

ABSTRACT

Title of dissertation: SUB-NYQUIST SENSING AND SPARSE
RECOVERY OF WIDE-BAND INTENSITY
MODULATED OPTICAL SIGNALS

Robert Lee
Doctor of Philosophy, 2018

Dissertation directed by: Professor Christopher Davis
Department of Electrical
and Computer Engineering

Intensity modulated optical transmitters, wide-bandwidth electro-optical receivers, highspeed digitizers, and digital matched-filters are being used in hybrid lidar-radar systems to measure the range and reflectivity of objects located within degraded visual underwater environments. These methods have been shown to mitigate the adverse effects of the turbid underwater channel due to the de-correlation of the modulated optical signal after undergoing multiple scattering events. The observed frequency-dependent nature of the underwater channel has driven the desire for wider bandwidth waveforms modulated at higher frequencies in order to improve range accuracy and resolution. While the described system has shown promise, the matched filter processing scheme, which is also widely used in the fields of radar and sonar, suffers from inherent limitations. One limitation is based on the achievable range resolution as dictated by the classical time-frequency uncertainty principle, where the bandwidth dictates the measurable resolution. The side-lobes generated

during the matched filtering process also present a challenge when trying to detect multiple targets. These limitations are further constrained by currently-available analog-to-digital conversion technologies which restrict the ability to directly sample the wide-band modulated signals. Even in cases where the technology exists that can operate at sufficient rates, often it is prohibitively expensive for many applications and high data rates can pose processing challenges.

This research effort addresses both the restrictions imposed by the available analog-to-digital conversion technologies and the limited resolution of the existing time-frequency methods for wide-band signal processing. The approach is based on concepts found within the fields of compressive sensing and sparse signal recovery and will be applied to the detection of objects illuminated with wide-band intensity modulated optical signals. The underlying assumption is that given the directive nature of laser propagation, the illuminated scene is inherently sparse and the limited number of reflecting objects can be treated as point sources. The main objective of this research is to provide results that show, when sampling at rates below those dictated by the traditional Shannon-Nyquist sampling theorem, it is possible to make more efficient use of the samples collected and detect a limited number of reflecting targets using specialized recovery algorithms without reducing system resolution. Through theoretical derivations, empirical simulations, and experimental investigation, it will be shown under what conditions the sub-Nyquist sampling and sparse recovery techniques are applicable, and how the described methods influence resolution, accuracy, and overall performance in the presence of noise.

SUB-NYQUIST SENSING AND SPARSE RECOVERY OF
WIDE-BAND INTENSITY MODULATED OPTICAL SIGNALS

by

Robert Lee

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2018

Advisory Committee:
Professor Christopher Davis, Chair/Advisor
Professor Behtash Babadi
Professor Rama Chellappa
Professor Douglas Currie
Dr. Linda Mullen

© Copyright by
Robert Lee
2018

Table of Contents

List of Tables	iv
List of Figures	v
1 Introduction	1
1.1 Underwater Optics	1
1.1.1 Absorption and Scattering	2
1.1.2 Underwater Optical Impulse Response	3
1.2 Hybrid Lidar-Radar	5
1.2.1 Benefits and Limitations	7
1.3 Research Goals	8
1.3.1 Paper Organization	10
2 Sub-Nyquist Sensing and Sparse Recovery	12
2.1 The Sparse Approximation	12
2.2 Sensing Matrix Formulation	14
2.3 Recovery of Sparse Vectors	16
2.3.1 L-1 Minimization	17
2.3.2 Necessary Conditions for Recovery	18
2.4 Discussion	20
3 Waveform Design Consideration for Sub-Nyquist Sensing	22
3.1 Coherence Investigation	23
3.1.1 Pseudo-Noise Code	24
3.1.2 Frequency Modulated Continuous Wave	27
3.2 Discussion	31
4 Sparse Recovery Algorithms	34
4.1 Convex Optimization	35
4.2 Greedy Algorithms	36
4.2.1 Orthogonal Matching Pursuit	37
4.2.2 Hard Thresholding Pursuit	39
4.3 Analysis of Recovery of Sparse Vectors	40

4.3.1	Simulation Design	40
4.3.1.1	Performance Metrics	41
4.3.2	Simulation Results	42
4.3.2.1	Results for Convex Optimization Algorithm	43
4.3.2.2	Results for OMP and HTP Algorithms	48
4.3.2.3	Discussion and Algorithm Selection	52
5	Analysis of Noise Susceptibility	54
5.1	Recovery of Sparse Vectors in the Presence of Noise	54
5.1.1	Simulation Method	55
5.1.2	Simulations Results	57
5.2	Experimental Verification	61
5.2.1	Experimental Design	62
5.2.2	Experimental Results	63
5.3	Summary	65
6	Resolution of Sparse Recovery Techniques	67
6.1	Resolution of Traditional Methods	67
6.2	Resolution Enhancement of Sub-Nyquist Chirp Sensing Matrix	70
6.2.1	Recovery of Closely Spaced Returns	73
6.2.2	Recovery of Shadowed Returns	78
6.3	Summary	81
7	Application to Underwater Lidar	82
7.1	Verification of the Sparse Approximation	82
7.1.1	Modeling the Underwater Hybrid Lidar-Radar System	83
7.1.2	Simulation Results	85
7.2	Experimental Verification	90
7.2.1	Experimental Design, Equipment, and Details	91
7.2.2	Clean Water Baseline Measurements	95
7.2.3	Experimental Results for Turbid Environments	98
8	Conclusion	103
8.1	Summary	103
8.2	Future Research Directions	106
	Bibliography	108

List of Tables

5.1	Simulations Results comparing the Matched Filter and Sub-Nyquist Sensing methods in the presence of noise	61
7.1	Max number of attenuation lengths caused by scattering that still result in successful detection during simulation and modeling of the underwater impulse response	90
7.2	Experimental results showing max recoverable attenuation lengths for the matched filter detector, and the OMP algorithm given various system sampling rates	102

List of Figures

1.1	Example of the turbid underwater channel frequency response	4
1.2	Pseudo mono-static laser ranging setup [1]	6
1.3	Example of Hybrid Lidar Radar system utilizing a matched filter receiver. (a) highlights the time delay due to propagation, (b) points to the dc offset caused by the de-correlated backscatter, (c) highlights the frequency dependent amplitude attenuation caused by forward scatter, and (d) points to the matched filter output which has a peak location proportional to the target distance and a resolution proportional to the waveform bandwidth.	7
1.4	Block Diagram of the sub-Nyquist Hybrid lidar-radar system. The red highlighted portion shows the changing hardware/processing when compared to figure 1.3	9
2.1	Pictorial representation of the sub-Nyquist sensing system. The sparse vector s has a non-zero elements at the targets location and multiplies the waveform matrix Ψ which has columns that correspond to all possible time shifts. The time shifted signal is then multiplied by the sampling matrix Φ which generates the sub-Nyquist rate measurements, y , of the time shifted waveform.	15
3.1	Auto-Correlation Function for a PN-Coded sequence	25
3.2	Empirically measured coherence values for the sub-Nyquist sensing matrices derived from PN-coded sequences. (a) shows the coherence as a function of the down-sampling parameter $\delta = m/N$, (b) is the estimated max recoverable support for a given sensing matrix. . . .	26
3.3	Autocorrelation function of a frequency modulated continuous wave signal.	28
3.4	Empirically measured coherence values for the sub-Nyquist sensing matrices derived from FMCW signals. (a) shows the coherence as a function of the down-sampling parameter $\delta = m/N$, (b) is the estimated max recoverable support for a given sensing matrix. . . .	28
3.5	Comparison of coherence results for sensing matrices derived from PN-coded and FMCW waveforms, $N = 512$. Shown on a log-log plot to highlight the PN-Coded waveforms logarithmic dependence on δ .	30

3.6	Empirically measured coherence values for the sub-Nyquist sensing matrices derived from FMCW signals with relatively prime m and N . (a) shows the coherence as a function of the down-sampling parameter δ , (b) is the estimated max recoverable support for a given sensing matrix.	31
4.1	Average NSE value of the convex optimization algorithm L_1 -magic [2] mapped to a color value and displayed as a function of the vector support K and the down-sampling parameter $\delta = m/N$. Simulation has $N = 199$	44
4.2	Probability of exact recovery as a function of the Support K and the down-sampling parameter $\delta = m/N$ where $N = 199$	45
4.3	Maximum number of recoverable sparse elements within a single vector as a function of the down-sampling parameter.	45
4.4	Max recoverable sparsity of a vector using convex optimization techniques.	47
4.5	Average computation time necessary to recover a signal with maximal support at a given down-sampling ratio parameter $\delta = m/N$	47
4.6	Average NSE value of the greedy optimization algorithms HTP and OMP mapped to a color value and displayed as a function of the vector support K and the down-sampling parameter $\delta = m/N$. Simulation has $N = 199$	49
4.7	Max recoverable sparsity ρ of a vector using greedy algorithms.	49
4.8	Computation time needed to obtain the optimally sparse solution using greedy algorithms.	50
4.9	Max recoverable sparsity using the greedy OMP and HTP algorithms. The support estimate was set such that $K_{EST} = m$	51
5.1	Simulation results for various SNR levels. Results show that even when sampling at rates far less than Nyquist it is possible to recover a sparse vector with unit support.	58
5.2	Minimum required down-sampling parameter necessary for 99% probability of detection Vs N for various SNR levels	58
5.3	Simulation results showing the Matched Filter performance as a function of chirp duration and SNR	60
5.4	Block diagram for preliminary experiments	63
5.5	Experimental Vs Simulated Results for $N \approx 250$ with a SNR of 0 and -10 dB.	64
5.6	Experimental Vs Simulated Results for $N \approx 1000$ with a SNR of 0 and -10 dB.	64
6.1	Example of the LFM waveform and the ambiguity function arising from pulse compression	68
6.2	Example of the inherent resolution limit of the pulse compression techniques.	69

6.3	Example of the shadowing effect from large returns near weaker returns.	70
6.4	Coherence shown as a function of the down sampling parameter for the over complete sub-Nyquist sensing matrix.	72
6.5	Recovery of closely spaced sparse vectors using the resolution enhanced sensing matrix compared to matched filter detection.	73
6.6	Detection method results for the resolution test. Two closely spaced targets of equal magnitude with varying separation were used and the magnitude of the detection method outputs were mapped to a color value. Plotted as a function of recovered index and target separation. The solid and dashed red lines represent the actual target indexes. Yellow values are the estimated target location for the given method.	75
6.7	Results of the simulations where two targets are spaced at varying distances, and multiple sub-Nyquist rates are utilized. The OMP and Convex Optimization algorithms were used to recover the sparse vector, and the magnitude of the recovered results was mapped to a color.	77
6.8	Results from Nyquist rate simulations of the shadowing effect. Shows that sparse recovery methods are not effected by side-bands and shadowing like the matched filtering methods.	80
6.9	Results from Sub-Nyquist rate simulations of the shadowing effect. Shows that sparse recovery methods are not effected by side-bands like the matched filtering methods are.	80
7.1	Example of the underwater channel impulse response model output for clean and turbid water.	84
7.2	Example of the underwater channel impulse response model output for clean and turbid water.	85
7.3	Block diagram for modeling the underwater hybrid lidar-radar sub-Nyquist sensing system.	86
7.4	Simulation results for the modeled sub-Nyquist underwater hybrid lidar-radar system. The range error was calculated for various test parameters and plotted as a function of attenuation length.	88
7.5	Block diagram of experimental setup	93
7.6	Example digitizer output waveforms at various sampling rates.	94
7.7	Initial experimental results of the sub-Nyquist sensing and sparse recovery of a target range situated in underwater environments. Errors in detection due to dictionary mismatch.	96
7.8	Magnitude Spectrum of the waveform used to generate the sensing matrix compared to the magnitude spectrum of the PMT signal.	97
7.9	Initial experimental results of the sub-Nyquist sensing and sparse recovery using a sensing matrix that is derived from a Nyquist sampled version of the PMT reference signal. Effects of dictionary mismatch have been suppressed	99
7.10	Experimental ranging results for $C_{532} = 1.30m^{-1}$. Detection error is evident when sampling at rates less than Nyquist.	100

7.11 Experimental ranging results for $C_{532} = 2.10m^{-1}$. While the correct range measurements are reported for closer in targets, all detection methods display range errors at further target positions. 101

Chapter 1: Introduction

Maintaining situational awareness in underwater environments has been a subject of interest for many years. Traditionally, acoustic wave based technologies have been the primary methods used to measure the distance to objects located in underwater environments. Due to the propagation characteristics of acoustic waves and speed of sound in water, these methods are limited in terms of achievable resolution and update rates [3]. While electrically steered arrays and synthetic aperture techniques can be used to improve system resolution [4, 5], these techniques cannot be utilized for all applications. Furthermore, traditional sonar methods allow for identification of the source's location which is often undesirable. These limitations have spurred interest in active electro-optical systems which could provide measurements of the range and reflectivity of illuminated objects. Specifically, laser based technologies are of interest [6, 7] due to their highly directive nature and potential for high speed range measurements with increased resolution and improved accuracy.

1.1 Underwater Optics

Laser-based methods for underwater optical ranging are attractive due to their covertness as well as the potential for improvements in resolution, accuracy, and

detection speed. However, there exist significant challenges imposed by the underwater environment that must be addressed. This section will introduce the basics of underwater optics and will provide details on how the underwater optical channel adversely affects optical ranging methods. An optical ranging system based on a hybrid lidar-radar architecture will be introduced, and signal processing methods to determine the range to targets will be discussed. The limitations of these method will then be highlighted.

1.1.1 Absorption and Scattering

Optical signals propagating through underwater environments suffer from two fundamental adverse effects: absorption and scattering. These inherent optical properties are predominately a consequence of the composition of the water sample and are typically reported as a function of wavelength. In open ocean environments, blue light is absorbed least while in turbid coastal waters, green wavelengths propagate the furthest. The shift in the absorption spectrum is due to the presence of particulates and dissolved material [8,9] in the water sample. The effect of absorption can be minimized by choosing wavelengths that are tailored to the specific underwater environment. Scattering occurs due to the interaction between the transmitted light and particulates in the water column. The angular distribution of scattered light is described by the volume scattering function [10], which is the probability that light will scatter into a certain angle. For typical natural waters, the volume scattering function is highly peaked at small angles (<10 degrees) due to the fact that the par-

ticles are larger than the blue-green wavelengths used for illumination [11, 12]. This forward-scattered light is that which interacts with the target of interest. A small percentage of light is scattered in the backward direction (near 180 degrees) and contains no information about the target.

The optical power transmitted through the underwater channel decays exponentially as a function of both scattering and absorption. This relationship is given by (1.1), where P_o is the initial optical power, a and b are the absorption and scattering coefficients respectively in units of m^{-1} . The absorption and scattering coefficients can be combined to create c , the total beam attenuation coefficient.

$$P(z) = P_o e^{-(a+b)z} = P_o e^{-cz} \quad (1.1)$$

The propagation distance z is commonly multiplied by the total beam attenuation coefficient to give a unit-less parameter, cz , the number of attenuation lengths. At one attenuation length, the received power would decrease by a factor of e^{-1} .

1.1.2 Underwater Optical Impulse Response

The scattering caused by the turbid underwater environment also gives rise to temporal dispersion of optical signals propagating through the channel [13]. Lasers with short pulse durations have been applied to ranging and imaging applications. However, the temporal dispersion caused by scattering limits the range resolution and accuracy due to broadening of the transmitted pulse [14]. These techniques also suffer from the effects of backscatter which generates a large, distributed return that arrives before target-reflected light, cluttering the signal of interest and limiting

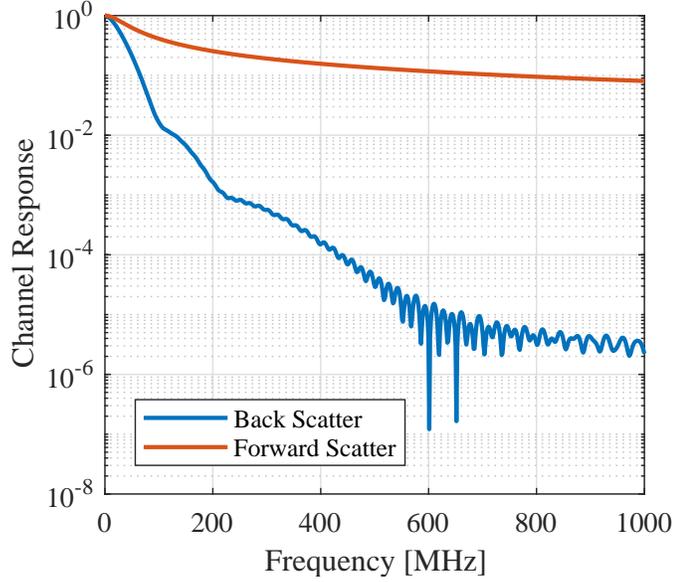


Figure 1.1: Example of the turbid underwater channel frequency response

detection performance. To overcome these effects, range gated receivers have been used to reduce the observed backscatter signal [15]. However, these methods create undesirable blind spots for ranges outside of the gate 'on' time.

By measuring the dispersion of short pulses through the underwater optical channel, the channel's frequency response can be determined via a Fourier Transform. An example of the frequency response of backscattered and forward-scattered light is shown in figure 1.1. The example frequency response shows that when scattering is observed in the underwater optical channel, the light scattered in the backwards direction has a low-pass nature, while light scattered in the forward direction observes minimal variation across the entire spectrum. This implies that backscattered and forward-scattered light have different spectral content. In the next section, it will be described how the frequency-dependent nature of the un-

derwater channel's response to modulated light can be leveraged to discriminate between forward and backward scattered light.

1.2 Hybrid Lidar-Radar

The frequency-dependent nature of the turbid underwater channel has led to the exploration of hybrid lidar-radar schemes [1] where the optical carrier is intensity modulated with a radar waveform. The hybrid lidar-radar ranging system has four main components: the radar waveform generator, the laser transmitter, the electro-optic detector, and the radar waveform receiver. The system geometry targeted in this research is a pseudo mono-static architecture, highlighted in figure 1.2. This type of geometry is considered pseudo mono-static because the transmitter and receiver are on the same platform but do not have the same optical axis or share the same aperture. Systems with this geometry are affected by light backscattered within the overlap region between the transmitted laser beam and the receiver field of view. Light that is scattered back towards the receiver without reaching the target will contain no information of the object's range and reflectivity. In addition to the effects of backscattered light, photons that scatter multiple times in the forward direction will cause blurring of the system point spread function and can increase the measured propagation delay due to the effective increased path length. Systems utilizing this architecture can suppress the observed backscatter and forward scatter by reducing the receiver/transmitter overlap region via minimizing the receiver field of view. While this is beneficial for scatter suppression, it creates blind spots where

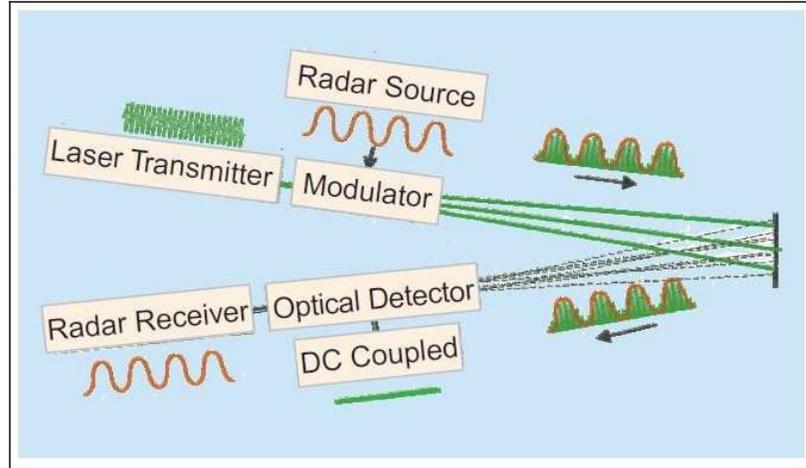


Figure 1.2: Pseudo mono-static laser ranging setup [1]

the receiver/transmitter field of view does not overlap.

Using the hybrid lidar-radar methods, it has been shown that modulation frequencies over 100 MHz can provide discrimination between backscattered and forward scattered light [16–18]. This claim is also supported by the frequency characteristics of the underwater channel shown in figure 1.1 Research also supports that when operating at higher frequencies, the effect of forward scatter can be minimized [19,20]. The hybrid lidar-radar technologies rely on the de-correlation of the modulated signal as light is randomly scattered throughout the round trip propagation to the target. Meanwhile, the modulation on the light that is minimally scattered remains correlated. In addition to the scatter discrimination feature of intensity modulated optical signals, the use of wide-band radio frequency waveforms allows for the application of radar processing techniques after the optical to electrical conversion [21–24]. This feature provides a way to apply the well established radar

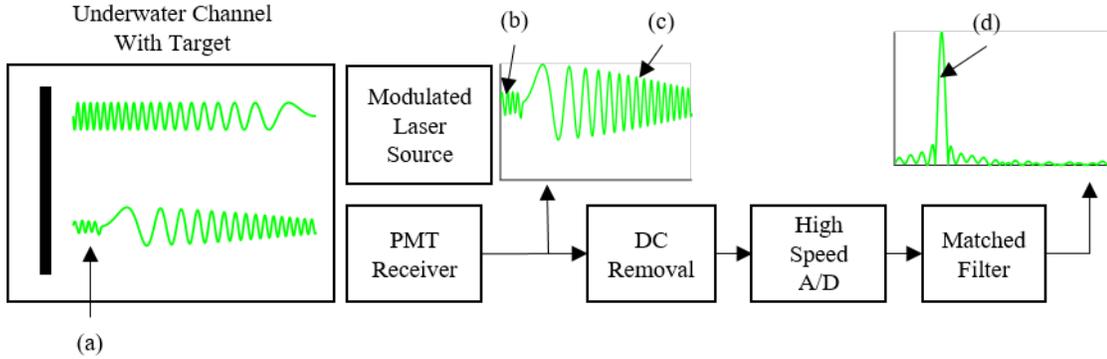


Figure 1.3: Example of Hybrid Lidar Radar system utilizing a matched filter receiver. (a) highlights the time delay due to propagation, (b) points to the dc offset caused by the de-correlated backscatter, (c) highlights the frequency dependent amplitude attenuation caused by forward scatter, and (d) points to the matched filter output which has a peak location proportional to the target distance and a resolution proportional to the waveform bandwidth.

processing techniques to an underwater ranging application where typical radar methods can not be used directly due high attenuation. The detailed operation of this method is highlighted in figure 1.3. In this figure, it is illustrated how scattering affects the intensity modulated waveforms, and how the matched filtering technique can be used to recover the target’s range.

1.2.1 Benefits and Limitations

The use of wide-band waveforms and traditional radar processing methods such as matched filtering has shown substantial improvements in the field of underwater optical ranging and imaging [20, 22, 23]. These benefits include the suppression of temporal dispersion caused by forward scattering and the ability to remove backscatter through selective filtering. Additionally, system resolution is improved

due to the lack of pulse broadening after the matched filter [14]. These observed benefits have prompted the desire for higher modulation frequencies and waveforms with increasing bandwidths. Similar to wide-band technologies in radar and sonar, these detection methods are fundamentally resolution-limited by the classical time-frequency uncertainty principles [25]. Additionally, currently-available analog to digital conversion technology is a limiting factor in the design of wide-band detection systems. Following similar trends in software defined radios [26], it is desirable to put the analog to digital conversion process as close to the receiver element as possible to directly digitize the RF waveforms. Using this approach, receiver flexibility can be maximized, and the number of noise sources can be reduced. Given the high speed and wide bandwidth nature of optical sources, and the current state of analog to digital conversion technologies, it is possible to transmit much more bandwidth than a single wide-band receiver can digitize. Even if components exist that can operate at the exceedingly high frequencies and bandwidths, they are prohibitively expensive for many applications [27]. Additionally, the necessary high data rates make real time processing a challenging effort.

1.3 Research Goals

The objective of this research effort is to address both the limited resolution of the existing time-frequency methods for wide-band signal processing and the restrictions imposed by the available analog to digital conversion technologies. The studied approach is based on concepts within the theory of compressed sensing and

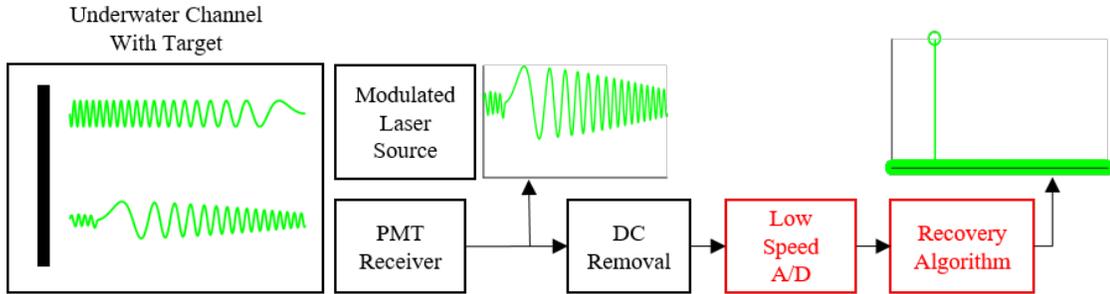


Figure 1.4: Block Diagram of the sub-Nyquist Hybrid lidar-radar system. The red highlighted portion shows the changing hardware/processing when compared to figure 1.3

sparse signal reconstruction. The research aims to show that for the detection of a limited number of reflecting targets, specialized recovery algorithms can be used to improve system resolution and reduce the rate at which samples must be collected. In turn, this will prove that it is possible to reduce the necessary sampling rate of the analog to digital components to below the traditional Nyquist sampling rate. As related to the Hybrid lidar-radar sensing system, this will be achieved by replacing the the high speed Analog to Digital (A/D) converter and the matched filter with a low rate A/D and a specialized recovery algorithm. This change is highlighted in red in figure 1.4. The hybrid lidar-radar system will be used as the target test bed for theoretical investigations and experimental verification. The research objectives will be met by first providing theoretical and empirical evidence that sub-Nyquist rates can be used successfully. Simulations and experiments will also be used to prove that these methods can be applied to optical ranging systems.

1.3.1 Paper Organization

The dissertation is organized as follows. In chapter 2, the sub-Nyquist sensing framework will be introduced and the linear algebraic formulation of the hybrid lidar-radar sensing system will be derived. Theoretical conditions that must be met for successful recovery will also be described. Chapter 3 will introduce two different types of waveforms commonly used in ranging systems, and their application to the sub-Nyquist sensing framework will be discussed. Empirical metrics will be measured and a trade-off analysis will be made. Chapter 4 will introduce three different recovery algorithms and will report on the results of simulations conducted to measure performance as a function sampling rate. Chapter 5 will investigate how noise affects the performance of the sub-Nyquist sensing and sparse recovery methods, and the relationships between noise susceptibility and sensing system parameters will be made. The matched filter detector typically used in Nyquist rate wide-band sensing systems will be used as a baseline comparison of performance. Chapter 5 will also provide the initial experimental proof that accurate sparse vectors can be recovered from noise corrupted sub-Nyquist measurements. Chapter 6 will investigate the resolution of the sparse recovery techniques and will describe how the sensing system can be re-derived to achieve higher resolution than those of typical matched filter detection methods. The resolution study will address two common problems observed in not only lidar systems but also those of radar and sonar. The first difficulty is the recovery of targets spaced within the ambiguity function, and the second is weak reflectors near larger more reflective targets. Chapter 7 will provide

model and experimental proof that sub-Nyquist sensing and sparse recovery can be used to measure the range to targets located in turbid underwater environments. Lastly, chapter 8 will discuss the potential future directions for this research.

