MIMO RADAR: SIGNAL PROCESSING, WAVEFORM DESIGN, AND APPLICATIONS TO SYNTHETIC APERTURE IMAGING

A Dissertation
Presented to
The Academic Faculty

By

Michael S. Davis

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
in
Electrical and Computer Engineering

School of Electrical and Computer Engineering
Georgia Institute of Technology
May 2015

MIMO RADAR: SIGNAL PROCESSING, WAVEFORM DESIGN, AND APPLICATIONS TO SYNTHETIC APERTURE IMAGING

Approved by:

Dr. Aaron D. Lanterman, Advisor
Associate Professor, School of ECE
Georgia Institute of Technology

Dr. Andrew F. Peterson
Professor, School of ECE
Georgia Institute of Technology

Dr. Mary Ann Weitnauer
Professor, School of ECE
Georgia Institute of Technology

Dr. Yao Xie
Assistant Professor, School of ISyE
Georgia Institute of Technology

Dr. Justin K. Romberg
Associate Professor, School of ECE
Georgia Institute of Technology

Date Approved: 10 March 2015

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

This dissertation analyzes the capability of multiple-input, multiple-output (MIMO) radar techniques to improve the image quality and area-coverage rate of synthetic aperture imaging systems. A signal processing architecture for MIMO radar is used to understand the applicability of MIMO for synthetic aperture radar (SAR) and synthetic aperture sonar (SAS) systems. MIMO SAR/SAS is shown to be a natural extension of standard multi-channel synthetic aperture imaging techniques to exploit transmit degrees of freedom in addition to those used on receive. Degradation in range sidelobe performance and the associated impact on image quality is identified as a key impediment to MIMO SAR/SAS. A novel mismatched filtering approach is presented to mitigate this issue. New results in sampling theory are derived that allow the aliasing that occurs when a wide-sense stationary random process is non-uniformly sampled to be quantified. These results are applied to the case of recurrent sampling and used to quantify the impact of azimuth ambiguities on MIMO SAR/SAS image contrast.

CHAPTER 1

INTRODUCTION

Radar systems have been actively developed since the 1930s with early systems focusing on the detection of aircraft using a ground-based system. In the 1950s, the concept of synthetic aperture radar (SAR) was demonstrated, which used an airborne radar to form high-resolution imagery of the surface of the earth. Today, spaceborne systems like TerraSAR-X provide images of the earth's surface to allow scientists to understand how the planet is continuously reshaped by natural and anthropogenic forces. SAR has even been used to image other planets; in the early 1990s, the Magellan spacecraft mapped the surface of Venus using a radar that was able to penetrate its optically-opaque atmosphere of sulfuric acid. While a more recent development than SAR, synthetic aperture sonar (SAS) systems are now used to map the seafloor.

A more recent development in radar research is the concept of multiple-input, multiple-output (MIMO) radar. The novelty of this idea is to employ multiple radiating elements on transmit that each emit an independent waveform. While traditional SAR and SAS systems may use multiple elements on receive, a MIMO SAR/SAS would also simultaneously transmit multiple independent waveforms. The goals of this dissertation are to create a general theory of MIMO radar and to apply it to analyze the capability of MIMO techniques to improve the performance of synthetic aperture imaging systems.

1.1 Organization and Summary of Contributions

The remainder of this chapter reviews some fundamental ideas of MIMO radar. In Chapter 2, an approach to understanding MIMO radar is presented. This is applied to the case of synthetic aperture imaging in Chapter 3. It will be shown that MIMO radar techniques can improve the area-coverage rate and/or image quality for a SAR/SAS system by improving

along-track sampling. This requires that multiple waveforms be simultaneously transmitted, which will lead to cross-correlation energy that degrades image contrast. A technique for combating this using mismatched filter theory is developed in Chapter 4.

MIMO SAR/SAS is a natural extension of multichannel systems that employ multiple elements on receive to also use multiple elements on transmit. The motivation for multichannel synthetic aperture imaging is to improve along-track sampling, but using multiple channels may lead to nonuniform sampling. Chapter 5 introduces the concept of recurrent sampling, which is a particular nonuniform sampling structure that can occur in multichannel SAR/SAS. A method for computing the impact of aliasing when sampling a wide-sense stationary random process is developed, which is applied to a number of interesting SAR/SAS cases in Chapter 6.

This dissertation documents a number of original contributions:

- A theoretical framework for coherent MIMO radar is developed in Chapter 2, which unifies a number of concepts that have appeared in the MIMO radar literature by extending standard concepts in radar array signal processing [Davis et al., 2014].

- The application of MIMO radar techniques to synthetic aperture imaging is detailed in Chapter 3. While this has been advocated in the literature by other authors, this chapter presents the benefits (and costs) of MIMO in SAR/SAS in terms of standard SAR performance metrics by extending the concept of multiplicative noise to the MIMO case [Davis and Cook, 2011; Davis et al., 2011].

- A new method for controlling the range sidelobes in MIMO radar is presented in Chapter 4. The technique is developed by extending the concept of mismatched filtering to the case of multiple waveforms [Davis and Lanterman, 2015b].

- Contributions to sampling theory are presented in Chapter 5. These include a rigorous method for computing the impact of aliasing in the case of sampling a wide-sense

stationary random process as well as an application to a particular case of nonuniform sampling called recurrent sampling [Davis and Lanterman, 2015a].

- New results in SAR/SAS performance analysis involving the ambiguity-to-signal ratio (ASR) are documented in Chapter 6. This work builds on standard results for ASR in single-channel stripmap SAR/SAS by first using the techniques of Chapter 5 to solidify their theoretical foundation [Davis and Cook, 2014]. Results are then developed for spotlight mode as well as multichannel systems.

1.2 Coherent and Noncoherent MIMO Radar Systems

Since the turn of the century, a great deal of research has investigated the utility of applying MIMO techniques to enhance the performance of radar systems. Although MIMO has been demonstrated to dramatically improve the capacity of communications systems in multipath-rich environments [Foschini and Gans, 1998], and MIMO technology has been included in multiple communications standards, the adoption of MIMO in operational radar systems has been limited. As will be discussed, this may be attributed to radar systems using waveforms and antennas to interact with their environments in a dramatically different way than communications systems. Still, some specific radar applications have been identified where a MIMO radar may outperform a more traditional system.

An active radar system emits electromagnetic energy to probe its environment. A MIMO radar transmits independent waveforms from a number of spatially separated radiating elements and observes the returns from a set of spatially diverse receive elements. The environment is considered as a system where the inputs are the transmitted waveforms, and the outputs are the echos observed by the receivers. This is, in general, a MIMO system. There has been a long history of exploiting multiple degrees of freedom on receive [Southworth, 1930]; the novelty of MIMO is to exploit similar degrees of freedom on transmit.

Of course, the idea of using multiple radiating elements on transmit is not new [Silver, 1949]. Phased array antennas have been an enabling technology for many systems in support of a variety of radar missions. Rather than construct a prohibitively large antenna to support long range detection of small targets, many small antennas are employed, each of which transmits an identical signal (up to a phase shift). This phase progression is chosen to steer a high-gain beam in a particular direction. This beam may be rapidly resteered through electronic scanning, which enables multiple target tracking performance that would be impossible for a mechanically scanned antenna.

Although a phased array uses multiple elements on transmit, the transmitted waveforms are perfectly correlated and therefore provide no additional degrees of freedom. In contrast to the general MIMO radar, phased arrays operate as a single-input, multiple-output (SIMO) system. Phased array radars typically provide receive degrees of freedom by digitizing multiple, spatially-diverse receive channels, hence they are multiple-output. The special case of a radar with a single transmit waveform and a single receive channel is referred to as a single-input, single-output (SISO) radar system.

The idea of transmitting multiple, uncorrelated waveforms has been explored over the years. For example, in [Messer et al., 1996], it was observed that a radar system transmitting two orthogonal waveforms and using a single receiver provides bearing estimates that are equivalent to a traditional radar transmitting a single waveform and using two receivers. (An analogous result in the communications literature [Alamouti, 1998] demonstrated that the diversity order of the same two configurations is equivalent if the transmitted waveform is an orthogonal space-time block code.) An experimental radar that transmitted "orthogonal" waveforms was described in [Dorey and Garnier, 1989].

It was not until the benefits of MIMO for communications were clear that the idea of transmitting uncorrelated waveforms in radar received consistent attention. Using the language of MIMO communications, researchers have described two broad classes of MIMO radars: those with widely-separated antennas that, as in the communications application,

4

seek to exploit spatial diversity to mitigate target fading [Fishler et al., 2006] as opposed to those with closely spaced antennas that permit coherent processing [Bliss and Forsythe, 2003; Li and Stoica, 2007].

The class of MIMO radars that use widely-separated antennas is often referred to as "statistical MIMO radar" because such systems seek to exploit the random fluctuation of target reflectivity as a function of aspect angle. This fluctuation causes spatial decorrelation of the target returns, which precludes coherent processing. As a result, statistical MIMO is often referred to as non-coherent MIMO, while the term coherent MIMO is restricted to systems where the antennas are sufficiently close (perhaps on the same platform) to limit target decorrelation. The following discussion is limited to the latter case of coherent MIMO radar. The case of non-coherent MIMO with widely-separated antennas has been discussed extensively in the radar community in the context of multistatic radar systems [Chernyak, 1998]. Just as coherent MIMO may be considered a natural extension of the phased array, statistical MIMO generalizes the concept of bistatic radar [Willis, 2005].

CHAPTER 2

MIMO RADAR SIGNAL PROCESSING

2.1 Overview

This chapter seeks to establish a framework for understanding the performance benefits and potential drawbacks of MIMO radar systems. MIMO radar is realized by transmitting independent waveforms, but in many ways the waveforms are simply a means to an end. The view presented is that MIMO radar is essentially an antenna technology because it is a generalization of the phased array architecture that has been used in radar and other systems for decades. With this philosophical approach, a remarkable number of concepts described in the MIMO radar literature can be illuminated.

A methodology will be established for evaluating the potential of a suite of waveforms to enable a MIMO radar to effectively perform its mission. This provides the capability of determining if MIMO techniques are appropriate for a given radar system. A tremendous variety of radars are operational. MIMO can dramatically improve the performance of some, but others will benefit from traditional phased array configurations. Further, this framework can inform the design of waveforms, which are critical to the realization of a MIMO radar system.

The first task is to characterize the gain of the antenna employed by a MIMO radar by determining the MIMO radar beampattern. A similar approach is taken in [Bekkerman and Tabrikian, 2006], though the present derivation explicitly deals with the possibility of spatially colored interference and proposes an alternative formulation for dealing with correlated transmit signals. The present approach also explicitly identifies a signal processing architecture and spatial weights that can be applied to form beams in desired directions.

A key result described in the MIMO radar literature is that the transmit beampattern is characterized by the correlation between the transmitted signals [Fuhrmann and San Antonio, 2008]. This correlation impacts the MIMO radar ambiguity function, which was

presented in [San Antonio et al., 2007]. Indeed, the correlation properties of the waveforms transmitted by a MIMO radar determine the characteristics of the system. These characteristics are captured by considering the quantities presented in this chapter.

This chapter is organized as follows. Although radar and communications systems are based on similar physical phenomena and are described by similar mathematics, there are fundamental differences, particularly the role of the antenna subsystem. A discussion of this is presented in Section 2.2. The utility of MIMO radar is motivated by the concept of the virtual array in Section 2.3. This is followed by a derivation of a MIMO radar signal processing chain, described in Section 2.4. These observations are used in Section 2.5 to compare the traditional phased array to a MIMO radar that is transmitting orthogonal waveforms. In Section 2.6, some quasi-orthogonal waveforms are described and their performance is examined.

2.2 MIMO Systems: Radar and Communications

In both radar and communications systems, electromagnetic energy is radiated. This signal interacts with the environment, and the resulting electromagnetic field is observed by a receiver. The goal of a communications system is to estimate the parameters of the input signal, e.g., the sequence of message symbols used to generate it in spite of any environmental effects. In radar, the goal is to infer some property of the environment based on knowledge of the transmitted waveform.

In either case, an input signal is applied to some system, and the response of this system to this input is observed. Radar and communications systems are typically modeled as linear. First, consider a single-input, single-output (SISO) system. Let $x(t)$ be the complex-baseband representation of the input signal that is transmitted on a carrier frequency, ω_c. If the system is linear, the output of the system, $y(t)$, may be written in terms of the input signal and the system impulse response, $h(t)$, as

$$y(t) = \int_0^\infty h(\tau) e^{-i\omega_c \tau} x(t - \tau) \, d\tau + v(t), \tag{2.1}$$

7

where $v(t)$ represents receiver noise.

In communications, the system impulse response, $h(t)$, describes the channel between the transmitter and receiver, which may involve a direct path contribution and/or multipath. For a monostatic radar system, where the transmitter and receiver are colocated, the relevant impulse response describes the scatterers in the environment and may be called the range profile.

For both communications and monostatic radar cases, the impulse response at a particular lag, τ, may be formed from returns from a number of scatterers. The returns from these scatterers all arrive after the same delay, but they may correspond to different angles-of-arrival, θ. In this case, the impulse response may be written in terms of the angle/delay reflectivity profile, $\alpha(\tau, \theta)$, as

$$h(\tau) = \int_{-\pi}^{\pi} \alpha(\tau, \theta)\, d\theta. \tag{2.2}$$

In many radar applications, the goal is to estimate this angle/delay (or, equivalently, angle/range profile). However, in communications, the only relevant parameter is the aggregate channel response, $h(t)$. The fundamental difference between the communications problem and the radar problem is the role of the channel impulse response: in communications it is a nuisance factor that must be estimated to establish a link, but in radar, estimating the "channel response" and, further, the underlying reflectivity profile, $\alpha(\tau, \theta)$, is essential.

In the MIMO case, a set of input signals, $\mathbf{x}(t) \in \mathbb{C}^M$, is used, and a set of output signals, $\mathbf{y}(t) \in \mathbb{C}^N$, are observed where M is the number of transmitted signals, and N is the number of received signals. The model in (2.1) can be extended to the MIMO case,

$$\mathbf{y}(t) = \int_0^{\infty} \mathbf{H}(\tau)\, e^{-i\omega_c \tau} \mathbf{x}(t - \tau)\, d\tau + \mathbf{v}(t), \tag{2.3}$$

where $\mathbf{H}(t)$ is the $N \times M$ MIMO channel matrix that describes the impulse response of the MN channels of the MIMO system.

If the multipath environment described by the channel matrix, $\mathbf{H}(t)$, is suitably rich

and sufficient information about the channel is available, then tremendous gains in the capacity of the communications channel are provided by MIMO. These MN channels may be exploited by sending redundant information over independent channels via space-time coding. Alternatively, if channel state information is available to the transmitter, diversity may be exploited through precoding, which is analogous to beamforming on transmit [Biglieri et al., 2007].

In either approach to MIMO communications, some level of diversity is provided by having access to the MN channels. Indeed, capacity gains are limited if the MIMO channel matrix is not well conditioned [Bliss et al., 2002]. In a sense, by using multiple channels on transmit, a MIMO communications system can exploit available spatial diversity. In many cases, two communications antennas placed on the order of a wavelength apart can observe completely independent channel realizations. Now, consider the case of a coherent MIMO radar and a single target. By definition, the antenna elements are spaced so closely that they observe the same target reflectivity and the only difference will be a phase shift that is related to the target angle.

Statistical MIMO radar, in which the elements are separated widely enough to provide independent realizations of target fading, is based on the same observation as MIMO communications. However, the connection between coherent MIMO radar techniques and MIMO communications is more elusive. Transmit precoding approaches to MIMO communications may seem related to coherent MIMO radar, but they are more closely related to the traditional phased array approach. Just as the transmit precoder of a MIMO communications system uses its knowledge of the multipath environment to maximize the signal-to-noise ratio (SNR) at the receiver, the phased array forms a beam to maximize SNR for a particular target location [Li and Stoica, 2010]. The approach of a coherent MIMO radar is to optimize other properties of the radar antenna at the cost of SNR.

While it would seem natural that MIMO radar and MIMO communications possess a great deal of similarity, the parallels in the context of coherent MIMO radar are limited. A

coherent MIMO radar is not seeking to exploit spatial diversity to ensure that at least one of the "links" is available. Instead, just like the phased array, the goal is to use coherent processing to provide a more capable antenna for the radar system. The following analysis of MIMO radar signal processing will demonstrate how this is accomplished and provide a method to assess its performance.

2.3   The MIMO Virtual Array

The advantage of a MIMO radar transmitting orthogonal waveforms over a traditional phased array is often explained by considering the virtual array [Li et al., 2008]. If an array of M elements is used on transmit and an array of N elements is used on receive, the corresponding virtual array of the system consists of MN virtual elements if the M waveforms are perfectly orthogonal. The additional degrees of freedom provided by transmitting orthogonal waveforms can potentially improve the performance of a radar system [Bliss and Forsythe, 2003]. The MIMO virtual array, which consists of these MN virtual phase centers is introduced. Following this, the MIMO virtual array is compared to the coarray [Hoctor and Kassam, 1990].

2.3.1   Monostatic-Bistatic Equivalence and the MIMO Virtual Array

Suppose that a single element is used to transmit a waveform, and the returns are observed by a single receive element. This system may be considered as a pseudo-bistatic pair, which allows the monostatic-bistatic equivalence theorem to be applied: the signal observed by the bistatic pair is well-approximated by the signal observed in the monostatic case where the transmit and receive elements are located directly between the bistatic elements [Kell, 1965]. As the bistatic angle increases, variation in radar cross section (RCS) between the bistatic pair and the monostatic equivalent will become apparent, but this effect is negligible in the coherent MIMO case where the elements are assumed to be closely spaced. Note that the term "pseudo-bistatic" is used for the case of coherent MIMO, where the bistatic angle is assumed to be small, while the term "bistatic" is reserved for when the transmitter and

receiver are separated by great distances, and the bistatic angle is large [Willis, 2005].

Consider a one-dimensional array where the location of each element is described by an offset, x, from some arbitrary phase center. If x_T is the location of a transmit element, and x_R is the location of a receive element, then the monostatic-equivalent phase center is located between the two at $(x_T + x_R)/2$. The MIMO virtual array is the set of monostatic-equivalent phase centers generated by considering each transmit element and each receive element as a pseudo-bistatic pair. If the locations of the M transmit elements are $x_{T/1}, \ldots, x_{T/M}$ and the locations of the N receive elements are $x_{R/1}, \ldots, x_{R/N}$, then the corresponding MIMO virtual array is

$$\left\{ \frac{x_{T/m} + x_{R/n}}{2} : m = 1, \ldots, M; n = 1, \ldots, N \right\}. \tag{2.4}$$

Note that each transmit element must employ an orthogonal waveform to contribute to the MIMO virtual array.

A comparison of the virtual arrays corresponding to the phased array and to two configurations using orthogonal waveforms is presented in Figure 2.1. In the phased array case, where the waveforms are perfectly correlated from element to element, only data corresponding to a single transmit phase center can be processed, providing only N virtual phase centers. This is illustrated in Figure 2.1(a). If orthogonal waveforms are transmitted, then each of the MN virtual phase centers may be processed independently, which results in the virtual array shown in Figure 2.1(b). Observe that the MN virtual phase centers are not all distinct. In Figure 2.1(c), a large, contiguous array of virtual phase centers is provided by separating the transmit elements.

By transmitting orthogonal waveforms, additional virtual phase centers are available as shown in Figure 2.1. This suggests two possible applications of MIMO radar. First, by simultaneously transmitting M orthogonal waveforms from different spatial locations, the spatial sampling rate can be improved by a factor of M compared to the phased array. This is useful in dealing with along-track sampling challenges in synthetic aperture imaging. Another advantage is the potential improvement in angular resolution that accompanies an
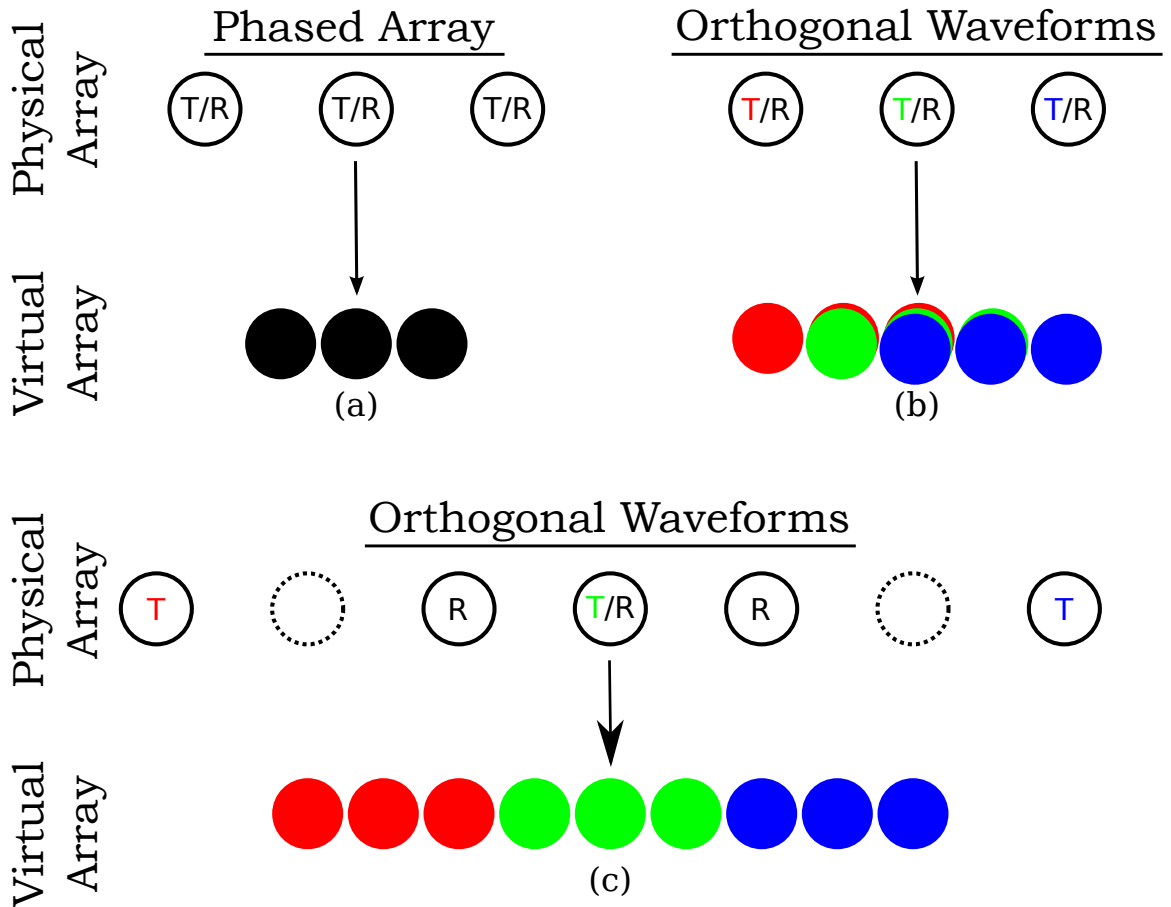
Figure 2.1: Examples of monostatic-bistatic equivalent virtual arrays.

increase in the size of the virtual array. The discussion of MIMO synthetic aperture imaging will be deferred until Chapter 3, but the following discusses the relationship between the MIMO virtual array and angular resolution. The related concept of the MIMO coarray is also introduced.

2.3.2 The Angular Point-Spread Function and The MIMO Coarray

The MIMO virtual array is an extension of the coarray concept for coherent imaging described in [Hoctor and Kassam, 1990], which considers an active, coherent imaging system that scans a transmit and receive beam simultaneously to estimate a reflectivity profile over a number of transmit-receive events. A phased array-like system was assumed to transmit a narrow beam in a particular direction, and linear beamforming was used on receive to steer

a narrow, digital beam in the same direction. This process was repeated for each direction of interest to form an estimate of the reflectivity profile.

By changing the weighting applied on transmit to steer the phased-array beam from dwell to dwell, angular resolution is provided on both transmit and receive. However, for a single phased-array dwell where the transmit beam is stationary, angular resolution is only provided by the receive array. Alternatively, if each element transmits an orthogonal waveform, then the transmit beam can be effectively resteered digitally as described in the next section. The consequence of this is that the standard coarray of [Hoctor and Kassam, 1990] can be applied to the MIMO case, but it corresponds to a single dwell with orthogonal waveforms. The result that the angular point-spread function is the Fourier transform of the coarray can also be applied. This is reviewed in the remainder of this section.

Before proceeding, a distinction between the beampattern of an aperture and the angular point-spread function of an array is made. The beampattern of a transmit aperture describes the power that is radiated in each direction. Similarly, the beampattern of a receive aperture describes the power that is received from a source as a function of angle of arrival. With a single aperture, it is difficult to determine the angle of arrival because it is impossible to distinguish between a strong source located at the peak of the beam and a weak source in the sidelobes. However, if an array of receive elements is available, the beam can be digitally resteered. If a single source is present, the result is the angular point-spread function, which provides a measure of the angular resolution capability of the receive array.

A well-known result of antenna theory is that the aperture illumination function and the antenna beampattern are Fourier duals of one another [Johnson and Dudgeon, 1993]. The two-way beampattern of a radar system is the product of the transmit beampattern and the receive beampattern. Because of the duality between aperture and beampattern and because multiplication in one domain is equivalent to convolution in the other domain, the two-way beampattern is related to the convolution of the transmit and receive apertures by Fourier transform. The convolution of the transmit and receive apertures can be thought of as a

13

"coaperture" that combines the contribution of the transmit aperture and receive aperture to the two-way beampattern. For example, if the same uniformly-illuminated, rectangular aperture is used on transmit and receive, then the two-way beampattern has a sinc-squared shape because the coaperture is the convolution of a rectangular function with itself, which is a triangular function, and the Fourier transform of a triangular function is a sinc-squared function.

An array of elements can be used on receive to determine a signal's angle of arrival. This array may consist of elements that have some spatial extent, e.g., reflector antennas or subarrays of a phased-array antenna. However, the beampattern provided by the elements does not affect the inherent angular resolution capability of the array. If the element patterns have low sidelobes, they may mitigate contributions of targets from angles of interest, but strong signals may overcome this attenuation. As a result, the angular point-spread function of an array does not depend on the beampattern of its constituent elements. However, if an array of perfectly omnidirectional elements is employed, then the beampattern and point-spread function coincide.

The two-way beampattern is the Fourier transform of the coaperture. Similarly, the two-way point-spread function is the Fourier transform of the coarray, which is the convolution of the transmit and receive array-weighting functions. For example, if $\{x_{T/m}\}$ is the set of locations of the elements in the transmit array, then the array function for the unweighted transmit array is

$$g_T(x) \triangleq \sum_{m=1}^{M} \delta(x - x_{T/m}). \tag{2.5}$$

Similarly, the array-weighting function for the unweighted receive array is

$$g_R(x) \triangleq \sum_{n=1}^{N} \delta(x - x_{R/n}), \tag{2.6}$$

where $\{x_{R/n}\}$ is the set of element locations for the receive array. If the array elements are uniformly spaced, then the Fourier transform of the array function is the familiar Dirichlet

kernel of Fourier analysis, which is also called the "aliased" or "periodic" sinc function:

$$D_K(x) \triangleq \sum_{k=-K}^{K} e^{ikx} = \frac{\sin((2K+1)x/2)}{\sin(x/2)},$$ (2.7)

where the number of elements is assumed to be of the form $2K + 1$ for an integer, K.

The MIMO coarray is obtained by convolving the transmit and receive array-weighting functions, which yields

$$g(x) = \left(g_T * g_R\right)(x) = \sum_{m=1}^{M} \sum_{n=1}^{N} \delta(x - (x_{T/m} + x_{R/n})).$$ (2.8)

The MIMO coarray is seen to consist of elements located at

$$\{x_{T/m} + x_{R/n} : m = 1, \ldots, M; n = 1, \ldots, N\}.$$ (2.9)

Compare the virtual array of (2.4) with the coarray defined in (2.9). Observe that they are equivalent up to a scaling by a factor of two. The elements of the virtual array describe an equivalent set of independent, monostatic systems that could collect the equivalent data of the MIMO system. For the single-transmitter/phased-array case, the virtual array is half the length of the physical receive array because of the two-way effects. On the other hand, the coarray is the Fourier dual of the angular point-spread function.

The MIMO coarray can be used to predict the performance of the arrays in Figure 2.1. Recall that the figures present virtual arrays, but the corresponding MIMO coarrays are simply scaled versions of these. The array of Figure 2.1(b) that uses orthogonal waveforms is both longer than the phased array case in Figure 2.1(a), but it also provides a triangular taper to the coarray. In addition to providing improved angular resolution with the larger coarray, the taper provides lower angular sidelobes. In the case of the large contiguous coarray of Figure 2.1(c) where the elements are distinct, by using M transmit elements with orthogonal waveforms, the length of the coarray increases by a factor of M relative to the phased array, which predicts an commensurate improvement in angular resolution.

15

2.4 A MIMO Radar Signal Processing Architecture

The MIMO virtual array and the coarray are useful for comparing the expected performance of a phased array that transmits perfectly correlated waveforms and an array that transmits orthogonal waveforms. However, it does not provide a signal processing architecture to achieve this nor does it handle intermediate cases where the waveforms have some correlation but are not perfectly correlated. In this section, such an architecture is developed, which is used to quantify performance of a MIMO radar that employs an arbitrary suite of waveforms. A signal model and beamforming approach will be developed. This is followed by an extension of standard antenna metrics to the MIMO case, which will elucidate the relation between SNR, resolution, and beampattern for a specified set of waveforms.

2.4.1 MIMO Radar Signal Model

Consider a MIMO radar that transmits M (generally) independent waveforms. Each signal may be emitted by a separate radiating element such as a reflector antenna or a horn antenna. Alternatively, an array divided into a number of subarrays (possibly overlapped) could be used where each subarray acts as an independent radiating element. The only requirement is that each signal be emitted from a distinct phase center. The signals will reflect off of scatterers in the environment, and the echos will be observed by N receive elements. The elements used on receive may or may not be the same as those used on transmit. The operation of a MIMO radar is illustrated in Figure 2.2.

First, a model is constructed for the data observed by a MIMO radar that is due to a single point scatterer. This can be extended to more complex scenarios because of linearity. Without loss of generality, assume that the return is observed with zero delay; equivalently, consider the output of the radar signal processor in the center of the range bin containing a target. This assumption will be revisited in Section 2.6, which investigates its impact. For clarity, it is also assumed that the signal is narrowband relative to the size of the array. Consequently, the data observed by each array element will be identical up to a phase shift.

Figure 2.2: MIMO radar signal model. In this illustration, $M = N = 3$.

The methodology developed can be naturally extended to the wideband case by averaging over frequencies within the radar's operating bandwidth.

Let $x_m(t)$ be the waveform emitted by element m of the transmit array. The signal observed by element n of the receive array from a target at an angle, θ_0, with a (complex-valued) backscatter coefficient, α, may be written as

$$y_n(t; \theta_0) = \alpha b_n(\theta_0) \sum_{m=1}^{M} a_m(\theta_0) x_m(t) + v_n(t), \tag{2.10}$$

where $a_m(\theta_0)$ and $b_n(\theta_0)$ are the phase shifts on transmit and receive, respectively, which correspond to the target angle, θ_0, and $v_n(t)$ is the observation noise. If, for example, the array used on transmit is a uniform linear array (ULA) with an interelement spacing, d, the transmit phase shifts corresponding to an angle, θ, would be

$$a_m(\theta) = \exp\left\{i\frac{\omega_c}{c}(m-1)d\sin\theta\right\} \tag{2.11}$$

for $m = 1, \ldots, M$, where c is the speed of waveform propagation and ω_c is the waveform center frequency expressed in radians per unit time. Of course, steering vectors may be specified for arbitrary array configurations,

$$\exp\left\{i\frac{\omega_c}{c}(R_m - R_1)\right\}, \tag{2.12}$$

where R_m is the range to the target from element m.

These phase shifts may be arranged into the transmit and receive steering vectors, $\mathbf{a}(\theta)$ and $\mathbf{b}(\theta)$. Note that $\mathbf{a} \in \mathbb{C}^M$ and $\mathbf{b} \in \mathbb{C}^N$ are considered to be column vectors. While (2.10) is the data observed by a single receiver, steering-vector notation can be used to compactly write the data observed by the MIMO radar, again corresponding to a single target at angle θ_0, as

$$\mathbf{y}(t; \theta_0) = \alpha \mathbf{b}(\theta_0) \mathbf{a}(\theta_0)^{\mathrm{T}} \mathbf{x}(t) + \mathbf{v}(t), \tag{2.13}$$

where $\mathbf{y}(t; \theta_0) \in \mathbb{C}^N$ is a column vector in which each element is the signal observed by one of the N receivers and $\mathbf{x}(t) \in \mathbb{C}^M$ contains the signals transmitted by each of the M transmitters. The noise observed by each receiver is represented by $\mathbf{v}(t) \in \mathbb{C}^N$. Note that each element of the vector defined by (2.13) is equivalent to (2.10) for $n = 1, \ldots, N$.

The signal model in (2.13) suggests the definition of the MIMO channel matrix given by

$$\mathbf{H}(\theta) \triangleq \mathbf{b}(\theta) \mathbf{a}(\theta)^{\mathrm{T}}. \tag{2.14}$$

This provides the interpretation that the data observed by a MIMO radar is a linear combination of the transmitted signals. This linear combination is described by the MIMO channel matrix, $\mathbf{H}(\theta)$, which depends explicitly on the angle, θ.

A MIMO radar may apply matched filters to resolve targets in range. This is motivated by the assumption that the waveforms are perfectly orthogonal as well as the historical ubiquity of the matched filter in radar signal processing [Turin, 1960; North, 1963]. In the case of orthogonal waveforms, each matched filter will select exactly one signal and reject the rest. Filters for each of the M transmitted signals are applied to each of the N receive signals. A block diagram of this operation is shown in Figure 2.3. In practice, receive beams may be formed first rather than applying the matched filters separately to each receive channel, but this is mathematically equivalent to the diagram.

The result of matched filter processing, when applied to data with a target at angle θ_0 at
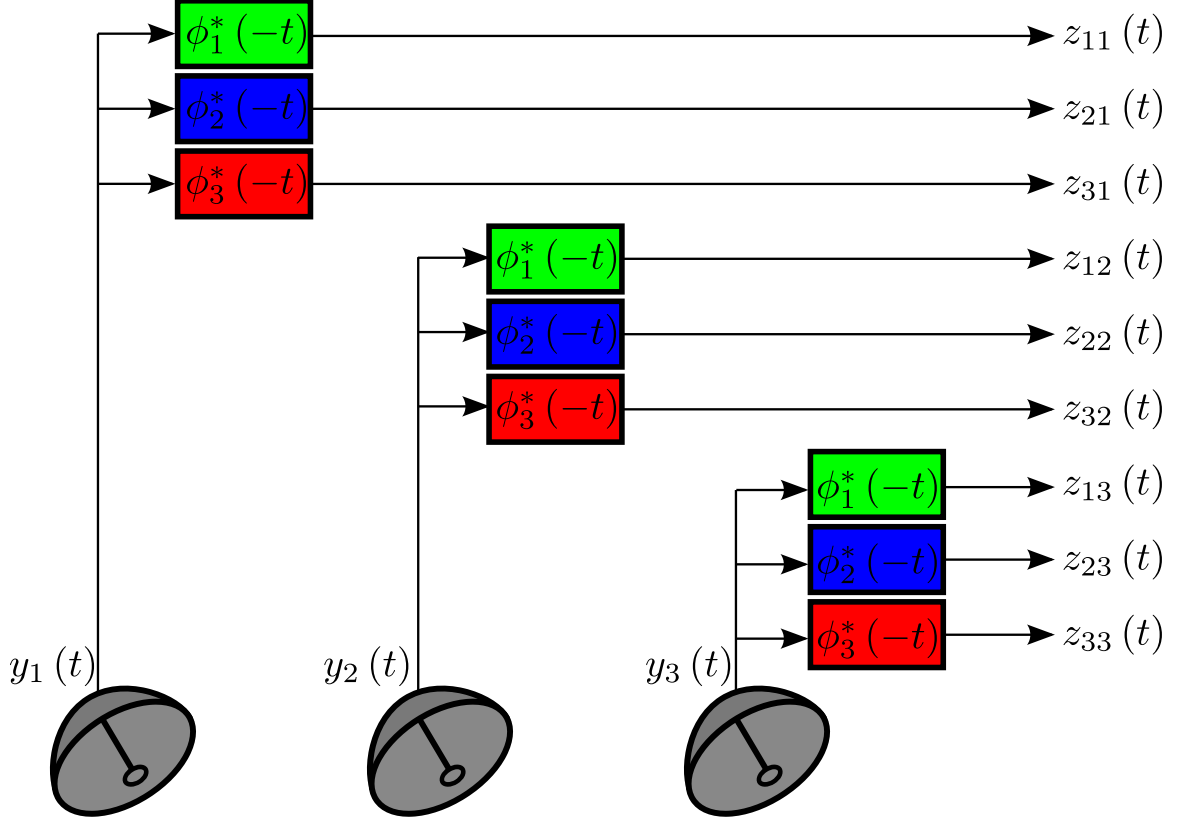
Figure 2.3: Notional MIMO radar signal processor.

the target range, is the $N \times M$ matrix given by

$$
\begin{aligned}
\mathbf{Z}(\theta_0) &\triangleq \int_{-\infty}^{\infty} \mathbf{y}(t;\theta_0)\,\mathbf{x}(t)^{\mathrm{H}}\;dt \\
&= \alpha \mathbf{H}(\theta_0)\,\mathbf{R}_x + \mathbf{E},
\end{aligned}
\tag{2.15}
$$

where $\mathbf{E} \triangleq \int \mathbf{v}(t)\,\mathbf{x}(t)^{\mathrm{H}}\,dt$ is the filtered noise, and $\mathbf{R}_x$ is the $M \times M$ MIMO signal correlation matrix, which describes the correlation among the transmitted waveforms. This matrix is given by

$$
\mathbf{R}_x \triangleq \int_{-\infty}^{\infty} \mathbf{x}(t)\,\mathbf{x}(t)^{\mathrm{H}}\;dt.
\tag{2.16}
$$

The integral of the matrix is understood to be computed element-wise. Each element of this matrix is recognized as the inner product of two of the transmitted signals. Each element of the matrix in (2.16) is the zero-lag term of the cross-correlation function of two of the transmitted signals. The diagonal elements correspond to the autocorrelation functions of the waveforms.

