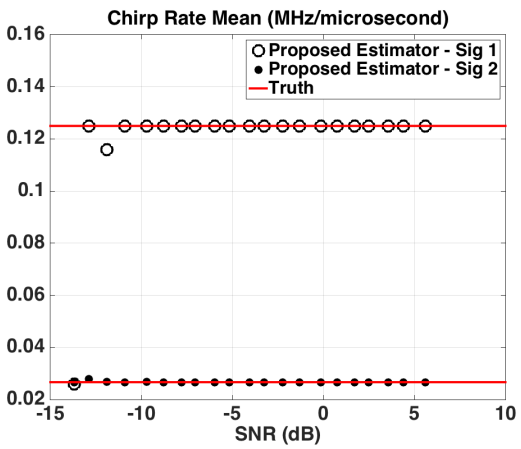

(a)

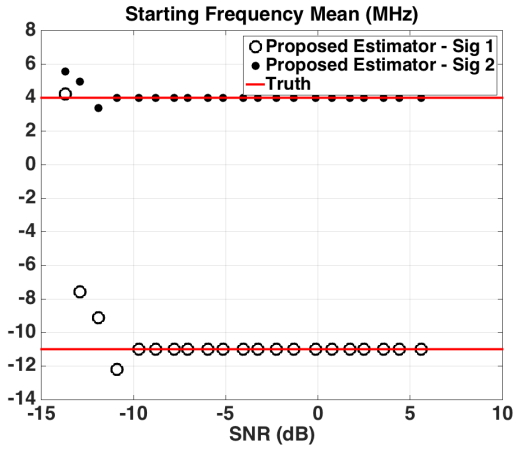

(b)

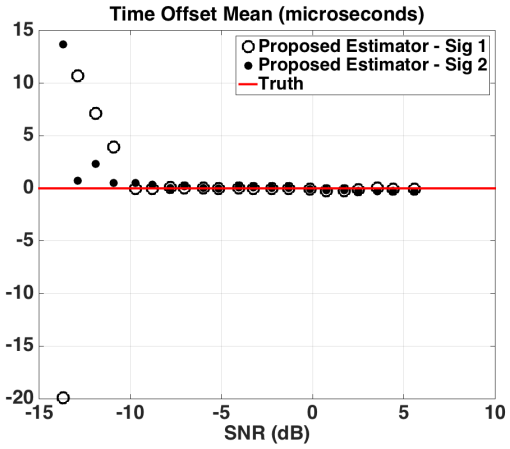

(c)

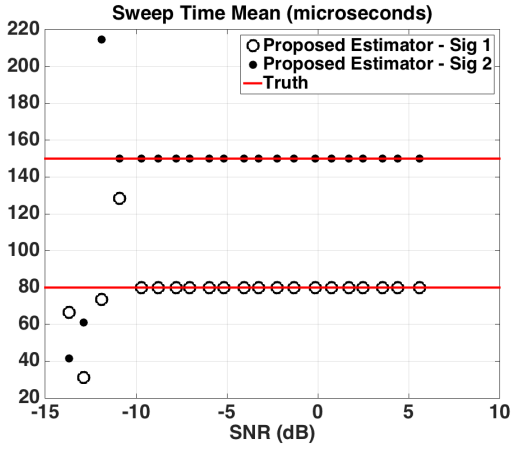

(d)

Figure 7: Monte Carlo Simulation Results for Mean of Parameter Values for $M = 2$ Based on s_1 and s_2

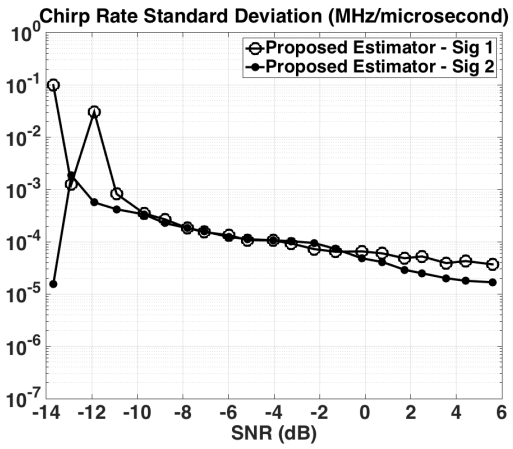

(a)

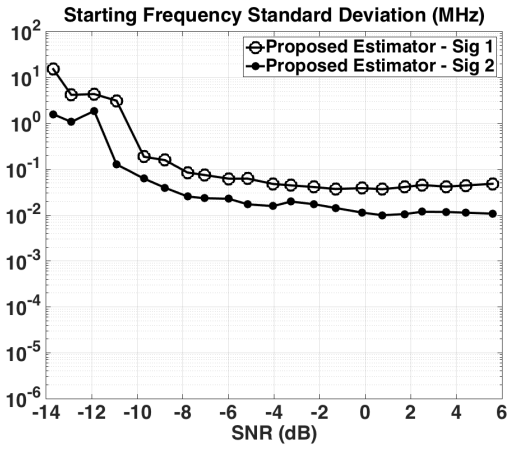

(b)

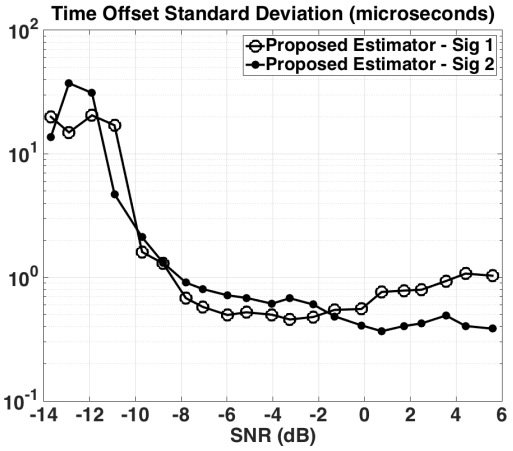

(c)

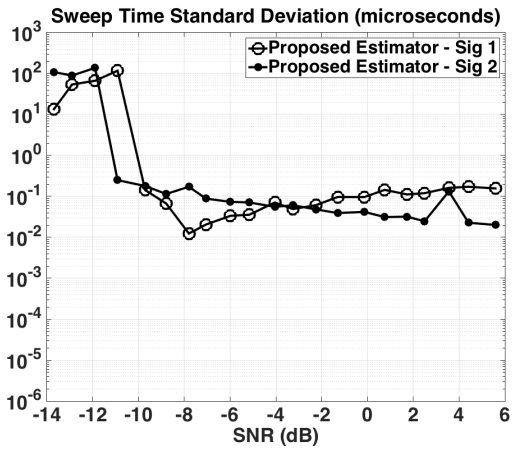

(d)

Figure 8: Monte Carlo Simulation Results for Standard Deviation of Parameter Values for $M = 2$ Based on s_1 and s_2

each experiment is close to the threshold region defined previously observed in Figures 6 and 8. That is, when SNR gets low enough that $P_C \leq 90\%$ the algorithm begins to operate in its threshold region.

To obtain Figure 10 we mapped the SNR to range through [36, pg. 24]

$$R = \sqrt{\frac{\text{EIRP} \cdot G_{TX} \cdot \lambda_c^2}{(4\pi)^2 k \cdot T_0 \cdot NF \cdot BW \cdot SNR}} \tag{1.86}$$

with the parameters used to implement (1.86) are summarized in Table 7 and it is assumed that the parameters specified in dB are converted to linear units before being used to evaluate (1.86). The range specified by this equation can be interpreted as the range at which an ELINT receiver passively collecting a radar signal transmitting with a fixed EIRP, center frequency and antenna gain, will observe the specified level of SNR assuming fixed operating bandwidth BW and noise factor NF of the receiver. The motivation for this example was to emphasize the ability of our method to operate at moderate standoff ranges with good characterization performance for an actual radar that is considered LPI, since the parameters given in Table 7 were derived from [46].

Table 7: Parameters Used to Compute Range from SNR

Variable Description	Variable Name	Value	Units
Effective Isotropically Radiated Power of LFMCW Source	EIRP	0.165	Watts
Gain of Transmitter	G_{TX}	25	dBi
Wavelength of LFMCW Center Freq.	λ_c	0.032	m
Standard Noise Temperature of Receiver	T_0	290	K
Noise Factor of Receiver	NF	5	dB
Bandwidth of Receiver	BW	25	MHz

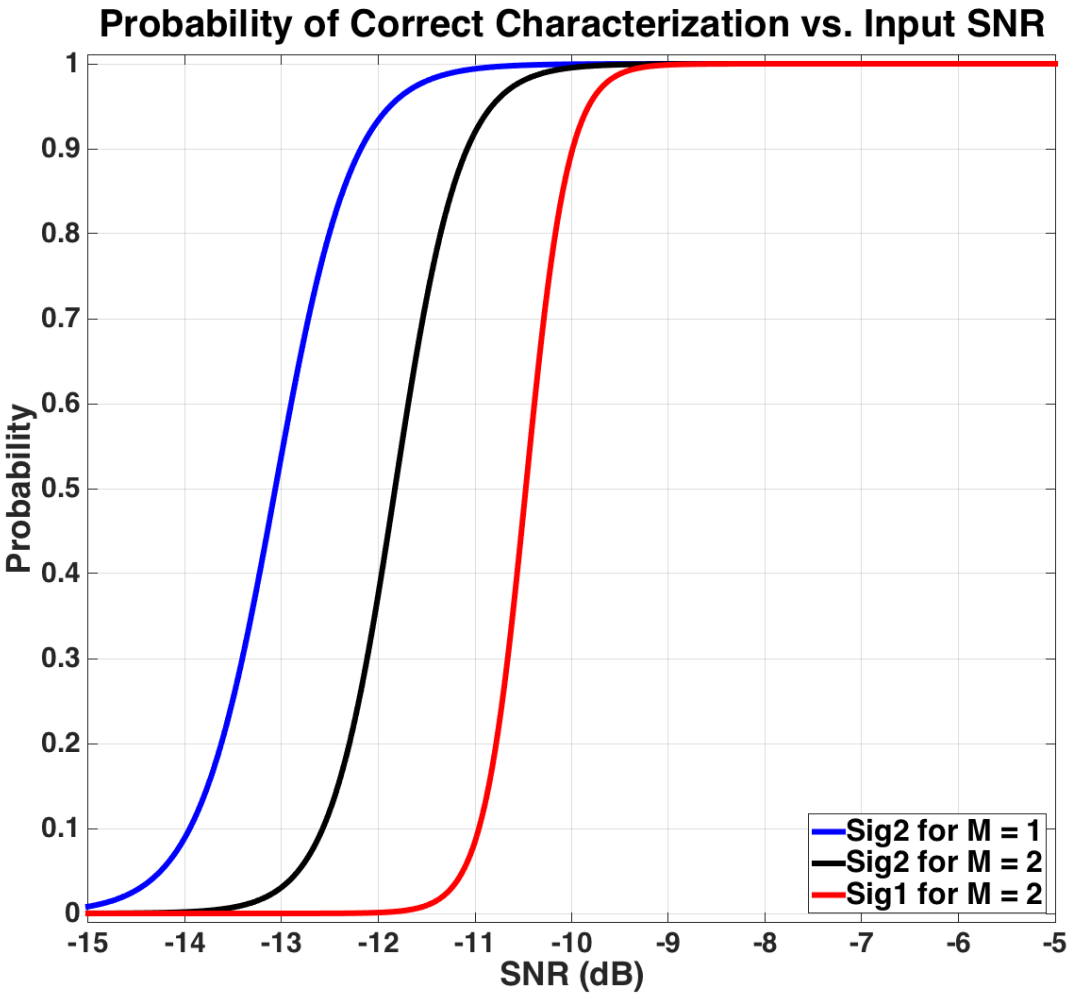

Figure 9: Characterization Probability vs. Input SNR

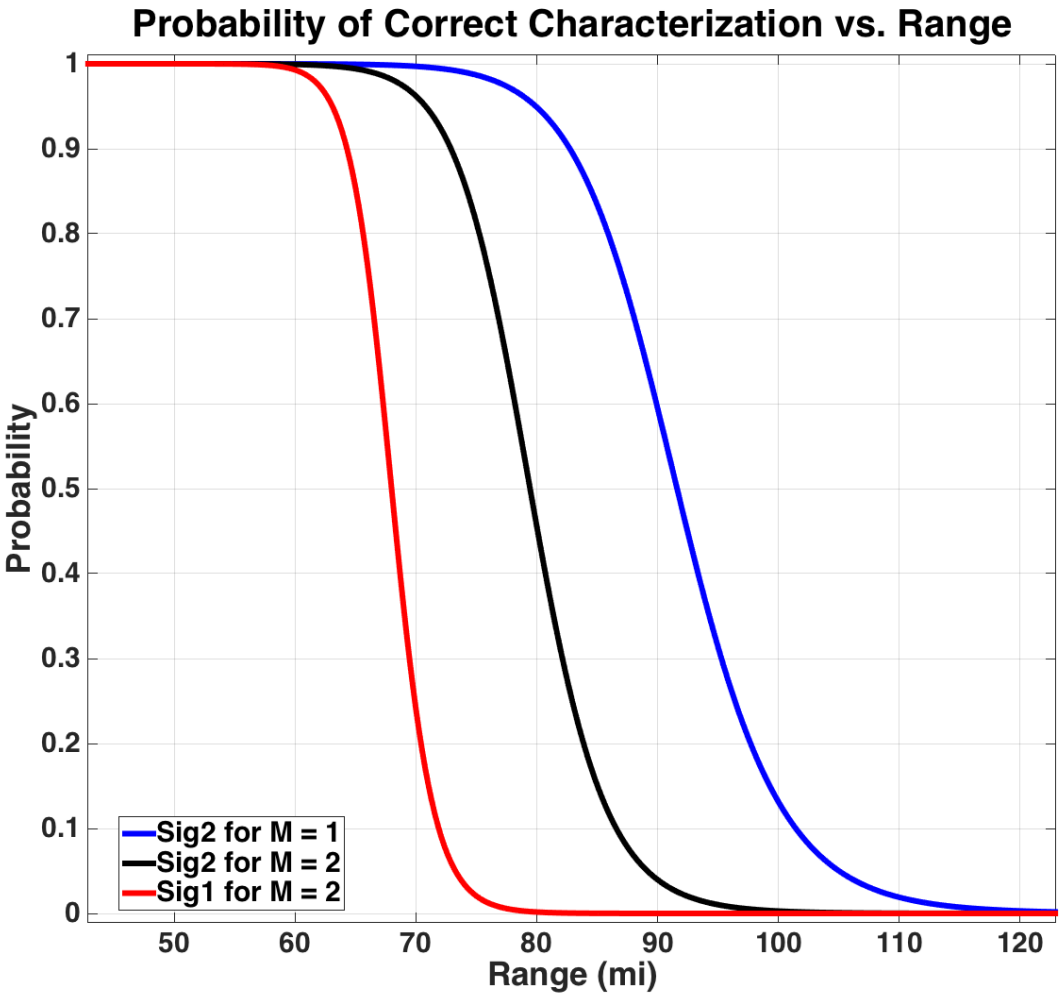

Figure 10: Characterization Probability vs. Range

1.6 CONCLUSION

In this chapter we developed the theory associated with a method that exhibits features that do not exist in any one approach otherwise available in the literature. The performance of the method is assessed via simulation and found to perform well when SNR ≥ -10 dB. In contrast to classical approaches to signal detection, such as those found in [22, 25], our approach is *sequential*. The benefit to such an approach is that the virtuous aspects of methods from a wide variety of disciples, such as signal processing [21, 22], statistics [40, 57], optimization [31], and optimal state estimation/geo-location [15], can be brought to bear on an old problem and yield new fruit that is not only of theoretical interest, but also has promise for practical implementation. The practicality of the method is by virtue of the facts that it is computationally realizable, does not require detailed assumptions about or prior knowledge on the statistics of the noise and interference, and can operate in multi-signal environments. The downside is that we give up the optimality that the MLE boasts in order to gain computational benefits and generality in terms of the environments it can operate.

In the next chapter we assess the best case estimation accuracy achievable via the Cramer-Rao Lower Bound. In Chapter 3 we asses the performance of the proposed method via laboratory-based hardware-in-the-loop and outdoor, over-the-air experiments using Software Defined Radios to transmit and receive LFMCW signals like those analyzed in this work and in Chapter 2.

1.7 FUTURE DIRECTIONS

There are at least six areas where future effort should be devoted.

- The first is in developing new or applying existing spectral search methods that are capable of determining the frequency bands where these signals exist. In other words, up to this point we have assumed that the receiver knows the appropriate center frequency and bandwidth that it should tune to in order to be able to sense the signals modeled in (1.17)-(1.23), if present. Prior information may be available to inform the receiver in practice, but if it is not then one must potentially search a large range of frequencies in order to determine what sub-band to sample and apply the methods developed in this chapter. Since a large range of frequencies can be several gigahertz, efficient search methods are required.
- The second area relates to speeding up the clustering approach found in Section 1.3.1. This is by far the most time consuming step in the proposed sequential method. More specifically, recall from Section 1.5.2 that the algorithm took 51 seconds to compute estimates of the parameters of both signals. Approximately two-thirds of this execution time was devoted to the search described in Section 1.3.1.
- The third area relates to extending the parameters estimated in the STHM of the data given in Section 1.2 from primarily just frequency to frequency *and* chirp-rate. Doing so could inform the clustering step in a way that could decrease overall execution speed. We take preliminary steps in this direction in Appendix C.
- The fourth is on a complexity study that determines the number of floating point operations required to produce parameter estimates, such as that provided in [34].
- The fifth is on extending the methods developed here to other LPI signal types. For example, signals that have LFM constituents, but are not CW or signals whose constituents are nonlinear frequency modulations. Such signals are also of practical interest.
- The sixth is on developing the localization component of this work more fully.

2.0 PERFORMANCE ANALYSIS

In this chapter we focus on deriving an approximate Cramer-Rao Lower Bound (CRLB) for the parameters of a multi-component Linear Frequency Modulated Continuous Wave (LFMCW) signal corrupted by complex additive white Gaussian noise. The approximation is necessary due to the discontinuities inherent in the mathematical model of the instantaneous phase of each LFMCW signal model. By comparing our approximate bound to a simulation of the Maximum Likelihood Estimator (MLE) of the LFMCW parameters, we confirm our analysis. In general, the CRLB is a useful tool for feasibility studies or in evaluating the degree of sub-optimality that non-MLE methods exhibit. For passive detection and estimation of LFMCW signals, the Generalized Likelihood Ratio Test and the associated MLE are difficult to implement in practice, primarily due to their large computational requirements. So, lower bounds on performance, such as those provided by the CRLB, are necessary to evaluate sub-optimal methods that are more suited for practical implementations.

2.1 INTRODUCTION

At the heart of estimation theory is the derivation of algorithms (*i.e.*, estimators) that operate on measured data with the goal being to estimate the value of a set of parameters that are related to the measurements in some way. A complementary point-of-view to the derivation of estimators is the derivation of performance bounds that predict the accuracy of a *class* of estimators. One popular approach to deriving performance bounds is based on the Cramer-Rao Theorem [21, Ch. 3], the application of which yields the so-called Cramer-

Rao Lower Bound (CRLB). The CRLB quantifies the best possible performance that can be achieved by the class of unbiased estimators in terms of their estimation error variance. In practice, the CRLB is a useful tool in performing feasibility studies or in determining the degree to which an algorithm is suboptimal.

In the context of ELectronic INTelligence (ELINT), feasibility studies requiring the use of the CRLB might be those that attempt to determine the minimum Signal-to-Noise Ratio (SNR) that can be tolerated while still achieving acceptable estimation accuracy. Based on this SNR, a link-budget analysis [45] can be performed to determine the maximum range from the radar source that an ELINT receiver can operate, which would aid in mission planning. Another use of the CRLB is in assessing whether an estimator achieves the best possible performance and, if not, the degree to which the performance of the estimator is deficient relative to the bound. Utilizing the bound in this way is particularly useful in practice since one is often faced with modifying optimal algorithms to account for practical issues or limitations encountered by the application. In particular, optimal algorithms can sometimes require computationally intractable numerical methods to implement them precisely or be derived based on simplified noise models that are mathematically convenient, but practically limited. Such is the case with the problem of estimating the parameters of an unknown number of LFMCW signals in noise with incompletely known properties. So, in order to solve the problem in a computationally effective way we, sought a suboptimal approach that allows us to trade estimation accuracy for computational tractability and the ability to operate under more general noise assumptions. The details of this method were the focus of Chapter 1 [19] of this dissertation. While the focus of this chapter [18] is on the derivation of the bound, Chapter 3 [20] of this dissertation, in part, looks at evaluating performance loss by comparing experimental and simulation results to the CRLB.

2.1.1 LFMCW Signal Model

The main element that differentiates one CRLB from another is the signal model and the parameters that define it. Hence, the CRLB can be thought of as a mechanism for quantifying the amount of information that exists in noisy observations of the signal model about the parameters that define it. For the case of estimating (*i.e.*, characterizing) the parameters of multiple LFMCW signals in noise, our *observation model* is

$$r\left(t\right) = s\left(t;\bar{\theta}\right) + w\left(t\right) \tag{2.1}$$

where the M component LFMCW *signal model* is

$$s\left(t,\bar{\theta}\right) = \sum_{m=1}^{M} s_m\left(t;\bar{\theta}_m\right) \tag{2.2}$$

$$\bar{\theta} = \begin{bmatrix}\bar{\theta}_1^T & \dots & \bar{\theta}_M^T\end{bmatrix}^T \tag{2.3}$$

with each individual LFMCW signal expressed as [1]

$$s_m\left(t;\bar{\theta}_m\right) = b_m e^{j\Phi\left(t;\bar{\theta}_m\right)} \tag{2.4}$$

The instantaneous phase of each LFMCW component is

$$\Phi\left(t;\bar{\theta}_m\right) = \varphi_m + 2\pi f_m^{\text{start}} t + \pi\beta_m \mod \left(t+\Delta_m, T_{\text{swp},m}\right)^2 \tag{2.5}$$

with the vector of unknown parameters defining the m^{th} LFMCW component given as

$$\bar{\theta}_m = \begin{bmatrix} b_m & f_m^{\text{start}} & \beta_m & \Delta_m & T_{\text{swp,m}} & \varphi_m \end{bmatrix}^T \tag{2.6}$$

with of the parameters in (2.1)-(2.6) defined in Table (8). For simplicity in our derivation we assume that $w\left(t\right)$ is white and Gaussian such that $w\left(t\right) \sim CN\left(0,\sigma_w^2\right)$. [2] The goal of this chapter is to derive the CRLB associated with the signal model given in (2.1)-(2.6). Other researchers [39] have studied the CRLB for signals similar to this signal model, but sufficiently

[1] Compared to [19], the signal model used here is lacking parameters to define start- and stop-time. We omitted them to simplify the analysis and focus on the main parameters that serve to characterize the structure of LFMCW signals and the radar systems that employ them.

[2] We take $CN\left(0,\sigma_w^2\right)$ to denote a complex normal distribution with mean zero and variance σ_w^2

different that a new approach is warranted. In particular, in [39] the author derives the CRLB for the polynomial coefficients of constant amplitude signals with polynomial phase. Since (2.5) can be viewed as a piecewise polynomial, one might be tempted to apply the bounds derived in [39]. Unfortunately, doing so would require a very high order polynomial to capture the periodicity in (2.5) and lead to difficultly in translating the CRLB of these numerous polynomial coefficients to a CRLB for the parameters of interest found in (2.6).

An example of the signal modeled by (2.1)-(2.6), for $M = 2$ is shown in Figure 11 with parameters defined and specified in Table 8. Signal $m = 2$ is annotated in Figure 11 in order to show the relationship between the variables that define the signal and the instantaneous frequency of each component. The basic signal structure given in Figure 11 and Table 8 is used in the examples we treat in this chapter.

With the problem framed, the signal model defined, and the motivation for and use of the CRLB stated, the remainder of this chapter is devoted to developing the CRLB mathematically. To do so the remainder of this chapterma is organized as follows. In Section 2.2 we state the Cramer-Rao theorem in the context of our problem and outline the mathematical components necessary to approximate the bound. The method for approximating (2.5) when taken as a function of Δ_m and $T_{\text{swp,m}}$ is the key contribution of this chapter and the focal point of this section. In Section 2.3 we compute the CRLB and compare it to the associated MLE for the case when $M = 1$ using the signal defined by $m = 2$ in Table 8. Additionally, we provide the bound for the case when $M = 2$ under the assumption that both signals, for $m = 1$ and $m = 2$, are present simultaneously. Since even the most restrictive implementation of the MLE associated with this case was computationally prohibitive[3], we provide only the bound in this example to demonstrate the CRLB result for the $M \geq 1$ case. Finally, in Section 2.4 we summarize the main elements of the chapter and provide some suggestions for future directions.

[3]A partial simulation of the MLE for $M = 2$ using Monte-Carlo settings similar to those described for the $M = 1$ case was found to take in excess of 72 hours to complete a single trial for a single SNR.