Chapter 2: Sub-Nyquist Sensing and Sparse Recovery

The Shannon-Nyquist sampling theorem dictates the need to sample signals at a rate greater than twice the signal bandwidth. This central principle of signal processing theory has been thought of as necessary for the error free reconstruction of band-limited signals [28, 29]. Recently, an alternative theory of compressive sensing has emerged and states that given certain conditions it is possible to recover signals that are sparse or compressible with far fewer measurements than what is traditionally considered necessary [30, 31]. These methods rely on non-adaptive linear projections/measurements that preserve signal information and allow for the reconstruction through some recovery process. The following section will describe what it means for a signal to be sparse and how that approximation is valid for the hybrid lidar-radar ranging application. The measurement process for such systems will be recast as a linear program, and the methods for signal recovery along with the necessary conditions required for a successful recovery process will be detailed.

2.1 The Sparse Approximation

A signal or vector $s \in \mathbb{R}^N$ is said to be K -sparse if its support, the number of non-zero elements, is far less than the total number of elements N . Similarly, a

signal is compressible if it has a K -sparse representation in an orthonormal basis or tight frame [32]. For example, a signal comprising a collection of Dirac pulses is inherently sparse, whereas a collection of sinusoidal signals is compressible in the Fourier Basis. For the hybrid lidar-radar ranging application, the first approximation is that the optically illuminated scene is only composed of a finite number of K point reflecting objects. For many cases, $K = 1$ due to the directive nature of laser source illumination. However, to maintain initial generality, that limit is not imposed here. The sparse approximation can be represented by (2.1) where the variables α_i and τ_i are the i^{th} object's reflectivity and propagation time delay, respectively. Sparse representations of this type belong to a broader set of signals with a finite rate of innovation [33].

$$s(t) = \sum_{i=1}^K \alpha_i \delta(t - \tau_i) \quad (2.1)$$

When considering the hybrid lidar-radar application, the sparse approximation is related to the optical impulse response of the underwater optical channel. As scattering increases, the point spread function broadens, and the probability increases that multiple reflections could occur. Additionally, as light propagates, the temporal dispersion can invalidate the representation of a target reflection as a delta function in (2.1). This dispersion could potentially break the sparse approximation. To better tailor the sparse approximation in (2.1), the impulse response of the underwater channel $\rho(t)$ and the attenuation coefficient c are taken into account (2.2):

$$s(t) = \sum_{i=1}^K \alpha_i e^{-c\nu\tau_i} \rho(t - \tau_i) \quad (2.2)$$

In this representation, ν is the speed of light in water, and the optical impulse response function $\rho(t)$ takes into account the pulse spreading caused by scatter as well as the time delay to the target. The goal of any ranging system is to recover the amplitude α and time delay τ for any and all reflecting objects. In the following sections, the methods for recovering this information will be discussed.

2.2 Sensing Matrix Formulation

Given a known periodic transmitted waveform $h(t)$ with period T , the sparse reflectivity scene can be convolved with the transmitted waveform to model the signal at the receiver, $x(t)$. This relationship is described by equation (2.3). The signal at the receiver is then sampled at a rate $F_D = F_N/D$ and is represented by $y[m]$ as the received signal post-digitization (2.4).

$$x(t) = (h * s)(t) = \int h(t - \tau)s(\tau)\delta\tau \quad (2.3)$$

$$y(m) = \sum_{n=1}^N h(mD - n)s(n) \quad (2.4)$$

In this formulation, F_N is the traditional Nyquist sampling rate, N is the number of total Nyquist samples given by $N = \lfloor F_N \times T \rfloor$, $M = \lfloor F_D \times T \rfloor$ is the number of measurements, and $D = N/M$ is a parameter that controls the system's sampling rate. It is important to note that M and N are strictly integer which causes D to be contained within the set of rational numbers, and that within the realm of sub-Nyquist sensing, $D > 1$ and is inversely proportional to the downsampling parameter $\delta = 1/D$. Given that the signal is periodic, the convolution operator in 2.4 can be represented by a matrix multiplication where the kernel of the integral

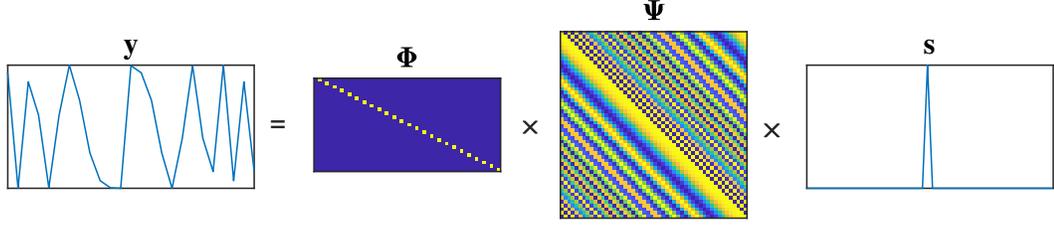


Figure 2.1: Pictorial representation of the sub-Nyquist sensing system. The sparse vector s has a non-zero elements at the targets location and multiplies the waveform matrix Ψ which has columns that correspond to all possible time shifts. The time shifted signal is then multiplied by the sampling matrix Φ which generates the sub-Nyquist rate measurements, y , of the time shifted waveform.

transform is the $N \times N$ circulant matrix Ψ . The columns of this matrix are circularly shifted versions of the transmit waveform stored at the Nyquist rate. The sampling operation of the low rate A/D can be modeled by a matrix multiplication where the $M \times N$ sampling function is the linear interpolation matrix Φ that maps the transmitted waveform's Nyquist rate to the system's sampling rate. The two operations are combined to create the $M \times N$ Sub-Nyquist sensing matrix Θ shown in equation (2.5) and demonstrated pictorially in figure 2.1.

$$y = \Phi\Psi s = \Theta s \quad (2.5)$$

To summarize the above, the hybrid lidar-radar ranging system was reformulated as a linear system of equations that relies on the transmitted waveform to spread the sparse target vectors information across the observed time domain. A uniform sampling process was then introduced that is linear and non-adaptive. In this formulation, the sparse vector s is the unknown impulse response of the illuminated scene and has non zero elements at the time bin locations related to the targets range.

The vector y is the output of the low rate digitizer, and can be referred to as the measurements. The question remains though if it is possible to recover the sparse vector s , and thus the range and reflectivity of the target, from the sub-Nyquist measurements y .

2.3 Recovery of Sparse Vectors

For the system described, detection and recovery methods traditionally use a matched filter which is the optimal linear filter for maximizing the SNR in the presence of noise [34], but it is only optimal in the sense of minimizing the L_2 norm of error. This method requires the signal of interest be sampled at or above the Nyquist rate where $M \geq N$, and is not applicable to under-determined systems. In addition, the matched filter suffers from limited resolution due to the ambiguity function of the waveform of choice. Given the sensing and sampling methods described previously, the sub-Nyquist measurement matrix is under-determined, and basic linear algebra principles state that there are infinitely many vectors s that could give the vector y . Methods to recover sparse vectors originate from optimization problems that are subject to additional linear equality constraints which converge to a finite number of solutions even when the system is under-determined. These methods are based on the minimization of cost functions that are derived from norms other than the L_2 norm.

2.3.1 L-1 Minimization

To address the issues due to the under-determined nature of the sensing matrix, theories of over-complete signal representations and sparse recovery techniques can be applied to the sub-Nyquist problem [35]. These alternative minimization methods exploit the sparsity of the signal to be recovered. Sparsity can be measured by the l_0 norm which measures the number of non-zero entries in a vector. While it is not proper norm, it has been used extensively within the fields of functional analysis, probability theory, and harmonic analysis [36]. Ideally one would hope to solve the primal problem by optimizing over the l_0 norm to obtain the minimum number of non-zero elements. This optimization problem is give in equation (2.6).

$$\min_s \|s\|_0 \text{ subject to } \Theta s = y \quad (2.6)$$

Unfortunately, this method has been proven to be NP-hard [37] and can not be used to effectively and efficiently recover sparse vectors. Basis pursuit (BP) [38] is a method that aims to solve an alternate optimization problem. This method utilizes a different cost function, where the optimal solution is characterized by the minimum l_1 norm (2.7).

$$\min_s \|s\|_1 \text{ subject to } \Theta s = y \quad (2.7)$$

This relaxation is advantageous because it leads to a problem that is no longer NP-hard, and it has been shown [39], that the minimal l_1 solution is a valid approximate of the sparsest solution. Even further, under certain conditions to be discussed, the solution to the minimal l_1 norm problem is exactly the solution to the minimal l_0

norm problem.

The downside to L_1 minimization techniques is that the most common way to solve this problem is through the use of interior point methods which are computationally complex. While L_1 techniques are ideal, and have been proven to exactly recover a sparse vector with a minimal amount of measurements, there exist other algorithms for recovery as well. These alternative methods are not optimal, meaning that they require more measurements than L_1 minimization. This fact prompts the classification of these methods as "greedy" algorithms [40]. While they might require more measurements, they are less computationally expensive, and many have fast methods for implementation [41]. In a later section, these algorithms will be discussed in more detail. Regardless of algorithm, all of these methods require that certain constraints be met in order to achieve successful recovery of the sparse vector. In the following section, these necessary conditions for recovery will be discussed.

2.3.2 Necessary Conditions for Recovery

The previous section alluded to conditions on the measurement matrix that must be met in order to successfully recover a sparse vector from the described under-determined system of equations. There has been substantial work completed in the compressive sensing community in an attempt to theoretically prove the minimum number of measurements necessary for exact recovery. The most famous of which is the Restricted Isometry Property (RIP) which guarantees that the solution to BP and other sparse recovery methods are the exact solution [40, 42, 43]. The RIP

(2.8) states that given the sub-matrix θ_T with $T \subset (1, \dots, N)$ generated through the extraction of columns of Θ corresponding to the indexes of T , there is a K -restricted isometry constant δ_K which is the smallest quantity that satisfies (2.8) for all possible subsets $T \leq K$ and all possible sparse vectors.

$$(1 - \delta_K) \|s\|_2^2 \leq \theta_T s \|_2^2 \leq (1 + \delta_K) \|s\|_2^2 \quad (2.8)$$

The RIP implies that to recover a sparse vector with support K , all possible K combinations of the measurement matrix columns must be approximately orthonormal. While this is a very tight bound on the recovery of sparse vectors, it is also something that is theoretically difficult to prove for arbitrary or deterministic measurement matrices, and it is computationally difficult to empirically measure. To prove RIP for an arbitrary matrix one would need to investigate all $\binom{N}{K}$ possible sub-matrices which is combinatorially difficult. It is possible to loosen the bound by looking at the coherence of the measurement matrix. The coherence μ is the measure of the maximum absolute correlation between different column elements of the measurement matrix given by (2.9), and is a fundamental characteristic associated with compressive sensing.

$$\mu = \max_{i \neq j} \frac{|\langle \theta_i, \theta_j \rangle|}{\|\theta_i\|_2 \|\theta_j\|_2} \quad (2.9)$$

Given the coherence of a measurement matrix it has been proven that both BP and similar greedy algorithms will recover the sparse vector if the condition (2.10) is true [40], and this mutual coherence property is a sufficient condition for exact recovery.

$$K < \frac{1}{2}(\mu^{-1} + 1) \quad (2.10)$$

To extend the coherence property as applied to hybrid lidar-radar system, it can be stated that if the coherence of the measurement matrix is strictly less than one, it is possible to recover a sparse vector with a single non-zero element.

In compressive sensing literature, there are many other sparse recovery conditions that have been presented. For example the Null-Space property [44] was introduced and provides a sufficient condition for satisfying the RIP. Bounds on the spark and eigenvalues of sensing matrices have also been used theoretically but verifying that any general measurement matrix satisfies these condition has a computational complexity that is combinatorial [45]. Given the overwhelming popularity as well as the ease of calculation, this research effort will rely on the coherence to make predictions of sparse recovery characteristics of the derived sensing matrix. It is acknowledge that the bound on coherence is only a sufficient condition, and is one of the more pessimistic bounds which under-predicts sparse recovery performance.

2.4 Discussion

In the following sections, the coherence of sensing matrices derived from wide band modulated waveforms will be explored, and the performance of sparse recovery algorithms will be measured. Additionally, empirical evidence will be provided through simulations and experimental tests that these methods can be used to measure the range to an object. It is important to first state how the above sub-Nyquist formulation differs from previously-proposed compressive sensing and sampling methods.

The original work that showed it is possible to reduce the number of measurements using compressive sensing theory was introduced in the seminal papers by E. Cands, J. Romberg, T. Tao, and D. Donoho [30, 42, 46, 47]. They provided theoretical proofs for recovery that relied on random projections of the sparse vector and it was shown that these random projections were optimal. Given that fully random measurements are often unrealizable in practice, other sampling methods have been proposed such as Non-uniform Multicoset Sampling [48], X-sampling [49], and measurements through random convolutions [50]. These methods have all had substantial success within the fields of compressive sensing and address a wide array of applications. Most proposed methods involve complicated analog hardware, non-uniform sampling methods, or require multiple measurement vectors to obtain the necessary number of projections for successful recovery.

The research presented in this effort differs in that only a specific application is being targeted, and not the recovery of broad classes of signals. In turn, it is hypothesized that by using long duration wide-band signals to modulate laser sources, it is possible to recover the range and reflectivity of illuminated objects by directly sampling the RF signal using a single low rate A/D converter and utilizing sparse recovery algorithms. It is also interesting to highlight that while laser ranging systems are being targeted, these results could also potentially be applied to many different types of systems that use active sensors and wide-band modulation.

Chapter 3: Waveform Design Consideration for Sub-Nyquist Sensing

The original research and theory of compressive sensing typically relied on the use of a random measurement matrix to sample the sparse or compressible vector. To prove that exact recovery conditions such as the RIP can be met, sensing matrices were derived from iid Gaussian, Bernoulli, or other random distributions [42]. While these types of measurement matrices are useful in proofs and can be applied to certain applications like the single pixel camera, [51], they are not always realizable due to causality constraints.

For the hybrid lidar-radar system, causality must be maintained in the measurement process, and random sensing matrices are inherently non-causal. To make sub-Nyquist measurements while maintaining causality, the measurement process previously described utilized uniform under-sampling of a waveform that had been convolved with a sparse vector. This measurement process leads to a sensing matrix that is inherently deterministic. While there are no proofs of RIP satisfaction for deterministic sensing matrices, there has been substantial work which details how to empirically prove that a given sensing matrix satisfies the RIP [52]. One such method is to measure the coherence of a the sensing matrix and use the Mutual Coherence Property (2.10) to make prediction on the recoverable support through

L_1 minimization methods. In this section, the relationship between waveform selection, sampling rate, and sensing matrix coherence will be investigated, and an initial estimate on expected recoverable support will be reported.

3.1 Coherence Investigation

In the previous section, the measurement matrix for an arbitrary waveform was derived, and sufficient conditions for recovery of sparse vectors from sub-Nyquist measurements were presented. This section will investigate how waveform choice can affect the coherence of the measurement system and will provide the initial evidence that it is indeed possible to recover sparse vectors from sub-Nyquist sampled versions of the modulated waveforms. The motivation for waveform selection stems from the mutual coherence property and its relation to the correlation function [53]. When revisiting the equation for coherence, it can be observed that the inner product of two columns of the measurement matrix is the inner product between two time shifted waveforms. The coherence in this application can then be defined as the max correlation between two waveforms with different time shifts. Wide-band waveforms have been used in many application specifically because of their low correlation between time shifted versions of the original signal, which leads to large peak to side-lobe ratios. This property motivates the selection of two specific types of wide-band waveforms to be studied: a Pseudo-Noise (PN) Code, and a linearly frequency modulated chirp.

For this investigation, sensing matrices will be generated based on a given

waveform which has a fixed bandwidth dictated by $\Delta f = \frac{F_N}{2}$, where F_N is the Nyquist rate for the transmitted waveform. The number of columns of the sensing matrix, N , is proportional to the Nyquist rate and the duration of the waveform, T , shown in equation (3.1). The variable N is adjusted by changing the duration of the given waveform. The number of measurements, M , is equal to the number of rows of the sensing matrix and is given by equation (3.2). The number of measurements is controlled by altering the sampling rate F_D .

$$N = \lfloor F_N \times T \rfloor \quad (3.1)$$

$$M = \lfloor F_D \times T \rfloor \quad (3.2)$$

In the following section, the sensing matrix coherence μ will be measured as a function of the waveform's discrete duration, N and down-sampling ratio, $\delta = \frac{M}{N}$.

3.1.1 Pseudo-Noise Code

PN codes repeat a deterministic yet noise like sequence of bits, ± 1 , at a given interval, and can be generated in numerous ways [54]. They have been used in many applications such as radar and communications systems due to the low correlation between different time shifted returns [55]. PN codes can be implemented in a hybrid lidar-radar system by modulating the laser on and off in the same pattern as the given sequence and then DC coupling the received signal to obtain the ± 1 pattern. A typical correlation function for a PN code is shown in figure 3.1. The correlation function shows that when given the fully sampled sequence, the non-zero sample shifts are not well correlated. This implies that sensing matrices derived from these

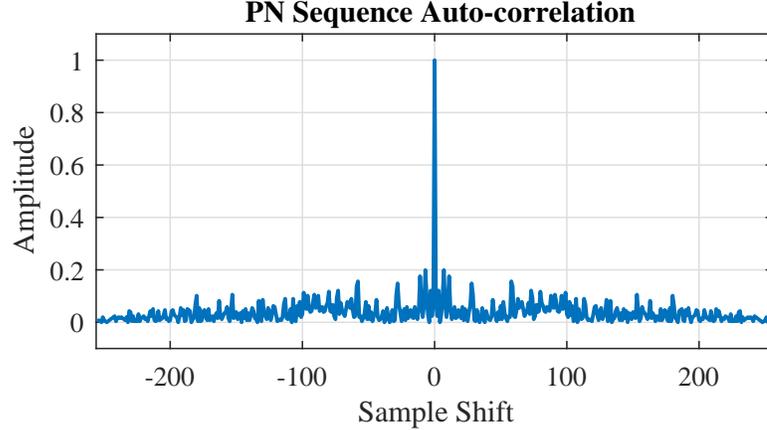


Figure 3.1: Auto-Correlation Function for a PN-Coded sequence

waveforms should have low coherence, but it remains to be shown that this is also the case for sub-Nyquist sampled versions of the PN coded waveforms.

Using the methods described by equations (2.3 - 2.5), a sensing matrix can be generated given the PN coded sequence. The number of measurements m are altered for various levels of N , and the coherence is measured given (2.9) over multiple independent trials. The number of measurements is adjusted by changing the system sampling rate, and the number of Nyquist samples is altered by changing the total number of bits in the PN sequence. Figure 3.2(a), shows the measured coherence as a function of the down-sampling parameter $\delta = m/N$, where the number of measurements was varied for three different values of N . As the down-sampling parameter approaches the Nyquist sampled case, $\delta = 1$, the coherence values are strictly less than one and decrease monotonically. It is important to note that the coherence also has a dependence on the total number of columns of the sensing matrix, which is related to the total number of bits in the original sequence. This

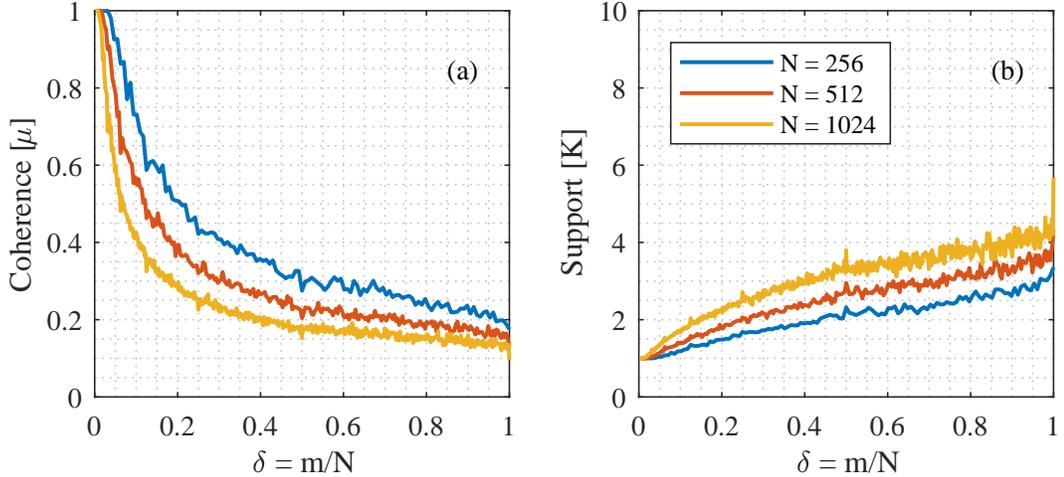


Figure 3.2: Empirically measured coherence values for the sub-Nyquist sensing matrices derived from PN-coded sequences. (a) shows the coherence as a function of the down-sampling parameter $\delta = m/N$, (b) is the estimated max recoverable support for a given sensing matrix.

can be described by the fact that for binary noise like sequences, the peak to side-lobe ratio increases as the number of bits increase [56]. These initial coherence measurements are promising because $\mu < 1$ for a majority of the generated sensing matrices implying that at least a single target can be recovered for the sub-Nyquist measurements of the transmitted waveform.

While these results imply that recovery of a single sparse vector is possible, they also imply that sparse vectors with support greater than one can also be recovered exactly. Figure 3.2(b) shows the prediction of the maximum recoverable support for the given sensing matrix calculated using equation (2.9) While this estimate ensures that the solution to the L1 minimization problem is the unique minimizer, it is important to note that it is only a sufficient condition and provides a loose bound on the recoverable support.

3.1.2 Frequency Modulated Continuous Wave

Frequency modulated continuous wave (FMCW) signals are generated by linearly sweeping through a range of frequencies for a given duration. Just like the PN coded sequences, FMCW modulation techniques have been used extensively within the field of radar, and there has a well known correlation function. The correlation function of a FMCW waveform is dictated by two main parameters: the bandwidth Δf and the duration T . The resolution is inversely proportional to the waveform bandwidth, and the time bandwidth product dictates the pulse compression gain. Figure 3.3 shows an example of the auto-correlation function for a FMCW signal. When compared to the auto-correlation of the PN-coded sequence with the same bandwidth and duration, figure 3.1, there is a clear benefit in terms of the reduced amplitude at the non-zero time shift positions. This implies that the coherence of sensing matrices derived from the FMCW signals might be lower than those derived from PN coded sequences.

Similar to the previous section, simulations were conducted to measure the coherence by keeping the bandwidth fixed, and changing the system sampling rate and waveform duration to alter m and N , respectively. Figure 3.4(a) shows the empirically measured values of coherence for the sensing matrices derived from the FMCW signal as a function of down-sampling parameter, $\delta = m/N$, given three values of N . Comparing the coherence values measured when investigating the PN coded sequence, there are three main differences. The first difference is that the coherence of sensing matrices derived from FMCW signals is largely independent of

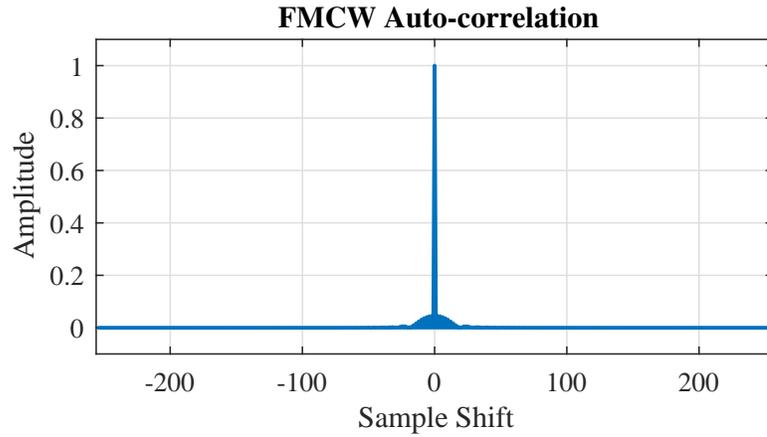


Figure 3.3: Autocorrelation function of a frequency modulated continuous wave signal.

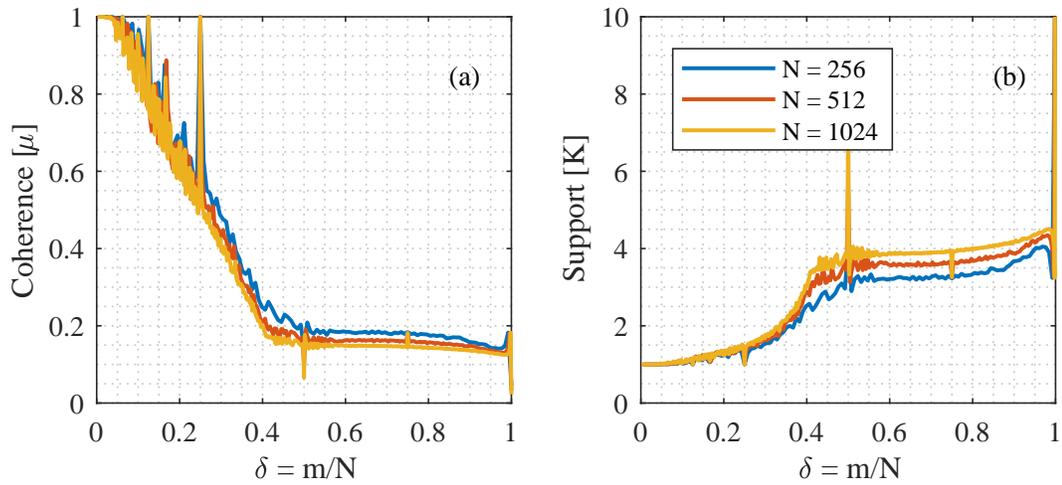


Figure 3.4: Empirically measured coherence values for the sub-Nyquist sensing matrices derived from FMCW signals. (a) shows the coherence as a function of the down-sampling parameter $\delta = m/N$, (b) is the estimated max recoverable support for a given sensing matrix.

the waveform duration. This stems from the fact that the peak to side-lobe ratio for the FMCW auto-correlation function is independent of the signal duration. This result implies that for a given bandwidth, the recoverable support characteristics should not change if the duration is altered. The second observed difference when comparing figure 3.2 and figure 3.4 is the relationship between the coherence and the down-sampling parameter. While the coherence for matrices derived from PN coded sequences had an approximately logarithmic dependence on δ , the coherence relationship for the FMCW-derived matrices can be described by a piecewise linear function. The coherence appears to linearly decrease up until $\delta = \frac{1}{2}$ and then remains relatively constant. Further analysis also shows that for certain parameters, the coherence for FMCW matrices is less than that of PN coded sequences when $\delta \geq .35$. These differences are highlighted in figure 3.5, where the coherence results for the two different types of sensing matrices with $N = 512$ are displayed on the same chart.

