SAMPLE REDUCTION DURING SIGNAL DETECTION USING DIFFERENCE SETS AND ALMOST DIFFERENCE SETS

by

Wade Raymond Smith

Submitted in partial fulfillment of the requirements for the degree
Master of Engineering (Electronic Engineering)

in the

Department of Electrical, Electronic and Computer Engineering
Faculty of Engineering, Built Environment and Information Technology

UNIVERSITY OF PRETORIA

9 July 2019

SUMMARY

SAMPLE REDUCTION DURING SIGNAL DETECTION USING DIFFERENCE SETS AND ALMOST DIFFERENCE SETS

by

Wade Raymond Smith

Supervisor(s): Professor W. P. du Plessis
Department: Department of Electrical, Electronic and Computer Engineering
University: University of Pretoria
Degree: Master of Engineering (Electronic Engineering)
Keywords: Signal detection, non-uniform sampling, difference sets, almost difference sets, sub-Nyquist sampling.

Wide signal bandwidths typically require receivers performing signal detection to collect very large quantities of data. However, in applications with limited size, weight and power (SWAP) requirements, reducing the amount of data becomes important for proper operation. Most existing sample reduction approaches rely on reconstruction algorithms to compensate for the missing data, but these are often computationally complex. Therefore, in this work sample reduction without reconstruction is considered.

This work proposes an approach to discarding samples prior to detection using difference sets (DSs) and almost difference sets (ADSs) – exploiting their sidelobe and cyclic properties – to minimise the negative impact on detection performance. Included are mathematical analyses, simulations, and experiments with practical data evaluating the effects of this technique on the detection performance. This work demonstrates that while the lack of a reconstruction algorithm does introduce interference, this is reduced when using DSs and ADSs compared to when samples are discarded at random, and the use of these sets allows predictions about performance to be made beforehand using only the set parameters. Additionally, the proposed technique performs much faster than detection with reconstruction, while having a reasonable decrease in detection performance.

ACKNOWLEDGEMENTS

I would like to acknowledge the support of GEW Technologies in financing my studies.

I would also like to sincerely thank the following people for all of their help and support, without whom this work could never have been completed:

Professor Warren du Plessis: For his unwavering support of both me and this research, especially when neither one of us were quite sure if it was actually a "good" idea or not.

Heinrich Laue and Nicholas Osner: For continuously inspiring me with their own work, and providing a nearly endless supply of ideas, distractions, and coffee-related information.

My friends: For helping to take my mind off of my work, diligently asking how my research was going, and politely listening to the answer, even when none of the words I was saying made any sense.

My brother, Grant: For persevering through his own struggles, which reminded me that I could persevere through mine.

My mother, Beverly, and my father, Edrien: For always supporting me in everything that I do (even when that meant spending far too much money on my education) and for never giving up on me, or letting me give up on myself.

LIST OF ABBREVIATIONS

ADC	analogue-to-digital converter
ADS	almost difference set
AWGN	additive white Gaussian noise
CDF	cumulative distribution function
CPU	central processing unit
CS	compressive sensing
DFT	discrete Fourier transform
DS	difference set
DS-CDMA	direct sequence code-division multiple access
EM	electromagnetic
ESPRIT	estimation of signal parameters via rotational invariance techniques
EW	electronic warfare
FFT	fast Fourier transform
FOCUSS	FOCal Undetermined System Solver
FRI	finite rate of innovation
GLRT	generalised likelihood ratio test
GNU	GNU's Not Unix!
GSM	Global System for Mobile Communications
IED	improvised explosive device
LRT	likelihood ratio test

M-FOCUSS	multiple measurement vector FOCal Undetermined System Solver
MLE	maximum likelihood estimate
MSE	mean square error
MUSIC	multiple signal classification
MWC	modulated wideband converter
NP	Neyman-Pearson
PDF	probability density function
PNS	periodic nonuniform sampling
RB	random binary
RD	random demodulator
RF	radio-frequency
RG	random Gaussian
RMPI	random modulation pre-integrator
ROC	receiver operating characteristic
RV	random variable
SDR	software-defined radio
SLL	sidelobe level
SNR	signal-to-noise ratio
SWAP	size, weight and power
UAS	unmanned aerial system
UoS	union of subspace
USRP	Universal Software Radio Peripheral

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION

1.1 INTRODUCTORY REMARKS

This chapter provides an introduction to this research work. Motivation for the technique considered in this work is given in Section 1.2. The research objectives and questions are given in Section 1.3, while the hypothesis and overall methodology are provided in Section 1.4. The goals of this research are stated in Section 1.5, and the expected contributions and research outputs are summarised in Section 1.6. The chapter concludes with a layout of the work in Section 1.7.

1.2 PROBLEM STATEMENT

The focus of this research work is the reduction of samples prior to performing signal detection, and analysis of the effects on systems that process data in both the time and frequency domains. The primary motivation for this research work is detailed in Section 1.2.1, and the resulting gap in the current research literature that this work seeks to fill is briefly summarised in Section 1.2.2.

1.2.1 Context of the Problem

The Shannon-Nyquist sampling theorem states that a signal should be sampled at a rate equal to twice its bandwidth to allow for perfect reconstruction from its samples [1]. Continued developments in fields such as wideband communications, electronic warfare (EW), and radio-frequency (RF) technology have led to ever increasing signal bandwidths, such that sampling wideband signals at the Nyquist rate is becoming increasingly difficult, and produces a large amount of samples required for perfect reconstruction.

While many systems are developed to cope with an increased number of data samples, there are certain scenarios where processing a large amount of data samples is simply not feasible. Consider a mobile platform, such as an electric unmanned aerial system (UAS) designed to survey a large area of land [2], fitted with a signal acquisition receiver that monitors a wide band of frequencies in order to detect the presence of certain signals. The on-board processing capability of the platform is limited by the available size, weight and power (SWAP), often making it infeasible to perform detection on a large amount of data samples on the platform. This problem can be overcome either by storing the data on the platform for later processing or by sending the data to a base station for processing. On-board storage is impractical in applications where real-time information is required (e.g. [2]) due to the need for the platform to return before processing of data can take place. However, high data rates and long ranges will be required to transmit the necessary data to a base station, again leading to difficulties as wideband, high-power transmitters have high SWAP requirements. The conclusion is that the amount of data to be processed and/or transmitted must be reduced in order to meet the SWAP requirements of such platforms.

1.2.2 Research Gap

The problem of reducing the number of samples taken during signal acquisition has existed for decades; since the initial formulation of the Shannon-Nyquist theorem [1] researchers have been looking for ways to surmount it. Numerous approaches have been developed to achieve what is known as sub-Nyquist sampling (sampling below the Nyquist rate), most of which are focused on perfect signal reconstruction. Unfortunately, within the confines of platforms with restrictive SWAP requirements, certain drawbacks preclude the use of most of them. The primary drawback is that in order to compensate for the missing samples, the reconstruction process becomes more complicated – and thus more computationally complex – compared to Nyquist rate sampling. Secondly, many approaches require additional analogue hardware, often repeated over multiple channels, which increases the power required as well as the cost. Finally, many of the approaches either require knowledge of, or make assumptions about, the signal being sampled in order to reconstruct it.

As mentioned previously, most of the existing sub-Nyquist sampling approaches are concerned with perfect signal reconstruction. However, many signal processing problems, including signal detection, do not require perfect signal reconstruction in order to operate correctly. Without this constraint, the reconstruction process can be removed entirely, as can any additional analogue hardware. The penalty associated with any data-rate decrease without reconstruction is a reduction in detection performance due to fewer samples being available [3]. However, reducing the samples which need to be processed/transmitted allows wider signal bandwidths and/or additional channels to be accommodated by a given data rate (and thus SWAP), leading to an overall improvement in system performance.

The problem now becomes how to effectively reduce the amount of samples that must be processed by a detection algorithm. Current research into this problem is focused on applying the framework of compressive sensing (CS) [4, 5] to detection without any reconstruction, an approach referred to in this work as compressive detection. Due to their prevalence in CS, almost all of the existing work involves the use of random sensing matrices [6], which discard

certain data samples and replace the remaining ones with random projections, effectively compressing the retained samples.

Replacing retained samples with random projections is an additional computational step that could potentially be avoided. In this work random sensing matrices are replaced with binary sensing matrices, which are equivalent to simply discarding certain samples within a sampling window. This effectively represents the least computationally complex way of discarding data samples (which is to simply discard them). The primary issue with this approach is that simply discarding samples is also equivalent to applying a binary window function to the samples. Since the purpose of a window function is to shape a frequency response, the way that samples are discarded will have an impact on the frequency information, which is important considering that typically wideband signal acquisition receivers (the kind that would be fitted to the mobile platform previously described) process samples in both the time and frequency domains in order to extract signal information. The question that arises is how to choose which samples to keep and which samples to discard, or equivalently how best to construct binary sensing matrices, in order to minimise any negative impact on detection performance.

Discarding large numbers of samples taken in the time domain can be viewed as similar to discarding a large number of antenna elements in a filled antenna array, a technique known as massive thinning. Historically, random placement of elements – equivalent to randomly discarding samples – has proven as effective as most deterministic approaches, but this typically results in a dramatic loss of sidelobe control [7]. It has been shown that the sidelobe level (SLL) can be reduced to an arbitrarily low level by sufficiently increasing the number of antenna elements [8], which in practice means that the effects of noise can be made to dominate the effects of sidelobes for a sufficiently large number of elements. However, deterministic element placement using difference sets (DSs) [7] and almost difference sets (ADSs) [9] has been shown to produce power patterns with well-controlled sidelobes, and a reduction in peak sidelobe level when compared to random element placement, without requiring a large number of elements. Similar improvements are to be expected when DSs and ADSs are applied to discarding samples, with the sidelobe benefits directly applicable to the window's effect on

frequency response. DSs and ADSs also have other properties that are expected to prove beneficial.

1.3 RESEARCH OBJECTIVES AND QUESTIONS

The primary objective of this research is to determine the effects that applying a binary window function to data samples will have on detection performance. It is well documented that performing detection with fewer samples results in reduced detection performance for an otherwise identical detection scheme [3]. However, in a system that processes data in both the time and frequency domains, using a binary window function to discard samples is expected to affect the frequency information, which should also have an effect on the achievable detection performance.

The secondary objective is to compare the effects of random binary windows (the closest equivalent to the random sensing matrices used in compressive detection approaches) with binary windows constructed using DSs and ADSs, in order to determine whether the well-controlled sidelobes typically encountered when using DSs/ADSs for constructing thinned antenna arrays translates to discarding samples for detection, and if this is indeed beneficial for detection compared to discarding samples randomly.

These research objectives are summarised into the following research questions:

- For systems that process samples in both the time and frequency domains prior to detection, how will the application of binary window functions affect signal detection performance?
- How will the use of random binary window functions affect signal detection performance?
- How will the use of DSs and ADSs to construct binary window functions affect signal detection performance?

- Will the use of binary window functions constructed using DSs and ADSs show superior performance compared to random binary windows?

1.4 APPROACH

The overarching hypothesis of this research work is that for a system that implements frequency domain data processing (like many wideband signal acquisition receivers), while the application of binary window functions to reduce the number of data samples used during signal detection is expected to reduce the achievable detection performance, binary window functions constructed using DSs/ADSs should have less of a negative impact on the performance than random binary windows of the same length and number of elements. This is based on the relatively well-controlled sidelobes observed when DSs/ADSs are used to construct massively thinned antenna arrays, compared to their random counterparts [7, 9].

The approach taken to test this hypothesis is given here. First, appropriate detectors (which includes the means of calculating the test statistics and the thresholds) with samples discarded using binary window functions will be derived based on existing classical detectors. One of these detectors will be used for the remainder of the research work. Appropriate detection performance metrics for all of the detectors will also be derived, giving a means of objectively evaluating the detection performance.

Appropriate DSs and ADSs of various sizes will then be chosen, and used to construct binary windows that are used to select data samples prior to applying the detector. For each DS/ADS, binary windows shall be constructed based on the set itself, as well as filled binary windows of the same length and random binary windows of the same length and with the same number of elements, in order to compare the detection performance achieved using each type of window.

The detection performance will then be evaluated for all of the constructed windows using two different types of signals: simulated signals and generated practical signals. The simulated signals are created in software, while the generated practical signals will be created by transmitting the output of a signal generator over-the-air before receiving the signal using a software-defined radio (SDR).

Additionally, because the potential application of this research is the detection of cellphone signals (e.g. for anti-poaching activities), real world off-the-air Global System for Mobile Communications (GSM) signals will be recorded using a SDR and then processed by the detection algorithm, in order to observe the resulting detections and verify that actual signals can still be detected using the binary windows.[1] A large frequency band should be recorded for a relatively long time period, to illustrate the effects on multiple simultaneously received signals with random arrival times.

Finally, performing detection without the use of a reconstruction algorithm will be compared to the case where there is reconstruction. Specifically, the achieved detection performance, as well as the computation time of both approaches will be compared to determine if the assumed reduction in computation time is worth the reduction in detection performance.

1.5 RESEARCH GOALS

The goal of this research work is to produce a deterministic means of discarding data samples prior to performing signal detection that is extremely computationally simple, for application in scenarios where SWAP requirements severely limit the available computational and data processing capabilities. While it is expected that the simplicity of this technique will result in degradation and detection performance loss, this is already the case in many existing compressive detection approaches except under specific assumptions.

[1]The derived detection performance metrics will not be run on these signals, because the transmissions will not be continuous and the parameters will be unknown.

The primary applications of this technique are for platforms with restrictive SWAP requirements performing signal detection in order to reduce processing complexity, data rate, and/or energy consumption at the cost of a predictable decrease in detection performance. Such platforms include UASs designed to search for communication signals for anti-poaching activities [2] or improvised explosive device (IED) identification. While this technique is designed with EW systems in mind, it could potentially be expanded to the field of cognitive radio as well.