The $N \times M$ data matrix, $\mathbf{Z}(\theta_0)$, given in (2.15) may be vectorized by stacking its columns. By exploiting properties of the Kronecker product operator, $\otimes$, this may be written as

$$\mathbf{z}(\theta_0) \triangleq \text{Vec}\{\mathbf{Z}(\theta_0)\} = \alpha \mathbf{s}(\theta_0) + \mathbf{e}, \tag{2.17}$$

where the MIMO steering vector, $\mathbf{s}(\theta)$, which corresponds to an angle, θ, is given by

$$\mathbf{s}(\theta) \triangleq \left(\mathbf{R}_x^{\text{T}} \otimes \mathbf{I}_N\right) \text{Vec}\{\mathbf{H}(\theta)\}, \tag{2.18}$$

and $\mathbf{e} \triangleq \text{Vec}\{\mathbf{E}\}$ is the vectorized noise matrix.

The MIMO steering vector in (2.18) may also be written in terms of the transmit and receive steering vectors,

$$\mathbf{s}(\theta) = \mathbf{R}_x^{\text{T}} \mathbf{a}(\theta) \otimes \mathbf{b}(\theta). \tag{2.19}$$

The elements of the noise vector, $\mathbf{e}$, will generally be correlated as a result of the matched filtering process. Correlation will be introduced if the transmitted signals are not orthogonal. If the original noise vector, $\mathbf{v}(t)$, in (2.13) is temporally white and wide-sense stationary but has spatial covariance matrix $\mathbf{R}_v$, then the interference covariance matrix of the data in (2.17) will be

$$\mathbf{R}_e \triangleq \text{E}\left[\mathbf{e}\mathbf{e}^{\text{H}}\right] = \mathbf{R}_x^{\text{T}} \otimes \mathbf{R}_v. \tag{2.20}$$

If the noise vector, $\mathbf{v}(t)$, consists only of thermal receiver noise and no external noise sources are present, then the receiver noise should be spatially white and $\mathbf{R}_v = \mathbf{I}_N$, where $\mathbf{I}_N$ is the $N \times N$ identity matrix.

The signal model in (2.17) describes the data observed by a MIMO radar given a target at an angle, θ_0. Each element of the length-MN vector is the voltage observed if the target is in the center of the range bin of interest after matched filtering is applied. This data vector is characterized by the MIMO steering vector in (2.19) and the MIMO interference covariance matrix in (2.20). Note the fundamental importance of the MIMO signal correlation matrix, $\mathbf{R}_x$, which figures prominently in the observed data. Indeed, this matrix characterizes the performance of a MIMO radar.

2.4.2 MIMO Spatial Beamforming

The goal of a spatial beamformer is to enhance signals from targets at some angle of interest, θ, while rejecting signals from other angles. A brief derivation of these weights was presented in [Davis and Cook, 2011], and a similar approach was used in [Bekkerman and Tabrikian, 2006].

The data observed by a MIMO radar that is due to a target at an angle, θ_0, was derived above and presented in (2.17). A linear beamformer will take a linear combination of the elements of the data vector, $\mathbf{z}$. This linear combination is described by the spatial weight vector, $\mathbf{w}$, which provides the (complex-valued) weight to be applied to each of the MN elements of the data vector, $\mathbf{z}$. If these weights are applied to data with a single target at a given angle, θ_0, the output is

$$\mathbf{w}^{\mathrm{H}}\mathbf{z}\left(\theta_0\right) = \underbrace{\alpha\mathbf{w}^{\mathrm{H}}\mathbf{s}\left(\theta_0\right)}_{\text{Signal}} + \underbrace{\mathbf{w}^{\mathrm{H}}\mathbf{e}}_{\text{Noise}}. \tag{2.21}$$

Observe that the average power of the output noise when the weights, $\mathbf{w}$, are applied is

$$P_{\mathrm{N}} \triangleq \mathrm{E}\left[\left|\mathbf{w}^{\mathrm{H}}\mathbf{e}\right|^2\right] = \mathbf{w}^{\mathrm{H}}\mathbf{R}_e\mathbf{w}, \tag{2.22}$$

where $\mathbf{R}_e$ is the spatial covariance matrix of the interference after matched filtering, which was given in (2.20). Similarly, the signal power in the beamformer output is

$$P_{\mathrm{S}} \triangleq \mathrm{E}\left[\left|\alpha\mathbf{w}^{\mathrm{H}}\mathbf{s}\left(\theta_0\right)\right|^2\right] = \sigma_\alpha{}^2\left|\mathbf{w}^{\mathrm{H}}\mathbf{s}\left(\theta_0\right)\right|^2, \tag{2.23}$$

where $\sigma_\alpha{}^2 \triangleq \mathrm{E}\left[|\alpha|^2\right]$.

To design a linear beamformer that preserves signals from targets at a particular angle, θ, the goal is to choose the weight vector, $\mathbf{w}\left(\theta\right)$, that preserves signals corresponding to the steering vector, $\mathbf{s}\left(\theta\right)$, while minimizing the contribution of noise to the beamformer output. The optimal spatial weights minimize the noise power, given in (2.22), yet preserve signals from the desired direction, θ. This leads to the constrained optimization problem

$$\mathbf{w}\left(\theta\right) = \underset{\mathbf{w}\in\mathbb{C}^{MN}}{\arg\min}\left\{\mathbf{w}^{\mathrm{H}}\mathbf{R}_e\mathbf{w} : \mathbf{w}^{\mathrm{H}}\mathbf{s}\left(\theta\right) = 1\right\}. \tag{2.24}$$

21

This optimization problem occurs frequently in beamforming applications. This should not be confused with the minimum variance distortionless response (MVDR) beamformer, which would replace $\mathbf{R}_e$ in (2.24) with $\mathbf{R}_z$.

The optimal weights also maximize the output signal-to-noise ratio (SNR) [Applebaum, 1976]. It is well-known that the optimal weights, $\mathbf{w}(\theta)$, satisfy the Weiner-Hopf equations (see, e.g., [Johnson, 1982]),

$$\mathbf{R}_e \mathbf{w}(\theta) = \gamma \mathbf{s}(\theta), \tag{2.25}$$

for some arbitrary scalar, $\gamma \neq 0$.

The weights will be unique if (and only if) the matrix, $\mathbf{R}_e$, is invertible. When this is the case, the optimal weights are given by $\mathbf{R}_e^{-1} \mathbf{s}(\theta_0)$. However, for the interference covariance matrix, $\mathbf{R}_e$, to be invertible, $\mathbf{R}_v$ and $\mathbf{R}_x$ must be invertible. The spatial noise covariance matrix, $\mathbf{R}_v$, will generally be invertible because of the inevitable presence of thermal noise in the receiver, but an example of a non-invertible signal correlation matrix is the phased array that employs multiple transmit elements that transmit waveforms that are perfectly correlated.

The system of equations in (2.25) cannot be solved precisely when the vector, $\mathbf{s}(\theta)$, does not lie in the column space of the matrix, $\mathbf{R}_e$. However, inspecting the structure of the MIMO steering vector in (2.19), reveals that valid steering vectors satisfy this requirement. So, a set of optimal weights exist, but they are not unique. The minimum-energy weights that achieve the maximum SNR are of the same form as those in the case where the interference covariance matrix is invertible. This can be derived via the pseudoinverse. The optimal weights for a direction, θ, are proportional to

$$\mathbf{w}(\theta) = \mathbf{a}(\theta) \otimes \left(\mathbf{R}_v^{-1} \mathbf{b}(\theta) \right). \tag{2.26}$$

The output of the optimal linear beamformer for a direction, θ, to a single target at an angle, θ_0, has an SNR of

$$\frac{P_{\mathrm{S}}}{P_{\mathrm{N}}} = \sigma_\alpha^2 \left(\frac{\left| \mathbf{a}(\theta_0)^{\mathrm{H}} \mathbf{R}_x^* \mathbf{a}(\theta) \right|^2}{\mathbf{a}(\theta)^{\mathrm{H}} \mathbf{R}_x^* \mathbf{a}(\theta)} \right) \left(\frac{\left| \mathbf{b}(\theta_0)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta) \right|^2}{\mathbf{b}(\theta)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta)} \right). \tag{2.27}$$

The quadratic forms in (2.27) are recognized as inner products with respect to the matrix $\mathbf{R}_x^*$ for transmit and $\mathbf{R}_v^{-1}$ for receive. This allows it to be expressed as

$$\frac{P_\mathrm{S}}{P_\mathrm{N}} = \left(\frac{\left|\langle \mathbf{a}\,(\theta_0)\,, \mathbf{a}\,(\theta)\rangle_{\mathbf{R}_x^*}\right|^2}{\|\mathbf{a}\,(\theta)\|_{\mathbf{R}_x^*}^2} \right) \left(\frac{\left|\langle \mathbf{b}\,(\theta_0)\,, \mathbf{b}\,(\theta)\rangle_{\mathbf{R}_v^{-1}}\right|^2}{\|\mathbf{b}\,(\theta)\|_{\mathbf{R}_v^{-1}}^2} \right),\qquad(2.28)$$

where $\langle \mathbf{x}, \mathbf{y}\rangle_\mathbf{A}$ denotes the inner product of two vectors, $\mathbf{x}$ and $\mathbf{y}$, with respect to a matrix, $\mathbf{A}$. The SNR gain for a target in a direction, θ_0, when beamforming in a direction, θ, is seen to be related to the "angle" between the steering vectors, i.e., the gain for a target at θ_0 is related to the similarity between its steering vector and the steering vector corresponding to the beamformed direction, θ.

2.4.3 MIMO Antenna Gain

The array factor describes the pattern of an array antenna if each subarray was omnidirectional. The MIMO array factor in a direction, θ_0, for an arbitrary set of beamforming weights, $\mathbf{w}$, is defined by

$$f\,(\theta_0) \triangleq \frac{\mathbf{w}^\mathrm{H}\mathbf{s}\,(\theta_0)}{\sqrt{\mathbf{w}^\mathrm{H}\mathbf{R}_e\mathbf{w}}}.\qquad(2.29)$$

This represents the (voltage) gain on a target at an angle, θ_0, relative to the noise as is evident from (2.21). If the optimal weights to steer the beam in a direction, θ, are employed, the resulting MIMO array factor is

$$f\,(\theta_0;\theta) = \left(\frac{\mathbf{a}\,(\theta_0)^\mathrm{H}\,\mathbf{R}_x^*\mathbf{a}\,(\theta)}{\sqrt{\mathbf{a}\,(\theta)^\mathrm{H}\,\mathbf{R}_x^*\mathbf{a}\,(\theta)}} \right) \left(\frac{\mathbf{b}\,(\theta_0)^\mathrm{H}\,\mathbf{R}_v^{-1}\mathbf{b}\,(\theta)}{\sqrt{\mathbf{b}\,(\theta)^\mathrm{H}\,\mathbf{R}_v^{-1}\mathbf{b}\,(\theta)}} \right),\qquad(2.30)$$

where the notation $f\,(\theta_0;\theta)$ is employed to imply that this is the array factor for a beamformer steered to a particular angle, θ. Observe that the first quotient in (2.30) is the transmit array factor, while the second quotient is the standard receive array factor. A similar expression for the array factor is derived in [Bekkerman and Tabrikian, 2006] for the case of spatially-white receiver noise ($\mathbf{R}_v = \mathbf{I}_N$) using an alternate but equivalent formulation of the optimal spatial weights.

The gain in SNR for a target at an angle, θ_0, when a beam is steered to a particular angle, θ, is often of interest. This gain is the magnitude-squared of the array factor and includes

the subarray gain. If, for a target at an angle, θ_0, each transmit subarray has identical gain, $G_{\text{TX}}(\theta_0)$, and each receive subarray has identical gain, $G_{\text{RX}}(\theta_0)$, the MIMO gain is

$$G(\theta_0; \theta) = \left(G_{\text{TX}}(\theta_0) \frac{\left| \mathbf{a}(\theta_0)^{\text{H}} \mathbf{R}_x^* \mathbf{a}(\theta) \right|^2}{\mathbf{a}(\theta)^{\text{H}} \mathbf{R}_x^* \mathbf{a}(\theta)} \right) \left(G_{\text{RX}}(\theta_0) \frac{\left| \mathbf{b}(\theta_0)^{\text{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta) \right|^2}{\mathbf{b}(\theta)^{\text{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta)} \right). \qquad (2.31)$$

Again, the first quotient is recognized as the transmit gain, while the second quotient is the receive gain.

This is a fundamental result. As in the traditional case of multiple degrees of freedom on receive (but not transmit), the response of an adaptive array on receive is characterized by the spatial covariance matrix of the interference, $\mathbf{R}_v$. In the MIMO case, this extends to the transmit gain through the MIMO signal correlation matrix, $\mathbf{R}_x$. Indeed, this matrix fully characterizes the behavior of a MIMO radar. This highlights the fact that, although MIMO radar is enabled by flexibility in waveform generation, it is inherently an antenna-based technology.

2.4.4 Antenna Patterns

The performance of an array antenna for use in a radar system is well-quantified by considering three gain patterns: the steered response, the beampattern, and the (angular) point-spread function [Johnson and Dudgeon, 1993]. These describe the ability of the data collected by the system to be used to digitally form beams in desired directions with desired properties in the following ways.

- The steered response, $G_1(\theta)$, quantifies the ability of the array to digitally steer a beam in a direction, θ.

- The beampattern, $G_2(\theta_0; \theta)$, quantifies the ability of an array to reject targets from an angle, θ_0, when the array is steered in a direction, θ.

- The point-spread function (PSF), $G_3(\theta; \theta_0)$, quantifies the angular response of an array to a target at an angle, θ_0, if the array is steered to an angle, θ.

Table 2.1: Description of antenna patterns. $G(\theta_0; \theta)$ is the gain for a target at an angle, θ_0, when the beam is digitally steered in a particular direction, θ.

Quantity	Describes ability to...	Definition
Steered Response	Digitally resteer antenna	$G_1(\theta) = \{G(\theta; \theta) : \theta \in \Theta\}$
Beampattern	Reject undesired targets	$G_2(\theta_0; \theta) = \{G(\theta_0; \theta) : \theta_0 \in \Theta_0\}$ for fixed θ
PSF	Resolve targets in angle	$G_3(\theta; \theta_0) = \{G(\theta_0; \theta) : \theta \in \Theta\}$ for fixed θ_0

The distinctions between these patterns are summarized in Table 2.1. Observe that the beampattern and point-spread function are equivalent if the subarrays are assumed to be omnidirectional.

These results consider the case where non-omnidirectional elements are used, which can be assumed to be subarrays of the larger array. The case of a MIMO radar that uses such elements on transmit has been discussed in the literature, e.g., in [Hassanien and Vorobyov, 2010], where it is referred to as "Phased-MIMO." Of course, the idea of using receive subarrays in phased array radars is a mature concept [Cheston, 1968].

Expressions have been derived for the antenna gain provided by a MIMO radar that uses a given set of waveforms. This allows three important gain patterns to be calculated. In the next section, these concepts will be employed to describe the characteristics of a radar using orthogonal waveforms and to contrast this with the phased array.

2.5 The Phased Array vs. Orthogonal Waveforms

As previously mentioned, the phased array may be considered as a MIMO radar that transmits waveforms that are perfectly correlated. At the other extreme is what springs to mind when one thinks of a MIMO radar: multiple elements that transmit orthogonal waveforms.

The MIMO signal correlation matrix was introduced in (2.16). In the following, the correlation matrices for these two extreme cases are presented and the resulting antenna performance is described. This will demonstrate the utility (and possible limitations) of transmitting orthogonal waveforms compared to the traditional phased array.

2.5.1 The MIMO Signal Correlation Matrix

In a phased array, each element ideally transmits an identical signal up to a phase shift that can vary from element to element. Let $x_0(t)$ be the (scalar-valued) radar waveform that is common to each element. To steer a beam in a direction, $\tilde{\theta}$, the appropriate phase progression is applied if the transmitted signals are

$$\mathbf{x}_{\text{PA}} \triangleq \mathbf{a}^*\left(\tilde{\theta}\right) x_0(t), \tag{2.32}$$

where $\mathbf{a}^*$ denotes the complex-conjugate (without transpose) of a vector, $\mathbf{a}$. The corresponding signal correlation matrix is

$$\mathbf{R}_{x/\text{PA}} = \frac{1}{M}\left[\mathbf{a}^*\left(\tilde{\theta}_0\right)\right]\left[\mathbf{a}^*\left(\tilde{\theta}_0\right)\right]^{\text{H}}, \tag{2.33}$$

where the steering vector, $\mathbf{a}^*$, was normalized so that $\|\mathbf{a}^*\| = \sqrt{M}$ and that the waveform, $x_0(t)$, was normalized to have total energy of

$$\int_{-\infty}^{\infty} |x_0(t)|^2 \, dt = \frac{1}{M}. \tag{2.34}$$

Note that the vector $[\mathbf{a}^*]^{\text{H}}$ in (2.33) is simply the transpose of the steering vector, $\mathbf{a}$, and could also be written as $\mathbf{a}^{\text{T}}$.

In the phased array case, the waveforms transmitted by each element are perfectly correlated. Now, consider the case where a set of orthogonal waveforms is used. Suppose that they are normalized so that

$$\int_{-\infty}^{\infty} x_m(t)\, x_{m'}^*(t) \, dt = \begin{cases} \frac{1}{M}, & \text{for } m = m' \\ 0, & \text{for } m \neq m'. \end{cases} \tag{2.35}$$

The MIMO signal correlation matrix corresponding to a set of mutually orthogonal waveforms, each with equal energy, is a scaled identity matrix,

$$\mathbf{R}_{x/\perp} = \frac{1}{M}\mathbf{I}_M. \tag{2.36}$$

Note that the signal correlation matrices, $\mathbf{R}_{x/\mathrm{PA}}$ and $\mathbf{R}_{x/\perp}$, have been normalized so that the trace (the sum of the diagonal elements) of the correlation matrix is unity in both cases,

$$\mathrm{Tr}\,(\mathbf{R}_x) = 1. \tag{2.37}$$

This is equivalent to requiring that the total transmitted energy be unity. Such a requirement allows suites of signals with different correlation matrices to be compared fairly.

These two cases, the phased array on one hand and the radar using orthogonal waveforms on the other, are the two extremes of the continuum on which a MIMO radar may operate. This is apparent from examining their MIMO signal correlation matrices: the phased array correlation matrix, $\mathbf{R}_{x/\mathrm{PA}}$, is rank-1 while the orthogonal waveform correlation matrix, $\mathbf{R}_{x/\perp}$, is full rank (rank-M).

As seen from the signal model given in (2.15), the rank of the signal correlation matrix determines how observable the channel matrix, $\mathbf{H}(\theta)$, is. Of course, in the degenerate case when no energy is transmitted ($\mathbf{R}_x = \mathbf{0}$), no portion of the channel matrix is observed. The other extreme is the full-rank case, where a full-rank observation is available. The phased array case, where the signal correlation matrix is rank-1, provides an observation of a single subspace of the channel matrix.

A third case will also be of interest, namely the phased array that uses a spoiled beam on transmit [Kinsey, 1997]. Instead of forming a narrow, high-gain beam, a phased array may trade some of its transmit beamforming gain in return for illuminating a larger set of angles. This may be accomplished by applying a nonlinear phase progression or an amplitude taper across the aperture. An alternate method is to transmit out of a single subarray. The term "spoiled" indicates that the phased array intentionally degrades coherence on transmit and accept a loss in peak gain to form a wider beam.

The MIMO signal correlation matrix, $\mathbf{R}_x$, characterizes the performance of the radar antenna in terms of its transmit gain for a given set of waveforms, $\{x_m(t)\}$. A phased array and a radar using orthogonal waveforms represent the two extremes of MIMO radars. The phased array has the rank-1 signal correlation matrix, $R_{x/\text{PA}}$, given in (2.33). The signal correlation matrix for orthogonal waveforms, $R_{x/\perp}$, given in (2.36), is a full-rank matrix. Using the framework developed above, the MIMO steering vector, interference covariance, and optimal spatial weights can be determined for each case. These quantities are summarized in Table 2.2.

Table 2.2: Comparison of a phased array and a radar using orthogonal waveforms.

Quantity	Phased Array	Orthogonal Waveforms
Signal Correlation, $\mathbf{R}_x$	$\frac{1}{M}\mathbf{a}^*(\tilde{\theta}_0)\mathbf{a}^*(\tilde{\theta}_0)^{\mathrm{H}}$	$\frac{1}{M}\mathbf{I}_M$
MIMO Steering Vector, $\mathbf{s}(\theta)$	$\left(\frac{1}{M}\mathbf{a}(\tilde{\theta}_0)^{\mathrm{H}}\mathbf{a}(\tilde{\theta}_0)\right)\left(\mathbf{a}(\tilde{\theta}_0)\otimes\mathbf{b}(\theta)\right)$	$\mathbf{a}(\theta_0)\otimes\mathbf{b}(\theta)$
Optimal Weights, $\mathbf{w}(\theta_0)$	$\mathbf{a}(\tilde{\theta}_0)\otimes\mathbf{b}(\theta_0)$	$\mathbf{a}(\theta_0)\otimes\mathbf{b}(\theta_0)$
Transmit Gain, $G_{\mathrm{TX}}(\theta;\theta_0)$	$\frac{1}{M}\left\|\mathbf{a}(\tilde{\theta}_0)^{\mathrm{H}}\mathbf{a}(\theta)\right\|^2$	$\frac{1}{M^2}\left\|\mathbf{a}(\theta_0)^{\mathrm{H}}\mathbf{a}(\theta)\right\|^2$
Receive Gain, $G_{\mathrm{RX}}(\theta;\theta_0)$	$\frac{1}{N}\left\|\mathbf{b}(\theta_0)^{\mathrm{H}}\mathbf{b}(\theta)\right\|^2$	$\frac{1}{N}\left\|\mathbf{b}(\theta_0)^{\mathrm{H}}\mathbf{b}(\theta)\right\|^2$
Peak Transmit Gain	M	1
Peak Receive Gain	N	N
Resteer Transmit Beam?	No	Yes
Resteer Receive Beam?	Yes	Yes

Recall that the phased array seeks to steer its transmit beam in some direction, $\tilde{\theta}$. This results in a transmit gain (neglecting subarray gain) for a target at angle, θ_0, of

$$G_{\mathrm{TX/PA}}(\theta_0) = \frac{1}{M}\left|\mathbf{a}(\tilde{\theta})^{\mathrm{H}}\mathbf{a}(\theta_0)\right|. \tag{2.38}$$

As expected, the phased array antenna does not have the capability of digital transmit beam steering. Because the waveforms are correlated from transmit subarray to subarray, there are no transmit degrees of freedom available to the radar.

Compare this to the radar that transmits orthogonal waveforms. The transmit gain for a

target in some direction, θ_0, when the transmit beam is (digitally) steered to an angle, θ, is

$$G_{\text{TX}/\perp}(\theta_0;\theta) = \frac{1}{M^2}\left|\mathbf{a}(\theta)^{\text{H}}\,\mathbf{a}(\theta_0)\right|. \tag{2.39}$$

Note that the transmit beam may be digitally resteered to any angle, θ, with the same set of received signals.

This notion of digitally resteering the transmit beam may seem nonsensical at first, but it is accomplished by using the transmit degrees of freedom provided by a radar transmitting orthogonal waveforms just as a phased array with receive degrees of freedom can digitally resteer a receive beam. The receive beam is digitally resteered by digitizing a number of spatially-diverse channels that are linearly combined to synthesize arbitrary receive beams. Exploiting the orthogonality of the transmit waveforms, a MIMO radar uses matched filters to isolate the contribution of each transmitter to provide a number of spatially-diverse channels that are linearly combined to synthesize arbitrary transmit beams.

Two things are evident when the transmit gains of these two extreme cases of MIMO radar in (2.38) and (2.39) are compared. First, the phased array provides additional transmit gain for targets with angle $\theta_0 = \tilde{\theta}$. Indeed, if M transmit elements are used, the phased array can provide an improvement in SNR by a factor of M when compared to the orthogonal waveform case.

However, this transmit beamforming gain provided by the phased array is only applicable to targets in the direction where the transmit beam was steered, $\tilde{\theta}$. Because the phased array lacks transmit degrees of freedom, it is unable to use digital processing to resteer its transmit beam. The phased array provides improved gain in a particular direction, but a MIMO radar provides a digital beamforming capability on transmit by using orthogonal waveforms.

2.5.2 A Comparison of Gain Patterns

These concepts are made more concrete by considering the beampattern, steered response, and angular point-spread function of the phased array and orthogonal waveforms. An example of each of these is presented in Figure 2.4. The (transmit) subarray pattern and the patterns of the spoiled phased array are also presented for reference.

From the steered response, it is seen that the phased array is able to provide superior gain for the direction in a which the beam was formed (in analog) on transmit. By transmitting M correlated waveforms, an improvement in gain in this direction is provided. The spoiled phased array or the radar using orthogonal waveforms has a steered response identical to the subarray pattern because no array gain is realized on transmit.

By considering the transmit beampattern, it is apparent that the radar using orthogonal waveforms is able to synthesize a beampattern with the same mainlobe width and same sidelobe performance as the phased array. Recall from the steered response that the overall gain is not as high as the phased array, but the beampattern is preserved.

The point-spread function demonstrates that the orthogonal waveform case possesses superior angular resolution performance compared to the phased array. In fact, neither the full phased array or the spoiled phased array are able to provide any resolution on transmit. Note that a similar point-spread function is presented in [Bekkerman and Tabrikian, 2006], where it is (somewhat misleadingly) referred to as a beampattern.

To summarize, by transmitting orthogonal waveforms, a MIMO radar is able to achieve the wide area coverage of a spoiled phased array while preserving its beampattern. It is also able to provide improved angular resolution. Note that, in this case, a uniform linear array was employed and the correlation among the waveforms was varied. More novel configurations may be used where orthogonal waveforms are able to operate in a sparse, irregular configuration.
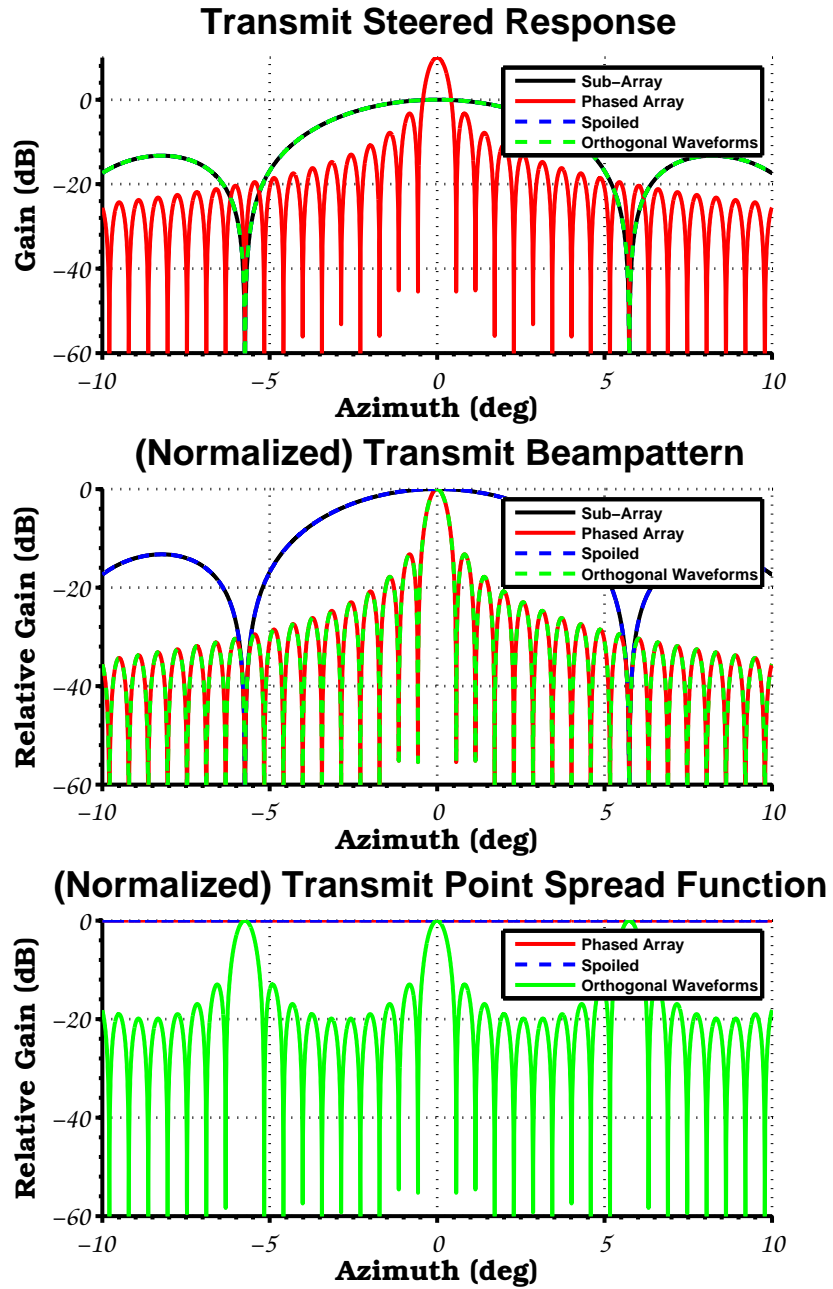
Figure 2.4: Phased array and orthogonal waveform gain patterns.

2.5.3 A Discussion of Search Rates

In many phased array radar systems, a high-gain, directional beam dwells in a particular location for some coherent processing interval (CPI). The length of the CPI is chosen to meet some requirement, e.g., a satisfactory probability of detection for a given class of targets. By changing the look angles, $\left(\tilde{\theta}, \tilde{\phi}\right)$, the phased array can steer this beam over the search volume.

Now, consider a radar that transmits M orthogonal waveforms. As was seen in the above discussion, for a single dwell, a phased array improves SNR by a factor of M over that of a radar using orthogonal waveforms, but the latter is able to illuminate an area that is M times as large in angular extent. Consequently, by using a CPI that is M times as long as that used by the phased array, the radar using orthogonal waveforms can obtain equivalent SNR. Also, since it illuminates a larger volume, the search rate is unaffected. This approach requires that the targets remain sufficiently coherent over the extended CPI, which may require compensation for non-constant velocity in some applications.

To summarize, consider the special case of omnidirectional elements (or elements that are matched to the specified surveillance volume) and for the moment ignore complications such as beam shape loss and beam overlapping strategies. The phased array would steer the beam through a series of M beam positions. On the other hand, the radar using orthogonal waveforms would use a single CPI that was M times as long as that used by each phased array dwell. In both cases, the search rate and SNR on any target will be equivalent.

This is consistent with the standard search-radar equation, which demonstrates that the search rate of a radar is fundamentally limited by its power-aperture product [Scheer, 2010]. Indeed, the SNR constraints on search rate are seen to be unaffected by the choice of operating an array as a phased array or using orthogonal waveforms.

This demonstrates that the "total" SNR of a coherent MIMO radar is unaffected by the correlation matrix in many cases. However, as was seen before, using orthogonal waveforms provides many benefits. This includes improved angular resolution and an improved

beampattern compared to the spoiled configuration.

Compare this with the case of statistical, i.e., non-coherent MIMO radar. For a non-coherent MIMO dwell, M radar systems simultaneously illuminate some area/volume. It is not clear that the CPI duration could be decreased by a factor of M. Also, for a volume-surveillance radar, the volumes illuminated by each radar may not completely overlap, representing an additional inefficiency. Careful bookkeeping is required when benchmarking the performance of a non-coherent MIMO radar and comparing it to the case where each radar operates independently and data fusion occurs at a higher-level, e.g., track-level fusion.

2.5.4 Gain Patterns in Applications of MIMO Radar

In many ways, gain patterns fully characterize the antenna performance of a MIMO radar. Consider, for example, the case of a MIMO ground-moving target indication (GMTI) radar, as described in [Forsythe and Bliss, 2010]. The first question to address is to assess the impact on the apparent loss of gain when moving from a phased array to transmitting orthogonal waveforms. It would seem that a MIMO GMTI radar, by transmitting M orthogonal waveforms would suffer a factor of M loss in signal-to-noise ratio (SNR). Of course, this could be compensated by dwelling M times as long for each CPI, but this would cause a commensurate decrease in area coverage rate. However, as was seen from the steered response, using orthogonal waveforms allows a larger area of the ground to be illuminated. Instead of using a phased array to scan a narrow high-gain beam on the ground from CPI to CPI, the MIMO radar is able to use a lower gain beam with a larger footprint and employ a longer CPI. So, if the targets of interest and the clutter remains sufficiently coherent over this long CPI, the SNR may be recovered.

As pointed out in [Forsythe and Bliss, 2010], these longer CPIs provide improved Doppler resolution. In addition, as seen by considering the point-spread function, MIMO radar also provides improved angular resolution. GMTI systems are able to detect very slow moving targets in the presence of strong clutter returns through the application of

space-time adaptive processing (STAP) [Melvin, 2004]. This is accomplished by exploiting the angle-Doppler coupling of stationary clutter. Clearly, a MIMO GMTI system that could transmit orthogonal waveforms could provide enhanced detection performance against slow moving targets when compared to a phased array. Flight test data has been collected to demonstrate improved detection performance that is provided by a MIMO GMTI [Kantor and Davis, 2010].

While the angular point-spread function of the array used by a GMTI system is critical, it is of less importance to a synthetic aperture radar (SAR) system, which seeks to form high resolution imagery of the ground by synthesizing a synthetic aperture much larger than its physical aperture through platform motion. In this case, the physical aperture does not provide the resolution necessary to resolve targets in the along-track dimension. Instead, the role of the antenna is to act as a spatial filter to reject returns that would otherwise be Doppler ambiguous for the radar's along-track sampling rate. In this case, the key figure of merit is the antenna beampattern.

This is an extension of the Vernier array approach that uses multiple receive channels to improve area coverage rates in along-track sampling limited applications [Kock, 1972]. This view of MIMO SAR was presented in [Davis et al., 2011]. An analogous approach applies to the related technology of synthetic aperture sonar (SAS) [Davis and Cook, 2011]. Other discussions of MIMO SAR are presented in [Krieger et al., 2008; Rennich, 2009; Ender and Klare, 2009]. MIMO SAR will be discussed in more detail in Chapter 3.

2.6 The MIMO Radar Range Response

In the previous discussion, the angular resolution properties of a MIMO radar have been discussed. This was accomplished by considering a single range bin where a target was assume to be present and examining the angular response of the beamformer. This discussion is now extended to consider the response in range. To accomplish this, the definition of the MIMO signal correlation matrix of (2.16) is redefined to include a dependence on time.

Up to this point, only the inner products of the transmitted waveforms with one another have been important. These correspond to the zero-lag terms of the autocorrelations and cross-correlations. Now, the full autocorrelations and cross-correlations of the suite of waveforms must be considered. This further reinforces the assertion that coherent MIMO radar is essentially an antenna concept rather than a waveform-based technique. Indeed, the waveforms (and their cross-correlations in particular) are the source of many of the problems facing the implementation of MIMO radar.