The third and most important difference is that the coherence is no longer strictly less than one. When $\mu = 1$, it is implied that there are two columns within the sensing matrix that are identical. This is problematic during sparse recovery because it means that there are two possible solutions to the sparse recovery problem. In figure 3.4, this case happens specifically when $\delta = \frac{1}{4}, \frac{1}{8}$. These discontinuities were investigated and it was shown that they are a side effect of the symmetric aliasing of the FMCW signal's spectral content [53]. When sub-Nyquist sampling a linear FMCW signal, the aliased frequencies are mirrored across the system sampling rate, $F_D/2$. When the down-sampling parameter is inversely proportional to a power of

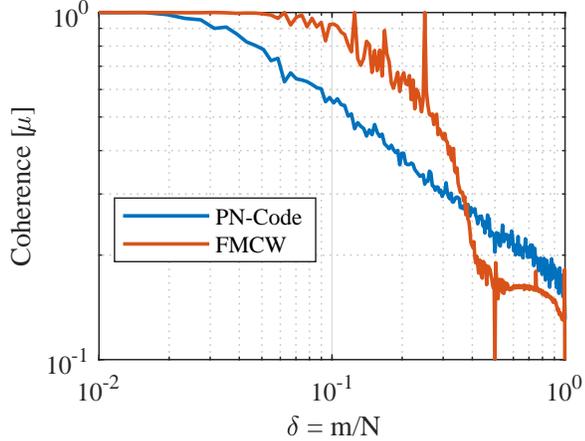


Figure 3.5: Comparison of coherence results for sensing matrices derived from PN-coded and FMCW waveforms, $N = 512$. Shown on a log-log plot to highlight the PN-Coded waveforms logarithmic dependence on δ

two, multiple columns of the sensing matrix, which correspond to different time shifts, can be identical. This leads to multiple possible solutions that can minimize the L_1 norm. This effect can be mitigated, and the coherence can be minimized by selecting sensing matrix generation parameters M and N to be relatively prime. By implementing this restriction, each Nyquist rate sample-shifted column ψ_n of the waveform matrix Ψ with $n \in 1, 2, \dots, N$, is sampled at a slightly different time by the sensing matrix Φ given that the least common multiple between M and N is their product. Figure 3.6(a) shows the mutual coherence as a function of the down-sampling parameter, and 3.6(b) shows the predicted max recoverable sparse vector support when using L_1 minimization and the tested sensing matrix derived from FMCW signals. The results show that by selecting parameters that require M and N to be relatively prime the coherence is strictly less than one, and that

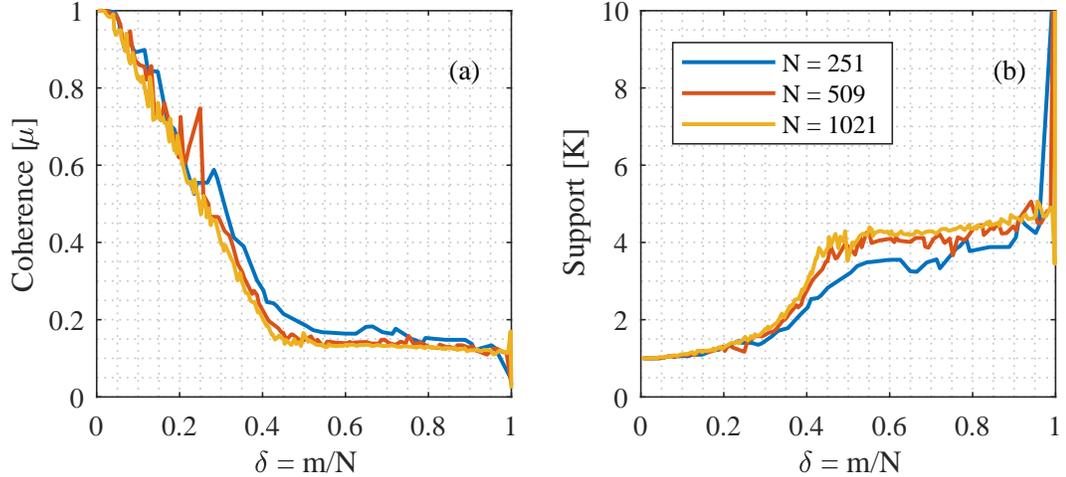


Figure 3.6: Empirically measured coherence values for the sub-Nyquist sensing matrices derived from FMCW signals with relatively prime m and N . (a) shows the coherence as a function of the down-sampling parameter δ , (b) is the estimated max recoverable support for a given sensing matrix.

recovery of sparse vectors with unit support should be possible.

3.2 Discussion

In this section, two different wide-band waveforms were explored and the initial bounds of sparse recovery performance were presented. Sensing matrices derived from PN coded sequences benefited from the non-symmetric and random nature of their modulation, and the coherence was dependent on the number of Nyquist rate samples N . When compared to the coherence of sensing matrices derived from FMCW signals, they suffered from overall higher coherence at down-sampling ratios $\delta \geq .35$. Additionally, sensing matrices that were based on FMCW signals were prone to aliasing due to the symmetry of their time-frequency content. Initially, it

seemed that there would be conditions where finding a unique solution to the L_1 minimization problem would not be possible. However, by using relatively prime numbers for the number of measurements M and number of Nyquist samples N , these adverse effects could be avoided. It is interesting to note that the use of relatively prime numbers is also useful when applying compressed sensing theory to wide-band spectrum sensing [57], and sub-Nyquist beam-forming [58].

In the following sections, the recovery of sparse vectors will be simulated with both noiseless and noisy observations. The recoverable support will also be investigated, and tighter bounds on support recovery will be explored. Although the coherence for sensing matrices derived from PN-Coded waveforms was lower when $\delta \leq .35$ only sensing matrices derived from FMCW waveforms will be further studied. This choice has been made for three main reasons. The first reason is the independent relationship between the coherence and the duration of the FMCW waveform. This is desirable because it implies that the sparse support recovery performance should be similar for short or long duration waveforms. The second reason is due to the deterministic nature of sensing matrices derived from FMCW waveforms. The deterministic nature leads to low storage requirements for the sensing matrix because it can be generated knowing only three parameters: the duration, bandwidth, and system sampling rate. The third reason is because FMCW signals allow for the transmission of pass-band waveforms more easily than the base-band binary PN-code which would require additional carrier modulation and demodulation steps. This is advantageous for underwater ranging in turbid environments to reduce the effects of the backscatter by transmitting outside of the backscatter's

frequency response region.

Chapter 4: Sparse Recovery Algorithms

Solving the under determined system of equations derived from compressive or sub-Nyquist sensing systems is a problem that cannot be solved using traditional methods of matrix inversion or least square estimation. As mentioned previously, the notion of sparsity is a powerful constraint and can be used to guarantee the recovery of sparse vectors. While the $\|s\|_0$ cost function is NP-hard and cannot be minimized directly; relaxing the cost function to $\|s\|_1$ has been proven to be acceptable provided RIP is satisfied [40]. While the sparse recovery problem can be reformulated in many ways, there are three main optimization problems that are targeted in this effort. The first is L_1 minimization with equality constraints shown in equation (4.1), and is commonly referred to as basis pursuit.

$$\min_s \|s\|_1 \text{ subject to } \Theta s = y \quad (4.1)$$

However, when noise is present in the measurement process, shown in equation (4.2), it would be impossible to obtain the equality constraint. This realization requires the use of different constraints.

$$y = \Theta s + n \quad (4.2)$$

There are two main ways to address this problem of measurement noise. The first involves minimizing the L_1 cost function while utilizing quadratic constraints on the

residual error. This formulation is shown in equation (4.3) where $\epsilon > 0$ and can be related to the power of the measurement noise.

$$\min_s \|s\|_1 \text{ subject to } \|\Theta s - y\|_2^2 \leq \epsilon \quad (4.3)$$

$$\min_s \frac{1}{2} \|\Theta s - y\|_2^2 + \lambda \|s\|_1 \quad (4.4)$$

The second method, equation (4.4), introduces a regularization parameter on the sparsity of the final solution, and is commonly referred to basis pursuit denoising. In this formulation, the variable λ is the regularization parameter and controls the trade off between sparsity and reconstruction accuracy in the L_2 sense. While the addition of noise will be discussed later in the text, the question remains: how does one actually solve the minimization problem to recover the sparse vector?

4.1 Convex Optimization

Generally, optimization problems are very difficult to solve and most methods involve compromises between computation time, stability, and convergence. The minimization problems presented previously can be recast as linear programs or second order cone programs, and are convex optimization problems. While there are no analytical solutions to convex optimization problems, there exist many reliable algorithms to solve them. Methods such as the primal-dual interior-point methods and log-barrier algorithms [59, 60] have been used traditionally within the fields of convex optimization, and the application of these methods to compressive sensing and sparse recovery has been extensively explored. To test the application of convex optimization to the sub-Nyquist hybrid lidar-radar system, the L_1 -Magic software

package [2] was used. This software package can solve many different types of L_1 minimization problems, but in this research effort, the Basis Pursuit and Basis Pursuit Denoising optimization problems were studied. The software package utilizes a primal dual interior point method to solve the system of equations generated when the basis pursuit problem is recast as a linear program. Alternatively for the Basis Pursuit Denoising problem, the quadratically constrained L_1 minimization problem is recast as a second-order cone program and solved via a log barrier algorithm.

While these methods are quite robust and require no a priori knowledge of the sparse vector to be recovered, they are computationally expensive and do not scale well for larger problems. The computation complexity of convex optimization problems are roughly proportional to $\mathcal{O}(N^2M)$, where N is the number of variables and M is the number of constraints [61]. When applied to our hybrid lidar-radar system, N is proportional to the waveform duration and the Nyquist rate and M is the number of sub-Nyquist samples. While other algorithms exist that are faster, there is a trade-off to be made between speed and the total number of constraints/measurements needed. In the following sections, these algorithms will be introduced, and results from simulations that were completed to test and compare the different recovery methods will be presented.

4.2 Greedy Algorithms

Greedy methods are generally defined as algorithms that attempt to solve for a globally optimal solution by finding the locally optimal solution at each iteration [62].

While greedy algorithms do not generally guarantee the globally optimal solution, they are significantly less computationally complex and typically lend themselves to fast implementations. Despite the iterative local optimization approach, it has been shown that when applied to sparse recovery, it is provable that these methods will recover the optimally sparse solution when the mutual coherence condition is met [40, 63]. Greedy algorithms are a broad class of algorithms, and there are many different types that have been created in various fields to solve different types of optimization problems. Specifically, in this section, the Orthogonal Matching Pursuit (OMP), and Hard Thresholding Pursuit (HTP) will be presented.

4.2.1 Orthogonal Matching Pursuit

OMP is a recursive algorithm that was originally designed and applied to wavelet decomposition [64], and it was an alteration of the original Matching Pursuit (MP) algorithm [65]. It has been shown that when compared to MP, OMP converges to a solution much faster than MP. It was also shown that OMP can exactly solve the optimization problem as long as the Mutual Coherence condition (2.10) is satisfied [40]. OMP belongs to a broader class of algorithms called residual correlation based algorithms. These types of algorithms use the correlation between the measurement matrix and the residual estimate error to update the estimate support set. The OMP algorithm is shown in Algorithm 1. This algorithm takes the measurement matrix Θ , the measurements y , an estimate of the sparse vector's support K_{est} , and a parameter ϵ as inputs. The outputs of the algorithm are the estimated sparse

vector \hat{s} , and the estimated support indexes Λ . The algorithm as presented here utilizes both the estimated sparsity as well as a condition that stops the algorithm if the residual error drops below the threshold parameter ϵ .

Algorithm 1 Orthogonal Matching Pursuit

Inputs: $\Theta, y, K_{est}, \epsilon$

Outputs: \hat{s}, Λ

Initialize: $r_0 = y, \Lambda_0 = \emptyset$

Iterate: $j = 1, 2, \dots, K_{est}$

(1) Choose $\lambda_j = \arg \max_i |\langle r_{j-1}, \theta_i \rangle|$ subject to $i \in (1, \dots, N)$

(2) Update $\Lambda_j = \Lambda_{j-1} \cup \lambda_j$

(3) Update $\hat{s}_j = \arg \min_s \|y - \Theta s\|_2^2$ subject to $s_i = 0 \quad \forall \quad i \notin \Lambda_j$

(4) Update $r_j = y - \Theta \hat{s}_j$

(5) Stop if $\|r_j\|_2 \leq \epsilon$ or $\|r_j\|_2 > \|r_{j-1}\|_2$

The OMP algorithm is iterative and only completes up to K_{est} iterations. At each step, the algorithm identifies the column of the measurement matrix λ_j that is maximally correlated with the residual r_j and adds that column to the support set Λ . It then solves a linear least squares error estimation to approximate \hat{s} but limits the minimizer to those elements within the support set Λ . Lastly, the residual is updated, and the stopping condition is checked. The stopping criteria used here is based on the idea of a minimal residual error condition, as well as a bounded residual. The residual is bounded from growing which implies error should constantly decrease, and if measurement noise is present, ϵ should be proportional to the noise power.

4.2.2 Hard Thresholding Pursuit

HTP is a method for recovering sparse vectors which is based on an iterative thresholding approach [66]. These types of algorithms are typically much more efficient than residual correlation based algorithms because the total number of matrix inversions is reduced. This efficiency comes at a cost because such algorithms tend to require more measurements in order to converge to the optimally sparse solution [67]. The HTP algorithm is shown in Algorithm 2. The variables are similar to the OMP algorithm with the addition of J , the maximum number of iterations. It is important to note that both greedy algorithms require an estimate of the sparsity, and this feature will be discussed further in later sections.

Algorithm 2 Hard Thresholding Pursuit

Inputs: $\Theta, y, K_{est}, \epsilon, J$

Outputs: \hat{s}, Λ

Initialize: $r_0 = y, \hat{s}_0 = 0, \Lambda_0 = H_{K_{est}}(\Theta^*y)$

Iterate: $j = 1, 2, \dots, J$

- (1) Calculate: $\mu = \frac{\|(\Theta^*(y - \Theta\hat{s}_{j-1}))_{\Lambda_{j-1}}\|_2^2}{\|\Theta((\Theta^*(y - \Theta\hat{s}_{j-1}))_{\Lambda_{j-1}})\|_2^2}$
 - (2) Threshold: $\Lambda_j = H_{K_{est}}(\hat{s}_{j-1} + \mu\Theta^*r_{j-1})$
 - (3) Update: $\hat{s}_j = \arg \min_s \|y - \Theta s\|_2^2$ subject to $\text{supp}(s) \subseteq \Lambda_j$
 - (4) Update: $r_j = y - (\Theta\hat{s}_j)_{\Lambda_j}$
 - (5) Stop if $\|r_j\|_2 \leq \epsilon$ or $\|r_j\|_2 > \|r_{j-1}\|_2$
-

One of the main difference between the OMP and HTP algorithms is that HTP updates the sparse estimate across the entire index set at every iteration whereas OMP only selects one sparse element at a time. The subset to estimate over is calculated using a hard threshold defined by $H_S(f(x))$, which is a non-linear operator that returns the indexes of the S largest entries of $f(x)$. This feature can potentially decrease the convergence time, but it can lead to errors when the

coherence of the sensing matrix is high. In this implementation of HTP, the same stopping condition is included which ends the algorithm once the error has been reduced past a certain threshold ϵ , or if the residual power begins to increase.

4.3 Analysis of Recovery of Sparse Vectors

To determine the performance of the aforementioned sparse recovery algorithms, simulations were designed to compare their performance. These simulations will help predict how sub-Nyquist sensing and sparse recovery can be expected to perform when applied to the hybrid lidar-radar system. The performance metrics that will be tracked include the computation time, δt , Normalized Square Error, NSE, between the actual sparse signal and the estimate, and the Probability of Exact Recovery P_{ER} . How these performance metrics are calculated will be explained in the following sections along with a description of the simulation design.

4.3.1 Simulation Design

The goal of these simulations is to track the performance of the sparse recovery algorithms for a noiseless measurement system. While this is not realizable in practice given that noise will always be present, it is an important step to empirically determine the conditions that must be satisfied in order to exactly recover sparse vectors. There are four main parameters that must be studied in order to exhaustively compare the different sparse recovery algorithms and measure the exact recovery conditions. The first parameter is the number of columns in the sensing

matrix, N , which is proportional to the waveform duration and the Nyquist rate. The second parameter is the down-sampling ratio $\delta = m/N$, and is controlled by changing the desired number of measurements, m . The third parameter is the support of the sparse vector, K , which equal to the total number of non-zero elements. By adjusting these three parameters, and by conducting multiple independent trials, the limits of the sparse recovery algorithms can be tested.

4.3.1.1 Performance Metrics

To measure performance, three key metrics will be tracked for comparison. The first performance metric is the normalized square error between the known sparse vector s , and the recovered sparse estimate \hat{s} . This metric is calculated by equation (4.5) and is calculated for each random trial t .

$$\text{NSE}(t) = \frac{\|s - \hat{s}\|_2^2}{\|s\|_2^2} \quad (4.5)$$

This metric provides an indicator for the "correctness" of the sparse estimate. Using the NSE measurement, the second performance metric, P_{ER} , can be calculated. By using multiple independent trials T , this metric can provide a measure of the likelihood of exact recovery given certain simulation parameters. This metric is first calculated by checking the condition (4.6). If this condition is true, a success is reported for trial t , and the Exact Recovery Probability can be calculated by equation (4.7).

$$\text{NSE}(t) \leq \epsilon_s \quad (4.6)$$

$$P_{ER} = \frac{\text{Total Number of Success}}{\text{Total Number of Trials}} \quad (4.7)$$

The choice of ϵ_s is a tolerance parameter for the acceptable correctness. This parameter is used often in literature, and the typical tolerance has been studied [68] and is set to 10^{-3} . The probability of exact recovery can then be used to find the maximum recoverable support K_{max} for a given down-sampling parameter $\delta = m/N$.

4.3.2 Simulation Results

To test the sparse recovery performance as a function of the previously mentioned metrics, Monte Carlo simulations were run to generate an understanding of the necessary conditions for recovery. These simulations involved parameter sweeps and multiple trials to measure the average NSE, P_{ER} , K_{max} , and computation time. The simulation steps are listed below and are the same for all of the algorithms.

Simulation Steps

- (1) Select recovery algorithm
 - (2) Set N, M and K values
 - (3) Generate sensing matrix, Θ
 - (4) For each independent trial
 - (4.1) Generate the sparse vector s
 - (4.2) Generate the measurements $y = \Theta s$
 - (4.3) Implement recovery algorithm
 - (4.4) Record Performance Metrics
 - (5) Calculate Average Performance Metrics over all Trials
-

The simulation was tailored to accurately represent the hybrid lidar-radar ranging system by making one important design choice. This choice restricts the amplitude values of the sparse vector during step 4.1 to be strictly positive. This restriction was based on the fact that there is no physical explanation for a negative reflection coefficient, and this implies that the amplitude values of s must be strictly positive. During these initial simulations the parameter N is fixed and the parame-

ters M and K are swept from 1 to N in a linear fashion. The following sections will present the results of these simulations.

4.3.2.1 Results for Convex Optimization Algorithm

When attempting to solve the Basis pursuit problem, Convex Optimization algorithms require only the measurement matrix and the measurements as input parameters. The algorithm then attempts to solve the minimization problem and only exits when either the error gap between the primal dual problem has been sufficiently reduced, or the algorithm iterates a maximum number of times. The first parameter to be investigated is the average NSE as a function of the changing parameters K and $\delta = m/N$. To visualize the joint effect that the two independent parameters have on the metric, the average NSE is mapped to a color value and the metric is displayed in figure 4.1 as a two dimensional image.

When analyzing the NSE results, it is noted that there is a clear transition region between high error and low error. The transition is approximately at $\text{NSE} = 10^{-3}$, and this type of feature is expected for sensing matrices that satisfy the exact recovery conditions dictated by the RIP. As mentioned before, a threshold on the NSE for the independent trials can be used to generate a measure of the probability of exact recovery. Similarly, the measured probability can be mapped to a color value for each K and $\delta = m/N$, and the results are shown in figure 4.2. A $P_{ER}(\delta, k) = 1$ implies that every non-zero element of the sparse vector was found in the right location with the right amplitude. This also implies that given the sensing matrix

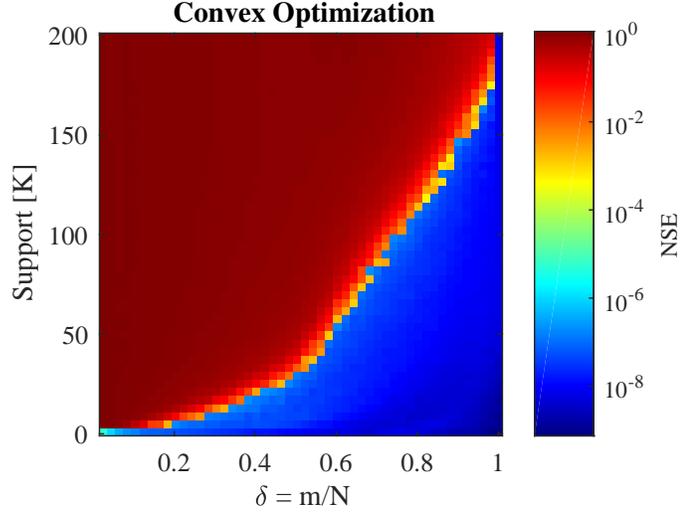


Figure 4.1: Average NSE value of the convex optimization algorithm L_1 -magic [2] mapped to a color value and displayed as a function of the vector support K and the down-sampling parameter $\delta = m/N$. Simulation has $N = 199$.

designed by the parameters δ , and N , one would expect to recover a vector with support k through L1 minimization. In the hybrid lidar-radar sense, this provides a limit for the maximum number of reflecting objects that can be detected. The probability of exact recovery results can be used to empirically measure the bound on the maximum recoverable sparsity as a function of the down-sampling parameter. This bound will provide an upper limit on the total number of detectable targets and is achieved by finding the points where $P_{ER}(\delta, k) > .99$. Additional simulations were run for various values of N in order to investigate the effect that waveform duration has on exact recovery when using convex optimization. The simulation results are shown in figure 4.3.

These results provide the initial empirical proof that exact recovery of a sparse

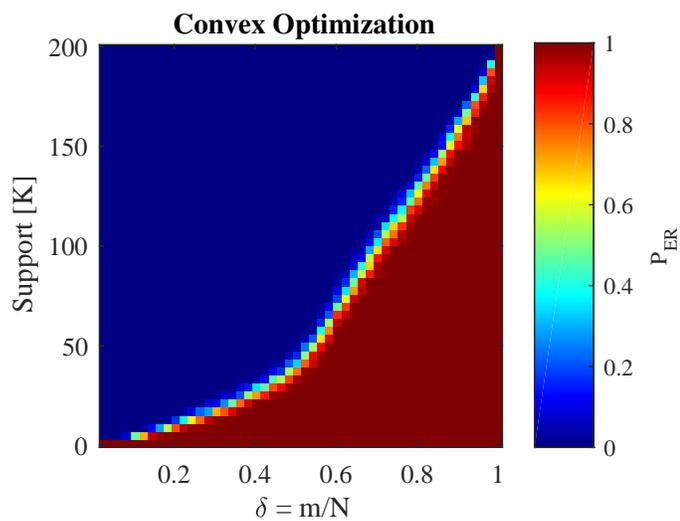


Figure 4.2: Probability of exact recovery as a function of the Support K and the down-sampling parameter $\delta = m/N$ where $N = 199$

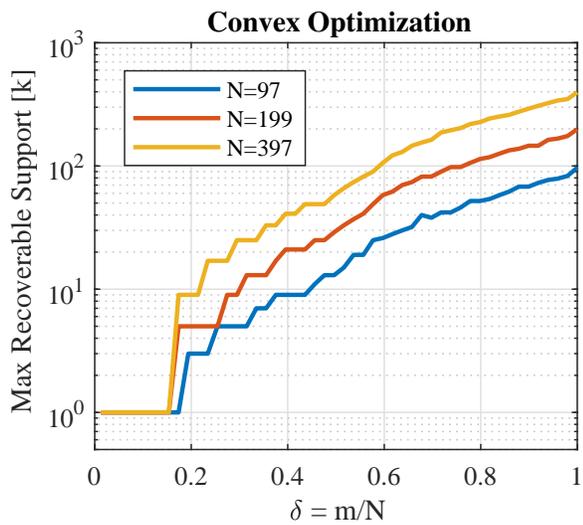


Figure 4.3: Maximum number of recoverable sparse elements within a single vector as a function of the down-sampling parameter.

vector using a sub-Nyquist sensing and wide band waveforms is possible. From figure 4.3, it can be seen that a vector with unit support can be exactly recovered given only 10% of the measurements traditionally necessary as dictated by the Nyquist sampling theorem. Additionally, it is shown that multiple sparse elements can be exactly recovered with high probability as the number of measurements increases. For a given down-sampling parameter, the maximum recoverable support is also much higher than that predicted by the mutual coherence bound. Another interesting feature observed in figure 4.3 is that the max recoverable support seems to grow proportionally with the parameter N , the total number of columns in the sensing matrix. To prove this point, the ratio $\rho = k/N$ is introduced as a measure of sparsity. This parameter is the ratio between the total number of non-zero elements in a sparse vector and the total number of possible elements. When $\rho = 1$ this implies that every element in the vector is non zero. When the max recoverable sparsity is plotted as a function of the down-sampling parameter, figure 4.4, it can be observed that the recoverable sparsity $\rho = k/N$ is independent of the total number of sensing matrix columns, and it only depends on the down-sampling parameter.

Figure 4.5 shows the average computation time necessary to exactly recover a given sparse vector. This metric is a measure of the amount of time it would take to recover a signal with maximal support at a given down-sampling ratio. This result is highly dependent on computation environment, and it should not be taken as the exact recovery time necessary, but as a useful way to compare algorithms given that all were implemented on the same machine. This plot displays the downside of convex optimization and how the computational complexity grows quickly as the

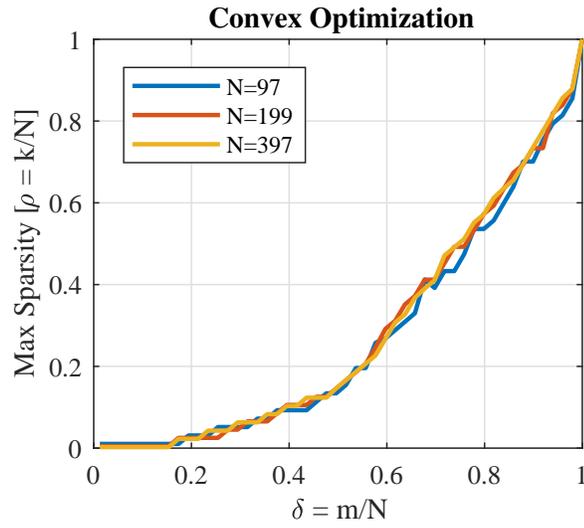


Figure 4.4: Max recoverable sparsity of a vector using convex optimization techniques.

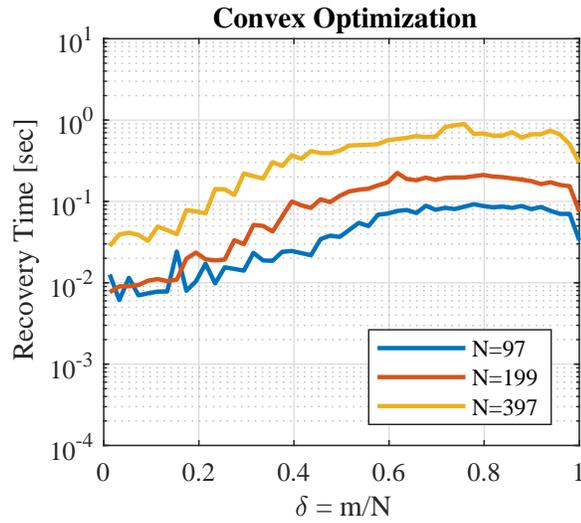


Figure 4.5: Average computation time necessary to recover a signal with maximal support at a given down-sampling ratio parameter $\delta = m/N$

size of the minimization problem grows.

4.3.2.2 Results for OMP and HTP Algorithms

While L_1 minimization through convex optimization provided reliable and robust recovery of sparse vectors, the computation time and its dependence on the size of the measurement matrices makes it an intractable problem to be applied to real time systems. The greedy algorithms can operate in fast modes, but there are difficulties with implementation. One difficulty arises from the fact that both OMP and HTP require a prior knowledge of the support of the sparse vector as an input to the algorithm. This presents a complication because often times there is no oracle that tells how many sparse elements one should expect to find. One solution is to first understand that it is impossible to recover a sparse vector with support greater than the number of measurements. Using this fact, the total number of measurements can be used as the support estimate, implying that there is the possibility of recovering one sparse element for every measurement made. To compare the two greedy algorithms, they are implemented first assuming a-priori knowledge of the sparsity, and then removing that assumption and using the number of measurements as a limit on the recoverable support.