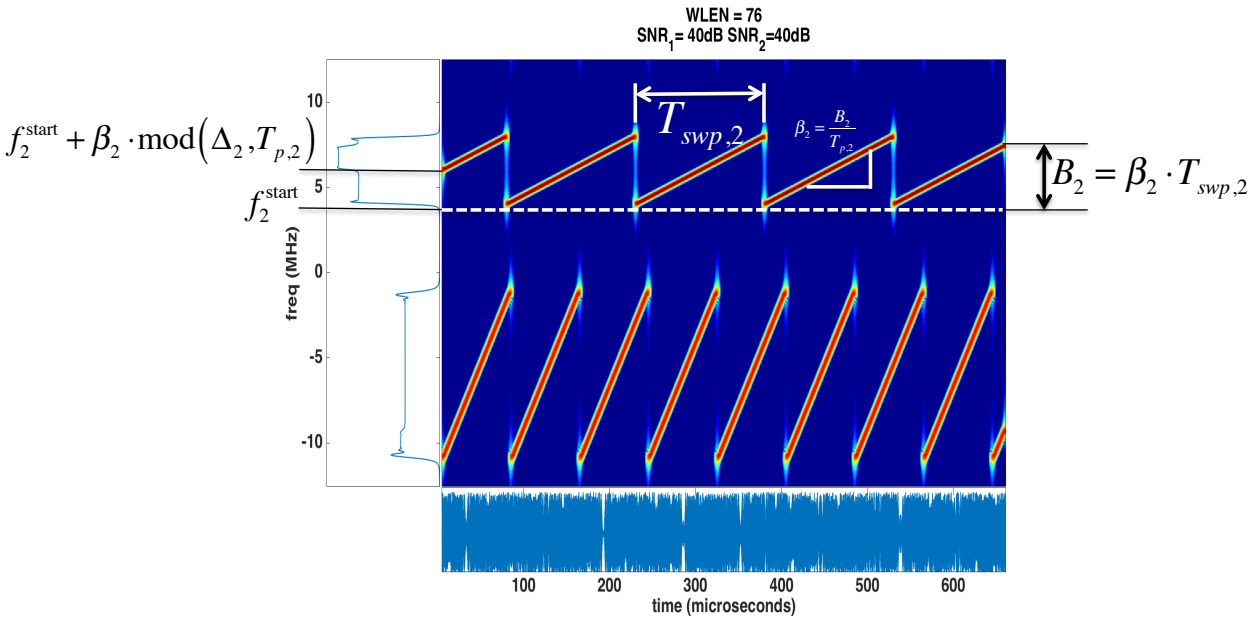

Figure 11: (main panel) Spectrogram (20 dB dynamic range) of two LFMCW signals based on signal model given in (2.1)-(2.6) (left panel) Frequency marginal of spectrogram (bottom panel) Real part of time-series $r\,(t)$

Table 8: Signal Parameters for Figure 11

Variable Description	Variable Name	Values (m=1, m=2)	Units
Amplitude	b_m	1,1	$\sqrt{J/s}$
Starting Frequency	f_m^{start}	-11,4	MHz
Chirp-Rate	β_m	0.125,0.0267	MHz/μs
Time-Offset	Δ_m	0,75	μs
Sweep Time	$T_{\text{swp,m}}$	80,150	μs
Phase	φ_m	0,0	rad
No. of LFMCW Signals	M	2	—
Sampling Rate	F_s	25	MHz
Observation Duration	T_{obs}	655.36	μs

2.2 CRAMER-RAO LOWER BOUND

The observation model given by (2.1)-(2.6) can be classified as a deterministic signal with unknown parameters $\bar{\theta}$ in Complex Additive White Gaussian Noise (CAWGN). As a result, the Cramer-Rao Theorem can be stated in the following form [21, pgs. 44-46, 530-531, 565]

Theorem 1 (Cramer-Rao) *If the PDF* $p\left(\bar{r};\bar{\theta}\right)$ *satisfies the regularity conditions* [4]

$$E\left[\frac{\partial \ln p\left(\bar{\mathbf{r}};\bar{\theta}\right)}{\partial \theta_i}\right] = 0 \qquad \forall\ \theta_i \tag{2.7}$$

where the expectation is taken with respect to $p\left(\bar{\mathbf{r}};\bar{\theta}\right)$*, then the covariance matrix of any unbiased estimator* $\hat{\theta}$ *satisfies*

$$\mathbf{R}_{\hat{\theta}} - \mathbf{I^{-1}}\left(\theta\right) \geq \mathbf{0} \tag{2.8}$$

[4] In the remainder of this chapter we use the notation θ_i to indicate indexing directly into (2.3). For example, $\theta_1 = b_1$, $\theta_2 = f_1^{\text{start}}$, $\theta_8 = f_2^{\text{start}}$, etc.

where $\geq \mathbf{0}$ *is interpreted as meaning that the matrix is positive semidefinite and* $\mathbf{R}_{\hat{\bar{\theta}}}$ *denotes the covariance matrix of* $\hat{\bar{\theta}}$*. The Fisher Information Matrix (FIM)* $\mathbf{I}\left(\theta\right)$ *is given as*

$$\left[\mathbf{I}\left(\bar{\theta}\right)\right]_{ij} = \frac{2}{\sigma_w^2}\text{Re}\left\{\sum_{n=0}^{N-1}\frac{\partial s^*\left(nT_s;\bar{\theta}\right)}{\partial\theta_i}\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\theta_j}\right\} \tag{2.9}$$

where the derivatives are evaluated at the true value of $\bar{\theta}$*.*

We see from (2.9) that the CRLB connects the accuracy of an estimator, which is manifest in $\mathbf{R}_{\hat{\bar{\theta}}}$, to how sensitive the signal model $s\left(t;\bar{\theta}\right)$ is to the parameters θ_i that define it. We can see this by recalling that derivatives, in general, reflect how small changes in an independent variable are reflected in the value of the dependent variable. For example, consider $\left[\mathbf{I}\left(\bar{\theta}\right)\right]_{ii}$. With reference to (2.9), if a small change in θ_i yields a large change in $s\left(t;\bar{\theta}\right)$ then, qualitatively speaking, the contribution of θ_i to $\left[\mathbf{I}\left(\bar{\theta}\right)\right]_{ii}$ will be large since $\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\theta_i}$ will be large. This, in turn, will contribute to making $\left[\mathbf{I}^{-1}\left(\bar{\theta}\right)\right]_{ii}$ small. That says that the more sensitive a signal model is to one of its parameters the more accurately it can be estimated since the associated CRLB will be small. Conversely, parameters that, when varied, yield little or no change in the signal make it more difficult for an estimator to distinguish between one parameter value and another, which results in a larger associated CRLB.

The statement of the Cramer-Rao Theorem makes it clear that determining the CRLB hinges on the computation of the inverse of the FIM, which in turn requires the computation of $\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\theta_i}$. So, for the CRLB to exist the FIM must also exist, which requires the existence of the derivatives $\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\theta_i}$. Upon examination of (2.5) it becomes clear that since neither $\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\Delta_m}$ nor $\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial T_{\text{swp,m}}}$ exist due to their placement as arguments in the mod $(\cdot,\cdot)$ operation, the FIM does not exist if it is based strictly on the signal model defined via (2.5). However, all is not lost if we realize that the mod $(\cdot,\cdot)$ operation is only a mathematical convenience used to impose the periodic sawtooth structure of the instantaneous frequency, an example of which can be observed in the spectrogram shown in Figure 11. So, approximations of the mod $(\cdot,\cdot)$ operation that seek to make it differentiable are not only a mathematical necessity for the derivation of the FIM, but are also practically justified since real hardware devices that generate these signals will not reflect the instantaneous changes in frequency that are

implied by (2.5). Specifically, they are likely to exhibit a more gradual transition between maximum and minimum frequencies, which we build into the approximation we develop in Section 2.2.1.

The remainder of this section is devoted to developing the approximations necessary to make (2.5) a piecewise differentiable function of Δ_m and $T_{\text{swp,m}}$. Once developed, they are implemented, for the $M = 1$ case, to compute the FIM via (2.9) and its inverse as called for by (2.8). These calculations culminate in the CRLB associated with estimating the parameters that define a single LFMCW signal. In Appendix E we generalize the results to the $M > 1$ case and show, in Appendix F, that the regularity conditions specified in (2.7) are satisfied.

For simplicity, in the remainder of this section we drop the subscripts m and 'swp' from Δ_m and $T_{\text{swp,m}}$ and let Δ and T denote time-offset and sweep time, respectively, for the single signal case. As we will see in the next section, we reintroduce subscripts on Δ and T to indicate the points where $\text{mod}\,(t + \Delta, T)$ is discontinuous when taken as a function of either of these parameters.

2.2.1 Approximating $\text{mod}\,(t + \Delta, T)$

In this section we motivate our approach for approximating $\text{mod}(t + \Delta, T)$ when taken as a function of Δ and T, outline the key aspects of the derivation of the approximation, then state the main results. Figure 12 is an examples of $\text{mod}(t + \Delta, T)$ plotted as a function of Δ for some specific values of t and T. Similarly, Figure 13 is an example of of the same function, but plotted as a function of T for some specific values of t and Δ. These figures are referenced extensively to aid in motivating and deriving the approximations and validate these approximations for the example chosen.

2.2.1.1 Approximating $\text{mod}\,(t + \Delta, T)$ wrt Δ

Figure 12 is a plot of $\text{mod}\,(t + \Delta, T)$ for $t = 50\mu s$ and $T = 80\mu s$. The first key take-away from this figure is that for fixed t and T, $\text{mod}\,(t + \Delta, T)$ is discontinuous when $t + \Delta$ is an integer multiple of T. This translates

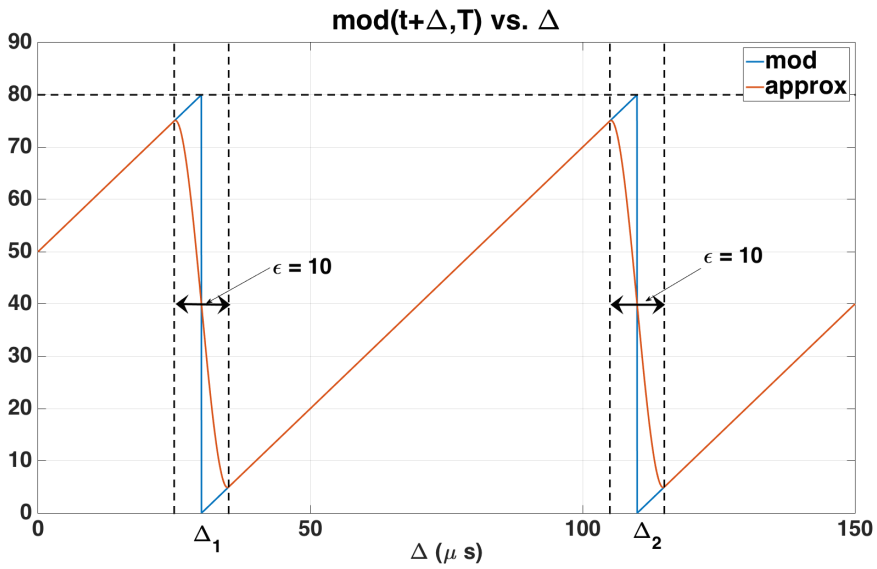

Figure 12: Diagram of mod $(t + \Delta, T)$ vs. Δ and Approximation Highlighting Impact of ε

to the condition

$$t + \Delta_d = d \cdot T \tag{2.10}$$

where $d = 0, \pm 1, \pm 2, \ldots$. Solving for Δ_d yields

$$\Delta_d = d \cdot T - t \tag{2.11}$$

which are the locations in the Δ-dimension where mod $(t + \Delta, T)$ is discontinuous. Since a discontinuous function is not differentiable [27, pg. 116], we are motivated to find a differentiable approximation of mod $(t + \Delta, T)$ by alleviating the discontinuity through a piecewise continuous approximation to mod $(t + \Delta, T)$ that is also differentiable $\forall t$ and $\forall T > 0$ on $\Delta \in [0, T]$.

One approach for developing such an approximation is to specify a small approximation region about each Δ_d and approximate the discontinuous transition with a more gradual transition using a 3^{rd} order polynomial. A 3^{rd} order polynomial is chosen because it is the smallest order polynomial whose coefficients can be chosen to match both the value and derivative of mod $(t + \Delta, T)$ at the endpoints of an approximation region given by

$$\Delta_d - \varepsilon/2 \leq \Delta \leq \Delta_d + \varepsilon/2 \tag{2.12}$$

where the parameter ε controls the width of the transition region, which we constraint to satisfy $0 < \varepsilon \ll T$. Hence, in each approximation region we have that

$$\operatorname{mod}(t+\Delta, T) \approx f(\Delta) \stackrel{\Delta}{=} \sum_{p=0}^{3} a_p(d)\,\Delta^p \tag{2.13}$$

which yields the following piecewise polynomial approximation of $\operatorname{mod}(t+\Delta, T)$

$$\operatorname{mod}(t+\Delta, T) \approx \begin{cases} t-(d-1)\,T+\Delta, & \text{if } \Delta_{d-1}+\frac{\varepsilon}{2} \leq \Delta \leq \Delta_d - \frac{\varepsilon}{2} \\ \sum_{p=0}^{3} a_p(d)\,\Delta^p, & \underbrace{\text{if } \Delta_d - \frac{\varepsilon}{2} \leq \Delta \leq \Delta_d + \frac{\varepsilon}{2}}_{\text{ApproximationRegion}} \end{cases} \tag{2.14}$$

To specify the coefficients $a_p(d)$ that define $f(\Delta)$ we must satisfy two requirements:

1. the *value* of $\operatorname{mod}(t+\Delta, T)$ at $\Delta = \Delta_d \pm \varepsilon/2$ must equal $f(\Delta_d \pm \varepsilon/2)$
2. the *slope* of $\operatorname{mod}(t+\Delta, T)$ at $\Delta = \Delta_d \pm \varepsilon/2$ must equal $\frac{\partial f}{\partial \Delta}\big|_{\Delta=\Delta_d \pm \varepsilon/2}$

Satisfying these two conditions yields the following system of four equations with four unknowns

$$f\left(\Delta_d + \frac{\varepsilon}{2}\right) = T + \frac{\varepsilon}{2} \tag{2.15}$$

$$f\left(\Delta_d - \frac{\varepsilon}{2}\right) = T - \frac{\varepsilon}{2} \tag{2.16}$$

$$\left.\frac{\partial f}{\partial \Delta}\right|_{\Delta=\Delta_d+\varepsilon/2} = 1 \tag{2.17}$$

$$\left.\frac{\partial f}{\partial \Delta}\right|_{\Delta=\Delta_d-\varepsilon/2} = 1 \tag{2.18}$$

Inverting this system of equations yields $a_p(d)$ given by

$$a_0(d) = \frac{T}{2} + t - T\,d - \frac{3\,T\,(t-T\,d)}{2\,\varepsilon} + \frac{2\,T\,(t-T\,d)^3}{\varepsilon^3} \tag{2.19}$$

$$a_1(d) = \frac{\frac{3\,T\,(2\,t-2\,T\,d)^2}{2} - \frac{3\,T\,\varepsilon^2}{2}}{\varepsilon^3} + 1 \tag{2.20}$$

$$a_2(d) = \frac{6\,T\,(t-T\,d)}{\varepsilon^3} \tag{2.21}$$

$$a_3(d) = \frac{2\,T}{\varepsilon^3} \tag{2.22}$$

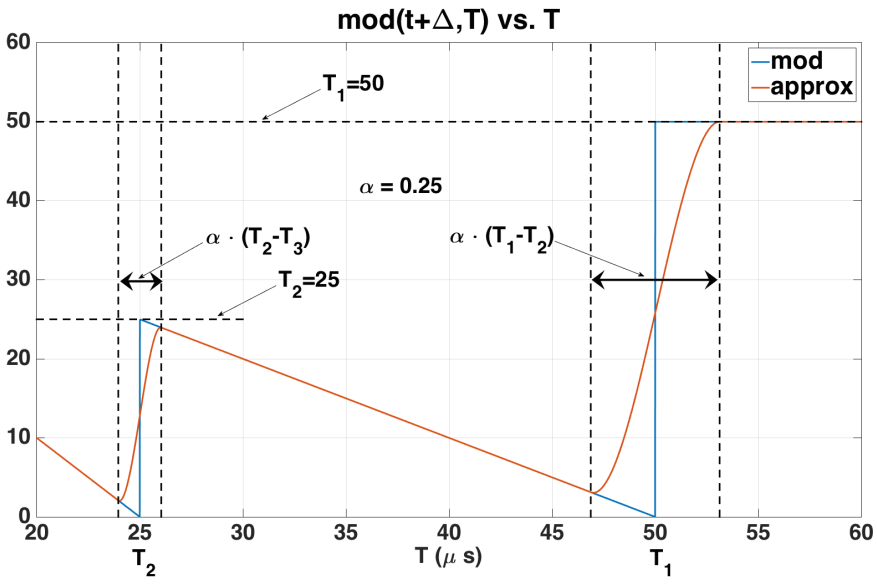

Figure 13: Diagram of mod $(t + \Delta, T)$ vs. T and Approximation Highlighting Impact of α

2.2.1.2 Approximating mod $(t + \Delta, T)$ **wrt** T We approximate mod $(t + \Delta, T)$ when taken as a function of T in a similar way as when taken as a function of Δ, with one main difference. In contrast to Figure 12, we see that in Figure 13 for $t = 50$ and $\Delta = 0$ the interval between discontinuities when mod $(t + \Delta, T)$ is taken as a function of T is not constant. The reason for this is that the condition given in (2.10), applied when T is an independent variable yields

$$t + \Delta = d \cdot T_d \tag{2.23}$$

Solving for T_d

$$T_d = \frac{t + \Delta}{d} \tag{2.24}$$

we see that as d increases, the distance between successive discontinuities decreases. To model the transition that occurs at T_d with a 3^{rd} order polynomial, this time the transition region must be a fraction of each $T_d - T_{d+1}$. We do this since if we were to fix the width of the transition region to, say, $\delta > 0$, then we could always find a d sufficiently large so that $\delta > T_d - T_{d+1}$, which would lead to a very poor approximation. Hence, we define each

transition region to be

$$T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2} \leq T \leq T_d + \frac{\alpha\left(T_d - T_{d+1}\right)}{2} \tag{2.25}$$

so that the width is $\alpha\left(T_d - T_{d+1}\right)$ with $0 < \alpha \ll 1$. So, similar to the last section, in each transition region we have

$$\operatorname{mod}\left(t + \Delta, T\right) \approx f\left(T\right) \triangleq \sum_{p=0}^{3} b_p\left(d\right) T^p \tag{2.26}$$

which yields a piecewise polynomial approximation of $\operatorname{mod}\left(t + \Delta, T\right)$, which is specified as follows

$$\operatorname{mod}\left(t + \Delta, T\right) \approx \begin{cases} \frac{-T_{d+1}}{T_d - T_{d+1}}\left(T - T_d\right), & \text{if } T_{d+1} + \frac{\alpha\left(T_{d+1} - T_{d+2}\right)}{2} \leq T \leq T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2} \\ \sum_{p=0}^{3} b_p\left(n\right) T^p, & \text{if } \underbrace{T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2} \leq T \leq T_d + \frac{\alpha\left(T_d - T_{d+1}\right)}{2}}_{\text{ApproximationRegion}} \\ t + \Delta, & \text{if } t + \Delta \leq T \end{cases} \tag{2.27}$$

To solve for $b_p\left(d\right)$ we again satisfy the boundary conditions at the edge of each approximation regions, which leads to the following system of four linear equations in four unknown

$$f\left(T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2}\right) = \frac{\alpha d \cdot \left(T_d - T_{d+1}\right)}{2} \tag{2.28}$$

$$f\left(T_d + \frac{\alpha\left(T_d - T_{d+1}\right)}{2}\right) = -\left(d-1\right)\left(T_d - \right. \tag{2.29}$$

$$\left. T_{d-1} + \frac{\alpha}{2}\left(T_d - T_{d+1}\right)\right) \tag{2.30}$$

$$\left.\frac{\partial f}{\partial T}\right|_{T = T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2}} = -d \tag{2.31}$$

$$\left.\frac{\partial f}{\partial T}\right|_{T = T_d + \frac{\alpha\left(T_d - T_{d+1}\right)}{2}} = -\left(d-1\right) \tag{2.32}$$

Inverting these equations and solving for $b_p(d)$ yields

$$b_0(d) = \frac{(\Delta+t)\,(\alpha^3-\alpha^2+2\,d^2+6\,d+6)}{\alpha^3} - \frac{\alpha\,(\Delta+t)}{8\,(d+1)} + \frac{(\Delta+t)\,(\alpha^2-4)^2}{8\,\alpha^3\,d} \tag{2.33}$$

$$b_1(d) = \frac{d+1}{2\,\alpha} - \frac{6\,(d+1)^3}{\alpha^3} - d + \frac{1}{2} \tag{2.34}$$

$$b_2(d) = \frac{d\,(d+1)\,(\alpha^2+12\,d^2+24\,d+12)}{2\,\alpha^3\,(\Delta+t)} \tag{2.35}$$

$$b_3(d) = -\frac{2\,d^2\,(d+1)^3}{\alpha^3\,(\Delta+t)^2} \tag{2.36}$$

One should select values of ε and α that are sufficiently small to produce small approximation regions, but not so small that numerical difficulties are encountered when inverting the FIM. Numerical instabilities are more likely as ε and α approach zero since they appear in the denominators of coefficients used to approximate $\mathrm{mod}\,(t+\Delta,T))$ in the approximation regions.

2.2.2 Evaluating (2.9)

To facilitate in evaluating (2.9) two properties of the FIM are important to note

1. $\frac{\partial s^*}{\partial\theta_i}\frac{\partial s}{\partial\theta_j} = \left(\frac{\partial s^*}{\partial\theta_j}\frac{\partial s}{\partial\theta_i}\right)^* \Rightarrow$ the FIM is symmetric
2. If a product $\frac{\partial s^*}{\partial\theta_i}\frac{\partial s}{\partial\theta_j}$ is purely imaginary then the corresponding FIM entry is zero

Also, we define input SNR

$$\mathrm{SNR} = \frac{b^2}{\sigma_w^2} \tag{2.37}$$

We now derive each element of the FIM. First consider the FIM entries such that $i=1$ and $j=1,\ldots,6$. For $i=j=1$, we have that

$$\frac{\partial s^*}{\partial\theta_1}\frac{\partial s}{\partial\theta_1} = \left|\frac{\partial s}{\partial\theta_1}\right| = 1 \tag{2.38}$$

Substituting this into (2.9) yields the $i=j=1$ entry of the FIM

$$[\mathbf{I}(\theta)]_{1,1} = \frac{2\cdot\mathrm{SNR}\cdot N}{b^2} \tag{2.39}$$

For $i = 1$ and $j = 2, \ldots, 6$ we have that

$$\frac{\partial s^*}{\partial \theta_i}\frac{\partial s}{\partial \theta_j} = j \cdot b\frac{\partial \Phi}{\partial \theta_j} \tag{2.40}$$

which is purely imaginary. Hence, by properties 1 and 2 given at the beginning of this section we have that

$$[\mathbf{I}(\theta)]_{1,i} = [\mathbf{I}(\theta)]_{i,1} = 0 \tag{2.41}$$

Next, for $i, j = 2, \ldots, 6$ we have that

$$\frac{\partial s^*}{\partial \theta_i}\frac{\partial s}{\partial \theta_j} = b^2 \cdot \frac{\partial \Phi}{\partial \theta_i}\frac{\partial \Phi}{\partial \theta_j} \tag{2.42}$$

Substituting this result in (2.9) and recalling property 2 above yields,

$$[\mathbf{I}(\theta)]_{i,j} = [\mathbf{I}(\theta)]_{j,i} \tag{2.43}$$

$$= 2 \cdot \mathrm{SNR} \sum_{n=0}^{N-1} \frac{\partial \Phi}{\partial \theta_i}\frac{\partial \Phi}{\partial \theta_j} \tag{2.44}$$

Computing (2.44) then amounts to determining $\frac{\partial \Phi}{\partial \theta_i}$ for $i = 2, \ldots, 6$, which are given below

$$\frac{\partial \Phi}{\partial \theta_2} = \frac{\partial \Phi}{\partial f^{\text{start}}} = 2\pi t \tag{2.45}$$

$$\frac{\partial \Phi}{\partial \theta_3} = \frac{\partial \Phi}{\partial \beta} = \pi \cdot \bmod (t + \Delta, T)^2 \tag{2.46}$$

$$\frac{\partial \Phi}{\partial \theta_6} = \frac{\partial \Phi}{\partial \varphi} = 1 \tag{2.47}$$

and below in (2.48) and (2.49).

$$\frac{\partial \Phi}{\theta_4} = \frac{\partial \Phi}{\partial \Delta} \approx \begin{cases} 2\pi\beta\left(t - (d-1)T + \Delta\right), & \text{if } \Delta_{d-1} + \frac{\varepsilon}{2} \leq \Delta \leq \Delta_d - \frac{\varepsilon}{2} \\ 2\pi\beta\left(\sum_{p=0}^{3} a_p(d)\,\Delta^p\right)\left(\sum_{p=1}^{3} p \cdot a_p(d)\,\Delta^{p-1}\right), & \text{if } \underbrace{\Delta_d - \frac{\varepsilon}{2} \leq \Delta \leq \Delta_d + \frac{\varepsilon}{2}}_{\text{ApproximationRegion}} \end{cases} \tag{2.48}$$

$$\frac{\partial \Phi}{\theta_5} = \frac{\partial \Phi}{\partial T}$$

$$\approx \begin{cases} 2\pi\beta d^2 \left(T - T_d\right), & \text{if } \; T_{d+1} + \frac{\alpha(T_{d+1} - T_{d+2})}{2} \leq T \leq T_d - \frac{\alpha(T_d - T_{d+1})}{2} \\ 2\pi\beta \left(\sum_{p=0}^{3} b_p\left(d\right) T^p\right) \left(\sum_{p=1}^{3} p \cdot b_p\left(d\right) T^{p-1}\right), & \text{if } \underbrace{T_d - \frac{\alpha\left(T_d - T_{d+1}\right)}{2} \leq T \leq T_d + \frac{\alpha\left(T_d - T_{d+1}\right)}{2}}_{\text{ApproximationRegion}} \\ 0, & \text{if } \; t + \Delta \leq T \end{cases} \tag{2.49}$$

To obtain $\frac{\partial \Phi}{\partial \theta_4}$ and $\frac{\partial \Phi}{\partial \theta_5}$ we substitute the approximations for $\text{mod}\left(t + \Delta, T\right)$ given in (2.14) and (2.27) into (2.5) and compute the partial derivatives with respect to θ_4 and θ_5. After the FIM is populated, it is inverted numerically. Since numerical matrix inversion is sensitive to scaling it is advisable to evaluate the components of the FIM in units of MHz and μs, and combinations thereof as specified in Table 8, to avoid producing entries that vary by many orders of magnitude. This would otherwise be the case since the units of some components are seconds (on the order of 10^{-6}) and others are in units of frequency (on the order of 10^6) , which would yield values that nominally differ in scale by 12 orders of magnitude.

2.3 RESULTS

To test our derivations we consider two examples, which are the focus of this section. The first example, given in Section 2.3.1 computes the bound for signal $m = 2$ from Table 8 and compares it to the result of a simulation of the MLE of f^{start}, β, Δ, and T. The second example, given in Section 2.3.2, evaluates the CRLB for both $m = 1$ and $m = 2$ from Table 8 for the case when $M = 2$. In both examples the bounds were computed with $\alpha = \varepsilon = 0.25$.