The preceding analysis has tacitly assumed that the range to the target was known so that the outputs of the matched filters in (2.15) may be sampled at the peak of the target response. In practice, a radar signal processor generates the sampled response of the matched filter at a rate corresponding to the expected range resolution. Each sample is the (complex) voltage associated with a range bin. The model developed above in (2.17) for $\mathbf{z}(\theta_0)$, which corresponds to the peak of the matched filter response to a target at an angle, θ_0, is now augmented to include a lag term, τ, which is relative to the peak of the response.

Observe that each of the elements of the matrix defined by (2.16) is the inner product of a pair of transmitted waveforms. This describes the response of each waveform to the matched filter constructed for each of the other waveforms sampled at the peak of the autocorrelation. A lag term, τ, is introduced to capture the response of each matched filter to a shifted version of each of the other waveforms. The MIMO signal correlation matrix for a lag, τ, is given by

$$\mathbf{R}_x(\tau) \triangleq \int_{-\infty}^{\infty} \mathbf{x}(t)\,\mathbf{x}(t-\tau)^{\mathrm{H}}\,dt. \tag{2.40}$$

This can obviously be extended to include mismatches in Doppler frequency [San Antonio et al., 2007; Friedlander, 2012]. Following a similar development to that above, the response of the M matched filters to the signals from the N receivers given a target at an

angle, θ_0, may be written as

$$
\begin{aligned}
\mathbf{Z}(\tau; \theta_0) &\triangleq \int_{-\infty}^{\infty} \mathbf{y}(t; \theta_0) \mathbf{x}(t - \tau)^{\mathrm{H}} \, dt \\
&= \alpha \mathbf{H}(\theta_0) \mathbf{R}_x(\tau) + \mathbf{E}(\tau),
\end{aligned}
\tag{2.41}
$$

where the filtered noise matrix is

$$
\mathbf{E}(\tau) \triangleq \int_{-\infty}^{\infty} \mathbf{v}(t) \mathbf{x}(t - \tau)^{\mathrm{H}} \, dt.
\tag{2.42}
$$

The $N \times M$ matrix $\mathbf{Z}$ may be vectorized, which yields

$$
\mathbf{z}(\tau; \theta_0) \triangleq \mathrm{Vec}\{\mathbf{Z}(\tau; \theta_0)\} = \alpha \mathbf{s}(\tau; \theta_0) + \mathbf{e}(\tau),
\tag{2.43}
$$

where the MIMO steering vector for a lag, τ, corresponding to a target at an angle, θ, is

$$
\mathbf{s}(\tau; \theta) \triangleq \mathbf{R}_x(\tau)^{\mathrm{T}} \mathbf{a}(\theta) \otimes \mathbf{b}(\theta),
\tag{2.44}
$$

and $\mathbf{e}(\tau)$ is the vectorized version of the filtered noise matrix, $\mathbf{E}(\tau)$.

The array factor corresponding to lag τ is found to be

$$
f(\theta_0; \tau, \theta) = \left(\frac{\mathbf{a}(\theta_0)^{\mathrm{H}} \mathbf{R}_x^*(\tau) \mathbf{a}(\theta)}{\sqrt{\mathbf{a}(\theta_0)^{\mathrm{H}} \mathbf{R}_x^*(\tau) \mathbf{a}(\theta_0)}} \right) \left(\frac{\mathbf{b}(\theta_0)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta)}{\sqrt{\mathbf{b}(\theta_0)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta_0)}} \right).
\tag{2.45}
$$

Compare this with the array factor in (2.30). This is interpreted as the (voltage) gain on a target at an angle, θ_0, provided by a beamformer steered to (τ, θ). The response of the linear beamformer that is steered to an angle, θ, and a lag, τ, to a single target at an angle, θ_0, is

$$
G(\theta_0; \tau, \theta) = \left(G_{\mathrm{TX}}(\theta_0) \frac{|\mathbf{a}(\theta_0)^{\mathrm{H}} \mathbf{R}_x^*(\tau) \mathbf{a}(\theta)|^2}{\mathbf{a}(\theta)^{\mathrm{H}} \mathbf{R}_x^*(\tau) \mathbf{a}(\theta)} \right) \left(G_{\mathrm{RX}}(\theta_0) \frac{|\mathbf{b}(\theta_0)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta)|^2}{\mathbf{b}(\theta)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta)} \right).
\tag{2.46}
$$

For fixed angles, θ and θ_0, this may be considered the range response of the waveforms. This is analogous to the waveform autocorrelation, which describes the range response of the waveform to the matched filter that is applied by a standard radar system. Observe that the contribution of the receive array is equivalent regardless of which range bin is considered. However, the transmit array factor varies as a function of range as well as angle.

For a moment, consider the single-waveform case where the "matrix," $\mathbf{R}_x$, is simply the scalar-valued autocorrelation function of the transmitted waveform. This allows the variation in range, which is due to the signal autocorrelation function, to be separated from the variation in angle, which is related to the similarity between the steering vectors. However, in the MIMO case, unless the waveforms are orthogonal, the range response can vary as a function of both angle and range in a non-separable manner. This is alluded to in [Friedlander, 2012], where the impact of range straddle loss is also discussed. Of course, this latter effect is a problem encountered by any radar system [Cann, 2002].

An interpretation of the result in (2.45) is that the cross-correlation functions of the transmitted waveforms interfere with each other in different ways depending on the angle of arrival. This impacts beamforming and direction finding [Friedlander, 2012] and the cancellation of clutter for GMTI radar [Rabideau, 2012].

2.7 Summary

Coherent MIMO radar is a natural extension of the phased array concept. Just as digital beamforming on receive provides additional degrees of freedom that improve radar performance, transmitting independent waveforms provides further flexibility. Like any new idea proposed to improve a decades-old technology, MIMO radar has not been without its critics [Daum and Huang, 2009]. Indeed, many radar applications will not benefit from MIMO and performance may be degraded by transmitting uncorrelated waveforms. In situations where performance is limited only by thermal noise, little improvement will likely be provided by MIMO. However, when performance is limited by other factors, e.g., multiplicative noise in SAR [Davis et al., 2011], clutter in GMTI [Forsythe and Bliss, 2010], ionospheric phenomenology [Frazer et al., 2009], or even synchronization challenges [Steyskal et al., 2003], a MIMO radar may outperform a traditional system.

The goal of this chapter has been to present a framework for understanding the appropriateness of a suite of MIMO waveforms for a particular radar application. The primary

contribution has been a careful analysis of the antenna performance of a MIMO radar that includes the impact on SNR as well as degradation of range sidelobe performance. This was illustrated by providing a fair comparison of the phased array to a MIMO radar that transmits orthogonal waveforms.

CHAPTER 3

MIMO SYNTHETIC APERTURE IMAGING

3.1 Overview

Synthetic aperture radars (SAR) provide fine-resolution images of the earth's surface, and synthetic aperture sonars (SAS) map the sea floor. Instead of resolving targets in angle using a large physical aperture, synthetic aperture systems use a more modest aperture and use platform motion to synthesize a large one. As the SAR or SAS system moves, it collects spatial samples using a number of pulses or pings that are coherently processed to resolve targets in the cross-range/along-track direction.

As the platform moves, it transmits pulses and receives the echoes from the illuminated scene. Each of these transmit/receive events occurs at a particular point in space and provides a spatial sample. Essentially, the synthetic aperture imaging system provides a sampled version of the desired continuous, physical aperture that is used to provide angular resolution. Pulses must be transmitted frequently enough so that the aperture is not spatially undersampled, which can lead to angle ambiguities in a manner analogous to temporal sampling. However, the time between pulses limits the maximum range extent of the image. Along-track sampling concerns constrain the performance of synthetic aperture imaging systems and are especially stressing for spaceborne SAR [Freeman et al., 2000] and for SAS [Cutrona, 1975].

Multichannnel systems that use multiple antennas on receive have been proposed to mitigate azimuth ambiguities without decreasing the area-coverage rate [Kock, 1972; Currie and Brown, 1992]. Receive array configurations where a wide transmit beam is used and the returns are collected by multiple receivers are ubiquitous in SAS. By simultaneously observing the returns from spatially separated receivers, multiple along-track samples are collected per pulse, which allows a lower pulse repetition frequency (PRF) to be used while still maintaining the required along-track sampling rate.

A more recent development is the proposed used of simultaneously transmitting independent waveforms from a number of transmit elements in an analogous manner [Krieger et al., 2008]. By simultaneously transmitting waveforms from spatially separated elements, multiple along-track samples are collected per pulse by each receive channel. This has the potential to allow for an even lower PRF to be used and enabling a further improved area coverage rate. This idea has been discussed in the SAR literature [Rennich, 2009; Ender and Klare, 2009; Wu et al., 2009; Correll, 2010; Davis et al., 2011; Wang, 2013; Krieger, 2013] as well as in the context of SAS [Davis and Cook, 2011]. Essentially, the idea is to extend the standard approach of using a SIMO system to exploit additional spatial diversity provided by MIMO.

The goal of this chapter is to present a methodology for assessing the performance of a MIMO SAR or MIMO SAS. To accomplish this, the standard techniques of SAR image quality analysis are used. In particular, the concept of multiplicative noise [Carrara et al., 1995] is extended to the MIMO case. While the average power of additive noise, like thermal noise or external interference, is independent of the desired signal power, a noise source is said to be multiplicative if the ratio of the noise power to the signal power is constant. This ratio is called the multiplicative noise ratio (MNR).

Transmitting multiple waveforms can modify the multiplicative noise in two ways. First, it can improve the ambiguity-to-signal ratio (ASR) by mitigating azimuth ambiguities through enhanced spatial sampling. However, it will also degrade the integrated sidelobe ratio (ISR), which lowers image contrast.

In this chapter, MIMO SAR/SAS is motivated in Section 3.2 as an extension of standard multichannel synthetic aperture systems that use multiple receive channels. A review of ASR is presented in Section 3.3, which leads to two MIMO SAR/SAS array configurations described in Section 3.4. Finally, the degradation of image quality through increased ISR is described in Section 3.5.

3.2 Multichannel and MIMO SAR/SAS

A synthetic aperture image is formed by collecting a number of spatial samples and processing them coherently to provide fine cross-range resolution. This is accomplished by transmitting a series of pulses from a moving platform. Each pulse corresponds to a spatial sample. These samples are also referred to as along-track samples because they are collected as the platform moves in the along-track direction.

In many ways, a synthetic aperture collection is conceptually similar to an array antenna. To unambiguously measure the angle of arrival of a signal from a far-field source, an array of uniformly-spaced, omnidirectional elements must be spaced by no more than a distance $\lambda/2$ where λ is the signal's wavelength. This half-wavelength spatial-sampling requirement is analogous to the Nyquist sampling rate required to unambiguously sample a signal with a particular maximum frequency. A synthetic aperture imaging system will employ a directional element on receive, which relaxes the sampling requirement to $D/2$ where D is the along-track extent of the receive aperture [Brown, 1967].

The system transmits pulses at the PRF, which may be referred to as the ping rate in sonar. For a fixed platform velocity, v, the PRF, f_p, determines the along-track sampling interval,

$$\delta_\mathrm{x} \triangleq \frac{v}{f_\mathrm{p}}, \tag{3.1}$$

which is the distance traveled between pulses. The PRF must be sufficiently high to provide an acceptable along-track sampling rate, i.e.,

$$\delta_\mathrm{x} \leq D/2. \tag{3.2}$$

However, the PRF must be low enough to provide an acceptable area-coverage rate (ACR). The extent of the range swath is limited by the ability of the sensor to unambiguously measure range. The ACR is maximized in stripmap operation, where the system continuously maps as it flies without reorienting the antenna. In this case, the ACR is

related to the platform speed, v, and the range extent of the image, R_{swath} by

$$\text{ACR} \triangleq vR_{\text{swath}}. \tag{3.3}$$

If the sensor transmits pulses at a PRF, f_{p}, the range swath is limited to $R_{\text{swath}} < (c/2)/f_{\text{p}}$, where c is the speed of waveform propagation. This leads to a bound,

$$\text{ACR} < \delta_{\text{x}} (c/2). \tag{3.4}$$

From (3.4), it is evident that there is a fundamental limitation imposed on ACR by the along-track sampling rate, $1/\delta_{\text{x}}$. There is also a dependence on c, which suggests that SAS systems are more challenged than SAR systems because the speed of light is about five orders of magnitude higher than the speed of sound in water.

A low PRF is desirable to maximize ACR, but sufficient along-track sampling is still required. The aperture size, D, can be increased to relax the along-track sampling requirement of (3.2), but this is undesirable for a number of reasons. First, it is counter to the motivation of synthetic aperture imaging, which is to provide the benefits of a large, physical aperture with a smaller one. Also, increasing the aperture size leads to a decreased beamwidth. In stripmap synthetic aperture imaging, this results in coarser cross-range resolution because of limited integration angle; in spotlight synthetic aperture imaging, it results in a smaller image.

Instead of requiring larger and larger apertures, the traditional solution is to employ an array of receive elements [Kock, 1972; Cutrona, 1975]. Each element is small enough to provide acceptable stripmap cross-range resolution or spotlight image size. By using multiple receive elements, the along-track sampling rate can be improved without increasing the PRF. This is accomplished because multiple along-track samples are simultaneously collected for each pulse. In this configuration, a single element is used on transmit. On receive, an array of N elements is used. If D is the length of an element, then the effective array length is ND. This array requires the along-track sampling rate corresponding to

an element of length *ND* but can still provide the cross-range resolution of an element of length *D*.

A synthetic aperture system that uses multiple elements on receive can be considered as a SIMO system. The scene is illuminated by a single transmit source, but the returns are observed by multiple receivers. This concept of a SIMO system can be extended to the MIMO case by also simultaneously transmitting independent waveforms from multiple transmitters. In either case, multiple along-track samples are collected with a single pulse. The spatial locations of these along-track samples can be considered as the virtual array corresponding to the physical transmit and/or receive arrays.

3.3   MIMO SAR/SAS Ambiguity-to-Signal Ratio

In (3.2), an upper bound on the along-track sampling interval was provided, but it does not guarantee that azimuth-ambiguous returns will be completely mitigated. A synthetic aperture system synthesizes a large, continuous aperture by collecting spatial samples. As is always the case when sampling a signal, the spatial sampling rate must be sufficiently high to prevent aliasing. In temporal sampling, out-of-band signals alias into the desired frequency band. In spatial sampling, targets displaced in cross-range can alias into the image. The beampattern acts as a spatial antialiasing filter, which can suppress these undesired returns, but this filtering is far from ideal. The impact of ambiguous returns is quantified by the ambiguity-to-signal ratio (ASR). The ASR depends on the beampattern as well as the along-track sampling rate [Bayma and McInnes, 1975; Mehlis, 1980; Hawkins, 1996].

A simplified interpretation of synthetic aperture imaging is that it resolves targets in the cross-range direction by applying Doppler processing. For a stationary target, the Doppler frequency, f_D, that is due to platform motion is related to the angle to the target, θ, by

$$f_D = \frac{2v}{\lambda} \sin \theta. \tag{3.5}$$

As the platform moves, it samples the slow-time signal at a rate that is precisely the PRF, f_p. Azimuth ambiguities are introduced by targets with Doppler frequencies, f_D, that are

43

Doppler ambiguous, which occurs when $|f_D| > f_p/2$. Range ambiguities (and their azimuth aliases) are also introduced. For a well-designed SAR or SAS system, these ambiguities will be far from the peak of the beampattern and will be attenuated by the antenna/hydrophone.

A standard approach to calculating SAR and SAS ASR is to compute the aliased Doppler power spectrum [Mehlis, 1980]. For a given range, the Doppler power spectrum, $P(f_D)$, quantifies the relative power from targets at a given Doppler frequency, f_D. For range-unambiguous targets, this depends primarily on the transmit and receive beampatterns. The contribution of range ambiguities to the Doppler power spectrum also depends on the collection geometry, signal attenuation from spherical spreading, and (especially for SAS) the attenuation from propagation of the waveform through a medium.

The ASR depends not only the PRF, but also the cross-range resolution. If the platform flies in a straight line, the Doppler frequency corresponding to a point on the ground will vary over the synthetic aperture. If the integration angle is increased, the degree of variation will increase as well. For a stripmap synthetic aperture imaging system, the finest cross-range resolution is provided by processing the full Doppler bandwidth that corresponds to the mainlobe of the beampattern. A coarser resolution image can be formed using a subset of this Doppler bandwidth. This can potentially improve the ASR by excluding Doppler frequencies that may include significant aliased energy.

The azimuth ambiguity-to-signal ratio (AASR) is the ratio of the power from azimuth ambiguous returns to the power of the desired return. If $P(f_D)$ is the Doppler power spectrum for a given range, then the AASR that results when a PRF, f_p, is used is

$$\text{AASR} \triangleq \frac{\int_\Omega \left(\sum_{n \neq 0} P\left(f_D + n f_p\right) \right) df_D}{\int_\Omega P(f_D) \, df_D}, \tag{3.6}$$

where Ω is the set of Doppler frequencies that are processed to form the image. The AASR formula in (3.6) can be modified in a straightforward manner to include range ambiguities

to provide the ASR. Note that the Doppler PSD is assumed to be zero for Doppler frequencies that exceed the maximum Doppler frequency, $2v/\lambda$, so the apparently infinite sum of (3.6) is only over a finite number of Doppler ambiguities.

The methodology summarized by (3.6) allows the AASR to be computed for a given array design. It assumes that the along-track samples are uniformly spaced and is based on a stripmap imaging interpretation. This result is extended to the multichannel case where nonuniform sampling is possible as well as to the spotlight case in Chapter 6. In any case, a low PRF is desirable to maximize ACR, increasing along-track sampling can improve image quality by providing a lower ASR.

3.4 Array Design for MIMO SAR and SAS

Two MIMO array configurations suggest themselves as being useful in a SAR system: a dense MIMO array that seeks to improve the density of along-track samples and a sparse MIMO array that allows a PRF than the corresponding SIMO array [Davis and Cook, 2011]. These configurations and their corresponding virtual arrays are presented in Figure 3.1.

The properties of the virtual arrays for these designs are given in Table 3.1. Note that SIMO is a special case of either a dense MIMO array or a sparse MIMO array. The results in the table for the MIMO configurations coincide with the SIMO case when $M = 1$.

For the same PRF and cross-range resolution, the dense MIMO configuration using M transmit elements provides an increase in the spatial sampling rate by a factor of M, which can potentially improve image quality by lowering the ASR. Similarly, the sparse MIMO configuration using M transmit elements can potentially increase the ACR by a factor of M while preserving cross-range resolution.

First, consider the dense MIMO array of Figure 3.1. This configuration uses a receive array that consists of N receive elements that are of length D and spaced at an interval of D. The system uses M transmit elements that are spaced by D/M. The corresponding virtual
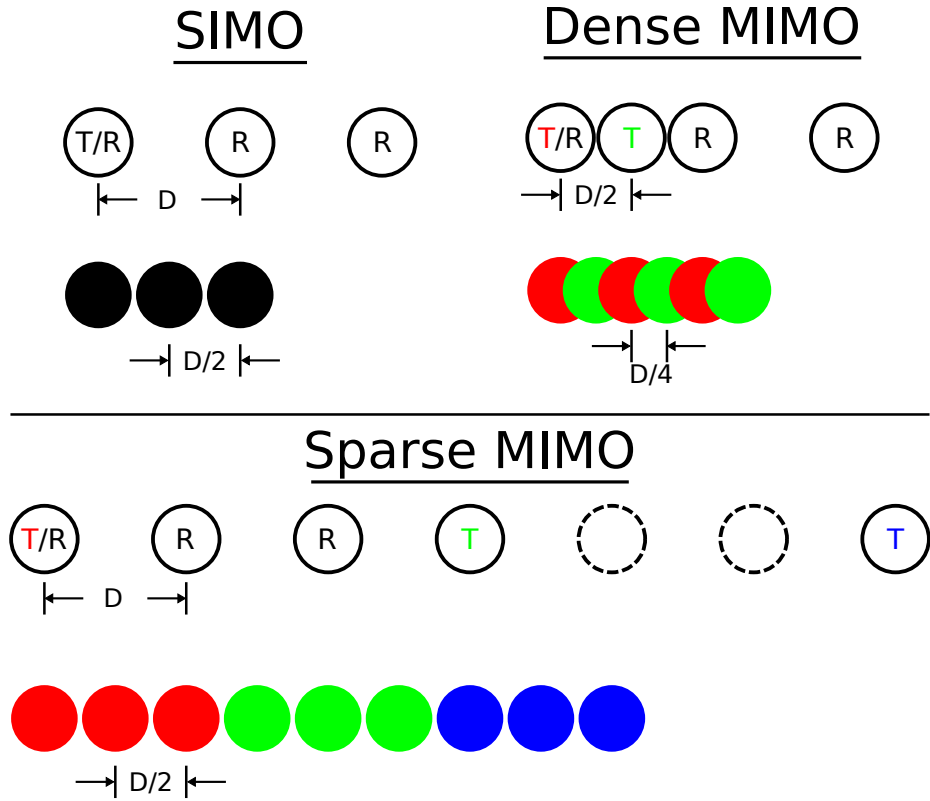
Figure 3.1: Examples of SAR/SAS physical arrays and corresponding virtual arrays.

array is sampled at an interval of $1/M \times D/2$. Note that the effective length of the dense MIMO array is the same as the corresponding SIMO array. However, a higher along-track sampling rate can be achieved by the dense MIMO array if it uses the same PRF as the SIMO array.

Now, consider the sparse MIMO array. Once again, begin with a receive array of N elements. Now, distribute the M transmitters such that there is a separation of ND between

Configuration	Length	Spacing
SIMO	$ND/2$	$D/2$
Dense MIMO	$ND/2$	$D/(2M)$
Sparse MIMO	$MND/2$	$D/2$

Table 3.1: Properties of the virtual arrays.

transmit elements. In this case, the spacing between virtual phase centers is $D/2$, as in the SIMO array case, but the resulting virtual array is M times as long as that of the SIMO virtual array. This allows the sparse MIMO array to use a PRF that is M times lower than the SIMO case, which provides a commensurate increase in area coverage rate. This configuration is called a Nyquist array in [Forsythe and Bliss, 2010] because it is critically sampled in the spatial sense.

3.5 MIMO SAR/SAS Integrated Sidelobe Ratio

As described above, a MIMO synthetic aperture imaging system will simultaneously transmit multiple waveforms to provide improved along-track sampling. However, one challenge is that this will lead to an increase in range sidelobes because of the inevitable cross-correlation between the waveforms, which will lead to a degraded integrated sidelobe ratio (ISR).

3.5.1 Quantifying the Impact of Sidelobes

Range sidelobes result from undesired energy that appears up range and down range of any target return. Sidelobes of strong targets can limit detection of smaller targets nearby and can cause extraneous false detections. The impact of sidelobes on detection is quantified by the peak sidelobe ratio (PSR) of the waveform, which is the power of the strongest sidelobe relative to the peak of the mainlobe.

Sidelobes can also degrade the quality of radar images. To form an image, the radar resolves targets in down range and cross range to provide an estimate of the reflectivity of the scatterers that are contained within each resolution cell. Range sidelobes from terrain and targets up and down range of a particular resolution cell will contribute undesired energy. For example, consider a low return area, e.g., a shadowed region or a road. A high-quality image will have significant contrast between these areas and higher-return areas that are nearby, but this contrast will be degraded by range sidelobes.

While a low PSR is desirable, imaging radar systems require low sidelobe levels over

the entire range response. Moving target indication (MTI) radars typically operate in environments where targets are relatively sparse, but a synthetic aperture radar (SAR) images distributed targets including terrain.

Just as PSR is a useful metric for predicting the impact of range sidelobes on MTI performance, the ISR is used to predict their impact on SAR image quality. ISR can be treated as component of the SAR MNR, and the integrated sidelobes can severely degrade image contrast [Carrara et al., 1995]. Historically, integrated sidelobe levels have not driven SAR image quality because a linear frequency modulated (LFM), "chirped" waveform has been used, which permits aggressive windowing to drive down the sidelobes. For example, a Taylor weighting can be applied to drive the ISR to about -30 dB. This is well below the multiplicative noise floor set by other contributors such as azimuth ambiguous returns.

The impact of a waveform with poor ISR on SAR image quality is shown in Figure 3.2. Results from an unweighted LFM are compared to a random phase-coded waveform. In both cases, a no-return area (NRA) is simulated between two areas of terrain. This occurs, for example, if a road, which has low radar backscatter, crosses a grassy area. Ideally, there should be significant contrast between the regions. However, the range sidelobes of terrain up and down range of the NRA will fill in the NRA and degrade contrast. The degree to which this impacts image quality is captured by the ISR. In this example, we see that the LFM preserves the NRA better than the phase-coded waveform because the ISR of the LFM is much lower.

3.5.2 Calculating the MIMO ISR

A MIMO radar/sonar signal processor was presented in Figure 2.3. The signal from each of N receive channels is processed by applying M filters. Each of these filters is matched to one of the transmitted waveforms. If the waveforms were orthogonal, the output of each filter would consist of the contribution from a single transmit element. The contribution of a single target at an angle, θ_0, to an angle bin corresponding to an angle, θ, and a range bin corresponding a delay, τ, was given in (2.46). Neglecting the terms that are constant with

48

Figure 3.2: Comparison of LFM and phase-coded waveform range profiles. The dotted green lines indicate a no-return area (NRA) with range extent equal to the pulse width. Ideally, there should be no energy in this region, but the sidelobes from terrain up and down range of the NRA will be present and degrade contrast. The expected sidelobe levels are indicated by the superimposed lines. The NRA is better preserved by using the LFM waveform, which has superior ISR compared to the phase-coded waveform.

range, this range response can be written as

$$P(\tau, \theta; \theta_0) = \left| \mathbf{a}(\theta)^{\mathrm{H}} \mathbf{R}_x^*(\tau) \mathbf{a}(\theta_0) \right|^2 . \tag{3.7}$$

This expression for the MIMO impulse response can be used to calculate the ISR for the MIMO case.

The lag, τ, in the impulse response of (3.7) is related to a (slant) range displacement from the true target location. If the output is sampled at set of lags, then each of these lags corresponds to the return in particular range bin. Ideally, all of the energy in the impulse response will be in the bin corresponding to $\tau = 0$, and there will be no energy in bins corresponding to $\tau \neq 0$.

ISR is the ratio of the undesired sidelobe energy to the energy in the mainlobe of the range response,

$$\mathrm{ISR} \triangleq \frac{\displaystyle\int_{-\infty}^{-T_0/2} P(\tau, \theta; \theta_0) \, d\tau + \int_{T_0/2}^{\infty} P(\tau, \theta; \theta_0) \, d\tau}{\displaystyle\int_{-T_0/2}^{T_0/2} P(\tau, \theta; \theta_0) \, d\tau} . \tag{3.8}$$

where T_0 is the null-to-null width of the mainlobe of the range response, which is related to the bandwidth of the transmitted waveforms. In general, the MIMO ISR will vary with target angle, θ_0, and steered angle, θ.

The ISR is of the same form as the ASR, which is the ratio of the energy from the undesired ambiguous returns to the energy from the desired unambiguous returns. The sum of these two terms is added to similar ratios for other multiplicative noise sources, e.g., quantization noise-to-signal ratio (QNSR), to determine the overall system MNR. Any benefits to the ASR gained by transmitting multiple waveforms may be canceled out if it comes at the cost of increased ISR.

3.5.3 ISR of Chirp-Slope Multiplexing

Many radar systems have historically employed linear frequency modulated (LFM) waveforms. These waveforms are easy to generate, their sidelobe characteristics may be improved through aggressive windowing, and they allow radar designers to relax receiver bandwidth requirements by employing stretch processing. LFM waveforms are ubiquitous in radar especially in SAR and GMTI radar systems.

A natural extension of LFM to multiple waveforms is to change the slope from signal to signal. To obtain two approximately orthogonal waveforms, one might consider an up chirp and a down chirp. Figure 3.3 shows their time-frequency representations. The autocorrelation and cross-correlation properties of these two waveforms are shown in Figure 3.4. At the zero-lag, the cross-correlation is down by almost 30 dB from the peak, which is well below the first sidelobe. From the range response, the PSR is seen to not increase by transmitting the second LFM. Unfortunately, further away from the mainlobe, the cross-correlation prevents the sidelobes from decaying to low levels. Consequently, transmitting the second LFM greatly degrades the ISR performance.

While these waveforms may be acceptable for an MTI application, they are likely unsuitable for SAR or other radar applications that seek to estimate the reflectivity profile of a continuum of scatterers. The ISR approaches 0 dB, meaning that the contrast of an image with distributed clutter will have very low contrast. It is likely that MIMO operation will require a departure from LFM waveforms and the adoption of phase-coded or noise-like waveforms for these cases.

3.6 Summary

An analysis of a SAR and SAS system using multiple transmitters has been presented. Just as the SIMO configuration common to most SAS systems uses multiple spatial channels on receive to increase area coverage rate, a MIMO SAR or MIMO SAS system adds multiple spatial channels on transmit to provide even more substantial performance gains. A SIMO system using N receive channels is able to effectively collect N times as much data per pulse, while a MIMO system with M transmit channels and the same receive array can potentially capture $M \times N$ as much data.

Two MIMO SAR/SAS configurations were proposed. By using M transmit elements that emit orthogonal waveforms, the dense MIMO configuration allows the effective along-track sampling rate to be increased by a factor of M relative to a SIMO system with the same area coverage rate and the same cross-range resolution. The sparse MIMO configuration allows the area coverage rate to be increased by a factor of M relative to a SIMO array with the same azimuth ambiguity-to-signal ratio and cross-range resolution. Note that the
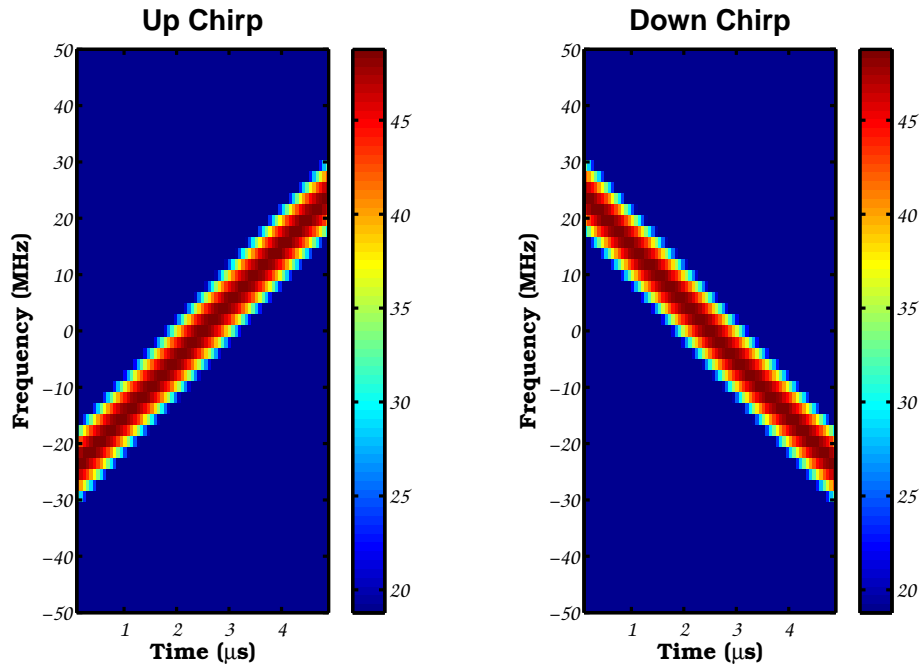
Figure 3.3: Spectrograms of an up chirp and a down chirp.

Figure 3.4: Correlation properties and range response of an up chirp and a down chirp. The upper figure presents the autocorrelation of each chirp and their cross-correlation. The bottom figure presents the range response if these two waveforms are transmitted simultaneously.

MIMO SAR/SAS and comparable SIMO array transmit the same amount of power.

It is important to note that while *M* orthogonal waveforms are necessary to achieve these full factor of *M* improvements, performance benefits can also be gained by using

"nearly" orthogonal waveforms. The framework established in Chapter 2 provides the tools for analyzing the performance characteristics of a MIMO SAR/SAS system given its MIMO signal correlation matrix even if the waveforms are not perfectly orthogonal.

A key challenge to the realization of a MIMOSAS system is the generation of quasi-orthogonal waveforms. Because a suite of waveforms can never be mutually orthogonal for all delays, the transmission of multiple waveforms will undoubtedly degrade the ISR of the SAR/SAS system. Detailed analysis must weigh the benefit, for example, of decreasing ASR at the cost of increasing ISR. Still, a MIMO SAR/SAS system will be able to exploit this tradespace more effectively than a system with a single transmit element.

CHAPTER 4

CONTROLLING RANGE SIDELOBES IN MIMO RADAR

4.1 Overview

Multiple-input, multiple-output (MIMO) radars [Bliss and Forsythe, 2003; Li and Stoica, 2008; Davis, 2012] promise to outperform traditional systems in many applications by simultaneously transmitting a number of waveforms from spatially-distinct radiating elements. While these multiple waveforms provide additional degrees of freedom that may be exploited to improve radar performance, one drawback is the increase in range sidelobes from the unavoidable cross-correlation between the radar waveforms.

Traditional radar signal processors employ a matched filter. This maximizes the signal-to-noise ratio (SNR) in the target range bin, but it does not consider the sidelobe energy that will be present in nearby range bins. A popular approach to range-sidelobe suppression is to use a so-called mismatched filter [Ackroyd and Ghani, 1973], which reduces the sidelobes at the cost of SNR. This chapter extends the concept of optimal filtering, which has been developed for a single waveform [Keel and Baden, 2012], to the case of multiple waveforms. In particular, the minimum integrated sidelobe ratio (ISR) filter is derived for MIMO radar.

Radar waveforms are carefully designed so that their matched filter responses have low range sidelobes. However, there are limits to which their peak sidelobe ratio (PSR) and ISR can be reduced, especially when waveforms are constrained to be constant amplitude. For a single waveform, these limitations are imposed by bounds on autocorrelation functions, and they are further exacerbated in the MIMO case where cross-correlation functions must also be considered [Sarwate, 1979; Welch, 1974]. This motivates a departure from a matched filter, which maximizes SNR, to a mismatched filter, which additionally suppresses sidelobes.

A simple approach to mismatched filtering that is typically used with linear frequency

modulated (LFM) waveforms is to apply a windowing function, which lowers sidelobe but also results in degraded SNR and degraded range resolution through a broadening of the mainlobe. There has also been extensive research into mismatched filtering techniques for phase-coded waveforms [Rihaczek and Golden, 1971; Ackroyd and Ghani, 1973; Zoraster, 1980; Griep et al., 1995]. In either case, the matched filter is either followed by or replaced by a filter that is designed to implement range compression while minimizing the PSR or ISR.

Just as in the single-waveform case, mismatched filtering techniques can be applied to MIMO radar. Approaches to this have been presented in the literature before [Li et al., 2008; Hu et al., 2010; Ma et al., 2010; Zou et al., 2011; Hua and Abeysekera, 2013], but all of these are motivated by the assumption that the waveforms are orthogonal. In this case, a natural signal processing architecture consists of a set of filters that are designed to extract the contribution of a single waveform while rejecting the response from the other waveforms. These filters can be independently designed to optimize the response for a single waveform. The outputs of the filters are then combined to apply beamforming. Like [Ma et al., 2010; Zou et al., 2011; Hua and Abeysekera, 2013], the focus of this discussion is on controlling the integrated sidelobe level while [Li et al., 2008; Hu et al., 2010] aim to control the peak sidelobes.

The proposed approach is based on the observation that the MIMO range response varies with angle. This fact motivates the idea of jointly performing range compression and beamforming rather than applying these processes separably as is typically advocated. This new method provides a significant advantage in minimizing ISR without significant SNR loss when compared to other approaches that are found in the literature.

4.2 An Architecture for MIMO Radar Processing
4.2.1 The MIMO Range Response as a Function of Angle

Consider a MIMO radar that uses M radiating elements that each transmit an independent waveform. The signal that is incident upon a target at an angle, θ_0, in the far field of the

radar will be a linear combination of these waveforms,

$$s(t; \theta_0) \triangleq \mathbf{a}(\theta_0)^{\mathrm{T}} \mathbf{x}(t) \tag{4.1}$$

where $\mathbf{a}(\theta)$ is the steering vector for an angle, θ, which describes the relative phase of the signals as the arrive at the target, and $\mathbf{x}(t)$ is the vector of transmitted waveforms. A portion of this incident waveform will be reflected by the target and observed by each of the N receive elements.