The NSE values achieved for both greedy algorithms OMP and HTP are shown in figure 4.6. These initial results were achieved by assuming that there was a priori knowledge of the sparse vector's support. Through similar methods as discussed during the previous section, the max recoverable sparsity can be measured using

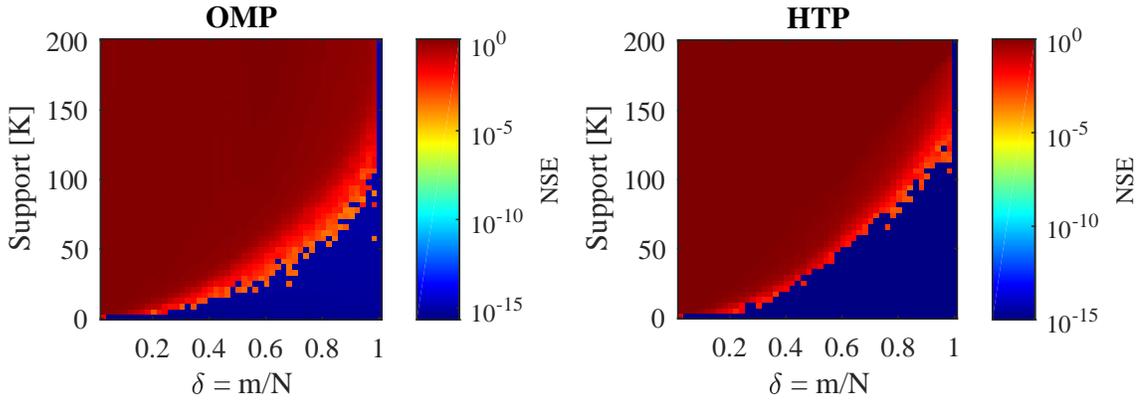


Figure 4.6: Average NSE value of the greedy optimization algorithms HTP and OMP mapped to a color value and displayed as a function of the vector support K and the down-sampling parameter $\delta = m/N$. Simulation has $N = 199$.

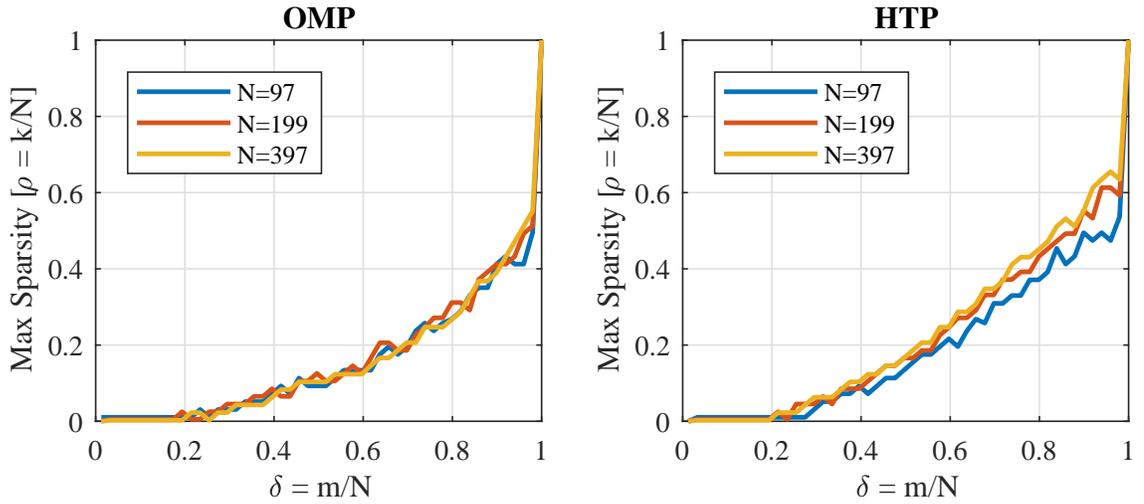


Figure 4.7: Max recoverable sparsity ρ of a vector using greedy algorithms.

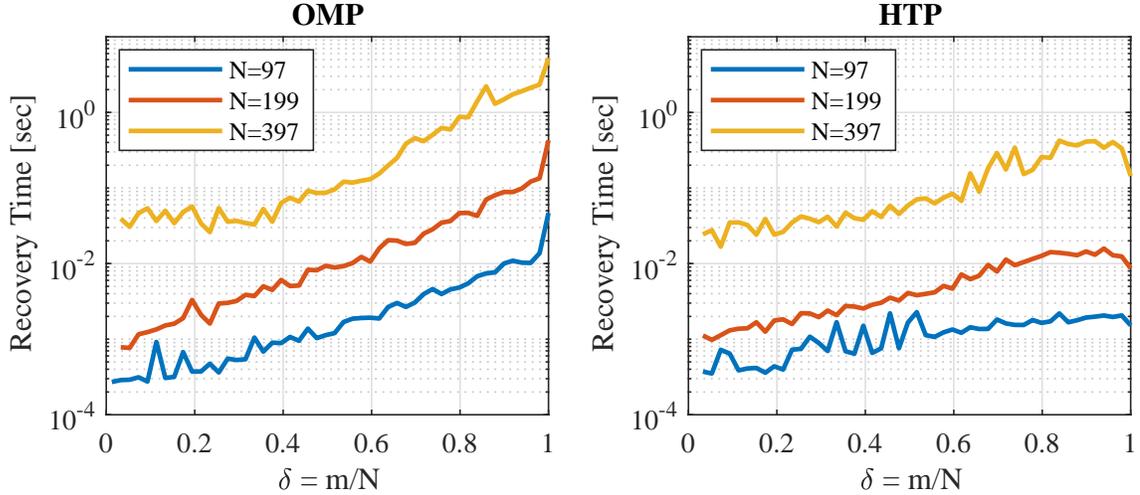


Figure 4.8: Computation time needed to obtain the optimally sparse solution using greedy algorithms.

the results of the simulations and is presented in figure 4.7. When compared to the results presented in figure 4.4, it is evident that the convex optimization algorithms recover signals with greater support than the greedy methods. However, greedy methods can still recover sparse vectors with appreciable support. Another interesting result of these simulations is that, similar to the convex optimization methods, the recoverable sparsity of greedy algorithms is independent of the number of samples traditionally needed as dictated by the Nyquist rate. When comparing the greedy algorithms to each other, HTP has slightly better recovery performance, and as shown in figure 4.8, achieves the optimal solution significantly faster as the size of the problem gets larger.

While the results shown for the greedy algorithms are promising, they rely on an often unrealistic assumption that there is a-prior knowledge of the support. For most applications, the details of the sparse vector are completely unknown. What

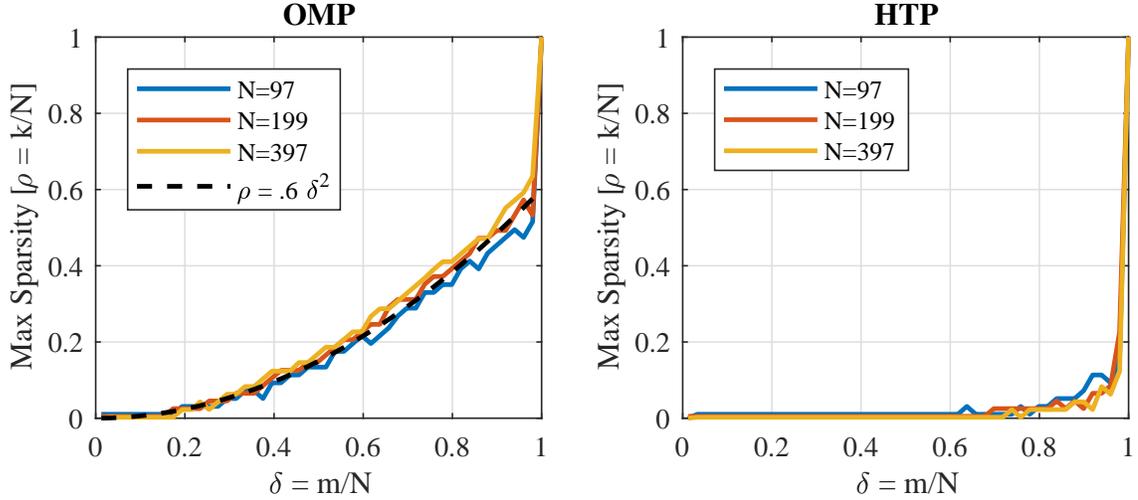


Figure 4.9: Max recoverable sparsity using the greedy OMP and HTP algorithms. The support estimate was set such that $K_{EST} = m$

is known is that the max recoverable support must be less than the total number of measurements. With this knowledge, the sparse support estimate can be limited to be equal to the total number of measurements M . This limitation drastically affects the results of recovery for the HTP algorithm. Using the same methods as described previously, the max recoverable sparsity can be measured as a function of the down-sampling parameter, and the results are shown in figure 4.9. These results show that the HTP algorithm is highly susceptible to the sparsity estimate whereas the OMP algorithm is not. The differences between the data presented in figures 4.7 and 4.9 is because of the different way each algorithm updates the sparse solution. For OMP, a sparse element is added at each iteration one by one according to the highest remaining residual correlation. Once the proper support is reached, the algorithm terminates because the error is below the threshold. The HTP algorithm however attempts to solve the problem using the entire estimated

support set and then updates the solution based on the stepping and threshold parameter. This leads to a potential solution that might minimize the error, but it is not the unique sparsest solution.

These simulations show that the OMP algorithm leads to a uniquely sparse solution more reliably than the HTP algorithm despite the initial support estimate. In figure 4.9, an empirically measured relationship between the recoverable sparsity $\rho = \frac{k}{N}$ and the down sampling ratio $\delta = \frac{m}{N}$ was found using a least squares fit. This relationship was measured to be $\rho \approx \mathcal{O}(\frac{3}{5}\delta^2)$, and provides a rough estimate on the number of measurements necessary to recover a sparse vector with a given support. This relationship between the number of measurements M , number of traditional Nyquist samples N , and support k is shown in equation (4.8).

$$M > \mathcal{O}\left(\sqrt{\frac{5}{3}}\sqrt{Nk}\right) \quad (4.8)$$

4.3.2.3 Discussion and Algorithm Selection

In this section three different methods for recovering sparse vectors from sub-Nyquist sampled versions of a chirp waveform were explored. While the convex optimization techniques performed the best in terms of max recoverable sparsity, the computational complexity was far higher than the alternative greedy algorithms. The two greedy matching pursuit algorithms, Orthogonal Matching pursuit and Hard Thresholding pursuit, both performed well when provided with an exact support estimate. However, HTP was unable to recover multiple sparse elements when

that ideal estimate was removed. Due to its relatively low computational burden and its exceptional sparse recovery characteristics, the OMP algorithm will be used for the remainder of this work to recover sparse vectors.

At this point, it has been shown empirically that the coherence of the the sub-Nyquist sensing matrix derived from linear frequency modulated waveforms satisfies the mutual coherence property for successful recovery. Through Monte Carlo simulations, an estimate for the max recoverable sparsity has been provided as a function of the down-sampling parameter. It was also shown that the max recoverable sparsity is independent of the total number of columns of the sensing matrix, and non-ideal estimates of the support of the sparse vector can be used. In the following chapters, the noise susceptibility of sub-Nyquist sensing systems will be explored.

Chapter 5: Analysis of Noise Susceptibility

In real world sensing systems, noise is an important factor to consider, and incorporating a noise model into a compressive sensing framework is necessary to predict potential system performance. The main source of noise is identified as the optical to electrical conversion process at the PMT, which introduces shot noise [69]. Before sampling, the noise from the photo-cathode is amplified over the multiple dynode amplification stages and increases total noise. To properly simulate this, the noise is added according to equation (5.1).

$$y = \Phi(\Psi s + n) \tag{5.1}$$

In this model, the noise is added prior to the measurement processes, and this method is considered noise folding [70]. This processes has effects that are very different from the case when only measurement noise is assumed.

5.1 Recovery of Sparse Vectors in the Presence of Noise

The goal of this section is to provide a robust analysis of the noise susceptibility of the sub-Nyquist measurement and the OMP algorithm as applied to sparse recovery. In the previous section, it was shown that these methods can recover multiple sparse elements with high probability given sampling rates much lower than

those dictated by the traditional Nyquist sampling theorem. While recovery of multiple sparse elements is an interesting problem, for the hybrid lidar-radar problem often it is only necessary to recover a sparse vector with a single non zero value. Due to the fact that these methods are targeting optical ranging systems where only a single reflection is expected, the following simulations will be limited to the recovery of a single sparse vector.

5.1.1 Simulation Method

To test the susceptibility of the sub-Nyquist sensing system to noise, three independent variables were targeted. The first two variables are similar to the simulations conducted before, and they are the total number of sensing matrix columns, N , which is related to the waveform duration, and the down-sampling ratio $\delta = m/N$ which is related to the sub-Nyquist sample rate. The third parameter is the Signal to Noise Ratio (SNR) where the signal is defined as $x = \Psi s$, the noise is n from equation (5.1) and the SNR is measured in dB.

While the exact recovery conditions provide useful insight into the recovery characteristics of the chosen algorithms, exact recovery is not a necessary condition for the hybrid lidar-radar ranging system. For ranging applications, typically the only necessary information is the range to the illuminated object. This is equatable to the concept of support set recovery where the proper index set for the sparse vector is found versus exactly recovering the index and amplitude of each non-zero element in the sparse vector [71]. For these simulations, the average probability of

detection can then be calculated by (5.2) which averages the total number of correct detections over the number of trials. For these simulations because it is assumed that the sparse vector has unit support, the detection can only be a correct detection or missed detection as defined by binary hypothesis testing [72].

$$P_D = \frac{\text{Number of Correct Detections}}{\text{Total Number of Trials}} \quad (5.2)$$

The simulation steps are outlined in 5.1.1 and define the independent trials that were used to calculate the statistics on the sparse recovery performance of sub-Nyquist measurement systems. The sparse vector had a single non-zero element with unit amplitude and a random range index $\lambda \in [1 \cdots N]$. The noise was drawn from a random Gaussian distribution and scaled so that the noise power provided the desired SNR. The OMP algorithm was used for sparse recovery, and the sparsity estimate was set to one given the nature of the ranging application. The probability of detection performance metric was calculated using 100 independent trials.

Simulation Steps

- (1) Set SNR and N
 - (2) Adjust M to change δ
 - (3) Generate the Chirp Matrix Ψ and Sampling Matrix Φ
 - (4) For each trial
 - (4.1) Generate the sparse vector s
 - (4.2) Generate the noisy signal $x = \Psi s + n$
 - (4.3) Generate the Sub-Nyquist Measurements $y = \Phi x$
 - (4.4) Implement the OMP Algorithm
 - (4.5) Determine Correct/Failed Detection
 - (5) Calculate the Probability of Detection
-

5.1.2 Simulations Results

Simulations varied the SNR from 0 to -14 dB, and the recovery probability was measured as a function of N and δ . Figure 5.1 shows the results of these simulation and in this figure, the probability of detection was mapped to a color value and displayed as a function of δ and N . These initial results show that it is indeed possible to correctly recover the support of the sparse elements and, thus, an object's range even in the presence of significant amounts of noise. A deviation from the noiseless results is observed, where in the noiseless simulation there was no observable dependence on the waveform duration, and recovery was solely dependent on the vector's sparsity and the down sampling ratio. In the presence of noise, the probability of detection is observed to be highly dependent on the waveform duration N .

To further investigate the relationship between SNR and the waveform duration N the minimum down-sampling parameter was measured which guaranteed a probability of detection greater than 99%. This measured parameter provides a bound on the minimum number of measurements necessary to recover a sparse vector with support one as a function of waveform duration and SNR. The results of these simulations are summarized in figure 5.2, and a more detailed analysis can be seen in [53]. These results show that sub-Nyquist sensing methods can overcome noise limitations by increasing the duration of the waveform. This implies that given a system designed to operate at a sample rate some ratio lower than Nyquist, the only change necessary to improve performance in the presence of noise would be

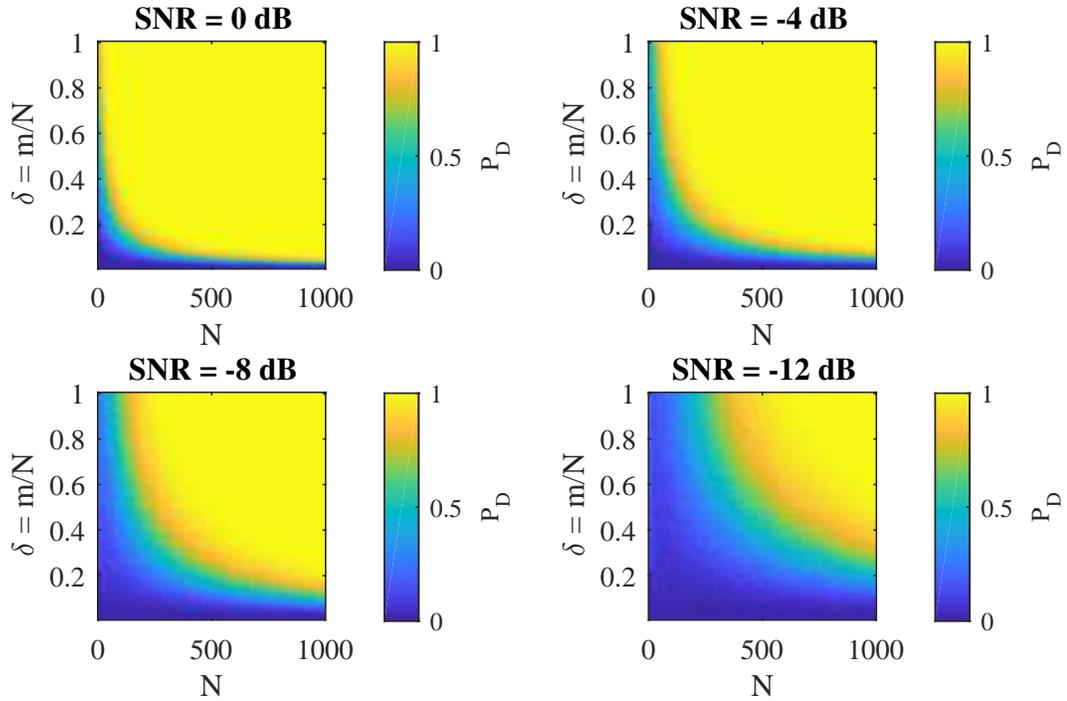


Figure 5.1: Simulation results for various SNR levels. Results show that even when sampling at rates far less than Nyquist it is possible to recover a sparse vector with unit support.

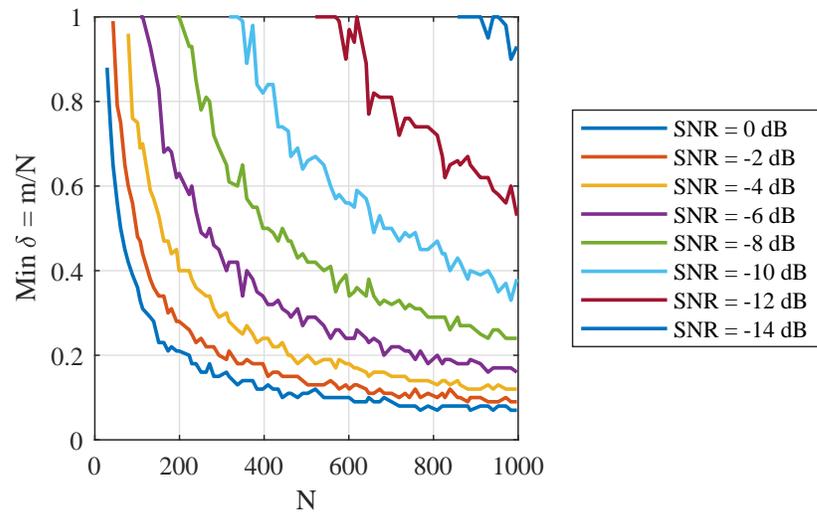


Figure 5.2: Minimum required down-sampling parameter necessary for 99% probability of detection Vs N for various SNR levels

to change the waveform duration and generate a new sub-Nyquist sensing matrix. A detailed example of this from the provided results is as follows. Given a system operating at $\delta \approx .2$, or at 20% the Nyquist rate, the system could recover a target with an SNR as low as 0 dB using a waveform only 223 samples long. If the SNR was reduced to -6 dB, the same system could still operate at the same rate but the duration would need to be increased to 769 samples. This is easily accomplished using software defined radios which can adjust waveform parameters without changing any hardware.

This result is very closely related to the concept of the time bandwidth product gain observed when using matched filtering detection methods [73]. Matched filters are the ideal filter in the presence of noise and have been used extensively in radar and sonar to maximize SNR when determining the range to objects. To compare the sparse recovery methods to this traditional detection scheme, the previous simulations are re-run using the Nyquist sampled version of the FMCW waveform and a matched filter instead of the OMP algorithm. The results of these simulations are shown in figure 5.3 where the probability of detection is plotted as a function of waveform duration, N , and SNR. The results show that when using matched filtering methods, there is also a minimum waveform duration needed to recover the correct range in the presence of noise.

To compare the results of the matched filter and sub-Nyquist simulations, the results from figure 5.3, and 5.2 are used. For a given SNR, the minimum waveform duration N resulting in a 99% probability of detection for the matched filter detector method is found. This parameter is obtained from the data in figure 5.3 and reported

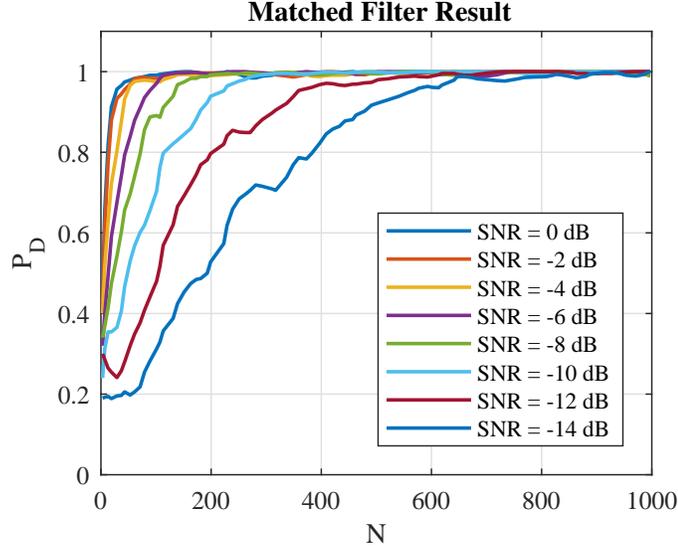


Figure 5.3: Simulation results showing the Matched Filter performance as a function of chirp duration and SNR

in table 5.1 in the second column, where the first column is the SNR. To quantify the sub-Nyquist results for a given SNR, the minimum waveform duration was found for a number of different downsampling parameters that led to successful reconstruction from figure 5.2. The minimum waveform duration necessary for successful recovery is reported in 5.1 as a function of the downsampling parameter and SNR.

The results presented in table 5.1 show that the sub-Nyquist sensing system can recover signals at the same SNR as matched filtering methods while using sampling rates much lower than the Nyquist sampling rate, $F_D \leq .5 \times F_N$. However, the results from these simulation also show that the OMP recovery algorithm typically requires a longer waveform duration to recover a signal in the presence of noise when compared to the matched filter under similar noise levels. These results highlight the trade-off between sampling rate and noise performance when attempting to use

Table 5.1: Summary of Matched Filter and Sub-Nyquist sensing methods in the presence of noise. The first column is the SNR. The second column is the minimum waveform duration necessary to recover a target using matched filter processing. The subsequent columns are the minimum waveform duration for a given $\delta = \frac{m}{N}$ used to successfully recover a target at the same SNR using OMP

SNR	Matched Filter	Sparse Recovery				
		$\delta = .5$	$\delta = .4$	$\delta = .3$	$\delta = .2$	$\delta = .1$
0	79	61	89	131	223	541
-2	113	101	127	173	293	719
-4	131	157	199	311	457	>997
-6	181	251	359	457	743	>997
-8	251	397	541	787	>997	>997
-10	293	659	881	>997	>997	>997

sub-Nyquist sensing and sparse recovery. As the sampling rate is reduced more measurements are necessary to successfully recover the sparse vector, which implies that the duration must be increased. It is important to note though that while the sparse recovery methods are more susceptible to noise, there is relatively little cost to increase waveform duration to improve the detection of sparse vectors in the presence of noise.

5.2 Experimental Verification

Experimental verification was completed to provide certification that the simulation results are able to be recreated using available hardware and in the presence of real noise. The experiment was designed in a way to test different simulated system setups so that comparisons could be made between the simulated results and the experimental verification data. The results were achieved using FMCW

waveforms, variable rate A/D converters, and analog noise sources.

5.2.1 Experimental Design

For these preliminary tests, the AD9914 digital synthesizer board was used as the source to generate FMCW waveforms with a bandwidth that matched the simulated waveform parameters, and a variable chirp duration was used to increase the number of traditionally necessary Nyquist samples. A constant delay line was used to approximate the propagation delay between the source and transmitter, and an Agilent noise source was used to generate signals with various SNRs. A SP-Devices variable rate digitizer was used to implement a range of sampling rates, and the OMP sparse recovery algorithm was implemented to determine the measured time delay using the sensing matrix generated by the known chirp parameters. Prior to the experiment, the matched filter detection method was used to measure the constant delay to provide a ground truth for calculation of the recovery probability. Recovery probability was measured by comparing the previously measured time delay to that of the recovered sparse index across multiple trials. This definition of recovery probability is more closely related to that of sparse support set recovery as defined in the compressive sensing literature [71]. The experimental block diagram is shown in figure 5.4.

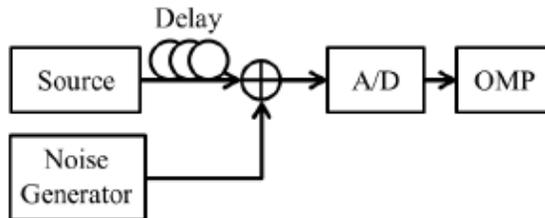


Figure 5.4: Block diagram for preliminary experiments

5.2.2 Experimental Results

Given that the goal of this experiment is to validate the simulated results, similar parameters to those used in prior simulations were tested. The experimental results are displayed alongside those found during simulations. Figures 5.5 and 5.6 show the results of these experiments and simulated results for Nyquist sample values $N \approx 250$ and 1000 respectively, and the noise source was set to achieve SNRs of 0 and -10dB for both data sets. For the given scenarios, the recovery probability was measured and plotted as a function of the downsampling ratio which was adjusted by changing the digitizer sampling rate.

It is observed that when using the experimental equipment and under-sampling the recovery statistics displayed trends similar to those achieved during simulation. However, the experimental data had slightly worse recovery probability and lacked features found in the simulation data when the down-sampling ratio was close to even integer multiples of the duration, N . The slightly decreased performance was studied [53, 74] and can be attributed to dictionary mismatch [75]. The dictionary

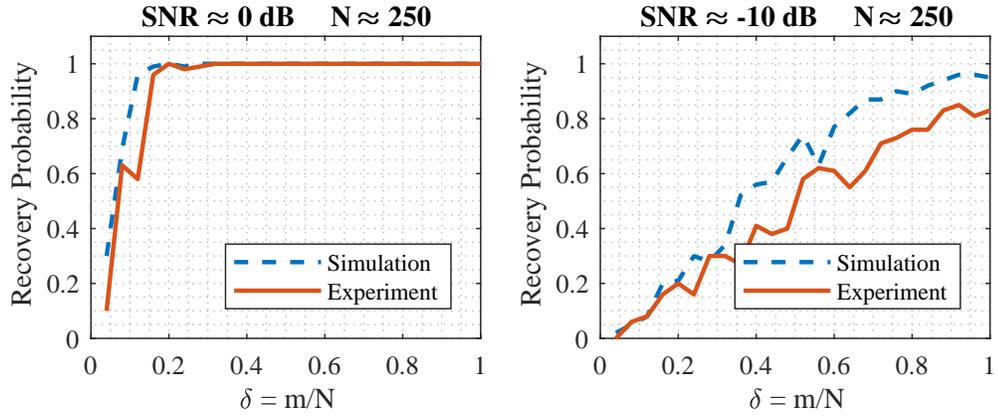


Figure 5.5: Experimental Vs Simulated Results for $N \approx 250$ with a SNR of 0 and -10 dB.