2.3.1 Comparison Between MLE and CRLB for $M = 1$

In some cases, comparing the MLE to the CRLB is a useful way to test the derivation of the CRLB. In general, the CRLB is known to reflect the minimum estimation error variance that can be achieved by any parameter estimation algorithm that is unbiased [21, Ch. 3]. It is also known that if there exits an algorithm that can achieve the CRLB, then the MLE will also achieve it asymptotically (*i.e.*, as the number of samples of $r(t)$ increases) [21, pg. 164]. Hence, if an estimator exists that can achieve the CRLB then the CRLB predicts the asymptotic performance of the MLE. So, if the simulated performance of the MLE attains the CRLB, we can conclude that the CRLB we derived is correct. For the signal model given in (2.1)-(2.6) with $w(t) \sim CN(0, \sigma_w^2)$, the MLE is derived in [25] and reproduced below in a form relevant to our problem

$$\begin{pmatrix} \widehat{f^{\text{start}}} \\ \hat{\beta} \\ \hat{\Delta} \\ \hat{T} \\ \hat{\varphi} \end{pmatrix} = \underset{f^{\text{start}},\beta,\Delta,T}{\arg\max}\ L\left(f^{\text{start}}, \beta, \Delta, T, \varphi\right) \tag{2.50}$$

where

$$\mathcal{L}\left(f^{\text{start}}, \beta, \Delta, T, \varphi\right) = \left| \sum_{n=0}^{N_s-1} r[n]\, s^*\left[n; f^{\text{start}}, \beta, \Delta, T, \varphi\right] \right| \tag{2.51}$$

The MLE simulation was performed with 50 trials per SNR using a grid search centered on the true values of the parameters with 11 equally spaced grid points. The width of

the search region in each of the five dimensions was equal to four standard deviations, with a standard deviation being that predicted by the CRLB of the parameter defining the associated dimension for the specified SNR. In other words, we specified the width of the search region for each dimension by using the standard deviation predicted by the CRLB for the associated SNR. We took this approach to specifying the search in order to design an efficient experiment since a grid that is too fine and a search region that is too large can easily lead to simulations that would require multiple *weeks* to complete. The idea was that four standard deviations were large enough to allow the MLE to exhibit its true Mean Squared Error (MSE) performance. The estimated MSE [5] is plotted along with the CRLB in Figure 14. The close correspondence between the MLE and the CRLB suggest our approximations, derivation, and implementation for the $M = 1$ case are correct.

2.3.2 CRLB for $M = 2$

Figure 15 shows the result of evaluating the CRLB when both signals $m = 1$ and $m = 2$ are included in the data model. In this example the CRLB was evaluated by following the analysis for $M > 1$ discussed in Appendix E. Since comparing these results to a simulation of the MLE would have been too time consuming, we justify them qualitatively by comparison to the results shown in Figure 14. With reference to Table 9, the standard deviations predicted by the CRLB for the $M = 1$ and signal $m = 2$ case is compared to the standard deviations predicted by the CRLB for the $M = 2$ case. Columns 2 and 3 of Table 9 correspond to the CRLBs of signal $m = 2$ for the $M = 1$ and $M = 2$ cases, respectively. The fact that their values are identical suggests two things. First, the derivation of bound for $M = 2$ is accurate to within the same approximation error as in the $M = 1$ case. Second, since the values did not change between the $M = 1$ and $M = 2$ cases for signal $m = 2$, the parameter estimates between signal $m = 1$ and $m = 2$ are uncorrelated.

[5] It is well known that the CRLB is overly optimistic in predicting the estimation error variance at sufficiently low SNR levels. The values of SNR for which the CRLB are grossly incorrect for a particular estimator is known as the estimators 'threshold region'. All estimators, including the MLE, exhibit a threshold region [3]. The simulation that yielded the performance results shown in Figure 14 does not show the MLE's threshold region because of the restrictive search range used to implement the MLE.

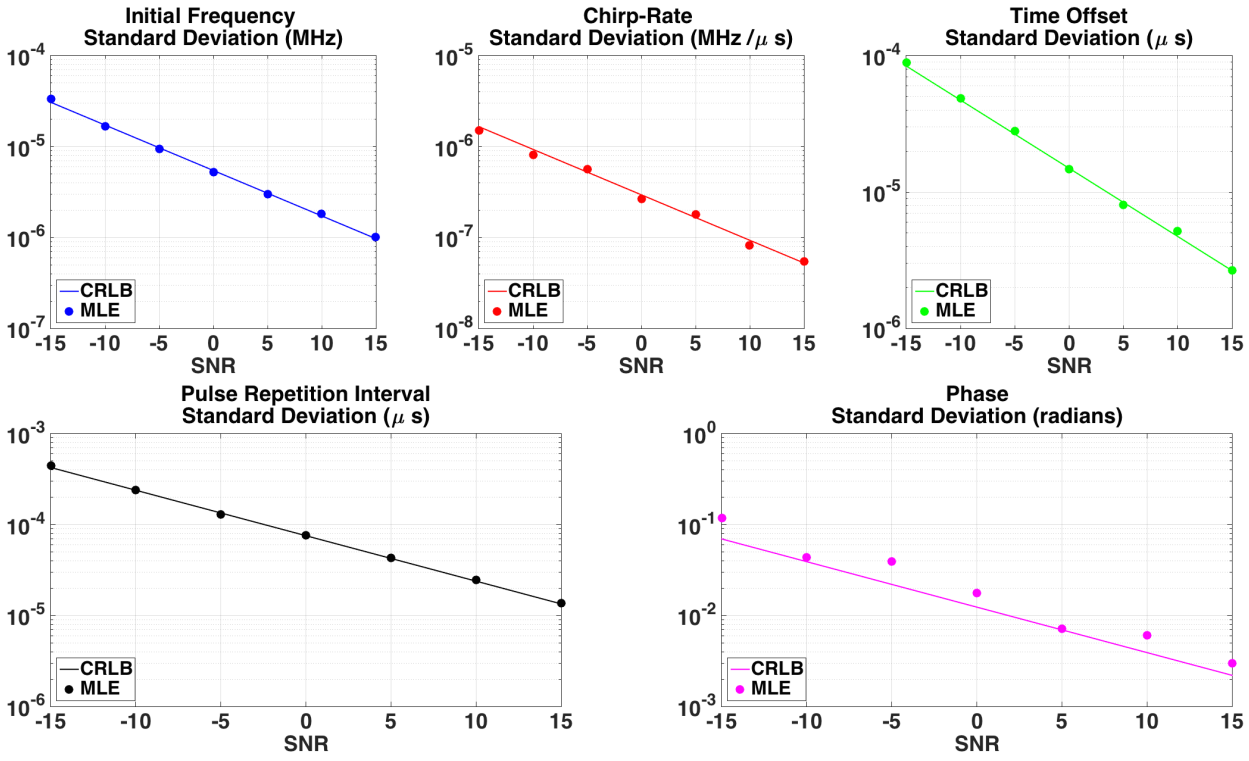

Figure 14: Comparison of MLE Estimation Error Variance to CRLB

Table 9: Comparison of Standard Deviation Predicted by CRLB for $M = 1$ and $M = 2$ Cases at SNR = 0 dB

Variable Name	$M = 1$, $m = 2$	$M = 2$, $m = 2$	$M = 2$, $m = 1$	Units
f_m^{start}	$5 \cdot 10^{-6}$	$5.5 \cdot 10^{-6}$	$5 \cdot 10^{-6}$	MHz
β_m	$3 \cdot 10^{-7}$	$3 \cdot 10^{-7}$	$1 \cdot 10^{-6}$	MHz/μs
Δ_m	$1.5 \cdot 10^{-5}$	$1.5 \cdot 10^{-5}$	$8 \cdot 10^{-6}$	μs
$T_{\text{swp,m}}$	$7.5 \cdot 10^{-5}$	$7.5 \cdot 10^{-5}$	$1.75 \cdot 10^{-5}$	μs
φ_m	$1.25 \cdot 10^{-2}$	$1.25 \cdot 10^{-2}$	$1.25 \cdot 10^{-2}$	rad

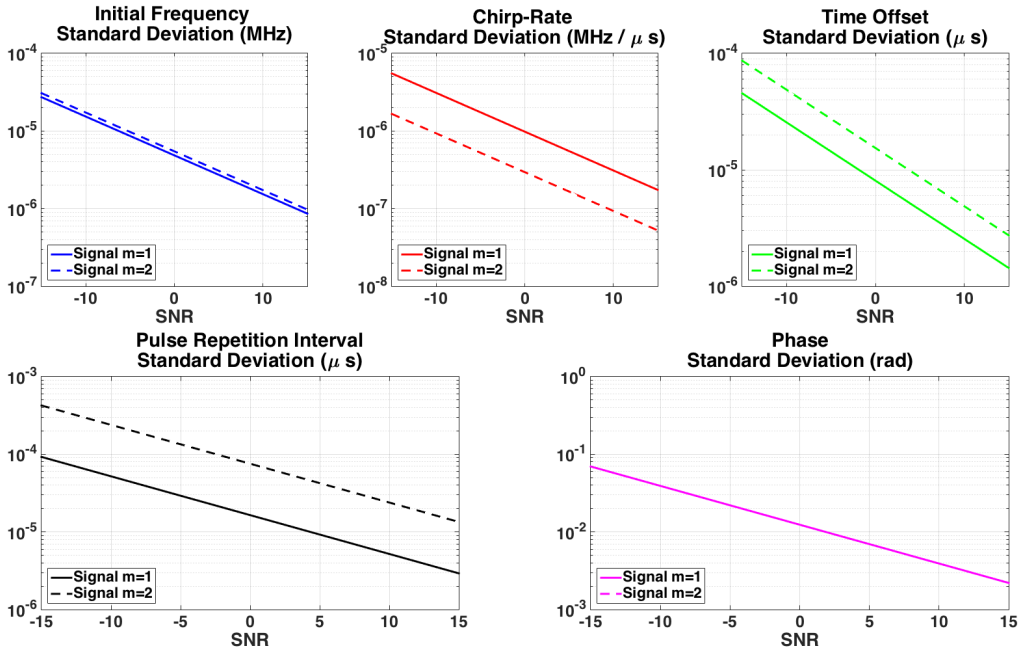

Figure 15: CRLB Computed for $m = 1$ and $m = 2$ from Table 8

2.4 CONCLUSION

In this Chapter we derived an approximate CRLB associated with the LFMCW signal model. Since the DCL approach [19] is inherently sub-optimal, relative to the MLE, the CRLB is a useful tool for quantifying the associated performance deficit, in addition to its typical application as a tool for use in trade studies. Since the instantaneous phase of the LFMCW signal model is not differentiable, we developed an approximation to the phase. Using this approximation, we developed an approximate CRLB, which we verified through simulation by comparing it to the MLE. Our approximation leads to a bound expressed directly in terms of the parameters of interest, rather than indirectly as would be required if previous work [39] on this topic were applied.

In the next chapter we assess the performance of the method based on laboratory hardware-in-the-loop and outdoor over-the-air experiments using Software Defined Radios to transmit and receive LFMCW signals like those analyzed in this chapter. The performance observed from these experiments is compared to the bounds just derived.

2.5 FUTURE DIRECTIONS

In Section 2.2.1 we give broad guidance for selecting the values of ε and α. However, specific analysis that provides more direct guidance on selecting these parameters so that the bound exists is of practical interest.

3.0 EXPERIMENTAL RESULTS

In this chapter we focus on testing the method developed in Chapter 1 [19] and comparing these experimental results to the bounds derived in Chapter 2 [18]. The experimental results are products of three experimental setups, which include 1) simulation, 2) hardware-in-the-loop and 3) over-the-air configurations. Each of these experimental setups was based on the same two test signals, but differ in the means by which the signals are generated, the medium over which they propagate, and the way they are sampled. Overall, the goal of these experiments was to incrementally introduce complicating factors that are inevitable in the practice. Specifically, the simulation analyzed performance in the presence of Gaussian system noise, the hardware-in-the-loop experiment added the potential for deviations from the ideal signal models assumed of the transmitter and receiver, and the over-the-air configurations experiment introduced interference sources whose properties are not only unknown, but also non-stationary. The data from these three experiments are analyzed in detail in this chapter, but, in brief, indicate that two simultaneous LFMCW signals can be detected and their parameters estimated to within 10% of their true values with probability greater than 90% for signal to noise ratios that are greater than -10 dB.

3.1 INTRODUCTION

Early in the history of radar is was known that low power signals existed that could simultaneously yield high probabilities of detection and accurate range/Doppler estimates [58] by increasing the signal duration and bandwidth through frequency modulation. These signals became known as Frequency Modulated Continuous Wave (FMCW). Some examples of applications and design architectures that are based FMCW signals were discussed in the literature as far back as the 60's (see [47, Ch. 16] and references therein), but the realization of small, low cost implementations have only recently become available [37, 46]. The hardware advancements that are enabling the practical realization of transmitters capable of generating low power, continuous wave (CW) signals with sophisticated frequency modulations is principally owed to the advancements in solid-state amplifier technology, while the ability to successfully process returns is a result of high speed digital signal processing hardware. These advancements in active radar systems have naturally perpetuated the ongoing battle that exists between the active radar technology and ELectronic Intelligent (ELINT) gathering methods, giving rise to the term Low Probability of Intercept (LPI) radar. The goal of this chapter is to present the experimental results associated with a new method for gathering ELINT on a specific LPI signal structure; the Linear Frequency Modulated Continuous Wave (LFMCW) .

The focus of Chapter 1 [19] of this dissertation was on the development of a computationally tractable method for Detecting, Characterizing and Localizing (DCL) low power, LFMCW [1] signals when the number of LFMCW signals and the autocorrelation properties of the noise corrupting them are both unknown. The method broke the signal down into small time intervals, estimated the frequency content in each, and through reconstructing the instantaneous frequencies that define the LFMCW signals, estimated the parameters that characterize them. In that work we evaluated the performance of the approach, which we termed a sequential method, via Monte Carlo simulation only. The focus of Chapter 2 [18] was on predicting the performance of any unbiased estimator of the parameters of the

[1] FMCW typically refers to radar signals whose amplitude is constant with arbitrary frequency modulation. In this work we focus on a specific type of frequency modulation that yields an instantaneous frequency that exhibits a periodic sawtooth structure. We refer to this structure as LFMCW.

LFMCW signals, via The Cramer-Rao Lower Bound (CRLB) [23, pgs. 35-48]. The focus of this chapter is on extending the performance evaluation to data generated and collected using Software Defined Radios (SDRs) [50]. These experiments demonstrate that this technique can reliably detect and characterize multiple low power LFMCW signals in the presence of noise processes with unknown autocorrelation properties and time-varying power spectra (*i.e.*, non-stationary) using low cost hardware. These results should be thought of as paving the way for higher cost, more capable receivers to attain even better performance.

3.1.1 LFMCW Signal Model

We demonstrate this capability by performing three types of experiments: (1) Simulation, (2) Hardware-in-the-Loop (HIL), and (3) Over-the-Air (OTA) on received signals that adhere to the following signal $s\left(t;\bar{\theta}\right)$ plus noise $w\left(t\right)$ model [9, 16, 17, 19, 34]

$$r\left(t\right) = s\left(t;\bar{\theta}\right) + w\left(t\right) \tag{3.1}$$

where the M component LFMCW signal is given by

$$s\left(t,\bar{\theta}\right) = \sum_{m=1}^{M} s_m\left(t;\bar{\theta}_m\right) \tag{3.2}$$

$$\bar{\theta} = \begin{bmatrix} \bar{\theta}_1^T & \dots & \bar{\theta}_M^T \end{bmatrix}^T \tag{3.3}$$

with each individual LFMCW signal expressed as

$$s_m\left(t;\bar{\theta}_m\right) = \Lambda_m\left(t;\bar{\theta}_m\right) \cdot b_m e^{j\Phi\left(t;\bar{\theta}_m\right)} \tag{3.4}$$

where

$$\Lambda_m\left(t;\bar{\theta}_m\right) = u\left(t - t_m^{\text{start}}\right) - u\left(t - t_m^{\text{stop}}\right) \tag{3.5}$$

is the rectangular function used to embed start-time t_m^{start} and stop-time t_m^{stop} of the signal into the model, with $u\left(t\right)$ denoting the unit step function. The instantaneous phase of each LFMCW component is given by

$$\Phi\left(t;\bar{\theta}_m\right) = \varphi_m + 2\pi \int_0^t f_m\left(\nu\right) d\nu \tag{3.6}$$

$$f_m\left(t\right) = f_m^{\text{start}} + \beta_m \bmod \left(t + \Delta_m, T_{\text{swp},m}\right) \tag{3.7}$$

We express $\Phi\left(t, \bar{\theta}_m\right)$ in (3.6) as a function of the integrated instantaneous frequency defined by (3.7) because our detection and characterization procedures are better suited to a definition of instantaneous frequency that does not require a time-derivative of instantaneous phase, due to the discontinuities in (3.7). Finally, the vector of unknown parameters defining the m^{th} LFMCW component is

$$\bar{\theta}_m = \begin{bmatrix} b_m \\ \varphi_m \\ t_m^{\text{start}} \\ t_m^{\text{stop}} \\ f_m^{\text{start}} \\ \beta_m \\ \Delta_m \\ T_{p,m} \end{bmatrix} \tag{3.8}$$

Unlike in [9, 34] we do not restrict $w(t)$ to be white, stationary, or have completely known autocorrelation properties. The autocorrelation function of a general zero mean noise process $w(t)$ is defined by

$$E\left\{w(t_1)\, w^*(t_2)\right\} = R(t_1, t_2) \tag{3.9}$$

for some function of t_1 and t_2. Accommodating a non-stationary noise process [8] means we **do not** require that $R(t_1, t_2) = R(t_1 - t_2) = R(\tau)$. Additionally, the fact that we can handle unknown (*i.e.*, possibly colored) noise processes means that we do not assume specific knowledge of the functional relationship defining the autocorrelation function [24, pg.581], [38, pg. 376] $R(t_1, t_2)$. One does need to assume that $w(t)$ is Gaussian, locally wide-sense stationary over short time intervals [8], and that the Power Spectral Density (PSD) associated with each short time interval does not change abruptly in frequency intervals of width $2W$. Both the width of a short time interval and the resolution bandwidth $2W$ are defined in Section 3.2. For the purposes of this chapter, noise processes that do not change abruptly on over frequency intervals of $2W$ are referred to as "slowly-varying" [40, pgs. 370, 498]. However, if any or all of these assumptions do not hold for some short time intervals,

we believe the overall method can recover[2]. We support this claim with a specific example, given by way of Figure 20 in Section 3.3.3, and with a more complete analysis in Section 3.4. Allowing for such unrestrictive assumptions on $R(t_1, t_2)$ is necessary in order to ensure successful operation the presence of real-world interference sources.

A numerical example of this signal model for $M = 2$ is shown using a spectrogram (*i.e.*, magnitude squared of the short-time Fourier Transform) in Figure 16 with parameters defined and specified in Table 10. Signal $m = 2$ is annotated in Figure 16 in order to show the relationship between the variables that define the signal and the instantaneous frequency of each component. The basic signal structure given in Figure 16 and Table 10 is used in each of the three experimental setups that we treat in this chapter. The following three sections describe the motivation for each of these experiments.

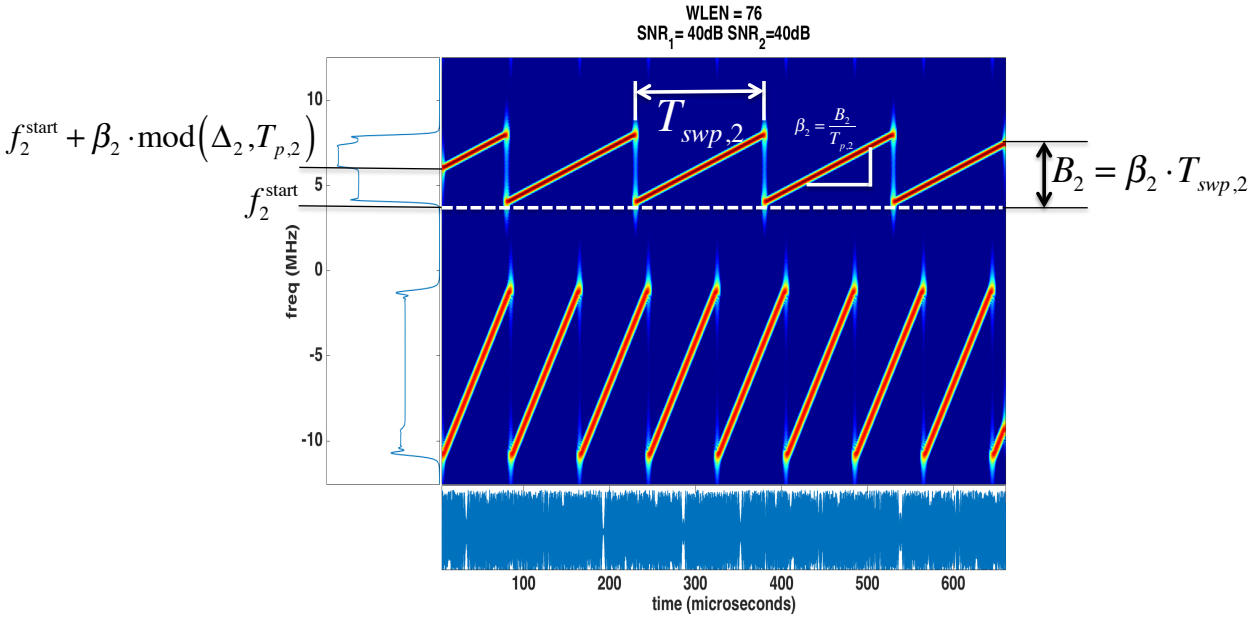

Figure 16: (main panel) Spectrogram (20 dB dynamic range) of two LFMCW signals based on signal model given in (3.1)-(3.8) (left panel) Frequency marginal of spectrogram (bottom panel) Time-series $r(t)$

[2]We believe that the assumptions for $w(t)$ to have a Gaussian distribution, locally stationary, and its power spectrum to be slowly-varying in intervals of width $2W$ can be broken some short time intervals. We describe the theoretical basis for this belief in briefly in Section 3.2, more fully in Section II of Part 1 [19], and test it more fully with the experimental setup given in Section 3.3 and results in Section 3.4

Table 10: Signal Parameters for Figure 16

Variable Description	Variable Name	Values (m=1, m=2)	Units
Amplitude	b_m	1, 1	$\sqrt{J/s}$
Phase	φ_m	0, 0	rad
Start Time	t_m^{start}	0, 0	μs
Stop Time	t_m^{stop}	655.36, 655.36	μs
Starting Frequency	f_m^{start}	-11, 4	MHz
Chirp-Rate	β_m	0.125, 0.0267	MHz/μs
Time-Offset	Δ_m	0, 75	μs
Sweep Time	$T_{\text{swp,m}}$	80, 150	μs
No. of LFMCW Signals	M	2	—
Sampling Rate	F_s	25	MHz
Observation Duration	T	655.36	μs

3.1.2 Simulation

In addition to aiding in the development of the theory, the simulation experiments are intended to establish a performance baseline in a completely controlled fashion. In particular, we controlled the experiment by generating signals in software that adhere to (3.1)-(3.7) and corrupting them with errors using a random number generator to model $w(t)$ for $R(t_1, t_2) = \sigma_w^2 \delta(t_1 - t_2)$ and $w(t) \sim CN(0, \sigma_w^2)$ (*i.e.*, Complex Additive White Gaussian Noise (CAWGN)).

The results of these experiments are used in two main ways. The first way, which was reported on in [19], is to demonstrate that reasonably good performance was achievable at input Signal to Noise Ratios (SNRs) low enough to support reception of signals emitted by transmitters with low Effective Isotropically Radiated Power (EIRP) at moderate ranges between transmitter and receiver, and a capacity for doing so in a computationally tractable way. Specifically, in [19] we showed that we could estimate the parameters of two simultaneous LFMCW signals to within 10% of their true values greater than 90% of the time when SNR is greater than -11 dB and the signals are sampled at 25 MHz over an observation period of 655.36 μs. Calculation of the signal parameters based on these samples required, at most, 35 seconds. These results supported the operational viability of the concept.

The second way we use these results is to assess the impact of error sources that are not reflected in a simple CAWGN model of $w(t)$. In particular, these error sources include

1. the non-ideal nature of real hardware used to generate and transmit LFMCW signals
2. extraneous interference sources that are superimposed on (3.2), and
3. multi-path reflections that are superimposed on (3.2).

More specifically, real hardware devices that generate Radio Frequency (RF) signals are likely to deviate from the mathematical models given in (3.1)-(3.5) and (3.6)-(3.7). Hence, software simulations that adhere to these models will not fully reflect errors that are likely present in practice. Next, since the RF spectrum is crowded in almost all bands, especially from the viewpoint of a receiver operating at high altitudes, extraneous in-band interference sources are likely to complicate $w(t)$ in ways that will tax our Gaussian assumptions and exercise our claims that neither stationarity nor knowledge of the autocorrelation function

is required. Finally, since LFMCW based radar systems are attractive in maritime [46] and in personnel-detection [47, pg. 16-35] applications, multi-path effects are likely to be faced by passive receivers situated near the ground. The goal of the HIL and OTA experiments is to quantify the performance loss, relative to simulation and the CRLB, from these sources of error.

3.1.3 Hardware-in-the-Loop

The primary goal of the HIL experiment is to determine whether a laboratory experiment using real RF hardware and an RF noise generator could approach the performance observed in simulation. The results of this experiment answer the question of whether the ideal model of the instantaneous frequency given (3.7), and upon which we develop our methods, are reasonable approximations for what can be generated and received by real hardware.

3.1.4 Over-the-Air

The primary goal of the OTA experiment is to determine the performance loss incurred by the presence of interference and multi-path [3] error sources that make the characteristics of $w(t)$ deviate significantly from $R(t_1, t_2) = \sigma_w^2 \delta(t_1 - t_2)$ and $w(t) \sim CN(0, \sigma_w^2)$.

3.1.5 Chapter Organization

With the intent of the chapter clearly stated, the remainder of this work is devoted to describing the experimental results. To do so the chapter is organized as follows. In Section 3.2 we overview the key aspects of the method reported on in [19] for the sake of motivating the experimental methods and results, which are found in Sections 3.3 and 3.4, respectively. More specifically, Section 3.3 details the approaches used for gathering the results, which

[3] Because multi-path is a practical concern for some operational environments, we have stressed that, in principle, it may have a performance impact. However, because of the proximity of the transmitter, receiver, and multi-path reflectors in the experimental setup used, if present at all, multi-path components were not clearly evident in any of the recorded data. In other words, if multi-path was present in the data, it was either too weak to be observed or the reflectors were at such short ranges that they were indistinguishable from the direct path. In any case, multi-path did not appear to affect performance. Further simulation analysis is planned to determine the performance impact of multi-path reflections.

are described in Section 3.4. Finally, we conclude the chapter with a summary of the main accomplishments and some suggestions for future directions.