1.6 RESEARCH CONTRIBUTION AND OUTPUTS

The contribution of this research work will be a deterministic technique for discarding data samples prior to performing signal detection that requires very little computational power, ideal for applications with restrictive SWAP requirements, or where data needs to be transmitted over long distances. Specifically, this technique is beneficial towards systems that process data in the frequency domain, such as wideband signal acquisition receivers. The use of DSs and ADSs to construct sampling windows means that the effects of the window sidelobes on the detection performance are calculable *a priori* using only the set parameters, which is not possible with random binary windows. In addition, the cyclic properties of DSs and ADSs allow for sample-by-sample window overlap – which is not possible for windows optimised for a single configuration – and therefore allows windows to completely overlap with signals, potentially compensating for the negative impact on detection performance (the details of these properties are provided in Chapter 3). This technique was designed for EW applications, but could be expanded to the field of cognitive radio as well.

This research has led to the submission of one journal publication, the details of which are as follows:

Author(s): Wade R. Smith and Warren P. du Plessis

Title: Sample Reduction During Signal Detection Using Difference Sets and Almost Difference Sets

Journal: IEEE Transactions on Signal Processing

Abstract: Signal detection over wide bandwidths typically requires collection of a very large quantities of data. However, reducing the amount of data is important in applications with limited size, weight and power (SWAP) requirements. Most data-reduction approaches rely on reconstruction algorithms to compensate for the missing data, but these algorithms are computationally complex. Therefore, sample reduction without reconstruction is proposed, and the resulting impact on signal detection performance is evaluated. The cyclic and sidelobe properties of difference sets (DSs) and almost difference sets (ADSs) are exploited to reduce the interference caused by discarding samples, resulting in significant improvements over randomly discarding samples.

1.7 OVERVIEW OF THE STUDY

The content of the chapters in this study are outlined below.

Chapter 1: An introduction to the research work is presented, which includes motivation based on the potential application of this work and the limitations of the existing research. The objectives, goals, contributions, and outputs of this work are also provided.

Chapter 2: A background study of various techniques for performing reduced-sample detection on a platform with SWAP limitations is presented. Existing approaches based on sub-Nyquist sampling and compressive detection are evaluated for the given platform. The effects of performing detection when discarding samples using binary windows are discussed, and a potential solution based on DSs and ADSs is presented.

Chapter 3: A background study on DSs and ADSs is presented, including an overview of the theory of the sets and various properties that make them attractive to this research work.

Chapter 4: An extended analysis of signal detection is provided. This analysis includes an overview of the formation of test statistics, the signal models that are considered in this study, and means of quantifying the performance of a detector. For the chosen signal model, the test statistics and performance metrics for a detector where samples are discarded using binary windows are derived. An approach for applying DSs and ADSs to the derived detector is also presented.

Chapter 5: Empirical results to verify the impact of various binary windows on the performance of the detector derived in Chapter 4 are presented. Results for simulated signals, generated practical signals, and recorded practical signals are provided. Comparisons between the effects of random binary windows and those based on DSs/ADSs are presented, as well as between the proposed technique and detection using a reconstruction algorithm.

Chapter 6: A brief conclusion and some suggestions for future work are given.

Additionally, two addenda are provided that include additional relevant information, whose content is summarised below.

Addendum A: Plots of the frequency responses for the binary windows used in the experiments conducted in Chapter 5 are provided. These are used in Chapter 5 to help evaluate the detection performance of the various windows.

Addendum B: The experimental results for four of the six DSs/ADSs evaluated in Chapter 5 are provided. These are referred to in Chapter 5, and provide insight into how detection performance varies between different DSs and ADSs.

CHAPTER 2 SAMPLE REDUCTION

2.1 INTRODUCTORY REMARKS

An extensive background study of the potential approaches to reducing the number of samples used prior to performing signal detection is provided in this chapter. These approaches are specifically evaluated within the context of a platform with limited size, weight and power (SWAP) capabilities.

Details of existing sample reduction techniques – both general and those specific to signal detection – as well as their limitations in the given context are given in Section 2.2. The potential solution presented in this research work is outlined in Sections 2.3 and 2.4: Section 2.3 summarises the effects of discarding samples using binary window functions, and Section 2.4 provides details on difference sets (DSs) and almost difference sets (ADSs). A summary is given in Section 2.5.

The work described in this chapter forms part of a submitted journal paper [10].[1]

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

2.2 EXISTING SAMPLE REDUCTION TECHNIQUES

There are numerous sample reduction techniques detailed in the existing literature; those applicable to wide bandwidth, multiband signals are summarised in this section.[2] The two most promising approach types are those that involve reducing the sampling rate below the required Nyquist rate – known as sub-Nyquist sampling – and those that involve the application of compressive sensing (CS) theory directly to the signal detection algorithm – referred to in this work as compressive detection. Multiple sub-Nyquist sampling strategies are detailed in Section 2.2.1, and compressive detection is summarised in Section 2.2.2.

2.2.1 Sub-Nyquist Sampling of Wideband Signals

One solution to reduce the number of samples that need to be collected would be to implement some form of sub-Nyquist sampling, so that less samples are taken overall. There are numerous approaches that have been developed to achieve sub-Nyquist sampling of wideband signals; this section deals with some of the more prominent techniques that are applicable to wide bandwidth, multiband signals. The interested reader is directed to [11] for a more comprehensive review.

There are two general types of sub-Nyquist sampling strategies: those that require either full or partial knowledge of the signal/s being sampled, known as non-blind and semi-blind approaches, respectively, and those that do not require any prior knowledge, known as blind approaches. While the nature of the application of this work precludes non-blind and semi-blind strategies because it is assumed that almost nothing is known about the signals being sampled, some approaches are mentioned for completeness.

[2] Multiband signals are signals that consist of several known frequency bands [11].

Two classic, well known non-blind approaches are: demodulation [12], where a band of interest is shifted from high frequencies to lower ones, and undersampling (also known as direct bandpass sampling) [13], which involves selecting a sampling rate that manipulates aliasing to sample a band at a rate closer to its information content (also known as the Landau rate [14]). However, demodulation requires knowledge of the carrier frequency of the bands of interest, and undersampling is not applicable to multiband signals, so neither are considered in this work.

Another classic approach is periodic nonuniform sampling (PNS), where instead of sampling at uniform time intervals, for signals with M frequency bands with individual widths B, samples are collected from a series of M channels with relative time shifts

$$y_i(n) = x(nT_s + \psi_i), \quad i = 1, \ldots, M \tag{2.1}$$

such that the total sampling rate M/T_s is lower than the Nyquist rate [15]. Early approaches were non-blind, but further extensions have been made that resulted in semi-blind approaches [16–18], with samplers that operate independently of band positions with reconstruction algorithms that require knowledge of the frequency support. These semi-blind approaches specifically make use of mulicoset sampling [16–18], which is a PNS method where the grid is a subset of the uniform Nyquist grid.

From this point, all of the approaches that are discussed are blind sub-Nyquist sampling techniques. The principle difference between the non-blind/semi-blind and blind approaches is the signal model considered: the non-blind/semi-blind approaches are based on models consisting of a single subspace, while the blind approaches are based on models consisting of a union of multiple subspaces, referred to as the union of subspace (UoS) model [19]. The details of these models are not expanded upon here (the reader is directed to [11]); a very simple summary is that in the UoS model, the extra subspaces grant an additional degree of freedom to be used for sample rate reduction, allowing resulting approaches to operate without knowledge of the frequency support.

The blind approaches based on the UoS model can be further categorised as being based on one of three model types: multiband signals with unknown carrier positions, variants of finite rate of innovation (FRI) models, and a sparse sum of harmonic sinusoids. For the first model type, a fully blind framework based on multicoset sampling [20] as well as reconstruction algorithms [21] has been developed. This work has been subsequently extended into an approach known as the modulated wideband converter (MWC) [22]. The MWC essentially combines multiple sampling channels, as in PNS/multicoset sampling, with radio-frequency (RF) demodulation by mixing the input of each channel with a periodic waveform, in order to spread the frequency bands across the entire spectrum, including baseband, where each of the desired bands has a different weight applied. A lowpass filter then removes the spectrum outside of baseband before sampling, with a reconstruction algorithm using knowledge of the periodic waveforms to reconstruct the spectrum. A hardware-efficient realisation has also been developed [23].

The second model type includes variants of FRI models [24, 25], where signals are determined by a finite number of unknowns, or innovations, for a finite time interval. The simplest and most pervasive FRI model considers attenuated and time delayed copies of a pulse, where the attenuation and delay values are the innovations to be determined. Instead of determining these unknowns directly, sub-Nyquist sampling is achieved by instead determining the amplitudes and frequencies of the resulting Fourier series components, and then relating back to attenuations and delays using existing tools for array processing and spectral estimation [24]. There are numerous sampling approaches to obtain the Fourier components (once again, the interested reader is directed to [11]), and reconstruction is then be carried out by annihilating filter [24], or potentially techniques like multiple signal classification (MUSIC) [26] or estimation of signal parameters via rotational invariance techniques (ESPRIT) [27].

The final model type is that of a sparse sum of harmonic sinusoids, which is an attempt to bridge the concepts of CS [4, 5] and sampling analog signals [28]. This model led to the development of the random demodulator (RD), which operates similarly to the MWC, by multiplying the input signal with a pseudorandom binary sequence in order to spread the bands across the spectrum, before integrating and dumping at a reduced rate [28]. A

direct extension of the RD is the random modulation pre-integrator (RMPI), which consists of multiple parallel RD channels [29]. Both of these approaches make use of CS recovery algorithms for reconstruction [30].

It is obvious that there are numerous methods for implementing non-blind, semi-blind, and fully blind sub-Nyquist sampling, for a variety of different signal models. Unfortunately, within the confines of platforms with restrictive SWAP requirements, each of the blind approaches has certain drawbacks, that limit their application to this particular problem. Firstly, noise can become an issue for many of the approaches. The FRI-based approaches cannot exactly recover a noisy input signal regardless of sampling rate: instead the signal is estimated from its noisy samples. PNS, multicoset sampling, the MWC, the RD, and the RMPI are all susceptible to accumulation of wideband noise (for most of them due to the RF mixing process).

Secondly, most of the approaches require additional analogue hardware, often repeated over multiple channels, which increases the SWAP requirements as well as the cost. PNS and multicoset sampling require extremely precise time-delay elements and Nyquist-rate analogue-to-digital converters (ADCs), which typically makes these approaches practically infeasible for wideband signals. The MWC, RD, and RMPI all require additional RF front ends to perform the input mixing. On a platform where data reduction is being considered to reduce power requirements, adding additional hardware is counter-intuitive.

The final, but arguably primary, drawback for many of the approaches is that in order to compensate for the missing samples, the reconstruction process becomes more complicated – and thus more computationally complex – compared to the well-known interpolation method of Nyquist-rate sampling. FRI-based approaches use array processing and spectrum estimation techniques like MUSIC and ESPRIT. The processes for the RD and the RMPI make use of CS recovery methods, which can become highly non-linear, intensifying the problem even further. For a mobile platform, even if the reconstruction process could be applied at a base station following sample transmission (effectively removing the limitation of the platform), the complexity of the reconstruction algorithm coupled with the transmission time may limit real-time operation.

2.2.2 Compressive Signal Detection

A potential means of reducing computational complexity would be to remove the reconstruction process entirely. While this may appear to be unreasonable at first, consider the following: many signal processing problems, such as estimation, classification, filtering, and detection do not require perfect signal recovery in order to operate correctly. In fact, considering that most practically received signals are corrupted by noise and interference, most of these processes have been developed to operate using imperfect signals. Therefore, removing the reconstruction process can be seen as quite sensible, at least in the given context.

The concept of reduced-sample detection without reconstruction is already being investigated using the framework of CS, an approach referred to in this work as compressive detection. Works have focused on constructing test statistics in the compressed domain for a variety of signal models, from deterministic signals with known [31] and unknown [32] parameters, to Gaussian signal models [33–35]. Even Bayesian (as opposed to the more common Neyman-Pearson (NP)) detection [36] and extensions to multiple sensors [37] have been investigated.

The approaches to compressive detection typically fall into two categories: those that exploit signal sparsity [32, 33, 36–38],[3] and those that consider signals that are not necessarily sparse [31, 34, 35, 40]. The consensus is that without exploiting sparsity prior, a specific compressive detector will operate at an expected performance loss compared to the equivalent uncompressed detector, which is understandable considering less samples are available in the compressed case. However, if signal sparsity can be assumed, which can be seen as providing additional information to replace that which has been lost, a compressive detector can achieve performance and robustness comparable to its uncompressed counterpart [32, 40].

Since compressive detection is based on the principles of CS, in most approaches the received

[3] Sparsity refers to a sampled signal having only a small number of significant values [39]. For a vector $\mathbf{x}$ with N values or coefficients, the vector is considered K-sparse if only $K \ll N$ of those coefficients are significant.

samples are compressed and discarded using random sensing matrices, due to the fact that random measurements are considered universal [6, 38]. This potentially presents another means of reducing computational complexity, since converting to random projections involves an additional computational step to be performed on the sampled data.

The other primary issue with using random sensing matrices for this particular application is how they affect the frequency information of the signal, something that is not addressed in the existing compressive detection literature. If CS is used to reduce the number of samples prior to frequency domain processing, then it is expected that the discarded samples and the retained random measurements will result in a distorted signal spectrum. In applications of CS to the field of spectrum sensing and cognitive radio, this is typically overcome by reconstructing the frequency spectrum [41], however this would defeat the purpose of removing the original reconstruction algorithm.