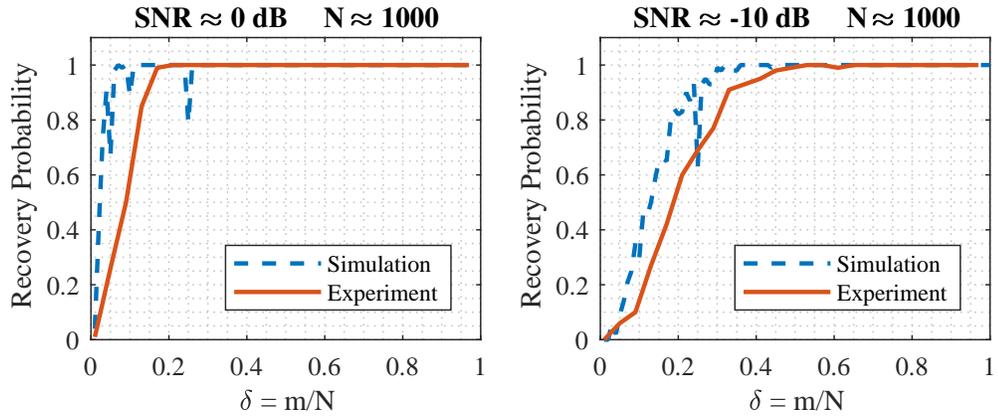


Figure 5.6: Experimental Vs Simulated Results for $N \approx 1000$ with a SNR of 0 and -10 dB.

mismatch stems from the fact that the generated sub-Nyquist sensing system is based off of an ideal linear FMCW waveform, whereas the DDS board generates a waveform based on discrete frequency steps. This mismatch results in measurements that are not as well correlated with the sensing matrix elements. The missing features in the experimental data when the down-sampling ratio is close to even integer multiples can be described by the fact that for a real sensing system, it is highly unlikely for a sparse element to fall directly into a single range bin. This is often referred to as sampling off the grid [76]. This feature implies that the system is never sampling the linear FMCW signal on the same grid that the dictionary is stored at, and it is highly unlikely that it will be perfectly correlated with more than one column of the sensing matrix. Despite the differences between experimental and simulated data, these results prove that it is possible to recover the accurate time of flight information from sub-Nyquist measurements of wide-band waveforms.

5.3 Summary

In this section, simulated and experimental evidence for the sub-Nyquist hybrid lidar-radar system was provided showing that the recovery of sparse vectors in the presence of noise is possible. The results from the initial simulations showed that while the sub-Nyquist methods perform worse than matched filter methods at high levels of down-sampling, performance in the presence of noise can be improved by increasing N , the waveform duration. The experimental verification of the conducted simulations provides increased confidence in the simulation results, and it

provides evidence that it is possible to improve performance in the presence of noise by increasing the waveform duration.

These results are promising because it implies that it is possible for a sub-Nyquist ranging system to improve its performance in the presence of noise in-situ. If the environment causes reduced SNR, the system could increase the duration of the chirp waveform that is modulating the laser and the sensing matrix could be updated given the new values of N . If the sampling rate of the digitizer remains constant, the downsampling ratio remains the same. Given the results from figure 5.2 and the described changes, it has been shown that the noise performance would be improved.

Chapter 6: Resolution of Sparse Recovery Techniques

In the previous section, it was shown that it is possible to recovery sparse vectors with various levels of sparsity, and even possible to recover sparse vectors in the presence of noise. Furthermore, it was shown that the sub-Nyquist sensing framework and the OMP algorithm is capable of performing in similar noise environments as traditional matched filters at rates lower than the traditional Nyquist rate. In this section, the achievable resolution of the sub-Nyquist sensing and sparse recovery methods will be investigated, and the differences between these techniques and traditional matched filter methods will be described.

6.1 Resolution of Traditional Methods

Traditional methods for determining the range to an object involve the use of wide-band waveforms and matched filters. The gain of these methods was described in chapter 5 where the time bandwidth product provides amplification for signal detection in noisy environments. Wide-bandwidth techniques also utilize pulse compression, where the modulated waveform with duration proportional to N is compressed in time to be inversely proportional to the bandwidth Δf [25,77]. The output of the pulse compression process has a well know ambiguity function for the

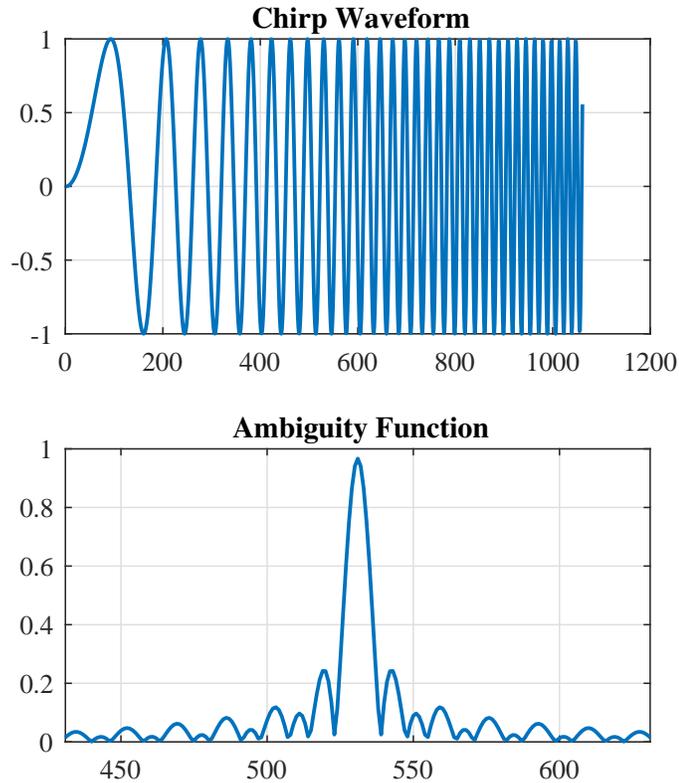


Figure 6.1: Example of the LFM waveform and the ambiguity function arising from pulse compression

linear frequency modulated waveform (LFM). An example of the LFM waveform and its associated ambiguity function are shown in figure 6.1.

Radar pulse compression techniques have been applied to various hybrid lidar-radar type systems [14, 78]. Similar to radar applications, these methods suffer from two main limitations in terms of resolution. The first limitation is that for targets spaced such that they both reside within the ambiguity function, it is impossible to resolve them. This effect can be seen in figure 6.2. Another common problem is when large returns can shadow nearby weaker signals. This effect is shown in figure 6.3 where the weaker target is shown to be irrecoverable from the side-bands of

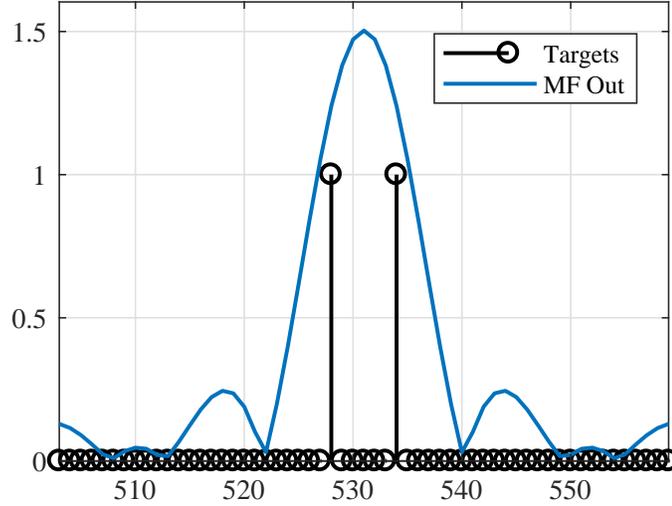


Figure 6.2: Example of the inherent resolution limit of the pulse compression techniques.

the stronger return. Methods have been studied to circumvent these issues [79, 80]. However, the implementation often requires multiple observations of the scene which is not often possible.

The following section will explore how the Sub-Nyquist sensing formulation and sparse recovery techniques can be used to improve achievable resolution using an alteration to the structure and design of the measurement matrix. These techniques will be applied to the two previously discussed problems of closely spaced targets and shadowing, and empirical evidence will be provided that the sparse recovery techniques do not suffer from the same ambiguity function limitation. The theory stems from the work within the fields of super resolution and over-complete dictionary representations [81, 82].

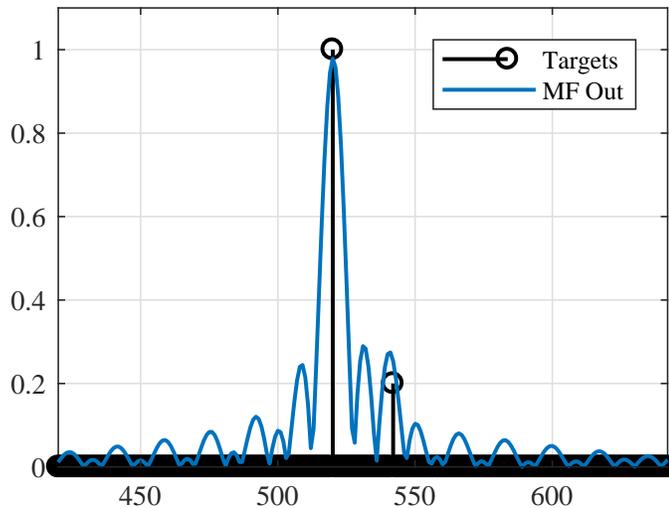


Figure 6.3: Example of the shadowing effect from large returns near weaker returns.

6.2 Resolution Enhancement of Sub-Nyquist Chirp Sensing Matrix

When the sub-Nyquist sensing matrix was original defined, the columns of the sensing matrix represented all possible time shifts that could potentially be recovered. The resolution for these times shifts was defined by the Nyquist rate of the transmitted waveform and given by $\delta t = 1/F_N$ where F_N is the Nyquist rate for the given waveform. Instead of generating a sensing matrix defined by the Nyquist rate, the sensing matrix can be defined where the columns represent more highly resolved time shifts. This leads to an over complete representation of the original waveform matrix, Ψ , where the elements are given by equations (6.1) and (6.2). In this representation, the discrete duration N_S is now given by $N_S = \lfloor T \times F_S \rfloor$, where T is the waveform duration, and F_S is the new sampling rate of the stored FMCW waveform such that $F_S > F_N$. The normalized frequencies are represented

by $\omega_0, \omega_1 \in [0, \frac{F_N}{F_S}]$ and the indexes $m, n \in [0, N]$.

$$h[m] = \sin \left[\frac{F_N}{F_S} \left(\omega_0 m + \frac{\omega_1 - \omega_0}{N} m^2 \right) \right] \quad (6.1)$$

$$\psi_n = \sum_{m=1}^N h[m] \delta[m - n] \quad (6.2)$$

For this framework, what has now been described is a resolution enhanced chirp dictionary $\Psi = [\psi_1, \psi_2, \dots, \psi_N]$ where the columns are time shift versions of an over-sampled FMCW waveform. The measurement matrix Φ would be generated using the same method as previously described and used to implement a model for the uniform under-sampling of the signal from a rate of F_S to the desired analog to digital conversion rate, F_D . When using this enhanced resolution sensing matrix, the sparse vector to be recovered now has a time bin spacing inversely proportional to the column sampling rate F_S , while the measurements are made at the rate F_D .

Increasing the resolution of the sensing matrix comes at a cost. For the same duration waveform it increases the size of the problem thus increasing the computational complexity. Additionally by increasing the time resolution, the coherence of the sensing matrix is significantly altered. The higher resolution causes an over complete dictionary which leads to an increased coherence. The effect that over complete representations have on coherence is shown in figure 6.4. Here, the coherence is plotted as a function of the down sampling parameter for various levels of column wise time resolution enhancement measured as the number of times over Nyquist. There is a slight change in representation where the down sampling parameter is now the ratio between the sub-Nyquist sampling rate, F_D , and the Nyquist rate, F_N . These results show that as the time resolution of the sensing matrices columns

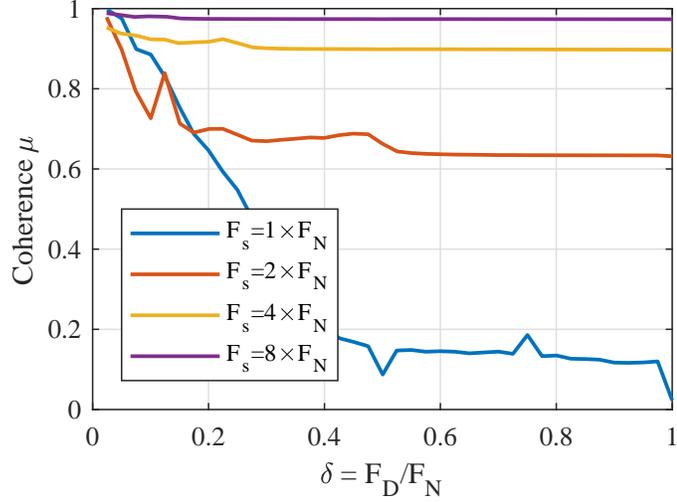


Figure 6.4: Coherence shown as a function of the down sampling parameter for the over complete sub-Nyquist sensing matrix.

increase so does the coherence. The increasing coherence as a function of column wise resolution enhancement implies that the recovery of multiple sparse vectors is unlikely. However, as mentioned previously, the mutual coherence condition is not a necessary condition. While OMP will likely fail for highly coherent sensing matrices, it has been shown that L1 minimization techniques such as the previously described interior point methods can be used to recover sparse vectors even when the sensing matrix is highly coherent [83,84]. In the following sections, the enhanced resolution sub-Nyquist sensing system and sparse recovery methods will be applied to the scenarios where traditional matched filter-based detectors have limited resolution due to the ambiguity function.

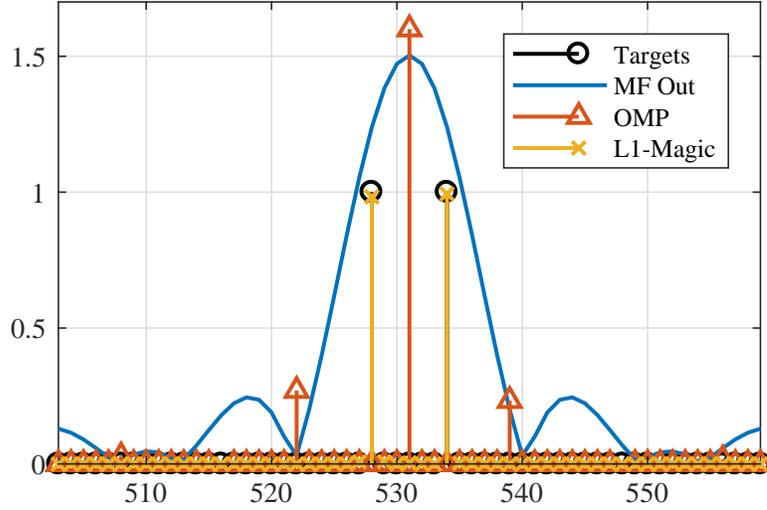


Figure 6.5: Recovery of closely spaced sparse vectors using the resolution enhanced sensing matrix compared to matched filter detection.

6.2.1 Recovery of Closely Spaced Returns

In the previous section, the first scenario presented was the simulation of two targets of equal amplitude spaced within the ambiguity function of the matched filter output. By simulating the same scenario and implementing the sparse recovery methods, the potential benefit of these methods can be observed in figure 6.5. In this simulation, the sensing matrix resolution was enhanced to $F_S = 8 \times F_N$ and the measurement matrix was generated using a sampling rate equal to the Nyquist rate, $F_D = F_N$. This data provides a comparison between the sparse recovery techniques and matched filter detection methods.

Figure 6.5 shows that, as predicted, the OMP algorithm fails to recover both closely spaced returns and provides a maximal output similar to that of the matched filter. This is due to the highly coherent nature of the resolution enhanced sensing

matrix and the fact that OMP depends on the minimization of the residual correlations. While OMP was unsuccessful, the data shows that if convex optimization methods for sparse recovery replaced the matched filter, the detectable resolution could potentially be far higher than what is achievable with the traditional matched filter approach.

To provide further analysis of the resolution of the sparse recovery methods, a simulation was designed where the spacing between two equal amplitude targets was varied, and the different detection methods were implemented. The amplitudes of the recovered outputs using the matched filter, OMP, and convex optimization algorithms were mapped to a color value and plotted as a function of the recovered target index and the target spacing. The recovered target index and target spacing, both measured in samples, are equivalent to the measured range and range separation respectively in a ranging application. The actual target indexes are superimposed onto the resulting image as solid and dashed red lines. The results of this simulation can be seen in figure 6.6, where the sampling rate of the sensing system is set to be the Nyquist rate of the transmitted waveform, allowing for a direct comparison of the Matched Filter and Sparse Recovery algorithms. For the OMP algorithm, similar assumptions to the previous sections were used, where the sparse estimate K_{est} was set to equal the total number of measurements M .

The results presented in figure 6.6 provide an insight into the minimal separation between targets necessary for successful recovery of their individual locations. For the matched filter method, the minimum spacing necessary is roughly twice the ratio $\frac{E_S}{E_N}$ which is as expected given the ambiguity function. These results show that

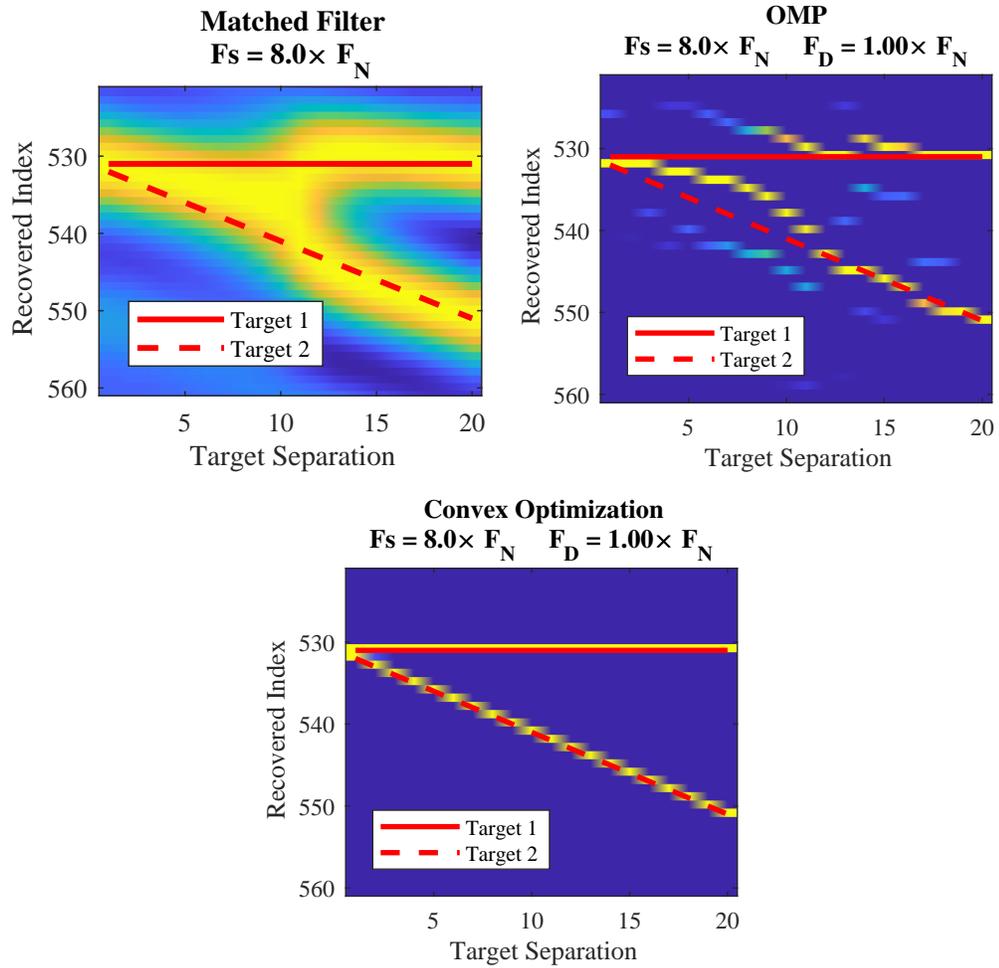


Figure 6.6: Detection method results for the resolution test. Two closely spaced targets of equal magnitude with varying separation were used and the magnitude of the detection method outputs were mapped to a color value. Plotted as a function of recovered index and target separation. The solid and dashed red lines represent the actual target indexes. Yellow values are the estimated target location for the given method.

by simply increasing the time resolution of the sampling method, there is no benefit in terms of recoverable target separation for the matched filter method. The OMP detection method provides similar results to the matched filter and only tracks the two separate targets reliably when they are spaced ≈ 12 samples/range bins apart, a slight improvement over the matched filter method. Aside from the slight improvement, one major difference between OMP and the matched filter method is that, by increasing time resolution of the sensing matrix, the OMP algorithm results in an output that has better temporal accuracy. This highlights the improved accuracy of the OMP detection method and shows how the output is not affected by the same time-frequency uncertainty as the matched filtering methods. When applying the convex optimization schemes, specifically L1-Magic [2], a substantial improvement in achievable resolution can be obtained. From these results, it is shown that two separate targets with equal amplitude can be recovered at a spacing as close as two samples, or equivalently two range bins, apart. This result implies that by simply increasing the column wise time resolution of the stored chirp matrix, sampling at the Nyquist rate, and using convex optimization, it is possible to achieve a time/range resolution that is proportional to the resolution of the stored sensing matrix, not the bandwidth of the transmitted waveform.

The previous results proved that resolution enhancement can be achieved by increasing the column wise time resolution of the sensing matrix and using convex optimization instead of matched filtering to recover the assumed sparse vector. These simulations utilized a sampling rate equal to the Nyquist rate of the transmitted waveform. To test the achievable resolution of the sub-Nyquist methods, the

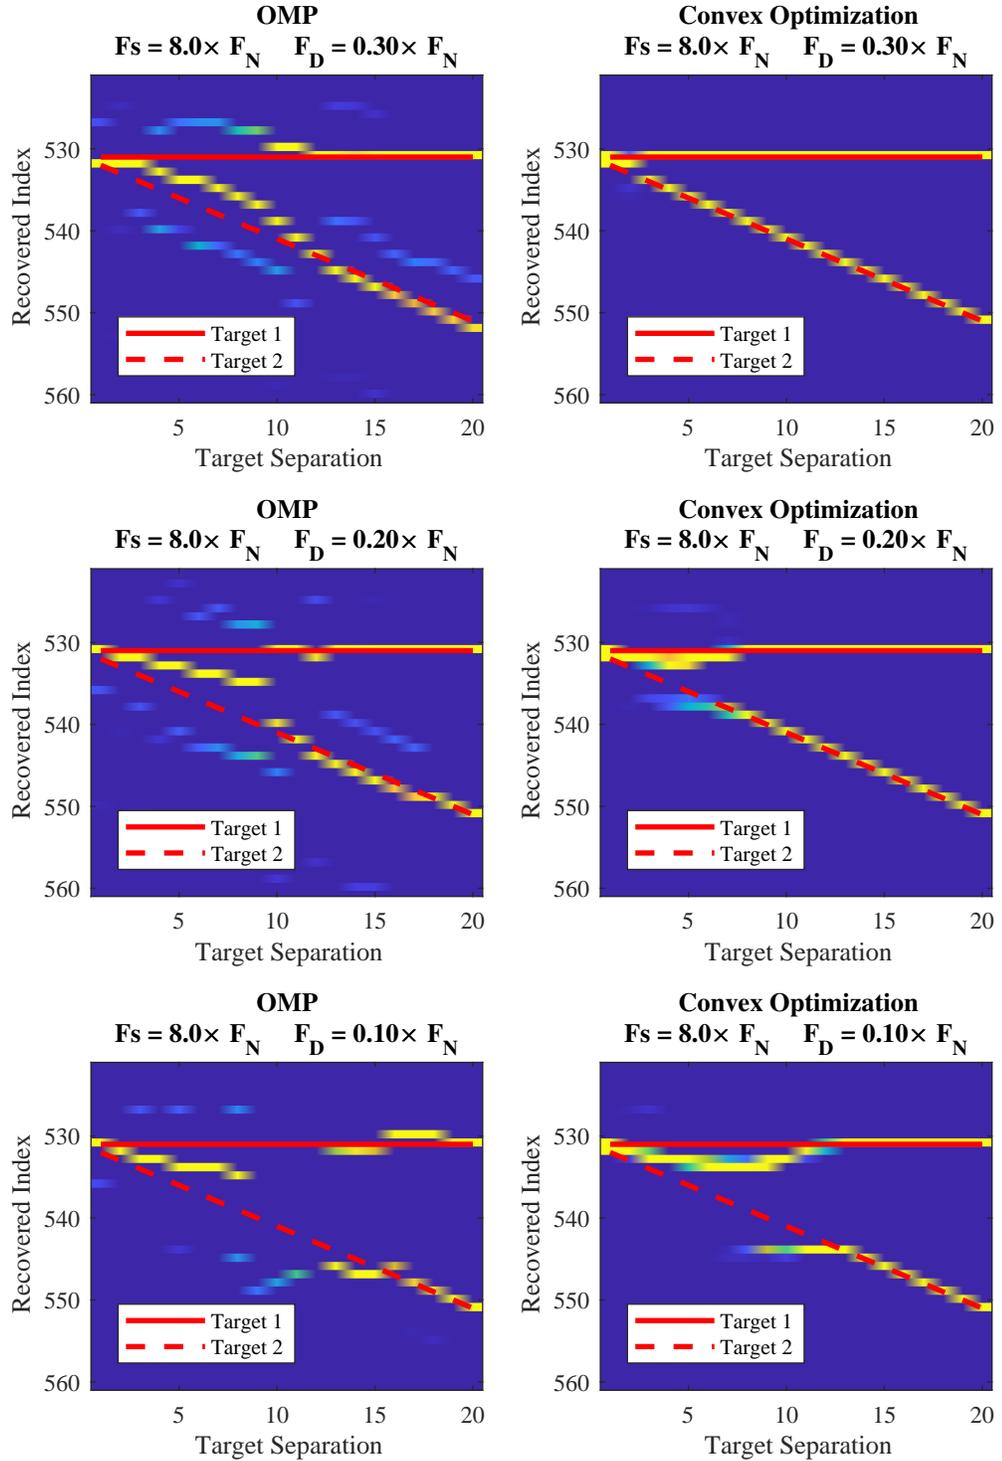


Figure 6.7: Results of the simulations where two targets are spaced at varying distances, and multiple sub-Nyquist rates are utilized. The OMP and Convex Optimization algorithms were used to recover the sparse vector, and the magnitude of the recovered results was mapped to a color.

simulations were conducted again while implementing a sampling rate 70%, 80%, and 90% slower than the transmitted waveforms Nyquist rate. The results are presented in figure 6.7 using similar conventions as described previously. These results can be interpreted by inspecting each graph from left to right and noting when the two separate targets, represented by yellow pixels, match the solid and dashed red lines. It can be seen that even when sampling at a rates 80% lower than the Nyquist rate, the convex optimization techniques can realize and recover the target given a higher resolution sensing matrix. These results show that even when sampling at rates much lower than the Nyquist rate, it is possible to improve the range resolution by utilizing a sensing matrix that is derived from a waveform stored at rates higher than Nyquist. It has also been shown that sparse recovery techniques do not suffer from the same time-frequency uncertainty principles as the matched filter methods, and that resolution enhancement allows for the detection of objects spaced within the traditional ambiguity function produced by the matched filter.