3.2 BACKGROUND

The purpose of this section is to summarize the main ideas behind the algorithms developed in Chapter 1 [20], but only enough to so that the approach can be understood from a high level. For a detailed description of the method the reader is encouraged to consult [19] directly. Since the experimental results discussed in this chapter are primarily focused on detection and characterization performance, we focus our background discussion on these aspects of the method.

The problem of detecting and characterizing LFMCW signals is most commonly formulated as [9, 25, 34]

$$\mathcal{H}_0: \quad r(t) = w(t) \tag{3.10}$$

$$\mathcal{H}_1: \quad r(t) = s\left(t;\bar{\theta}\right) + w(t) \tag{3.11}$$

where $s\left(t;\bar{\theta}\right)$ and $w(t)$ are as defined previously. When $w(t)$ is complex, white, and Gaussian the optimal detection and characterization method, in the Generalized Likelihood Ratio Test (GLRT) sense, is given in [25] as

$$\mathcal{L} = \arg\max_{\bar{\theta}} \left| \int_0^T r(t)\, s^*\left(t;\bar{\theta}\right) dt \right|^2 > \eta \tag{3.12}$$

In accordance with the theory associated with the GLRT, if the threshold η is exceeded by $\left| \int_0^T r(t)\, s^*\left(t;\bar{\theta}\right) dt \right|^2$ then a detection is declared and the value of $\bar{\theta}$ corresponding to the maximum is the Maximum Likelihood Estimate (MLE) of $\bar{\theta}$. While optimal, methods that implement (3.12) directly, or thorough surrogates such as in [9], face computational difficulty when $M = 1$ and are overwhelmed by computational burdens when M is unknown, as is likely the case in practice. Furthermore, since deriving the probability density function of $\mathcal{L}$ under $\mathcal{H}_0$ appears to be an unsolved problem, these methods give no guidance for selecting η, which puts false alarm probability (P_{FA}) outside of the direct control of the designer [4].

[4]The detection threshold can be selected using simulation or laboratory experiments that attempt to model the worst-case environment, or through built-in test equipment that senses the noise environment when it is known that no signal is present. While practical, these methods can be error prone or costly. Hence, it would desirable to have detection and characterization methods that do not require detailed knowledge of the noise properties to *automatically* select η.

These and other methods designed for tracking instantaneous frequency [26] suffer from one or more of the following shortcomings

1. their computational requirements limit their applicability
2. unable to handle case for $M > 1$
3. unable to automatically and efficiently determine M when unknown
4. unable to set a meaningful detection threshold the without detailed knowledge or assumptions on the statistics of the additive noise process
5. unable to operate at input SNR levels characteristic of LPI radar signals
6. unable to provide parameter estimates to enable localization, such as TDOA, without access to high speed data links

The detection problem framed in (3.10) and (3.11) and the time series defined by (3.1)-(3.8) are useful starting points to formulate the detection and characterization problems. However, the complexity of the signal model (*i.e.*, the large number of parameters necessary to define it) and the number of samples required for a LFMCW signal to show its full structure in a received time series makes the optimal correlation-based procedures, like that given in (3.12), prone to computational limitations. In an effort to retain the essence of this problem formulation but redefine it so that some of these short comings could be addressed, we adopted an idea common in the field of *spectral estimation* [40,49] that attempts to extend methods developed for stationary processes to processes with time-varying spectral content. In particular, our approach breaks the signal down by analyzing the received time-series over short, non-overlapping, contiguous time intervals. Within each interval the received signal is modeled as a sum of harmonic components with unknown order and, within this interval, zero or more frequency components are detected and estimated. The new hypothesis test, based on the redefined signal model is [5]

$$\mathcal{H}_0 : \; r_{n,k} = w_{n,k} \tag{3.13}$$

$$\mathcal{H}_1 : \; r_{n,k} = s_{n,k} + w_{n,k} \tag{3.14}$$

[5] We reference the samples of a continuous time signal $x(t)$ as $x_n \equiv x(t)\,|_{t=nT_s}$

where $r_{n,k}$ is the n^{th} sample, for $n = 1, \ldots, N$, of $r(t)$ in the k^{th} time interval. The k^{th} time interval is given by

$$\delta_k = \{t : (k-1) \cdot NT_s \leq t \leq k \cdot NT_s)\} \tag{3.15}$$

where $k = 1, \ldots, K = \lfloor \frac{T}{N \cdot T_s} \rfloor$ and $T_s = \frac{1}{F_s}$. That is,

$$r_{n,k} = r\left((k-1) \cdot NT_s + nT_s\right) \tag{3.16}$$

Within this short time interval we model the received time-series as a sum of harmonic components, which gives rise to the so-called Short Time Harmonic Model (STHM) of the data

$$s_{n,k} \approx s_{n,k}^{\text{STHM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi f_{c,k} n T_s} \tag{3.17}$$

where $L_k \leq M$ is the number of constant frequencies $f_{c,k}$ present in δ_k. The complex amplitude is given by $C_{c,k} = b_{c,k} e^{j\phi_{c,k}}$ where $b_{c,k}$ is as defined in (3.4) and $\phi_{c,k}$ is the phase of the c^{th} harmonic component in the k^{th} short-time interval. We emphasize that L_k and each $b_{c,k}$, $f_{c,k}$, and $\phi_{c,k}$ are assumed to be unknown, but deterministic. From this point forward we take the approximation in (3.17) to be an equality. Furthermore, in δ_k we assume that $w_{n,k}$ is a band-limited, zero mean, complex, wide-sense stationary Gaussian random process with unknown PSD $P_{w_{n,k}}(f)$ that is not necessarily constant (*i.e.*, $\{w_{n,k}\}$ can be colored noise), but is slowly-varying. [6] The ability to handle an unknown colored noise process is what justifies the generality with which we can treat (3.9) , and is a key feature for a practical implementation since the PSD of the noise is rarely known *a-priori*. Hence, the STHM of the received multicomponent LFMCW signal in δ_k is

$$r_{n,k} = s_{n,k}^{\text{STHM}} + w_{n,k} \tag{3.18}$$

So, within δ_k selection between $\mathcal{H}_0$ and $\mathcal{H}_1$ can be achieved using Thompson's multi-taper based method for harmonic analysis [51]. Collecting the frequency estimates from each δ_k

[6] The assumption that the PSD is slowly-varying is a mathematical necessity for the development of the test statistic used to decide between (3.13) and (3.14). However, in practice we only need the PSD to not exhibit dramatic variations in subintervals of the PSD having width $2W$, where $2W$ is the resolution bandwidth defined by Thomson's method. If these assumptions are met in most of the short-time intervals then the small percentage of missed or erroneous detections that are a result of an instability in the detector can be tolerated by the later steps devoted to characterization, which are outlined in detailed in Section III of [19].

over the full observation duration T amounts to measuring the instantaneous frequency of each LFMCW signal. Once these frequency estimates are obtained they are clustered by chirp-segment by finding those frequency estimates that fall along the same straight line. The method estimates the number of unique line segments and the slope and frequency intercept of of each segment. The slope, frequency intercept, and corresponding frequency estimates are compared to each of the other clusters in order to determine which chirp segments correspond to the same LFMCW signal. Two statistical test are applied to these data to make the associations. Once each cluster is assigned to an LFMCW signal, the samples for each signal are used to estimate the parameters that define (3.7) by solving the following optimization problem

$$\underset{f_m^{\text{start}},\beta_m,\Delta_m,T_{\text{swp,m}}}{\arg\min} \left\{ \sum_{i=1}^{N_m} \left| \hat{f}_{m,i} - f_{m,i} \right| \right\} \tag{3.19}$$

where $f_{m,i}$ denotes the i^{th} sample of the instantaneous frequency of the m^{th} LFMCW signal and $\hat{f}_{m,i}$ are the corresponding frequency estimates that resulted from the preceding STHM, clustering, and association steps.

This simplification of the signal model and the algorithms developed to use the frequency samples to estimate LFMCW parameters led to a sequential method with the following features

1. no calibration is necessary, either on-line or off-line in the laboratory, to determine a detection threshold η that is a function of P_{FA};
2. one does not need to assume the system noise present in the received time-series is white or has a known autocorrelation function;
3. environments containing multiple LFMCW signals can be treated without adding a stifling computational burden;
4. automatic recognition of the number ($M = 0, 1, 2 \dots$) of LFMCW signals present in the environment is achieved
5. in addition to estimates of the parameters that define the instantaneous frequency of each LFMCW signal, our approach produces estimates of the variance of a subset of these parameters directly from the data;

However, with these benefits comes a loss in performance relative to the GLRT and the associated MLE, which we quantify in Section 3.4. We do so by comparing performance results from simulation, HIL, and OTA experiments to one another and to the CRLB, which was derived in [18]. Comparison to the CRLB is a useful benchmark since the CRLB is known to reflect the minimum estimation error variance that can be achieved by any parameter estimation algorithm that is unbiased [21, Ch. 3].

3.3 EXPERIMENTAL METHODS

In this section we describe the details associated with each experimental setup. In particular, we describe how signals were generated, transmitted, received, and processed for each of the three experiments. Figures 17-19 are block diagrams that describe how the components are connected for the simulation, HIL, and OTA experiments. By comparing Figures 17-19 we see that the common aspects between the three experiments are that each "transmit" the same two simultaneous signals (*i.e.*, $m = 1, 2$) , with parameters given in Table 10 and do all of the calculations necessary to estimate $\hat{\theta}_1$ and $\hat{\theta}_2$ via post-processing using MATLAB R015a. [7] The experiments differ in the means by which the signals are generated, the medium over which they propagate, and the way they are sampled. The following three sections elaborate on each of these three experimental setups, primarily in terms of 1) how SNR was computed or estimated and 2) how repeated trails were initiated.

3.3.1 Simulation

The goal of the simulation was to divorce the performance of the algorithms from errors sources that could not be completely controlled. Examples of such error sources are imperfections in the hardware that cause the transmitted and received signals, in the absence of additive noise like that modeled by $w(t)$, to deviate from (3.1)-(3.5) and (3.6)-(3.7). Additionally, in the real world we expect $w(t)$ to not only have an unknown noise auto-correlation function, but also exhibit non-stationarities resulting from extraneous interference sources.

3.3.1.1 Controlling SNR The blocks in Figure 17 that contain "lfmcw.m" and "cawgn.m" emphasize that we developed MATLAB functions and scripts to generate controlled signals and noise, with direct control over SNR. The parameters that define signals s_1 and s_2 are given in Table 10 and the variance of the noise samples produced by the random number

[7] All post processing and timing values reported on in this chapter were preformed in MATLAB using a MacPro 2.6 GHz Core i7 laptop wi th 16GB of 1600 MHz DDR3 RAM. The purpose of alerting the reader to the approximate computation times and the platform they were derived from is to simply give a ballpark sense for the computational burden to support the general claims that the method were evaluating is computationally feasible.

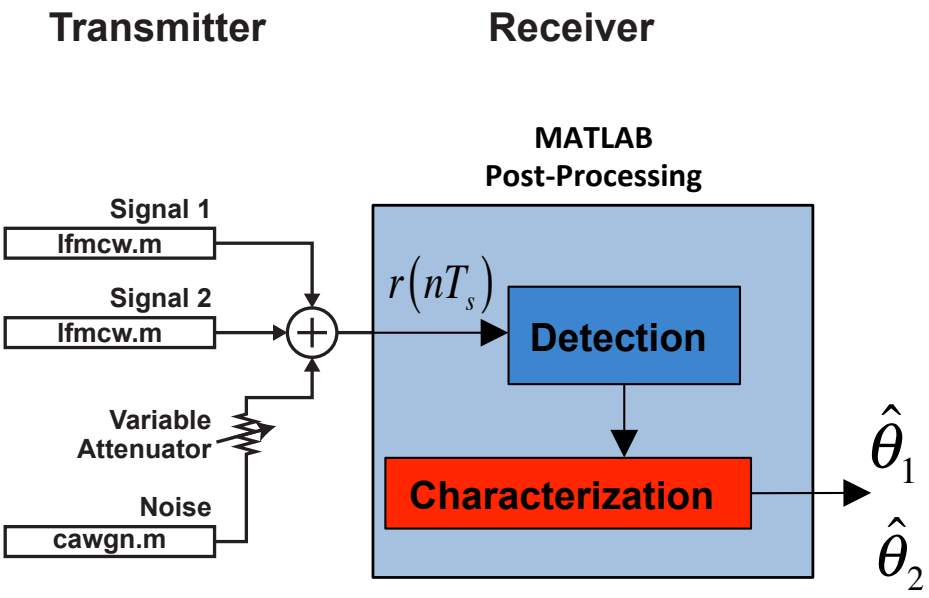

Figure 17: Block Diagram of Simulation Experimental Setup

generator to model $w(t)$ are scaled to yield a variance given by

$$\sigma_w^2 = \frac{b_m^2}{\text{snr}} \tag{3.20}$$

where $b_m^2 = 1$ and snr is the signal to noise ratio, in linear units[8], of $r(nT_s)$. For all three experiments, we vary SNR and observe various metrics, which are defined in Section 3.4, to assess performance. By scaling the output of the random number generator, which is signified by the block with the arrow though it denoted "ATTN," to abbreviate "attenuation," we vary SNR by keeping the signal amplitude constant.

3.3.1.2 Repeating Trials Since the simulation experiment was performed entirely in software, we were able to repeat trials for a specific SNR with a new set of random draws from MATLAB's random number generators. The results given for this experiment are based on 200 independent trials.

[8]We capitalized SNR when referencing it in units of decibels (dB).

3.3.2 Hardware-in-the-Loop

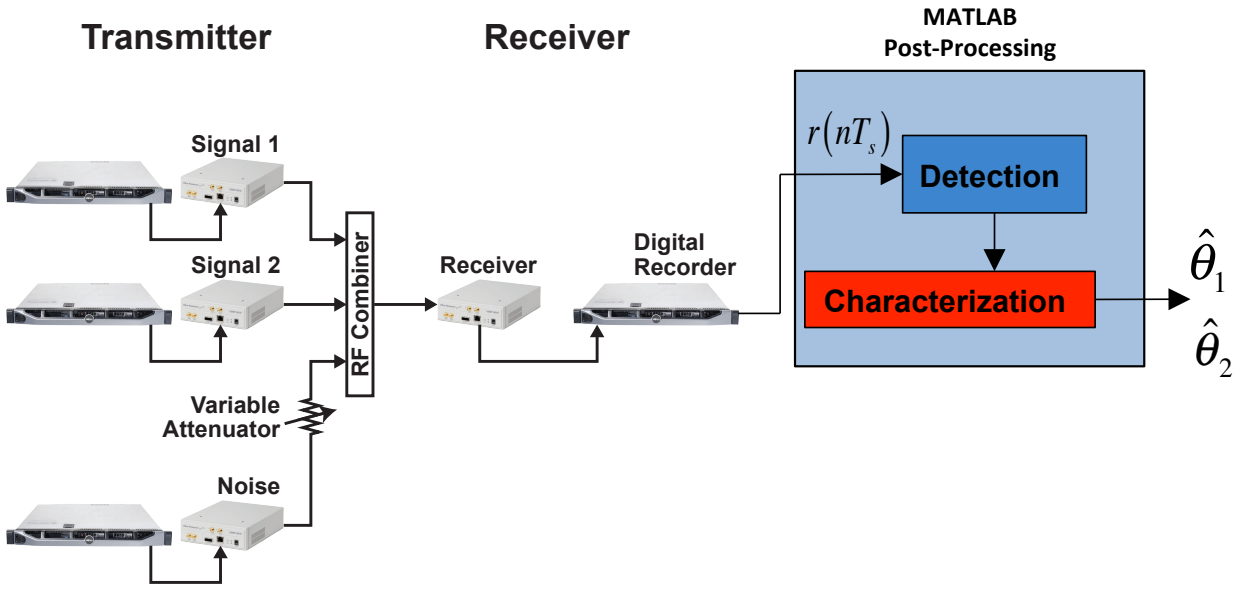

Figure 18: Block Diagram of HIL Experimental Setup

The goal of the HIL experiment was to determine whether the algorithms we developed in [19] could, in the absence of extraneous interference, perform nearly as well as in simulation when real hardware was employed for the transmission and reception of the LFMCW signals. We avoided most interference by designing the experiment so all devices were connected via shielded RF cabling and by adding controlled noise to corrupt the received signal. After sampling and saving the baseband signals to disk, we used MATLAB to post-process each $T = 655.36\mu s$ segment of sampled data using the same algorithms developed and tested in simulation.

Figure 18 shows how the SDR's [50] were connected and controlled in addition to how the data were received and corrupted. One SDR was programmed to produce $s_1(t)$ and the other to produce $s_2(t)$ on a 3 GHz carrier. These signals were then superimposed using an RF combiner and allowed to propagate over an RF cable to the input of another RF combiner. The output of this combiner, laboratory generated noise was superimposed and the total signal plus noise (as in (3.1))was received, down-converted to baseband, and sampled by a third SDR of the same type as those used to generated the LFMCW signals.

In order to ensure we observed a wide range of SNR's, the noise source was attenuated from 0 dB to 50 dB in steps of 1 dB. After each attenuation was applied the receiver SDR was commanded, by the Digital Recorder, to record $N_{total} = 24 \cdot 2^{20}$ samples, which is approximately 1 second at 25 Mega Samples Per Second (MSPS). We also note that both the in-phase and quadrature components of the complex baseband signal produced by the SDR were sampled at 25 MSPS, yielding 25 MHz of bandwidth. Each attenuation step corresponded to a unique 1 second data file, which was divided into 1536 segments of length $N_{T_{obs}} = 2^{14}$ in order to produce observations of duration $T_{obs} = 655.36\mu s$ to be post-processed in MATLAB. Hence, each attenuation level contains enough samples for a maximum of $P = 1536$ repeated trials for performance analysis.

3.3.2.1 Controlling SNR Since direct control over the SNR was not possible in the same way it was in the simulation, we needed to estimate it instead. To do so we took 1 second noise only recordings (*i.e.*, s_1 and s_2 turned off) at each attenuation level and estimated the noise variance at the output of the receiver's Analog to Digital Converter (ADC) in each of the P segments of length $N_{T_{obs}}$. In particular, for each attenuation level α we computed P estimates of noise variance

$$\hat{\sigma}^2_{w_\alpha,p} = \frac{1}{N_{T_{obs}} - 1} \sum_{n_p=0}^{N_{T_{obs}}-1} \left| r_\alpha \left(n_p T_s \right) - \hat{\mu}_{w_\alpha,p} \right|^2 \tag{3.21}$$

$$\hat{\mu}_{w_\alpha,p} = \frac{1}{N_{T_{obs}}} \sum_{n_p=0}^{N_{T_{obs}}-1} r_\alpha \left(n_p T_s \right) \tag{3.22}$$

where n_p is intended to denote the indices of p^{th} length $N_{T_{obs}}$ segment of r_α. To compute the uncertainty in these noise variance estimates, we computed the sample variance over the P noise variance estimates, for each attenuation. Specifically,

$$\Delta\hat{\sigma}^2_{w_\alpha} = \sqrt{\frac{1}{P-1} \sum_{p=0}^{P-1} \left(\hat{\sigma}^2_{w_\alpha,p} - \hat{\mu}_{\hat{\sigma}^2_{w_\alpha}} \right)^2} \tag{3.23}$$

$$\hat{\mu}_{\hat{\sigma}^2_{w_\alpha}} = \frac{1}{P} \sum_{p=0}^{P-1} \hat{\sigma}^2_{w_\alpha,p} \tag{3.24}$$

Similarly, to estimate signal power b_m^2, we disconnected the lab noise generator, transmitted only s_1 or s_2, and computed

$$\hat{b}_{m,p}^2 = \frac{1}{N_{T_{obs}} - 1} \sum_{n_p=0}^{N_{T_{obs}}-1} \left| r\left(n_p T_s\right) - \hat{\mu}_{s_m,p} \right|^2 \tag{3.25}$$

$$\hat{\mu}_{s_m,p} = \frac{1}{N_{T_{obs}}} \sum_{n_p=0}^{N_{T_{obs}}-1} r\left(n_p T_s\right) \tag{3.26}$$

Similar to (3.23) and (3.24), we estimated the uncertainties in these signal power estimates as

$$\Delta\hat{b}_m^2 = \sqrt{\frac{1}{P-1} \sum_{p=0}^{P-1} \left(\hat{b}_{m,p}^2 - \hat{\mu}_{\hat{b}_m^2} \right)^2} \tag{3.27}$$

$$\hat{\mu}_{\hat{b}_m^2} = \frac{1}{P} \sum_{p=0}^{P-1} \hat{b}_{m,p}^2 \tag{3.28}$$

The estimated SNR for each attenuation is then given by

$$\widehat{\mathrm{SNR}}_{m,\alpha} = 10 \cdot \log_{10} \left(\frac{\hat{\mu}_{\hat{b}_m^2}}{\hat{\mu}_{\hat{\sigma}_{w_\alpha}^2}} \right) \tag{3.29}$$

To determine the uncertainty associated with (3.29) due to $\Delta\hat{b}_m^2$ and $\Delta\hat{\sigma}_{w_\alpha}^2$, we employ the following approximation [21, pg. 39], which is based on a first-order Taylor expansion of $\widehat{\mathrm{SNR}}_{m,\alpha}$ about $\left(\hat{\mu}_{\hat{b}_m^2}, \hat{\mu}_{\hat{\sigma}_{w_\alpha}^2}\right)$

$$\Delta\widehat{\mathrm{SNR}}_{m,\alpha} = \sqrt{\begin{array}{l} \left(\Delta\hat{b}_m^2 \cdot \frac{\partial \mathrm{SNR}_{m,\alpha}}{\partial \mu_{\hat{b}_m^2}} \Big|_{\mu_{\hat{b}_m^2} = \hat{\mu}_{\hat{b}_m^2}} \right)^2 + \quad \cdots \\ \left(\Delta\hat{\sigma}_{w_\alpha}^2 \cdot \frac{\partial \mathrm{SNR}_{m,\alpha}}{\partial \mu_{\hat{\sigma}_{w_\alpha}^2}} \Big|_{\mu_{\hat{\sigma}_{w_\alpha}^2} = \hat{\mu}_{\hat{\sigma}_{w_\alpha}^2}} \right)^2 \end{array}} \tag{3.30}$$

We quantify algorithm performance for this experiment in Section 3.4 as a function of these estimates of SNR. The uncertainties in SNR are used to develop the error bars [9] (*i.e.*, confidence intervals) in Figure 24.

[9] The purpose of breaking up the signal and noise time series up and computing multiple estimates of signal and noise power was so that the variance estimates of signal and noise power could be computed and used to develop the error bars. Additionally, since the algorithms being tested are applied to intervals of equal duration as those used to compute each estimate of signal and noise power, variations in SNR levels that occur from interval to interval are captured on the same time scale as used to derive parameter estimates.

3.3.2.2 Repeating Trials In order to repeat trials for each SNR in this experiment we treated each of the 1536 segments for each of the 1 second recordings as a candidate trial. For the results given in 3.4 we randomly selected 200 segments per attenuation setting to derive performance results.

3.3.3 Over-the-Air

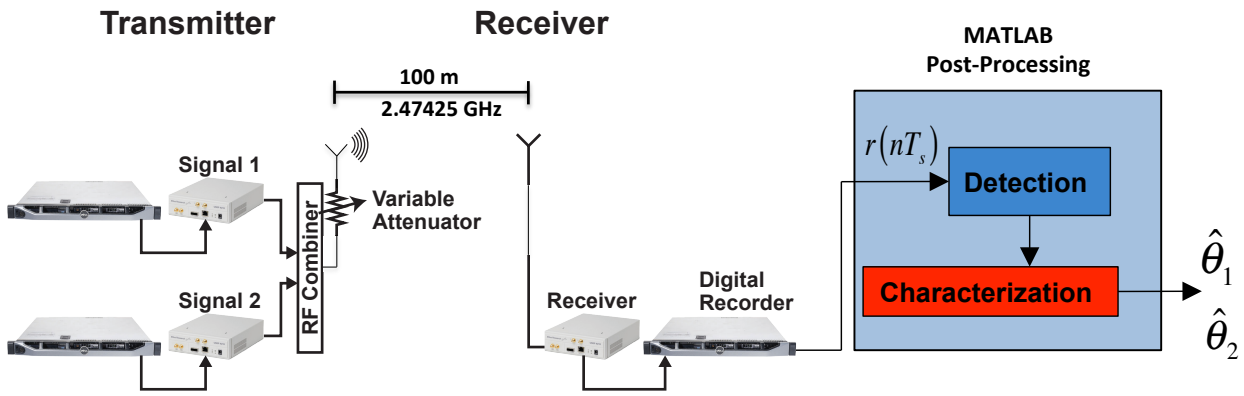

Figure 19: Block Diagram of OTA Experimental Setup

In contrast to the HIL experiment, the goal of the OTA experiment was to face all the same challenges present in the HIL in addition to those posed by a *wireless* link between unsynchronized transmitter and receiver. Consequently, two additional challenges followed

1. extraneous interference sources whose properties were unknown *a-priori* were intermittently present in $r(t)$
2. synchronization between the Local Oscillators (LO's) on the transmit and receive ADC's was not obtained

The first challenge was brought on by the general crowded nature of the RF spectrum in the 25 MHz band centered at 2.47425 GHz. The reason this band is particularly crowded is because it is allocated by the Federal Communications Commission (FCC) as an the Industrial, Scientific, and Medical (ISM) band. The ISM designation allows for RF energy to be radiated without a license if the Effective Isotropic Radiated Power (EIRP) [45, pg. 66, 80] is less than 1 Watt and the signal structure is such that the energy is spread, in

our application, over the 25 MHz of bandwidth that our signals occupy. Hence, a large number of devices (*e.g.*, Cordless Phones, Bluetooth, near field communications, and wireless computer networks) operate at low powers in these bands, which offers a rich set of potential interference sources against which we can pit our algorithms. An example of a 655.36 μs recording taken at an attenuation level such that the Signal to Interference Plus Noise (SINR) levels are $\text{SINR}_1 = -6.3$ dB, $\text{SINR}_2 = -9.6$ dB, and where heavy interference was present is given in Figure 20. It is interesting to note that this segment of data led to a correct characterization, which is defined in Section 3.4, for both signals. We emphasize that in this section that we are forced to deal with SINR rather than SNR, since interference and noise cannot be separated in the experiment.