2.3 SAMPLE REDUCTION USING BINARY WINDOWS

To save additional computational power and limit the potential degradation of frequency information, instead of discarding certain samples and replacing the remaining ones with random projections, samples could simply be discarded using equivalent binary sensing matrices, which represents the simplest approach to discarding samples. This would eliminate an additional computational step incurred in most existing compressive detection approaches, which would be beneficial to the proposed SWAP-limited system. A practically identical implementation of a binary sensing matrix is applying a binary window function to the data samples.

Unlike other data-reduction approaches, simply discarding samples does not make any assumptions about the data, requires only negligible processing on the platform itself, and does not require complex reconstruction before processing. The penalty associated with

any data-rate decrease is a reduction in detection performance due to fewer samples being available [3] (which is also the case in compressive detection where sparsity is not considered prior [31, 40]).

The question that arises is how to construct binary windows that will help limit the negative effects imposed on the detection process. This essentially involves choosing which samples to keep and which to discard, with the chosen sequence being influenced by certain considerations. In this work, the two primary considerations are as follows:

Sidelobe effects: The issue associated with simply discarding samples using a binary window function is that in the absence of a reconstruction process there will still be some distortion of the frequency information (even if the retained samples are not compressed using random measurements) that would also negatively impact detection performance. As mentioned previously, this is based on the fact that sampling with a binary sensing matrix is practically identical to applying a binary window function to the data samples, which will shape the frequency response in some way.

Based on this fact it becomes important to construct binary windows with good sidelobe properties. The sidelobes of a window function are a standard figure of merit [42], and will influence how the frequency spectrum is shaped: a large signal at one frequency that has high and/or uncontrolled sidelobes could obscure a smaller signal at a different frequency. It is therefore important to construct windows with low, or at least well controlled, sidelobes.

Window overlap and cyclicity: For a long stream of data samples, samples are collected using sequential finite length windows. Sampling windows can potentially be applied to non-overlapping partitions of the sequence, as shown in Figure 2.1(a), or to overlapped partitions, as shown in Figure 2.1(b). In the case of overlapped partitions, the window is cyclically shifted as it moves through time, ensuring that the same samples would be collected independent of the overlap.

When samples are discarded to optimise the sidelobe properties for only one sample

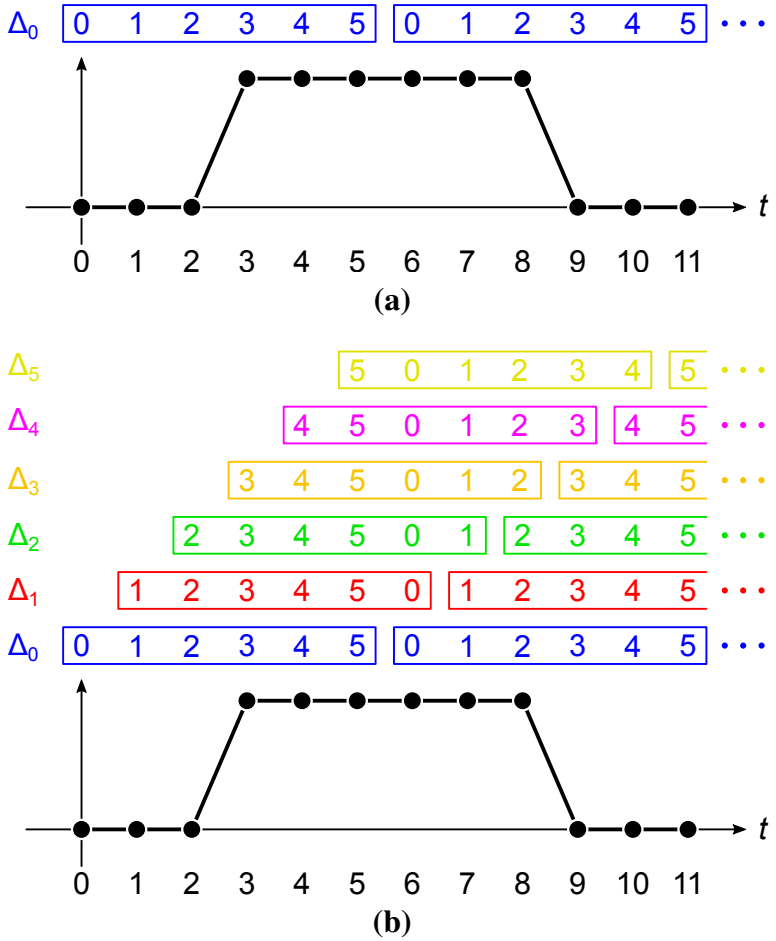

Figure 2.1. Collecting time samples using windows (a) without and (b) with cyclic properties. Samples within a block are processed together. (Adapted, with permission, from [10]. ©2018 IEEE.)

configuration, as in Figure 2.1(a), any cyclic shift of the samples could lead to poor sidelobe performance. Overlapping of the windows is not possible, which means that the shortest signal under the worst-case condition will only fill 50% of each window when the window and signal lengths are matched [43]. An ideal solution would be to construct windows with cyclic properties so that their performance is independent of cyclic shift, allowing for windows to be overlapped. While overlapping windows is well known to increase computational complexity [42], because samples are being discarded this can be somewhat mitigated, and assuming that windows lengths are chosen appropriately, large overlap could be achieved that would benefit overall performance by allowing for perfect overlap with a signal.

2.4 DIFFERENCE SETS AND ALMOST DIFFERENCE SETS

The question that arises is how to choose which samples to keep and which samples to discard, or equivalently how to design binary sensing matrices, in order to minimise any negative impact on detection performance. In Section 2.3 it was elucidated that binary sequences which are cyclic and have relatively good sidelobe properties would be desirable for this purpose.

One obvious approach would be to simply discard samples randomly, or equivalently to construct random binary sensing matrices. After all, random sensing matrices are prevalent in CS literature [4, 5] and by extension many approaches to compressive detection, so random binary sensing matrices are simply a further extension, removing the compression to reduce computational complexity.

Unfortunately, discarding samples at random is expected to lead to a dramatic loss of sidelobe control in the resulting frequency response. This assumption is based on the fact that discarding samples taken in the time domain can be viewed as similar to discarding antenna elements from a filled antenna array, a technique known as thinning.[4] Antenna arrays with random placement of elements – equivalent to randomly discarding samples – has been shown to have comparable antenna properties to filled arrays with similar parameters but apart from dramatic a loss of sidelobe control [8, 44].

However, while the field of antenna array thinning provides evidence of the problems associated with random sampling, it is also the source of a potential solution. Historically, few deterministic element placement techniques fared any better at retaining sidelobe control as so called “cut-and-try” random element placement [7]. Of these promising techniques, deterministic element placement using difference sets (DSs) [7] and almost difference sets (ADSs) [9] has been shown to produce power patterns with well-controlled sidelobes, and

[4]The only difference is the domain: discarding samples in the time domain versus discarding elements in the spatial domain.

a reduction in peak sidelobe level when compared to random element placement. Similar improvements are to be expected when DSs and ADSs are applied to discarding samples.

DSs are a concept from the field of mathematics known as combinatorics [45, 46], and have been studied extensively in combinatorial design theory [45–48]. As mentioned previously, DSs have been used in the design of thinned antenna arrays with desirable sidelobes that outperform equivalent randomly thinned arrays [7]. Additionally, they have been used in the construction of spreading codes for direct sequence code-division multiple access (DS-CDMA) systems [49], optimal codebooks for communications and related applications [50], and discrete Fourier transform (DFT) measurement matrices for subsampling in radar systems [51]. These are all applications that use DSs to construct sequences or matrices that achieve the Welch bound [52], which is a lower bound on the maximum cross-correlation, which in turn is a lower bound on the minimal matrix coherence. The importance of this is that CS uses low coherence as an indication of good sparse recovery capability [6], so if measurement matrices constructed using DSs produce similar results in other applications, their potential application to compressive detection is promising.

In addition to DSs, there are also ADSs, which as the name implies are "almost" DSs, possessing similar properties to DSs but having slightly different guidelines for their construction [53–56]. Consequently, this means that when there is not DS, there should be an ADS. While the definition for a DS is widely agreed upon, the exact definition for an ADS varies [53, 54].

Finally, in addition to their sidelobe properties, DSs and ADSs also have cyclic properties, which make them applicable to the construction of cyclically-shifting sample windows as shown in Figure 2.1. DSs and ADSs provide sidelobe control for all cyclic shifts, which means that all the cases seen in Figure 2.1(b) will have well-controlled sidelobes if the windows are constructed using a DS. The resulting cyclic shifts will always allow 100% overlap of the windows and the signal (shift Δ_3 in Figure 2.1(b)) as long as the window length is appropriately chosen. The cyclic properties of DSs and ADSs mean that some of the detection performance lost by discarding samples can be regained by ensuring full overlap with all signals while maintaining sidelobe control.

2.5 SUMMARY

Several approaches for sample reduction during signal detection have been considered. However, on a platform with limited SWAP capabilities, most existing sample reduction schemes are not feasible for implementation, primarily due to the high computational costs of reconstruction algorithms. In order to minimise computational complexity, discarding samples using binary sensing matrices is considered in this work, which represents the simplest approach to discarding samples (which is to discard them).

As a means of constructing binary sensing matrices to minimise negative impact on frequency information and detection performance, DSs and ADSs have been proposed due to their attractive properties. A detailed background study into DSs, ADSs, and their properties is given in Chapter 3.

CHAPTER 3 DIFFERENCE SETS AND ALMOST DIFFERENCE SETS

3.1 INTRODUCTORY REMARKS

This chapter summarises the concepts of difference sets (DSs) and almost difference sets (ADSs). Specifically, it provides a definition of the sets and details important properties that make the sets attractive to the problem of discarding samples prior to performing signal detection.

An overview of DSs and ADSs is presented in Section 3.2. The important properties of DSs and ADSs, specifically cyclicity, complementary sets, and frequency response characteristics, are summarised in Section 3.3. A summary is provided in Section 3.4.

The work described in this chapter forms part of a submitted journal paper [10].[1]

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

3.2 OVERVIEW

The presentation and notation below are based on [7, 9]. DSs and ADSs are concepts from the branch of mathematics known as combinatorics [45, 46]. A (V, K, Λ) DS given by

$$D = \{d_k \in [0, V-1], d_h \neq d_l \, \forall \, h \neq l, k, h, l = 0, \ldots, K-1\} \tag{3.1}$$

is a subset of an Abelian group G of order V such that the set of differences

$$M = \{m_j = (d_h - d_l), d_h \neq d_l \, \forall \, j = 0, \ldots, K \cdot (K-1) - 1\} \tag{3.2}$$

contain every nonzero element of G each exactly Λ times [7, 45, 46].[2] [3] Likewise a (V, K, Λ, t) ADS D is a subset of an Abelian group G of order V such that the set of differences M as defined in (3.2) contains t nonzero elements of G each exactly Λ times, and the remaining $V - 1 - t$ nonzero elements of G each exactly $\Lambda + 1$ times [9, 56]. Consequently, DSs are ADSs for which $t = V - 1$ or $t = 0$.

Two simple examples are given here: one for a DS, and one for an ADS. First, consider the DS (7, 4, 2)

$$D_1^{\mathrm{DS}} = \{0, 3, 5, 6\}; \qquad V = 7, \; K = 4, \; \Lambda = 2. \tag{3.3}$$

The difference property of the group is shown in Table 3.1, where the differences between every element of D_1^{DS} and every other element are given. Observe that a difference corresponding to each nonzero element of the set $G_1^{\mathrm{DS}} = [0, 7-1]$ appears exactly $\Lambda = 2$ times.

Now consider the ADS (13, 3, 0, 6)

$$D_1^{\mathrm{ADS}} = \{1, 3, 9\}; \qquad V = 13, \; K = 3, \; \Lambda = 0, \; t = 6, \tag{3.4}$$

[2] The notation used in this work is the same as that in [7] in order to prevent confusion regarding standard electromagnetic (EM) symbols, such as λ, which typically expresses wavelength.

[3] The differences are taken modulo V.

Table 3.1. The difference property for the DS (7, 4, 2). (Adapted from [7]. ©1999 IEEE.)

i	j	$(d_i - d_j) \bmod V$
0	1	$(0-3) \bmod 7 = 4$
0	2	$(0-5) \bmod 7 = 2$
0	3	$(0-6) \bmod 7 = 1$
1	0	$(3-0) \bmod 7 = 3$
1	2	$(3-5) \bmod 7 = 5$
1	3	$(3-6) \bmod 7 = 4$
2	0	$(5-0) \bmod 7 = 5$
2	1	$(5-3) \bmod 7 = 2$
2	3	$(5-6) \bmod 7 = 6$
3	0	$(6-0) \bmod 7 = 6$
3	1	$(6-3) \bmod 7 = 3$
3	2	$(6-5) \bmod 7 = 1$

whose difference property is shown in Table 3.2. For the ADS, differences corresponding to $t = 6$ nonzero elements of the group $G_1^{\text{ADS}} = [0, 13-1]$ appear exactly $\Lambda = 0$ times, while the remaining $V - 1 - t = 6$ nonzero elements appear exactly $\Lambda + 1 = 1$ times.

Table 3.2. The difference property for the ADS (13, 3, 0, 6). (Adapted from [7]. ©1999 IEEE.)

i	j	$(d_i - d_j) \bmod V$
0	1	$(1-3) \bmod 13 = 11$
0	2	$(1-9) \bmod 13 = 5$
1	0	$(3-1) \bmod 13 = 2$
1	2	$(3-9) \bmod 13 = 7$
2	0	$(9-1) \bmod 13 = 8$
2	1	$(9-3) \bmod 13 = 6$

3.3 IMPORTANT PROPERTIES

DSs and ADSs have several properties that make them attractive to the construction of binary windows for discarding samples prior to detection. Specifically, DSs and ADSs are cyclic, have complementary sets, and their frequency responses have relatively well-controlled sidelobes [7, 9]. Cyclicity is explored in Section 3.3.1, complementary sets in Section 3.3.2, and frequency response in Section 3.3.3.