6.2.2 Recovery of Shadowed Returns

The second major challenge for traditional matched filtering methods is shown in figure 6.3 where strong returns shadow nearby weaker returns. This effect is due to the relatively strong side-lobes of the chirp waveform ambiguity function. While methods exist to reduce the side-lobes [80,85,86], these come at various costs such as broadening the main lobe through pulse shaping, increasing integration time using adaptation, and reducing the pulse compression gain. In this section, the recovery of

weak returns in close proximity to strong returns will be tested using the OMP and convex optimization algorithms presented previously. The resulting sparse solutions of these algorithms will be compared to traditional matched filtering methods, and the results as a function of two different sampling ratios will be provided.

To test the application of sparse recovery algorithms to the shadowing problem, a sparse vector was simulated with two non-zero elements placed at indexes such that they would be spaced far enough apart to be outside of the traditional ambiguity function. One non-zero element was given a magnitude 10 times greater than the other, and a Nyquist sampled system was investigated. By initially using the Nyquist sampling rate, a direct comparison to the traditional matched filter detection process can be made and the advantages of using the sparse constraint in detection can be shown. The sensing matrix was generated using the resolution enhancement discussed previously and the columns were stored at a resolution $8 \times F_N$. Both the OMP and Convex Optimization algorithms were tested and the results, shown in figure 6.8, provide the sparse estimate of the different algorithms, the matched filter output, and the actual sparse input vector. Figure 6.9 shows the results when using a rate 90% lower than the waveform Nyquist rate.

Figures 6.8 and 6.9, show that both the OMP and convex optimization algorithms perform better than the matched filter in terms of recovering both the strong and weakly reflecting targets. These results also show that the sparse recovery techniques are not susceptible to the same shadowing effects as the matched filter method, and can be used to recover weak returns in the presence of larger signals.

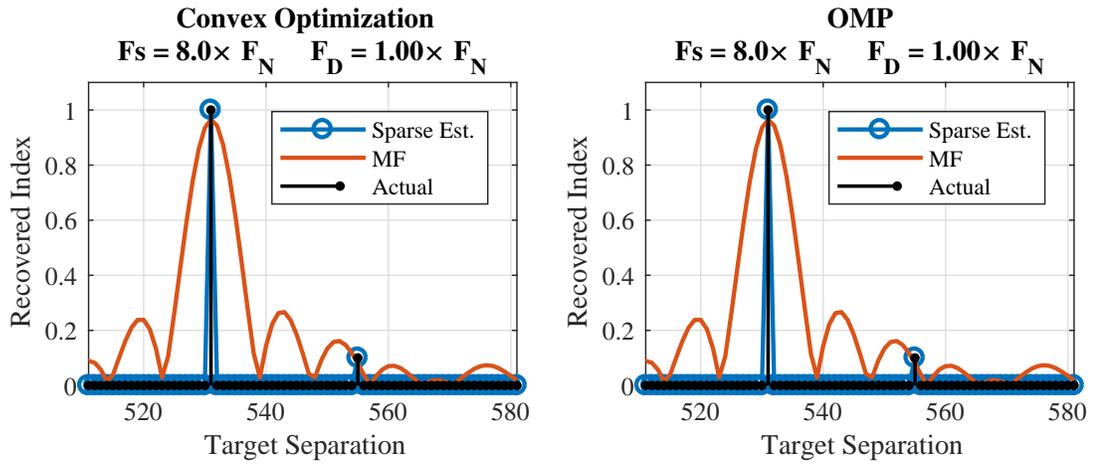


Figure 6.8: Results from Nyquist rate simulations of the shadowing effect. Shows that sparse recovery methods are not effected by side-bands and shadowing like the matched filtering methods.

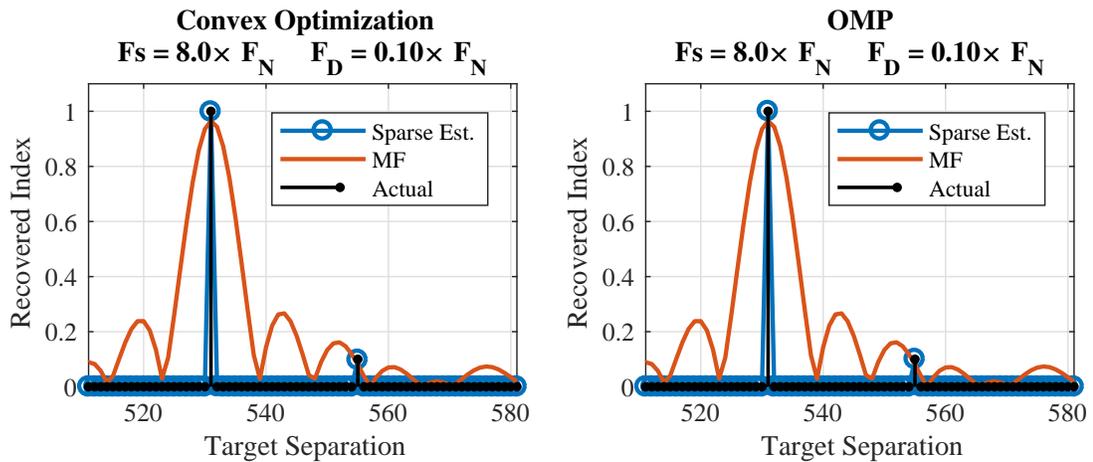


Figure 6.9: Results from Sub-Nyquist rate simulations of the shadowing effect. Shows that sparse recovery methods are not effected by side-bands like the matched filtering methods are.

6.3 Summary

The difficulties in detecting two closely spaced reflecting objects and weakly reflecting objects located in close proximity was explored in the section. It was shown that when using traditional matched filter detection schemes it is impossible to detect objects located within the ambiguity function. Additionally, it was shown that the side-bands of the ambiguity function can shadow weak returns. Through simulations, evidence was provided that by using the sub-Nyquist sensing methods and increasing the time resolution of the stored sensing matrix, it was possible to increase the achievable time resolution of the recovered sparse target vector. This results implies that the $L1$ minimization methods are not affected by the ambiguity function typically encountered when using wide-band modulation.

These shown results imply that sub-Nyquist sensing and sparse recovery when applied to hybrid lidar-radar ranging systems, can improve upon the resolution of current methods while still sampling at rates much lower than the Nyquist rate. This result is important for the cases when the environment is highly scattering, causing multiple target features to be illuminated which leads to multiple reflections being observed by the receiver.

Chapter 7: Application to Underwater Lidar

In the previous sections, simulations were conducted to study the application of sub-Nyquist sensing and sparse recovery algorithms to the hybrid lidar-radar system. The sensing matrix was derived from the hybrid lidar-radar measurement methods, and the coherence of the sensing matrix was measured. Initially, bounds on the recoverable sparsity were determined through Monte-Carlo simulations, and the susceptibility to noise was investigated through similar simulations. A preliminary bench top experiment was conducted showing that the sub-Nyquist sensing and sparse recovery algorithms can be applied when using actual sampling hardware, and the experimental results were comparable to those produced by simulations. Throughout this section, the application of sub-Nyquist sensing and sparse recovery algorithms to the specific underwater ranging application will be explored, and the effect of the degraded visual environment will be studied.

7.1 Verification of the Sparse Approximation

The major underlying assumption in the discussed sub-Nyquist sensing is the notion of sparsity. This is the constraint that must be true for the different recovery algorithms to be successful. For the application to underwater lidar, the sparsity

stems from the highly directive nature of laser propagation. As discussed previously in chapter 1, the underwater optical channel causes scattering and spreading of the system’s point spread function, which leads to a broadening of the optical impulse response and the addition of a backscatter clutter component. These effects can degrade the inherent sparsity of the sensing system. In this section, the effect that the underwater channel has on the sparse approximation will be studied using an optical impulse response model.

7.1.1 Modeling the Underwater Hybrid Lidar-Radar System

To model the underwater impulse response, the complex mathematical radiative transfer equations must be solved given a number of different system variables [10]. These equations are solved with inputs such as system geometry, receiver characteristics, and the absorption and scattering properties of the underwater channel. While solving the radiative transfer equations is a complex problem, analytical techniques to approximate the solution have been developed, and tools have been created to model the underwater impulse response [87]. These tools are part of the iterative work by Zege et al. [88–90] where different approximations are made to estimate the effect that backscatter and forward scatter has on the optical impulse response. The output of the model is the number of photons incident on the receiver as a function of time given a delta function-like pulse of photons that have propagated to and from a given target of interest. Figure 7.1 shows the model outputs for both clean ($c_{532} = 0.1m^{-1}$) and turbid environments ($c_{532} = 2.0m^{-1}$). For the

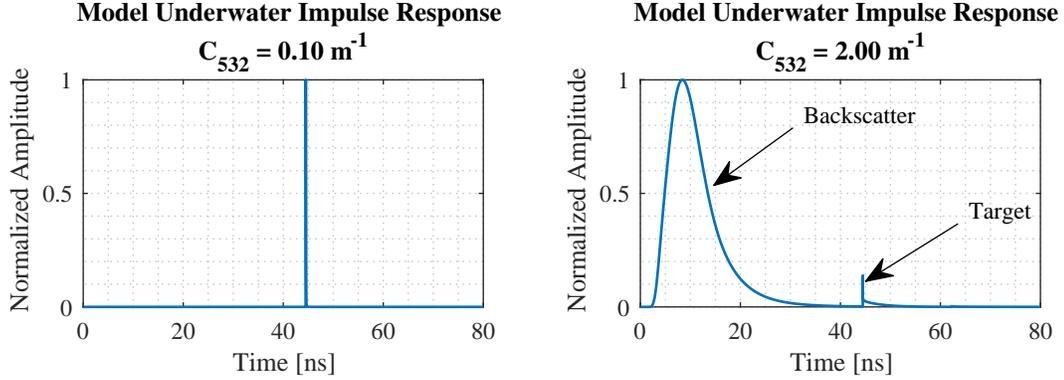


Figure 7.1: Example of the underwater channel impulse response model output for clean and turbid water.

clean environment, the delta like impulse response shows the inherent sparsity of the underwater lidar system. The modeled turbid water optical impulse response details how the effects of scattering in the backward and forward directions can affect the sparse approximation. However, as mentioned in chapter 1, the use of high frequency modulation and AC coupling, the effects of scatter is suppressed and the sparse approximation can be used.

Given the optical impulse response generated by the model, the transmission of wide-band waveforms can be simulated by convolving the output of the model with a given waveform of interest. For these simulations, FMCW waveforms introduced in the previous chapters are targeted. These waveforms can be generated with a given duration T_c and bandwidth ΔF . The optical to electrical conversion process can be simulated given the optical receiver characteristics [91], and the shot noise introduced by this process can be generated using the known photon counts and Poisson statistics. The DC component of the modeled analog waveform can be

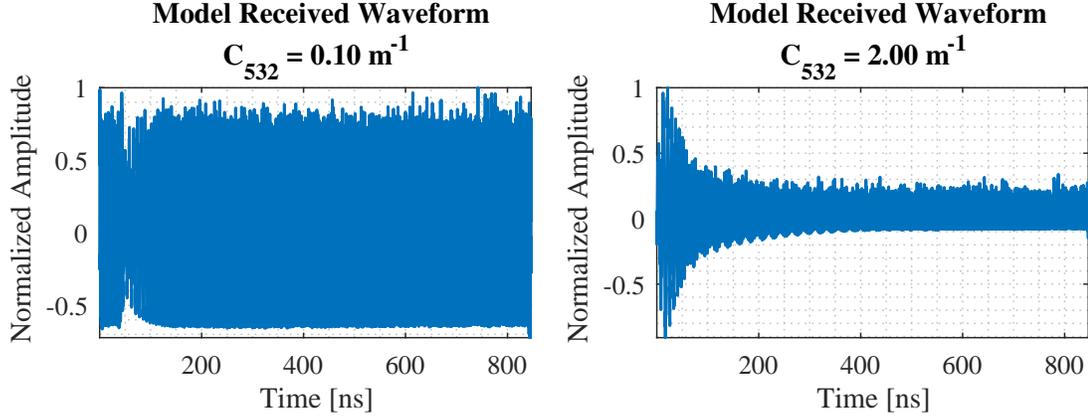


Figure 7.2: Example of the underwater channel impulse response model output for clean and turbid water.

removed to simulate an AC coupled receiver. Figure 7.2 shows the modeled AC coupled return signal prior to digitization for both clean and turbid environments. The increased noise due to both the backscatter response and the decreased amplitude of the target return signal is the most notable difference between the modeled clean and turbid water scenarios. Sub-Nyquist sensing can be simulated by re-sampling the received waveform at a given rate, and the sparse recovery OMP algorithm highlighted in Chapter 4 can be tested as a function of the underwater optical properties. The functional block diagram for the simulation steps can be seen in figure 7.3

7.1.2 Simulation Results

The goal of these simulations is to study the relationship between the underwater channel characteristics, the optical detector, and the sub-Nyquist sensing framework. To accomplish this goal, the impulse response for the underwater chan-

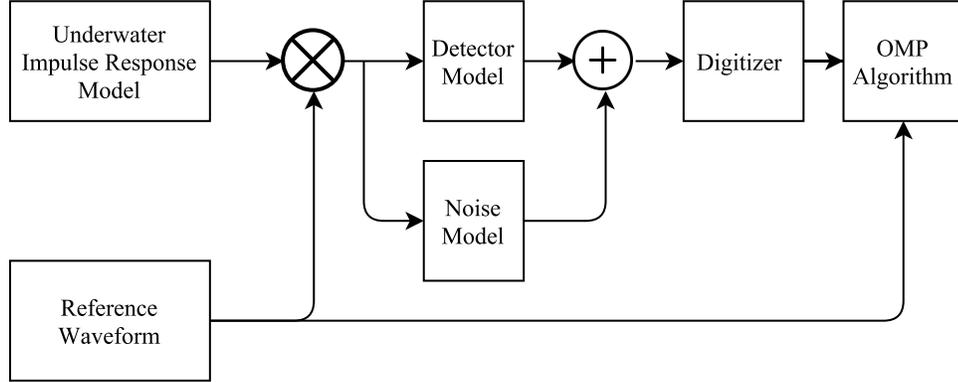


Figure 7.3: Block diagram for modeling the underwater hybrid lidar-radar sub-Nyquist sensing system.

nel can be modeled given a diffuse flat gray target situated a nominal distance away from the transmitter/receiver platform. In these simulations, transmitted beam and receiver field of view are centered on the target, and different water clarities are simulated by varying the beam attenuation coefficient, $c_{532}[m^{-1}]$.

A FMCW waveform with a variable duration T_c and constant bandwidth Δf is then convolved with the modeled impulse response, and the optical detector as well as the inherent noise of the optical detection process is simulated. The simulated analog signal is then sampled at a given rate F_D to produce a total number of $M = F_D \times T_c$ measurements. The OMP recovery algorithm is implemented using an assumed sparsity of one, and the dictionary is generated from the known transmitted chirp waveform and the given sampling rate F_D . Given typical Nyquist sampling rules, the simulated return signal should be sampled at a rate $F_{Nyq} = 2 \times \Delta f$, and would require $N = T_c \times F_{Nyq}$ samples. The Sub-Nyquist system is simulated using the given waveform duration and number of measurements leading to a down-

sampling parameter of $\delta = M/N$. For comparison, the matched filter detection process was also simulated to provide a baseline of performance that could be expected as a function of water conditions.

To quantify performance, the average error between the range measured from the OMP algorithms output, and the actual target range is calculated over multiple trials and plotted as a function of attenuation lengths $C_{532} \times Z$ where Z is the range to the target. The number of measurements and chirp duration was varied, and the results of these simulations are shown in figure 7.4 for all the simulated water conditions, waveform durations, and various levels of down-sampling.

When processed using Nyquist sampling and a matched filter detector (black dashed line in figure 7.4) the error is minimal until approximately 9.5 attenuation lengths. The matched filter data in figure 7.4 also shows that increasing the duration does not improve the detection as a function of attenuation length. At the point where the error begins to increase ($C_{532} \times Z = 9.5$), the residual noise from the backscatter signal is much larger than the target return, and the detected range error is systematically biased. These results are similar to those found in previously conducted studies [78,92,93] that investigated the application of different methods to suppress the adverse effects of scattering.

It is also clear from figure 7.4 that the effects of scattering play a role in the sparse recovery of a target's range. For the case when $\delta \approx 1$, there is no down-sampling, and the OMP algorithm is being used on the full Nyquist rate waveform. Due to the nature of the OMP algorithm as discussed in Chapter 4, one would expect the outcome to be similar to the matched filter result. This is

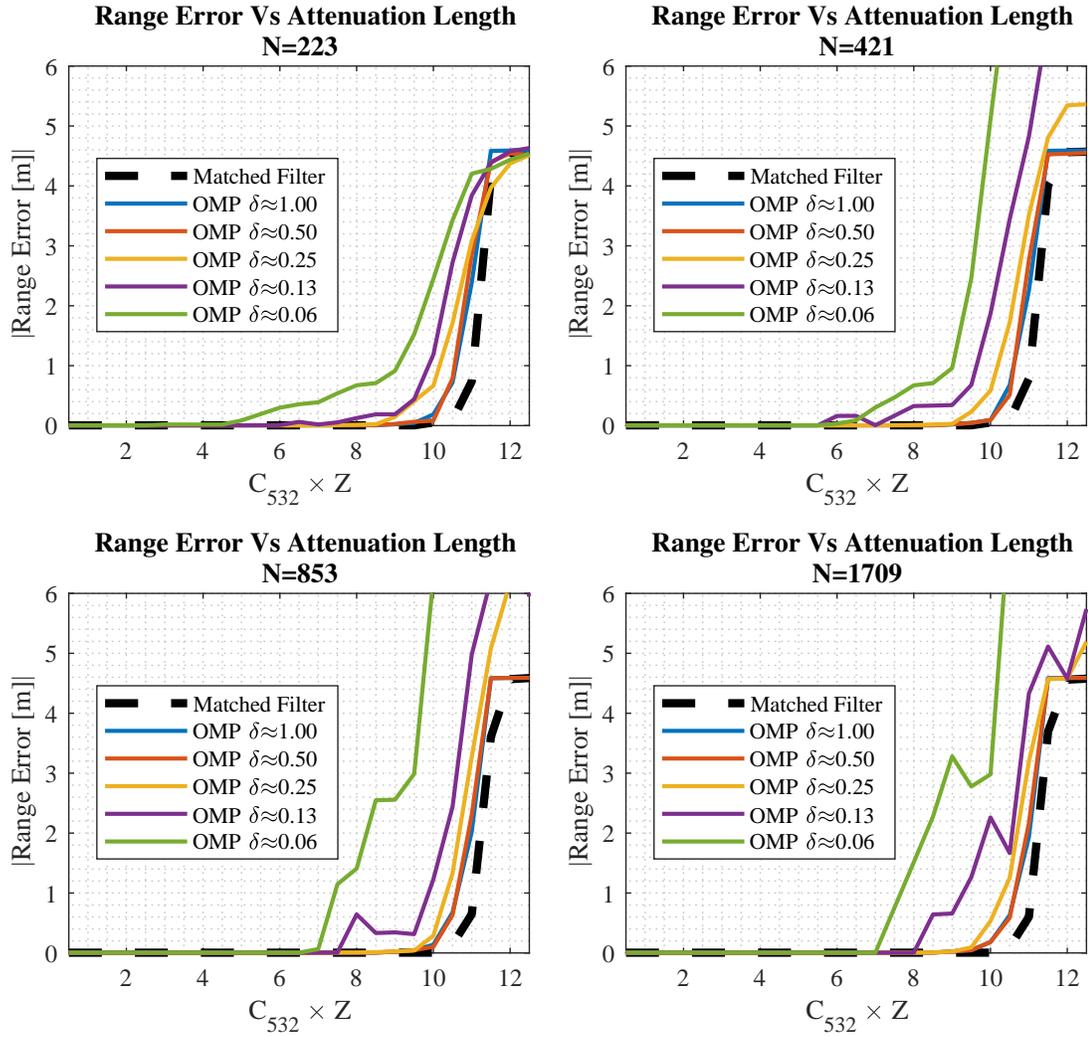


Figure 7.4: Simulation results for the modeled sub-Nyquist underwater hybrid lidar-radar system. The range error was calculated for various test parameters and plotted as a function of attenuation length.

because both methods are based on finding the positions of maximum correlation. This expectation is substantiated in the modeled results and the OMP algorithm performs similarly to the matched filter. When the down-sampling parameter is set to $\delta \approx 0.50$, similar performance to the matched filter is also observed, especially as the waveform duration is increased. Both the $\delta \approx 1$ and $\delta \approx 0.50$ cases display an increase in range error at 9.5 attenuation lengths that is similar to the matched filter results. This implies that at this point, for the given simulated system, the effects of scatter have become so severe that the target can no longer be recovered.

For higher levels of down-sampling, it is shown that measurement noise is affecting the performance more severely than the effects of the channel. This is evident given that the recovered range error can be reduced by increasing the waveform duration. This result can also be substantiated from the simulation and experimental observations in chapter 5 where it was shown that in the presence of noise, sub-Nyquist sensing and sparse recovery through the OMP algorithm have similar performance to the matched filter as long as the down-sampling parameter and waveform duration lead to a sufficient number of measurements. To highlight the results from these simulations, the highest level of scattering and noise that led to minimal range error can be tracked and presented as a function of the waveform duration and down-sampling parameter. These results are presented in table 7.1

Table 7.1 shows that as N increases, it is possible to recover targets in environments with high levels of scattering and resulting receiver noise. Once the effects of scattering become too severe, all detection methods fail to recover the target irrespective of the waveform duration and down-sampling parameter. This result

Table 7.1: Summary of modeling results detailing the max number of attenuation lengths caused by scattering that results in successful detection of the underwater target in the modeled environment. Results presented for both the matched filter and the OMP algorithms at different levels of down-sampling $\delta = m/N$

Duration N	Max Detectable $C_{532} \times Z$ [A.L]					
	Matched Filter	$\delta \approx 1$	$\delta \approx 0.50$	$\delta \approx 0.25$	$\delta \approx 0.13$	$\delta \approx 0.06$
223	9.5	8.5	8.0	7.5	6.0	2.5
421	9.5	9.0	8.5	7.5	7.0	5.5
853	9.5	9.5	9.0	8.0	7.5	6.5
1709	9.5	9.5	9.5	8.5	8.0	7.0

provides the initial verification that the sparse approximation is valid, and that sub-Nyquist sensing and sparse recovery can be applied to the hybrid lidar-radar system as long as the amount scattering has not reduced the signal amplitude below a certain SNR level. Through this modeling effort it has also been shown that the sparse recovery techniques can approach the performance of the matched filter detector as long as the waveform duration is long enough to produce the necessary number of measurements in the presence of noise.

7.2 Experimental Verification

The previous sections detailed simulations that were conducted, and how the underwater channel was modeled. This modeling and simulation effort was used to determine if the sparse approximation is valid for underwater lidar systems, and to measure the performance limits for sparse recovery. In this section, the implementation of the sub-Nyquist sensing and sparse recovery methods will be detailed, and an experiment that was designed to test the methods will be discussed. This experiment

was designed to investigate the real world implementation of sub-Nyquist sensing and sparse recovery, and to test the performance of the methods in environments with different levels of scattering. First, the experimental setup and equipment will then be described and the measured values will be detailed. Processing and performance metrics will be highlighted, and the results will be discussed.

7.2.1 Experimental Design, Equipment, and Details

The goal of these experiments is to test the application of sub-Nyquist sensing and sparse recovery as applied to hybrid lidar-radar systems. As discussed in chapter 1, the targeted system is an active ranging system that utilizes a modulated laser transmitter and an optical receiver co-located on the same platform. The transmitter and receiver were separated by 30cm, and the optical receiver had a field of view of 4.8 degrees. These settings were used to obtain a transmitter receiver overlap that covered the entire range of target positions. After the PMT, a bias T was used to AC couple the electrical signal, and a analog filter low pass filter was used to reduce noise prior to digitization. After this analog signal conditioning, a digitizer was used to sample the received electrical signal. The system was placed near a window of a circular test tank, and the transmitted beam/receiver field of view were directed at a flat gray target suspended in the tank and attached to a translation stage above the tank. The translation stage has a total variable range of 2.5 meters and was placed approximately 3.88 meters away from the transmitter and receiver platform. The turbidity of the underwater channel was adjusted using a known

scattering agent (Equate antacid), and the scattering coefficient was measured using a transmissometer. At each water clarity, the target was swept through the entire translation stage range, and the received signal was digitized at multiple sampling rates.

The FMCW source used to modulate the transmitter is an Analog Devices AD9914 direct digital synthesizer (DDS) board [94]. This board was used to produce a FMCW signal with a bandwidth of 1 GHz and a duration of $T_c \approx 2\mu s$. The laser transmitter is a custom 532nm CW laser source manufactured by AdvR, Inc. The laser consists of an infrared laser diode at 1064nm which is modulated by a Mach-Zehnder modulator, amplified via fiber amplification stages, and frequency doubled by a second harmonic generating crystal to 532nm. The optical receiver is comprised of collection optics, and a 1Ghz bandwidth 5-stage Photo-Multiplier Tube (PMT) from Photonis [95]. A converging $f/2$ lens and variable aperture was used to control the receiver field of view, and a 532nm interference filter was used to reduce the collection of ambient light. A 12 bit, 2 channel digitizer with a 2 GHz analog bandwidth from SP-Devices [96] was used to directly sample the RF waveform. The digitizer has the option for variable sampling rates, and various clocks were used to generate measurements at rates both over and under the Nyquist sampling rate. The experimental block diagram can be seen in figure 7.5.

After digitization, the sampled waveforms were stored and the detection algorithms were applied. To test performance, the matched filter detector was implemented on the Nyquist sampled data and used as a baseline comparison for the results generated using the OMP algorithm. In this experiment, it was assumed

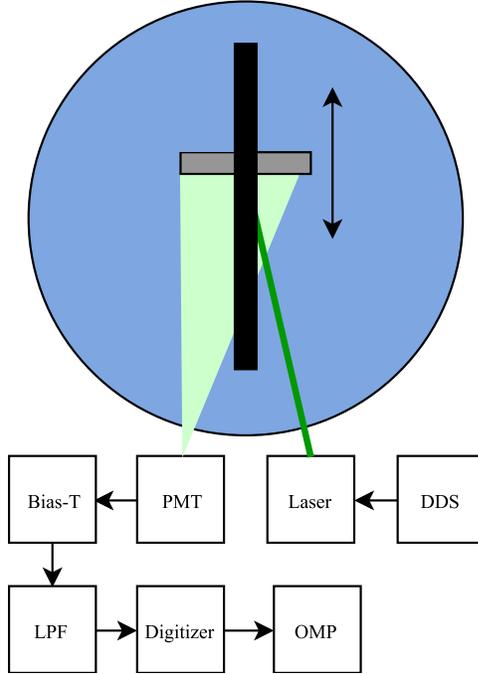


Figure 7.5: Block diagram of experimental setup

that there was only a single reflecting object so the OMP algorithm was implemented with a sparse estimate of $K_{est} = 1$. After the sparse estimate was determined, its index was used to calculate the range given the sample rate of the stored sensing matrix. For each water clarity and stage position, the measured range was reported and compared to the known stage position. For these experiments, the level of under sampling is defined as $\delta = \frac{F_S}{F_N}$, where F_N is the transmitted waveform Nyquist Rate, and F_S is the system sampling rate. For $\delta \geq 1$, this represents signals that have been sampled at or above the Nyquist rate, and for $\delta < 1$ the signals have been sampled at sub-Nyquist rates. Figure 7.6 shows the digitizer output at various sample rates. As the sample rate is reduced, the effects of aliasing are evident in the sub-Nyquist sampled waveforms.