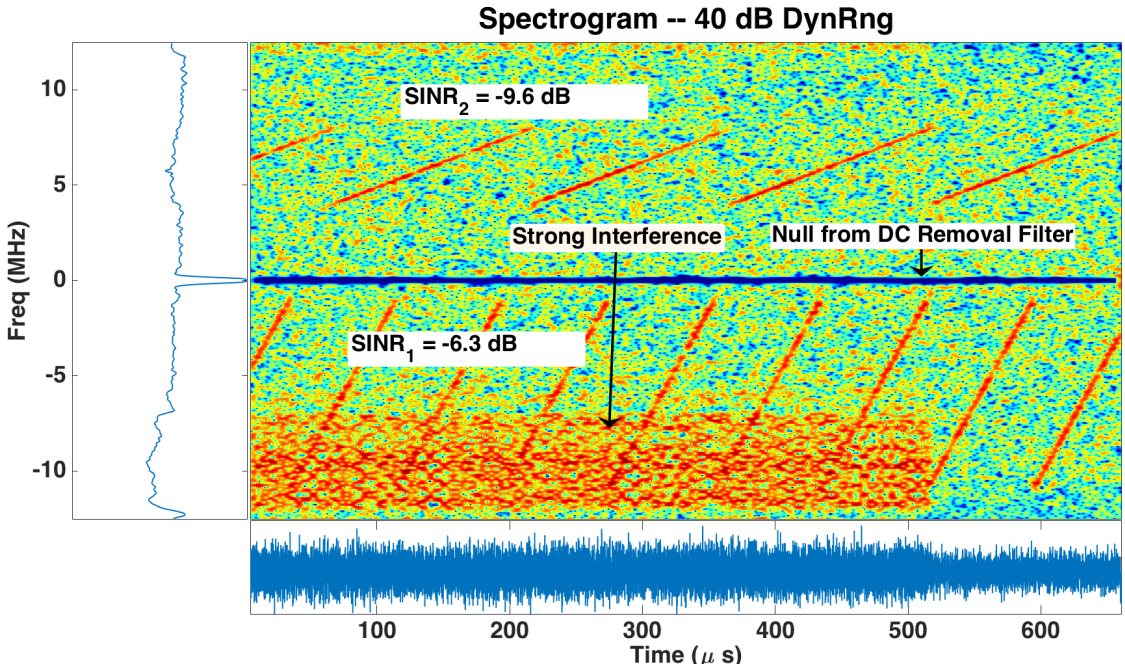

Figure 20: (main panel) Spectrogram of s_1 and s_2 observed during OTA experiment at low SNR levels in the presence of strong interference and ambient noise (left panel) Magnitude spectrum of associated $r\,(t)$ (bottom panel) Real part of time-series $r\,(t)$

The second challenge is related to the intentional lack of synchronization between the LO's on the transmitter and receiver. The reason we intentionally did not synchronize the LO's was because doing so would have produced received signals whose structure would have been unrealistically close to that of the transmitted signals since synchronization would have meant virtually no discrepancy, for example, between the tuned center frequency of the transmitter and the receiver. Since it's unlikely that a practical implementation of our method will be able to synchronize with the clock of the transmitter, we avoided doing so in our test. The result is difficult to perceive visually, but manifests in an instability in the center frequencies of each signal. As a result, parameters like starting frequency f_m^{start} will be affected by this systematic error source in a way that's not directly reflected in a CAWGN model, but likely present in practice.

3.3.3.1 Controlling SINR In contrast to the data collection strategy used in the HIL experiment, the OTA experiment affects SINR by attenuating the transmitted signal rather than the additive interference plus noise. Doing so is intended to emulate power loss that occurs due to propagation over long distances, as would be the case in real world scenarios where the receiver is a long distance away from the transmitter. In these scenarios, there is no control over the noise power once the operating bandwidth is selected. Since the physical distance between the transmitter and receiver, both of which were on rooftops to increase the likelihood of observing interference sources, was only approximately 100 meters, power levels were reduced by 0 to 30 dB in steps of 1 dB with a variable attenuator at the transmitter. We sampled the signal environment at a center frequency of 2.47425 GHz at a rate of 25 MHz. In the OTA data collection we opted to lengthen the duration of the recordings for each attenuation level from approximately 1 second to approximately 5 seconds, which increased the number of samples we collected to $N_{total} = 119 \cdot 2^{20}$. The reason for doing so was to increase the likelihood of observing interference in each recording.

To estimate interference plus noise variance in this experiment, $Q = 14$ individual 5 second recordings were taken intermittently over a 1 hour period during intervals of time where the transmitter was turned off. These interference plus noise recordings were taken over the span of 1 hour in order to capture interference events that were as diverse as possible, without collecting a burdensome amount of data. The calculations used to estimate interference plus noise variance were similar to those used in the HIL experiment. In particular, since we planned to apply our detection and characterization algorithms to observation intervals such that $T_{obs} = 655.36\mu s$, we partitioned the total N samples corresponding to 5 seconds of noise into $N_{total}/N_{T_{obs}} = P = 7616$ segments. Hence, for $q = 1, \ldots, Q$ we computed noise power estimates as

$$\hat{\sigma}^2_{w_q,p} = \frac{1}{N_{T_{obs}} - 1} \sum_{n_p=0}^{N_{T_{obs}}-1} \left| r_q\left(n_p T_s\right) - \hat{\mu}_{w_q,p} \right|^2 \tag{3.31}$$

$$\hat{\mu}_{w_q,p} = \frac{1}{N_{T_{obs}}} \sum_{n_p=0}^{N_{T_{obs}}-1} r_q\left(n_p T_s\right) \tag{3.32}$$

where n_p again indexes the p^{th} segment, but this time of the q^{th} interference plus noise recording, r_q. The uncertainty in the interference plus noise power estimate is then computed

over all PQ individual estimates as [10]

$$\Delta\hat{\sigma}_w^2 = \sqrt{\frac{1}{PQ-1}\sum_{i=0}^{PQ-1}\left(\hat{\sigma}_{w,i}^2 - \hat{\mu}_{\hat{\sigma}_w^2}\right)^2} \tag{3.33}$$

$$\hat{\mu}_{\hat{\sigma}_w^2} = \frac{1}{PQ}\sum_{i=0}^{PQ-1}\hat{\sigma}_{w_\alpha,p}^2 \tag{3.34}$$

To estimate signal power $\hat{b}_m^2$ we take a very similar approach to that taken to estimate the signal powers in the HIL experiment via (3.25)-(3.28). In particular, we transmit each signal individually then compute the signal power $\hat{b}_m^2$ as in (3.25) and (3.26). To determine the signal power associated with each of the attenuated values we scale these estimates by the appropriate amount. Similarly, we compute the uncertainty in the estimates of $\hat{b}_m^2$ as in (3.27), which we assume is the same for each attenuation level. Finally, the uncertainty in the SINR estimate is computed as in (3.29).

3.3.3.2 Repeating Trials In order to repeat trials for each SINR in this experiment we treated each of the 7616 segments for each of the 5 second recordings as a candidate trial. For the results given in 3.4 we randomly selected 200 segments per attenuation setting to derive performance results.

[10] We now index noise power estimates as $\hat{\sigma}_{w,i}^2$ from $i = 1, \ldots, PQ$ to reflect the PQ independent estimates obtained over all Q recordings.

3.4 RESULTS

Performance[11] is quantified in two ways for each of the three experiments using parameter estimation error variance and probability of correct characterization (P_C). In general, as reflected in Table 10 and in (3.8), there are 8 parameters that define each LFMCW signal. Though each parameter is estimated at each trial for each experiment, we only evaluate the performance of the algorithm in terms of chirp-rate, starting frequency, and sweep time. We limit our performance analysis to this subset of parameters because it is difficult to know truth for the others. Furthermore, these three parameters are the main identifiers that would most likely be of interest to an ELINT operation. Estimation error variance is computed for each of the three parameters, for each of the two signals, at each SNR[12], over the 200 trials. In trials where s_1 or s_2 is not detected, there is no contribution to the associated estimation error variance.

Each of the three panels in Figures 21-23 show the estimated standard deviations derived from each of the three experiments along with the associated CRLB, which is the same for each experiment since the true parameter values are the same. As mentioned previously, comparison to the CRLB is a useful benchmark since it is known to reflect the minimum estimation error variance that can be achieved by any parameter estimation algorithm that is unbiased [21, Ch. 3]. It is also known that if there exits an algorithm that can achieve the CRLB, then the MLE will also achieve it asymptotically (*i.e.*, as the number of samples of $r(t)$ increases) [21, pg. 164]. Hence, if an estimator exists that can achieve the CRLB then the CRLB predicts the asymptotic performance of the MLE and serves as a means to assess the degree of sub-optimality we suffer compared to the more computationally intensive, but optimal, MLE. If an estimator does not exist that can achieve the CRLB, then the CRLB still provides a useful benchmark for comparison that is widely accepted in the literature.

[11] In order to implement the algorithms that produce these various results, a number of parameter settings are required. Since a detailed description of each of these parameters would require that we provide contextual details that are outside of the scope of this chapter and would serve only to shift focus from the results, we refer the reader to Part 1 [19] where these parameters were first defined. We also note that the parameter settings given in [19] are the same settings used to derive the results in this section for each of the three experiments.

[12] For simplicity, in this section we refer to performance as a function of SNR. However, the reader should take SNR to mean SINR when considering performance results from the OTA experiment.

The CRLB for the signals treated in this chapter was derived in Chapter 2 [18]. In general, the more samples of a process that are observed the better the estimation error achieved. The CRLBs reflected in Figures 21-23 are computed using all $N_T = T \cdot F_s = 16384$, while the sequential method is computed using a series of non-overlapping intervals of length $N = 76$. Hence, the majority of the performance gap that is apparent in these figures is a consequence of the disparity between the number of samples used to form the estimates. However, the short time harmonic model that requires N to be small is also the main mechanism for the all of the benefits yielded by the sequential method; in particular its computational tractability compared to the MLE.

Probability of correct characterization, per signal, is defined as the average number of trails per SNR level that yield parameter estimates that are within 10% of their true values. Figure 24 shows the estimated values of P_C for each experiment, with the left panel corresponding to s_1 and the right panel to s_2. The results from the HIL and OTA experiments include confidence intervals, centered at each estimated SNR, that are derived from the uncertainties in SNR computed according (3.29) and described in Sections 3.3.2.1 and 3.3.3.1. The confidence intervals widths are selected, assuming the errors in SNR are Gaussian, to have a 95% probability of containing the true SNR. Hence, we are able to quantify experimental uncertainty, thereby giving a basis from which to make inferences about whether one experiment yielded statistically different performance than another. Each solid dot in the this figure corresponds to an experimental data point, about which 95% confidence intervals are centered for the HIL and OTA experiments. The solid lines are least-squares fits of these data points to a logistic curve of the form [55]

$$f\left(\text{SNR}\right) = \frac{1}{1 + e^{-\kappa(\text{SNR}-\text{SNR}_{0.5})}} \tag{3.35}$$

where κ is a scale parameter and $\text{SNR}_{0.5}$ denotes the SNR where $P_C = 0.5$. We included these fits to the data for ease of comparison between experiments that did not occur at the same levels of SNR. Figure 24 is the most concise graphical summary of the performance of the algorithms in each of the three experiments.

A common method for quantifying the performance of an algorithm in terms of a single number is by defining the Minimum Detectable Signal (MDS). We define MDS as the SNR

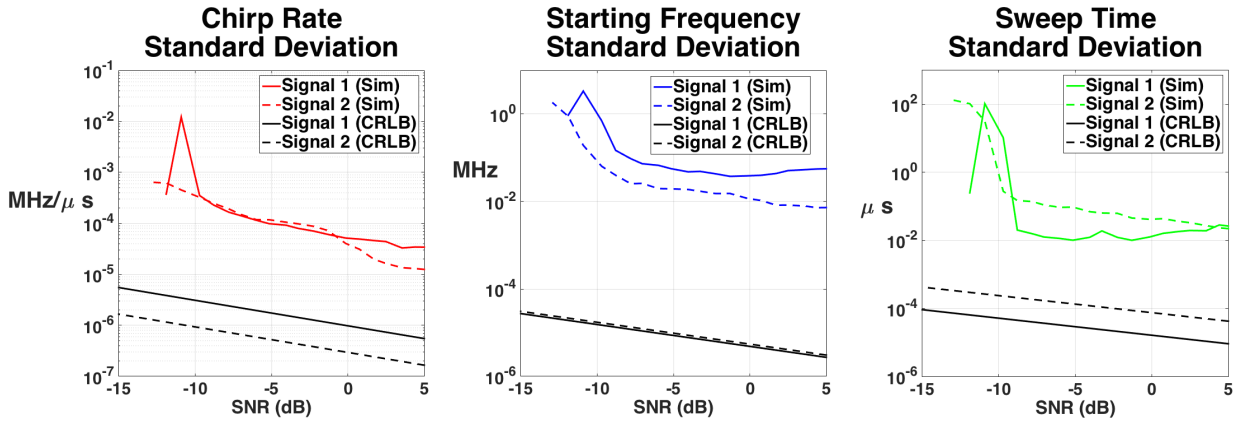

Figure 21: Simulation Results

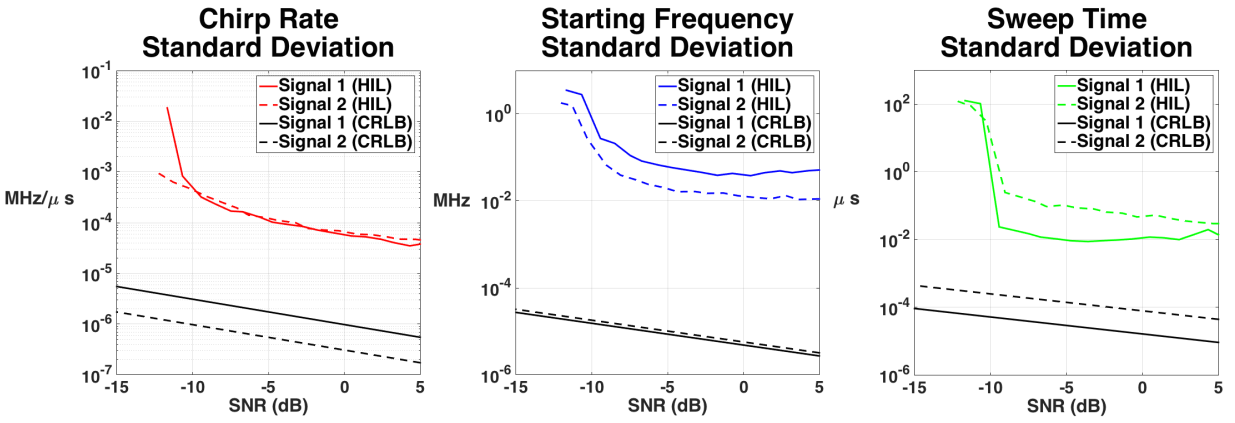

Figure 22: HIL Results

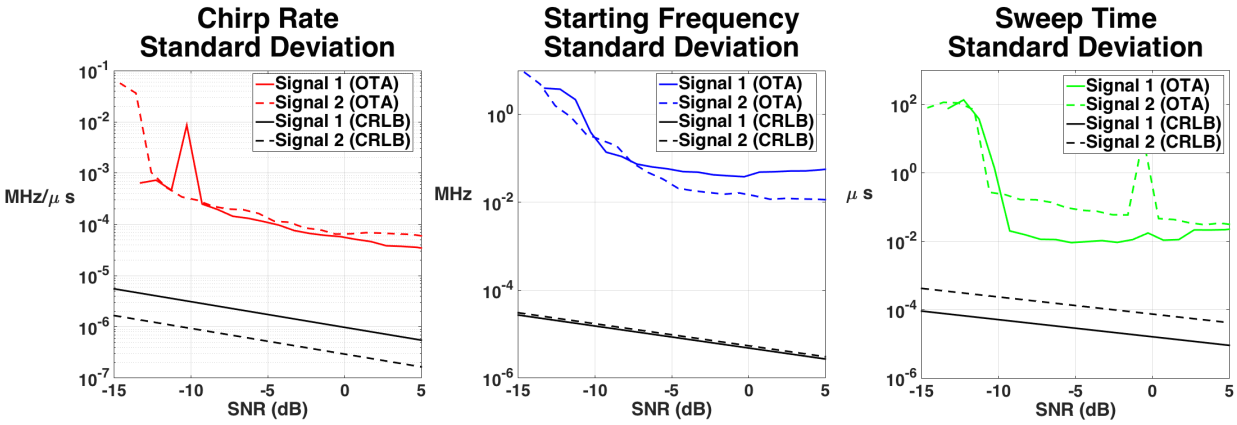

Figure 23: OTA Results

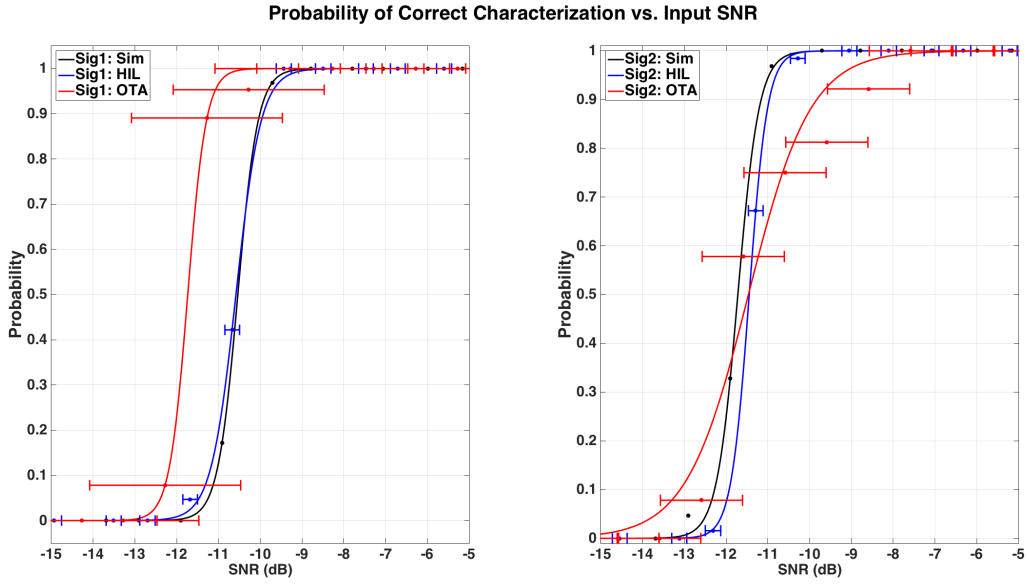

Figure 24: Comparison of Probability of Correct Characterization Between Three Experiments

at which $P_C \geq 90\%$. MDS is commonly used in link budget analyses that determine, under assumptions on the transmitter, receiver, and link characteristics, the range between the transmitter and receiver where the associated SNR would be observed. Table 11 summarizes the MDS associated with each signal for each of the three experiments, along with the 95% confidence interval ranges for the HIL and OTA experiments. The relatively large error bars in the OTA experiment are likely due to fluctuations in the received signal power or the presence of powerful, transient interfering sources.

Table 11: MDS for each Experiment

	Simulation	HIL	OTA	units
Signal No. 1	-10.0	-10.0 $\pm$ 0.2	-11.2 $\pm$ 1.8	dB
Signal No. 2	-11.1	-10.6 $\pm$ 0.2	-9.78 $\pm$ 1.0	dB

3.5 CONCLUSIONS

In this chapter performance of the method was evaluated by comparing the results from three experiments. We found that the simulation, hardware-in-the-loop, and over-the-air experiments each indicate, to within their associated confidence intervals, that the parameters of two simultaneous signals can be estimated to within 10% of their true values $\geq 90\%$ of the time as long as SNR≥ -10 dB. This is a powerful result given that the OTA experiments were performed outdoors in the Industrial, Scientific, and Medical (ISM) band [54], which is a band allocated by the FCC for unlicensed low-power use, with low-cost Software Defined Radios [50] acting as the transmitters and receivers. Examples of some devices that operate in this band are Bluetooth, cordless phones, wireless networks, short-range push-to-talk communication devices, and other experimental equipment. Hence, the band is likely to contain frequent transient interference sources whose noise characteristics are unknown *a-priori*.

3.6 FUTURE DIRECTIONS

Future work related to testing the DCL method originally developed in [19] should focus on evaluating its ability to operate in more challenging environments. In particular, environments where multi-path is clearly present and interference is even stronger and more frequent should be considered.

4.0 SUMMARY OF CURRENT WORK

In this work we developed an approach to Detecting, Characterizing, and Localizing (DCL) multiple Low Probability of Intercept Linear Frequency Modulated Continuous Wave sources and demonstrated features of the approach that do not exist in any one approach in the literature. We demonstrated the performance using a combination of simulation, Hardware-In-the-Loop, and Over-the-Air testing, which we compared to the Cramer-Rao Lower Bound. In contrast to classical approaches to signal detection, such as in [22, 25], our approach is in large part *ad-hoc*, and implemented by breaking the DCL problem up into parts and solving them sequentially. The upside to such an approach is that the virtuous aspects of methods from a wide variety of disciplines, such as signal processing [21,22], statistics [40,57], optimization [31], and optimal state estimation/geo-location [15], can be brought to bear on a challenging problem.

APPENDIX A

THE COMPLEX VALUED SIGNAL MODEL

In this appendix we describe some aspects related to the complex representation of real valued signals. In most physical situations where signals are produced, such as through vibrations, electromagnetic radiation, or neural transmissions, the signals that results are most naturally treated as a real-valued quantities. It is only to this point that most signal analysts would likely remain in agreement. Beyond this point the motives of the analyst dictate how they go about relating a real-valued quantity to a complex valued one.

The motives of an analyst interested in uncovering theoretical properties of the signal tend to avoid making assumptions, since doing so increases the chances of clouding the emergence of new, fundamental signal properties. This group tends to appeal to the so-called *Analytic Signal* [4, Ch. 2] [32, Ch. 3]. In contrast, the analyst whose interest is in solving an engineering problem is at ease in making assumptions and applying them to enable simplifications. This group, particularly those working in the areas of radar and communications, tend to appeal to the so-called *Complex Envelope* [21, pgs. 493-496], [32, Ch. 3] representation. Since this work is developed in context of a radar application, our development of the complex signal will focus on the complex envelope representation. However, the appendix concludes by showing its relationship to the analytic signal.

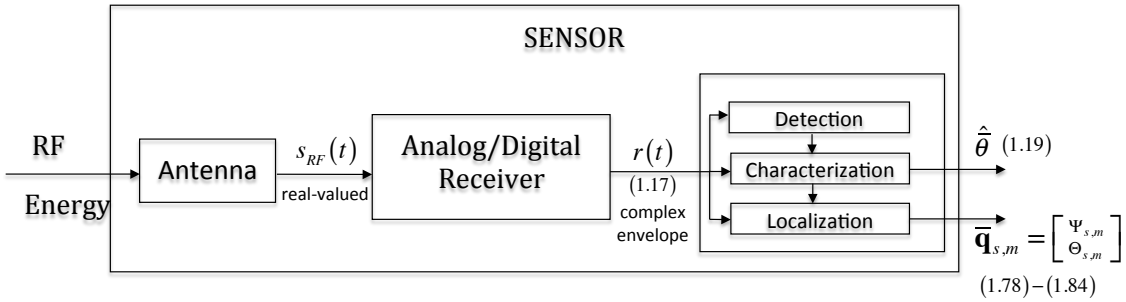

Figure 25: Sensor Block Diagram

A.1 THE COMPLEX ENVELOPE SIGNAL MODEL

With reference to Figure 25, this appendix develops the block labeled 'Analog/Digital Receiver' (ADR), whose task is to translate the real signal $s_{RF}(t)$ into its complex envelop representation $r(t)$. The figure includes references to equations in order to put the block diagram into the context of Chapter 1.

The output of the antenna $s_{RF}(t)$ is the real valued signal we translate, via the ADR, into the complex envelope $r(t)$ (1.17) which, under $\mathcal{H}_1$ is processed by the methods developed in Chapter 1 to produce estimates of $\bar{\theta}$ (1.19) and $\mathbf{q}_{s,m}$ (1.78)-(1.84). Our motives are to construct a complex representation of $s_{RF}(t)$ that can be related to the LFMCW signals model and facilitates efficient receiver designs.

The development of the complex envelope begins by assuming that is meaningful to express $s_{RF}(t)$ as

$$s_{RF}(t) = A(t)\cos(\omega_C t + \phi(t)) \tag{A.1}$$

where ω_C is the carrier frequency, which we assume is known, and $A(t)$ and $\phi(t)$ are unknown functions, commonly referred to as the amplitude and phase modulations, respectively. We can see, by inspection, that solving (A.1) for the unknowns $A(t)$ and $\phi(t)$ based on observations of $s_{RF}(t)$ is an under determined problem since, for each t, we have two unknowns, namely, $A(t)$ and $\phi(t)$. So, for say N_T observations of $s_{RF}(t)$ we have $2N_T$ unknowns. Hence, we need to apply constraints to (A.1) in order to make the problem well-defined.

The development of the complex envelope model hinges two assumptions, which may also be interpreted as constraints. In particular, we assume

1. $s_{RF}(t)$ is narrowband, with known bandwidth $2\pi B$
2. ω_C is known.