3.3.1 Cyclicity

Given a (V, K, Λ) DS D, the set

$$D' = \{d_0 + n, d_1 + n, \ldots, d_{K-1} + n\} = D + n \tag{3.5}$$

where n is an integer and each element is taken modulo V, is also a (V, K, Λ) DS; D' is called a cyclic shift of D [7]. A similar cyclic property holds for ADSs [9]. As an example consider the DS D_1^{DS} in (3.3), whose cyclic shifts are illustrated in Figure 3.1 (which is modified from Figure 2.1(b)). Note that the windows are cycled in reverse (i.e. the first element in the current window is placed at the end of the subsequent window), so each shift n is actually achieved by adding $V - n$ to D_1^{DS}. The shift Δ_4 overlaps completely with the desired signal; to verify, observe that

$$\Delta_4 = \left(D_1^{\text{DS}} + V - 4\right) \bmod V \tag{3.6}$$

$$= \left(D_1^{\text{DS}} + 3\right) \bmod 7 \tag{3.7}$$

$$= \{3, 6, 8, 9\} \bmod 7 \tag{3.8}$$

$$= \{3, 6, 1, 2\} \tag{3.9}$$

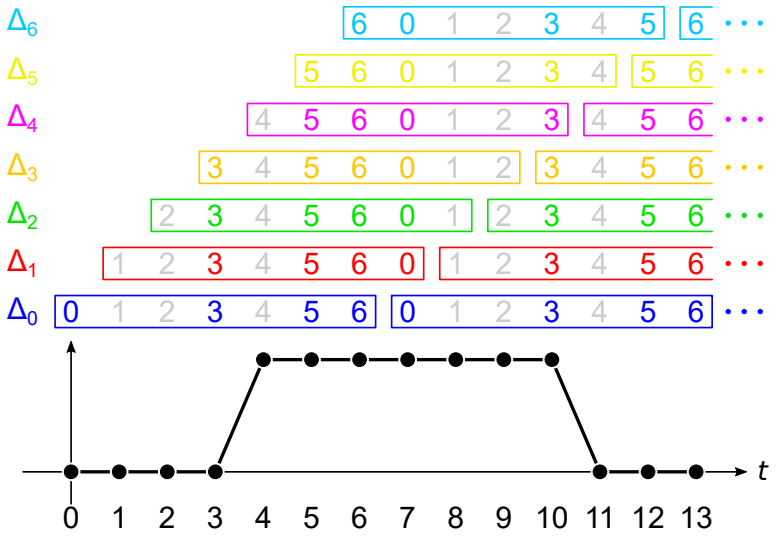

Figure 3.1. Cyclically shifting windows based on the DS (7, 4, 2). (Adapted, with permission, from [10]. ©2018 IEEE.)

which once rearranged gives

$$\Delta_4 = \{1, 2, 3, 6\} \tag{3.10}$$

meaning that the shift Δ_4 takes the 1st, 2nd, 3rd, and 6th elements from the window, as can be seen in Figure 3.1.

The cyclic property means that each of the $V - 1$ cyclic shifts of a DS /ADS will have the exact same parameters, and since the sidelobe performance of a DS/ADS is dictated by these parameters (as will be seen later on), the performance of a window whose samples are selected using a DS or ADS will not change significantly as that window cyclically shifts in time.

3.3.2 Complementary Sets

Each DS and ADS has a complementary set [7, 9], given by

$$D^* = \begin{cases} (V, V-K, V-2K+\Lambda) & \text{for a DS} \\ (V, V-K, V-2K+\Lambda, t) & \text{for an ADS.} \end{cases} \tag{3.11}$$

The complementary set is formed from the original set by removing all the values of d_k in (3.5) and replacing them by the values in $[0, V-1]$ which were not originally present, expressed as

$$D^* = \{d_k \in [0, V-1], d_k \notin D, \forall\, k = 0, \ldots, V-K-1\}\,. \tag{3.12}$$

For example, the complement of the DS D_1^{DS} in (3.3) is

$$D_1^{*\mathrm{DS}} = \{1, 2, 4\}\,; \qquad V = 7,\ K = 3,\ \Lambda = 1 \tag{3.13}$$

and the complement of the ADS D_1^{ADS} in (3.4) is

$$D_1^{*\mathrm{ADS}} = \{0, 2, 4, 5, 6, 7, 8, 10, 11, 12\}\,; \qquad V = 13,\ K = 10,\ \Lambda = 7,\ t = 6. \tag{3.14}$$

The principle benefit of this property is that DSs and ADSs with a wider range of parameters can be constructed by simply taking one set and using the elements that are not present. Effectively, one set offers two possible configurations.

3.3.3 Autocorrelation and Frequency Response

The final useful property relates to sidelobe control and provides the key motivation for the use of DSs and ADSs in this work. DSs and ADSs have seen successful application in the construction of massively thinned antenna arrays by dictating the locations of antenna elements [7, 9], a problem that is analogous to discarding samples from a sample window. A DS or ADS D can be used to construct a binary sequence

$$A_V = \{a_0, a_1, \ldots, a_{V-1}\} \tag{3.15}$$

where $a_j = 1$ if $j \in D$ and $a_j = 0$ if $j \notin D$. In antenna array thinning, this binary sequence dictates the location of elements in the array; a one indicates that an element should be placed at a location, and a zero indicates that it should not. In sampling, this sequence dictates whether or not a sample should be discarded (for clarity, a one indicates that a sample should be taken, and a zero indicates that it should be discarded).

By periodically repeating A_V to create an infinite sequence A_∞, the autocorrelation function can then be defined as

$$C_\infty(\tau) = \sum_{n=0}^{V-1} a_n a_{n+\tau}. \tag{3.16}$$

If and only if A_∞ is formed from a DS, the autocorrelation function reduces to [7]

$$C_\infty^{\mathrm{DS}}(\tau) = \begin{cases} K & \text{if } \tau(\mathrm{mod}\ V) = 0 \\ \Lambda & \text{otherwise.} \end{cases} \tag{3.17}$$

Likewise, if and only if A_∞ is formed from an ADS, then

$$C_\infty^{\mathrm{ADS}}(\tau) = \begin{cases} K & \text{if } \tau(\mathrm{mod}\ V) = 0 \\ \Lambda + 1 & \text{if } \tau(\mathrm{mod}\ V) \in L \\ \Lambda & \text{otherwise} \end{cases} \tag{3.18}$$

where $L = \{l_p \in \mathbb{Z};\ p = 1, \ldots, V - 1 - t\}$, i.e. L represents the $V - 1 - t$ elements that occur in the set of differences each exactly $\Lambda + 1$ times [9].

The power pattern of an antenna array, which is equivalent to the frequency response of a sample window, can be determined by taking the Fourier transform of the autocorrelation function. For a DS, the Fourier transform of (3.17) can be expressed as [7]

$$ff_\infty^{\mathrm{DS}}(u) = \Lambda \left[\frac{1}{x_0} \sum_{n=-\infty}^{\infty} \delta\left(u - \frac{n}{x_0}\right) \right] + (K - \Lambda) \times \left[\frac{1}{V x_0} \sum_{n=-\infty}^{\infty} \delta\left(u - \frac{n}{V x_0}\right) \right], \tag{3.19}$$

where x_0 is the temporal spacing between samples. Making use of the fact that $K(K-1) = \Lambda(V-1)$ [7], a normalised version of (3.19) can be expressed as

$$\overline{ff_{\infty}^{\mathrm{DS}}(u)} = (1-\rho)\left[\frac{1}{x_0}\sum_{n=-\infty}^{\infty}\delta\left(u-\frac{n}{x_0}\right)\right] + \rho\left[\frac{1}{Vx_0}\sum_{n=-\infty}^{\infty}\delta\left(u-\frac{n}{Vx_0}\right)\right] \tag{3.20}$$

where

$$\rho = \frac{1}{K}\left[1-\frac{(K-1)}{(V-1)}\right]. \tag{3.21}$$

This normalised power pattern has a mainlobe impulse with an area of 1 at $u = 0, \pm 1/x_0, \pm 2/x_0, \ldots$ and identical sidelobe impulses with area ρ at $u = \pm 1/Vx_0, \pm 2/Vx_0, \ldots$ [7]. This means that the peak sidelobe level (SLL) of the pattern is

$$\mathrm{PSL}_{\infty}^{\mathrm{DS}} = \rho. \tag{3.22}$$

For an ADS, the Fourier transform of (3.18) is given by [9]

$$ff_{\infty}^{\mathrm{ADS}}(u) = \Lambda\left[\frac{1}{x_0}\sum_{n=-\infty}^{\infty}\delta\left(u-\frac{n}{x_0}\right)\right] + \left(K-\Lambda+\sum_{p=1}^{V-1-t}e^{j\,2\pi n\,l_p/V}\right)\times \left[\frac{1}{Vx_0}\sum_{n=-\infty}^{\infty}\delta\left(u-\frac{n}{Vx_0}\right)\right]. \tag{3.23}$$

Unfortunately, (3.23) cannot be further simplified as is the case with (3.19) for DSs due to the fact that the term l_p from the set L depends on the ADS at hand, and thus (3.23) cannot be further evaluated in a general fashion [9]. However, it is still possible to provide an estimate of the peak SLL, which will be limited by the following upper and lower bounds [9]

$$\mathrm{PSL}_{\infty,\max}^{\mathrm{ADS}} = \frac{K-\Lambda-1+\sqrt{t(V-t)}}{(V-1)\Lambda+K-1+V-t} \tag{3.24}$$

$$\mathrm{PSL}_{\infty,\min}^{\mathrm{ADS}} = \frac{K-\Lambda-1-\sqrt{\frac{t(V-t)}{(V-1)}}}{(V-1)\Lambda+K-1+V-t}. \tag{3.25}$$

For both DSs and ADSs, the Fourier transform of an infinite length window can be expressed in terms of a finite length window as

$$ff_\infty(u) = ff_V(u) \cdot \left[\frac{1}{Vx_0}\sum_{n=-\infty}^{\infty}\delta\left(u - \frac{n}{Vx_0}\right)\right] \tag{3.26}$$

where $ff_V(u)$ is the Fourier transform of the autocorrelation function of the finite length window. This shows that $ff_V(u)$ forms a form of envelope for the $ff_\infty(u)$ impulse train, and must necessarily pass through the same fixed points [7], prescribed by (3.19) for a DS and (3.23) for an ADS. This means that for a DS, the peak SLL is

$$\mathrm{PSL}_V^{\mathrm{DS}} = \mathrm{PSL}_\infty^{\mathrm{DS}} = \rho \tag{3.27}$$

and for an ADS, the peak SLL can be estimated by the inequalities

$$\mathrm{PSL}_{V,\min}^{\mathrm{ADS}} \leq \mathrm{PSL}_V^{\mathrm{ADS}} \leq \mathrm{PSL}_{V,\max}^{\mathrm{ADS}} \tag{3.28}$$

where

$$\mathrm{PSL}_{V,\min}^{\mathrm{ADS}} = \mathrm{PSL}_{\infty,\min}^{\mathrm{ADS}} \tag{3.29}$$

$$\mathrm{PSL}_{V,\max}^{\mathrm{ADS}} = \mathrm{E}\left[\Phi_V^{\min}\right]\mathrm{PSL}_{\infty,\max}^{\mathrm{ADS}} \tag{3.30}$$

and $\mathrm{E}\left[\Phi_V^{\min}\right] \approx 0.8488 + 1.128\log_{10}(V)$ for sufficiently large values of V [9].

To illustrate, the frequency responses for windows based on the DSs (37,9,2) and (40,13,4) and the ADSs (37,10,2,18) and (40,20,9,10) are presented; specifically, the frequency responses of windows based on the DS and ADS themselves, their complements, and corresponding random binary windows. These sets were selected due to the fact that they are short, so the reader can more easily see the responses, and of the shorter sets available the DS (40,13,4) and the ADS (40,20,9,10) are two of the only sets that have an even value of V, with the two remaining sets being chosen for having a similar length V that is also odd (this is done to explore the effects of symmetry versus asymmetry). Random binary windows are used as they are the closest comparison to the random sensing matrices used in compressive detection works.

Each random binary window was selected by generating 100 random binary windows with the appropriate values for V and K and determining the maximum SLL of the resulting frequency response (calculated over $10 \times V$ points) for all V possible cyclic shifts of the window, and then selecting the window which had the lowest maximum SLL for a cyclic shift.

For each window the cyclic shifts corresponding to the lowest, median, and highest SLL values are plotted against u, which represents normalised frequency in this work. The parameter u is borrowed from the field of antenna arrays, where it refers to angular range. This stems from the use of DSs and ADSs in constructing thinned antenna arrays [7, 9], which is where the concept of applying them to compressive detection originated from.

Figure 3.2 shows the frequency responses for filled binary windows based on the chosen sets. These serve as a reference, as they represent the best case scenario where no samples are discarded. Note how the sidelobes of the frequency responses roll off, as in a typical rectangular window.

Figures 3.3 and 3.5 show the frequency responses for the windows based on the DS (37,9,2) and the ADS (37,10,2,18) and their complements, respectively, and Figures 3.4 and 3.6 show the frequency responses for the windows based on the DS (40,13,4) and the ADS (40,20,9,10) and their complements, respectively. For each plotted frequency response, three patterns are shown, corresponding to the cyclic shifts that produce the lowest, median, and highest maximum SLL values for the window. The black markers in each plot indicate the level at the sample points given in (3.26). In all of the plots the black line indicates the average power for any binary window with length V and K elements, which also happens to be the peak SLL value for a DS given by (3.21). For the ADS-based windows, the blue and red lines indicate the estimates for the minimum and maximum SLL values given in (3.29) and (3.30), respectively.