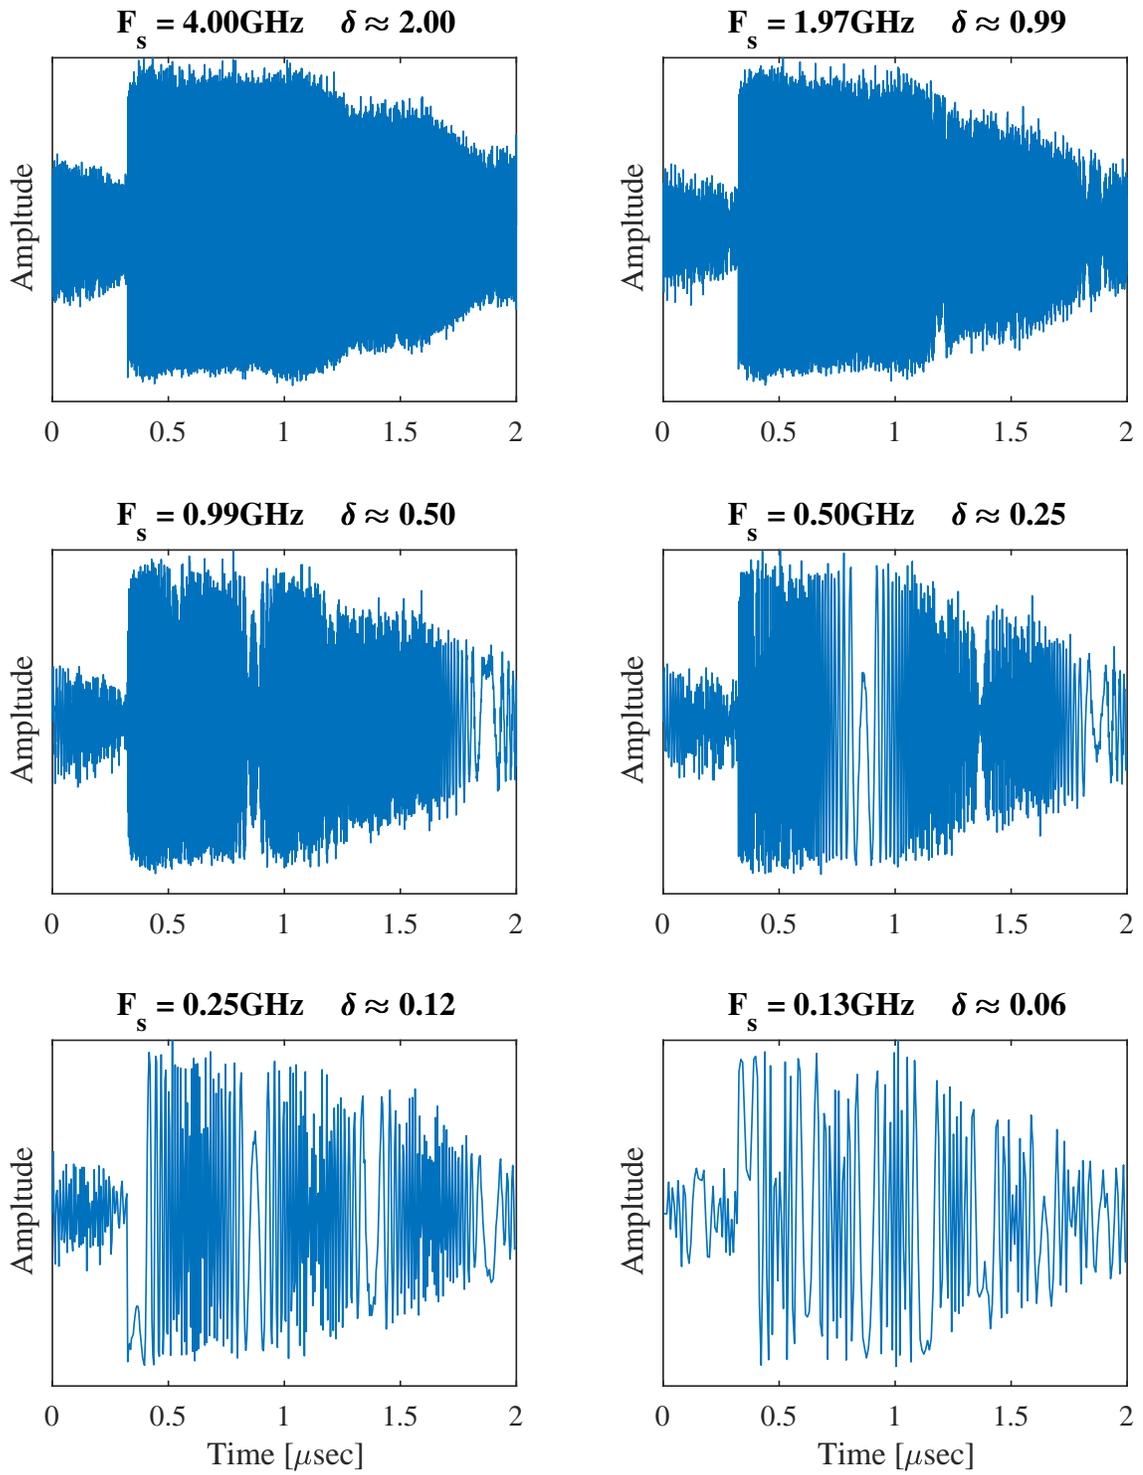


Figure 7.6: Example digitizer output waveforms at various sampling rates.

7.2.2 Clean Water Baseline Measurements

Initial testing of the sub-Nyquist sensing and sparse recovery occurred in clean water environments, $C_{532} = 0.14m^{-1}$. In order to implement the OMP algorithm for sparse recovery, the chirp dictionary and sampling matrix must be generated. This was accomplished initially as described in chapter 3 by using the known parameters of the transmitted FMCW signal and the known sampling rate of the digitizer. Results from this initial testing can be seen in figure 7.7 where the measured range is plotted as a function of the actual target range for the matched filter detector as well as the OMP algorithm at various levels of down-sampling.

These initial results show significant error between the target's actual range and the detected range when using sparse recovery methods at sub-Nyquist sampling rates. This error can be attributed to dictionary mismatch that arises from the non-ideal signal generation and detection process. The OMP algorithm relies on a dictionary generated from an ideal linearly frequency modulated signal. In practice, the output of the DDS board is not an ideal chirp. For example, the analog waveform is generated from a stepped frequency ramp vs the assumed continuous linear ramp, and the output power is not constant across all bands. An additional cause of dictionary mismatch is that the measured signal has been affected by the frequency response of the PMT, as well as the analog components in the receiver chain. These facts are highlighted in figure 7.8 where the magnitude spectrum for the waveform used to generate the sensing matrix is compared to the magnitude spectrum of the signal digitized from the PMT. Here it can be seen that there is a substantial

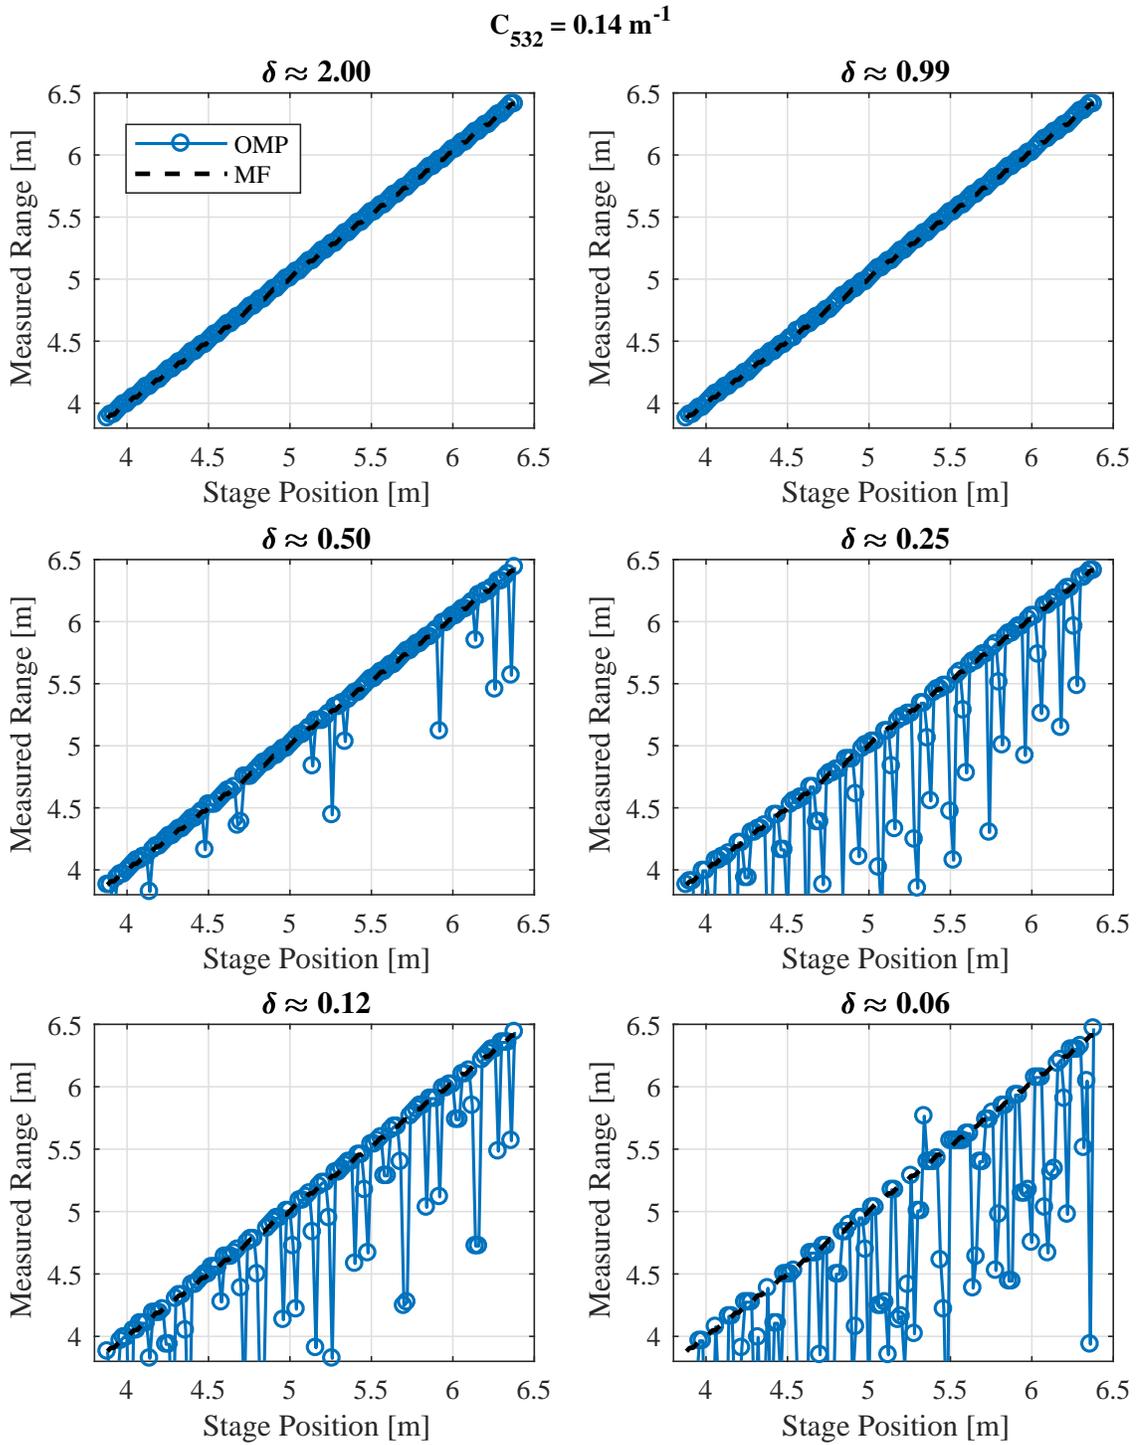


Figure 7.7: Initial experimental results of the sub-Nyquist sensing and sparse recovery of a target range situated in underwater environments. Errors in detection due to dictionary mismatch.

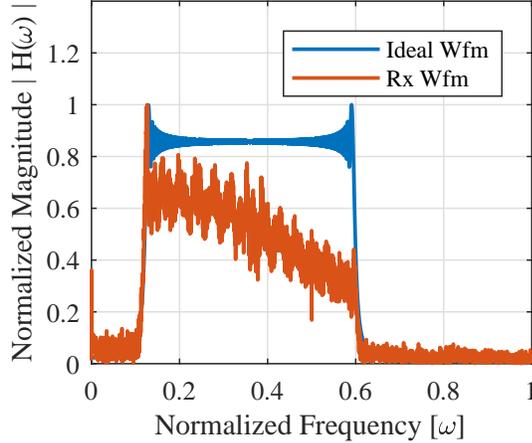


Figure 7.8: Magnitude Spectrum of the waveform used to generate the sensing matrix compared to the magnitude spectrum of the PMT signal.

discrepancy between the two waveforms.

To alleviate the adverse effects of dictionary mismatch, the signal from the laser transmitter was redirected into the PMT receiver and a reference copy was sampled and stored at $F_S \approx 2 \times F_{Nyq}$. This new reference waveform was then used as the basis function to generate the column wise time shifts for the chirp dictionary, and the sensing matrix was created given the chosen sampling rate of the digitizer. Figure 7.9 presents the detected target range as a function of the stage position using the new sensing matrix for all the various down-sampling parameters used. These results confirm that the previous errors were in fact due to dictionary mismatch, and that by using a stored reference copy to generate the sensing matrix, the adverse effects of dictionary mismatch can be alleviated. Here it can also be seen that the measured range using sub-Nyquist sensing and sparse recovery is effectively equal to the matched filter detector and is not adversely affected by the lower sampling

rates.

7.2.3 Experimental Results for Turbid Environments

In the previous section, it was shown that sub-Nyquist sensing and sparse recovery algorithms can be used to measure the range to an underwater target. This was demonstrated using the described hybrid lidar-radar system, uniform sampling at sub-Nyquist rates, and the OMP algorithm to recover sparse vectors. While initial dictionary miss-match caused failed detection even in high SNR and clean water environments, this effect was mitigated by using a stored reference copy of the known transmitted waveform. In this section, the effect of turbid underwater environments will be explored.

The turbidity of the underwater environment was increased incrementally until errors in detection were observed. Figure 7.10, shows the measured range results as a function of actual target range and down-sampling parameter for the case when the measured extinction coefficient was equal to $C_{532} = 1.30m^{-1}$. At this given environmental condition, the reduced performance at lower sampling rates is caused by the limited SNR at further stand-off ranges. For example, the reliability of the range measurements begins to be limited at approximately 7.4 attenuation lengths for the lowest sampling rate. When the sampling rate is higher and $\delta \approx 0.25$, it is still possible to recover the target range out to approximately 8 attenuation lengths. Therefore, the sparse approximation is still valid and detection is possible with enough measurements.

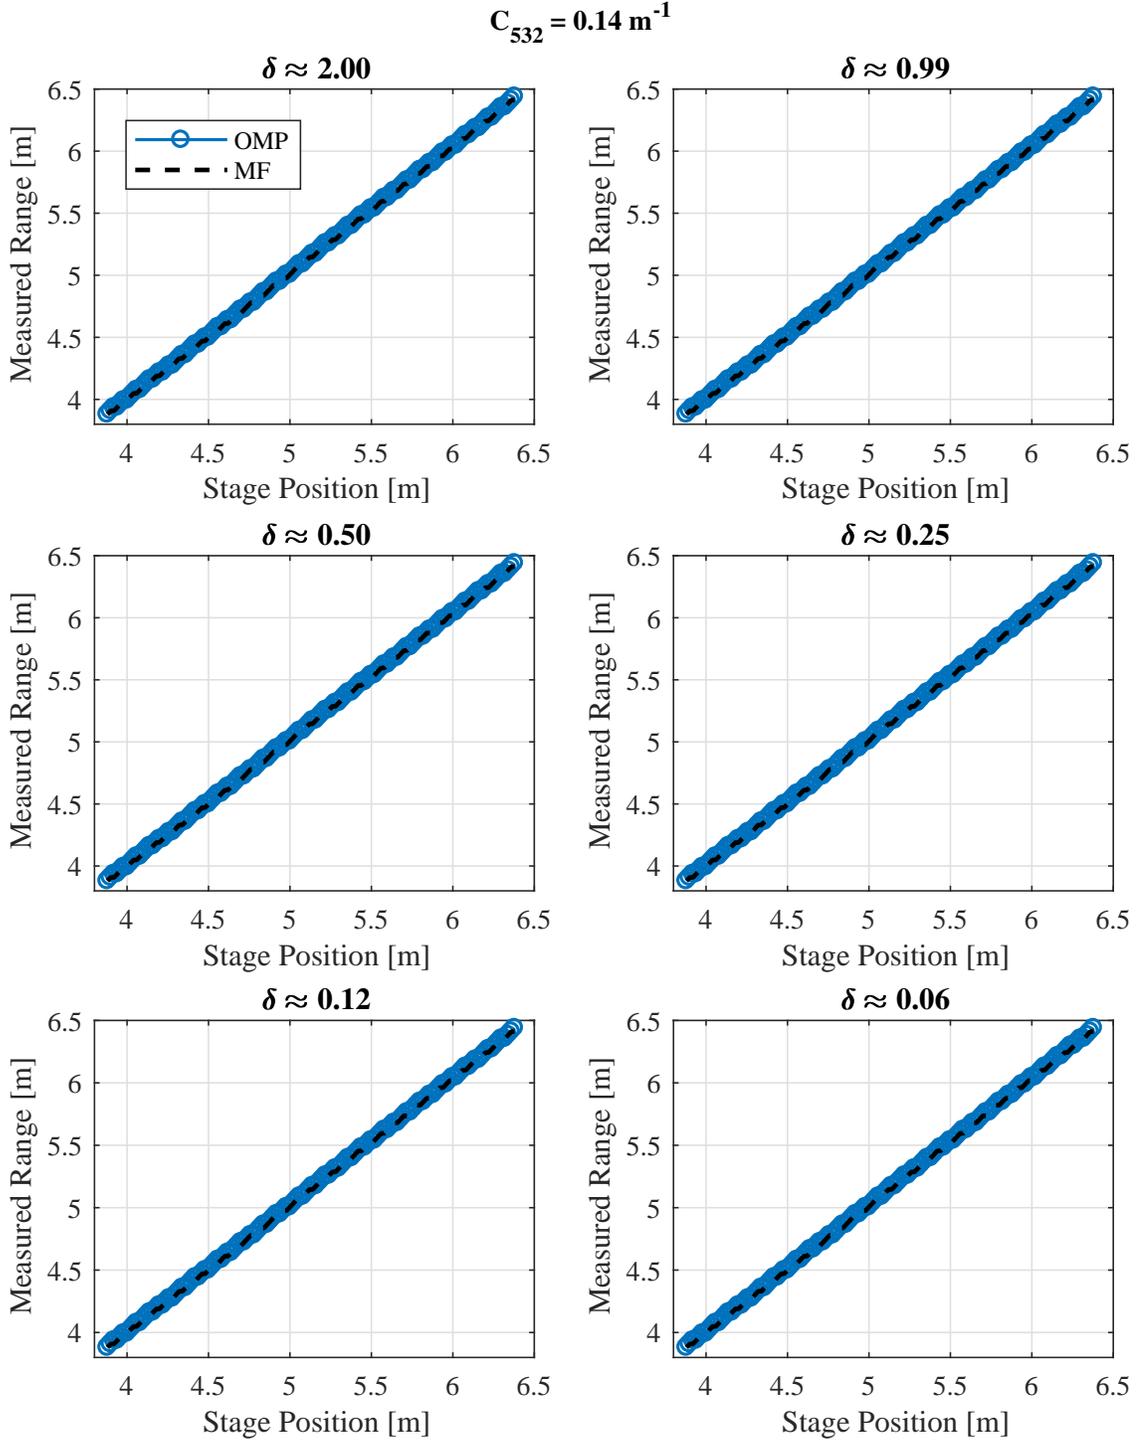


Figure 7.9: Initial experimental results of the sub-Nyquist sensing and sparse recovery using a sensing matrix that is derived from a Nyquist sampled version of the PMT reference signal. Effects of dictionary mismatch have been suppressed

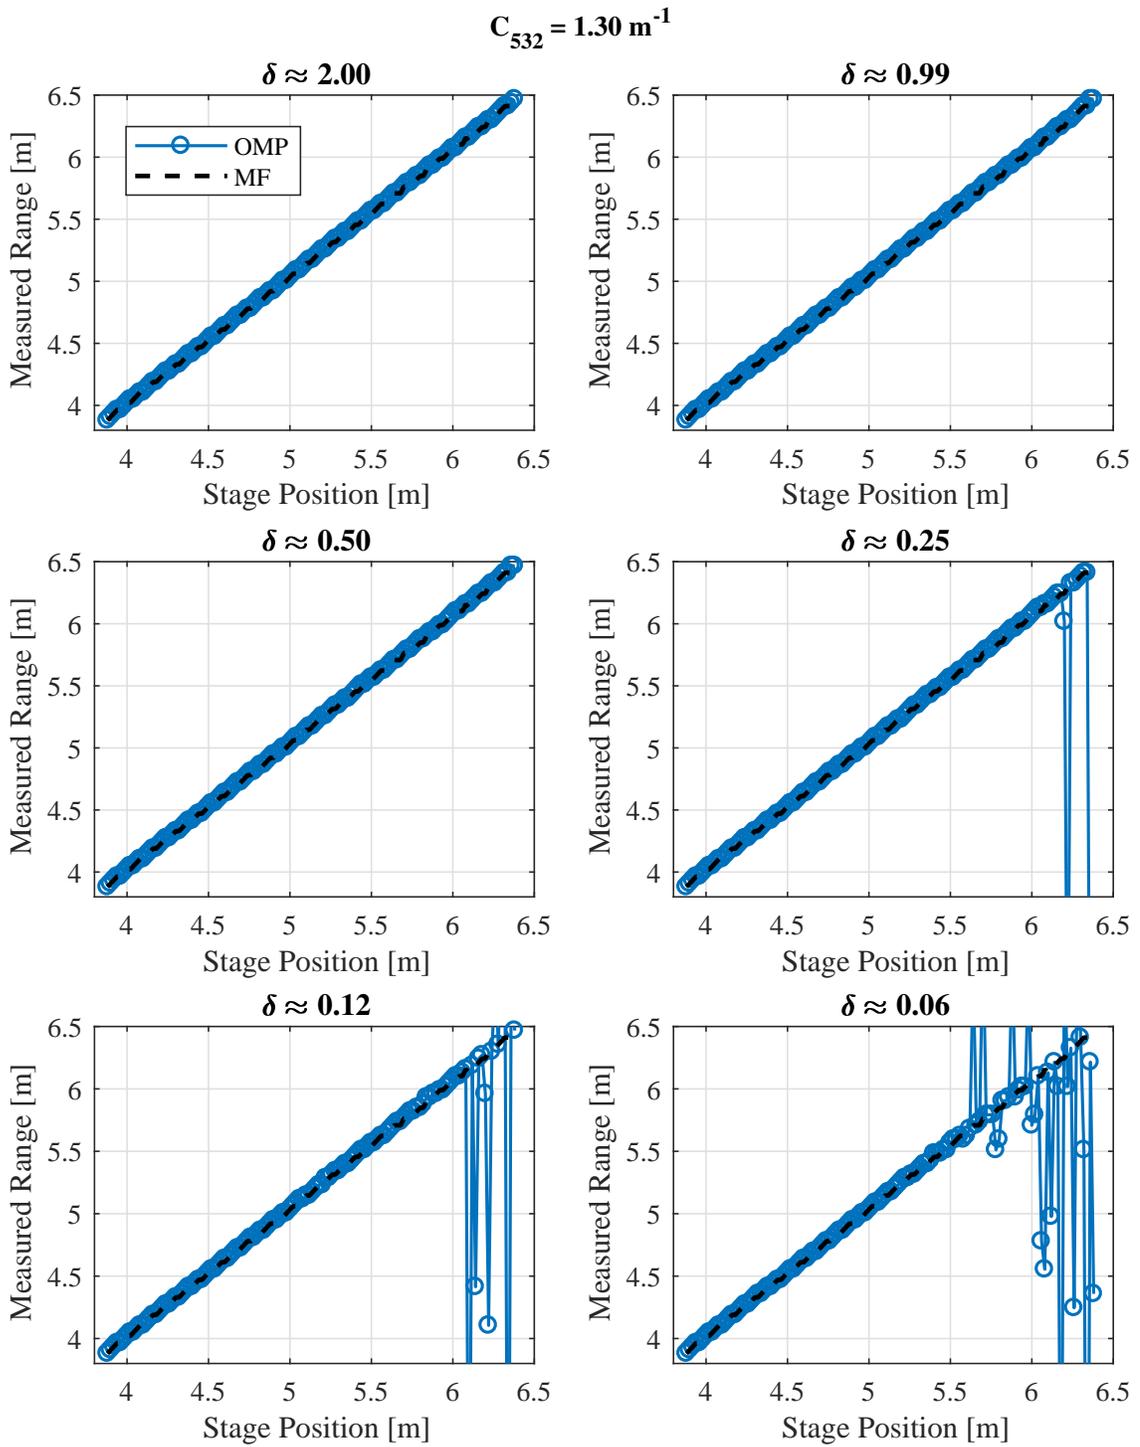


Figure 7.10: Experimental ranging results for $C_{532} = 1.30 \text{ m}^{-1}$. Detection error is evident when sampling at rates less than Nyquist.