Assuming $s_{RF}(t)$ is narrowband amounts to assuming that the lowest frequency in $s_{RF}(t)$ is much greater than zero. Mathematically, this is equivalent to $\omega_C - \frac{2\pi B}{2} \gg 0$, where bandwidth $2\pi B$ is defined to be the range of frequencies about ω_C where the majority [1] of the signal energy is confined. In radar and communications applications the narrowband assumption is standard since both types of transmissions are limited to certain regions of the frequency spectrum either by law, such as those imposed by the Federal Communications Commission (FCC), or by limitations on the transmitter. Assuming knowledge of $2\pi B$ and ω_C are somewhat less certain in passive applications, such as those treated in this dissertation. Some typical methods for determining them though are through *a-priori* knowledge or through wide-band search techniques that queue specialized processing methods (*i.e.*, Chapter 1) to focus on a particular center frequency and bandwidth. As we will see, these two assumptions give rise to the complex envelope model which takes the familiar form

$$s(t) = A(t)\, e^{j\phi(t)} \tag{A.2}$$

which is no longer under determined with respect to $A(t)$ and $\phi(t)$. This is the case since, now for each t we have two equations and two unknowns. Specifically,

$$\mathrm{Re}\{s(t)\} = s_I(t) = A(t)\cos(\phi(t)) \tag{A.3}$$

$$\mathrm{Im}\{s(t)\} = s_Q(t) = A(t)\sin(\phi(t)) \tag{A.4}$$

where $A(t)$ and $\phi(t)$ can be solved for by

$$A(t) = \sqrt{s_I(t)^2 + s_Q(t)^2} \tag{A.5}$$

$$\phi(t) = \arctan\left(\frac{s_Q(t)}{s_I(t)}\right) \tag{A.6}$$

[1] A precise definition of 'majority' is typically application dependent.

This solution to (A.2) implicitly constraints $A(t)$ to be non-negative. In radar and communications applications additional constraints on $s(t)$ arise naturally from knowledge of signal used by the transmitter. The main benefits of expressing (A.1) as (A.2) is that it reduces the rate at which the signal needs to be sampled, since the carrier ω_C is removed, in order to retain the information conveyed in $A(t)$ and $\phi(t)$. This is beneficial since the higher the sampling rate of the Analog to Digital Converter (ADC) the greater the cost. Furthermore, slower sampling rates mean fewer samples need to be stored and therefore storage requirements are minimized.

Developing the procedure that turns (A.1) into (A.2) begins by expanding (A.1) as

$$s_{RF}(t) = A(t)\cos(\phi(t))\cos(\omega_C t) - A(t)\sin(\phi(t))\sin(\omega_C t) \tag{A.7}$$

using the trigonometric identity

$$\cos(\alpha+\beta) = \cos(\alpha)\cos(\beta) - \sin(\alpha)\sin(\beta) \tag{A.8}$$

From (A.7) we see that the components we seek, namely (A.3) and (A.4), are embedded in the right hand side. The remainder of this derivation is devoted to extracting them. Since the steps in the derivation mirror the receiver architecture used to implement the conversion of (A.1) to (A.2) an example block diagram of the receiver, which is commonly referred to a quadrature demodulator, is given in Figure 26. This figure is adapted from [21, Figure 15.2].

First, we manipulate (A.7) to extract (A.3). Multiplying s_{RF} by $\cos(\omega_C t)$ yields

$$\begin{aligned}
a(t) &= s_{RF}(t)\cdot 2\cos(\omega_C t) && \text{(A.9)}\\
&= 2A(t)\cos(\phi(t))\cos(\omega_C t)\cos(\omega_C t) - 2A(t)\sin(\phi(t))\sin(\omega_C t)\cos(\omega_C t) && \text{(A.10)}\\
&= 2A(t)\cos(\phi(t))\frac{1}{2}(1+\cos(2\omega_C t)) - 2A(t)\sin(\phi(t))\frac{1}{2}(0+\sin(2\omega_C t)) && \text{(A.11)}\\
&= A(t)\cos(\phi(t)) + A(t)\cos(\phi(t))\cos(2\omega_C t) - A(t)\sin(\phi(t))\sin(2\omega_C t) && \text{(A.12)}\\
&= A(t)\cos(\phi(t)) + \underbrace{A(t)\cos(\phi(t)+2\omega_C t)}_{\text{narrowband signal centered at } 2\omega_C} && \text{(A.13)}
\end{aligned}$$

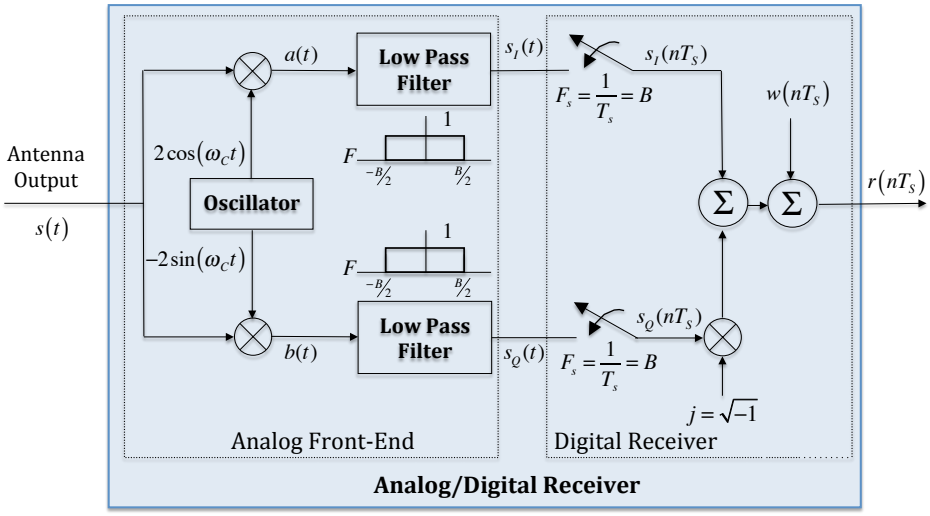

Figure 26: Analog/Digital Receiver Block Diagram [21, pg. 495]

where we go from (A.10) to (A.11) by applying the identities

$$\cos(\alpha)\cos(\beta) = \frac{1}{2}(\cos(\alpha-\beta)+\cos(\alpha+\beta)) \tag{A.14}$$

$$\sin(\alpha)\cos(\beta) = \frac{1}{2}(\sin(\alpha-\beta)+\sin(\alpha+\beta)) \tag{A.15}$$

As indicated, the second term in the right hand side of (A.13) has a spectrum whose lowest frequency component, which is at $2\omega_C - \frac{2\pi B}{2}$, is much greater than zero. Hence, filtering $a(t)$ with a low pass filter with cutoff frequency $2\pi B/2$ yields

$$[a(t)]_{LPF} = A(t)\cos(\phi(t)) \tag{A.16}$$

$$= s_I(t) \tag{A.17}$$

$$= \mathrm{Re}\{s(t)\} \tag{A.18}$$

where $[\cdot]_{LPF}$ denotes the low pass filtering operation. In a similar fashion we get can manipulate (A.7) to get (A.4)

$$[b(t)]_{LPF} = [s_{RF}(t)\, 2\sin(\omega_C t)]_{LPF} \tag{A.19}$$

$$= A(t)\sin(\phi(t)) \tag{A.20}$$

$$= s_Q(t) \tag{A.21}$$

$$= \mathrm{Im}\{s(t)\} \tag{A.22}$$

Obtaining the complex envelope in this way allows $s_I(t)$ and $s_Q(t)$ to be sampled at B Hz rather than requiring that $s_{RF}(t)$ be sampled at $2(\omega_C + 2\pi B)/2\pi$. Additionally, this receiver architectures gives direct access to the amplitude and phase modulation functions that contain the information imparted on the carrier by the transmitter. In our application we extract that information by relating (A.2) to (1.18) and processing the noisy observations $r(t)$ using the methods described in Chapter 1. Applying our model from (1.18) to (A.2) is equivalent incorporating another constraint.

A.2 RELATIONSHIP TO THE ANALYTIC SIGNAL

To those used to dealing with the analytic signal, the form of the complex envelope given by (A.2) is familiar. The key difference between the analytic signal and the complex envelope is that the analytic signal is derived without assumptions on bandwidth or center frequency. This is the case because analysis that is based on the analytic signal, such as that found throughout [4], seeks to keep such aspects of the signal unknown so that they may be uncovered. To define the analytic signal [4, pg. 37] we begin with only the real-valued time series $s_{RF}(t)$), free of assumptions, and let

$$s_a(t) = \frac{2}{\sqrt{2\pi}} \int_0^\infty S_{RF}(\omega)\, e^{j\omega t} d\omega \tag{A.23}$$

where $S_{RF}(\omega)$ is the spectrum of the real signal

$$S_{RF}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} s_{RF}(t)\, e^{-j\omega t} dt \tag{A.24}$$

From [4, Eq. 2.67] we have that

$$|s_a(t) - s(t)| \leq \frac{2}{\sqrt{2\pi}} \int_{-\infty}^{0} |S(\omega)|\, d\omega \tag{A.25}$$

where $S(\omega)$ is the spectrum of the complex envelope

$$S(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} s(t)\, e^{-j\omega t} dt \tag{A.26}$$

So, if we replace $s(t)$ by $s(t)\, e^{j\omega_C t}$ then $S(\omega)$ is replaced by $S(\omega - \omega_C)$. Then (A.25) becomes

$$\left|s_a(t) - s(t)\, e^{j\omega_C t}\right| \leq \frac{2}{\sqrt{2\pi}} \int_{-\infty}^{0} |S(\omega - \omega_C)|\, d\omega \tag{A.27}$$

By the narrowband assumption made on $s_{RF}(t)$ we have that the lowest frequency component in $S(\omega - \omega_C)$ is much greater than zero, which drives the right hand side of (A.27) to zero. So, we can conclude that

$$s_a(t) = s(t)\, e^{j\omega_C t} \tag{A.28}$$

APPENDIX B

THOMSON'S MULTITAPER METHOD FOR HARMONIC ANALYSIS

Recall from Section 1.2, the hypothesis test that Thomson's method for harmonic analysis is designed to decide on is

$$\mathcal{H}_0 : \; r_{n,k} = w_{n,k} \tag{B.1}$$

$$\mathcal{H}_1 : \; r_{n,k} = s_{n,k} + w_{n,k} \tag{B.2}$$

where

$$s_{n,k} \approx s_{n,k}^{\text{STHM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi f_{c,k} n T_s} \tag{B.3}$$

and $L_k \leq M$ is the number of constant frequencies $f_{c,k}$ present in each short time interval δ_k and M is the number of LFMCW signals present. The complex amplitude is given by $C_{c,k} = b_{c,k} e^{j\phi_{c,k}}$ where $b_{c,k}$ is as defined in (1.20) and $\phi_{c,k}$ is the phase of the c^{th} harmonic component in the k^{th} short-time interval. Under $\mathcal{H}_1$, $\{r_{n,k}\}$ is a set of samples with a time-varying mean. Thus, the decision procedure is designed to decide between (B.1) and (B.2) by determining whether the mean of $\{r_{n,k}\}$ is statistically different from zero and consistent with the harmonic model of the mean given in (B.3). As we will see, the key idea behind making this determination is to express (B.2) in the frequency domain and, upon application of Thomson's multi-taper method, recognize the regression problem that results. The decision procedure then amounts to determining the statistical significance of the computed regression coefficients.

To simply the derivations that follow, we develop the decision procedure for the case when the goal is to determine whether one harmonic component is present (*i.e.*, $L_k = 1$) in $\{r_{n,k}\}$ or whether it is noise only. Subsequently we show how the derivation generalizes for $L_k > 1$.

The expected value of $r_{n,k}$ is

$$E[r_{n,k}] = C_{1,k} e^{j2\pi f_{1,k} n T_s} \tag{B.4}$$

where $E[\cdot]$ is the statistical expectation operator. So,

$$r_{n,k} = E[r_{n,k}] + w_{n,k} \tag{B.5}$$

Next, we begin the development of the frequency domain representation of (B.2) by considering the Discrete Time Fourier Transform (DTFT) of $\left\{h_{q,n} r_{n,k}/\sqrt{T_s}\right\}$

$$J_{q,k}(f) = T_s \sum_{n=1}^{N} \left(h_{q,n} r_{n,k}/\sqrt{T_s}\right) e^{-j2\pi f n T_s} \tag{B.6}$$

where $\{h_{q,n}\}$ is the order q Discrete Prolate Spheroidal Sequence (DPSS) [40, 378-390]. These sequences have the property that for a given number, N, of time-domain samples, sampled at a rate F_s, their DTFT has its energy maximally concentrated in the normalized frequency range $[-W, W]$ (*i.e.*, $0 \leq W \leq 0.5$), where $2W$ is defined as the resolution bandwidth of the spectral estimator. This property allows one to select the resolution bandwidth in terms of N, or vice-versa. For our problem, having a resolution bandwidth of $2W$ means that we cannot resolve the instantaneous frequencies of multiple LFMCW signals if they are closer to one another than $2WF_s$ in a given δ_k. Hence, $J_{q,k}(f)$ is the DTFT of $\{r_{n,k}\}$ tapered (a/k/a windowed) by $\{h_{q,n}/\sqrt{T_s}\}$. Taking the expected value of $J_{q,k}(f)$ and using (B.4) yields

$$\begin{aligned} E[J_{q,k}(f)] &= \sqrt{T_s} \sum_{n=1}^{N} h_{q,n} C_{1,k} e^{j2\pi f_{1,k} n T_s} e^{-j2\pi f n T_s} \\ &= \frac{C_{1,k}}{\sqrt{T_s}} H_q(f - f_{1,k}) \end{aligned} \tag{B.7}$$

where

$$H_q(f) \equiv T_s \sum_{n=1}^{N} h_{q,n} e^{-j2\pi f n T_s} \tag{B.8}$$

is the DTFT of $\{h_{q,n}\}$. Similar to (B.5) we express $J_{q,k}(f)$ in terms of its mean as

$$J_{q,k}(f) = T_s \sum_{n=1}^{N} \left(h_{q,n} r_{n,k} / \sqrt{T_s} \right) e^{-j2\pi f n T_s} \tag{B.9}$$

$$= T_s \sum_{n=1}^{N} \left(\left(h_{q,n} s_{n,k}^{\text{STHM}} + h_{q,n} w_{n,k} \right) / \sqrt{T_s} \right) e^{-j2\pi f n T_s} \tag{B.10}$$

$$= T_s \sum_{n=1}^{N} \left(h_{q,n} s_{n,k}^{\text{STHM}} / \sqrt{T_s} \right) e^{-j2\pi f n T_s} + \tag{B.11}$$

$$T_s \sum_{n=1}^{N} \left(h_{q,n} w_{n,k} / \sqrt{T_s} \right) e^{-j2\pi f n T_s} \tag{B.12}$$

$$= E[J_{q,k}(f)] + T_s \sum_{n=1}^{N} \left(h_{q,n} w_{n,k} / \sqrt{T_s} \right) e^{-j2\pi f n T_s} \tag{B.13}$$

where we note that the magnitude-squared of the term containing $w_{n,k}$ is an estimator of $P_w(f)$ in δ_k using $\{h_{q,n}\}$ as the data taper. This is the frequency domain representation of (B.2) we described as a key element to the method.

We now begin to develop a test statistic for deciding between the hypotheses given in (B.1) and (B.2). The first step is to evaluate (B.6) at $f = f_{1,k}$

$$J_{q,k}(f_{1,k}) = E[J_{q,k}(f_{1,k})] + \sqrt{T_s} \sum_{n=1}^{N} h_{q,n} w_{n,k} e^{-j2\pi f_{1,k} n T_s} \tag{B.14}$$

where from (B.7)

$$E[J_{q,k}(f_{1,k})] = \frac{C_{1,k}}{\sqrt{T_s}} H_q(0) \tag{B.15}$$

By substituting (B.15) into (B.14) we arrive at the following complex-valued, first order regression model for $J_{q,k}(f_{1,k})$

$$J_{q,k}(f_{1,k}) = C_{1,k} \frac{H_q(0)}{\sqrt{T_s}} + \tilde{\epsilon}_{q,k}, \quad q = 0, \ldots, Q-1, \tag{B.16}$$

where

$$\tilde{\epsilon}_{q,k} \equiv \sqrt{T_s} \sum_{n=1}^{N} h_{q,n} w_{n,k} e^{-j2\pi f_{1,k} n T_s} \tag{B.17}$$

and Q is the number of DPSS data-tapers used (typically $Q < 2NW$), $C_{1,k}$ is the complex-valued unknown to be estimated, $H_q(0)/\sqrt{T_s}$ is the q^{th} independent-variable, and $\tilde{\epsilon}_{q,k}$ is the error term. It is interesting to note that the magnitude-squared of the error term in (B.17) is

equivalent to the estimated PSD of $P_{w_{n,k}}(f)$ using the q^{th} data taper and evaluated at $f_{1,k}$. Intuitively, the requirement that $P_{w_{n,k}}(f)$ is slowly varying in $[f_{1,k}-W, f_{1,k}+W]$ is necessary to ensure that $P_{w_{n,k}}(f)$ has structure different from a harmonic, which ideally appears as a spike, at $f_{1,k}$.

Before invoking the equation that solves for the least-squares estimate of $C_{1,k}$ [33, Eq. (3.12)], we must establish two conditions on $\tilde{\epsilon}_q$:

Condition 1: $\tilde{\epsilon}_{q,k}$ is zero-mean, complex-valued, and Gaussian distributed with uncorrelated real and imaginary components, with each component having equal semi-variances $\sigma^2_{\tilde{\epsilon}_{q,k}}/2$.

Condition 2: For each q, $\tilde{\epsilon}_{q,k}$ must be mutually uncorrelated and have equal variance $\sigma^2_{\tilde{\epsilon}_k}$.

Condtions 1 and 2 are established in [40, pg. 498]. Therefore, the least-squares estimator of $C_{1,k}$ is given by [33, Eq. (3.12)]

$$\hat{C}_{1,k} = \sqrt{T_s} \cdot \frac{\sum_{q=0}^{Q-1} J_{q,k}(f_{1,k}) H_q(0)}{\sum_{q=0}^{Q-1} H_q^2(0)} \tag{B.18}$$

In addition to the result given in (B.18), Theorem 8.1 of [33] establishes four useful properties related to Equation (B.16), which will be necessary later when we eliminate the dependence of the noise power from the calculation of the detection threshold.

Property 1: $\hat{C}_{1,k}$ is a complex Gaussian random variable with mean $C_{1,k}$ (*i.e.*, unbiased) and variance $\sigma^2_{\hat{C}_{1,k}} = \sigma^2_{\tilde{\epsilon}_k} T_s / \sum_{q=0,2,\ldots}^{Q-1} H_q^2(0)$.

Property 2: An estimator of $\sigma^2_{\tilde{\epsilon}_k}$ is given by

$$\hat{\sigma}^2_{\tilde{\epsilon}_k} = \frac{1}{Q} \sum_{q=0}^{Q-1} \left| J_{q,k}(f_{1,k}) - \hat{J}_{q,k}(f_{1,k}) \right|^2 \tag{B.19}$$

where $\hat{J}_{q,k}(f_{1,k})$ is the value of $J_{q,k}(f_{1,k})$ predicted by the estimate $\hat{C}_{1,k}$

$$\hat{J}_{q,k}(f_{1,k}) \equiv \hat{C}_{1,k} \frac{H_q(0)}{\sqrt{T_s}} \tag{B.20}$$

Property 3: The random variable $2Q\hat{\sigma}^2_{\tilde{\epsilon}_k}/\sigma^2_{\tilde{\epsilon}_k} \sim \chi^2_{2Q-2}$, where $\sim$ is read as 'is distributed according to' and χ^2_{2Q-2} denotes a chi-squared distribution with $2Q-2$ degrees of freedom.

Property 4: $\hat{C}_{1,k}$ and $2Q\hat{\sigma}^2_{\epsilon_k}/\sigma^2_{\tilde{\epsilon}_k}$ are independent of each other.

The development between (B.4) and (B.20) forms the basis for the test statistic that determines whether periodicity exists in $\{r_{n,k}\}$ that adheres to (B.2) (*i.e.*, $\mathcal{H}_1$) or whether *only* the background random process $\{w_{n,k}\}$ exists in δ_k, in which case $\{r_{n,k}\}$ adheres to (B.1) (*i.e.*, $\mathcal{H}_0$). To decide between these two hypotheses we recognize that, for $L_k = 1$, (B.1) is equivalent to the case where $J_{q,k}(f_{1,k}) = \tilde{\epsilon}_{q,k}$. Likewise, (B.2) is equivalent to the case where $J_{q,k}(f_{1,k}) = C_{1,k}\frac{H_q(0)}{\sqrt{T_s}} + \tilde{\epsilon}_{q,k}$. *Thus our approach will be to formulate a test that rejects the null hypothesis if $\hat{C}_{1,k}$ is statistically different from zero.*

To do so first recall from **Property 1** that under $\mathcal{H}_0$ the random variable $\hat{C}_{1,k}$ has a complex Gaussian distribution with zero mean and variance $\sigma^2_{\hat{C}_{1,k}} = \sigma^2_{\tilde{\epsilon}_k} T_s / \sum_{q=0}^{Q-1} H_q^2(0)$. Next, recall that under the null hypothesis $J_{q,k}(f_{1,k}) = \tilde{\epsilon}_{q,k}$ and, by **Conditions 1** and **2**, we have that for each q, $\tilde{\epsilon}_{q,k}$ are mutually uncorrelated with independent real and imaginary components. Therefore, since $\hat{C}_{1,k}$ is the weighted sum of $\tilde{\epsilon}_{q,k}$ for $q = 0, \ldots, Q-1$ implies that $\hat{C}_{1,k}$ has independent real and imaginary components as well. Hence, under the null hypothesis $\hat{C}_{1,k}/(\sigma_{\hat{C}_{1,k}}/\sqrt{2})$ has uncorrelated real and imaginary components that are zero mean and unit variance (*i.e.*, independent standard normal random variables). Since the sum of the squares of two standard normal random variables is chi-squared distributed with two degrees-of-freedom we have that

$$\left|\hat{C}_{1,k}/(\sigma_{\hat{C}_{1,k}}/\sqrt{2})\right|^2 = \frac{2\left|\hat{C}_{1,k}\right|^2 \sum_{q=0}^{Q-1} H_q^2(0)}{\sigma^2_{\tilde{\epsilon}_k} T_s} \tag{B.21}$$

$$\sim \chi^2_2 \tag{B.22}$$

If one knew the value of $\sigma^2_{\tilde{\epsilon}_k}$, which would in turn require knowledge of the variance of $w_{n,k}$, then Equations (B.21) and (B.22) could be used as the basis for a statistical test that determines whether $\hat{C}_{1,k} = 0$. However, this would contradict the claims made in Items 1

and 2 of Subsection 1.1.3. To hold to these claims we seek a test statistic whose distribution is not only obtainable under $\mathcal{H}_0$ but also free of $\sigma^2_{\tilde{\epsilon}_k}$. To obtain such a test statistic recall from **Properties 3** and **4** that the random variables $\hat{C}_{1,k}$ and $2Q\hat{\sigma}^2_{\tilde{\epsilon}_k}/\sigma^2_{\tilde{\epsilon}_k}$ are independent. This implies that the right-hand side of Equation (B.21) is independent of $2Q\hat{\sigma}^2_{\tilde{\epsilon}_k}/\sigma^2_{\tilde{\epsilon}_k}$ as well. Next recall that if $a \sim \chi^2_\nu$ and $b \sim \chi^2_\kappa$, then

$$\frac{a/\nu}{b/\kappa} \sim F_{\nu,\kappa} \tag{B.23}$$

where $F_{\nu,\kappa}$ is an F-Distribution with ν and κ degrees-of-freedom [7]. Thus, if we let

$$D_k\left(f_{1,k}\right) \equiv \frac{\left(2\left|\hat{C}_{1,k}\right|^2 \sum\limits_{q=0,2,\ldots}^{Q-1} H_q^2\left(0\right) \Big/ \sigma^2_{\tilde{\varepsilon}_k} T_s\right) \Big/ 2}{\left(2Q\hat{\sigma}^2_{\tilde{\varepsilon}_k}/\sigma^2_{\tilde{\varepsilon}_k}\right) \Big/ (2Q-2)} \tag{B.24}$$

then substitute Equation (B.19) from **Property 2** into Equation (B.24) and simplify, we get

$$D_k\left(f_{1,k}\right) = \frac{\left(Q-1\right)\left|\hat{C}_{1,k}\right|^2 \sum\limits_{q=0,2,\ldots}^{Q-1} H_q^2\left(0\right)}{T_s \sum\limits_{q=0}^{Q-1} \left|J_q\left(f_{1,k}\right) - \hat{J}_q\left(f_{1,k}\right)\right|^2} \tag{B.25}$$

$$\sim F_{2,2Q-2} \tag{B.26}$$

At last, we have a test statistic that, under the null-hypothesis, has a known distribution and is a function of all of the relevant known quantities.

To get the detection threshold as a function of P_{FA}, let

$$P_{FA} \equiv \Pr\left(D_k\left(f_{1,k}\right) > \eta | \mathcal{H}_0\right) \tag{B.27}$$

where $\Pr\left(A|B\right)$ is the conditional probability that event A occurs given that the conditions defined by B are true. Then

$$P_{FA} = 1 - \int_0^\eta F_{2,2Q-2}\left(\xi\right) d\xi \tag{B.28}$$

It can be shown [40, pg. 501] that the solution to this equation for η is

$$\eta = \frac{(Q-1)\left(1 - {P_{FA}}^{1/(Q-1)}\right)}{P_{FA}^{1/(Q-1)}} \tag{B.29}$$

which gives a detection threshold that is a function of P_{FA}, as desired. Thus, if $D_k(f) > \eta$ for $f = f_{1,k}$ then $f_{1,k}$ is significant, the null hypothesis is rejected, and the resulting value of $\hat{C}_{1,k}$ can be used to solve for $\hat{b}_{1,k}$ and $\hat{\phi}_{1,k}$ since $C_{1,k} \equiv b_{1,k}e^{\phi_{1,k}}$. Finally, suppose that $L_k > 1$. If this is the case and $f_{i,k} \notin [f_{j,k} - WF_s, f_{j,k} + WF_s]$ for $i \neq j$, then $D_k(f_{i,k})$ is influenced very little by $f_{j,k}$, $D_k(f_{i,k})$ is a $F_{2,2Q-2}$ random variable, and $D_k(f_{i,k})$ exceeds η for the same value of P_{FA}. The influence is small because the leakage properties of the Slepian sequences used are very good [40, pg. 331, 334-335].

At this point we highlight the fact that this approach is equipped to deal with harmonic components in colored noise because in the preceding development we treated $P_{w_{n,k}}(f)$ as an unknown PSD corresponding to a Gaussian random process. Moreover, due to the short-time model we impose, non-stationary random processes are naturally admissible as well as long as N is small enough. That is, $P_{w_k}(f)$ can be different from one time interval to the next, but, within each time-interval δ_k, $\{w_{n,k}\}$ must be wide-sense stationary and have a slowly-varying spectrum [40, pgs. 496-501]. This is a common feature of window-based spectral analysis methods.

We close the appendix with Figure 27, which is intended to walk the reader through the implementation of Thomson's method as it relates to the equations developed in this appendix.