The signal observed by each receiver will be identical up to an amplitude scaling (if the elements are not identical) and a delay. Because of this, exploiting multiple receive channels provides no value in terms of range-sidelobe mitigation and only provides angular resolution. Instead of jointly processing the data from all of the receive channels, the process of range compression and receive beamforming should be performed separately.

Also, because all targets at an angle, θ_0, are illuminated by the same linear combination of waveforms, $s(t; \theta_0)$, given by (4.1), this is equivalent to a radar that transmitted the single waveform, $s(t; \theta_0)$. This effective waveform changes as a function of target angle, but the radar antenna will suppress targets from other angles. This fact motivates the processing architecture of Figure 4.1(a). A receive beam is formed in a desired direction, which combines the N receive channels into a single channel that contains the contributions of the M transmitters. Then, a filter is applied to jointly perform range compression and transmit beamforming. This filter can be designed to optimize the range response for targets that are present at the angle corresponding to the beam.

4.2.2 Single-Filter vs. Multiple-Filter Architectures

An alternative approach that is based on the standard MIMO radar signal processor of Figure 2.3 is shown in Figure 4.1(b). As in Figure 4.1(a), the N channels are first combined into a single channel by applying receive beamforming. This signal contains the contributions of the M transmitted waveforms. To extract each contribution, a bank of filters, $h_0(t), \ldots, h_{M-1}(t)$, is applied. Each filter, $h_m(t)$, is designed so that it provides a
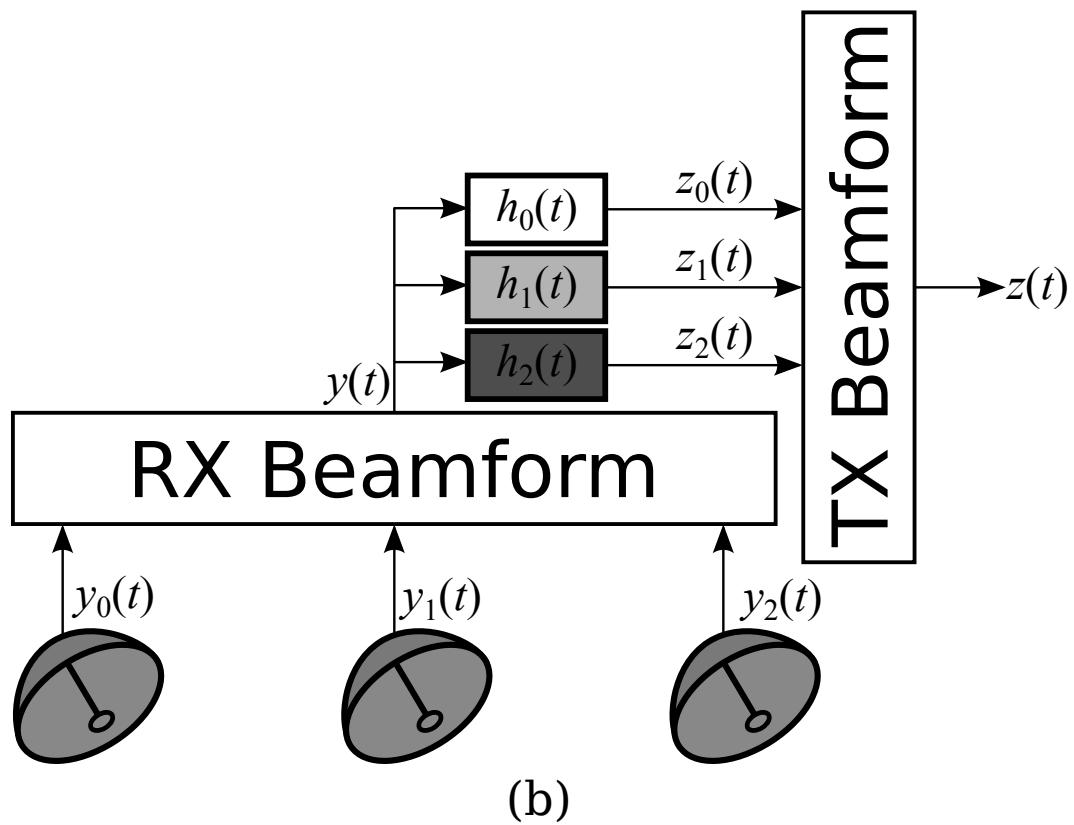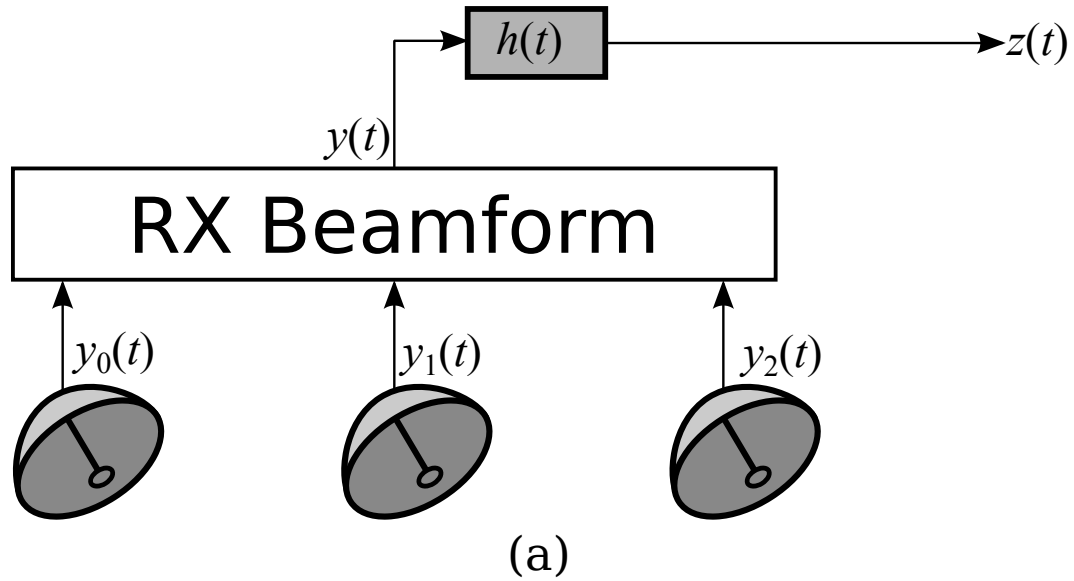
Figure 4.1: (a) Proposed MIMO radar signal processing architecture. The filter, $h(t)$, is designed to optimize the response for a target at a given angle and jointly implements transmit beamforming and range compression. (b) Standard MIMO radar signal processor. The range compression filters, $h_m(t)$, are independent of target angle.

low-sidelobe range response to a particular waveform, $x_m(t)$, and also has a low cross-correlation with the other transmitted waveforms, $x_{m'}(t)$, with $m \neq m'$. The outputs of the M filters are combined to apply transmit beamforming. Both architectures in Figure 4.1 provide an output signal, $z(t)$, that is a range profile for a particular angle.

The MIMO radar signal processors of Figure 4.1(a) and Figure 4.1(b) can be referred to as a single range compression filter and a multiple range compression filter architecture, respectively. It is expected that the single range compression filter implementation will be able to provide better sidelobe control because it is only required to optimize the response to a single waveform, $s(t)$, but the multiple filter approach requires that M filters be designed and that each of these filters provides a desired response to M waveforms. A potential drawback of the single filter approach is that there will be degraded performance for targets at other angles. This may be mitigated by designing the filter over a range of angles.

Another distinction to be made between the two approaches is that the single minimum-ISR range compression filter of Figure 4.1(a) must be redesigned for each beam that is formed, i.e., for each set of spatial weights that is used to steer a beam to a particular angle. Because the filters can be designed offline and only a relatively small number of beams are required, this requirement should not challenge the implementation of this mismatched filtering approach. The standard MIMO radar processor in Figure 4.1(b) requires that M range compression filters be applied. The proposed architecture in Figure 4.1(a) requires that only a single range compression filter be applied, but it must be applied for each beam. In most radar applications, the number of digital beams formed does not exceed the number of channels. For multibeam SAR, the number of beams typically equals the number of channels [Currie and Brown, 1992]. For beamspace STAP, the number of beams is at most equal to the number of channels [Wang and Cai, 1994], but fewer beams are formed in many implementations. For a MIMO system with M transmit elements and N receive elements, potentially MN virtual channels can be synthesized. If all MN beams are formed, then the number of range compression filters that must be applied increases by a factor of N with

the approach of Figure 4.1(a) compared to that of Figure 4.1(b).

Each of the filters in Figure 4.1(b) is designed to extract the response from one of the transmitted waveforms. These filters are independent of target angle because the transmit beamforming and range compression operations are separated. In [Li et al., 2008], a method is proposed to minimize SNR loss while satisfying a peak sidelobe criterion. This approach is extended in [Ma et al., 2010], where an integrated sidelobe criterion is used. The ability to get extremely low sidelobes over a limited range extent at the expense of elevated sidelobes over other regions is also included. As noted in [Ma et al., 2010], a restrictive limit of optimizing over about $(K/M - 1)/2$ of the K range bins constrains this approach. Further contributions were made in [Hua and Abeysekera, 2013] with an extension to include Doppler effects.

The approach to generating minimum ISR filters developed in this chapter fundamentally differs from those in [Ma et al., 2010] and [Hua and Abeysekera, 2013] in two primary ways. First, the proposed method is based on the key observation that, for a given angle, the range response for a particular linear combination of transmitted waveforms needs to be optimized, while the other approaches in the literature develop filters that simultaneously minimize the sidelobes for all linear combinations of the transmitted waveforms. By combining the operations of beamforming and range compression, better sidelobe control is possible. A second difference is that the technique developed here allows an arbitrarily long filter to be used, which provides further sidelobe control. For both of these reasons, the single-filter technique provides a powerful tool for controlling range sidelobes in MIMO radar.

4.2.3 The MIMO Mismatched Filter Range Response

The range response of the single filter approach presented in Figure 4.1(a) is now derived, which is analogous to the matched filter range response of (2.46). Suppose that a target is present at an angle, θ_0, and that a receive beam is formed for an angle, θ, and a single filter, $h(t; \theta)$, is used to jointly perform range compression and transmit beamforming. The

power of the contribution from this target in the mismatched filter processor output at a lag, τ, and a Doppler shift, v, is

$$
\begin{aligned}
P(\tau, v, \theta; \theta_0) &\triangleq \left| \int_{-\infty}^{\infty} \left(\mathbf{b}(\theta)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{y}(t) \right) h(\tau - t; \theta) e^{-ivt} \, dt \right|^2 \\
&= \left| \int_{-\infty}^{\infty} \left(\mathbf{a}(\theta_0)^{\mathrm{T}} \mathbf{x}(t) \right) h(t - \tau; \theta) e^{-ivt} \, dt \right|^2 \left| \mathbf{b}(\theta)^{\mathrm{H}} \mathbf{R}_v^{-1} \mathbf{b}(\theta_0) \right|^2,
\end{aligned}
\tag{4.2}
$$

where $\mathbf{a}(\theta)$ and $\mathbf{b}(\theta)$ are the transmit and receive steering vectors, respectively, and $\mathbf{R}_v$ is the spatial covariance matrix of the interference and noise.

The range sidelobes depend on the result of convolving the filter impulse, $h(t; \theta)$, which is designed for a particular angle, θ, with a Doppler-shifted, linear combination of waveforms, where the linear combination depends on the target angle, θ_0, through the transmit steering vector, $\mathbf{a}(\theta_0)$. Compare this to the matched filter response in (2.46) where the sidelobes depend on the MIMO signal correlation matrix, $\mathbf{R}_x(\tau, v)$, through

$$
\left| \mathbf{a}(\theta)^{\mathrm{H}} \mathbf{R}_x^*(\tau, v) \mathbf{a}(\theta_0) \right|^2.
\tag{4.3}
$$

While the output of the standard processor of Figure 2.3, which uses matched filters, depends on the MIMO signal correlation matrix, the filter, $h(t)$, of Figure 4.1(a) can be freely to designed to minimize the sidelobes. If matched filters are used, sidelobes can only be controlled through waveform design. By removing this constraint and introducing mismatched filters, sidelobes can be mitigated through filter design as well.

4.3 Minimum-ISR Filtering for a Single Waveform

First, the filter response is defined, and the relevant figures of merit, ISR and SNR loss, are defined. This is followed by a brief review of an approach to determine the optimal filter for the single waveform case. The development of the direct method for determining the minimum ISR filter for a single waveform follows [Keel and Baden, 2012], but an iterative method and a new SNR loss constraint are also added.

4.3.1 Signal Model

Let $s(t)$ be the radar waveform, and let $\{s_k\}$ be the discrete-time representation of this signal with $s_k \triangleq s(kT)$, for a sampling interval, T. These fast-time samples of the waveform are assumed to correspond to a single pulse. Suppose that the waveform can be represented by K samples so that $s_k = 0$ for $k \notin \{0, 1, \ldots, K-1\}$. A digital filter, $\{w_\ell\}$, with L taps ($L \geq K$) is applied to implement range compression. The goal is to find the filter taps that minimize the resulting ISR.

While the matched filter is the same length as the waveform ($L = K$), long mismatched filters ($L > K$) can be designed that provide very low ISR [Levanon, 2005]. While there is no theoretical limitation on the length of the mismatched filter, there are some practical concerns. These include the computational costs of applying long digital filters and edge effects to account for filter roll in. Also, the number of sidelobes will be $K + L - 1$. While long mismatched filters can strongly suppress the sidelobes, strong target returns can create sidelobes significantly displaced from their true location.

The response of the length-L filter, $\{w_\ell\}$, to the length-K waveform, $\{s_k\}$, is the convolution of the two sequences,

$$z_k \triangleq \sum_{\ell=0}^{L-1} w_\ell s_{k-\ell} \tag{4.4}$$

for $k = 0, 1, \ldots, K + L - 1$. Each sample of the output, z_k, is referred to as the output for a particular range bin. The filter output can be written as a length-$(L + K - 1)$ vector, $\mathbf{z}$. This is generated from the length-L filter vector, $\mathbf{w}$, by

$$\mathbf{z} = \mathbf{\Phi}\mathbf{w}, \tag{4.5}$$

where $\mathbf{\Phi}$ is the $(K + L - 1) \times L$ convolution matrix of $\mathbf{s}$. For example, if $K = 3$ and $L = 4$,

$$\begin{pmatrix} z_0 \\ z_1 \\ z_2 \\ z_3 \\ z_4 \\ z_5 \end{pmatrix} = \begin{pmatrix} s_0 & 0 & 0 & 0 \\ s_1 & s_0 & 0 & 0 \\ s_2 & s_1 & s_0 & 0 \\ 0 & s_2 & s_1 & s_0 \\ 0 & 0 & s_2 & s_1 \\ 0 & 0 & 0 & s_2 \end{pmatrix} \begin{pmatrix} w_0 \\ w_1 \\ w_2 \\ w_3 \end{pmatrix}. \tag{4.6}$$

The convolution matrix, $\mathbf{\Phi}$, is Toeplitz.

4.3.2 Figures of Merit

The ISR is defined as the ratio of the energy in the sidelobes to the energy in the peak of the filter response. The matched filter maximizes SNR in the target range bin, but a minimum-ISR filter seeks to control the sidelobes, which comes at the expense of SNR. This tradeoff is characterized by evaluating the ISR and the SNR loss of a mismatched filter.

The ISR characterizes the total sidelobe energy of the filter response, $\{z_k\}$, which was defined by (4.4). To calculate the ISR, the ratio of the the power in the sidelobes to the power in the peak of the response is calculated,

$$\text{ISR} \triangleq \frac{\displaystyle\sum_{\substack{k=0 \\ k \neq p}}^{K+L-2} |z_k|^2}{|z_p|^2}, \tag{4.7}$$

where the desired location of the peak, p, is

$$p \triangleq \left\lfloor \frac{K + L}{2} \right\rfloor - 1. \tag{4.8}$$

This choice of the peak location is convenient because the peak will be centered within the range response, which consists of $K + L - 1$ samples.

A filter's SNR loss, L_s, is the loss in SNR relative to the matched filter,

$$L_\mathrm{s} \triangleq \frac{\text{SNR}}{\text{SNR}_{\max}} = \frac{|z_p|^2}{\|\mathbf{w}\|^2 \, \|\mathbf{s}\|^2}. \tag{4.9}$$

The SNR loss is limited to $0 \leq L_s \leq 1$ because $\text{SNR} \leq \text{SNR}_{\text{max}}$. An SNR loss of $L_s = 1$ indicates that there is no SNR loss compared to the matched filter. A low ISR and an SNR loss as close as possible to 1 (= 0 dB) are desirable.

4.3.3 A Direct Method for Designing the Optimal Filter

To minimize ISR, a filter should provide a response that is as close as possible to the ideal response where all but one sample is zero. This desired response may be written as

$$d_k \triangleq \begin{cases} 1, & k = p \\ 0, & k \neq p \end{cases} \tag{4.10}$$

where p is the desired location of the peak, which was defined as (4.8).

The sum squared error between the response, $\mathbf{z}$, and the desired response, $\mathbf{d}$, is

$$\epsilon \triangleq \|\mathbf{z} - \mathbf{d}\|^2 = \|\mathbf{\Phi}\mathbf{w} - \mathbf{d}\|^2 . \tag{4.11}$$

Ideally, the weights would satisfy $\mathbf{\Phi}\mathbf{w} = \mathbf{d}$ for $\mathbf{w}$ so that $\epsilon = 0$. This will require solving an overdetermined system of equations, but the least-squares solution can be used, which minimizes the error, ϵ. The optimal weights are

$$\mathbf{w} = \left(\mathbf{\Phi}^{\text{H}}\mathbf{\Phi}\right)^{-1} \mathbf{\Phi}^{\text{H}}\mathbf{d}. \tag{4.12}$$

The $L \times L$ matrix, $\mathbf{\Phi}^{\text{H}}\mathbf{\Phi}$, is invertible unless the waveform is zero for all samples ($s_k = 0$ for $k = 0, 1, \ldots, K - 1$).

The optimal weight vector of (4.12) is recognized as column p of the pseudoinverse of the convolution matrix, where the pseudoinverse of a matrix, $\mathbf{A}$, with linearly independent columns is

$$\mathbf{A}^+ = \left(\mathbf{A}^{\text{H}}\mathbf{A}\right)^{-1} \mathbf{A}^{\text{H}}. \tag{4.13}$$

Because $\mathbf{\Phi}^{\text{H}}\mathbf{\Phi}$ is a Toeplitz matrix when $\mathbf{\Phi}$ is a convolution matrix, the minimum ISR filter may be efficiently computed using, e.g., the Levinson algorithm [Golub and Van Loan, 1996], by solving the system of equations,

$$\left(\mathbf{\Phi}^{\text{H}}\mathbf{\Phi}\right)\mathbf{w} = \mathbf{\Phi}^{\text{H}}\mathbf{d}, \tag{4.14}$$

for the optimal weight vector, $\mathbf{w}$.

4.3.4 An Iterative Method for Designing the Optimal Filter

A direct method for designing the optimal filter is given by (4.12). Alternatively, an iterative approach can be used. This may be more numerically stable than the direct approach and also provides access to intermediate filters, which allows tradeoffs between ISR and SNR loss to be made. A simple approach is to apply gradient descent [Boyd and Vandenberghe, 2004]. The matched filter can be used as the initial filter, $\mathbf{w}_0$, which will minimize SNR loss but provide no control of the ISR. If $\mathbf{w}_n$ is the filter after the n^{th} iteration, then the filter after the next iteration is

$$\mathbf{w}_{n+1} = \mathbf{w}_n - \mu \left(\mathbf{\Phi}^{\text{H}} \mathbf{\Phi} \mathbf{w}_n - \mathbf{\Phi}^{\text{H}} \mathbf{d} \right), \tag{4.15}$$

where $\mu > 0$ is the step size and for $n = 0, 1, \ldots$. If the steps are sufficiently small, then the iteration should converge to the optimal, minimum ISR filter. One should expect the ISR to generally decrease as iteration continues. At the same time, the SNR loss will likely increase.

While the direct form of (4.12) provides immediate access to the minimum ISR filter, the iterative approach of (4.15) allows a tradeoff between ISR and SNR loss. By choosing an intermediate result, a filter that provides acceptable ISR can be found that has less SNR loss than the minimum ISR filter. While this approach of "early stopping" applied to an iterative algorithm is a common form of regularization [Veklerov and Llacer, 1987], it provides no guarantee that a filter with a given ISR achieves the minimum SNR loss possible. However, this can be accomplished by augmenting the cost function of (4.11) with an SNR loss constraint.

4.3.5 Minimum ISR Filter with an Explicit SNR Loss Constraint

The minimum ISR filter was found by minimizing the sum squared error defined by (4.11), which is an unconstrained optimization problem. The following adds a constraint that

allows the SNR loss to be controlled. The resulting filter will be the minimum ISR filter for a given SNR loss.

The SNR loss, which was defined in (4.9), depends on the filter response at the desired peak location, z_p, the total power of the signal, $\|\mathbf{s}\|^2$, and the total power of the weights, $\|\mathbf{w}\|^2$. The cost function of (4.11) already penalizes filters that provide a peak, z_p, that is different from $d_p = 1$, and the scaling of the signal vector, $\mathbf{s}$, is arbitrary. This leaves the norm of the weight vector, $\mathbf{w}$, which suggests the constrained optimization problem,

$$\arg\min_{\mathbf{w}} \left\{ \|\mathbf{\Phi}\mathbf{w} - \mathbf{d}\|^2 : \|\mathbf{w}\|^2 = \gamma \right\}, \tag{4.16}$$

for some scalar, γ. This is a quadratically-constrained, quadratic program, which can be solved by the method of Lagrange multipliers [Boyd and Vandenberghe, 2004].

The Lagrangian of (4.16), for a Lagrange multiplier, $\lambda \geq 0$, is

$$
\begin{aligned}
L(\mathbf{w}; \lambda) &\triangleq \|\mathbf{\Phi}\mathbf{w} - \mathbf{d}\|^2 + \lambda \left(\|\mathbf{w}\|^2 - \gamma \right) \\
&= \mathbf{w}^{\mathrm{H}} \left(\mathbf{\Phi}^{\mathrm{H}} \mathbf{\Phi} \right) \mathbf{w} - \mathbf{w}^{\mathrm{H}} \left(\mathbf{\Phi}^{\mathrm{H}} \mathbf{d} \right) - \left(\mathbf{d}^{\mathrm{H}} \mathbf{\Phi} \right) \mathbf{w} + \lambda \mathbf{w}^{\mathrm{H}} \mathbf{w} + \mathbf{d}^{\mathrm{H}} \mathbf{d} - \lambda \gamma,
\end{aligned}
\tag{4.17}
$$

and the gradient of the Lagrangian is

$$\nabla_{\mathbf{w}^*} L(\mathbf{w}; \lambda) = \left(\mathbf{\Phi}^{\mathrm{H}} \mathbf{\Phi} + \lambda \mathbf{I} \right) \mathbf{w} - \mathbf{\Phi}^{\mathrm{H}} \mathbf{d}. \tag{4.18}$$

For a given Lagrange multiplier, which corresponds to some SNR loss, the minimum ISR filter is found by solving $\nabla L(\mathbf{w}; \lambda) = \mathbf{0}$, which yields

$$\mathbf{w}(\lambda) = \left(\mathbf{\Phi}^{\mathrm{H}} \mathbf{\Phi} + \lambda \mathbf{I} \right)^{-1} \mathbf{\Phi}^{\mathrm{H}} \mathbf{d}. \tag{4.19}$$

There is no straightforward way to relate the value of the Lagrange multiplier, λ, and the resulting SNR loss, but it is clear that increasing λ leads to decreased SNR loss. If $\lambda = 0$, then no SNR loss constraint is applied and the solution reduces to the unconstrained result of (4.12) as expected. As λ increases, the result converges to the matched filter.

Equation (4.19) can be interpreted as applying diagonal loading proportional to λ to prevent the algorithm from being overly aggressive in suppressing sidelobes at the expense

of SNR. A similar result is encountered in the communications literature where a Bayesian interpretation may be applied, which leads to a minimum mean square error (MMSE) estimator that resembles (4.19).

4.4 Minimum-ISR Filtering for Multiple Waveforms

The minimum ISR filter is derived for the case where multiple waveforms are transmitted simultaneously by the radar. The methods described in Section 4.3 for designing a filter for a single waveform is extended to the case of multiple waveforms, which provides a new architecture for MIMO radar signal processing.

Consider the case of a MIMO radar that transmits M waveforms. Let

$$\{x_{m,k} : k = 0, 1, \ldots, K - 1\} \tag{4.20}$$

denote the samples of the waveform used by transmitter m. The signal observed from a single zero-delay, zero-Doppler target by receiver n is

$$s_{n,k} = b_n \sum_{m=0}^{M-1} a_m x_{m,k}, \tag{4.21}$$

where a_m and b_n are elements of the transmit steering vector, $\mathbf{a}$, and receive steering vector, $\mathbf{b}$, respectively. Note that these steering vectors depend on the angle to the target, θ_0.

Without loss of generality, the following discussion assumes that a single receive channel ($N = 1$) is used. In this case, the signal observed will be proportional to

$$s_k = \sum_{m=0}^{M-1} a_m x_{m,k}. \tag{4.22}$$

Recall that this is a linear combination of the transmitted signals, where the combination is described by the elements of the transmit steering vector; hence, this linear combination depends on the angle to the target. For example, in the case of a planar array of identical elements, targets at broadside will lead to an effective waveform that is just the sum of the MIMO waveforms.

If these data samples are grouped into a length-K vector, $\mathbf{x}$, then

$$\mathbf{s} = \mathbf{X}^{\mathrm{T}}\mathbf{a}, \tag{4.23}$$

where $\mathbf{X}$ is the $M \times K$ waveform matrix,

$$\mathbf{X} \triangleq \begin{pmatrix} x_{0,0} & \cdots & x_{0,K-1} \\ \vdots & \ddots & \vdots \\ x_{M-1,0} & \cdots & x_{M-1,K-1} \end{pmatrix}. \tag{4.24}$$

For a given target angle/transmit steering vector, a MIMO radar is equivalent to a single waveform radar that transmits this signal, $\mathbf{s}$. This allows the method for designing a minimum ISR filter for a single waveform to be used to handle the MIMO case.

4.5 Numerical Results

4.5.1 The Kasami Codes

The Kasami codes [Kasami, 1966] are a set of binary sequences with good cross-correlation functions. Their periodic correlation functions meet the Sarwate bound [Sarwate and Pursley, 1980]. Let m be a non-negative, even integer. The small set of Kasami codes consists of $2^{m/2}$ binary sequences, and the length of each sequence is $2^m - 1$. The sidelobe levels that result from minimum ISR filtering will clearly depend on the choice of waveforms. Although there may be better codes, the Kasami codes will be sufficient to demonstrate the utility of the proposed filtering approach and to compare performance with other approaches.

Because a large number of sequences with good cross-correlation properties are available for each code length, the Kasami codes are useful for spreading sequences in multi-user communications systems. MIMO radar applications will likely require only a few waveforms. For this analysis, the $M = 3$ case is the initial focus, and three arbitrarily-selected Kasami codes are used. The autocorrelation of a Kasami code and its cross-correlation with two other Kasami codes is shown in Figure 4.2. For a target at broadside, the matched filter response is the sum of the three autocorrelation sequences and the six cross-correlation sequences, which is shown in Figure 4.3.
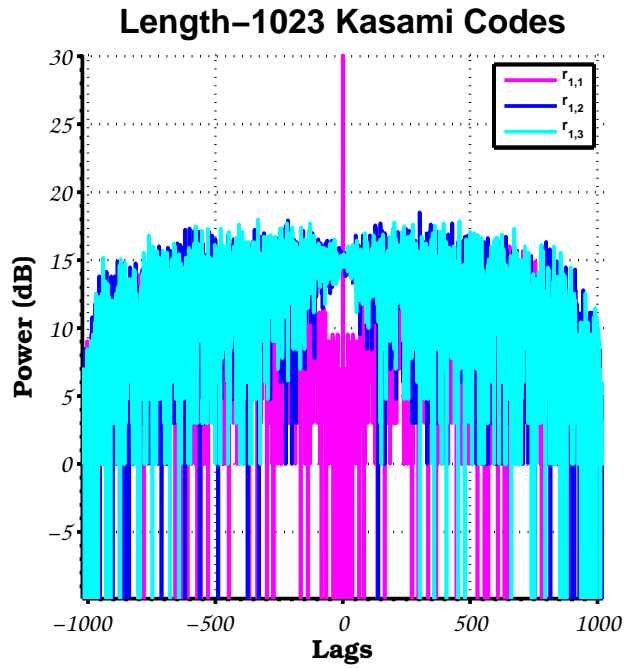
Length−1023 Kasami Codes

Figure 4.2: Autocorrelation sequence of a Kasami code and cross-correlation sequences of it with two other Kasami codes of the same length.

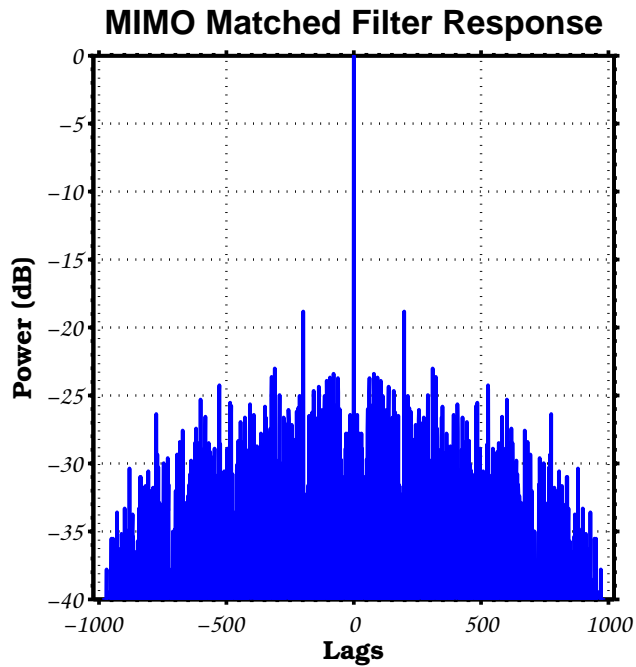

MIMO Matched Filter Response

Figure 4.3: Matched filter response of three length-1023 Kasami codes for a broadside target in the target angle/Doppler bin.

4.5.2 Minimum ISR Filtering for Multiple Kasami Codes

While the matched filter response has decent peak-sidelobe performance, the total sidelobe energy is high; the range ISR is near unity (-0.2 dB). The minimum ISR filter improves this by more than 5 dB, as shown in Figure 4.4. This can be further improved by using a longer filter. The result of designing a mismatched filter that is four times the length of the waveform ($L = 4K$) is presented in Figure 4.5.

The method of [Hua and Abeysekera, 2013] seeks to minimize the sidelobe levels over a certain range extent. In Figure 4.6, an example of minimizing the energy over the maximum number of sidelobes allowed by this method is presented. The proposed method accomplishes the same nearly-zero sidelobe performance but with less SNR loss. In fact, the proposed method allows more sidelobes to be effectively zero. This is shown in Figure 4.7, where the results are compared with the method of [Hua and Abeysekera, 2013], which fails to provide nearly-zero sidelobe levels over the region of interest.

4.5.3 Experiments with Additional Transmitted Waveforms

The results above have focused on the case of three transmitted waveforms ($M = 3$) to illustrate the ability of mismatched filters to dramatically improve the resulting ISR. In Figure 4.8, cases with more waveforms being simultaneously transmitted are considered. In addition to the Kasami codes, results for a set of randomly generated biphase codes are shown.

A number of interesting observations can be made regarding the results of Figure 4.8. First, observe that ISR performance is not significantly degraded by adding more waveforms. One exception to this, however, is going from the single-waveform case ($M = 1$) to multiple waveforms ($M > 1$) for the Kasami codes. Each Kasami code has good autocorrelation properties, so the ISR for $M = 1$ is usually better than a randomly generated biphase code. However, the integrated sidelobes increase to about the level of a randomly generated biphase code when multiple Kasami codes are linearly combined when $M > 1$. Even though the Kasami codes have relatively low peak cross-correlation levels, they still

Figure 4.4: Mismatched filter response of three Kasami codes for a broadside target in the target angle/Doppler bin with $L = K$.
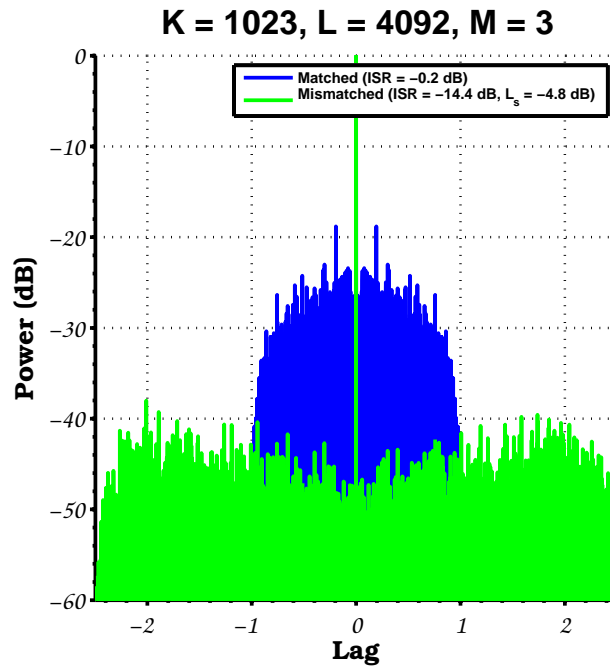

Figure 4.5: Mismatched filter response of three Kasami codes for a broadside target in the target angle/Doppler bin with $L = 4K$.

Figure 4.6: Mismatched filter response where the sidelobes are optimized over $K_0 = 170$ of the 1022 range bin on each side of the peak.


Figure 4.7: Mismatched filter response where sidelobes are optimized over $K_0 = 500$ of the 1022 range bin on each side of the peak.
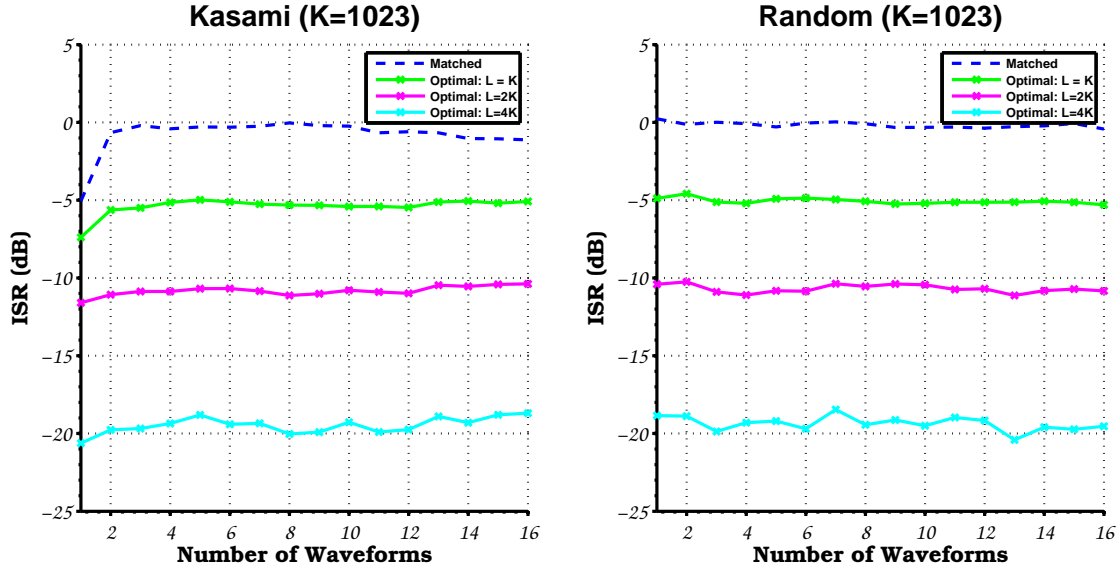
Figure 4.8: ISR as a function of number of transmitted waveforms. Results for Kasami codes (left) and random biphase codes (right) are presented.

provide poor ISR when the matched filter is used.

When considering Figure 4.8, one may be surprised to see that the ISR can apparently be improved by transmitting additional waveforms. Recall that the these results consider a single linear combination of the waveforms. While the ISR may be very slightly improved (fractions of a dB) by transmitting an additional waveform, at some other angle, the ISR will be degraded.