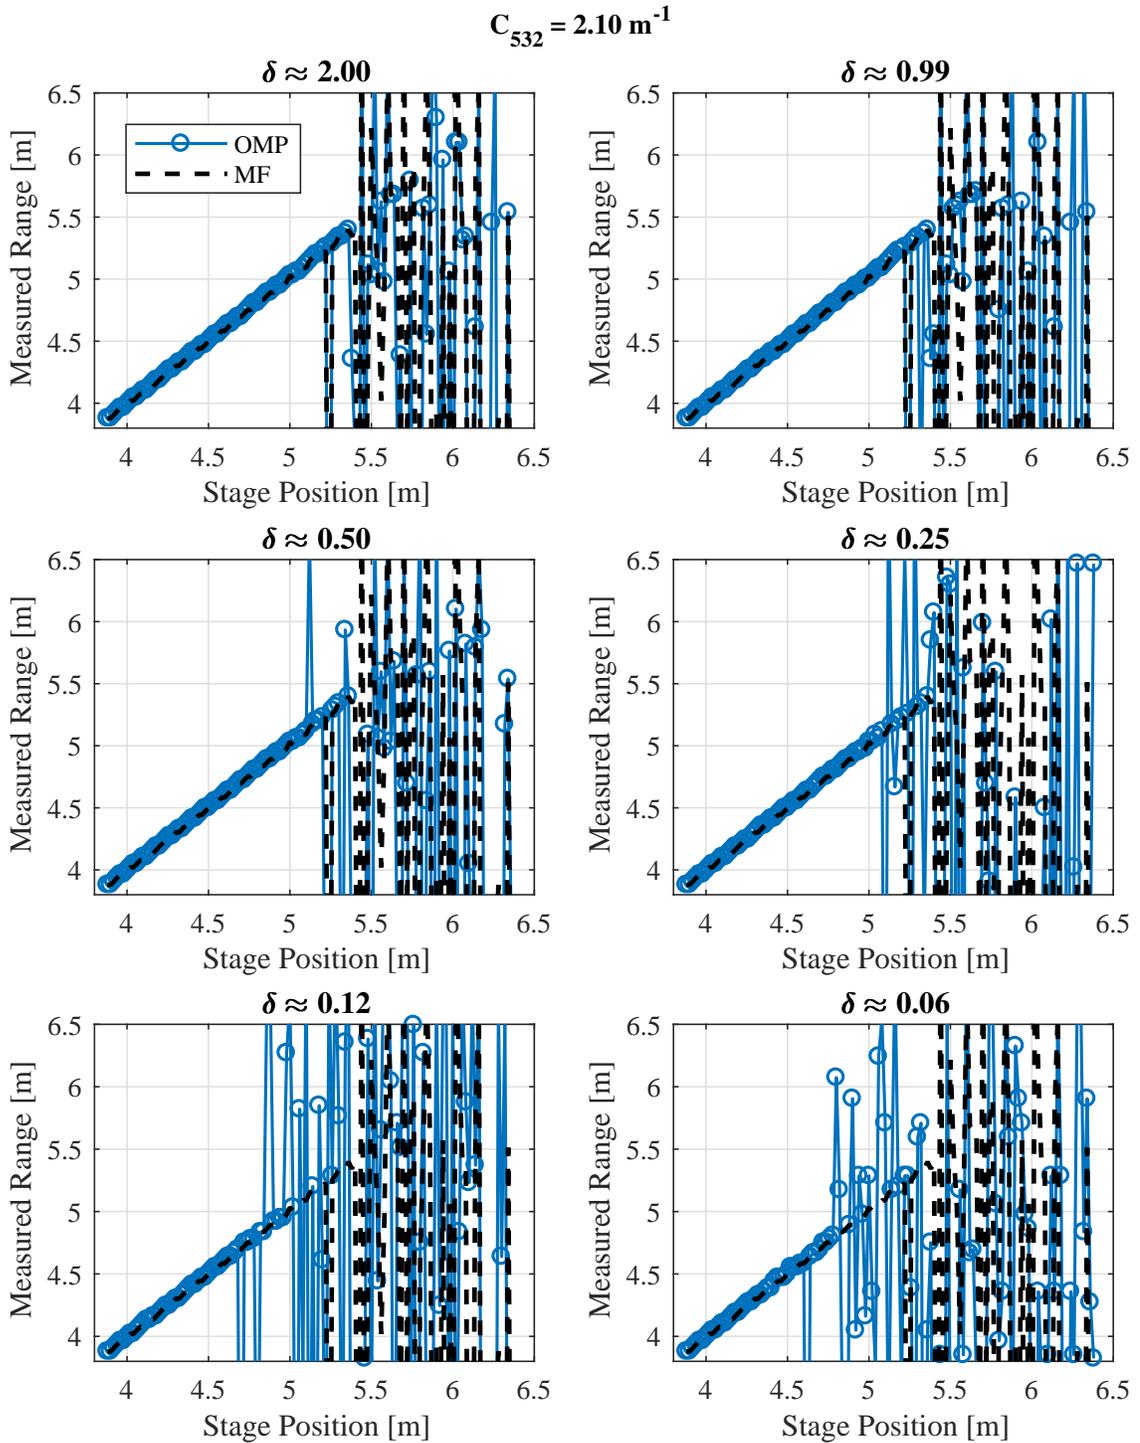


Figure 7.11: Experimental ranging results for $C_{532} = 2.10 \text{ m}^{-1}$. While the correct range measurements are reported for closer in targets, all detection methods display range errors at further target positions.

Table 7.2: Experimental results showing max recoverable attenuation lengths for the matched filter detector, and the OMP algorithm given various system sampling rates

Matched Filter	OMP					
	$\delta \approx 2.00$	$\delta \approx 0.99$	$\delta \approx 0.50$	$\delta \approx 0.25$	$\delta \approx 0.12$	$\delta \approx 0.06$
11	11	11	10.7	10.7	9.8	9.6

To further test the system performance limits, the extinction coefficient was increased until the matched filter detector also began to produce errant range measurements. This occurred when $C_{532} = 2.10m^{-1}$, and the results can be seen in figure 7.11. These experimental results show that when using a matched filter detector, it is possible to accurately measure the range to a target up to approximately 11 attenuation lengths (or 5.24m for $c=2.1/m$). When the signal is sampled at Nyquist, the OMP detection methods perform the same as the matched filter, but they degrade in performance as the sampling rate is reduced. These results prove that detection is still possible even in highly scattering environments. For example, when the sample rate is approximately 6% of the Nyquist rate, it is possible to accurately detect a target's range out to ≈ 9.6 attenuation lengths. The experimental results are summarized in table 7.2 where the max number of attenuation lengths that still resulted in a target recovery is presented as a function of the detection method and system sampling rate $\delta = F_S/F_N$. This was measured by finding the first stage position where the error between the target's range and the measured range was greater than 10cm. This target position was then used to calculate the max number of recoverable attenuation lengths.

Chapter 8: Conclusion

This research effort met the objective of providing evidence that sub-Nyquist rate digitizers and sparse recovery algorithms can be applied to ranging applications utilizing wide-band hybrid lidar radar technologies. Furthermore, it was discussed that the results from this work can be extended to other types of wide-band active ranging methods. This concept was substantiated using both simulated and experimental tests to measure the performance of sub-Nyquist sensing as a function of various system and environmental parameters. The system parameters include waveform characteristics such as modulation type and duration, digitizer sampling rates, sparse recovery algorithm choice, and sensing matrix resolution. Environmental parameters include the number of reflecting objects, SNR, and underwater channel characteristics. When applicable, the traditional matched filter method was used as a baseline to determine what advantages the sub-Nyquist sensing and sparse recovery method have over conventional processing techniques.

8.1 Summary

In chapter 3, it was shown that sensing matrices derived from wide-band waveforms and uniform sub-Nyquist samplers satisfies the Mutual coherence condition,

and theoretical predictions on the recoverable support were provided. It was also shown that by choosing system parameters such that the number of measurements M , and number of Nyquist samples N are relatively prime, the coherence is strictly less than one for the linear FMCW basis waveform. Results from testing various sparse recovery algorithms were presented in chapter 4. It was shown that both convex optimization and greedy algorithms are capable of exactly recovering sparse vectors from measurements made given sensing frameworks derived from wide-band FMCW waveforms. The trade-off between computation time and detection performance was presented, and it was shown that the recoverable support was much greater than that predicted by the coherence bound. Empirically it was found that when using the OMP algorithm to recover measurements made using the sub-Nyquist sensing framework and wideband linear FMCW waveforms, the relationship between the number of measurements and recoverable support was $M \approx \mathcal{O}(\sqrt{\frac{5}{3}}\sqrt{Nk})$

In chapter 5, it was shown that the sub-Nyquist sensing system and OMP recovery algorithm could still recover sparse vectors even in the presence of noise. When compared to the matched filter detection, it was shown that the sparse recovery techniques had the same noise performance as long as the downsampling ratio, $\delta = m/N$, was greater than 0.6. It was also shown that given a fixed system sampling rate, the noise performance of the sub-Nyquist sensing system could be improved by increasing the duration of the waveform. In chapter 5 the initial experimental evidence was also provided showing that these methods can be applied to real hardware under low SNR conditions. The experimental results agreed closely

with the previous simulations, confirming the extended simulation results. Chapter 6 presented results from additional simulations in order to provide insight into the resolution of the sub-Nyquist sensing system and the sparse recovery algorithms. It was shown that by increasing the column wise resolution of the sensing matrix, it was possible to increase the resolution of the sparse estimate. This resolution enhancement technique was applied to two common resolution/detection problems: two closely spaced targets, and one highly reflective target next to a weaker target.

Lastly in chapter 7, simulation and experimental results were provided showing that the range to a target can be measured in various water conditions while using sample rates approximately 15 times lower than the Nyquist rate. Initially the underwater channel was modeled for various levels of scattering to determine when the sparse approximation would no longer be valid. The optical detection process and digitizers with various sample rates were simulated, and the OMP algorithm was applied. Through these modeling and simulation efforts, it was shown that as the waveform duration increased, the sparse recovery detection performance would approach that of the Nyquist sampled matched filter. This result provided initial evidence that the sparse approximation maintained validity up to approximately 9.5 attenuation lengths, and that these methods could be applied to the underwater channel. Results from experimental tank tests were also presented in chapter 7 and it was shown that after addressing dictionary mismatch, sub-Nyquist sensing and sparse recovery could be applied to the hybrid lidar-radar based ranging system. Sampling rates approximately 16 time lower than the Nyquist rate were successfully used and were capable of detecting targets at 9.6 attenuation lengths.

8.2 Future Research Directions

Throughout this research effort, the only constraint imposed on the L_1 minimization problem was that the solution should be sparse and the error between the measurements and the projected sparse estimate should be minimized. In future work, it would be advantageous to investigate different constraints on the L_1 minimization problem. One powerful constraint that could be studied is the non-negativity constraint on the sparse solution [97, 98]. L_1 minimization with non-negativity constraints is shown in equation (8.1).

$$\min_s \|s\|_1 \text{ subject to: } \|\Theta s - y\| \leq \epsilon \text{ and } |s| \geq 0 \quad (8.1)$$

This constraint is inherently applicable to the hybrid lidar-radar sensing framework where the recovered sparse vector represents the optical impulse response given a certain propagation distance and the reflective nature of the target. Targets of interest are relatively flat diffuse reflectors which implies that a negative reflection coefficients would not be encountered in practice. Future work will look to investigate how adding the non-negativity constraint would effect the recovery performance of the L_1 minimization methods in terms of max support recovery and noise susceptibility. Within the field of compressive sensing, this constraint has proven useful in various applications [99, 100] and fast algorithms have been developed to solve the L_1 minimization problem [101]. However, it remains to be seen how it would effect the sparse recovery of wide-band signals given uniform sub-Nyquist sampling.

Another future area for research would be altering the sensing framework

slightly. Currently, the described sensing framework works by sampling the return waveform at a given rate. Once all the samples have been collected it attempts to recover the sparse vector using some L_1 minimization algorithm. This implies that sparse estimates and thus range measurements are generated at an approximate rate of $\frac{1}{T_c}$. Often multiple periods are measured before updating the range estimate, and averaging is used to improve SNR and detection performance. In the altered sensing framework, one could imagine taking the individually measured return signals of length $M = T_c \times F_s$, and creating a $M \times L$ matrix of measurements Y , where L is the number of full waveforms that have been sampled. In this framework, the sensing matrix is the same but the sparse vector is now a matrix where the columns represent the number of waveform records. This framework is known as a Multiple Measurement Vector model and can be solved using various algorithms which can exploit the joint sparsity or inter-correlations of the MMV's [102–105]. In the context of the hybrid lidar-radar ranging system, the correlation between the columns of the measurement matrix would be dependent on target movement speed, and the sparse approximation would be the location of the various targets at the given time the waveform records were taken. Under this framework, the effect MMV's have on recovery performance, and SNR remains to be studied.

Bibliography

- [1] Linda J Mullen and V Michael Contarino. Hybrid lidar-radar: seeing through the scatter. *IEEE Microwave magazine*, 1(3):42–48, 2000.
- [2] Emmanuel Candes and Justin Romberg. l1-magic: Recovery of sparse signals via convex programming. *URL: www.acm.caltech.edu/l1magic/downloads/l1magic.pdf*, 4:14, 2005.
- [3] J. L. Sutton. Underwater acoustic imaging. *Proceedings of the IEEE*, 67(4):554–566, April 1979.
- [4] L. Henriksen. Real-time underwater object detection based on an electrically scanned high-resolution sonar. In *Autonomous Underwater Vehicle Technology*, pages 99–104, Jul 1994.
- [5] Takuso Sato, Mitsuhiro Ueda, and Sadanao Fukuda. Synthetic aperture sonar. *The Journal of the Acoustical Society of America*, 54(3):799–802, 1973.
- [6] Mike DeLong. Underwater laser imaging system (uwlis). Technical report, Lawrence Livermore National Lab., CA (United States), 1994.
- [7] Georges R Fournier, Deni Bonnier, J Luc Forand, and Paul W Pace. Range-gated underwater laser imaging system. *Optical Engineering*, 32(9):2185–2190, 1993.
- [8] Raymond C Smith and Karen S Baker. Optical properties of the clearest natural waters (200–800 nm). *Applied optics*, 20(2):177–184, 1981.
- [9] Nils Gunnar Jerlov. *Marine optics*, volume 14. Elsevier, 1976.
- [10] CD Mobley. Radiative transfer in the ocean. *Encyclopedia of ocean sciences*, pages 2321–2330, 2001.
- [11] Curtis D Mobley. *Light and water: radiative transfer in natural waters*. Academic press, 1994.

- [12] John R Apel. *Principles of ocean physics*, volume 38. Academic Press, 1987.
- [13] B. Cochenour, L. Mullen, and J. Muth. Temporal response of the underwater optical channel for high-bandwidth wireless laser communications. *IEEE Journal of Oceanic Engineering*, 38(4):730–742, Oct 2013.
- [14] Linda Mullen, Robert Lee, and Justin Nash. Digital passband processing of wideband-modulated optical signals for enhanced underwater imaging. *Applied optics*, 55(31):C18–C24, 2016.
- [15] EA McLean, HR Burris, and MP Strand. Short-pulse range-gated optical imaging in turbid water. *Applied optics*, 34(21):4343–4351, 1995.
- [16] Linda Mullen, Alan Laux, Brian Concannon, Eleonora P Zege, Iosif L Katsev, and Alexander S Prikhach. Amplitude-modulated laser imager. *Applied optics*, 43(19):3874–3892, 2004.
- [17] Robert W Lee, Alan Laux, and Linda J Mullen. Hybrid technique for enhanced optical ranging in turbid water environments. *Optical Engineering*, 53(5):051404–051404, 2014.
- [18] David W Illig, Luke K Rumbaugh, Mahesh K Banavar, Erik M Bollt, and William D Jemison. Backscatter suppression via blind signal separation for a 532 nm underwater chaotic lidar rangefinder. In *OCEANS’15 MTS/IEEE Washington*, pages 1–6. IEEE, 2015.
- [19] Brandon Cochenour, Linda Mullen, and Alan Laux. Spatial and temporal dispersion in high bandwidth underwater laser communication links. In *Military Communications Conference, 2008. MILCOM 2008. IEEE*, pages 1–7. IEEE, 2008.
- [20] Brandon Cochenour, Alan Laux, and Linda Mullen. Temporal dispersion in underwater laser communication links: Closing the loop between model and experiment. In *Underwater Communications and Networking Conference (UComms), 2016 IEEE Third*, pages 1–5. IEEE, 2016.
- [21] Robert W Lee, Justin K Nash, Brandon M Cochenour, and Linda J Mullen. Pulse compression techniques to improve modulated pulsed laser line scan systems. In *SPIE Defense+ Security*, pages 94590C–94590C. International Society for Optics and Photonics, 2015.
- [22] Shawn O’connor, Robert Lee, Linda Mullen, and Brandon Cochenour. Waveform design considerations for modulated pulse lidar. In *Ocean Sensing and Monitoring VI*, volume 9111, page 91110P. International Society for Optics and Photonics, 2014.
- [23] Linda J Mullen, Brandon Cochenour, Alan Laux, Robert Lee, and Justin Nash. Application of radar technology to underwater detection, ranging and

- imaging. In *Applications of Lasers for Sensing and Free Space Communications*, pages LM4F–3. Optical Society of America, 2015.
- [24] Justin K Nash, Robert W Lee, and Linda J Mullen. Tailoring of rf coded optical pulses for underwater 3d imaging. In *OCEANS’15 MTS/IEEE Washington*, pages 1–8. IEEE, 2015.
- [25] Leon Cohen. Time-frequency analysis. *Wavelets and Signal Processing*, 2012.
- [26] Syed A Ahson and Mohammad Ilyas. *RFID handbook: applications, technology, security, and privacy*. CRC press, 2008.
- [27] Bin Le, Thomas W Rondeau, Jeffrey H Reed, and Charles W Bostian. Analog-to-digital converters. *IEEE Signal Processing Magazine*, 22(6):69–77, 2005.
- [28] Harry Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928.
- [29] Claude Elwood Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.
- [30] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [31] Emmanuel J Candès and Michael B Wakin. An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2):21–30, 2008.
- [32] David L Donoho. Unconditional bases are optimal bases for data compression and for statistical estimation. *Applied and computational harmonic analysis*, 1(1):100–115, 1993.
- [33] Martin Vetterli, Pina Marziliano, and Thierry Blu. Sampling signals with finite rate of innovation. *IEEE transactions on Signal Processing*, 50(6):1417–1428, 2002.
- [34] Merrill Ivan Skolnik. Radar handbook. 1970.
- [35] David L Donoho and Xiaoming Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7):2845–2862, 2001.
- [36] Stefan Banach. *Theory of linear operations*, volume 38. Elsevier, 1987.
- [37] Balas Kausik Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995.
- [38] Scott Shaobing Chen, David L Donoho, and Michael A Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.

- [39] David L Donoho. For most large underdetermined systems of linear equations the minimal 1-norm solution is also the sparsest solution. *Communications on pure and applied mathematics*, 59(6):797–829, 2006.
- [40] Joel A Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information theory*, 50(10):2231–2242, 2004.
- [41] Mohammad Gharavi-Alkhansari and Thomas S Huang. A fast orthogonal matching pursuit algorithm. In *Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference on*, volume 3, pages 1389–1392. IEEE, 1998.
- [42] Emmanuel J Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies. *IEEE transactions on information theory*, 52(12):5406–5425, 2006.
- [43] Emmanuel J Candes and Terence Tao. Decoding by linear programming. *IEEE transactions on information theory*, 51(12):4203–4215, 2005.
- [44] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore. Compressed sensing and best -term approximation. *Journal of the American mathematical society*, 22(1):211–231, 2009.
- [45] Marco F Duarte and Yonina C Eldar. Structured compressed sensing: From theory to applications. *IEEE Transactions on Signal Processing*, 59(9):4053–4085, 2011.
- [46] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.
- [47] Emmanuel J Candès et al. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.
- [48] Raman Venkataramani and Yoram Bresler. Optimal sub-nyquist nonuniform sampling and reconstruction for multiband signals. *IEEE Transactions on Signal Processing*, 49(10):2301–2313, 2001.
- [49] Moshe Mishali, Yonina C Eldar, and Asaf J Elron. Xampling: Signal acquisition and processing in union of subspaces. *IEEE Transactions on Signal Processing*, 59(10):4719–4734, 2011.
- [50] Justin Romberg. Compressive sensing by random convolution. *SIAM Journal on Imaging Sciences*, 2(4):1098–1128, 2009.
- [51] Marco F Duarte, Mark A Davenport, Dharmpal Takbar, Jason N Laska, Ting Sun, Kevin F Kelly, and Richard G Baraniuk. Single-pixel imaging via compressive sampling. *IEEE signal processing magazine*, 25(2):83–91, 2008.

- [52] Afonso S Bandeira, Matthew Fickus, Dustin G Mixon, and Percy Wong. The road to deterministic matrices with the restricted isometry property. *Journal of Fourier Analysis and Applications*, 19(6):1123–1149, 2013.
- [53] Robert Lee, Piya Pal, and Linda Mullen. Theoretical and experimental study of sub-nyquist fmcw lidar systems. In *Compressive Sensing V: From Diverse Modalities to Big Data Analytics*, volume 9857, 2016.
- [54] F Jessie MacWilliams and Neil JA Sloane. Pseudo-random sequences and arrays. *Proceedings of the IEEE*, 64(12):1715–1729, 1976.
- [55] Hans-Jürgen Zepernick and Adolf Finger. Binary pseudo random sequences. *Pseudo Random Signal Processing*, pages 87–125.
- [56] LászlóJ Naszódi. On digital filtration in correlation time-of-flight spectrometry. *Nuclear Instruments and Methods*, 161(1):137–140, 1979.
- [57] Hongjian Sun, Arumugam Nallanathan, Cheng-Xiang Wang, and Yunfei Chen. Wideband spectrum sensing for cognitive radio networks: a survey. *IEEE Wireless Communications*, 20(2):74–81, 2013.
- [58] PP Vaidyanathan and Piya Pal. Theory of sparse coprime sensing in multiple dimensions. *IEEE Transactions on Signal Processing*, 59(8):3592–3608, 2011.
- [59] Sanjay Mehrotra. On the implementation of a primal-dual interior point method. *SIAM Journal on optimization*, 2(4):575–601, 1992.
- [60] Arkadi Nemirovski. Interior point polynomial time methods in convex programming. *Lecture Notes*, 2004.
- [61] Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. *Foundations and Trends in Machine Learning*, 8(3-4):231–357, 2015.
- [62] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. Greedy algorithms. *Introduction to algorithms*, 1:329–355, 2001.
- [63] David L Donoho, Michael Elad, and Vladimir N Temlyakov. Stable recovery of sparse overcomplete representations in the presence of noise. *IEEE Transactions on information theory*, 52(1):6–18, 2006.
- [64] Yagyensh Chandra Pati, Ramin Rezaifar, and Perinkulam Sambamurthy Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, pages 40–44. IEEE, 1993.
- [65] Stéphane G Mallat and Zhifeng Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on signal processing*, 41(12):3397–3415, 1993.

- [66] Thomas Blumensath and Mike E Davies. Iterative thresholding for sparse approximations. *Journal of Fourier Analysis and Applications*, 14(5):629–654, 2008.
- [67] Simon Foucart. Hard thresholding pursuit: an algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49(6):2543–2563, 2011.
- [68] Arian Maleki and David L Donoho. Optimally tuned iterative reconstruction algorithms for compressed sensing. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):330–341, 2010.
- [69] Frank Neville Hosband Robinson. Noise and fluctuations in electronic devices and circuits. *Oxford, Clarendon Press, 1974. 252 p.*, 1974.
- [70] Ery Arias-Castro and Yonina C Eldar. Noise folding in compressed sensing. *IEEE Signal Processing Letters*, 18(8):478–481, 2011.
- [71] Galen Reeves and Michael Gastpar. Sampling bounds for sparse support recovery in the presence of noise. In *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, pages 2187–2191. IEEE, 2008.
- [72] Harry L Van Trees. *Detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [73] George Turin. An introduction to matched filters. *IRE transactions on Information theory*, 6(3):311–329, 1960.
- [74] Robert Lee, Linda Mullen, Piya Pal, and David Illig. Time of flight measurements for optically illuminated underwater targets using compressive sampling and sparse reconstruction. In *OCEANS’15 MTS/IEEE Washington*, pages 1–6. IEEE, 2015.
- [75] Daniel H Chae, Parastoo Sadeghi, and Rodney A Kennedy. Effects of basis-mismatch in compressive sampling of continuous sinusoidal signals. In *Future Computer and Communication (ICFCC), 2010 2nd International Conference on*, volume 2, pages V2–739. IEEE, 2010.
- [76] Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht. Compressed sensing off the grid. *IEEE transactions on information theory*, 59(11):7465–7490, 2013.
- [77] John R Klauder, AC Price, Sidney Darlington, and Walter J Albersheim. The theory and design of chirp radars. *Bell Labs Technical Journal*, 39(4):745–808, 1960.
- [78] David W Illig, Alan Laux, Robert W Lee, William D Jemison, and Linda J Mullen. Fmcw optical ranging technique in turbid waters. In *Ocean Sensing and Monitoring VII*, volume 9459, page 94590B. International Society for Optics and Photonics, 2015.

- [79] Justin Nash, Robert Lee, and Linda Mullen. Wide bandwidth optical signals for high range resolution measurements in water. In *Proc. SPIE*, volume 9827, page 98270A, 2016.
- [80] Robert W Lee, David W Illig, and Linda J Mullen. Adaptive underwater channel estimation for hybrid lidar/radar. In *Ocean Sensing and Monitoring IX*, volume 10186, page 101860F. International Society for Optics and Photonics, 2017.
- [81] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.
- [82] Ron Rubinstein, Alfred M Bruckstein, and Michael Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045–1057, 2010.
- [83] Emmanuel J Candes, Yonina C Eldar, Deanna Needell, and Paige Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2011.
- [84] Ronald A DeVore. Deterministic constructions of compressed sensing matrices. *Journal of complexity*, 23(4-6):918–925, 2007.
- [85] Martin H Ackroyd and F Ghani. Optimum mismatched filters for sidelobe suppression. *IEEE Transactions on Aerospace and Electronic Systems*, (2):214–218, 1973.
- [86] Jeffrey S Fu and Xin Wu. Sidelobe suppression using adaptive filtering techniques. In *Radar, 2001 CIE international conference on, proceedings*, pages 788–791. IEEE, 2001.
- [87] E Zege, I Katsev, and A Prikhach. Waveform response 3d model. 2014.
- [88] Iosif L Katsev, Eleonora P Zege, Alexander S Prikhach, and Igor N Polonsky. Efficient technique to determine backscattered light power for various atmospheric and oceanic sounding and imaging systems. *JOSA A*, 14(6):1338–1346, 1997.
- [89] Eleonora Petrovna Zege, Arkadi Petrovich Ivanov, and Iosif Lebovich Katsev. *Image transfer through a scattering medium*. Springer Verlag, 1991.
- [90] Eleonora P Zege, Iosif L Katsev, Alexander S Prikhach, Geoff D Ludbrook, and Piero Brusaglioni. Analytical and computer modeling of the oceanic lidar performance. In *12th International Workshop on Lidar Multiple Scattering Experiments*, volume 5059, pages 189–200. International Society for Optics and Photonics, 2003.

- [91] R Foord, R Jones, CJ Oliver, and ER Pike. The use of photomultiplier tubes for photon counting. *Applied optics*, 8(10):1975–1989, 1969.
- [92] David W Illig, William D Jemison, Luke Rumbaugh, Robert Lee, Alan Laux, and Linda Mullen. Enhanced hybrid lidar-radar ranging technique. In *Oceans-San Diego, 2013*, pages 1–9. IEEE, 2013.
- [93] David W Illig, William D Jemison, Robert W Lee, Alan Laux, and Linda J Mullen. Optical ranging techniques in turbid waters. In *Ocean Sensing and Monitoring VI*, volume 9111, page 91110R. International Society for Optics and Photonics, 2014.
- [94] Analog Devices. *3.5 GSPS Direct Digital Synthesizer with 12-Bit DAC*, 2014.
- [95] Photonis. *Optical Receiver Module*, 2014.
- [96] Teledyne SP Devices. *ADQ412 Datasheet*, 2017.
- [97] David L Donoho and Jared Tanner. Thresholds for the recovery of sparse solutions via l1 minimization. In *Information Sciences and Systems, 2006 40th Annual Conference on*, pages 202–206. IEEE, 2006.
- [98] Alfred M Bruckstein, Michael Elad, and Michael Zibulevsky. Sparse non-negative solution of a linear system of equations is unique. In *Communications, Control and Signal Processing, 2008. ISCCSP 2008. 3rd International Symposium on*, pages 762–767. IEEE, 2008.
- [99] Chunjie Zhang, Jing Liu, Qi Tian, Changsheng Xu, Hanqing Lu, and Songde Ma. Image classification by non-negative sparse coding, low-rank and sparse decomposition. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1673–1680. IEEE, 2011.
- [100] Simon Foucart and David Koslicki. Sparse recovery by means of nonnegative least squares. *IEEE Signal Processing Letters*, 21(4):498–502, 2014.
- [101] Mehrdad Yaghoobi, Di Wu, and Mike E Davies. Fast non-negative orthogonal matching pursuit. *IEEE Signal Processing Letters*, 22(9):1229–1233, 2015.
- [102] Jie Chen and Xiaoming Huo. Theoretical results on sparse representations of multiple-measurement vectors. *IEEE Transactions on Signal Processing*, 54(12):4634–4643, 2006.
- [103] Zhilin Zhang and Bhaskar D Rao. Sparse signal recovery with temporally correlated source vectors using sparse bayesian learning. *IEEE Journal of Selected Topics in Signal Processing*, 5(5):912–926, 2011.
- [104] Zhilin Zhang and Bhaskar D Rao. Sparse signal recovery in the presence of correlated multiple measurement vectors. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pages 3986–3989. IEEE, 2010.

- [105] Joel A Tropp, Anna C Gilbert, and Martin J Strauss. Simultaneous sparse approximation via greedy pursuit. In *Acoustics, Speech, and Signal Processing, 2005. Proceedings. (ICASSP'05). IEEE International Conference on*, volume 5, pages v-721. IEEE, 2005.