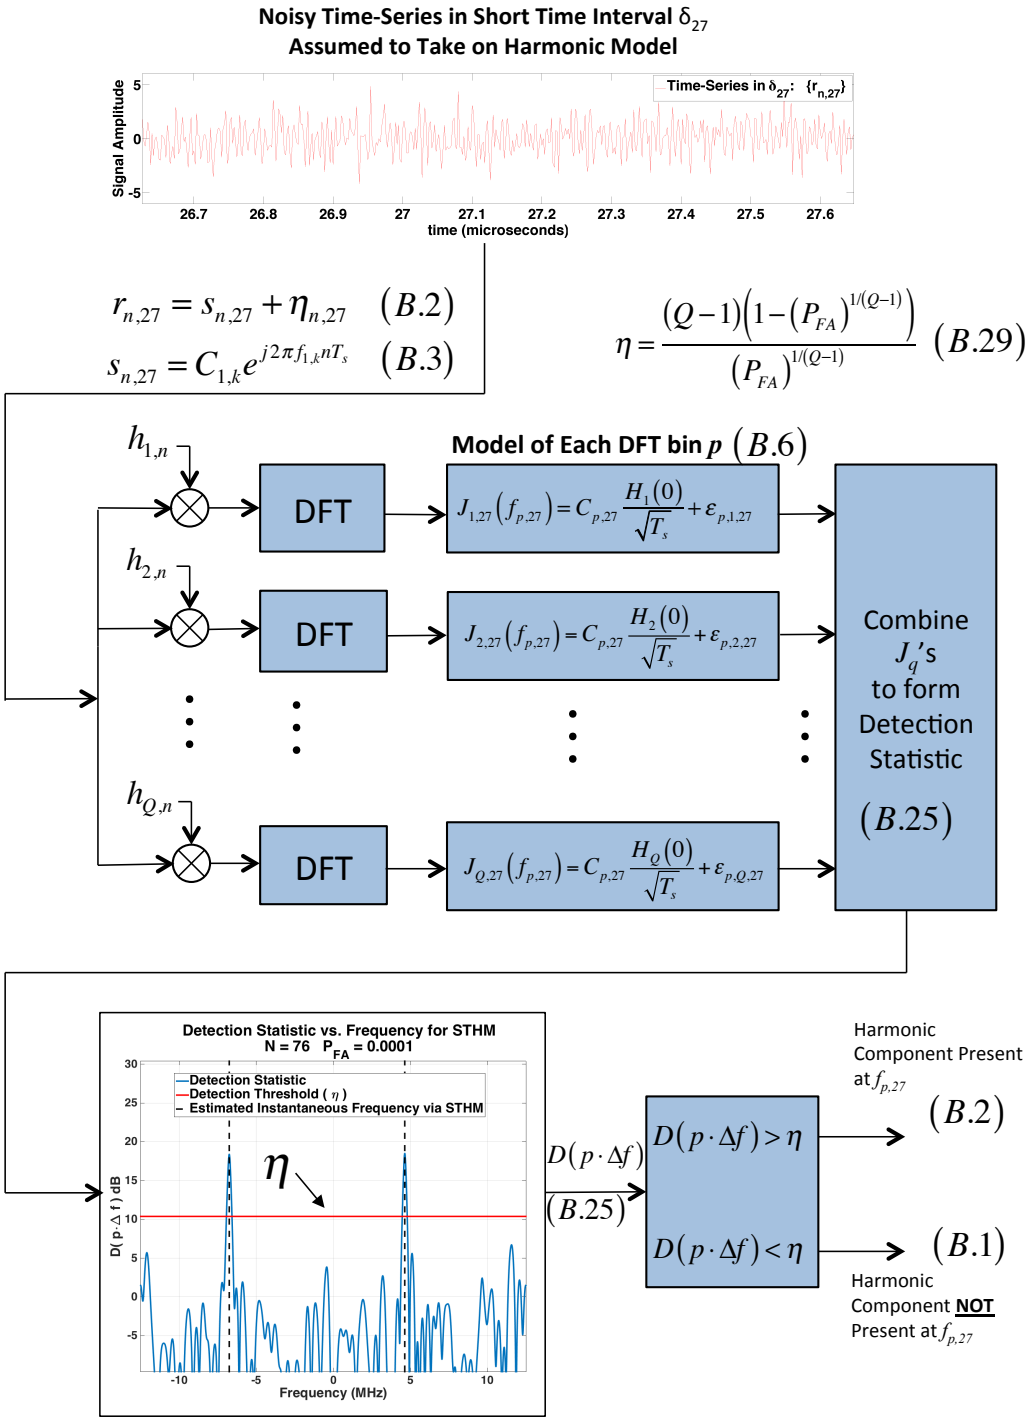

Figure 27: Block Diagram Showing Implementation of Thomson's Method for Harmonic Analysis

APPENDIX C

EXTENDING THE SHORT TIME HARMONIC MODEL

In this Appendix we develop an extension to the STHM developed in Chapter 1. The main idea behind the extension is to increase the duration of the short-time interval (C.3) that defines each STHM (C.5) so that within these intervals the frequency of each signal component has enough time increase or decrease according to the chirp-rate. Hence, increasing the duration of δ_k as such leads allows for the possibility to estimate chirp-rate, which results in an ability to jointly estimate the center frequency and chirp-rate of each signal within each short-time interval. With estimates of center frequency and chirp-rate we speculate that one or more of the following future directions can be pursued:

1. By jointly estimating center frequency and chirp-rate, false alarms that are consistent with a simplified harmonic model of the data could be eliminated on the basis that they have no appreciable linear frequency modulation local to the harmonic energy. In other words, one might expect environmental interference to appear as a sinusoidal component if observed over a short-time interval, but less likely to appear to have linear frequency modulation over longer time intervals. Hence, false alarms due to the appearance of interference components, in addition to false alarms induced by random noise fluctuations, will likely be reduced.

2. By jointly estimating center frequency and chirp-rate in each short-time interval, additional information becomes available to later processing steps; namely, clustering (see Section 1.3.1). It is speculated that by adding estimates of chirp-rate to (1.44)-(1.45)

or (1.52)-(1.53) the current process of clustering frequencies by chirp segment may be accelerated and its performance improved. Accelerating this step would be a major improvement since it is currently the most time consuming of all the algorithm steps.

3. Including an estimate of chirp-rate that results from each short-time interval may also make tracking frequency with an estimator like a Kalman Filter (KF) or Extended Kalman Filter (EKF) viable. These ideas were pursued in [1, pgs. 200-202] and by others, which are reviewed in [26]. However, after a preliminary literature review it appears the ideas were pursued when either samples of the complex time-series or estimates of frequency are treated as the measurement input to the KF or EKF. Formulating a KF or EKF with estimates of frequency and chirp-rate as measurement inputs could make tracking multi-component signals, which would be equivalent to multi-target tracking [2], more feasible.

4. Finally, the most clear benefit of developing a capability to jointly estimate chirp-rate and center frequency is that constraints on making the width in samples N of the short-time interval (see 1.2.2) can be relaxed. This has two main benefits. First, one does not need to rely as heavily on *a-priori* assumptions on the maximum chirp-rate to determine the maximum N that will ensure the signal adheres to the STHM given by (C.5). Second, allowing N to increase will enable more accurate estimates of center frequency since, in general, more samples means increased estimation accuracy of estimators.

In the remainder of this Appendix we develop one approach to estimating chirp-rate in addition to center frequency and test its performance compared to the CRLB [39]. We found that the method did not substantially increase the computational requirements of Thomson's method and that for sufficiently high SNR the CRLB is achieved. Additionally, a fundamental limitation of the methods proposed in this dissertation is clearly brought to light, which joint estimation of center frequency and chirp-rate may possibly address.

C.1 A JOINT FREQUENCY, CHIRP-RATE ESTIMATOR

The approach is an extension of the method developed in Section 1.2 for estimating the frequency of each harmonic component, if present, in a short-time interval.

For convenience we reproduce some of the relevant equations from Section 1.2. Consider the following hypothesis test [1]

$$\mathcal{H}_0 : \; r_{n,k} = w_{n,k} \tag{C.1}$$

$$\mathcal{H}_1 : \; r_{n,k} = s_{n,k} + w_{n,k} \tag{C.2}$$

where $r_{n,k}$ is the n^{th} sample, for $n = 1, \ldots, N$, of $r(t)$ in the k^{th} time interval. The k^{th} time interval is given by

$$\delta_k = \{t : (k-1) \cdot NT_s \leq t \leq k \cdot NT_s)\} \tag{C.3}$$

where $k = 1, \ldots, K = \lfloor \frac{T}{N \cdot T_s} \rfloor$. That is,

$$r_{n,k} = r\left((k-1) \cdot NT_s + nT_s\right) \tag{C.4}$$

Within this short time interval we model the received time-series as a sum of harmonic components, which gives rise to the following Short Time Harmonic Model (STHM) of the data

$$s_{n,k} \approx s_{n,k}^{\text{STHM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi f_{c,k} n T_s} \tag{C.5}$$

where L_k is the number of constant frequencies $f_{c,k}$ present in δ_k. The complex amplitude is given by $C_{c,k} = b_{c,k} e^{j\phi_{c,k}}$ where $b_{c,k}$ is as defined in (1.20) and $\phi_{c,k}$ is the phase of the c^{th} harmonic component in the k^{th} short-time interval. In this model of the data N is selected (see Section 1.2.2) so that

$$N \leq N_{max} = \left\lceil \frac{2F_s}{\sqrt{2\,|\beta_{max}|}} \right\rceil \tag{C.6}$$

which determines each δ_k.

[1] Samples of a continuous time signal $x(t)$ are denoted by $x_n \equiv x(t)\,|_{t=nT_s}$

Suppose that we let $N = p \cdot N_{max}$ for some $p > 1$. Then the more appropriate Short Time **Chirp** Model (STCM) of $s_{n,k}$ in the associated δ_k is

$$s_{n,k} \approx s_{n,k}^{\text{STCM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi\left(\left(f_{c,k}-\beta_{c,k}\cdot t_{0,k}\right)nT_s+\frac{1}{2}\beta_{c,k}(nT_s)^2\right)} \tag{C.7}$$

where L_k is now the number of short time chirp segments in δ_k. The complex amplitude of each chirp is similarly defined as $C_{c,k} = b_{c,k}e^{j\phi_{c,k}}$, except that $\phi_{c,k}$ is now interpreted as the initial phase of the c^{th} short time chirp segment in the k^{th} time-interval. The amplitude, $b_{c,k}$ is still as defined in (1.20). There are two main differences in the STCM compared to the STHM. First, $f_{c,k}$ is now interpreted as the center frequency of the c^{th} short time chirp segment at the center time $t_{0,k}$ in δ_k. That is, $t_{0,k} = (k - 1/2)(N-1)T_s$. Second, the extended duration of δ_k means that chirp-rate is now observable in the data. Hence, $\beta_{c,k}$ is the chirp-rate of the c^{th} short time chirp segment in the k^{th} short time interval.

The joint frequency/chirp-rate estimator is a three step process.

1. Solve the hypothesis testing problem from (C.1)-(C.2) assuming (C.5) holds in $r_{n,k}$ by selecting N such that (C.6) is satisfied. The methods from Section 1.2 apply directly.
2. If $L_k > 0$ then we assume (C.7) holds by letting $N = p \cdot N_{max}$ for $p > 1$. Each $\beta_{c,k}$ is estimated by maximizing the discretized form of (1.3) for $s\left(t;\bar{\theta}=\beta_{c,k}\right)\big|_{t=nT_s} = s_{n,k}$, which becomes

$$\hat{\beta}_{c,k} = \arg\max_{\beta_{c,k}} \left|\mathcal{L}_1\left(\beta_{c,k}\right)\right|^2 \tag{C.8}$$

   where

$$\mathcal{L}_1\left(\beta_{c,k}\right) = \sum_{n=0}^{N-1} r_{n,k}e^{-j2\pi\left(\left(\hat{f}_{c,k}-\beta_{c,k}t_{0,k}\right)nT_s+\frac{1}{2}\beta_{c,k}(nT_s)^2\right)} \tag{C.9}$$

   Note that $C_{c,k}$ is omitted from the objective function because including it would not change the point in the $\beta_{c,k}$ dimension where the argument is maximized.
3. With estimates of center frequency $\hat{f}_{c,k}$ and chirp-rate $\hat{\beta}_{c,k}$ we can refine our estimates of each by jointly estimate these parameters by solving the following optimization problem with $\hat{f}_{c,k}$ and $\hat{\beta}_{c,k}$ as initial conditions

$$\begin{bmatrix} \hat{f}_{c,k} \\ \hat{\beta}_{c,k} \end{bmatrix} = \arg\max_{f_{c,k},\beta_{c,k}} \left|\mathcal{L}_2\left(f_{c,k},\beta_{c,k}\right)\right|^2 \tag{C.10}$$

where

$$\mathcal{L}_2\left(f_{c,k},\beta_{c,k}\right)=\sum_{n=0}^{N-1}r_{n,k}e^{-j2\pi\left(\left(f_{c,k}-\beta_{c,k}t_{0,k}\right)nT_s+\frac{1}{2}\beta_{c,k}(nT_s)^2\right)} \tag{C.11}$$

One could start by jointly estimating $\hat{f}_{c,k}$ and $\hat{\beta}_{c,k}$ in Step 2, but the objective function defined by (C.10) is prone to numerous local extrema. So, reliable initial conditions on each parameter are necessary in order to reliably find the global minimum using a numerical minimization technique such as the Nelder-Mead simplex direct search algorithm [29,31].

C.2 EXAMPLE

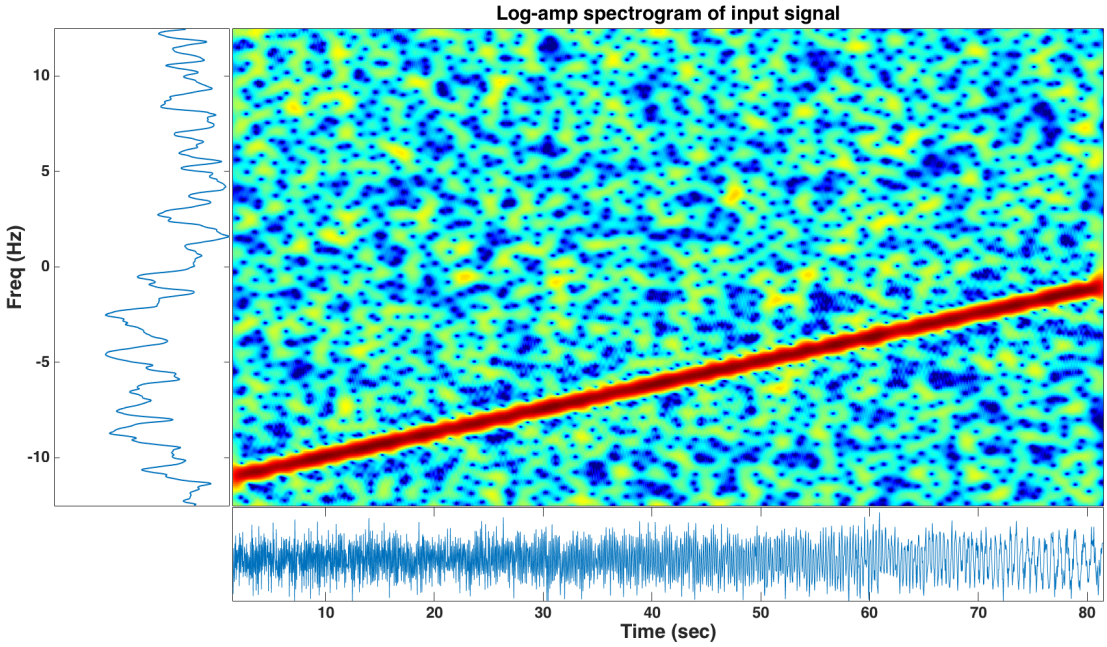

Figure 28: Spectrogram of Single Chirp ($T = 80\mu s$) of $r_{n,k}$ Under $\mathcal{H}_1$ at SNR = 5 dB and 40 dB Dynamic Range

In this section we describe an example where the above three steps are implemented to jointly estimate $f_{c,k}$ and $\beta_{c,k}$. The signal used in the example is given by $m = 1$ in Table 2 for $T = 80\mu s$. A spectrogram of $r_{n,k}$ under $\mathcal{H}_1$ at SNR = 5 dB is given in Figure 28. At this same SNR and example plot of $\mathcal{L}_1(\beta)$ is given in Figure 29, showing the result of a grid based evaluation of (C.9) and the associated maximum value leading to the initial $\widehat{\beta}_{c,k}$. Figure 30 shows the surface corresponding to (C.9) evaluated over $f_{c,k}$ and $\beta_{c,k}$. The main take-away from this figure is that even over a limited range of $f_{c,k}$ and $\beta_{c,k}$ the objective function exhibits a number of relative maxima. Hence, good initial conditions are required in order to make convergence to the global maximum likely when implementing a numerical optimization method.

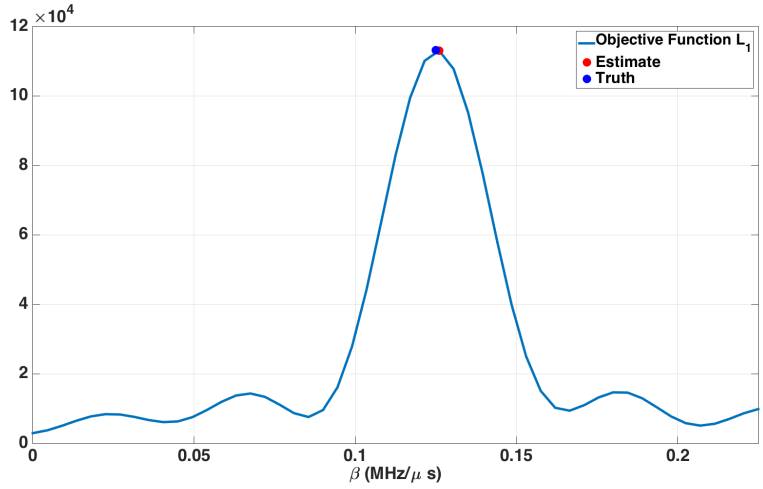

Figure 29: Objective Function $\mathcal{L}_1$ for Estimating $\beta_{c,k}$ Only

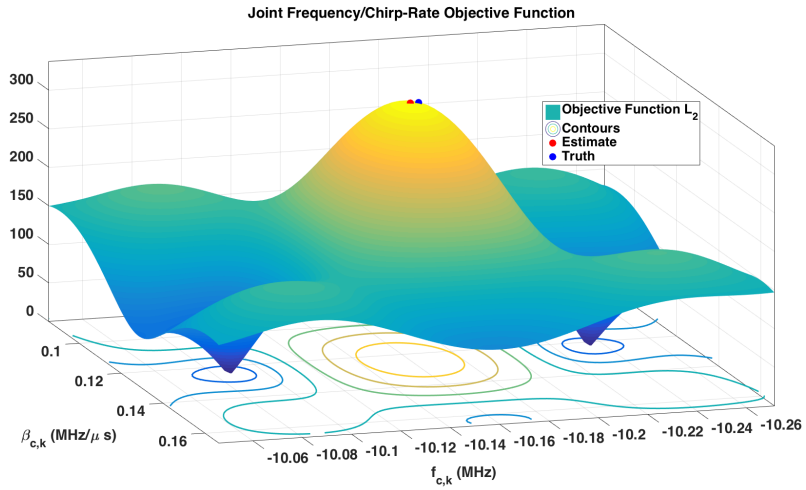

Figure 30: Objective Function $\mathcal{L}_2$ for Jointly Estimating $f_{c,k}$ and $\beta_{c,k}$

C.3 SIMULATION RESULTS

In this section we show the results of a simulation developed to test the method just described for joint estimation. The same test signal described in Section C.2 is used in this example with $N_1 = 113$ for Step 1 and $N_2 = 3 \cdot N_1$ for Step 2. The full $T = 80\mu s$ contained $N_T = 2000$ samples. Figure 31 shows the Monte Carlo simulation results for jointly estimating $f_{c,k}$ and $\beta_{c,k}$ compared to the CRLB associated with a single LFM chirp signal model, such as in our STCM from (C.7), which was obtained from [39]. The settings required to implement Step 1 are the same as those given in Table 3 from Section 1.2.3. We see that estimator's Mean Squared Error (MSE) associated with both center frequency and chirp-rate achieves the CRLB associated with a short-time interval having N_2 samples when input SNR$> -2.5dB$, where input SNR is as defined in (1.15). Since the estimator has only N_2 samples upon which to form each estimate it makes sense that this bound is achieved rather than the bound corresponding to the actual total number of samples available, N_T. The fact that the estimator does not achieve the CRLB associated with N_T is a manifestation of the fundamental trade-off that we accept for the sake of computational tractability. However, by jointly estimating $f_{c,k}$ and $\beta_{c,k}$, which requires $p > 1$, we now have the possibility of achieving lower and lower bounds as computational resources allow.

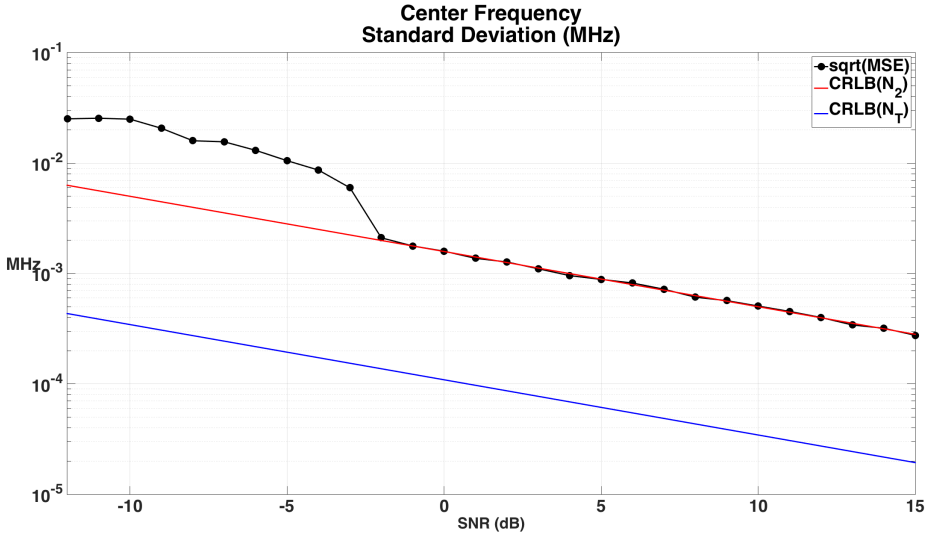

(a) Center Frequency Standard Deviation

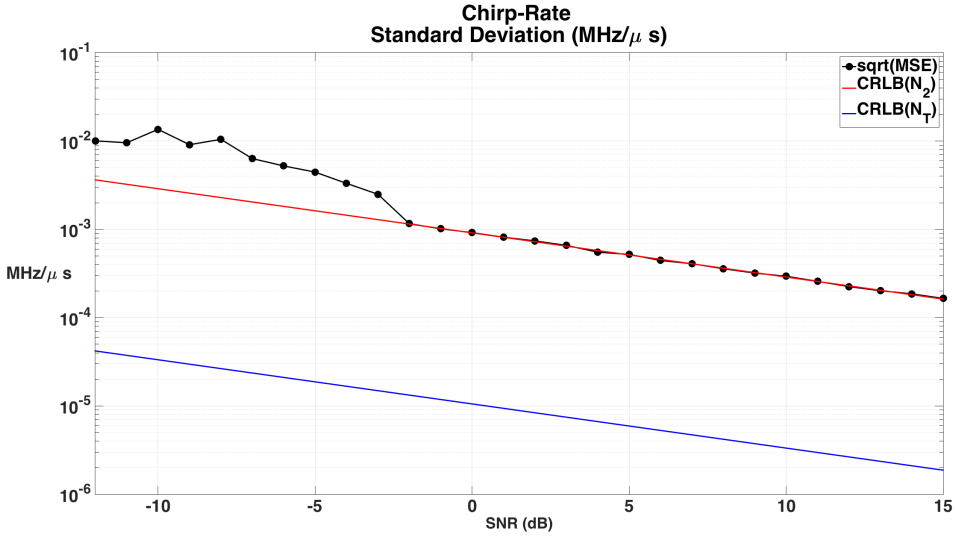

(b) Chirp-Rate Standard Deviation

Figure 31: Joint Frequency/Chirp-Rate Performance Results

APPENDIX D

AN APPROXIMATE ANALYSIS FOR PREDICTING PROBABILITY OF CORRECT CHARACTERIZATION

In this appendix an approximate analysis that predicts the Probability of Correct Characterization (P_C) is provided. Examples of P_C derived from simulation, hardware-in-the-loop, and over-the-air experiments are given in Figures 10 and 24. The main result relates P_C to SNR, Probability of False Alarm (P_{FA}) per short time interval (see Section 1.2 and associated Step 6), number of samples per short time interval N (see (1.27) and Section 1.2.2), and duration of the full signal observation T (see Table 2). In addition to these parameters, which have already been defined, we connect P_C to the percentage of short time intervals wherein a harmonic component is correctly detected. The analysis is approximate for several reasons, which we will highlight during the development, but primarily because the result is independent of the specific signal parameters that define each LFMCW signal, such as chirp-rate, sweep-time, initial frequency, etc. A full analysis would likely depend on the true values of these parameters as well. Nevertheless, the approximate analysis provides a tool for predicting P_C as a function of key system parameters which can aid the designer in performing trade studies or in making parameter selections.