4.5.4 Tradeoffs Between ISR and SNR Loss

The proposed method for designing mismatched filters is based on solving the optimization problem that minimizes the squared error between the filter response and an idealized desired response. A closed-form solution was presented in (4.12), which involves the pseudoinverse. Instead of finding the optimal weights directly, the iterative approach of (4.15) can be employed, which is based on the method of gradient descent. This will provide a sequence of filters that provide increasingly better ISR. If the matched filter is used as a starting point to initialize the algorithm, any refinements to the filter will result in SNR loss. Results using the gradient descent algorithm are shown in Figures 4.9 and 4.10.

K = 1023, L = 1023, K_0 = 1022

Figure 4.9: Results of iterative improvement of filter for $M = 3$ with gradient descent.
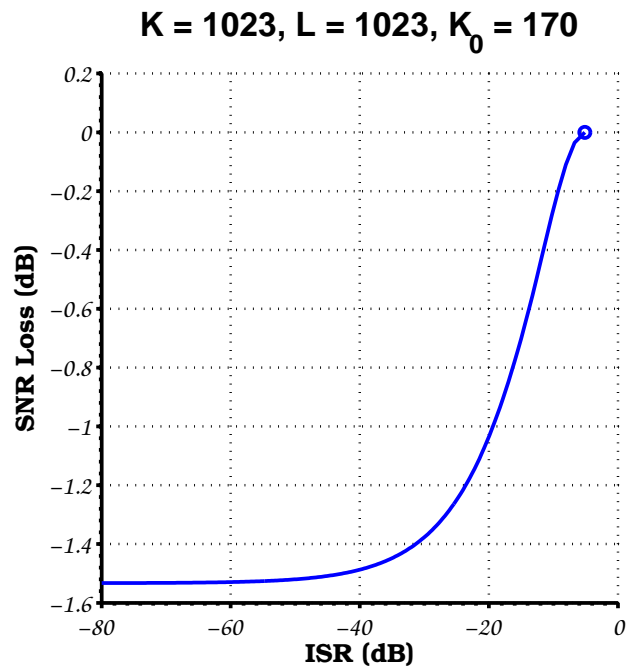

K = 1023, L = 1023, K_0 = 170

Figure 4.10: Results of iterative improvement of filter for $M = 3$ that minimizes $K_0 = 170$ of the 1022 sidelobes on either side of the peak with gradient descent.

The iterative approach is especially attractive for the case where nearly-zero sidelobes are required over a limited region ($K_0 < K$). This is clear from Figure 4.10, where arbitrarily low sidelobes are provided over the region of interest but with much less SNR loss compared to the optimal filter. This is because this case results in a poorly conditioned problem, which is amenable to an iterative solution. The iterative approach can provide low ISR of -80 dB with an SNR loss of only about -1.5 dB. The direct approach, shown in Figure 4.6, provides an extremely low ISR of -286 dB, but this comes at the cost of significant SNR loss, $L_s = -17$ dB.

When sidelobes are optimized only over a limited number of range bins, the iterative approach can be used to provide acceptably low ISR with much less SNR loss than the direct approach. In most radar applications, the system designer can decide at what level ISR is no longer the limiting factor in radar performance, e.g., when noise from range sidelobes is pushed below the noise floor set by some other source. Consequently, a principled decision on when to stop the iteration can be made based on a well-defined maximum acceptable ISR.

The performance of the gradient-descent based approach with the constrained-optimization approach are presented in Figure 4.11 and Figure 4.12. As is expected, the constrained-optimization approach provides improved SNR loss compared to the gradient-descent approach, but the improvement seems to be limited.

4.5.5 Robustness to Angle and Doppler Offsets

Mismatched filters are so successful in improving sidelobe performance because they are designed specifically for the waveform, and their performance degrades when used with a different waveform. The classic example of this is when an uncompensated Doppler shift is applied to the waveform. While the mismatched filter response to the waveform may be quite good, the sidelobe levels and SNR loss may be significantly degraded when applied to the Doppler-shifted waveform. Results showing the degraded SNR and increase in ISR with Doppler shift are shown in Figure 4.13. Note that SNR loss is computed relative to
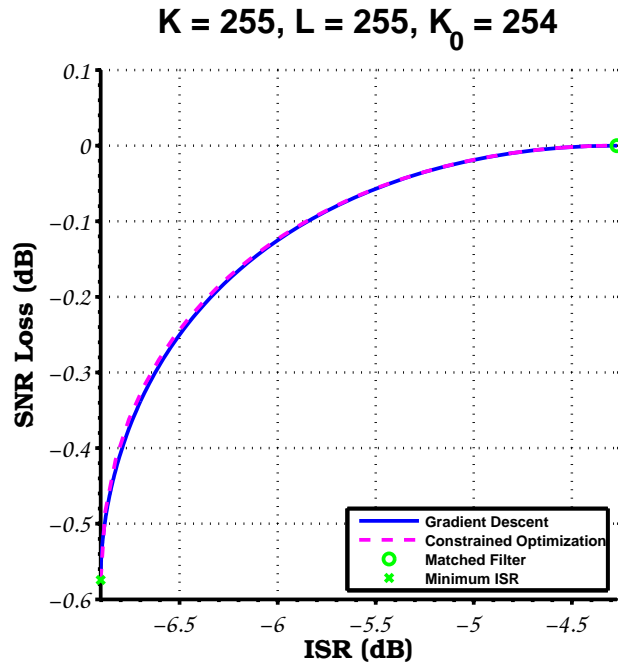
Figure 4.11: Results of iterative improvement of filter for $M = 1$ with gradient descent and constrained optimization approaches.
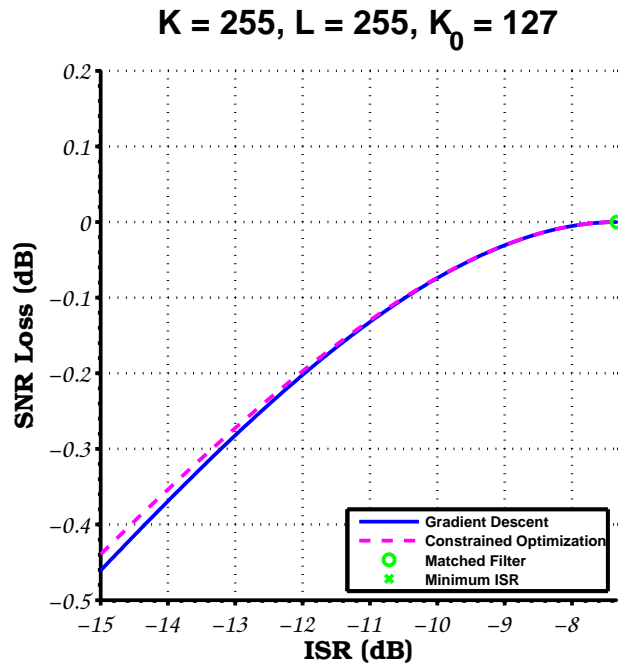

Figure 4.12: Results of iterative improvement of filter for $M = 3$ that minimizes $K_0 = 127$ of the 255 sidelobes on either side of the peak.
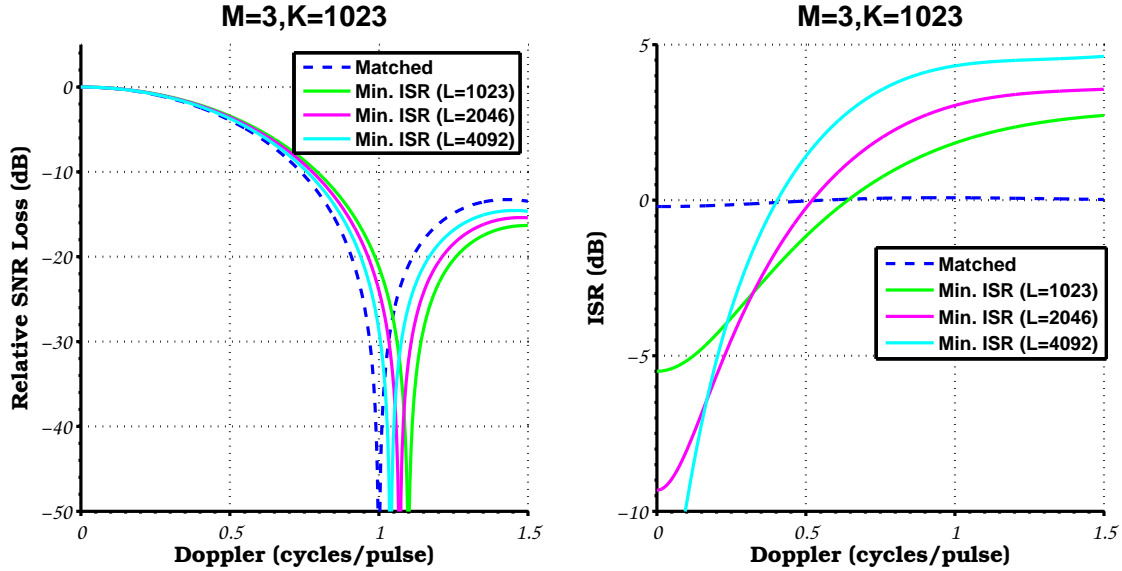
Figure 4.13: SNR loss and ISR as a function of Doppler offset.

the case of no Doppler offset for each filter. Also, the ISR is computed relative to the peak of the zero-Doppler case because the concern is that the range sidelobes of Doppler-shifted returns will interfere with returns from targets at the expected Doppler.

The filters could be made more robust to Doppler by optimizing the response over some Doppler extent. For applications like SAR, however, this is probably not necessary. The Doppler shift imposed by the motion of the platform can be removed, which may be necessary for highly-squinted collections. Targets on the ground are likely not moving fast enough towards or away from the radar to impose a significant Doppler shift for a reasonable SAR pulse width.

Another concern is that the sidelobes may be degraded for off-angle targets. The proposed approach is to design optimal filters that assume a particular target angle. As has been shown, these filters provide improved performance when the target is in the center of the beam that is formed. However, degraded performance is expected for targets at other angles. An example is shown in Figure 4.14. Again, ISR is presented relative to the case of no angle offset. Longer mismatched filters, which can provide improved ISR performance, are more sensitive to angle offsets.

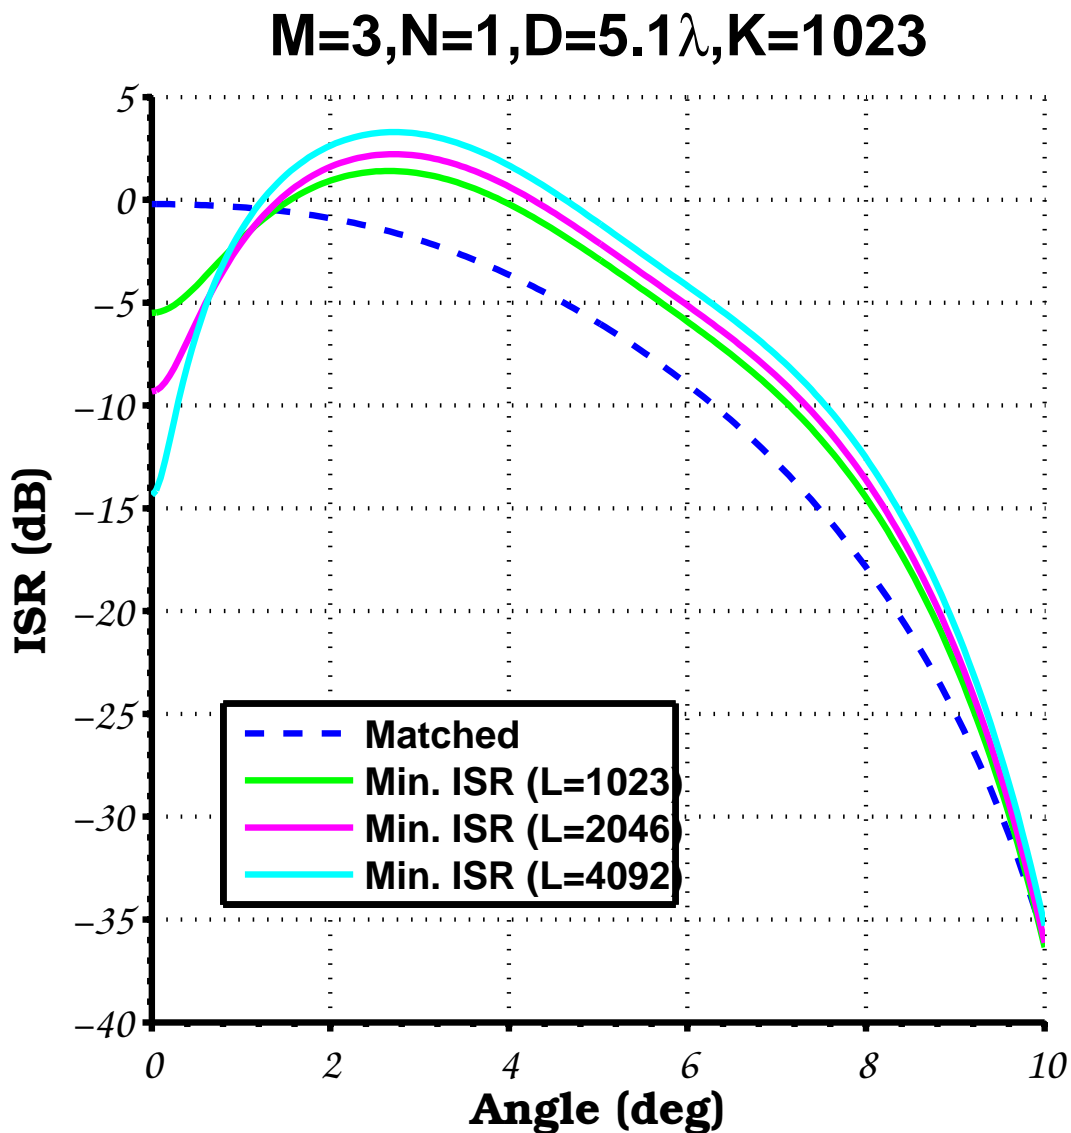Figure 4.14: ISR as a function of angle offset. The mismatched filters are designed for a target at $\theta = 0$. ISR is computed relative to the peak of the response of the target at $\theta = 0$. The transmit and receive arrays are uniform linear arrays composed of subarrays with a 10° beamwidth. The matched filter ISR trend matches the two-way subarray pattern. The transmit aperture is three times as long as the receive aperture.

While targets in the center of the beam will have low sidelobes, targets that are offset in angle will have elevated sidelobes when the proposed mismatched filtering approach is used. The proposed approach optimizes the response for a single target angle. This could be made more robust to angle offset by optimizing the response over a range of target angles. Indeed, this is the approach used by other MIMO mismatched filtering approaches, which optimize over all angles. An avenue for future work is to explore optimizing over a range of angles to provide good ISR over the entire beam.

Fortunately, radar systems have other methods for suppressing off-angle targets. Unlike a phased array, a MIMO radar transmitting nominally orthogonal waveforms cannot exploit array gain on transmit, but both systems can take advantage of array gain on receive. The example of Figure 4.14 assumed a relatively small receive aperture. Because of this, off-angle targets are not suppressed, and their range sidelobes will be high. Figure 4.15 presents cases where the receive aperture is the same size as the transmit aperture and where a larger receive aperture is used. The receive beampattern effectively suppresses the targets sufficiently so that their sidelobes do not interfere with the desired target.

4.5.6 Cumulative ISR

Traditionally, the ISR has been the primary metric to capture the contribution of range sidelobes to the multiplicative noise in a SAR image. The ISR is the sum of the contributions of all of the sidelobes in the impulse response. As can be seen from the example of an NRA in Figure 3.2, the ISR only captures the impact at the edge of the NRA closest to a high-return area. The impact of sidelobes decreases further into the NRA. This can be quantified by adding the contribution of the appropriate sidelobes. An example is shown in Figure 4.16. The cumulative ISR is shown as a function of range into the NRA. Note that this assumes that this only includes the sidelobes from one side of the NRA. The matched filter responses are the same as those highlighted by the superimposed lines in Figure 3.2.

Figure 4.15: ISR as a function of angle offset with larger receive apertures. In the left figure, the transmit and receive apertures are the same size, and in the right figure, the receive aperture is three times the size as the transmit aperture. The receive aperture is able to suppress off-angle targets thereby mitigating their poor sidelobe performance.
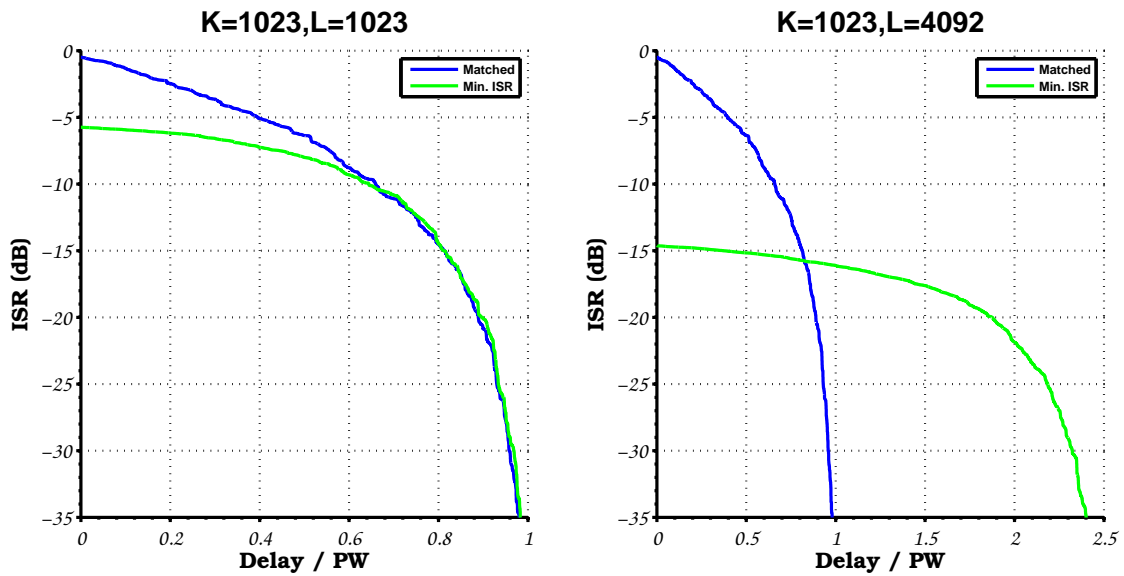
Figure 4.16: The cumulative ISR for the matched filter and two mismatched filters using three Kasami codes ($M = 3$).

4.6 Summary

A new method for designing filters that control range sidelobes when multiple waveforms are simultaneously transmitted, as in MIMO radar, has been presented. These filters minimize the ISR, which is critical for applications such as MIMO SAR. This technique allows arbitrarily long filters to be applied to aggressively control the sidelobes. It can also be modified to control sidelobes over only a specified region in range. The proposed method outperforms other techniques presented in the literature by jointly applying range compression and transmit beamforming rather than applying these operations in a separable manner.

An iterative approach to filter design was also presented, which suggests the potential for trading off ISR and SNR loss. It also allows numerical instabilities to be avoided when designing a filter that optimizes sidelobes over a limited range extent. This simple, gradient-descent based iterative approach provides similar performance to a more complicated constrained-optimization approach that was also introduced.

The results described here focused on a particular angle/Doppler bin, i.e., the filter is optimized for the target angle and Doppler frequency. Like all filter optimization approaches in radar, performance will degrade if a significant, uncompensated Doppler offset is present. While the antenna will provide some suppression of off-angle targets, it is expected that sidelobe performance may degrade for sidelobe targets. The proposed algorithm can be extended to optimize over angle and/or Doppler in a natural way analogous to the method for optimizing over a set of range bins. This is a topic for future work.

CHAPTER 5

RECURRENT AND MULTICHANNEL SAMPLING

5.1 Overview

Numerous versions of the sampling theorem have been proposed over the years that state conditions under which a signal can be perfectly reconstructed from a sampled version of itself. This fact, along with revolutionary advances in computing, enabled the current ubiquity of digital signal processing. These sampling theorems provide assurance that if a signal is sampled at a sufficiently high rate relative to the bandwidth of the signal, then the signal can be perfectly reconstructed from these samples. However, no practical signal is perfectly bandlimited, and degradation from aliasing is inevitable. This chapter presents a method for quantifying this impact in a number of interesting cases. These include the general case of multichannel sampling and a special case of nonuniform sampling called recurrent sampling. Sampling of deterministic signals as well as random processes are considered.

The sampling theorem discussed by Whittaker [Whittaker, 1915], Nyquist [Nyquist, 1928], and Shannon [Shannon, 1949] asserts that a bandlimited signal can be perfectly reconstructed from a set of uniformly-spaced samples if the sampling rate is at least twice the highest frequency of the signal. Higgins provides conditions on the sampling instants under which a signal can be recovered from nonuniform samples [Higgins, 1976]. Lloyd extended the sampling theorem to the case of wide-sense stationary (WSS) random process [Lloyd, 1959] where the spectral content of the random process is described by its power spectral density (PSD). Lee develops a sampling theorem for nonstationary processes [Lee, 1978], which is based on the definition of a bandlimited nonstationary random process of Zakai [Zakai, 1965]. Extensions of the sampling theorem to the case of nonuniform sampling have also been provided [Yen, 1956].

These versions of the sampling theorem are all based on the assumption of a perfectly

bandlimited signal. This chapter analyzes the consequences of applying the reconstruction formulas derived for bandlimited signals to signals that are not bandlimited. The case of uniformly sampling a deterministic signal is understood by the familiar concept of aliasing, where a sinusoid at a frequency higher than the sampling rate becomes indistinguishable from its aliased frequency. A similar result is presented for the case of recurrent sampling where the aliases are weighted depending on the spacing of the samples within a recurrence.

5.2   Aliasing in Uniform and Recurrent Sampling

Before turning to the case of a random process, sampling of a deterministic signal is considered. The notion of recurrent sampling is introduced, and it is compared to the case of uniform sampling.

5.2.1   The Spectrum of a Sampled Deterministic Signal

Suppose that the signal, $x(t)$, is sampled at a set of instants, $\{T_n\}$, to generate a discrete-time sequence. Practically, sampling converts the continuous-time signal into a sequence of samples. Conceptually, however, the process of sampling is equivalent to multiplying the continuous-time signal, $x(t)$, by a sampling signal, $s(t)$, which is a collection of Dirac delta functions located at the sampling instants,

$$s(t) \triangleq \sum_{n=-\infty}^{\infty} \delta(t - T_n). \tag{5.1}$$

The resulting signal after sampling is

$$\tilde{x}(t) \triangleq x(t) s(t) = \sum_{n=-\infty}^{\infty} x(T_n) \delta(t - T_n). \tag{5.2}$$

Because the sampled signal is the product of the signal, $x(t)$, and the sampling signal, $s(t)$, its spectrum is the convolution of their Fourier transforms,

$$\tilde{X}(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(\xi) X(\omega - \xi) \, d\xi, \tag{5.3}$$

where $X(\omega)$ and $S(\omega)$ are the spectra of the signal and the sampling signal respectively. The convolution in (5.3) provides the interpretation that the action of sampling a signal

alters the signal's spectrum, $X(\omega)$, by passing it through a system with a frequency response corresponding to the spectrum of the sampling signal, $S(\omega)$. An alternate interpretation is that the spectrum of the sampled signal is the nonuniform discrete-time Fourier transform (DTFT) of the sampled signal,

$$\tilde{X}(\omega) = \sum_{n=-\infty}^{\infty} x(T_n) e^{-i\omega T_n}. \tag{5.4}$$

5.2.2 Uniform and Recurrent Sampling Schemes

In uniform sampling, a signal is sampled at a regular interval, T. Equivalently, the signal is sampled at a constant rate, $1/T$. Recurrent sampling is a generalization of uniform sampling where sets of N samples are collected at a uniform rate. While these N samples need not be uniformly spaced, each of these samples recurs at the same uniform rate, $1/T$. The sampling instants for uniform and recurrent sampling can be written as

$$\textbf{Uniform: } T_m \triangleq mT \text{ for } m = 0, \pm 1, \pm 2, \ldots \tag{5.5}$$

$$\textbf{Recurrent: } T_{nm} \triangleq mT + T_n \text{ for } m = 0, \pm 1, \pm 2, \ldots \text{ and } n = 1, \ldots, N, \tag{5.6}$$

where $T_1, \ldots, T_N$ are the samples collected during the recurrence corresponding to $m = 0$. Uniform and recurrent sampling are illustrated in Figure 5.1.
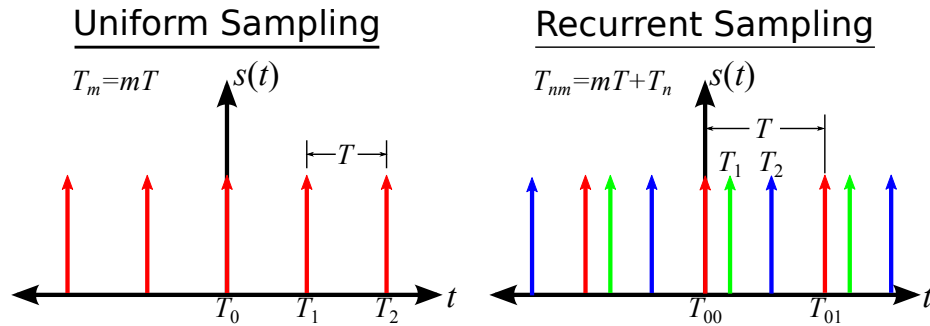
Figure 5.1: Sampling functions, $s(t)$, for uniform and recurrent sampling. In uniform sampling, all samples are spaced by an interval, T. In recurrent sampling, an additional $M - 1$ samples are collected for each uniform sample.

If the samples within one recurrence are uniformly spaced, i.e., $T_n = (n-1)T_0$ for some interval, T_0, then there is a particular recurrence interval, T, where recurrent sampling reduces to uniform sampling. This is the case when $T_0 = T/N$, which is equivalent

to uniformly sampling the signal at a rate, N/T. Recurrent sampling trivially reduces to uniform sampling for $N = 1$.

If the samples within a recurrence are not uniformly spaced, or if they are uniformly spaced but with $T_0 \neq T/N$, then the recurrent sampling will be nonuniform. To quantify this deviation from uniform sampling, define the average spacing of the samples within a recurrence as

$$\bar{T}_0 \triangleq \frac{1}{N-1} \sum_{n=1}^{N-1} (T_{n+1} - T_n). \tag{5.7}$$

The uniformity, κ, of a recurrent sampling scheme is found by comparing the average spacing, $\bar{T}_0$, to the spacing required for uniform sampling, T/N, which yields

$$\kappa \triangleq \frac{\bar{T}_0}{T/N}. \tag{5.8}$$

For the trivial case of $N = 1$, the average spacing, $\bar{T}_0$, is not well defined, but such sampling is always uniform so that $\kappa = 1$ by definition for this case. Examples of recurrent sampling with varying degrees of uniformity are presented in Figure 5.2.

Figure 5.2: Varying uniformity, κ, in recurrent sampling for $N = 3$. Note that $\kappa = 1.5$ is a degenerate case where two samples coincide for each recurrence.

5.2.3 Uniform and Recurrent Sampling of a Deterministic Signal

The uniform sampling signal is the sum of evenly spaced Dirac delta functions, which is called an impulse train or a Dirac comb. The spectrum of the uniform sampling signal is also an impulse train. In Appendix B, it is shown that the spectrum of the recurrent sampling signal is also a set of impulses but with a weighting that depends on the spacing

	Uniform	Recurrent
Sampling Signal, $s(t)$	$T \sum\limits_{m=-\infty}^{\infty} \delta(t - mT)$	$\frac{T}{N} \sum\limits_{n=1}^{N} \sum\limits_{m=-\infty}^{\infty} \delta(t - mT - T_n)$
Sampling Signal Spectrum, $S(\omega)$	$2\pi \sum\limits_{m=-\infty}^{\infty} \delta\left(\omega + \frac{2\pi}{T}m\right)$	$2\pi \sum\limits_{m=-\infty}^{\infty} \alpha_m \delta\left(\omega + \frac{2\pi}{T}m\right)$
Spectrum of Sampled Signal, $\tilde{X}(\omega)$	$\sum\limits_{m=-\infty}^{\infty} X\left(\omega + \frac{2\pi}{T}m\right)$	$\sum\limits_{m=-\infty}^{\infty} \alpha_m X\left(\omega + \frac{2\pi}{T}m\right)$

Table 5.1: Sampling signals, spectra of sampling signals, and spectra of sampled signals for uniform and recurrent sampling. The weights, α_m, for recurrent sampling are given in (5.9)

of the samples within a recurrence,

$$\alpha_m \triangleq \frac{1}{N} \sum_{n=1}^{N} \exp\left\{i\frac{2\pi m}{T}T_n\right\}.$$ (5.9)

The sampling signals and their corresponding spectra are presented in Table 5.1.

5.3 Recurrent Sampling of a WSS Random Process

A zero-mean WSS random process is characterized by its power spectral density (PSD). In this section, the PSD of a sampled random process is related to the PSD of the continuous-time random process. The definition of the PSD is briefly reviewed before proceeding to the case of sampling.

5.3.1 The PSD of a WSS Random Process

Let $x(t)$ be a zero-mean, wide-sense stationary (WSS) random process. In general, the Fourier transform of a realization of a random process does not converge, but this problem can be remedied by only considering the signal over an interval of a finite duration, D. Define this windowed signal by

$$x_D(t) \triangleq x(t)\,\text{rect}\left(\frac{t}{D}\right),$$ (5.10)

where $\text{rect}(t)$ is the rectangular function that is zero for $t \notin (-1/2, 1/2)$, and let $X_D(\omega)$ be the Fourier transform of the windowed signal. PSD of the random process is the average

power of this spectrum as D is allowed to grow arbitrarily large,

$$P_x(\omega) \triangleq \lim_{D \to \infty} \frac{1}{D} \mathrm{E}\left[|X_D(\omega)|^2\right]. \tag{5.11}$$

The Wiener-Khinchin theorem asserts that the PSD of a WSS random process is the Fourier transform of its autocorrelation function, $R_x(\tau)$, which is defined as

$$R_x(\tau) \triangleq \mathrm{E}\left[x(t)\, x^*(t - \tau)\right]. \tag{5.12}$$

Because the random process is WSS, this is independent of t.

5.3.2 The PSD of a Recurrently Sampled WSS Random Process

Suppose that the WSS random process, $x(t)$, is sampled at instants, $\{T_n\}$. The resulting discrete-time random process, $\tilde{x}[n]$ is not generally WSS, so its PSD is not the discrete-time Fourier transform (DTFT) of its autocorrelation sequence. However, the nonuniform DTFT of (5.4) can still be computed. In the following, the formulation of the spectrum of a recurrently sampled signal in Table 5.1 is used with the definition of the PSD given by (5.11) to find the PSD of a recurrently sampled WSS random process.

The spectrum of the truncated random process after recurrent sampling is

$$\tilde{X}_D(\omega) \triangleq \int_{-D/2}^{D/2} \tilde{x}(t)\, e^{-i\omega t}\, dt = \sum_{-\infty}^{\infty} \alpha_m X_D\left(\omega + \frac{2\pi}{T}m\right), \tag{5.13}$$

where $X_D(\omega)$ is the spectrum of the truncated random process, $x_D(t)$, defined in (5.10). In general, a realization of the random process, $x(t)$, will not be absolutely integrable, but it is assumed that $|\tilde{x}(t)|^2 < \infty$ for all t so that the series converges because the integral is computed over a interval of finite length, D.

The power spectrum corresponding to (5.13) is

$$\left|\tilde{X}_D(\omega)\right|^2 = \sum_{m=-\infty}^{\infty} \sum_{m'=-\infty}^{\infty} \alpha_m \alpha_{m'}^* X_D\left(\omega + \frac{2\pi}{T}m\right) X_D^*\left(\omega + \frac{2\pi}{T}m'\right), \tag{5.14}$$

and the average periodogram is

$$\frac{1}{D}\mathrm{E}\left[\left|\tilde{X}_D(\omega)\right|^2\right] = \sum_{m=-\infty}^{\infty} \sum_{m'=-\infty}^{\infty} \alpha_m \alpha_{m'}^* \frac{1}{D}\mathrm{E}\left[X_D\left(\omega + \frac{2\pi}{T}m\right) X_D^*\left(\omega + \frac{2\pi}{T}m'\right)\right]. \tag{5.15}$$

To determine the PSD, the limit as the length of the interval, D, increases arbitrarily in (5.15),

$$
\begin{aligned}
P_{\tilde{x}}(\omega) &= \lim_{D\to\infty} \frac{1}{D} \sum_{m=-\infty}^{\infty} \sum_{m'=-\infty}^{\infty} \alpha_m \alpha_{m'}^* \mathrm{E}\left[X_D\left(\omega + \frac{2\pi}{T}m\right) X_D^*\left(\omega + \frac{2\pi}{T}m'\right)\right] \\
&= \sum_{m=-\infty}^{\infty} \sum_{m'=-\infty}^{\infty} \alpha_m \alpha_{m'}^* \lim_{D\to\infty}\left[\frac{1}{D}\mathrm{E}\left[X_D\left(\omega + \frac{2\pi}{T}m\right) X_D^*\left(\omega + \frac{2\pi}{T}m'\right)\right]\right].
\end{aligned}
\tag{5.16}
$$

Interchanging the order of the infinite summations and the limiting operation is justified by appealing to Lebesgue's dominated convergence theorem. Because $\alpha_m \alpha_{m'}^* \leq |\alpha_m| \leq 1$, a suitable dominating function, is $g_m = P_x(\omega + 2\pi m/T)$. In (A.15), it was shown that the summands with $m \neq m'$ are equal to zero. The PSD of a recurrently sampled WSS random process is

$$
P_{\tilde{x}}(\omega) = \sum_{m=-\infty}^{\infty} |\alpha_m|^2 \, P_x\left(\omega + \frac{2\pi}{T}m\right),
\tag{5.17}
$$

where the weights, α_m, are as given in (5.9). Compare the PSD of a recurrently sampled random process, which is given by (5.17) to the spectrum of a recurrently sampled deterministic signal in Table 5.1. In both cases, the aliased are weighted according to the weights, α_m. Because the PSD is related to the power spectrum, the aliases are weighted by the power terms, $|\alpha_m|^2$, rather than the generally complex-valued weights, α_m.

5.4 Multichannel Sampling and Reconstruction

The minimum sampling rate that permits reconstruction of a bandlimited signal is called the Nyquist rate. A similar result exists for a multichannel system where the signal is passed through N linear systems and the outputs are uniformly sampled at a rate that is slower than the Nyquist rate by a factor of N [Papoulis, 1977a]. This generalized sampling theorem asserts that a bandlimited signal can be perfectly reconstructed from its samples as long as the average sampling rate is sufficiently high, which suggests that reconstruction error is independent of the structure of the sampling. In the following, it will be shown that this is only a valid conclusion if the signal is perfectly bandlimited.

A multichannel sampling and reconstruction scheme is illustrated in Figure 5.3. In

[Brown, 1981], a general method for determining the frequency responses of the recon-struction filters, $G_n(\omega)$, in terms of the channel frequency responses, $F_n(\omega)$, is derived that guarantees perfect reconstruction of a suitably bandlimited signal.

Recurrent sampling can be considered as a special case of this multichannel sampling structure where $F_n(\omega) = e^{-i\omega T_n}$. The general method of [Brown, 1981] provides a nu-merical method for determining the frequency responses of the reconstruction filters for a general multichannel sampling system, i.e., for arbitrary $F_n(\omega)$ in Figure 5.3. In [Yen, 1956], a general closed-form interpolation formula is provided for reconstruction from re-current samples. This method was extended in [Eldar and Oppenheim, 2000] to implement this reconstruction as a filterbank using the same structure as [Brown, 1981].

In this section, a model is first derived for the reconstructed signal given the frequency response of the N reconstruction filters. Then, the conditions that the frequency responses must satisfy to provide perfect reconstruction of an appropriately band-limited signal are derived. This follows the development of [Brown, 1981]. After reviewing multichannel sampling and reconstruction, a method for calculating the impact of aliasing is presented. The closed-form expressions for reconstruction from recurrent samples of [Yen, 1956] and [Eldar and Oppenheim, 2000] are also discussed and placed in the general framework of [Brown, 1981].

5.4.1 Signal Model

In multichannel sampling, the continuous-time signal, $x(t)$, is not directly sampled, but instead the responses of N linear time-invariant systems to the signal are sampled. To reconstruct $x(t)$ from samples of the system outputs, $y_n(t)$, a (digital) reconstruction filter is applied to each channel output, and the results are summed. The process of multichannel sampling and reconstruction is presented in Figure 5.3. In the following, a model for the spectrum of the reconstructed signal is derived.