For all of the plots in Figures 3.3, 3.4, 3.5, and 3.6, the sidelobes do not roll off as in Figure 3.2 for the filled windows, and the patterns are seemingly uncontrolled whether the window is based on a DS/ADS or a random equivalent. However, in this work it is the pattern value

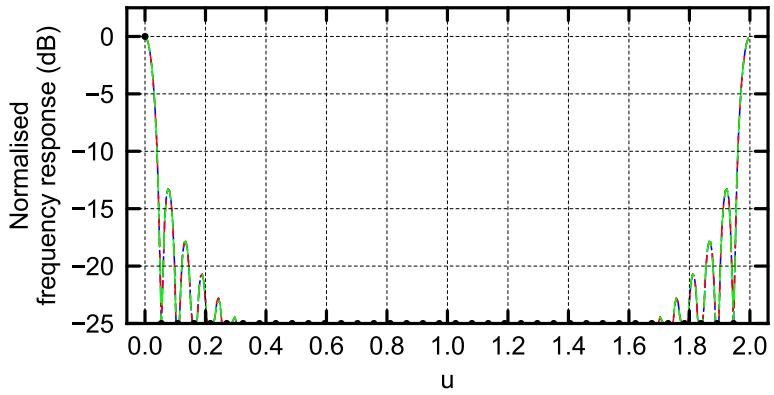

(a) Filled window based on (37,9,2)

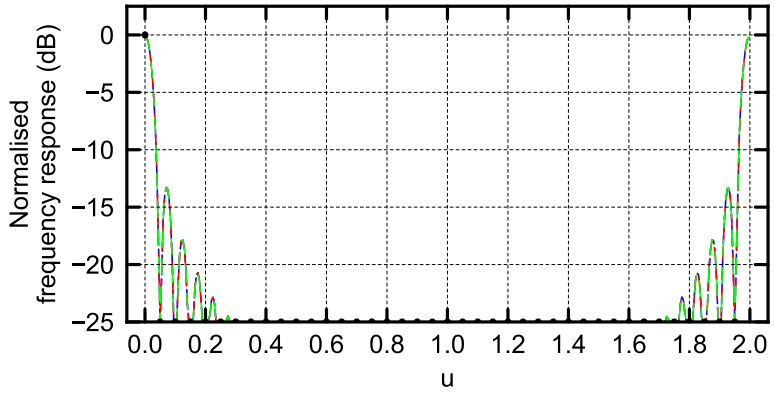

(b) Filled window based on (40,13,4)

Figure 3.2. Frequency responses resulting from filled windows based on the DSs (37,9,2) and (40,13,4).

at each of the sample points (the black markers) that is important, rather than the values in between. Figures 3.3(a) and 3.4(a) show that the sample points follow no apparent pattern for a random binary window. However, the sample points are much more controlled for the windows based on DSs and ADSs. Figures 3.3(b) and 3.4(b) show that for DS-based windows the sample points are constrained to a fixed level whose value can be calculated using (3.21) with the appropriate values of V and K.

Figures 3.3(c) and 3.4(c) show that for ADS windows the results are similar to those for DS windows, although not as consistent. In Figure 3.3(c) it can be seen that for ADS with V odd, the sample points converge to two distinct levels, as expected from (3.18) and (3.23), and while those levels cannot be directly calculated they do fall between the extremes estimated

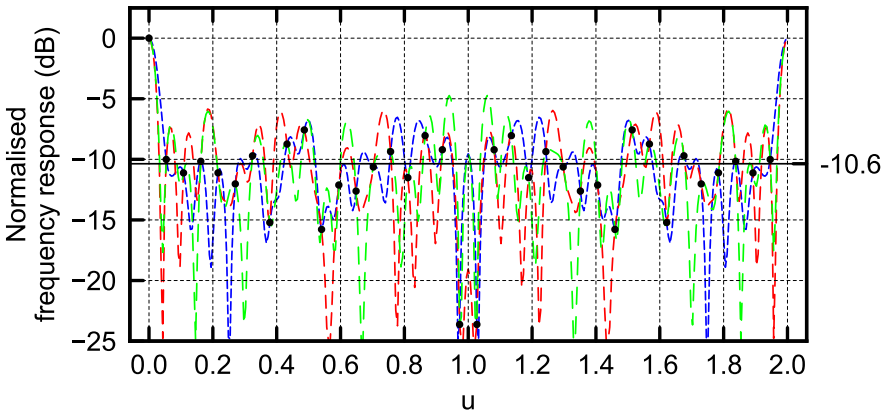

(a) Random window based on (37,9,2)

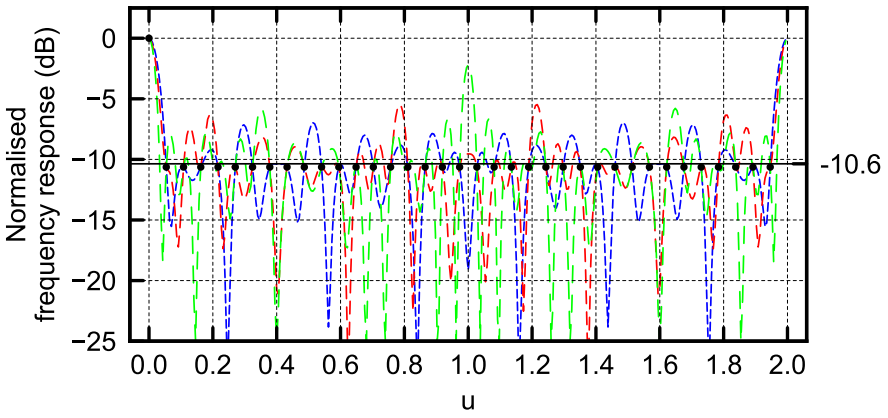

(b) DS window based on (37,9,2)

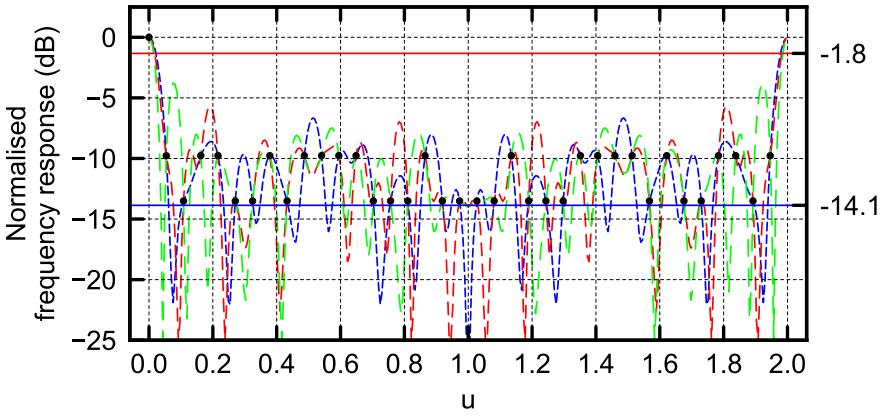

(c) ADS window based on (37,10,2,18)

Figure 3.3. Frequency responses resulting from windows based on the DS (37,9,2) and the ADS (37,10,2,18).

by (3.29) and (3.30). However, in Figure 3.4(c) it can be seen that for ADS with V even, the levels are not distinct, although they are still more controlled than those of the equivalent

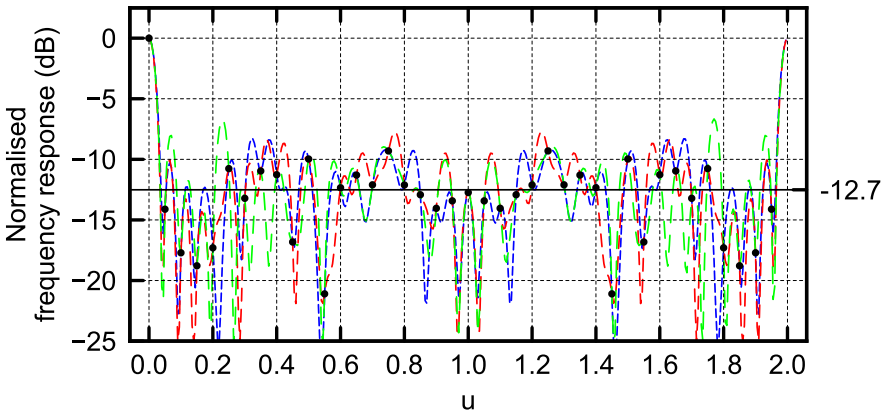

(a) Random window based on (40,13,4)

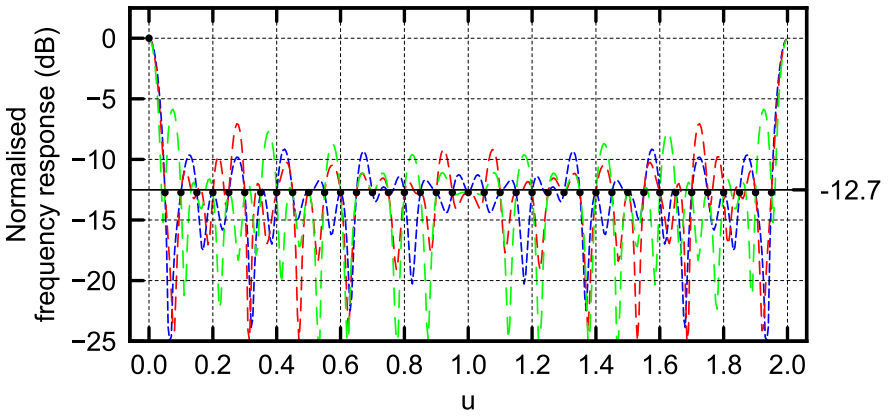

(b) DS window based on (40,13,4)

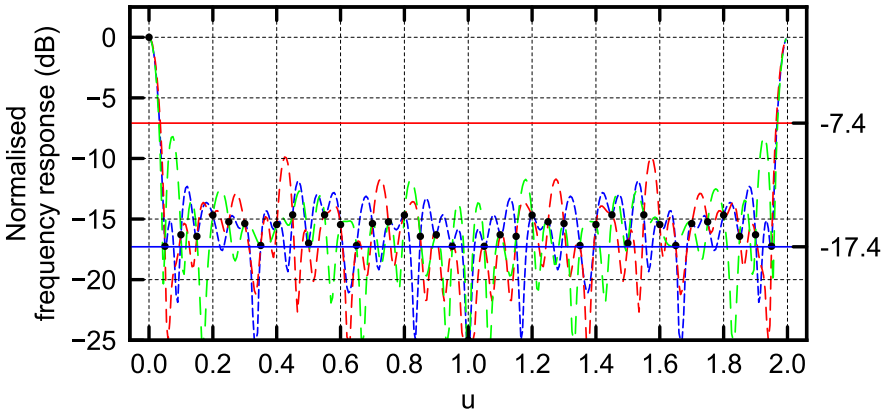

(c) ADS window based on (40,20,9,10)

Figure 3.4. Frequency responses resulting from windows based on the DS (40,13,4) and the ADS (40,20,9,10).

random binary window, and the middle most sample point of the frequency response has a value far below the minimum extreme estimated by (3.29). The value is so far below that it

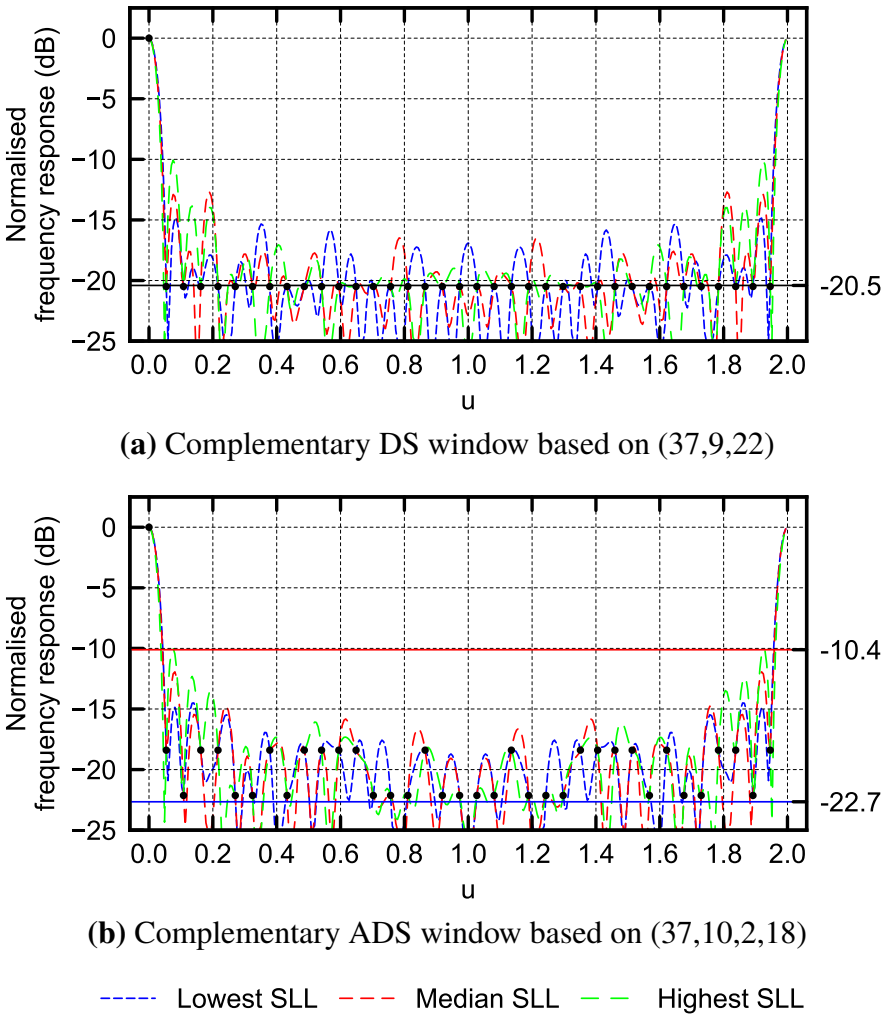

(a) Complementary DS window based on (37,9,22)

(b) Complementary ADS window based on (37,10,2,18)

Figure 3.5. Frequency responses resulting from windows based on the complements of the DS (37,9,22) and the ADS (37,10,2,18).

is assumed that it is actually a null in the frequency response. The reason for this deviation from the distinct levels expected (and achieved when V is odd) is unknown at this point, but is beyond the scope of this work.