D.1 ANALYSIS

The analysis assumes that there is one signal present so that $M = 1$ and $L_k = 1\ \forall k$. So, (1.29) becomes

$$s_{n,k} \approx s_{n,k}^{\text{STHM}} = \sum_{c=1}^{L_k} C_{c,k} \cdot e^{j2\pi f_{c,k} n T_s} \tag{D.1}$$

$$= C_{1,k} \cdot e^{j2\pi f_{1,k} n T_s} \tag{D.2}$$

for $n = 0, \ldots, N-1$ and $k = 1, \ldots, K = \left\lfloor \frac{T}{N \cdot T_S} \right\rfloor$. Under these assumptions (1.28) becomes

$$r_{n,k} = s_{n,k}^{\text{STHM}} + w_{n,k} \tag{D.3}$$

$$= C_{1,k} \cdot e^{j2\pi f_{1,k} n T_s} + w_{n,k} \tag{D.4}$$

where $C_{1,k}$ is the complex amplitude, which is assumed unknown, of the harmonic component with frequency $f_{1,k}$, which is assumed known. Additionally, $w_{n,k}$ is assumed to be complex, white, and Gaussian with known variance $\sigma_{w_k}^2$. These assumptions on $C_{1,k}$, $f_{1,k}$, and $w_{n,k}$ for each δ_k allows for the detection performance of the Generalized Likelihood Ratio Test (GLRT) for the following hypothesis test

$$\mathcal{H}_0 : \ r_{n,k} = w_{n,k} \tag{D.5}$$

$$\mathcal{H}_1 : \ r_{n,k} = s_{n,k} + w_{n,k} \tag{D.6}$$

to be derived in closed form [22, pgs. 484-485]. In particular, it can be shown that the GLRT is given by

$$T(\mathbf{r}_k) = \frac{\mathbf{r}_k^H \mathbf{r}_k}{\sigma_{w_k}^2 / 2} > \gamma' \tag{D.7}$$

where $\mathbf{r}_k$ is an $N \times 1$ vector of samples $r_{n,k}$ and γ' is specified using the distribution of $T(\mathbf{r}_k)$ under $\mathcal{H}_0$. In particular, $T(\mathbf{r}_k)$ adheres to the following distributions under $\mathcal{H}_0$ and $\mathcal{H}_1$

$$T(\mathbf{r}_k) \sim \begin{cases} \chi_2^2 & \text{under } \mathcal{H}_0 \\ \chi'^2_2(\lambda_k) & \text{under } \mathcal{H}_1 \end{cases} \tag{D.8}$$

where χ_2^2 is a chi-square distribution with two degrees-of-freedom and $\chi_2'^2$ is a noncentral chi-square distribution with two degrees of freedom and noncentrality parameter λ_k [22, pgs. 24-26]. For our problem

$$\lambda_k = \frac{N|C_{1,k}|^2}{\sigma_{w_k}^2/2} \tag{D.9}$$

$$= 2N \cdot \mathrm{SNR}_k \tag{D.10}$$

Assuming the signal power $|C_{1,k}|^2$ and noise power $\sigma_{w_k}^2$ are constant in each δ_k allows us to drop the subscript k from λ_k and SNR_k, which we do in the remainder of this appendix.

Given knowledge of the PDF's of $T(\mathbf{r}_k)$ we can determine P_{FA} and the probability of detection P_D associated with a single harmonic in each δ_k

$$P_{FA} = 1 - F_{\chi_2^2}(\gamma') \tag{D.11}$$

$$P_D = 1 - F_{\chi_2'^2(\lambda)}(\gamma') \tag{D.12}$$

where $F_{\chi_2^2}(\gamma')$ and $F_{\chi_2'^2(\lambda)}(\gamma')$ are the Cumulative Distribution Functions (CDF) of chi-square and non-central chi-square random variables, respectively. Solving (D.11) for γ' and substituting the result into (D.12) yields

$$P_D = 1 - F_{\chi_2'^2(\lambda)}\left(F_{\chi_2^2}^{-1}(1 - P_{FA})\right) \tag{D.13}$$

where $F_{\chi_2^2}^{-1}(p)$ denotes the inverse CDF of a chi-square random variable with two degrees-of-freedom. When evaluated at $\lambda = 2N \cdot \mathrm{SNR}$, (D.13) yields the probability of detecting a single harmonic component in δ_k as a function of N, SNR, and P_{FA}.

If we then interpret a detection in each δ_k as a Bernoulli random variable with probability of success P_D, then detecting a harmonic correctly in some fraction of the total number of short time intervals is a binomial random variable. Let α be the percent of K total intervals, where $K = \left\lfloor \frac{T}{N \cdot T_s} \right\rfloor$, then the probability of detecting a harmonic in $\lfloor \alpha \cdot K \rfloor$ or more short time intervals is

$$P_C = 1 - F_{\mathrm{bin}}(\lfloor \alpha \cdot K \rfloor, K, P_D) \tag{D.14}$$

where $F_{\mathrm{bin}}(\lfloor \alpha \cdot K \rfloor, K, P_D)$ is the CDF of a binomial random variable associated with K trials and P_D probability of success. To relate this probability to P_C as represented in

Figures 10 and 24 requires that we assume that α percent of the K intervals must yield correct detections to yield estimates of the LFMCW parameters to within 10% of their true values. The difficulty in accepting this assumption is that we don't know, without first performing a simulation, what α yields estimates accurate to within this tolerance. So, we use this result to bound the true P_C for fixed values of N, T, T_s, P_{FA}, and SNR rather than using it to predict P_C exactly.

We develop this bound by making two arguments, one for the lower bound and one for the upper bound on P_C. First, the lower bound. If we let α be sufficiently small so that $\lfloor \alpha \cdot K \rfloor = 1$ then interpreting (D.14) as the probability of observing at least one success in K trials will yield the highest possible probability for every SNR, with all other parameters fixed. Interpreting P_C as the probability of estimating the parameters of an LFMCW signal to with 10% of their true values makes it clear that doing so with, possibly, only one frequency estimate is impossible. In other words, this 'one-detection event' reflects the best possible harmonic detection performance that can *never* be achieved as a probability of correct characterization. So, evaluating (D.14) as a function of SNR for $\lfloor \alpha \cdot K \rfloor = 1$ provides a lower bound on the probability of correctly estimating the LFMCW parameters. See the blue curve in Figure 32.

The argument for the upper bound is similar to that for the lower bound. In particular, suppose we choose α to be such that $\lfloor \alpha \cdot K \rfloor = K - 1$, then observing correct detections in $K - 1$ or more time intervals becomes an extremely rare event[1], which makes the probability of observing it at each SNR level low. Since the estimation accuracy increases with increasing number of correct frequency estimates, the probability of estimating the parameters correctly with a large number of correct frequency estimates must be greater then each associated binomial probability. See the red curve in Figure 32. Finally, to test the bounds we plotted the results form the simulation performed in Section 1.5.1 for $M = 1$ and found the performance does in fact fall within these bounds.

In the example we used to facilitate the analysis, the definition of P_C was accompanied by a 10% tolerance on the parameter estimates. This was for illustrative purposes and can be

[1]Note that observing $K - 1$ or more successes is equivalent to observing $K - 1$ or K successes out of K trails. Hence, it is a rare event.

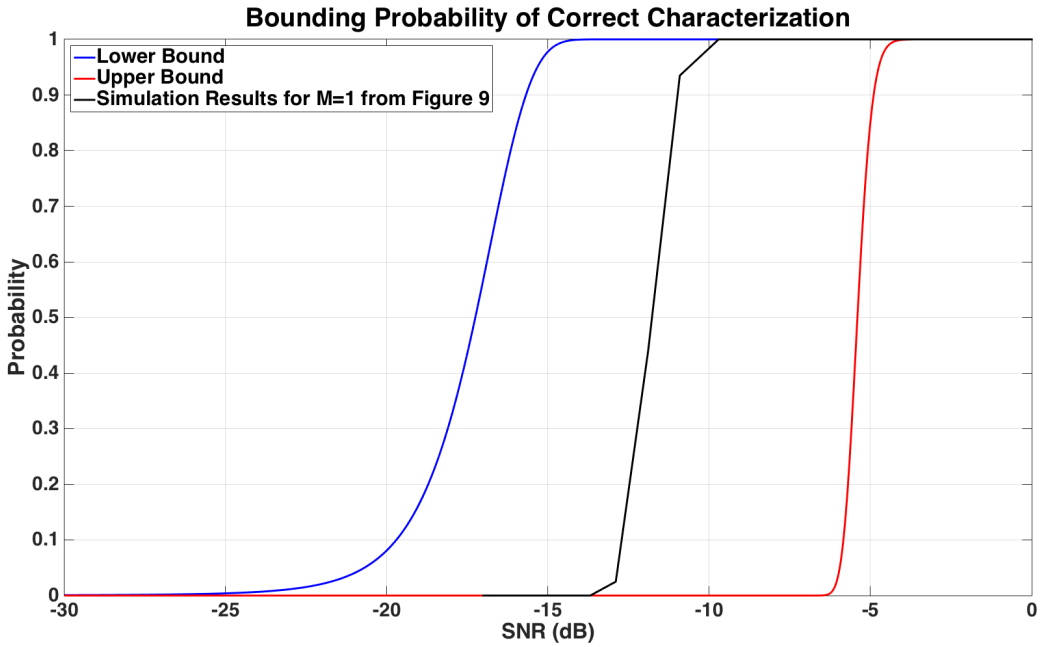

Figure 32: Bounding Probability of Correct Characterization

varied. With reference to Figure 32, for all other parameters fixed, increasing the tolerance will shift the simulated P_C vs. SNR curve to the left, but still to the right of the lower bound. Similarly, decreasing the tolerance would shift the simulated results to the right, but still to the left of the upper bound.

From Figure 32 we see that, though reliable, the gap between the upper and lower bounds is large. To decrease the width of the gap one can manipulate α. However, doing so causes our interpretation of the lower and upper bounds to be less reliable. Figure 33 includes to the P_C vs. SNR curve (in green) the value of α that yielded $P_C = 0.5$ at the same SNR as simulation. It was found to be $\alpha = 0.17$, which suggests that 17% of the total intervals needed to produce a correct frequency detection in order to achieve the performance of the sequential algorithm tested.

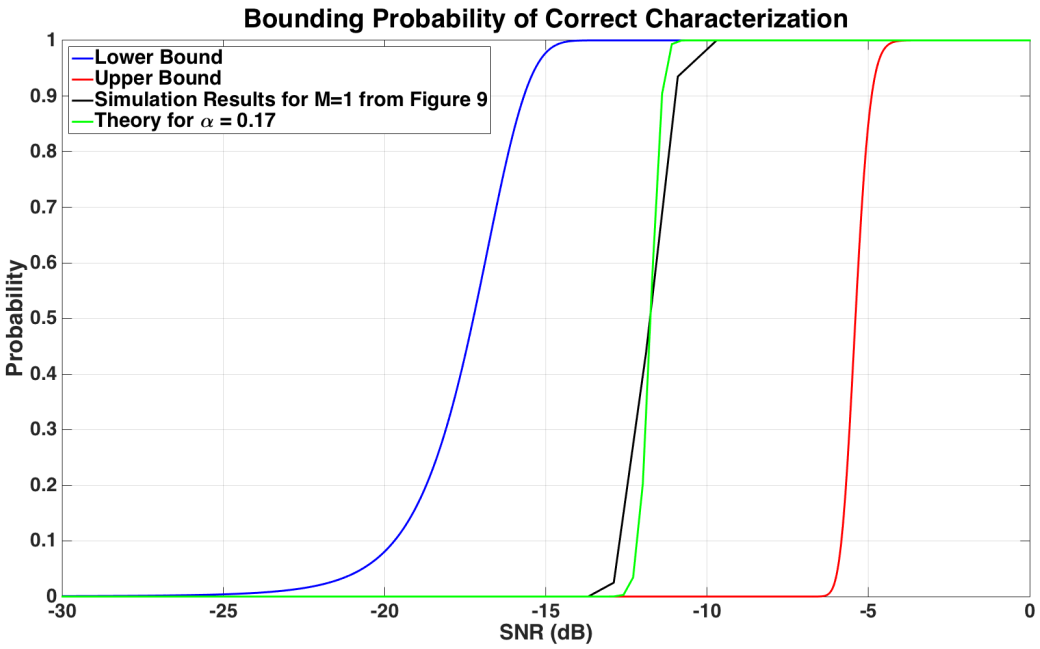

Figure 33: Result for $\alpha = 0.17$

APPENDIX E

CRLB GENERALIZATION TO $M > 1$

Recall that the Fisher Information Matrix (FIM) is

$$\left[\mathbf{I}\left(\bar{\theta}\right)\right]_{h,q} = \frac{2}{\sigma_w^2}\mathrm{Re}\left\{\sum_{n=0}^{N-1}\frac{\partial s^*\left(nT_s;\bar{\theta}\right)}{\partial\theta_{m,i}}\frac{\partial s\left(nT_s;\bar{\theta}\right)}{\partial\theta_{p,j}}\right\} \tag{E.1}$$

where rows of the FIM are now indexed by $h = 6\left(m-1\right)+i$, columns by $q = 6\left(p-1\right)+j$, individual variables by $i, j = 1, \ldots, 6$, and individual signals by $m, p = 1, 2, \ldots$. Also, we have that

$$\bar{\theta} = \left[\begin{array}{cccc} \bar{\theta}_1^T & \bar{\theta}_2^T & \cdots & \bar{\theta}_M^T \end{array}\right]^T \tag{E.2}$$

$$\bar{\theta}_m^T = \left[\begin{array}{cccccc} b_m & f_m^{\mathrm{start}} & \beta_m & \Delta_m & T_{\mathrm{swp},m} & \phi_m \end{array}\right]^T \tag{E.3}$$

$$= \left[\begin{array}{cccccc} \theta_{m,1} & \theta_{m,2} & \theta_{m,3} & \theta_{m,4} & \theta_{m,5} & \theta_{m,6} \end{array}\right]^T \tag{E.4}$$

So,

$$\frac{\partial s}{\partial\theta_{m,1}} = \frac{\partial s}{\partial b_m} = e^{j\Phi_m} \tag{E.5}$$

$$\frac{\partial s}{\partial\theta_{m,2}} = \frac{\partial s}{\partial f_m^{\mathrm{start}}} = j\frac{\partial\Phi_m}{\partial f_m^{\mathrm{start}}}s_m\left(t;\bar{\theta}_m\right) \tag{E.6}$$

$$\frac{\partial s}{\partial\theta_{m,3}} = \frac{\partial s}{\partial\beta_m} = j\frac{\partial\Phi_m}{\partial\beta_m}s_m\left(t;\bar{\theta}_m\right) \tag{E.7}$$

$$\frac{\partial s}{\partial\theta_{m,4}} = \frac{\partial s}{\partial\Delta_m} = j\frac{\partial\Phi_m}{\partial\Delta_m}s_m\left(t;\bar{\theta}_m\right) \tag{E.8}$$

$$\frac{\partial s}{\partial\theta_{m,5}} = \frac{\partial s}{\partial T_m} = j\frac{\partial\Phi_m}{\partial T_m}s_m\left(t;\bar{\theta}_m\right) \tag{E.9}$$

$$\frac{\partial s}{\partial\theta_{m,6}} = \frac{\partial s}{\partial\phi_m} = j\frac{\partial\Phi_m}{\partial\phi_m}s_m\left(t;\bar{\theta}_m\right) \tag{E.10}$$

where our results for the single signal case, found in (2.45)-(2.49), apply directly for populating (E.1). In contrast to the $M = 1$ case, the term inside the sum in (E.1) does not simplify to $\frac{\partial \Phi}{\partial \theta_{m,i}} \frac{\partial \Phi}{\partial \theta_{p,j}}$.

APPENDIX F

DEMONSTRATION THAT CRLB REGULARITY CONDITIONS ARE SATISFIED

In this Appendix we show that for an LFMCW signal with unknown parameters in CAWGN, the regularity condition

$$E\left[\frac{\partial \ln p\left(\bar{\mathbf{r}};\bar{\theta}\right)}{\partial \theta_i}\right] = 0 \qquad for\ all\ \theta_i \tag{F.1}$$

is approximately satisfied based on the development in Section 2.2.1.

Let $\bar{\mathbf{r}}$ be the Nx1 vector of samples defined by (2.1) so that

$$\bar{\mathbf{r}} = \bar{\mathbf{s}} + \bar{\mathbf{w}} \tag{F.2}$$

where $\bar{\mathbf{w}} \sim CN\left(0, \sigma_w^2\right)$. Hence, the Probability Density Function (PDF) of $\bar{\mathbf{r}}$ is given by

$$p\left(\bar{\mathbf{r}};\bar{\theta}\right) = \frac{1}{\pi^N \sigma_w^{2N}} \exp\left\{-\left(\bar{\mathbf{r}}-\bar{\mathbf{s}}\right)^H \frac{1}{\sigma_w^2}\left(\bar{\mathbf{r}}-\bar{\mathbf{s}}\right)\right\} \tag{F.3}$$

So,

$$\frac{\partial \ln p\left(\bar{\mathbf{r}};\bar{\theta}\right)}{\partial \theta_i} = \left(\bar{\mathbf{r}}-\bar{\mathbf{s}}\right)^H \frac{1}{\sigma_w^2}\frac{\partial \bar{\mathbf{s}}}{\theta_i} \tag{F.4}$$

$$= \bar{\mathbf{w}}^H \frac{1}{\sigma_w^2}\frac{\partial \bar{\mathbf{s}}}{\theta_i} \tag{F.5}$$

Through our approximations of $\bar{\mathbf{s}}$ from Section 2.2.1 for Δ_m and $T_{swp,m}$, (F.5) exits $\forall \theta_i$ and since $E\left\{\bar{\mathbf{w}}\right\} = \mathbf{0}$, we have that

$$E\left\{\frac{\partial \ln p\left(\bar{\mathbf{r}};\bar{\theta}\right)}{\partial \theta_i}\right\} = \mathbf{0} \tag{F.6}$$

which completes the demonstration that the regularity condition is satisfied up to our approximation of the signal.

BIBLIOGRAPHY

[1] Brian DO Anderson and John B Moore. *Optimal Filtering. 1979*. Prentice-Hall, Englewood Cliffs, NJ, 1st edition, 1979.

[2] Yaakov Bar-Shalom, Peter K Willett, and Xin Tian. *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.

[3] Kristine L Bell, Yossef Steinberg, Yariv Ephraim, and Harry L Van Trees. Extended ziv-zakai lower bound for vector parameter estimation. *Information Theory, IEEE Transactions on*, 43(2):624–637, 1997.

[4] L. Cohen. *Time-Frequency Analysis*. Prentice Hall Signal Processing Series. Prentice Hall, 1995.

[5] Keith L Davidson and Patrick J Loughlin. Compensating for window effects in the calculation of spectrographic instantaneous bandwidth. *Biomedical Engineering, IEEE Transactions on*, 47(4):556–558, 2000.

[6] BL Decker. World geodetic system 1984. Technical report, DTIC Document, 1986.

[7] J Devore. *Probability and statistics for engineers and scientists*. CengageBrain.com, 4th edition, 2004.

[8] Lorenzo Galleani, Leon Cohen, and Bruce Suter. Local stationarity and time-frequency distributions. In *SPIE Optics+ Photonics*, pages 63130Z–63130Z. International Society for Optics and Photonics, 2006.

[9] Francis G Geroleo and Maite Brandt-Pearce. Detection and estimation of lfmcw radar signals. *Aerospace and Electronic Systems, IEEE Transactions on*, 48(1):405–418, 2012.

[10] Michael T. Grabbe. A correlation algorithm for geo-location position measurements, 2011.

[11] Michael T. Grabbe and Brandon M. Hamschin. *Passive Emitter Geo-Location: Measurements and Algorithms*. The Johns Hopkins University Applied Physics Lab, 11100 Johns Hopkins Rd, Laurel, MD, 20723-6099, 1st edition, September 2012.

[12] Michael T. Grabbe and Brandon M. Hamschin. Time difference of arrival estimation for pulsed signals. Memo KVW-0-12-U-001, The Johns Hopkins University Applied Physics Laboratory, Laurel, MD, February 2012.

[13] Michael T Grabbe and Brandon M Hamschin. Geo-location using direction finding angles. *Johns Hopkins APL Technical Digest*, 31(3):254–262, 2013.

[14] Michael T Grabbe and Brandon M. Hamschin. Geo-location of lpi signal signal sources. Technical Report KVW-0-15U-001, The Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Rd., Laurel, MD 20723, April 2015.

[15] Michael T Grabbe, Brandon M Hamschin, and Andrew P Douglas. A measurement correlation algorithm for line-of-bearing geo-location. In *Aerospace Conference, 2013 IEEE*, pages 1–8. IEEE, March 2013.

[16] B. Hamschin, J. Clancy, M. Grabbe, M. Fortier, and J. Novak. Detection, characterization, and localization of multiple lfmcw lpi signals. Radar Conference, May 2014.

[17] B. Hamschin, J. Clancy, M. Grabbe, M. Fortier, and J. Novak. Passive detection, characterization, and localization of multiple lfmcw lpi signals. SPIE Automatic Target Recognition, May 2014.

[18] B. M. Hamschin, Michael T. Grabbe, and John F. Ferguson. An approximate cramer-rao lower bound for multiple lfmcw signals. *IEEE Trans. Aerospace and Electronic Systems*, submitted.

[19] B. M. Hamschin, Michael T. Grabbe, and John F. Ferguson. A sequential method for passive detection, characterization, and localization of multiple lfmcw lpi signals. *IEEE Trans. Aerospace and Electronic Systems*, submitted.

[20] B. M. Hamschin, Michael T. Grabbe, and John F. Ferguson. A sequential method for passive detection, characterization, and localization of multiple lfmcw lpi signals - experimental results. *IEEE Trans. Aerospace and Electronic Systems*, submitted.

[21] S. Kay. *Fundamentals of Statstical Signal Processing: Estimation Theory*, volume 1 of *Prentice Hall Signal Processing Series*. Prentice Hall, 1993.

[22] S. Kay. *Fundamentals of Statstical Signal Processing: Detection Theory*, volume 2 of *Prentice Hall Signal Processing Series*. Prentice Hall, 1998.

[23] S. Kay. Optimal signal design for detection of gaussian point targets in stationary gaussian clutter/reverberation. *IEEE J. Sel. Topics in Sig. Proces. Mag.*, 1(1):31–41, June 2007.

[24] Steven Kay. *Intuitive probability and random processes using MATLAB®*. Springer Science & Business Media, 2006.

[25] Steven Kay and G Boudreaux-Bartels. On the optimality of the wigner distribution for detection. In *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP'85.*, volume 10, pages 1017–1020. IEEE, 1985.

[26] Peter J Kootsookos. A review of the frequency estimation and tracking problems. Technical report, Systems Engineering Department, Australian National University, 1993.

[27] R. Larson, R.P. Hostetler, and B.H. Edwards. *Calculus: Early Transcendental Functions*. Houghton Mifflin (Boston), 1999.

[28] D.G. Luenberger and Y. Ye. *Linear and nonlinear programming*. Springer Verlag, 2008.

[29] The Math Works, Inc., 24 Prime Park Way, Natick, MA 01760-1500. *Optimization Toolbox User's Guide Version 5*, 5.1 edition, September 2002.

[30] The Math Works, Inc., 24 Prime Park Way, Natick, MA 01760-1500. *Signal Processing Toolbox User's Guide Version 5 Toolbox User's Guide Version 5*, r2015b edition, September 2015.

[31] John H Mathews and Kurtis D Fink. *Numerical methods using MATLAB*, volume 31. Prentice hall Upper Saddle River, NJ, 3rd edition, 1999.

[32] Robert N McDonough and Anthony D Whalen. *Detection of signals in noise*. Academic Press, 1995.

[33] Kenneth S Miller. Complex linear least squares. *Siam Review*, 15(4):706–726, 1973.

[34] VK Nguyen. A low complexity parameter estimation technique for lfmcw signals. *Aerospace and Electronic Systems, IEEE Transactions on*, 50(4):2554–2563, 2014.

[35] NIMA. Department of defense world geodetic system 1984, its definition and relationships with local geodetic systems. Technical Report TR8350.2, 3 rd Ed., NIMA, January 2000.

[36] Phillip E Pace. *Detecting and classifying low probability of intercept radar*. Artech House, 685 Canton Street, Norwood, MA 02062, 1st edition, 2004.

[37] Phillip E Pace. *Detecting and classifying low probability of intercept radar*. Artech House, 685 Canton Street, Norwood, MA 02062, 2nd edition, 2009.

[38] A. Papoulis, S.U. Pillai, and S. Unnikrishna. *Probability, random variables, and stochastic processes*, volume 73660116. McGraw-Hill New York, 2002.

[39] S. Peleg and B. Porat. The cramer-rao lower bound for signals with constant amplitude and polynomial phase. *Signal Processing, IEEE Transactions on*, 39(3):749–752, Mar 1991.

[40] Donald B Percival and Andrew T Walden. *Spectral analysis for physical applications*. Cambridge University Press, 1993.

[41] Alfred Price. *The History of US Electronic Warfare:"The Years of Innovation-Beginnings to 1946"*, volume 1. Association of Old Crows, 1984.

[42] Alfred Price. *The History of US Electronic Warfare: "The Renaissance years, 1946 to 1964"*, volume 2. Association of Old Crows, 1989.

[43] Alfred Price. The history of us electronic warfare, volume iii, rolling thunder through allied force, 1964 to 2000, 1989.

[44] Mark A Richards. *Fundamentals of radar signal processing*. Tata McGraw-Hill Education, 2005.

[45] John S Seybold. *Introduction to RF propagation*. Wiley, 2005.

[46] SIMRAD. Broadband 4g radar, June 2014.

[47] M.I. Skolnik, editor. *Radar Handbook*. McGraw-Hill Book Co., Inc., 1970.

[48] Seymour Stein. Algorithms for ambiguity function processing. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 29(3):588–599, 1981.

[49] P. Stoica and R.L. Moses. *Spectral analysis of signals*. Pearson/Prentice Hall, 2005.

[50] Ettus Technology. Usrp n210, April 2015.

[51] D.J. Thomson. Spectrum estimation and harmonic analysis. *Proceedings of the IEEE*, 70(9):1055–1096, 1982.

[52] H. L. Van Trees. *Detection, Estimation and Modulation Theory, Part III*. Wiley-Interscience, 1971.

[53] Jingbo Wang, Wai Wan Tsang, and George Marsaglia. Evaluating kolmogorov's distribution. *Journal of Statistical Software*, 8(18), 2003.

[54] Wikipedia. Ism band — wikipedia, the free encyclopedia, 2015. [Online; accessed 1-June-2015].

[55] Wikipedia. Logistic function — wikipedia, the free encyclopedia, 2015. [Online; accessed 11-June-2015].

[56] Richard G Wiley. *ELINT: The interception and analysis of radar signals*. Artech House Boston, 2006.

[57] William T. Vetterling William H. Press, Saul A. Teukolsky and Brian P.Flannery. *Numerical Recipes in C: The art of scientific computing*. Cambridge university press, 2nd edition, 1988.

[58] P. M. Woodward. *Probability and Information Theory, with Applications to Radar.* McGraw-Hill Book Co., Inc., 1953.