If $f_n(t)$ is the impulse response of a system for $n = 1, 2, \ldots, N$, then the output of each
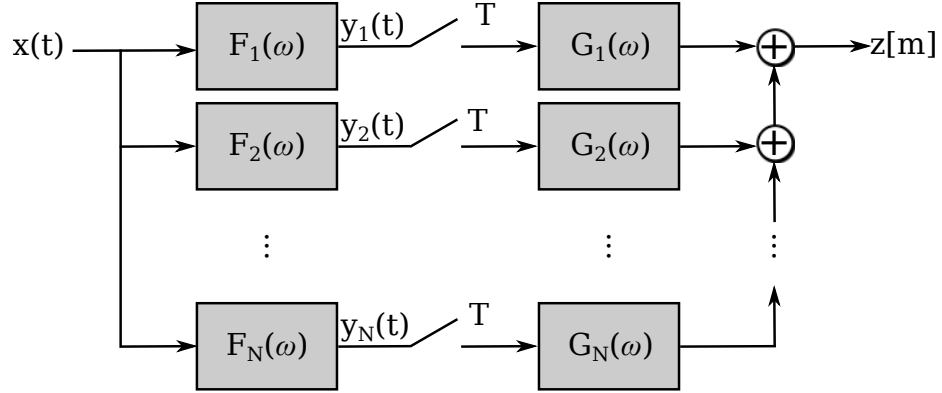
Figure 5.3: Multichannel sampling and reconstruction. The outputs of N linear time-invariant systems, $\{F_n(\omega)\}$, are sampled a rate, $1/T$. The samples from channel n are processed by reconstruction filters, $\{G_n(\omega)\}$. The results are summed to provide the reconstructed signal.

system, $y_n(t)$, is the convolution of the input signal and the system impulse response,

$$y_n(t) \triangleq (f_n * x)(t) = \int_{-\infty}^{\infty} f_n(\tau)\, x(t-\tau)\, d\tau. \tag{5.18}$$

The output of each system is sampled simultaneously at the same rate, $1/T$. The collected samples are

$$y_n[m] \triangleq y_n(mT) \tag{5.19}$$

for $n = 1, \ldots, N$ and $m = 0, \pm 1, \pm 2, \ldots$ where T is the sampling interval.

To reconstruct the input signal, $x(t)$, from the samples of (5.19), the output of each system will be processed by a reconstruction filter and the results summed. The sampled sequence, $y_n[m]$, will be upsampled by a factor of N before being digitally processed by a reconstruction filter with a frequency response, $G_n(\omega)$. These reconstruction filters will be implemented digitally at a sampling rate, N/T. Because of this, the frequency response of the reconstruction filters will be $2\pi N/T$-periodic.

The output of each reconstruction filter will be the discrete convolution of an upsampled version of the sampled signal and the samples of the impulse response, $g_n(t)$. These will be summed to provide the reconstructed signal,

$$z[m] = \sum_{n=1}^{N} \sum_{k=-\infty}^{\infty} y_n(kT)\, g_n\!\left((m-kN)\frac{T}{N}\right). \tag{5.20}$$

89

The spectrum of the reconstructed signal is

$$Z(\omega) \triangleq \sum_{m=-\infty}^{\infty} z[m] e^{-i\omega mT/N} = \sum_{n=1}^{N} G_n(\omega) \sum_{k=-\infty}^{\infty} y_n(kT) e^{-i\omega Tk}, \tag{5.21}$$

where $G_n(\omega)$ is the frequency response of the reconstruction filter applied to the output of system n,

$$G_n(\omega) \triangleq \sum_{m=-\infty}^{\infty} g[m] e^{-i\omega mT/N}, \tag{5.22}$$

which is periodic with period $2\pi N/T$. The spectrum of the reconstructed signal, $Z(\omega)$, is also $2\pi N/T$-periodic.

By Poisson's summation formula, the discrete Fourier transform of a uniformly sampled signal is equivalent to an aliased version of the Fourier transform of the original signal,

$$\sum_{m=-\infty}^{\infty} x(mT) e^{-imT\omega} = \frac{1}{T} \sum_{m=-\infty}^{\infty} X(\omega + 2\pi m/T), \tag{5.23}$$

which allows the spectrum of (5.21) to be rewritten as

$$Z(\omega) = \frac{1}{T} \sum_{n=1}^{N} G_n(\omega) \sum_{k=-\infty}^{\infty} Y_n\left(\omega + \frac{2\pi}{T}k\right), \tag{5.24}$$

where $Y_n(\omega)$ is the spectrum of each channel. These spectra can be written as the product of the spectrum of the input signal, $X(\omega)$, and the frequency response of each channel, $F_n(\omega)$,

$$Y_n(\omega) = X(\omega) F_n(\omega). \tag{5.25}$$

Substituting (5.25) into (5.24) yields

$$Z(\omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} \left(\sum_{n=1}^{N} G_n(\omega) F_n\left(\omega + \frac{2\pi}{T}k\right) \right) X\left(\omega + \frac{2\pi}{T}k\right). \tag{5.26}$$

The form of the reconstructed signal spectrum in (5.26) motivates the definition of

$$H_k(\omega) \triangleq \sum_{n=1}^{N} G_n(\omega) F_n\left(\omega + \frac{2\pi}{T}k\right), \tag{5.27}$$

so that (5.26) becomes

$$Z(\omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} H_k(\omega) X\left(\omega + \frac{2\pi}{T}k\right). \tag{5.28}$$

The frequency response, $H_k(\omega)$ is applied to a frequency that is offset from the reconstructed frequency, ω, by $(2\pi/T)k$.

In the above derivation, the order of summation of two infinite series was interchanged. This is permitted because under the condition that all of the series are absolutely convergent for reasonable signals of interest. Even if the input signal, $x(t)$, is not absolutely integrable, it is assumed that the impulse response of each channel is suitably well-behaved so that each $y_n[m]$ is absolutely summable. This guarantees the absolute convergence of the series in (5.20), which ensures the absolute convergence of the series in (5.21).

5.4.2 Multichannel Reconstruction Filter Derivation

Above, the spectrum of the reconstructed signal, given by (5.28), was derived in terms of the overall frequency responses, $H_k(\omega)$, which are given by (5.27). From this, observe that $Z(\omega) = X(\omega)$ if $H_0(\omega) = T$ and $H_k(\omega) = 0$ for $k \neq 0$. To eliminate reconstruction error, the N reconstruction filters, $G_n(\omega)$, should be designed to satisfy this condition. In general, this is impossible if the signal spectrum, $X(\omega)$, is nonzero over unbounded support, which would require infinitely-many constraints to be satisfied with finitely-many degrees of freedom. However, if $X(\omega)$ is appropriately bandlimited, then these N reconstruction filters will be sufficient.

The output of each system is sampled at a rate of $1/T$ samples per second. Because there are N systems, the effective sampling rate is N/T samples per second. Consequently, a bandwidth of

$$\beta \triangleq N/T \tag{5.29}$$

can be unambiguously reconstructed from the samples. Suppose that the bandwidth of $x(t)$ is $2\pi\beta$ so that $X(\omega) = 0$ for $\omega \notin \Omega$ with

$$\Omega \triangleq \{\omega : -\pi\beta \leq \omega < \pi\beta\}. \tag{5.30}$$

Even though the channel responses, $F_n(\omega)$, may be nonzero over any range of frequencies, ω, only frequencies $\omega \in \Omega$ are important because the signal passing through the

channels is bandlimited. Without loss of generality, assume that $F_n(\omega) = 0$ for $\omega \notin \Omega$.

Define the $2\pi\beta$-periodic extension of $F_n(\omega)$, which is denoted $\bar{F}_n(\omega)$, i.e., $\bar{F}_n(\omega) = F_n(\omega)$ for $0 \leq \omega < 2\pi\beta$, and $\bar{F}_n(\omega)$ is $2\pi\beta$-periodic. Similarly, denote the $2\pi\beta$-periodic extension of the signal spectrum, $X(\omega)$, by $\bar{X}(\omega)$. Each reconstruction filter is a digital filter that is implemented at a rate of β, so its frequency response, $G_n(\omega)$, is naturally $2\pi\beta$-periodic.

Write the periodic extension of the reconstructed spectrum for $0 \leq \omega < 2\pi\beta$ in (5.28) as

$$\bar{Z}(\omega) = \frac{1}{T}\bar{H}_0\bar{X}(\omega) + \frac{1}{T}\sum_{k=1}^{N-1}\bar{H}_k(\omega)\bar{X}\left(\omega + \frac{2\pi\beta}{N}m\right), \tag{5.31}$$

where $\bar{H}_k$ is the periodic extension of the frequency response in (5.27), which can be written as

$$\bar{H}_k(\omega) = \sum_{n=1}^{N} G_n(\omega)\bar{F}_n\left(\omega + \frac{2\pi\beta}{N}k\right). \tag{5.32}$$

There is no reconstruction error if $\bar{Z}(\omega) = Z(\omega)$ for all ω. Because $\bar{Z}(\omega)$ is $2\pi\beta$-periodic, this equivalence must only be established over an interval of length $2\pi\beta$. In fact, any interval of that length will be sufficient.

From (5.31), observe that the reconstructed spectrum satisfies this requirement if

$$\bar{H}_k(\omega) = \begin{cases} T, & \text{for } k = 0 \\ 0, & \text{for } k = 1, 2, \ldots, N-1 \end{cases} \tag{5.33}$$

for $0 \leq \omega < 2\pi\beta$. This is accomplished if the frequency responses of the reconstruction filters satisfy

$$\begin{cases} \sum_{n=1}^{N} G_n(\omega)\bar{F}_n(\omega) = T \\ \sum_{n=1}^{N} G_n(\omega)\bar{F}_n(\omega + 2\pi\beta k/N) = 0 \text{ for } k = 1, 2, \ldots, N-1. \end{cases} \tag{5.34}$$

This provides N equations that could be solved for the N desired filters for each frequency, ω. The equations can be written as a linear system, $\mathbf{Ax} = \mathbf{b}$ where the $N \times N$ system matrix

is

$$
\mathbf{A}(\omega) \triangleq \begin{pmatrix}
\bar{F}_1(\omega) & \bar{F}_2(\omega) & \cdots & \bar{F}_N(\omega) \\
\bar{F}_1(\omega + 2\pi\beta/N) & \bar{F}_2(\omega + 2\pi\beta/N) & \cdots & \bar{F}_N(\omega + 2\pi\beta/N) \\
\vdots & & \ddots & \vdots \\
\bar{F}_1(\omega + 2\pi\beta(N-1)/N) & \bar{F}_2(\omega + 2\pi\beta(N-1)/N) & \cdots & \bar{F}_N(\omega + 2\pi\beta(N-1)/N)
\end{pmatrix}
$$
(5.35)

and the length-N vector, $\mathbf{b}$, and the length-N vector of frequency responses, $\mathbf{x}$, are

$$
\mathbf{b} \triangleq \begin{pmatrix} T \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \mathbf{x}(\omega) \triangleq \begin{pmatrix} G_1(\omega) \\ G_2(\omega) \\ \vdots \\ G_N(\omega) \end{pmatrix}.
$$
(5.36)

This process provides the frequency responses of the reconstruction filters, $\{G_n(\omega)\}$ for frequencies with $\omega \in [0, 2\pi\beta)$.

5.4.3 Aliasing in Reconstruction from Multichannel Sampling

The above derivation was based on the assumption that the signal, $x(t)$, was bandlimited. Now, consider the impact of using these reconstruction filters when this assumption is invalid. To accomplish this, the input signal, $x(t)$, will be modeld as WSS random process with a PSD, $P_x(\omega)$, that is not necessarily bandlimited. This development will use some facts established in Appendix A, which includes results on WSS random processes and linear systems. The reconstruction filters are assumed to be designed to unambiguously recover signals that are bandlimited to $\omega \in \Omega$, which was defined in (5.30).

To quantify the impact of aliasing, the PSD of the reconstructed signal, $z(t)$, will be computed given the PSD of the input signal, $x(t)$. To accomplish this, the cross spectral densities (CSD) of each reconstruction filter's output will need to be computed.

The CSD of the signal from two channels, $y_n(t)$ and $y_{n'}(t)$, is related to the input PSD, $P_x(\omega)$, and the channel frequency responses, $F_n(\omega)$ and $F_{n'}(\omega)$, by

$$
P_{y_n y_{n'}}(\omega) = F_n^*(\omega) F_{n'}(\omega) P_x(\omega),
$$
(5.37)

93

and, after sampling and applying the reconstruction filters, the CSD of channel n and channel n' is

$$P_{z_n z_{n'}}(\omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} G_n^*(\omega) G_{n'}(\omega) P_{y_n y_{n'}}\left(\omega + \frac{2\pi}{T}k\right), \tag{5.38}$$

which is $2\pi\beta$-periodic. After the N channels are summed, the PSD of the reconstructed signal is found to be

$$\begin{aligned} P_z(\omega) &= \sum_{n=1}^{N} \sum_{n'=1}^{N} P_{z_n z_{n'}}(\omega) \\ &= \frac{1}{T} \sum_{k=-\infty}^{\infty} \left| \sum_{n=1}^{N} G_n(\omega) F_n\left(\omega + \frac{2\pi}{T}k\right) \right|^2 P_x\left(\omega + \frac{2\pi}{T}k\right). \end{aligned} \tag{5.39}$$

As in the deterministic case, this can be partitioned into the desired response and the contributions of the undesired aliases,

$$P_z(\omega) = \frac{1}{T} |H_0(\omega)|^2 P_x(\omega) + \frac{1}{T} \sum_{\substack{k=-\infty \\ k\neq 0}}^{\infty} |H_k(\omega)|^2 P_x\left(\omega + \frac{2\pi}{T}k\right) \tag{5.40}$$

for $\omega \in \Omega$. If the filters were successfully designed to reconstruct a signal that is bandlimited to $\omega \in \Omega$, then $H_0(\omega) = T$ and $H_1(\omega), \ldots, H_{N-1}(\omega) = 0$ for all $\omega \in \Omega$. In this case, the PSD of the reconstructed signal is

$$P_z(\omega) = T P_x(\omega) + \frac{1}{T} \sum_{\substack{k=-\infty \\ k\neq 0}}^{\infty} |H_k(\omega)|^2 P_x\left(\omega + \frac{2\pi}{T}k\right). \tag{5.41}$$

The average power of a random process is the integral of the PSD. In general, the alias-to-signal ratio (ASR) for a signal that is reconstructed over an interval, Ω, is

$$\text{ASR} \triangleq \frac{\sum\limits_{\substack{k=-\infty \\ k\neq 0}}^{\infty} \int_\Omega |H_k(\omega)|^2 P_x\left(\omega + \frac{2\pi}{T}k\right) d\omega}{\int_\Omega |H_0(\omega)|^2 P_x(\omega) \, d\omega}. \tag{5.42}$$

5.4.4 Special Case: Recurrent Sampling

Recurrent sampling is a special case of multichannel sampling, so the general method for deriving reconstruction filters of [Brown, 1981], which was described above, can also be used. A few results from the literature are briefly reviewed to show their relationship.

It is well-known that a uniformly-sampled, bandlimited signal can be reconstructed by using a sinc interpolator. A related result is an interpolation formula for a recurrently-sampled signal that employs a modified sinc interpolator [Yen, 1956]. The reconstruction formula of $x(t)$ from its recurrent samples, $\{x[n,m]\}$, is

$$\hat{x}(t) = \sum_{m=-\infty}^{\infty} \sum_{n=0}^{N-1} x[n,m] \, \Psi_{nm}(t) \tag{5.43}$$

where the interpolation kernel for a sample point, t, is

$$\Psi_{nm}(t) \triangleq \frac{(-1)^{mN}}{\prod\limits_{\substack{n'=0 \\ n' \neq n}}^{N-1} \sin\left(\pi\left(T_n - T_{n'}\right)/T\right)} \frac{\prod\limits_{n'=0}^{N-1} \sin\left(\pi\left(t - T_{n'}\right)/T\right)}{\pi\left(t - \left(T_n + mT\right)\right)/T} \tag{5.44}$$

with

$$a_n \triangleq \frac{1}{\prod\limits_{\substack{n'=0 \\ n' \neq n}}^{N-1} \sin\left(\pi\left(T_n - T_{n'}\right)/T\right)} \tag{5.45}$$

for $n = 0, 1, \ldots, N-1$.

Rewrite (5.44) as

$$\Psi_{nm}(t) = a_n \mathrm{sinc}\left(\left(t - \tau_{n,m}\right)/T\right) \prod_{\substack{n'=0 \\ n' \neq n}}^{N-1} \sin\left(\pi\left(t - \tau_{n',m}\right)/T\right) \tag{5.46}$$

where $\mathrm{sinc}(x) \triangleq \sin(\pi x)/(\pi x)$. Writing the recurrent-sampling interpolation kernel as (5.46) highlights its relationship to the sinc interpolator. Indeed, when recurrent sampling reduces to uniform sampling, the kernel of (5.46) reduces to the standard sinc interpolator for uniform samples.

The interpolation filter of [Yen, 1956] for a recurrently-sampled signal was implemented as a filterbank in [Eldar and Oppenheim, 2000]. The reconstruction of (5.43) can be written as

$$\hat{x}(t) = \sum_{n=0}^{N-1} \left(s_n * h_n\right)(t) \tag{5.47}$$

where $s_n(t)$ is the sampled signal from channel n,

$$s_n(t) \triangleq \sum_{m=-\infty}^{\infty} x[n,m] \, \delta\left(t - \tau_{n,m}\right), \tag{5.48}$$

95

and $h_n(t)$ is the impulse response of the reconstruction filter applied to channel n,

$$h_n(t) \triangleq a_n \text{sinc}\left(t/T\right) \prod_{\substack{n'=0 \\ n' \neq n}}^{N-1} \sin\left(\pi\left(t + T_n - T_{n'}\right)/T\right) \tag{5.49}$$

While the method of [Brown, 1981] provides a method for generating reconstruction filters for a general multichannel sampling problem, (5.49) is a closed form for the impulse response of the same reconstruction filter.

5.5 Summary

A rigorous method for calculating the impact of aliasing when sampling a WSS random process has been developed. This approach naturally handles uniform sampling as well as recurrent sampling. While the sampling theorem states that one can perfectly reconstruct a bandlimited signal from its samples, this result allows the aliasing that occurs when this assumption is not valid to be quantified.

CHAPTER 6

ALIASING IN MULTICHANNEL SAR/SAS

6.1 Overview

A synthetic aperture imaging system provides a sampled, discrete-space representation of the desired continuous-space aperture. The sampling theorem applies to spatial sampling just as it does to temporal, and the impact of aliases will inevitably degrade image quality. Their impact is quantified by the azimuth ambiguity-to-signal ratio (AASR), which is the ratio of the power within an image pixel from the undesired ambiguous returns to the desired return. In many systems, this term is the dominant contribution to the system multiplicative noise ratio (MNR) [Carrara et al., 1995], which provides a fundamental limit on the achievable image quality. Note that range ambiguities are also likely to be present because synthetic aperture imaging systems use a pulsed waveform, but the present discussion is limited to azimuth ambiguities that result from along-track sampling.

To understand the operation of synthetic aperture radar (SAR) and synthetic aperture sonar (SAS) systems, one can consider the problem in the following progression [Showman, 2010].

SAR/SAS as array processing The first approach is to consider the synthetic aperture simply as a sampled version of the desired large physical aperture. This idea provides a justification of the fact that synthetic aperture techniques can provide improved angular resolution and provides a bound on the required along-track sampling interval.

SAR/SAS as a Doppler processor Early SAR systems extended standard radar techniques of measuring range and Doppler to coarse imaging modes based on the observation that the Doppler frequency of a stationary target observed from a moving platform is related to the angle from the platform to the target. However, this simplified signal model becomes increasingly inaccurate as the resolution is improved. A more

sophisticated model is required to allow the design of effective image formation processors.

SAR/SAS as a general inverse problem (spatial-frequency sampling) The general approach allows the possibility of nonuniform sampling in the spatial frequency domain. Typically, the spatial frequency samples can be interpolated onto a uniform grid, which allows an image to be efficiently formed with the use of the fast Fourier transform (FFT). This motivates algorithms like the polar-format algorithm for spotlight and the range migration algorithm for stripmap.

These three approaches provide three increasingly accurate methods for understanding the along-track sampling requirements of synthetic aperture imaging and calculating the impact of ambiguous returns on image quality.

If the synthetic aperture is considered simply as an array of along-track samples, then a spatial sample must be collected at half-wavelength spacing, $\lambda/2$, to ensure Nyquist sampling if an omnidirectional beampattern is used for each pulse. If a directional element with an along-track extent, D, is used for transmit and receive, this spacing can be relaxed to $D/2$. The array processing interpretation of SAR is useful in understanding the worst-case along-track sampling requirements, but it leads to the so-called myth of the minimum SAR antenna constraint [Freeman et al., 2000].

The justification for the standard methods of calculating AASR is the Doppler-based approach to understanding synthetic aperture imaging. The idea is that a moving platform uses a number of pulses to construct Doppler filters where each Doppler frequency corresponds to a target at a particular azimuth. Because a signal with a particular Doppler frequency is sampled at a rate corresponding to the pulse repetition frequency (PRF), aliasing will occur for targets with a Doppler frequency greater than the PRF. This process is referred to as slow-time sampling.

A Doppler-based method for calculating the AASR has been discussed previously in the SAR literature [Curlander and McDonough, 1991]. In [Bayma and McInnes, 1975], a

method is provided for quantifying the impact of range and Doppler ambiguities, which described the relationships between antenna size, area-coverage rate, and ASR. This analysis was extended to generate curves of AASR as a function of along-track sampling rate (relative to antenna length) and of processed Doppler bandwidth/cross-range resolution [Mehlis, 1980]. A discussion in the context of SAS was presented in [Hawkins, 1996]. Note that all of these works have focused on stripmap synthetic aperture imaging, and they also assume that the along-track samples are uniformly spaced.

In the following, a number of contributions are presented. First, the standard Doppler-based method for stripmap AASR calculation is rigorously developed based on the concept of sampling a wide-sense stationary (WSS) random process. This method is used to calculate AASR for both stripmap and spotlight modes. A set of numerical results is presented to illustrate the utility of these methods in SAR/SAS system design and performance prediction. This analysis includes multichannel synthetic aperture imaging systems including the possibility of nonuniform along-track sampling.

6.2   Doppler-Based AASR Calculation

Let $x(t)$ be the slow-time signal for a particular range bin and neglect any range-walk effects. For distributed clutter, this can be treated as a zero-mean, WSS random process with a power spectral density (PSD), $P_x(\omega)$. For convenience, define the Doppler power spectrum, $P(f_D)$, with $f_D \triangleq \omega/(2\pi)$. This PSD describes the average power of signals with a Doppler frequency, f_D, which is in units of cycles per second.

Because the terrain is assumed to be a distributed target that contains an arbitrary number of unresolved targets, the slow-time signal should be temporally white. However, because the signal is observed using a directional antenna, which introduces correlation into the slow-time signal, the Doppler PSD is proportional to the two-way beampattern. An illustration of a simplified two-dimensional stripmap geometry is presented in Figure 6.1. The angle, θ, characterizes the point within the integration angle relative to the indicated

Figure 6.1: Two-dimensional stripmap collection geometry.

point. The point, $\theta = 0$, corresponds to the point of closest approach.

6.2.1 SAR as Doppler Processing

A simplified interpretation of SAR image formation processing is that it can resolve targets in azimuth by applying Doppler processing. For coarse cross-range resolutions, an image of acceptable quality may be formed using Doppler processing, but more sophisticated processing, such as the polar format algorithm, must be employed as the resolution improves. Still, thinking of SAR image formation as Doppler processing is a useful concept for understanding the impact of azimuth ambiguities on SAR image quality.

By transmitting a single pulse, a radar is able to measure the range to a target by calculating the time between transmission and reception of the pulse. By transmitting multiple pulses, a stationary radar can also determine the velocity of a moving target by measuring the rate of change in phase of the return over a number of pulses, which is proportional to the change in range. The rate of phase change is called the Doppler frequency, which is related to the target velocity by

$$ f_D = \frac{2v_r}{\lambda}, \tag{6.1} $$

where λ is the radar wavelength, and v_r is the radial velocity of the target, i.e., the component of the target velocity vector in the direction of the radar. The radar cannot measure the tangential component of the target velocity vector because it does not result in a change in

range to the target.

In synthetic aperture imaging, it is assumed that the targets are stationary and the radar is moving. Consider the two-dimensional geometry of Figure 6.1. If v is the platform velocity, then the Doppler frequency of a return from a stationary target with an angle, θ, is

$$f_D = \frac{2v}{\lambda} \sin \theta. \tag{6.2}$$

Because of this relationship between angle and Doppler, the angle to a target can be determined by applying Doppler processing. A range-Doppler map can be formed by applying the discrete Fourier transform across the pulses of range compressed data. By rearranging (6.2), Doppler frequency can be mapped to angle by

$$\theta = \arcsin\left(\frac{\lambda/2}{v} f_D\right). \tag{6.3}$$

As the radar moves, it transmits pulses at at a particular rate, namely the PRF. Because of this sampling, Doppler frequencies over an interval equal to the PRF can be unambiguously measured. If the PRF is f_p, then a set of unambiguous Doppler frequencies is defined by $|f_D| \leq f_p/2$. A target with a Doppler frequency, f_D, will have the same apparent Doppler frequency as a target with a Doppler frequency, $\tilde{f}_{D/n}$, given by

$$\tilde{f}_{D/n} = f_D + n f_p \tag{6.4}$$

for $n \neq 0$.

6.2.2 The Doppler PSD for Broadside Stripmap Collections

For a stationary target and a moving platform, the Doppler frequency, f_D, and angle, θ, are related according to (6.2). For a fixed platform velocity and center frequency, this establishes a mapping between an angle, θ, and a Doppler frequency, f_D. Because of this, the Doppler power spectrum, $P(f_D)$ is related to the average power of returns as a function of angle, θ.

For a given range bin, the relative power of returns as a function of angle is driven primarily by the antenna gain. Let D be the azimuthal extent of an idealized, uniformly

illuminated antenna. The one-way gain for an angle, θ, is proportional to

$$G(\theta) = \text{sinc}^2\left(\frac{D}{\lambda}\sin\theta\right), \tag{6.5}$$

where the sinc function is defined as

$$\text{sinc}(x) \triangleq \begin{cases} \frac{\sin(\pi x)}{\pi x}, & \text{for } x \neq 0 \\ 1, & \text{for } x = 0. \end{cases} \tag{6.6}$$

The antenna gain varies with angle, but (6.3) describes the relationship between an angle, θ, and a Doppler frequency, f_D. For the broadside geometry of Figure 6.1, where the antenna is oriented in a direction that is orthogonal to the velocity vector, the Doppler PSD is

$$P(f_D) = \text{sinc}^2\left(\frac{D_{\text{TX}}/2}{v}f_D\right)\text{sinc}^2\left(\frac{D_{\text{RX}}/2}{v}f_D\right)\text{rect}\left(\frac{f_D}{4v/\lambda}\right). \tag{6.7}$$

where D_{TX} and D_{RX} are the lengths of the transmit and receive antenna, respectively. For stripmap imaging, where the antenna is not reoriented during the collection, the Doppler PSD of (6.7) is constant and does not change as the platform moves.

6.2.3 AASR for Stripmap Synthetic Aperture Imaging

The Doppler PSD of clutter for a stripmap collection was given by (6.7). This can be used to calculate the AASR for a given along-track sampling. In the following, two cases are considered. In the first, it is assumed that the along-track samples are uniformly spaced. This occurs in a single-channel system when the PRF is constant, and it also can occur for a multichannel system when the PRF is constant and chosen appropriately. Next, the more general case of a multichannel system with a constant PRF but with an arbitrary channel spacing is considered. These two examples correspond to uniform and to recurrent sampling, respectively.

6.2.3.1 Single-Channel Stripmap: Uniform Sampling

If $P(f_D)$ is the Doppler PSD, the PSD after sampling with a constant PRF, f_p, is

$$\tilde{P}(f_D) = \underbrace{P(f_D)}_{\text{Unambiguous}} + \underbrace{\sum_{m \neq 0} P(f_D + mf_p)}_{\text{Ambiguous}} \tag{6.8}$$

for unambiguous Doppler frequencies, $f_D \in \left[-f_p/2, f_p/2\right)$. The aliased PSD of (6.8) demonstrates how a target at a Doppler frequency, f_D, is affected by targets at other Doppler frequencies.

To resolve a particular target in cross-range, the radar will process a number of pulses that correspond to a particular integration angle, θ_{int}, i.e., the pulses that are used correspond to points in the synthetic aperture with $|\theta| \leq \theta_{\text{int}}/2$ where θ is defined as in Figure 6.1.

As the platform flies the synthetic aperture, the target angle, θ, changes. Equivalently, because of the relationship between angle and Doppler frequency in (6.2), the Doppler frequency of the target changes. For a given platform velocity, v, and wavelength, λ, an integration angle, θ_{int}, can be related to a processed Doppler bandwidth, β, by

$$\beta = \frac{4v}{\lambda} \sin\left(\frac{\theta_{\text{int}}}{2}\right). \tag{6.9}$$

For a broadside stripmap collection, the target Doppler frequency will vary with $-\beta/2 \leq f_D \leq \beta/2$. As shown in (6.8), as the Doppler frequency of the target varies, the nature of the aliasing varies. To account for this, the unambiguous power will be computed by

$$P_{\text{unambig}} = \int_{-\beta/2}^{\beta/2} P(f) \, df, \tag{6.10}$$

and the ambiguous power is

$$P_{\text{ambig}} = \int_{-\beta/2}^{\beta/2} \sum_{m \neq 0} P(f + mf_p) \, df. \tag{6.11}$$

The corresponding AASR is

$$\text{AASR} \triangleq \frac{P_{\text{ambig}}}{P_{\text{unambig}}} = \frac{\int\limits_{-\beta/2}^{\beta/2} \sum\limits_{m \neq 0} P\left(f + mf_{\text{p}}\right) df}{\int\limits_{-\beta/2}^{\beta/2} P(f)\, df}. \tag{6.12}$$

6.2.3.2 Multichannel Stripmap: Recurrent Sampling

Just as in the single-channel case, the AASR for a multichannel system can be calculated by first determining the Doppler PSD after aliasing. If the PRF, f_{p}, and platform velocity, v, are both constant, then the synthetic aperture will be recurrently sampled. The resulting PSD after this sampling is

$$\tilde{P}\left(f_{\text{D}}\right) = P\left(f_{\text{D}}\right) + \sum\limits_{m \neq 0} \alpha_m P\left(f_{\text{D}} - mf_{\text{p}}\right), \tag{6.13}$$

where $\{\alpha_m\}$ are the recurrent sampling weights given by (5.9), which depend on the spacing of the receive elements within the physical array. The multichannel AASR is

$$\text{AASR} = \frac{\int\limits_{-\beta/2}^{\beta/2} \sum\limits_{m \neq 0} \alpha_m P\left(f + mf_{\text{p}}\right) df}{\int\limits_{-\beta/2}^{\beta/2} P(f)\, df}. \tag{6.14}$$

6.2.4 CNR and SNR in Stripmap SAR/SAS

In the previous section, AASR was computed by invoking a Doppler processing interpretation of synthetic aperture imaging. This involved the concept of the processed Doppler bandwidth that corresponds to a particular integration angle. The AASR varies as a function of the processed Doppler bandwidth, and the cross-range resolution becomes more coarse as less Doppler bandwidth is used. However, the clutter-to-noise ratio (CNR) and the signal-to-noise ratio (SNR) of a point target will also vary. In this section, the relationship between processed Doppler bandwidth and SNR/CNR will be explored for both the stripmap and spotlight cases.

The CNR describes the relative power of returns from the terrain relative to additive noise, while SNR is defined for a point target. Because the return from a point target will be coherent from pulse to pulse, integration gain allows the SNR to be improved as more pulses are collected. However, because the terrain is a distributed target, speckle prevents an increase in CNR.

The CNR can be used to define the system noise equivalent reflectivity, σ_N. A normalized reflectivity coefficient, σ_0, can be defined to characterize the reflectivity of a distributed target like terrain. Just as the radar cross section (RCS) can be defined for a point target, the reflectivity, σ_0, can be defined for a distributed target where the equivalent RCS of a clutter patch with a given area, A, is $\sigma = A\sigma_0$, i.e., σ_0 can be interpreted as the RCS in square meters of a one square meter patch of clutter.

The CNR will vary depending on the reflectivity of the underlying terrain. A system-level metric can be defined that is independent of the clutter type. This metric is called the noise equivalent backscatter coefficient and is denoted σ_N [Curlander and McDonough, 1991]. It is also called the noise equivalent sigma zero (NESZ). As suggested by the term NESZ, σ_N is the average backscatter coefficient in a noise-only pixel. This motivates the definition that

$$\text{CNR} = \frac{\sigma_0}{\sigma_N}. \tag{6.15}$$

In the following, the change in CNR as a function of processed Doppler bandwidth will be derived. One interpretation of this is a change in σ_N.

For stripmap, the average clutter power is computed by integrating the clutter Doppler PSD over the processed Doppler bandwidth.

$$P_C = \int_{-\beta/2}^{\beta/2} P(f)\, df. \tag{6.16}$$

The corresponding noise power is

$$P_N = \beta N_0, \tag{6.17}$$

where N_0 is the noise PSD. The clutter-to-noise ratio (CNR) is

$$\text{CNR} \triangleq \frac{P_\text{C}}{P_\text{N}} = \frac{1}{N_0}\frac{1}{\beta} \int_{-\beta/2}^{\beta/2} P(f)\, df. \tag{6.18}$$

As the bandwidth becomes arbitrarily small, the CNR converges to

$$\text{CNR}_0 \triangleq \lim_{\beta \to 0} \text{CNR} = \frac{1}{N_0}, \tag{6.19}$$

where it is assumed that the clutter Doppler PSD has been normalized so that $P(0) = 1$. The improvement in the CNR by integrating over a Doppler bandwidth, β, is

$$\frac{\text{CNR}}{\text{CNR}_0} = \frac{1}{\beta} \int_{-\beta/2}^{\beta/2} P(f)\, df. \tag{6.20}$$

The SNR for a point target as a function of processed Doppler bandwidth is calculated in a similar manner. The signal power is

$$P_\text{S} = \sigma \int_{-\beta/2}^{\beta/2} G^2(f)\, df \tag{6.21}$$

where $G(f)$ is the one-way antenna gain when the target has a Doppler frequency, f. Using this, the SNR for a point target is found to be

$$\text{SNR} \triangleq \frac{P_\text{S}}{P_\text{N}} = \beta \frac{\sigma}{N_0}. \tag{6.22}$$

The SNR increases linearly as the Doppler bandwidth is increased, and

$$\text{SNR}_0 \triangleq \lim_{\beta \to 0} \text{SNR} = 0. \tag{6.23}$$

6.2.5 AASR for Spotlight Synthetic Aperture Imaging

In stripmap synthetic aperture imaging, the antenna is not reoriented during the collection. Because of this, the Doppler PSD of (6.7) does not vary as the platform moves. For a broadside collection, the beam is always oriented orthogonal to the velocity vector so that the peak of the antenna beampattern will always correspond to zero Doppler. Even if the antenna pointing vector is not orthogonal to the velocity vector, the peak of the beam will always be aimed at a point with a constant Doppler frequency.

A spotlight system will continuously reorient the antenna to constantly aim the beam at a point on the ground, which is called the scene reference point (SRP). As a result, the Doppler PSD varies over the synthetic aperture. The peak of the antenna beampattern is always in the direction of the SRP, so the peak of the Doppler PSD will vary as the Doppler frequency of the SRP varies.

Another difference between AASR for spotlight and stripmap synthetic aperture imaging is a result of the fact that stripmap image quality does not vary in the along-track direction. In a stripmap collection, every point at a given range at the point of closest approach will have the same gain variation over its integration angle. However, for spotlight, points displaced from the SRP in the cross-range direction will experience difference gain variations. For example, the SRP will have no gain variation in spotlight, but it would come into and go out of the beam like any other point in a stripmap collection.

The Doppler PSD when the platform is at an angle, θ, relative to the SRP is

$$P(f_{\mathrm{D}};\theta) = G^2\left(\theta + \arcsin\left(\frac{\lambda f_{\mathrm{D}}}{2v}\right)\right),\tag{6.24}$$

for $|f_{\mathrm{D}}| \leq 2v/\lambda$ where $G^2(\theta)$ is the two-way antenna gain, which assumes that the same aperture is used on transmit and receive. This can be extended to the case where different apertures are used on transmit and receive in a straightforward manner.