Figures 3.5 and 3.6 show that complementary DSs/ADSs obey the same principles as typical DSs/ADSs, with the sample points converging to distinct levels. The benefit of complementary sets is also obvious from Figure 3.6(a): retaining more of the original samples lowers the peak SLL for the sample points, as would be expected.

There is one final observation to note. For each of the windows, the resulting frequency

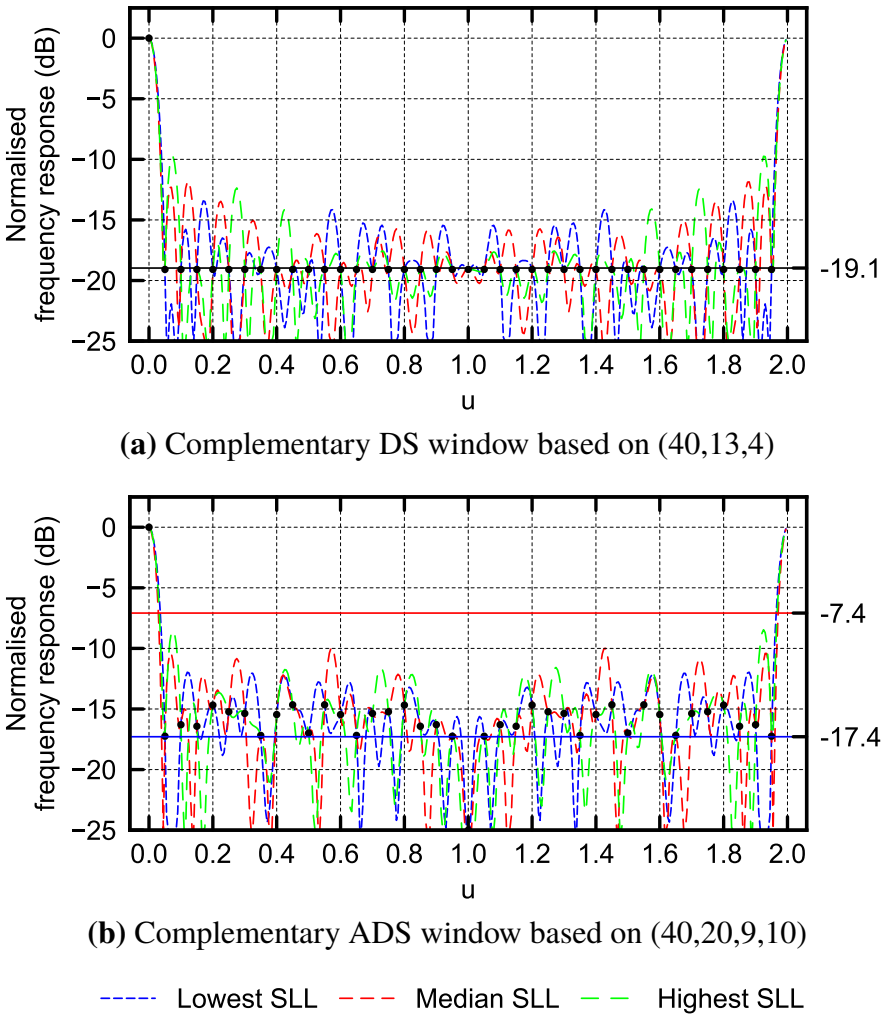

(a) Complementary DS window based on (40,13,4)

(b) Complementary ADS window based on (40,20,9,10)

Figure 3.6. Frequency responses resulting from windows based on the complement of the DS (40,13,4) and the ADS (40,20,9,10).

response for every cyclic shift passes through the sample points dictated by the black markers. While this is expected for the windows based on DSs and ADSs, due to their cyclic properties, this is not expected for the random binary windows. A reasonable explanation for this is that a cyclic shift of a binary window is equivalent to shifting the window in time, as seen in Figure 3.1. A basic property of the Fourier transform is that a time shift in the time domain is equivalent to a phase shift in the frequency domain. Since only the magnitude information of the frequency responses is kept, a phase shift should have no effect on the resulting patterns, at least at the sample points, which is why all of the responses for a given window pass through these points.[4]

[4]For a given window, the frequency response of each cyclic shift does vary in between the sample points, due to the interpolation that occurs as a result of the parameter u being much longer than V.

In summary, while binary windows based on DSs and ADSs have frequency responses that do not roll off like the response of a filled window, at the sample points the patterns are far more controlled than the equivalent random binary window, which means that the interference introduced by the lack of roll off can be mitigated somewhat, especially by correctly exploiting thinning percentages.

3.4 SUMMARY

An overview of DSs and ADSs has been provided. The sets have several properties that make them an attractive option for constructing binary sensing matrices. All DSs and ADSs are cyclic, meaning that as a binary sampling window moves through time its parameters should not change (although this also appears to be the case for random binary matrices). Each set has a complementary set, which effectively doubles the number of sets available (which is already considerable [57, 58]). Finally, the frequency response of a window based on a DS/ADS is well-controlled compared to that of the equivalent random binary window, with SLLs at specific sample points being predictable in the case of a DS (and estimable in the case of an ADS) using only the parameters of the set, which is not possible for a random binary window. Based on this, DSs/ADSs are applied to the detection problem in Chapter 4.

CHAPTER 4 DETECTION ANALYSIS

4.1 INTRODUCTORY REMARKS

A mathematical analysis of the process of discarding certain data samples and performing signal detection is presented in this chapter. It seeks to provide an understanding of how discarding certain data samples will affect the formation of test statistics and thresholds – which will invariably affect the detection performance – from a mathematical point of view.

An overview of signal detection theory, including hypothesis testing, potential signal models, and performance metrics is given in Section 4.2. A mathematical framework for representing the process of discarding data samples is explained in Section 4.3. For the signal model considered in this work, the rederived test statistics and performance metrics, for both real and complex data samples, are presented in Section 4.4, as well as practical means of implementing the test statistics. Application of difference sets (DSs) and almost difference sets (ADSs) to the derived detector is discussed in Section 4.5. A summary is provided in Section 4.6.

The work described in this chapter forms part of a submitted journal paper [10].[1]

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

4.2 DETECTION OVERVIEW

Signal detection refers to the process of identifying the presence of a desired signal within a received signal possibly consisting of multiple signals, noise, and interference [43]. The basics of detection and hypothesis testing are given in Section 4.2.1, while models for the signal and noise considered in this work are given in Section 4.2.2, and metrics for characterising detection algorithm performance are given in Section 4.2.3.

4.2.1 Hypothesis Testing

Assume that there are N samples of a received signal which have the following form

$$y(n) = s(n) + w(n), \quad n = 0, \ldots, N-1 \tag{4.1}$$

where $s(n)$ is the signal to be detected and $w(n)$ is the accompanying interference, referred to here as noise. The process of signal detection seeks to distinguish whether or not there is a signal present in $y(n)$ (i.e. whether $y(n) = s(n) + w(n)$ or $y(n) = w(n)$) [3, 43]. This process can be modelled as a binary hypothesis testing problem [3]

$$\begin{cases} \mathcal{H}_0 \ : \quad y(n) = w(n) \\ \mathcal{H}_1 \ : \quad y(n) = s(n) + w(n) \end{cases} \tag{4.2}$$

which involves distinguishing between two hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$, where $\mathcal{H}_0$ denotes the signal being absent and $\mathcal{H}_1$ denotes the signal being present. The objective of a detection algorithm is to make a decision between these two hypotheses based on the received signal $y(n)$ [3, 43]. This is typically achieved by comparing a test statistic $T(\mathbf{x})$ – computed for a

data vector $\mathbf{x}$ – to a decision threshold γ, expressed as

$$T(\mathbf{x}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.3}$$

where $\mathcal{H}_1$ is decided if the test statistic exceeds the threshold, and $\mathcal{H}_0$ is decided otherwise.

The formulation of the test statistic and threshold is dependant on the information about the received signal $y(n)$ that is available. Due to the fact that the desired signal $s(n)$ is corrupted by some unknown noise $w(n)$, the received signal $y(n)$ is then a random variable (RV) with a certain probability density function (PDF) under each hypothesis $\mathcal{H}_0$ and $\mathcal{H}_1$ [3]. The information available about the PDF under each hypothesis, specifically whether the parameters of the PDF are known or unknown, determines the types of hypotheses being dealt with, which in turn dictates the type of hypothesis testing that is performed. Hypotheses can be separated into two types: simple, and composite [3].

Simple Hypotheses

In the case where there is complete knowledge of the PDF under a hypothesis, that hypothesis is said to be simple. The Neyman-Pearson (NP) theorem states that for a binary hypothesis test involving two simple hypotheses, the test statistic that maximises the probability of detection is the likelihood ratio test (LRT), defined for a data vector $\mathbf{x}$ as [3]

$$L(\mathbf{x}) = \frac{p(\mathbf{x};\mathcal{H}_1)}{p(\mathbf{x};\mathcal{H}_0)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.4}$$

where $p(\cdot)$ denotes the PDF. The test decides $\mathcal{H}_1$ when $L(\mathbf{x})$ exceeds a threshold γ, and $\mathcal{H}_0$ otherwise.

Composite Hypotheses

In the case where certain parameters of the PDF under a hypothesis are unknown, that hypothesis is said to be composite. When certain parameters of the signal and/or noise are unknown, parts of the PDFs become unknown, and the related hypotheses become composite hypotheses [3]. There are two typical approaches to composite hypothesis testing: the Bayesian approach, which considers the unknown parameters as RVs and assigns a prior PDF to them, and the generalised likelihood ratio test (GLRT) approach, which estimates the unknown parameters for use in a LRT [3]. Since the Bayesian approach requires prior information about the unknown parameters while the GLRT does not, the GLRT is used in this work.

In general, for a data vector $\mathbf{x}$ the GLRT is defined as

$$L_G(\mathbf{x}) = \frac{p\left(\mathbf{x}; \hat{\theta}_1, \mathcal{H}_1\right)}{p\left(\mathbf{x}; \hat{\theta}_0, \mathcal{H}_0\right)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.5}$$

where $\hat{\theta}_i$ is the maximum likelihood estimate (MLE) of θ_i, or the value that maximises $p\left(\mathbf{x}; \hat{\theta}_i, \mathcal{H}_i\right)$, with θ_i being the parameter vector under the i^{th} hypothesis [3].

A simplified version of the GLRT exists for the case where the PDF under both hypotheses is the same, given by $p(\mathbf{x}; \theta)$, except that the value of the unknown parameter vector θ is different, and there are no nuisance parameters [3].[2] In this case the detection problem becomes

$$\begin{cases} \mathcal{H}_0 : & \theta = \theta_0 \\ \mathcal{H}_1 : & \theta \neq \theta_0 \end{cases} \tag{4.6}$$

[2] A nuisance parameter is one that is unknown but is of no interest, i.e. it does not factor into the detection problem.

and the GLRT is then defined as [3]

$$L_G(\mathbf{x}) = \frac{p\left(\mathbf{x}; \hat{\theta}_1, \mathcal{H}_1\right)}{p\left(\mathbf{x}; \hat{\theta}_0, \mathcal{H}_0\right)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.7}$$

where $\hat{\theta}_1$ is the MLE under $\mathcal{H}_1$ [3].

4.2.2 Signal Models

As mentioned previously, the information available and the assumptions made about the signal and the noise will heavily influence the formation of the test statistic and threshold of a detection algorithm. Potential models for the signal and the noise are laid out here.

For both the signal and the noise, either real or complex data samples can be used. The type of data samples used should have an effect on the performance of a detector, considering that complex-valued samples inherently have twice as much data as real-valued samples (although this may not necessarily translate to twice as much information) [59]. For this reason both cases are considered in this work, for both the signal and the noise.[3]

The noise can either be Gaussian or non-Gaussian, having known and/or unknown parameters. For simplicity, in this work all noise is assumed Gaussian with known parameters, where each noise sample is a zero-mean Gaussian RV, expressed as

$$\begin{cases} \tilde{w}(n) \sim \mathcal{CN}\left(0, \sigma_w^2\right), & \text{for complex data samples} \\ w(n) \sim \mathcal{N}\left(0, \sigma_w^2\right), & \text{for real data samples} \end{cases} \tag{4.8}$$

where $\mathcal{N}$ and $\mathcal{CN}$ represent real and complex Gaussian RVs, respectively, and σ_w^2 is the noise variance [59].

[3] In this work complex variables are represented using the tilde (e.g. $\tilde{x}$).

The signal can be either deterministic or random, having known and/or unknown parameters. Three potential signal cases are considered here:

Known signal case: The signal is assumed to be completely known. This represents the best case scenario, where a copy of the original signal is available to test against the received signal.

Unknown Gaussian signal case: The signal is assumed to be completely random, where each signal sample is a zero-mean Gaussian RV, expressed as

$$\begin{cases} \tilde{s}(n) \sim \mathcal{CN}\left(0, \sigma_s^2\right), & \text{for complex data samples} \\ s(n) \sim \mathcal{N}\left(0, \sigma_s^2\right), & \text{for real data samples} \end{cases} \tag{4.9}$$

where σ_s^2 is the signal variance. This represents the worst case scenario, where there is no prior knowledge about the signal.

Unknown deterministic signal case: The signal is assumed to have a known deterministic structure, with certain unknown parameters. This represents a scenario somewhere in between the first two cases: some knowledge about the types of signals that will be encountered can be assumed, but the specifics are still unknown.

In this work only one signal model is considered for simplicity: the unknown deterministic signal case. The primary reason for this choice is that this is the type of signal likely to be encountered when surveying an area looking for transmissions (e.g [2]) – a search will be conducted for a specific class of signal, but the specifics of those signals will be unknown. More specifically, a simple narrowband signal model is assumed, where the signal is of the form

$$\begin{cases} \tilde{s}(n) = Ae^{j\psi}e^{j2\pi f_c n T_s}, & \text{for complex data samples} \\ s(n) = \mathrm{Re}\left(Ae^{j\psi}e^{j2\pi f_c n T_s}\right), & \text{for real data samples} \end{cases} \tag{4.10}$$

where the amplitude A, phase ψ, and carrier frequency f_c are all assumed to be unknown. A narrowband signal model is assumed due to the fact that many modern receivers channelise wideband signals into separate narrowband channels [43].

4.2.3 Quantifying Detection Performance

There are two primary metrics for indicating the performance of a detection method: the probability of detection, and the probability of false alarm [3]. The probability of detection, P_D, defines the probability of detecting a signal when it is truly present, and can be written as

$$P_D = P\left(T\left(\mathbf{x}\right) > \gamma \mid \mathcal{H}_1\right) \tag{4.11}$$

where $P(\cdot)$ denotes an event probability. The probability of false alarm, P_{FA}, defines the probability of detecting a signal when it is not actually present, and can be written as

$$P_{FA} = P\left(T\left(\mathbf{x}\right) > \gamma \mid \mathcal{H}_0\right). \tag{4.12}$$

In addition to these two primary metrics, there are two remaining metrics, stated here for completeness: the probability of rejection and the probability of missed detection [3]. The probability of rejection, P_R, defines the probability of not detecting a signal when it is not present, written as

$$P_R = P\left(T\left(\mathbf{x}\right) < \gamma \mid \mathcal{H}_0\right) \tag{4.13}$$

and the probability of missed detection, P_M, defines the probability of not detecting a signal when it is present, written as

$$P_M = P\left(T\left(\mathbf{x}\right) < \gamma \mid \mathcal{H}_1\right). \tag{4.14}$$

Expressions for the probabilities of false alarm and detection can be determined from the PDFs of the test statistic under each hypothesis $\mathcal{H}_0$ and $\mathcal{H}_1$. Once the PDFs $p\left(T\left(\mathbf{x}\right);\mathcal{H}_0\right)$ and $p\left(T\left(\mathbf{x}\right);\mathcal{H}_1\right)$ are known, the probabilites of false alarm and detection can be determined from the appropriate right-tail probability, expressed here as

$$P_{FA} = \int_{\gamma}^{\infty} p\left(T\left(\mathbf{x}\right);\mathcal{H}_0\right) d\mathbf{x} \tag{4.15}$$

$$P_D = \int_{\gamma}^{\infty} p\left(T\left(\mathbf{x}\right); \mathcal{H}_1\right) d\mathbf{x} \tag{4.16}$$

From (4.15) and (4.16) it is clear that there are three parameters of importance: P_{FA}, P_D, and the threshold γ, all of which are linked together. In order to minimise the probability of false alarm, the threshold should be set rather high, but this will lower the probability of detection, and vice versa. A typical approach in detector design is to determine the threshold value required to achieve a certain set value of P_{FA} [3].

4.3 DISCARDING SAMPLES DURING DETECTION

A mathematical representation of the process of discarding certain data samples is presented here. The result is a detection problem that is cast in terms of the retained samples that can be used to modify existing derivations of test statistics, thresholds, and performance metrics for the chosen signal model.

While a test statistic can be formed using a single received sample, many detection algorithms use multiple samples during the formation of test statistics in order to improve detection performance [43]. Assume that the N samples of the received signal in (4.1) are collected into a $N \times 1$ vector $\mathbf{y}$ of the form

$$\mathbf{y} = \begin{bmatrix} y(0) & \cdots & y(N-1) \end{bmatrix}^T = \mathbf{s} + \mathbf{w} \tag{4.17}$$

where $\mathbf{s}$ and $\mathbf{w}$ are also $N \times 1$ vectors representing the N signal and noise samples, respectively.[4]

Discarding samples – or equivalently, applying a binary window function – can be thought of as forming a new $N \times 1$ vector $\mathbf{y}'$ from the original $N \times 1$ vector $\mathbf{y}$, where only K samples are retained and the remaining sample values are set to zero. This vector can be represented in

[4]For this section real data samples are assumed, but the analysis is identical for complex data samples.

matrix form as

$$\mathbf{y}' = \begin{bmatrix} y'(0) & \cdots & y'(N-1) \end{bmatrix}^T = \mathbf{\Phi}\,\mathbf{y} = \mathbf{\Phi}\,(\mathbf{s}+\mathbf{w}) \tag{4.18}$$

where $\mathbf{\Phi}$ is the $N \times N$ binary sensing matrix, a concept borrowed from the field of compressive sensing (CS). In CS, the sensing matrix can be thought of as a matrix representation of the sampling and compression procedures, discarding unwanted samples and suitably compressing the remaining ones [4, 5]. In this work, only the sampling process is required, and so $\mathbf{\Phi}$ is a binary matrix consisting solely of ones and zeros. This matrix can be formed by using the set of K integers

$$D = \{d_k \in [0, N-1],\ k = 0, 1, \ldots K-1\}, \tag{4.19}$$

which represent the indices of the desired samples in the window, to form a binary window

$$\mathbf{a}_N = \begin{bmatrix} a_0 & a_1 & \cdots & a_{N-1} \end{bmatrix}^T \tag{4.20}$$

where $a_j = 1$ if $j \in D$ and $a_j = 0$ if $j \notin D$, and then form the $N \times N$ matrix

$$\mathbf{\Phi} = \operatorname{diag}(\mathbf{a}_N) = \begin{bmatrix} a_0 & 0 & \cdots & 0 \\ 0 & a_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{N-1} \end{bmatrix}. \tag{4.21}$$

For the conventional case where all N samples are kept, $\mathbf{\Phi} = \mathbf{I}_N$. The formation of the $N \times 1$ vector $\mathbf{y}'$ can then be represented as

$$\mathbf{y}' = \overset{N\times 1}{\begin{bmatrix} y'(0) \\ y'(1) \\ \vdots \\ y'(N-1) \end{bmatrix}} = \overset{N\times N}{\begin{bmatrix} a_0 & 0 & \cdots & 0 \\ 0 & a_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_{N-1} \end{bmatrix}} \times \overset{N\times 1}{\begin{bmatrix} s(0)+w(0) \\ s(1)+w(1) \\ \vdots \\ s(N-1)+w(N-1) \end{bmatrix}}. \tag{4.22}$$

Additionally, there are two matrix products of interest that will be used later on: $\mathbf{\Phi}\mathbf{\Phi}^T$ and $\mathbf{\Phi}^T\mathbf{\Phi}$. Since $\mathbf{\Phi}$ is a diagonal matrix, $\mathbf{\Phi} = \mathbf{\Phi}^T$, and so $\mathbf{\Phi}\mathbf{\Phi}^T = \mathbf{\Phi}^T\mathbf{\Phi} = \mathbf{\Phi}^2$. Additionally, since

$\boldsymbol{\Phi}$ is a binary matrix, $\boldsymbol{\Phi}^2 = \boldsymbol{\Phi}$.

The detection problem in (4.2) can now be recast as

$$\begin{cases} \mathcal{H}_0 : & \mathbf{y} = \boldsymbol{\Phi}\,\mathbf{w} \\ \mathcal{H}_1 : & \mathbf{y} = \boldsymbol{\Phi}\,(\mathbf{s}+\mathbf{w}) = \boldsymbol{\Phi}\,\mathbf{s} + \boldsymbol{\Phi}\,\mathbf{w} \end{cases} \tag{4.23}$$

where $\mathbf{y}'$ has been replaced with $\mathbf{y}$. This is the detection problem that the derived test statistics in this work are based on.

4.4 DETECTOR ANALYSIS

The test statistics, threshold values, and associated detection performance metrics for the unknown deterministic signal case from Section 4.2.2 are presented here. These are derived for both complex-valued data samples as well as real-valued data samples, to illustrate the differences between the two. The test statistics are presented in Section 4.4.1, while the thresholds and performance metrics are presented in Section 4.4.2. All of the presented derivations are based on the associated derivations in [3], slightly modified to account for discarded samples by using the sensing matrix described in Section 4.3. Finally, validation of the derived detection performance metrics is provided in Section 4.4.3.

Before proceeding, a quick note on the assumptions made for all considered Gaussian RVs. For a vector of Gaussian RVs, the elements can either be uncorrelated – where the elements are statistically independent and identically distributed – or correlated. For each of the detectors that follow, the Gaussian RVs within each vector are assumed to be correlated, and thus each vector is assumed to be of the form $\mathbf{x} \sim (\mathbf{0}, \mathbf{C})$, where $\mathbf{C}$ is the covariance matrix. Following the derivation of each generalised test statistic, a special case is then presented where the Gaussian vectors are assumed to be uncorrelated, or white.

4.4.1 Test Statistics

For an unknown deterministic signal, unlike the known signal case and the unknown Gaussian signal case, there is no optimal test statistic due to the unknown parameters, for either real or complex data samples [3]. For this signal model $\mathcal{H}_1$ is now a composite hypothesis, which means that the LRT cannot be used to derive an optimal test statistic; instead, the GLRT is used to derive a sub-optimal test statistic.

The test statistics for complex-valued samples and real-valued samples are presented here, first for arbitrary noise covariance, and then for white noise.

Complex Samples

Assume that the complex signal vector $\widetilde{\mathbf{s}} = \widetilde{\mathbf{h}}\,\widetilde{\mathrm{C}}$, where $\widetilde{\mathrm{C}}$ is a deterministic complex scalar value with unknown parameters and

$$\widetilde{\mathbf{h}} = \begin{bmatrix} 1 & e^{j2\pi f_c T_s} & \dots & e^{j2\pi f_c (N-1) T_s} \end{bmatrix}^T \tag{4.24}$$

is the temporal steering vector, and that the complex noise vector $\widetilde{\mathbf{w}} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C}_{\widetilde{w}})$, where the noise covariance $\mathbf{C}_{\widetilde{w}}$ is known. The detection problem in (4.23) can then be reformulated as

$$\begin{cases} \mathcal{H}_0: & \widetilde{\mathbf{y}} = \mathbf{\Phi}\,\widetilde{\mathbf{w}} \\ \mathcal{H}_1: & \widetilde{\mathbf{y}} = \mathbf{\Phi}\left(\widetilde{\mathbf{h}}\,\widetilde{\mathrm{C}} + \widetilde{\mathbf{w}}\right) = \mathbf{\Phi}\,\widetilde{\mathbf{h}}\,\widetilde{\mathrm{C}} + \mathbf{\Phi}\,\widetilde{\mathbf{w}} \end{cases} \tag{4.25}$$

which can be further simplified to deciding $\mathcal{H}_0$ if $\widetilde{\mathrm{C}} = 0$ and deciding $\mathcal{H}_1$ if $\widetilde{\mathrm{C}} \neq 0$. Due to the unknown parameters in $\widetilde{\mathrm{C}}$ the GLRT in (4.7) is used to derive the test statistic, which in this

case decides $\mathcal{H}_1$ if

$$L\left(\widetilde{\mathbf{y}}\right) = \frac{p\left(\widetilde{\mathbf{y}}; \hat{\widetilde{\mathrm{C}}}_1\right)}{p\left(\widetilde{\mathbf{y}}; \widetilde{\mathrm{C}} = 0\right)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.26}$$

where $\hat{\widetilde{\mathrm{C}}}_1$ is the MLE of $\widetilde{\mathrm{C}}$ under $\mathcal{H}_1$, or the value that maximises

$$p\left(\widetilde{\mathbf{y}}; \mathcal{H}_1\right) = \frac{1}{\pi^N \det\left(\mathbf{C}_{\widetilde{w}}\right)} \exp\left(-\left(\widetilde{\mathbf{y}} - \mathbf{\Phi}\,\widetilde{\mathbf{h}}\,\widetilde{\mathrm{C}}\right)^H \mathbf{C}_{\widetilde{w}}{}^{-1}\left(\widetilde{\mathbf{y}} - \mathbf{\Phi}\,\widetilde{\mathbf{h}}\,\widetilde{\mathrm{C}}\right)\right), \tag{4.27}$$

where the superscript H denotes the Hermitian of a matrix. This has been shown to be given by [60]

$$\hat{\widetilde{\mathrm{C}}}_1 = \left(\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\mathbf{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1} \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\widetilde{\mathbf{y}}\,. \tag{4.28}$$