The Doppler frequency of a target at an angle, θ_0, relative to the SRP when the platform is at the center of its integration angle is

$$f_{\mathrm{D}}(\theta;\theta_0) = \frac{2v}{\lambda}\left(\sin\theta_0 - \sin\theta\right).\tag{6.25}$$

This can be used to determine the point, θ, when the target is at a given Doppler frequency, f_{D},

$$\theta(f_{\mathrm{D}}) = \arcsin\left(\sin\theta_0 - \frac{\lambda f_{\mathrm{D}}}{2v}\right).\tag{6.26}$$

The notation, $\theta(f_{\mathrm{D}})$, is used to highlight the fact that this angle varies as a function of Doppler frequency.

The unambiguous power is found by integrating the PSD of (6.24) over the Doppler frequencies that the target is observed at over the integration angle,

$$
\begin{aligned}
P_{\text{unambig}} &\triangleq \int_{-\beta/2}^{\beta/2} P(f; \theta(f)) \, df \\
&= \int_{-\beta/2}^{\beta/2} P\left(f; \arcsin\left(\sin\theta_0 - \frac{\lambda f}{2v}\right)\right) df \\
&= \int_{-\beta/2}^{\beta/2} G^2\left(\arcsin\left(\sin\theta_0 - \frac{\lambda f}{2v}\right) + \arcsin\left(\frac{\lambda f}{2v}\right)\right) df.
\end{aligned}
\tag{6.27}
$$

Similarly, the ambiguous power is found to be

$$
P_{\text{ambig}} = \sum_{m \neq 0} \int_{-\beta/2}^{\beta/2} G^2\left(\arcsin\left(\sin\theta_0 - \frac{\lambda f}{2v}\right) + \arcsin\left(\frac{\lambda(f + mf_{\text{p}})}{2v}\right)\right) df.
\tag{6.28}
$$

The AASR is

$$
\text{AASR} \triangleq \frac{P_{\text{ambig}}}{P_{\text{unambig}}} = \frac{\sum\limits_{m \neq 0} \int\limits_{-\beta/2}^{\beta/2} G^2\left(\arcsin\left(\sin\theta_0 - \frac{\lambda f}{2v}\right) + \arcsin\left(\frac{\lambda(f+mf_{\text{p}})}{2v}\right)\right) df}{\int\limits_{-\beta/2}^{\beta/2} G^2\left(\arcsin\left(\sin\theta_0 - \frac{\lambda f}{2v}\right) + \arcsin\left(\frac{\lambda f}{2v}\right)\right) df}
\tag{6.29}
$$

For small integration angles, (6.29) can be approximated by

$$
\text{AASR} \approx \frac{\sum\limits_{m \neq 0} G^2\left(\theta_0 + \frac{\lambda f_{\text{p}}}{2v}m\right)}{G^2(\theta_0)},
\tag{6.30}
$$

which is independent of integration angle and depends only on the target offset from scene center, θ_0.

6.3 Numerical Results: Impact of AASR on System Design

The Doppler PSD is shaped by the antenna, and the impact of the aliases depends on the along-track sampling. System performance is quantified as a function of along-track sampling relative to the antenna size. It is assumed that the apertures used on transmit and receive are uniformly-illuminated linear apertures with the sinc beampattern given by (6.5), which lead to the stripmap Doppler PSD of (6.7) and the spotlight Doppler PSD of (6.24). The calculations can be repeated for arbitrary beampatterns.

The results presented in this section assume that the platform and the targets are constrained to lie in the same two-dimensional plane, i.e., the grazing angle is zero. Only Doppler/azimuth ambiguities are considered, so the calculations are limited to AASR. The analysis can be extended to include range ambiguities as well to calculate the ASR.

6.3.1 The Along-Track Oversampling Factor and Area-Coverage Rate

In presenting results for both stripmap and spotlight AASR, it will be convenient to define the along-track oversampling factor, η, which normalizes the PRF, f_p, by the aperture length, D.

The Doppler extent of mainlobe clutter is approximately $v/(D/2)$ where D is the length of the aperture. To unambiguously sample this Doppler extent, the PRF must be at least $v/(D/2)$, which is equivalent to sampling the synthetic aperture with an along-track sampling interval of $D/2$. Because of this, $D/2$ along-track sampling corresponds to the critical "Nyquist" sampling rate. For the single-channel stripmap analysis, it is assumed that the transmit and receive apertures are the same size, i.e., $D_{TX} = D_{RX} = D$. Because critical along-track sampling is achieved when $v/f_p = D/2$, the along-track oversampling factor is given by

$$\eta \triangleq \frac{D/2}{v/f_p}. \tag{6.31}$$

Critical sampling corresponds to $\eta = 1$. For $\eta = 2$, $D/4$ sampling is provided.

For a fixed velocity and aperture size, increasing the PRF leads to an increased along-track oversampling factor, which provides improved AASR. However, this comes at the cost of area-coverage rate (ACR). The maximum (slant) range swath is constrained by the PRF according to

$$R_{\text{swath}} \leq \frac{c/2}{f_p}. \tag{6.32}$$

Note that other limitations, e.g., transmit power, can further constrain the imaged range swath, so the PRF provides an upper bound on range swath but no guarantee that this bound is achievable.

The maximum ACR is provided by stripmap operation where the platform moves at a velocity, v, and continuously forms imagery with a range extent, R_{swath}. The corresponding ACR is

$$\text{ACR}_{\text{max}} \triangleq vR_{\text{swath}}. \tag{6.33}$$

An upper bound is placed on ACR by the range swath limit of (6.32), which can be written in terms of the along-track oversampling factor as

$$\text{ACR}_{\text{max}} \leq \left(\frac{c}{2}\right)\left(\frac{v}{f_{\text{p}}}\right) = \left(\frac{c}{2}\right)\left(\frac{D}{2}\right)\frac{1}{\eta}. \tag{6.34}$$

Note that this ACR is relative to a slant-plane image, but it can be adjusted by the grazing angle if desired.

The bound of (6.34) suggests that increasing the aperture size, D, provides improved ACR, but this provides additional constraints on the image. The larger aperture corresponds to a narrower beamwidth. In stripmap, this leads to coarser cross-range resolution because $\Delta_{\text{CR}} \geq D/2$. In stripmap, this limits the cross-range extent of the image.

To summarize, the PRF can be increased to improve the along-track sampling factor, which can provide improved AASR by rejecting additional Doppler-ambiguous clutter. However, this can limit the ACR by limiting the extent of the unambiguous range swath.

6.3.2 Single-Channel Stripmap AASR

In stripmap mode, the system continuously images with the ACR given by (6.33). However, the finest cross-range resolution that can be provided by a stripmap synthetic aperture imaging system is limited by its beamwidth. While the finest cross-range resolution, $\Delta_{\text{CR}} = D/2$, corresponds to the case where the integration angle corresponds to the full beamwidth, a processor is free to use only a subset of this angle. This results in coarser cross-range resolution, but it will provide improved AASR.

Similar to the along-track oversampling factor, η, a cross-range resolution spoiling factor, ξ, can be defined. As ξ increases, the cross-range resolution degrades, but improved

AASR can be achieved. For a given spoiling factor, the corresponding cross-range resolution is

$$\Delta_{CR} = \xi D/2. \qquad (6.35)$$

In stripmap synthetic aperture imaging, AASR can be improved at the cost of ACR and cross-range resolution by η and ξ as shown in (6.34) and (6.35). The results of these tradeoffs are presented in Figure 6.2. A similar figure was included in [Hawkins, 1996]. Each curve in Figure 6.2 corresponds to a given cross-range resolution spoiling factor, ξ. For each curve, the AASR is given as a function of along-track oversampling factor, η. These results are extended to the case where the transmit aperture is smaller than the receive aperture in Figure 6.3.
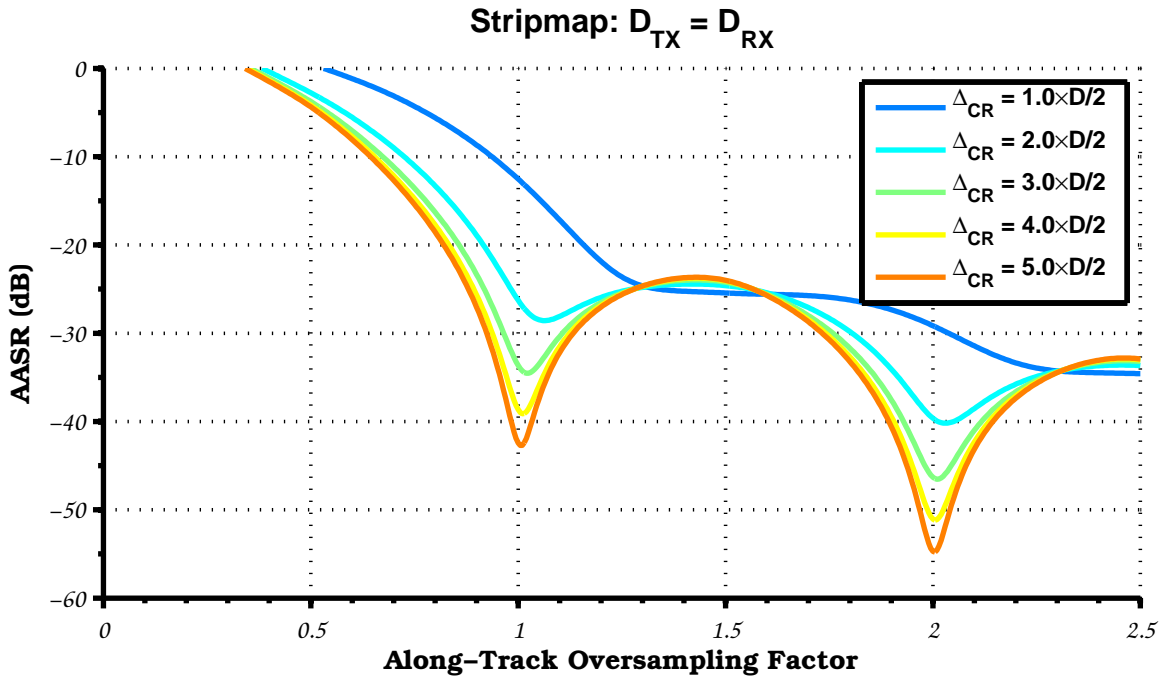

Figure 6.2: AASR for stripmap synthetic aperture imaging.

The previous analysis demonstrated how AASR performance can be improved by forming imagery with degraded cross-range resolution. According to (6.18), the CNR and, equivalently, σ_N, will also vary. This is illustrated in Figure 6.4. The cross-range resolution is degraded by integrating over a smaller angle. This means that the beampattern rolls
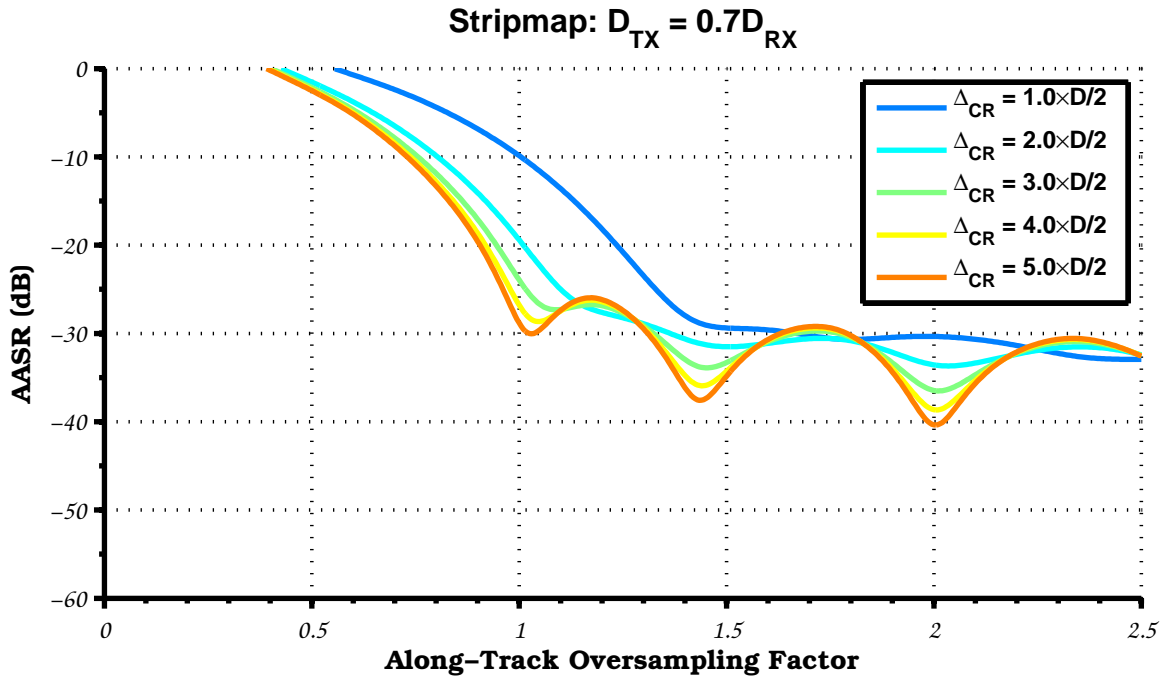
Figure 6.3: AASR for stripmap synthetic aperture imaging where the physical transmit aperture is larger than the receive aperture.

off less over the synthetic aperture, which leads to the slight improvement in σ_N. The effect is more pronounced if a larger transmit element is used.
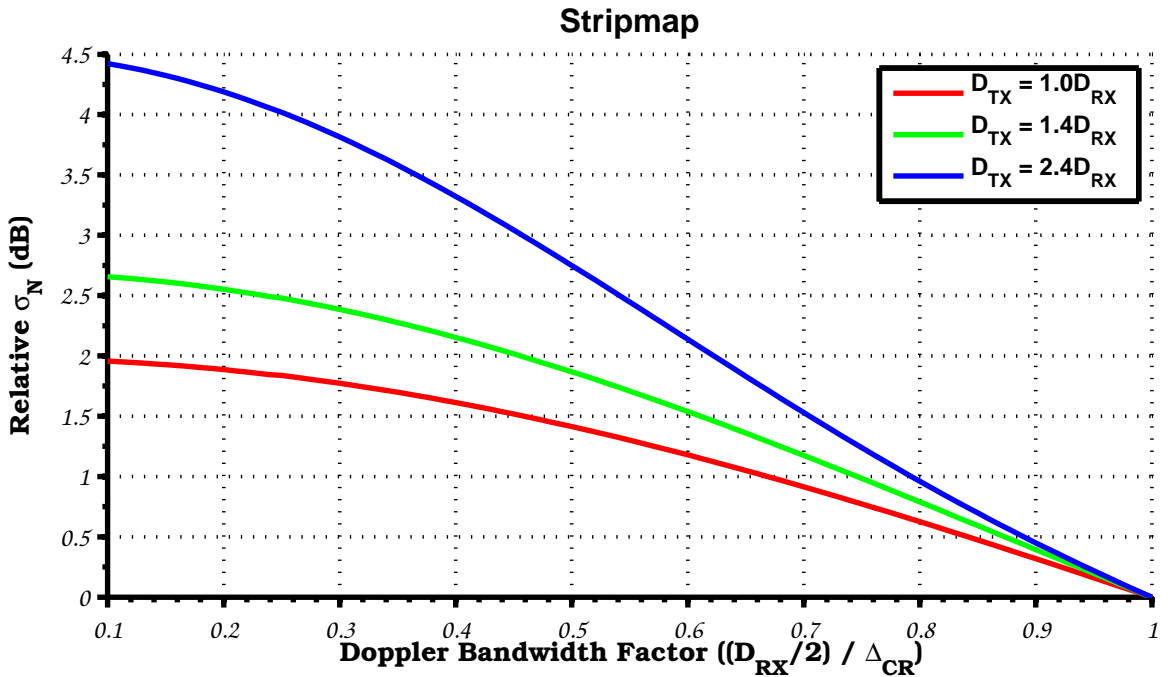

Figure 6.4: Variation in σ_N as a function of processed Doppler bandwidth for stripmap.

6.3.3 Multichannel Stripmap AASR: Uniform Sampling

Now, consider a multichannel synthetic aperture imaging system that uses N elements on receive, where each has an along-track extent, D_{RX}, and a single transmit element with an extent, D_{TX}. Assume that the receive elements are uniformly spaced and that the length of the full receive array is

$$L_{RX} \triangleq N D_{RX}. \tag{6.36}$$

For each pulse, N along-track samples are simultaneously collected with a spacing, $D_{RX}/2$. To provide a uniformly sampled synthetic aperture, the PRF is chosen so that $v/f_p = L_{RX}/2$, which limits the maximum ACR to

$$\text{ACR}_{max} = \frac{L_{RX}}{2} \frac{c}{2}. \tag{6.37}$$

The cross-range resolution is constrained by the maximum integration angle supported by the beamwidth, so the transmit aperture limits the finest achievable cross-range resolution. As before, a cross-range resolution spoiling factor, $\xi \geq 1$, can be introduced with a corresponding resolution of

$$\Delta_{CR} = \xi D_{TX}/2. \tag{6.38}$$

To ensure that each receive aperture is no larger than the transmit aperture, it is required that the number of receive channels, N, satisfy

$$N \geq \frac{L_{RX}}{D_{TX}}. \tag{6.39}$$

The along-track samples are spaced by $D_{RX}/2$, and the critical sampling interval is $D_{TX}/2$. The along-track oversampling factor for the multichannel case is

$$\eta_{multi} = \frac{D_{TX}}{D_{RX}}. \tag{6.40}$$

The length of the receive array, L_{RX}, the length of each receive element, D_{RX}, and the length of the transmit aperture, D_{TX}, as well as the cross-range resolution spoiling factor, ξ, are chosen to satisfy requirements on ACR, cross-range resolution, and AASR. A methodology for doing so follows.

- The length of the transmit aperture, D_{TX}, must be small enough to support the required cross-range resolution; see (6.38).

- The length of the receive array, L_{RX}, must be sufficient to support the required ACR; see (6.37).

- The number of elements in the receive array, N, must be suitably large so that the length of an individual receive element, D_{RX}, is small enough to provide a sufficient along-track sampling interval, $D_{\text{RX}}/2$, that yields an acceptable AASR; see (6.40). Also, it is required that $N \geq N_{\text{min}}$; see (6.39).

The ACR and cross-range resolution requirements determine the length of the receive and transmit apertures, respectively, but this procedure allows two methods for achieving the AASR requirement. The extent of each receive channel can be no larger than the minimum transmit aperture size without degrading the cross-range resolution. However, as the receive aperture is divided into more and more channels, the along-track sampling increases. This is equivalent to having a transmit aperture that is larger than a single receive element, i.e., $D_{\text{TX}} \geq D_{\text{RX}}$. Also, as in the single-channel case, only a portion of the integration angle supported by the transmit aperture can be processed. The resulting cross-range resolution corresponding to a spoiling factor, ξ, was given by (6.38).

These tradeoffs are explored in Figure 6.5. While the single-channel results of Figure 6.2 and Figure 6.3 showed results as a function of the along-track oversampling factor, η, Figure 6.5 corresponds to $\eta = 1$, which is required to provide uniform along-track sampling. In the upper plot, Figure 6.5a, the AASR is plotted as the ratio, $D_{\text{TX}}/D_{\text{RX}}$ increases for a number of cross-range resolution spoiling factors, ξ. Identical results are shown in Figure 6.5b, but the scale has been normalized by cross-range resolution.

Two important conclusions can be made from Figure 6.5 regarding the optimal ratio, $D_{\text{TX}}/D_{\text{RX}}$. First, for a given cross-range resolution requirement, the best AASR is achieved by using the largest possible transmit aperture and integrating over the entire beamwidth,

i.e., $\xi = 1$. Second, observe that AASR does not monotonically decrease as the ratio, D_{TX}/D_{RX}, increases. For the $\xi = 1$ case, AASR is minimized when $D_{TX} \approx (k + .5)\, D_{RX}$ for integers $k = 1, 2, \ldots$.

(a) AASR versus transmit aperture size and integration angle
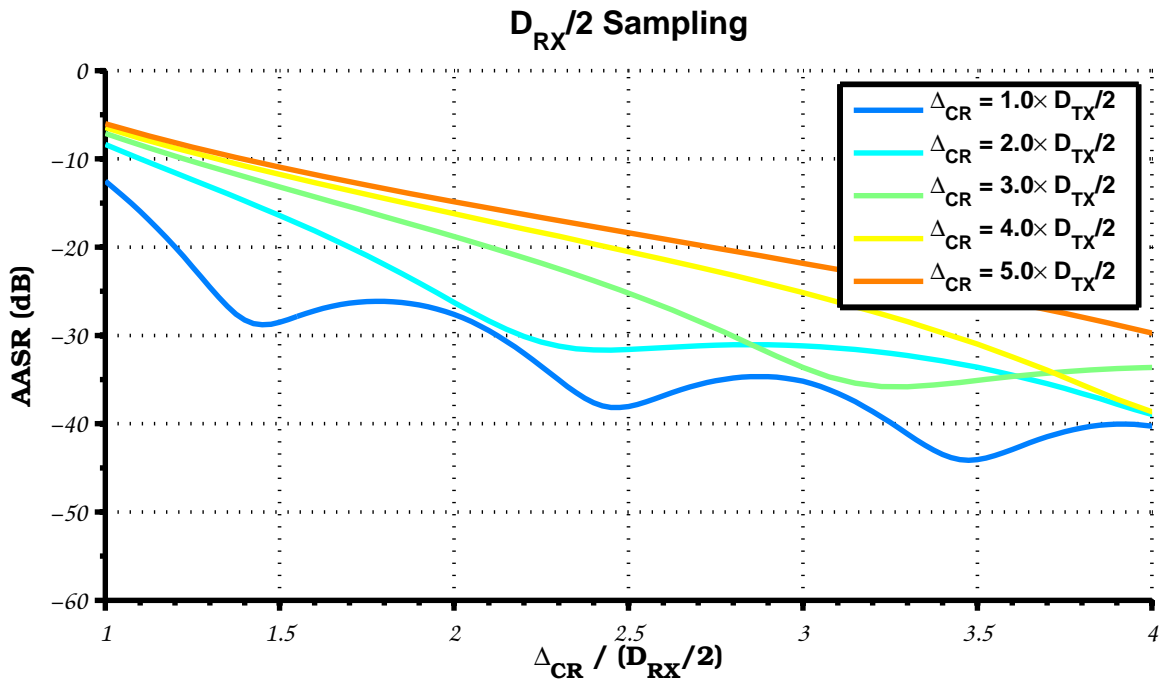
(b) Normalized by cross-range resolution

Figure 6.5: Stripmap AASR as a function of the relative transmit aperture size.

6.3.4 Single-Channel Spotlight AASR

As suggested by (6.30), spotlight AASR is not strongly dependent on integration angle. This is confirmed by examining Figure 6.6, which presents spotlight AASR for a target at scene center. The antenna is aimed at this target for the entire integration angle, and the beampattern apparently rotates about this point over the synthetic aperture. As is seen from Figure 6.6, the dominant effect of increasing integration angle is to prevent the nulls that the ideal beampattern provides at $D/2$ and $D/4$ sampling from perfectly canceling the ambiguous returns.

While spotlight AASR is not sensitive to integration angle, it does vary over the image. In Figure 6.7, AASR is plotted for a number of target offsets. This is parameterized as a percentage of the beamwidth. For targets near the edge of the scene, AASR is much worse for a fixed along-track oversampling factor. Consequently, while a certain PRF may be sufficient to provide acceptable AASR for a target at scene center, a higher PRF may be required to prevent degradation at the edge of the scene.

Compare the spotlight result of Figure 6.7 with the stripmap result of Figure 6.2. In spotlight, AASR can be improved by forming a smaller image. In stripmap, AASR can be improved by forming an image with coarser cross-range resolution. Similarly, spotlight AASR is largely independent of cross-range resolution/integration angle, and stripmap AASR does not vary with cross range in the final image.

6.3.5 Multichannel Spotlight AASR: Uniform Sampling

Just as in stripmap, a multichannel spotlight synthetic aperture imaging system must use a specific PRF to obtain uniformly-spaced along-track samples. This PRF depends on the platform velocity and the size of the receive elements. A larger transmit aperture can be used to allow finer along-track sampling, but this comes at the cost of a smaller image size.

Under this uniform-sampling constraint, the AASR can be plotted as a function of target offset as in Figure 6.8. Note that the maximum image size corresponding to each curve is different because the transmit beamwidth differs from curve to curve.

Figure 6.6: Scene center AASR for spotlight synthetic aperture imaging.

Figure 6.7: AASR for spotlight synthetic aperture imaging as a function of angular offset from scene center.

Figure 6.8: AASR for spotlight synthetic aperture imaging as a function of angular offset from scene center with $D_{RX}/2$ along-track sampling.

6.3.6 Nonuniform Sampling and Multichannel Stripmap AASR

The multichannel stripmap case is now reconsidered but without the requirement that the PRF is carefully chosen to ensure that the system provides uniform along-track sampling. To calculate AASR, the recurrent sampling approach summarized by (6.14) is employed.

The results are presented in Figure 6.9. In Figure 6.9a, the transmit aperture is assumed to be matched to the size of each aperture used on receive. In the single-channel case ($N = 1$), AASR always improves if a higher PRF is employed (at the cost of ACR). For the true multichannel case ($N > 1$), increasing the PRF to beyond the uniform-sampling case ($\eta = 1$) tends to lead to degradation in AASR. The expected AASR reduction from finer along-track sampling is overwhelmed by the impact of nonuniform sampling. The impact of nonuniform sampling is more significant as the number of channels increases.

The effect of using a larger transmit aperture is illustrated in Figure 6.9b. As expected, using a larger transmit aperture provides an improvement in AASR, but it comes at the cost of degraded cross-range resolution. However, note that this AASR benefit rapidly declines

as the sampling becomes increasingly nonuniform when a large number of channels are employed.
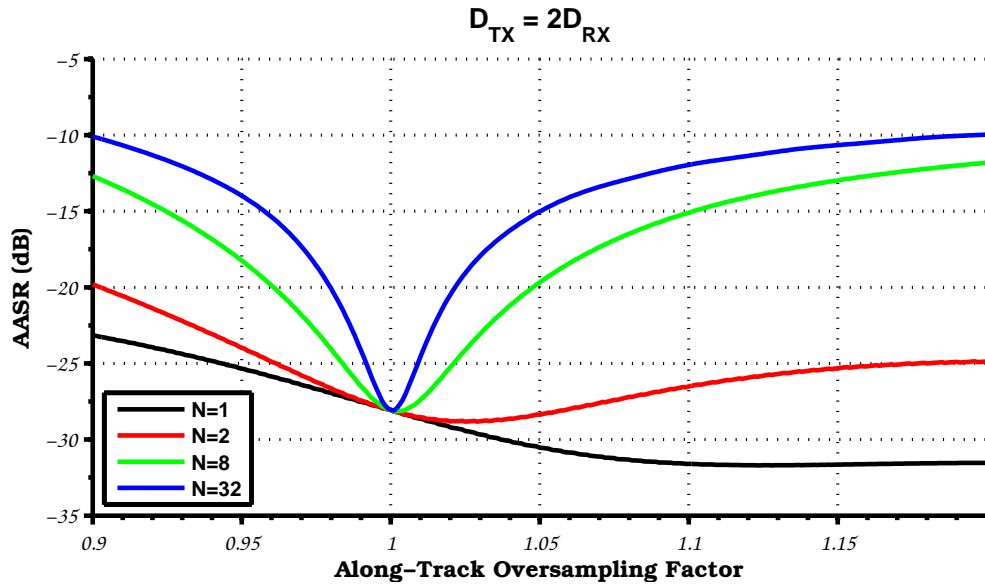

(a) Transmit aperture is matched to receive element size

(b) Transmit aperture is larger than receive element size
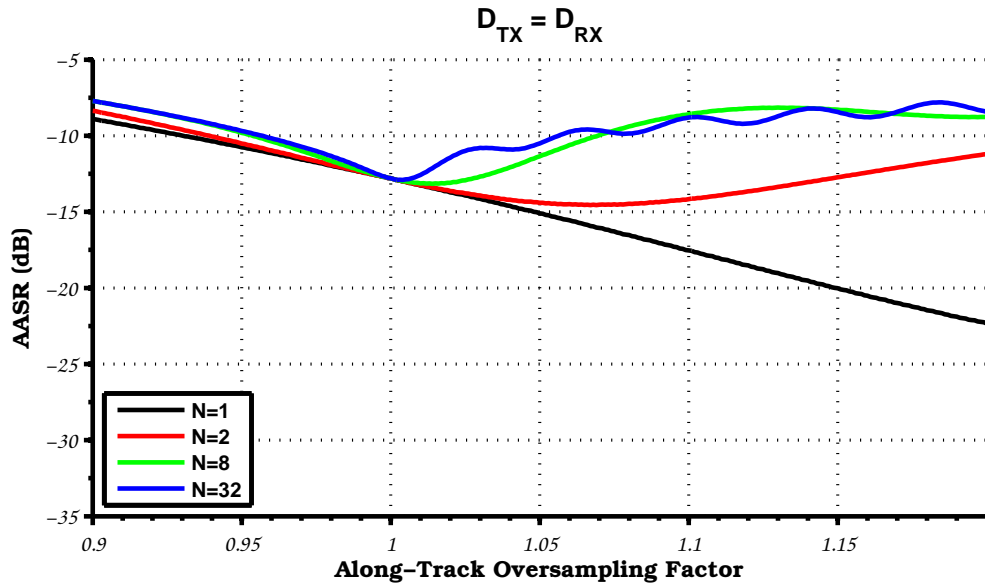
Figure 6.9: Multichannel stripmap AASR as a function of along-track sampling non-uniformity and number of channels.

6.4 Summary

This chapter presented the development of a rigorous method for calculating the impact of aliasing on synthetic aperture images. This was based on the concept of sampling a WSS random process that was developed in Chapter 5, which included the case of the recurrent sampling that can occur in multichannel SAR/SAS systems. Results were presented for both stripmap and spotlight synthetic aperture imaging including the multichannel case.

CHAPTER 7

CONCLUSIONS

This dissertation analyzed the benefits of using multiple-input, multiple-output (MIMO) techniques with synthetic aperture imaging systems. By transmitting multiple, independent waveforms, a MIMO system can take advantage of spatial diversity on transmit as well as receive. By exploiting these transmit degrees of freedom, a MIMO synthetic aperture radar (SAR) or synthetic aperture sonar (SAS) is able to improve its area-coverage rate and/or image quality by providing improved along-track sampling without an increase in pulse repetition frequency (PRF).

A significant contribution of this dissertation is the development of a comprehensive theory of coherent MIMO radar that is based on an array-processing approach. While the MIMO virtual array is an interesting concept to convey the basic motivation of MIMO radar, the beamforming approach developed in Chapter 2 provides a method for understanding the behavior of a MIMO-based radar antenna. This is critical for assessing the appropriateness of MIMO for a particular radar application.

A MIMO SAR/SAS is able to take advantage of multiple transmit phase centers for each pulse by simultaneously transmitting independent waveforms as discussed in Chapter 3. However, the primary challenge of MIMO synthetic aperture imaging is that the cross-correlation among these waveforms will increase the range sidelobes and degrade image quality by lowering contrast. This is quantified by the integrated sidelobe ratio (ISR), which is increased by transmitting additional waveforms.

While fundamental limits exist that prevent the autocorrelation sidelobes and the cross-correlation energy of a suite of waveforms from being made arbitrarily low, Chapter 4 presented a signal processing approach based on mismatched filtering seeks to improve ISR. This technique extends the concept of a minimum ISR filter that has been developed for a single waveform to the case of multiple waveforms. It is based on the observation

that the MIMO range response varies as a function of angle, which motivates the joint application of range compression and beamforming rather than the traditional approach of performing them in a separable fashion. By tailoring the filter to returns from a particular angle, acceptable ISR may be obtained that may enable MIMO SAR/SAS.

While radar systems have historically used multiple channels on receive, the novelty of a MIMO radar is essentially to operate multiple channels on transmit as well as using a number of independent waveforms. Consequently, a MIMO SAR/SAS is an extension of standard multichannel systems to now include multiple transmit elements. Especially in multichannel SAS systems, the signal processing is based on the assumption that the PRF is carefully controlled to ensure uniformly-spaced along-track samples. If this is not the case, then the sample spacing may be nonuniform. This dissertation presented results that quantify the impact of this nonuniform sampling in multichannel synthetic aperture imaging.

The analysis of nonuniform sampling in multichannel SAR/SAS was based on a rigorous analysis of aliasing in the sampling of a wide-sense stationary (WSS) random process. Chapter 5 described the concept of recurrent sampling, which may occur in multichannel SAR/SAS, and developed a method for calculating the impact of aliasing in this special case of nonuniform sampling. In Chapter 6, these techniques were applied to calculate the along-track ambiguity-to-signal ratio (AASR) to quantify the impact of ambiguities on image contrast. This included the case where methods were used to reconstruct uniform samples from the recurrent samples that were collected. Future work could explore alternative reconstruction methods. For instance, one can set up the task of reconstruction from sparse samples as a linear inverse problem and calculate the condition number (and related quantities) of the resulting transformation matrix (see [Ghaoui, 2002] and p. 83 of [Bertero and Boccacci, 1998].

Remaining future work includes further controlling the range sidelobes in MIMO radar

as well as dealing with the potential for nonuniform along-track sampling in multichannel SAR/SAS. While the present work focused on extending the standard matched filter processing to mismatched filtering to improve the ISR, indirect solutions based on sparse-reconstruction techniques may also be successful [Potter et al., 2010; Çetin and Lanterman, 2005]. Similarly, a filtering technique based on a continuous-time development was analyzed for reconstruction from nonuniform samples. This could be extended to a more practical discrete-time derivation that could exploit advanced reconstruction techniques.

The presented results focused on SAR or SAS systems for which the "real aperture" array elements were equally spaced. Future work could explore the potential advantages of designing systems with irregular spacing. This may involve adapting ideas from the literature on the design of nonuniform arrays [Oraizi and Fallanpour, 2008; Ridwan et al., 2011]. Similarly, it was assumed that distance traveled between pulses was equal. The results of this dissertation could be expanded to explore cases of staggered PRFs and other nonuniform slow-time sampling schemes.

APPENDIX A

WIDE-SENSE STATIONARY RANDOM PROCESSES

This appendix reviews results regarding wide-sense stationary (WSS) random processes, which are used some of the derivations presented in Chapter 5. Many of these facts can be found in standard references [Papoulis, 1977b]. Notation is established for dealing with WSS random processes as well as their autocorrelation functions and power spectral densities. A standard proof of the Wiener-Khinchin theorem is provided, which is extended to illustrate that the average power spectrum is uncorrelated from frequency to frequency as stated by (A.15). Also, while many references contain the result that the power spectral density of the output of a linear system is related to its input as in (A.19), a more obscure result related to the cross spectral density of the output of two linear systems is included as (A.20). Notation is also established for dealing with discrete random processes.

A.1   Preliminaries

A random process, $x(t)$, is said to be **wide-sense stationary** (WSS) if

1. The mean function of $x(t)$ is a constant,

$$\mu_x(t) \triangleq \mathrm{E}[x(t)] = \mu \text{ for all } t \tag{A.1}$$

2. The autocorrelation function, $R_x(t_1, t_2)$, only depends on the lag $t_2 - t_1$,

$$R_x(t_1, t_2) \triangleq \mathrm{E}[x(t_1) x^*(t_2)] = R_x(0, t_2 - t_1) \text{ for all } \tau \tag{A.2}$$

3. Two random processes, $x(t)$ and $y(t)$, are said to be jointly WSS if their cross-correlation function, $R_{xy}(t_1, t_2)$, only depends on the lag, $t_2 - t_1$,

$$R_{xy}(t_1, t_2) \triangleq \mathrm{E}[x(t_1) y^*(t_2)] = R_{xy}(0, t_2 - t_1) \text{ for all } \tau \tag{A.3}$$

124

A.2 The Autocorrelation Function of a WSS Random Process

For WSS processes, write the autocorrelation and cross-correlation functions in terms of the lag τ as

$$R_x(\tau) \triangleq R_x(t, t - \tau) = \mathrm{E}\left[x(t)\, x^*(t - \tau)\right] \tag{A.4}$$

$$R_{xy}(\tau) \triangleq R_{xy}(t, t - \tau) = \mathrm{E}\left[x(t)\, y^*(t - \tau)\right] \tag{A.5}$$

because these do not depend on the value of t. For the cross-correlation function to depend only on lag requires that the two random processes be jointly WSS, which is a stronger requirement than requiring that each random process be WSS independently.

The following are interesting properties of the autocorrelation function of a WSS random process. Let $X(t)$ be a WSS random process with autocorrelation function $R(\tau)$.

1. The zero-lag value of the autocorrelation function is the mean-squared value of the random process,

$$\mathrm{E}\left[|X(t)|^2\right] = R(0). \tag{A.6}$$

2. The autocorrelation function is conjugate symmetric,

$$R(\tau) = R^*(-\tau). \tag{A.7}$$

3. The autocorrelation function is maximized at $\tau = 0$,

$$R(0) \geq |R(\tau)| \tag{A.8}$$

for all τ.

A.3 The Power Spectral Density (PSD) of a WSS Random Process

The autocorrelation function characterizes the behavior of a random process in the time domain, but the spectral content of the process may also be of interest. In general, a realization of a random process, $x(t)$, does not have a well-defined Fourier transform. Instead,

begin by considering the Fourier transform of the random process that has been truncated to an interval of length T,

$$X_T(\omega) \triangleq \int_{-T/2}^{T/2} x(t) e^{-i\omega t} \, dt. \tag{A.9}$$

The power spectral density (PSD) characterizes the average behavior of the spectrum in (A.9) as the observation time, T, becomes arbitrarily long,

$$P_x(\omega) \triangleq \lim_{T \to \infty} \frac{1}{T} \mathrm{E}\left[|X_T(\omega)|^2\right]. \tag{A.10}$$

The Wiener-Khinchin theorem states that the limit in (A.10) converges, so the PSD exists. It also provides the important result that the autocorrelation function and the PSD are related by the Fourier transform. If the autocorrelation function, $R_x(\tau)$, is absolutely integrable, then the PSD, $P_x(\omega)$, defined in (A.10) exists. Further, the autocorrelation function and the power spectral density of a WSS random process are Fourier duals,

$$P_x(\omega) = \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau} \, d\tau. \tag{A.11}$$

A standard proof of the result of (A.11) is provided. The average periodogram with observation interval, T, of (A.9) can be written as

$$
\begin{aligned}
\mathrm{E}\left[|X_T(\omega)|^2\right] &= \mathrm{E}\left[\left(\int_{-T/2}^{T/2} x(\tau) e^{-i\omega\tau} \, d\tau\right)\left(\int_{-T/2}^{T/2} x(t) e^{-i\omega t} \, dt\right)^*\right] \\
&= \int_{-T/2}^{T/2} \int_{-T/2}^{T/2} R_x(\tau - t) e^{-i\omega(\tau - t)} \, d\tau \, dt \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_x(\tau - t) e^{-i\omega(\tau - t)} \mathrm{rect}\left(\frac{\tau}{T}\right) \mathrm{rect}\left(\frac{t}{T}\right) \, d\tau \, dt.
\end{aligned}
$$

Apply two changes of variable: $\tau \mapsto t + \tau$ and $t \mapsto t - \tau$, which yields

$$\mathrm{E}\left[|X_T(\omega)|^2\right] = \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau} \underbrace{\int_{-\infty}^{\infty} \mathrm{rect}\left(\frac{t}{T}\right) \mathrm{rect}\left(\frac{t - \tau}{T}\right) \, dt}_{=\mathrm{rect}\left(\frac{t}{T}\right) * \mathrm{rect}\left(\frac{t}{T}\right)} \, d\tau.$$

The convolution of two unit rectangle functions is a unit triangle function,

$$\Lambda(t) \triangleq \begin{cases} 1 - |t|, & \text{for } |t| < 1 \\ 0, & \text{otherwise.} \end{cases}$$

For two identical scaled rectangles,

$$\Lambda\left(\frac{\tau}{T}\right) = \frac{1}{T} \int_{-\infty}^{\infty} \mathrm{rect}\left(\frac{t}{T}\right) \mathrm{rect}\left(\frac{\tau - t}{T}\right) dt,$$

Using this result,

$$\frac{1}{T} \mathrm{E}\left[|X_T(\omega)|^2\right] = \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau} \Lambda\left(\frac{\tau}{T}\right) d\tau.$$

By definition, the PSD is the limit of this as $T \to \infty$,

$$P_x(\omega) = \lim_{T \to \infty} \int_{-\infty}^{\infty} R_x(\tau) e^{-i\omega\tau} \Lambda\left(\frac{\tau}{T}\right) d\tau. \tag{A.12}$$

Lebesgue's dominated convergence theorem provides a set of conditions under which the limit and integral operations can be interchanged. Define $f_T(\tau)$ to be the integrand for an interval, T. To satisfy the conditions of the dominated convergence theorem, it must be shown that there is some function, $f(\tau)$, such that $f_T(\tau) \to f(\tau)$ pointwise as $T \to \infty$. It also requires that an integrable function, g, such that $|f_T| \leq g$ for all $T > 0$. First, observe that

$$\lim_{T \to \infty} f_T(\tau) = R_x(\tau) e^{-i\omega\tau}.$$

The dominating function, g, can be chosen to be $g(\tau) \triangleq |R_x(\tau)|$ because

$$|f_T(x)| = \left| R_x(\tau) \Lambda\left(\frac{\tau}{T}\right) \right| \leq |R_x(\tau)|.$$

Note that g is integrable by definition because it was required that the autocorrelation function, $R_x(\tau)$, be absolutely integrable.

Interchanging the limit and integral and evaluating the limit in (A.12) yields the desired result, which proves that the autocorrelation function and PSD of a random process are related by Fourier transform as stated in (A.11).

Because the PSD is the Fourier transform of the autocorrelation function, the autocorrelation can be recovered from the PSD using the Fourier inversion theorem. The autocorrelation function can be computed from the power spectral density by

$$R_x(\tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} P_x(\omega) e^{i\omega\tau} d\omega. \tag{A.13}$$

Just as the autocorrelation function and PSD are Fourier duals, the cross spectral density (CSD) can be defined as the Fourier transform of the cross-correlation function, i.e., the CSD of two random processes, $X(t)$ and $Y(t)$, is related to their cross-correlation function, $R_{xy}(\tau)$, by

$$P_{xy}(\omega) \triangleq \int_{-\infty}^{\infty} R_{xy}(\tau) e^{-i\omega\tau} d\tau. \tag{A.14}$$

The result of the Wiener-Khinchin theorem, given by (A.11), provides the interpretation the Fourier transform of the autocorrelation function is the average power spectrum of the random process. It is now demonstrated that the average power spectrum is uncorrelated from frequency to frequency, i.e.,

$$\lim_{T\to\infty} \frac{1}{T} E\left[X(\omega) X^*(\omega_0)\right] = \begin{cases} P_x(\omega), & \text{if } \omega = \omega_0 \\ 0, & \text{if } \omega \neq \omega_0. \end{cases} \tag{A.15}$$

As mentioned, the result for $\omega = \omega_0$ follows from the Wiener-Khinchin formula. Now, suppose that $\omega \neq \omega_0$. Proceeding in a similar manner to that proof, it is found that

$$E\left[X_T(\omega) X_T(\omega_0)\right] = \int_{-\infty}^{\infty} R_x(t) e^{-i\omega_0 t} \underbrace{\int_{-\infty}^{\infty} e^{-i(\omega-\omega_0)t_0} \text{rect}\left(\frac{t_0}{T}\right) \text{rect}\left(\frac{t_0 - t}{T}\right) dt_0}_{f(t)} \, dt.$$

Rewrite the function, $f(t)$, as

$$f(t) = \int_{-T/2}^{T/2} e^{-i\xi t_0} \text{rect}\left(\frac{t_0 - t}{T}\right) dt_0,$$

where $\xi \triangleq \omega - \omega_0 \neq 0$. Note that $f(t) = 0$ for $|t| \geq T$. Otherwise, it is simply the integral of $e^{-i\xi t_0}$ over some interval of t_0. It can be shown that

$$f(t) = \frac{1}{\xi/2} e^{-i(\xi/2)t} \sin\left(\frac{D\xi}{2} \Lambda\left(\frac{t}{T}\right)\right).$$

Using this, it can be written that

$$\begin{aligned} \lim_{T\to\infty} \frac{1}{T} E\left[X_T(\omega) X_T(\omega_0)\right] &= \lim_{T\to\infty} \frac{1}{T} \frac{1}{\xi/2} \int_{-\infty}^{\infty} R_x(t) e^{-i\omega_0 t} e^{-i(\xi/2)t} \sin\left(\frac{D\xi}{2} \Lambda\left(\frac{t}{T}\right)\right) dt \\ &= \frac{1}{\xi/2} \int_{-\infty}^{\infty} R_x(t) e^{-i\omega_0 t} e^{-i(\xi/2)t} \left[\lim_{T\to\infty} \frac{1}{T} \sin\left(\frac{D\xi}{2} \Lambda\left(\frac{t}{T}\right)\right)\right] dt \\ &= 0. \end{aligned}$$

The argument for interchanging the order of the integral and the limit operations is justified in a similar manner to the proof of the Wiener-Khinchin theorem. The fact that the limit goes to zero follows from the fact that $\sin(ax)/x \to 0$ as $x \to \infty$.

A.4 Random Processes and Linear Systems

Recall that a linear time-invariant (LTI) system is completely characterized by its impulse response, $h(t)$, or, equivalently, its frequency response, $H(\omega)$. If the (deterministic) input signal is $x(t)$ with a spectrum, $X(\omega)$, then the output of the LTI system, $y(t)$, is the convolution of the input with the impulse response of the system,

$$y(t) \triangleq \int_{-\infty}^{\infty} x(\tau) h(t-\tau)\ d\tau, \tag{A.16}$$

and the spectrum of the output, $Y(\omega)$, is the product of the input spectrum and the frequency response of the system,

$$Y(\omega) = H(\omega) X(\omega). \tag{A.17}$$

This concept extends to the case where the input to the LTI system is a random process. Let $X(t)$ be a random process with autocorrelation function, $R_x(t_1, t_2)$. If the random process, $X(t)$, is the input to an LTI system with impulse response, $h(t)$, then the autocorrelation function of the output random process, $Y(t)$, has the autocorrelation function

$$R_y(t_1, t_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_x(t_1 - \tau_1, t_2 - \tau_2) h(\tau_1) h^*(\tau_2)\ d\tau_1 d\tau_2. \tag{A.18}$$

If the input is WSS, then the output is as well. Let $X(t)$ be a WSS process with mean function, μ_x, and autocorrelation function, $R_x(\tau)$. Consider a linear system with an impulse response, $h(t)$. If the random process $X(t)$ is the input to this system, then the output random process, $Y(t)$, is WSS with mean, μ_y, and autocorrelation function, $R_y(t)$, given by

$$\mu_y = \mu_x \int_{-\infty}^{\infty} h(\tau)\ d\tau,\ \text{and}$$

$$R_y(\tau) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} R_x(\tau - (\tau_1 - \tau_2)) h(\tau_1) h^*(\tau_2)\ d\tau_1\ d\tau_2$$

$$= h^*(-\tau) * h(\tau) * r_x(\tau).$$

The corresponding power spectral density of the output is

$$P_y(\omega) = |H(\omega)|^2 P_x(\omega), \tag{A.19}$$

where $H(\omega)$ is the frequency response of the system.

Suppose that there are two random processes, and that each is passed through its own LTI system. The CSD of the outputs can be written in terms of the CSD of the input processes and the system frequency responses. Let $X_1(t)$ and $X_2(t)$ be jointly WSS random processes with cross spectral density, $P_{x_1 x_2}(\omega)$. Let $X_1(t)$ be the input to a system with impulse response $h_1(t)$, and let $X_2(t)$ be the input to a system with impulse response $h_2(t)$. The cross spectral density of the output random processes, $Y_1(t)$ and $Y_2(t)$, is

$$P_{y_1 y_2}(\omega) = P_{x_1 x_2}(\omega) H_1^*(\omega) H_2(\omega), \tag{A.20}$$

where $H_1(\omega)$ and $H_2(\omega)$ are the frequency responses of the LTI systems. Similarly, the cross-correlation of the outputs is related to the cross-correlation function of the input random processes and the impulse responses of the systems,

$$R_{y_1 y_2}(\tau) = h_1^*(-\tau) * h_2(\tau) * R_{x_1 x_2}(\tau). \tag{A.21}$$

A.5   Linear Combinations of WSS Random Processes

Let $X_1(t), \ldots, X_N(t)$ be jointly WSS random processes. Define a new random process, $Y(t)$, that is a linear combination of these random processes by

$$Y(t) \triangleq \sum_{n=1}^{N} w_n^* X_n(t) \tag{A.22}$$

for complex weights $\{w_n\}$. This may also be written using vector notation by defining the weight vector, $\mathbf{w}$, and the vector of random processes, $\underline{X}(t)$, as

$$Y(t) = \mathbf{w}^H \underline{X}(t), \tag{A.23}$$

where

$$\mathbf{w} \triangleq \begin{pmatrix} w_1 \\ \vdots \\ w_N \end{pmatrix}, \underline{X}(t) \triangleq \begin{pmatrix} X_1(t) \\ \vdots \\ X_N(t) \end{pmatrix}. \tag{A.24}$$

The autocorrelation function of $Y(t)$ is then

$$R_y(\tau) = \mathbf{w}^{\mathrm{H}} \mathbf{R}_x(\tau) \mathbf{w} \tag{A.25}$$

where $\mathbf{R}_x$ is the covariance matrix of the vector, $\underline{X}$, given by

$$\mathbf{R}_x(\tau) \triangleq \mathrm{E}\left[\underline{X}(t)\,\underline{X}(t-\tau)^{\mathrm{H}}\right]. \tag{A.26}$$

A.6 Discrete Random Processes

The above analysis can be repeated for the case of discrete random processes. Attention will be restricted to discrete random processes that are WSS. Let $x[n]$ be a WSS discrete random process. It is characterized by the following quantities:

$$\textbf{Mean Sequence: } \mu[n] \triangleq \mathrm{E}[x[n]] \tag{A.27}$$

$$\textbf{Variance Sequence: } \sigma^2[n] \triangleq \mathrm{E}\left[|x[n]-\mu[n]|^2\right] \tag{A.28}$$

$$\textbf{Autocorrelation Sequence: } R[n] \triangleq \mathrm{E}[x[m]\,x^*[m-n]] \tag{A.29}$$

$$\textbf{Autocovariance Sequence: } C[n] \triangleq \mathrm{E}[(x[m]-\mu[m])\,(x[m-n]-\mu[m-n])^*] \tag{A.30}$$

The PSD of a discrete WSS random process is the discrete-time Fourier transform of its autocorrelation sequence,

$$P_x(\omega) \triangleq \sum_{n=-\infty}^{\infty} R_x[n]\,e^{-i\omega n}. \tag{A.31}$$

Because it is the result of applying a discrete-time Fourier transform, the PSD of a discrete random process is periodic with period 2π. If it is assumed that the discrete-time process is generated by uniformly sampling a continuous-time process at a rate $1/T$, the PSD may be alternatively defined by instead scaling ω by $1/T$ so that it becomes $2\pi/T$-periodic.

The autocorrelation sequence can be recovered from the PSD by the inverse discrete-time Fourier transform,

$$R_x[n] = \frac{1}{2\pi} \int_{-\pi}^{\pi} P_x(\omega)\,e^{i\omega n}\,d\omega. \tag{A.32}$$

Because $P_x(\omega)$ is 2π-periodic, the integral may be computed over any interval of length 2π.

A.7 Discrete WSS Random Processes and LTI Systems

Consider the case where a discrete WSS random process, $x[n]$, is applied to an LTI system with impulse response, $h[n]$. The output sequence, $y[n]$ is the convolution of the input and the impulse response,

$$y[n] \triangleq (h * x)[n] = \sum_{k=-\infty}^{\infty} h[k] x[n-k].$$ (A.33)

As before, the output is also a WSS random process. The output autocorrelation function is found by convolving the input autocorrelation with the impulse response, and the output PSD is found by multiplying the input PSD with the frequency response.

Let $x[n]$ be a discrete WSS random process with autocorrelation sequence, $R_x[n]$. If the random process, $x[n]$, is the input to an LTI system with impulse response, $h[n]$, then the output random process, $y[n]$, is WSS with autocorrelation sequence

$$R_y[n] = \sum_{k=-\infty}^{\infty} \sum_{\ell=-\infty}^{\infty} R_x[n-(k-\ell)] h[k] h^*[\ell]$$

$$= h^*[-n] * h[n] * R_x[n].$$ (A.34)

The PSD of the output is

$$P_y(\omega) = |H(\omega)|^2 P_x(\omega)$$ (A.35)

where $H(\omega)$ is the discrete-time Fourier transform of the impulse response,

$$H(\omega) \triangleq \sum_{n=-\infty}^{\infty} h[n] e^{-i\omega n}.$$ (A.36)

APPENDIX B

SPECTRUM OF A RECURRENT SAMPLING SIGNAL

This appendix derives the spectrum of the sampling signal that corresponds to recurrent sampling. This result is used in Chapter 5 to derive the spectrum of a recurrently sampled signal. Two special cases are also presented.

B.1 The Recurrent Sampling Signal and Its Spectrum

The sampling signal for recurrent sampling was given in Table 5.1 and is reproduced below:

$$s(t) \triangleq \frac{T}{N} \sum_{n=1}^{N} \sum_{m=-\infty}^{\infty} \delta(t - (T_n + mT)), \tag{B.1}$$

where T is the recurrence interval and $T_1, \ldots, T_N$ are the offsets of the N samples within a recurrence. An impulse train can be written as a Fourier series,

$$\sum_{k=-\infty}^{\infty} \delta(t - kT) = \frac{1}{T} \sum_{k=-\infty}^{\infty} \exp\left\{i\frac{2\pi k}{T}t\right\}. \tag{B.2}$$

This allows the sampling signal of (B.1) to be written as

$$\begin{aligned}
s(t) &= \frac{1}{N} \sum_{n=1}^{N} \sum_{m=-\infty}^{\infty} \exp\left\{-i\frac{2\pi m}{T}(t - T_n)\right\} \\
&= \sum_{m=-\infty}^{\infty} \left(\frac{1}{N} \sum_{n=1}^{N} \exp\left\{i2\pi\frac{T_n}{T}m\right\}\right) \exp\left\{-i\frac{2\pi m}{T}t\right\}.
\end{aligned} \tag{B.3}$$

This is a linear combination of complex sinusoids, so the spectrum is a linear combination of shifted Dirac delta functions,

$$S(\omega) = 2\pi \sum_{m=-\infty}^{\infty} \alpha_m \delta\left(\omega + \frac{2\pi m}{T}\right), \tag{B.4}$$

where the weight of each delta function is given by (5.9).

B.2 Special Case: Uniformly-Spaced Samples within a Recurrence

Consider the special case where the samples within each recurrence are uniformly spaced by T_0, i.e., $T_n \triangleq (n-1)T_0$ for $n = 1, 2, \ldots, N-1$. In this case, the weights of (5.9) can

be written as

$$\alpha_m = \frac{1}{N} \sum_{n=0}^{N-1} \exp\left\{i\left(2\pi\frac{T_0}{T}m\right)n\right\}. \tag{B.5}$$

This sum is related to the Dirichlet function, $D_N(x)$, by

$$\alpha_m = D_N\left(2\pi\frac{T_0}{T}m\right)\exp\left\{i\pi(N-1)\frac{T_0}{T}m\right\}, \tag{B.6}$$

where the Dirichlet function is given by

$$D_N(x) \triangleq \begin{cases} (-1)^{k(N-1)}, & \text{for } x = 2\pi k \text{ for an integer, } k \\ \frac{\sin(Nx/2)}{N\sin(x/2)}, & \text{otherwise.} \end{cases} \tag{B.7}$$

B.3   Special Case: Uniformly-Spaced Samples

Now, consider a related case where the recurrent sampling reduces to uniform sampling, i.e., $T_n \triangleq (n-1)/NT$. This is related to the above case with $T_0 = T/N$. In this case,

$$\alpha_m = D_N\left(2\pi\frac{m}{N}\right)\exp\left\{-i\pi\frac{N-1}{N}m\right\}. \tag{B.8}$$

The weights, α_m, are nonzero only for $m = 0, \pm N, \pm, 2N$ where $\alpha_m = \pm1$. As expected, the sampling spectrum reduces to the uniform sampling case.

REFERENCES

Ackroyd, M. H. and Ghani, F. (1973). Optimum mismatched filters for sidelobe suppression. *IEEE Transactions on Aerospace and Electronic Systems*, 9(2):214–218.

Alamouti, S. M. (1998). A simple transmit diversity technique for wireless communications. *IEEE Journal on Selected Areas in Communications*, 16(8):1451–1458.

Applebaum, S. (1976). Adaptive arrays. *IEEE Transactions on Antennas and Propagation*, 24(5):585–598.

Bayma, R. W. and McInnes, P. A. (1975). Aperture size and ambiguity constraints for synthetic aperture radar. In *Proceedings of the 1975 IEEE International Radar Conference*, pages 499–504, Arlington, VA.

Bekkerman, I. and Tabrikian, J. (2006). Target detection and localization using MIMO radars and sonars. *IEEE Transactions on Signal Processing*, 54(10):3873–3883.

Bertero, M. and Boccacci, P. (1998). *Introduction to Inverse Problems in Imaging*. CRC Press.

Biglieri, E., Calderbank, R., Constantinides, A., Goldsmith, A., Paulraj, A., and Poor, H. V. (2007). *MIMO Wireless Communications*. Cambridge University Press.

Bliss, D. W. and Forsythe, K. W. (2003). MIMO radar and imaging: Degrees of freedom and resolution. In *Conference Record of the Thirty-Seventh Asilomar Conference on Signals, Systems, and Computers*, pages 54–59, Pacific Grove, CA.

Bliss, D. W., Forsythe, K. W., Hero III, A. O., and Yegulalp, A. F. (2002). Environmental issues for MIMO capacity. *IEEE Transactions on Signal Processing*, 50(9):2128–2142.

Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Brown, J. L. (1981). Multi-channel sampling of low-pass signals. *IEEE Transactions on Circuits and Systems*, 28(2):101–106.

Brown, W. M. (1967). Synthetic aperture radar. *IEEE Transactions on Aerospace and Electronic Systems*, 3(2):217–229.

Cann, A. J. (2002). Range gate straddling loss and joint probability with partial correlation. *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):1054–1058.

Carrara, W. G., Goodman, R. S., and Majewski, R. M. (1995). *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*. Artech House, Norwood, MA.

Çetin, M. and Lanterman, A. D. (2005). Region-enhanced passive radar imaging. *IEE Proceedings - Radar, Sonar and Navigation*, 152(3):185–194.

Chernyak, V. S. (1998). *Fundamentals of Multisite Radar Systems*. CRC Press.

Cheston, T. C. (1968). Phased arrays for radars. *IEEE Spectrum*, 5(11):102–111.

Correll, B. (2010). Efficient spotlight SAR MIMO linear collection configurations. *IEEE Journal of Selected Topics in Signal Processing*, 4(1):33–39.

Curlander, J. C. and McDonough, R. N. (1991). *Synthetic Aperture Radar: Systems and Signal Processing*. Wiley Series in Remote Sensing and Image Processing. Wiley.

Currie, A. and Brown, M. A. (1992). Wide-swath SAR. *IEE Proceedings F: Radar and Signal Processing*, 139(2):122–135.

Cutrona, L. J. (1975). Comparison of sonar system performance achievable using synthetic-aperture techniques with the performance achievable by more conventional means. *Journal of the Acoustical Society of America*, 58(2):336–348.

Daum, F. and Huang, J. (2009). MIMO radar: Snake oil or good idea? *IEEE Aerospace and Electronic Systems Magazine*, 24(5):8–12.

Davis, M. and Cook, D. (2011). MIMOSAS: Improving SAS performance with transmit diversity. In *Proceedings of the 4th Underwater Acoustic Measurements Conference*, Kos, Greece.

Davis, M. S. (2012). MIMO Radar. In Melvin, W. L. and Scheer, J. A., editors, *Principles of Modern Radar: Advanced Techniques*, pages 119–145. SciTech Publishing, Edison, NJ.

Davis, M. S. and Cook, D. A. (2014). Quantifying the impact of range and azimuth ambiguities on multichannel SAR and SAS image quality. In *Proceedings of the 10th European Conference on Synthetic Aperture Radar*.

Davis, M. S. and Lanterman, A. D. (2015a). Aliasing in recurrently sampled signals with an application to synthetic aperture imaging. *IEEE Transactions on Signal Processing*. accepted January 2015.

Davis, M. S. and Lanterman, A. D. (2015b). Minimum integrated sidelobe ratio filters for MIMO radar. *IEEE Transactions on Aerospace and Electronic Systems Magazine*, 51(1). to appear in.

Davis, M. S., Showman, G., Pechner, D., and Dai, L. (2011). Improving SAR image quality and area coverage through MIMO operation. In *Proceedings of the 57th Annual Meeting of the MSS Tri-Service Radar Symposium*, Monterey, CA.

Davis, M. S., Showman, G. A., and Lanterman, A. D. (2014). Coherent MIMO radar: The phased array and orthogonal waveforms. *IEEE Aerospace and Electronic Systems Magazine: Tutorials*, 29(8):76–90.

Dorey, J. and Garnier, G. (1989). RIAS, Radar à impulsion et antenna synthétique. *L'Onde Electrique*, 69(6):36–44.

Eldar, Y. C. and Oppenheim, A. V. (2000). Filterbank reconstruction of bandlimited signals from nonuniform and generalized samples. *IEEE Transactions on Signal Processing*, 48(10):2864–2875.

Ender, J. H. G. and Klare, J. (2009). System architectures and algorithms for radar imaging by MIMO-SAR. In *Proceedings of the 2009 IEEE Radar Conference*, Pasadena, CA.

Fishler, E., Haimovich, A., Blum, R. S., Cimini, L. J., J., Chizhik, D., and Valenzuela, R. A. (2006). Spatial diversity in radars-Models and detection performance. *IEEE Transactions on Signal Processing*, 54(3):823–838.

Forsythe, K. W. and Bliss, D. (2010). MIMO radar waveform constraints for GMTI. *IEEE Journal of Selected Topics in Signal Processing*, 4(1):21–32.

Foschini, G. J. and Gans, M. J. (1998). On limits of wireless communications in a fading environment when using multiple antennas. *Wireless Personal Communications*, 6:311–335.

Frazer, G. J., Abramovich, Y. I., and Johnson, B. A. (2009). Multiple-input multiple-output over-the-horizon radar: experimental results. *IET Radar, Sonar, and Navigation*, 3(4):290–303.

Freeman, A., Johnson, W., Huneycutt, B., Jordan, R., Hensley, S., Siqueira, P., and Curlander, J. (2000). The "myth" of the minimum SAR antenna area constraint. *IEEE Transactions on Geoscience and Remote Sensing*, 38(1):320–324.

Friedlander, B. (2012). Effects of model mismatch in MIMO radar. *IEEE Transactions on Signal Processing*, 60(4):2071–2076.

Fuhrmann, D. R. and San Antonio, G. (2008). Transmit beamforming for MIMO radar systems using signal cross-correlation. *IEEE Transactions on Aerospace and Electronic Systems*, 44(1):171–186.

Ghaoui, L. E. (2002). Inversion error, condition number, and approximate inverses of uncertain matrices. *Linear Algebra and its Applications*, pages 171–193.

Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. JHU Press.

Griep, K. R., Ritcey, J. A., and Burlingame, J. J. (1995). Poly-phase codes and optimal filters for multiple user ranging. *IEEE Transactions on Aerospace and Electronic Systems*, 31(2):752–767.

Hassanien, A. and Vorobyov, S. A. (2010). Phased-MIMO radar: A tradeoff between phased-array and MIMO radars. *IEEE Transactions on Signal Processing*, 58(6):3137–3151.

Hawkins, D. W. (1996). *Synthetic aperture imaging algorithms with applications to wideband sonar*. PhD thesis, University of Canterbury, Christchurch, New Zealand.

Higgins, J. (1976). A sampling theorem for irregularly spaced sample points. *IEEE Transactions on Information Theory*, 22(5):621–622.

Hoctor, R. T. and Kassam, S. A. (1990). The unifying role of the coarray in aperture synthesis for coherent and incoherent imaging. *Proceedings of the IEEE*, 78(4):735–752.

Hu, L., Liu, H., Feng, D., Jiu, B., Wang, X., and Wu, S. (2010). Optimal mismatched filter bank design for MIMO radar via convex optimization. In *Proceedings of the 2010 International Waveform Diversity and Design Conference*, pages 126–131, Niagra Falls, ON.

Hua, G. and Abeysekera, S. S. (2013). Receiver design for range and doppler sidelobe suppression using MIMO and phased-array radar. *IEEE Transactions on Signal Processing*, 61(6):1315–1326.

Johnson, D. H. (1982). The application of spectral estimation methods to bearing estimation problems. *Proceedings of the IEEE*, 70(9):1018–1028.

Johnson, D. H. and Dudgeon, D. E. (1993). *Array Signal Processing: Concepts and Techniques*. Prentice-Hall.

Kantor, J. and Davis, S. K. (2010). Airborne GMTI using MIMO techniques. In *Proceedings of the 2010 IEEE Radar Conference*, pages 1344–1349, Arlington, VA.

Kasami, T. (1966). Weight distribution formula for some class of cyclic codes. Technical Report R-285, University of Illinois.

Keel, B. and Baden, J. M. (2012). Advanced pulse compression waveform modulations and techniques. In Melvin, W. L. and Scheer, J. A., editors, *Principles of Modern Radar: Advanced Techniques*, pages 19–85. SciTech Publishing.

Kell, R. E. (1965). On the derivation of bistatic RCS from monostatic measurements. *Proceedings of the IEEE*, 53(8):983–988.

Kinsey, R. (1997). Phased array beam spoiling technique. In *Digest of the IEEE Antennas and Propagation Society International Symposium*, volume 2, pages 698–701, Montreal, QC.

Kock, W. E. (1972). Extending the maximum range of synthetic aperture (hologram) systems. *Proceedings of the IEEE*, 60(11):1459–1460.

Krieger, G. (2013). MIMO-SAR: Opportunities and pitfalls. *IEEE Transactions on Geoscience and Remote Sensing*, 52(5):2628–2645.

Krieger, G., Gebert, N., and Moreira, A. (2008). Multidimensional waveform encoding: A new digital beamforming technique for synthetic aperture radar remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 46(1):31–46.

Lee, A. J. (1978). Sampling theorems for nonstationary random processes. *Transactions of the American Mathematical Society*, 242:225–241.

Levanon, N. (2005). Cross-correlation of long binary signals with longer mismatched filters. *IEE Proceedings - Radar, Sonar and Navigation*, 152(6):377–382.

Li, J. and Stoica, P. (2007). MIMO radar with colocated antennas. *IEEE Signal Processing Magazine*, 24(5):106–114.

Li, J. and Stoica, P. (2008). *MIMO Radar Signal Processing*. Wiley-IEEE Press, Hoboken, NJ.

Li, J. and Stoica, P. (2010). The phased array is the maximum SNR active array. *IEEE Signal Processing Magazine*, 27(2):143–144.

Li, J., Stoica, P., and Zheng, X. (2008). Signal synthesis and receiver design for MIMO radar imaging. *IEEE Transactions on Signal Processing*, 56(8):3959–3968.

Lloyd, S. P. (1959). A sampling theorem for stationary (wide sense) stochastic processes. *Transactions of the American Mathematical Society*, 92(1):1–12.

Ma, C., Yeo, T. S., Tan, C. S., Qiang, Y., and Zhang, T. (2010). Receiver design for MIMO radar range sidelobes suppression. *IEEE Transactions on Signal Processing*, 58(10):5469–5474.

Mehlis, J. G. (1980). Synthetic aperture radar range-azimuth ambiguity design and constraints. In *Proceedings of the 1980 IEEE International Radar Conference*, pages 143–152, Arlington, VA.

Melvin, W. L. (2004). A STAP overview. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):19–35.

Messer, H., Singal, G., and Bialy, L. (1996). On the achievable DF accuracy of two kinds of active interferometers. *IEEE Transactions on Aerospace and Electronic Systems*, 32(3):1158–1164.

North, D. O. (1963). An analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems. *Proceedings of the IEEE*, 51(7):1016–1027.

Nyquist, H. (1928). Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47.

Oraizi, H. and Fallanpour, M. (2008). Nonuniformly spaced linear array design for the specified beamwidth/sidelobe level or specified directivity/sidelobe level with coupling considerations. *Progress in Electromagnetics Research M*, 4:185–209.

Papoulis, A. (1977a). Generalized sampling expansion. *IEEE Transactions on Circuits and Systems*, 24(11):652–654.

Papoulis, A. (1977b). *Signal Analysis*. Wiley Series in Remote Sensing and Image Processing. McGraw-Hill, Inc.

Potter, L. C., Ertin, E., Parker, J. T., and Çetin, M. (2010). Sparsity and compressed sensing in radar imaging. *Proceedings of the IEEE*, 98(6):1006–1020.

Rabideau, D. J. (2012). MIMO radar waveforms and cancellation ratio. *IEEE Transactions on Aerospace and Electronic Systems*, 48(2):1167–1178.

Rennich, P. K. (2009). Four-platform distributed MIMO radar measurements and imagery. In *Proceedings of the 2009 IEEE Radar Conference*, Pasadena, CA.

Ridwan, M., Abdo, M., and Jorswieck, E. (2011). Design of non-uniform antenna arrays using genetic algorithm. In *Proceedings of the 13th International Conference on Technology*, pages 422–427.

Rihaczek, A. W. and Golden, R. M. (1971). Range sidelobe suppression for Barker codes. *IEEE Transactions on Aerospace and Electronic Systems*, 7(6):1087–1092.

San Antonio, G., Fuhrmann, D. R., and Robey, F. C. (2007). MIMO radar ambiguity functions. *IEEE Journal of Selected Topics in Signal Processing*, 1(1):167–177.

Sarwate, D. V. (1979). Bounds on crosscorrelation and autocorrelation of sequences. *IEEE Transactions on Information Theory*, 25(6):720–724.

Sarwate, D. V. and Pursley, M. B. (1980). Crosscorrelation properties of pseudorandom and related sequences. *Proceedings of the IEEE*, 68(5):593–619.

Scheer, J. A. (2010). The radar range equation. In Richards, M. A., Scheer, J. A., and Holm, W. A., editors, *Principles of Modern Radar: Basic Principles*, chapter 2. SciTech Publishing, Edison, NJ.

Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21.

Showman, G. A. (2010). An overview of radar imaging. In Richards, M. A., Scheer, J. A., and Holm, W. A., editors, *Principles of Modern Radar: Basic Principles*, pages 835–891. SciTech Publishing.

Silver, S. (1949). *Microwave Antenna Theory and Design*. Radiation Laboratory Series. McGraw-Hill Book Company, Inc.

Southworth, G. C. (1930). Certain factors affecting the gain of directive antennas. *Proceedings of the Institute of Radio Engineers*, 18(9):1502–1536.

Steyskal, H., Schindler, J. K., Franchi, P., and Mailloux, R. J. (2003). Pattern synthesis for TechSat21 - a distributed space-based radar system. *IEEE Antennas and Propagation Magazine*, 45(4):19–25.

Turin, G. L. (1960). An introduction to matched filters. *IRE Transactions on Information Theory*, 6(3):311–329.

Veklerov, E. and Llacer, J. (1987). Stopping rule for the MLE algorithm based on statistical hypothesis testing. *IEEE Transactions on Medical Imaging*, 6(4):313–319.

Wang, H. and Cai, L. (1994). On adaptive spatial-temporal processing for airborne surveillance radar systems. *IEEE Transactions on Aerospace and Electronic Systems*, 30(3):660–670.

Wang, W. (2013). MIMO SAR imaging: Potential and challenges. *IEEE Aerospace and Electronic Systems Magazine*, 28(8):18–23.

Welch, L. (1974). Lower bounds on the maximum cross correlation of signals. *IEEE Transactions on Information Theory*, 20(3):397–399.

Whittaker, E. T. (1915). On the functions which are represented by the expansions of the interpolation theory. *Proceedings of the Royal Society of Edinburgh: Section A*, 35:181–194.

Willis, N. J. (2005). *Bistatic Radar*. SciTech Publishing.

Wu, Q., Xing, M., Bao, Z., and Shi, H. (2009). Wide swath, high range resolution imaging with MIMO-SAR. In *Proceedings of the 2009 IET International Radar Conference*, Guilin, China.

Yen, J. L. (1956). On nonuniform sampling of bandwidth-limited signals. *IRE Transactions on Circuit Theory*, 3(4):251–257.

Zakai, M. (1965). Band-limited functions and the sampling theorem. *Information and Control*, 8(2):143–158.

Zoraster, S. (1980). Minimum peak range sidelobe filters for binary phase-coded waveforms. *IEEE Transactions on Aerospace and Electronic Systems*, 16(1):112–115.

Zou, B., Dong, Z., and Liang, D. (2011). Design and performance analysis of orthogonal coding signal in MIMO-SAR. *Science China Information Sciences*, 54(8):1723–1737.