Substituting (4.27) (with $\widetilde{\mathrm{C}}$ replaced by $\hat{\widetilde{\mathrm{C}}}_1$) and the PDF under $\mathcal{H}_0$, given by

$$p\left(\widetilde{\mathbf{y}}; \mathcal{H}_0\right) = \frac{1}{\pi^N \det\left(\mathbf{C}_{\widetilde{w}}\right)} \exp\left(-\widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\widetilde{\mathbf{y}}\right), \tag{4.29}$$

into (4.26) and then taking the logarithm of both sides multiplied by 2 results in

$$2\ln\left(L\left(\widetilde{\mathbf{y}}\right)\right) = -2\left[\left(\widetilde{\mathbf{y}} - \mathbf{\Phi}\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1\right)^H \mathbf{C}_{\widetilde{w}}{}^{-1}\left(\widetilde{\mathbf{y}} - \mathbf{\Phi}\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1\right) - \widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}}\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} 2\ln(\gamma) \tag{4.30}$$

$$\begin{aligned} &= -2\left[\widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}} - \widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\mathbf{\Phi}\,\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1 - \hat{\widetilde{\mathrm{C}}}_1{}^H \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}} + \right. \\ &\quad \left. \hat{\widetilde{\mathrm{C}}}_1{}^H \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\mathbf{\Phi}\,\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1 - \widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}}\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \end{aligned} \tag{4.31}$$

$$\begin{aligned} &= -2\left[-\widetilde{\mathbf{y}}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\mathbf{\Phi}\,\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1 - \hat{\widetilde{\mathrm{C}}}_1{}^H \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}} + \right. \\ &\quad \left. \hat{\widetilde{\mathrm{C}}}_1{}^H \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\mathbf{\Phi}\,\widetilde{\mathbf{h}}\hat{\widetilde{\mathrm{C}}}_1\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \end{aligned} \tag{4.32}$$

Noting that

$$\mathbf{C}_{\widetilde{w}}{}^{-1}\mathbf{\Phi}\,\widetilde{\mathbf{h}}\left(\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\mathbf{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1} \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H = \mathbf{\Phi}\,\widetilde{\mathbf{h}}\left(\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1}\,\mathbf{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1} \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \mathbf{C}_{\widetilde{w}}{}^{-1} = \mathbf{I}_N \tag{4.33}$$

and then reformatting (4.32) results in

$$2\ln\left(L\left(\widetilde{\mathbf{y}}\right)\right) = -2\left[-\widetilde{\mathbf{y}}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\left(\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\,\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1}\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1 - \hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\left(\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\,\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1}\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}} + \hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.34}$$

which, when making use of (4.28), is shown to reduce to

$$2\ln\left(L\left(\widetilde{\mathbf{y}}\right)\right) = -2\left[-\hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1 - \hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1 + \hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.35}$$

which finally results to

$$T\left(\widetilde{\mathbf{y}}\right) = 2\hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.36}$$

For white noise with $\mathbf{C}_{\widetilde{w}} = \sigma_w^2\,\boldsymbol{\Phi}\boldsymbol{\Phi}^H = \sigma_w^2\,\boldsymbol{\Phi}$, the test statistic becomes

$$T\left(\widetilde{\mathbf{y}}\right) = \frac{\hat{\hat{\mathrm{C}}}_1{}^H\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\widetilde{\mathbf{h}}\hat{\hat{\mathrm{C}}}_1}{\sigma_w^2/2} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.37}$$

For the specific deterministic structure in (4.24), $\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\widetilde{\mathbf{h}}$ is the sum of the magnitudes of the K retained elements in $\widetilde{\mathbf{h}}$, which are all equal to one, resulting in $\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\widetilde{\mathbf{h}} = K$. Combining this with the fact that $\hat{\hat{\mathrm{C}}}_1$ is a complex scalar results in

$$T\left(\widetilde{\mathbf{y}}\right) = \frac{K\left|\hat{\hat{\mathrm{C}}}_1\right|^2}{\sigma_w^2/2} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.38}$$

Substituting (4.28) into (4.38) results in

$$T\left(\widetilde{\mathbf{y}}\right) = \frac{2K}{\sigma_w^2}\left|\left(\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\,\boldsymbol{\Phi}\,\widetilde{\mathbf{h}}\right)^{-1}\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}}\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.39}$$

$$= \frac{2K}{\sigma_w^2}\left|\frac{\mathbf{C}_{\widetilde{w}}\widetilde{\mathbf{h}}^H\boldsymbol{\Phi}^H\mathbf{C}_{\widetilde{w}}{}^{-1}\widetilde{\mathbf{y}}}{K}\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.40}$$

$$= \frac{2}{K\,\sigma_w^2}\left|\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \widetilde{\mathbf{y}}\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.41}$$

$$= \frac{2}{K\,\sigma_w^2}\left|\sum_{n=0}^{N-1} a_n \tilde{y}(n) h^*(n)\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.42}$$

$$= \frac{2}{K\,\sigma_w^2}\left|\sum_{n=0}^{N-1} a_n \tilde{y}(n) e^{-j2\pi f_c n T_s}\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.43}$$

where a_n refers to the weights from the binary window in (4.20). This is the squared magnitude of the discrete Fourier transform (DFT) of the received samples at the carrier frequency f_c. This is a well documented test statistic, known as the short-time periodogram or the spectrogram.

Real Samples

Assume that the real signal vector $\mathbf{s} = \mathbf{h}\,\mathrm{C}$, where C is a real deterministic scalar value with unknown parameters and

$$\mathbf{h} = \mathrm{Re}\left(\begin{bmatrix} 1 & e^{j2\pi f_c T_s} & \cdots & e^{j2\pi f_c (N-1) T_s} \end{bmatrix}^T\right) \tag{4.44}$$

is the real version of the temporal steering vector in (4.24), and that the real noise vector $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_w)$, where the noise covariance $\mathbf{C}_w$ is known. As in the complex-valued sample case the detection problem reduces to deciding $\mathcal{H}_0$ if $\mathrm{C} = 0$ and deciding $\mathcal{H}_1$ if $\mathrm{C} \neq 0$. Once again the GLRT in (4.7) is used, which in this case decides $\mathcal{H}_1$ if

$$L(\mathbf{y}) = \frac{p\left(\mathbf{y}; \hat{\mathrm{C}}_1\right)}{p(\mathbf{y}; \mathrm{C} = 0)} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma. \tag{4.45}$$

where $\hat{\mathrm{C}}_1$ is the value that maximises

$$p(\mathbf{y}; \mathcal{H}_1) = \frac{1}{(2\pi)^{N/2} \det^{1/2}(\mathbf{C}_w)} \exp\left(-\frac{1}{2}(\mathbf{y} - \mathbf{\Phi}\,\mathbf{h}\,\mathrm{C})^T \mathbf{C}_w{}^{-1} (\mathbf{y} - \mathbf{\Phi}\,\mathbf{h}\,\mathrm{C})\right) \tag{4.46}$$

given by

$$\hat{\mathrm{C}}_1 = \left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1}\, \mathbf{y}\,. \tag{4.47}$$

Substituting (4.46) (with C replaced by $\hat{\mathrm{C}}_1$) and the PDF under $\mathcal{H}_0$, given by

$$p\left(\mathbf{y};\mathcal{H}_0\right) = \frac{1}{(2\pi)^{N/2} \det^{1/2}\left(\mathbf{C}_w\right)} \exp\left(-\frac{1}{2}\mathbf{y}^T \mathbf{C}_w{}^{-1}\, \mathbf{y}\right), \tag{4.48}$$

into (4.45) and then taking the logarithm of both sides multiplied by 2 results in

$$2\ln\left(L\left(\mathbf{y}\right)\right) = -\left[\left(\mathbf{y} - \mathbf{\Phi}\mathbf{h}\hat{\mathrm{C}}_1\right)^T \mathbf{C}_w{}^{-1} \left(\mathbf{y} - \mathbf{\Phi}\mathbf{h}\hat{\mathrm{C}}_1\right) - \mathbf{y}^T \mathbf{C}_w{}^{-1}\mathbf{y}\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} 2\ln\left(\gamma\right) \tag{4.49}$$

$$= -\left[\mathbf{y}^T \mathbf{C}_w{}^{-1}\mathbf{y} - \mathbf{y}^T \mathbf{C}_w{}^{-1}\, \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1 - \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1}\mathbf{y} + \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1 - \mathbf{y}^T \mathbf{C}_w{}^{-1}\mathbf{y}\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.50}$$

$$= -\left[-\mathbf{y}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1 - \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1}\mathbf{y} + \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.51}$$

As with the complex-valued case, noting that

$$\mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h} \left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T = \mathbf{\Phi}\, \mathbf{h} \left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} = \mathbf{I}_N \tag{4.52}$$

and then reformatting (4.51) results in

$$2\ln\left(L\left(\mathbf{y}\right)\right) = -\left[-\mathbf{y}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h} \left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1 - \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h} \left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1}\mathbf{y} + \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1\right] \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.53}$$

which is shown to reduce to

$$T\left(\mathbf{y}\right) = \hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\hat{\mathrm{C}}_1 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.54}$$

For white noise with $\mathbf{C}_w = \sigma_w^2 \mathbf{\Phi}\mathbf{\Phi}^T = \sigma_w^2 \mathbf{\Phi}$, the test statistic becomes

$$T(\mathbf{y}) = \frac{\hat{\mathrm{C}}_1{}^T \mathbf{h}^T \mathbf{\Phi}^T \mathbf{h} \hat{\mathrm{C}}_1}{\sigma_w^2} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.55}$$

Once again making use of the fact that $\mathbf{h}^T \mathbf{\Phi}^T \mathbf{h} = K$ and $\hat{\mathrm{C}}_1$ is a real scalar results in

$$T(\mathbf{y}) = \frac{K\left(\hat{\mathrm{C}}_1\right)^2}{\sigma_w^2} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma'. \tag{4.56}$$

Substituting in (4.47) results in

$$T(\mathbf{y}) = \frac{K}{\sigma_w^2}\left(\left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{\Phi}\, \mathbf{h}\right)^{-1} \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{y}\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.57}$$

$$= \frac{K}{\sigma_w^2}\left(\frac{\mathbf{C}_w \mathbf{h}^T \mathbf{\Phi}^T \mathbf{C}_w{}^{-1} \mathbf{y}}{K}\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.58}$$

$$= \frac{1}{K\sigma_w^2}\left(\mathbf{h}^T \mathbf{\Phi}^T \mathbf{y}\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.59}$$

$$= \frac{1}{K\sigma_w^2}\left(\sum_{n=0}^{N-1} a_n y(n) h(n)\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.60}$$

$$= \frac{1}{K\sigma_w^2}\mathrm{Re}\left(\sum_{n=0}^{N-1} a_n y(n) h^*(n)\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.61}$$

$$= \frac{1}{K\sigma_w^2}\mathrm{Re}\left(\sum_{n=0}^{N-1} a_n y(n) e^{-j2\pi f_c n T_s}\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma' \tag{4.62}$$

where a_n refers to the weights from the binary window in (4.20). This is the square of the real part of the DFT of the received samples at the carrier frequency f_c.

4.4.2 Detection Performance

For the test statistics presented in Section 4.4.1, the resulting detection performance is now quantified, using expressions for the probability of false alarm P_{FA} and the probability of

detection P_D. For each test statistic expressions for P_{FA} and P_D are given, and from those further expressions for the threshold γ required to achieve a desired value of P_{FA} and the receiver operating characteristic (ROC) (P_D in terms of P_{FA}) are derived. As with the test statistics, the detection performance for both complex-valued data samples and real-valued data samples is presented.

Complex Samples

Assume the test statistic from (4.41), where the noise is assumed white, as

$$T(\widetilde{\mathbf{y}}) = \frac{2}{K\sigma_w^2}\left|\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \widetilde{\mathbf{y}}\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.63}$$

$$= \frac{2}{K\sigma_w^2}\left|\sum_{n=0}^{N-1} a_n \tilde{y}(n) h^*(n)\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.64}$$

$$= \frac{2}{K\sigma_w^2}\left|\sum_{k=0}^{K-1} \tilde{y}(d_k) h^*(d_k)\right|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma, \quad d_k \in D. \tag{4.65}$$

$T(\widetilde{\mathbf{y}})$ is the squared magnitude of the sum of K complex Gaussian RVs, since $\widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \widetilde{\mathbf{y}}$ is a linear transformation of $\widetilde{\mathbf{y}}$, where

$$\widetilde{\mathbf{y}} \sim \begin{cases} \mathcal{CN}\left(\mathbf{0}, \sigma_w^2 \mathbf{\Phi}\mathbf{\Phi}^H\right) & \text{under } \mathcal{H}_0 \\ \mathcal{CN}\left(\mathbf{\Phi}\widetilde{\mathbf{h}}\widetilde{\mathrm{C}}, \sigma_w^2 \mathbf{\Phi}\mathbf{\Phi}^H\right) & \text{under } \mathcal{H}_1\,. \end{cases} \tag{4.66}$$

Since a sum of K Gaussian RVs is itself a Gaussian RV, the test statistic in (4.65) can alternatively be expressed in terms of the complex Gaussian RV $\tilde{z}$ as

$$T(\widetilde{\mathbf{y}}) = \frac{2}{K\sigma_w^2}|\tilde{z}|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.67}$$

$$= \frac{2}{K\sigma_w^2}\left(z_r{}^2 + z_i{}^2\right) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.68}$$

where z_r and z_i are the real and imaginary parts of $\tilde{z}$, which means that $T(\widetilde{\mathbf{y}})$ is now the sum of 2 squared real Gaussian RVs. Therefore, $T(\widetilde{\mathbf{y}})$ follows a central chi-squared distribution with 2 degrees of freedom under $\mathcal{H}_0$, and a non-central chi-squared distribution with 2 degrees of freedom under $\mathcal{H}_1$ [3], expressed as

$$T(\widetilde{\mathbf{y}}) \sim \begin{cases} \chi_2^2 & \text{under } \mathcal{H}_0 \\ \chi_2'^2(\eta) & \text{under } \mathcal{H}_1, \end{cases} \tag{4.69}$$

where η is the noncentrality parameter given by

$$\eta = \frac{\hat{\widetilde{\mathbf{C}}}_1^H \widetilde{\mathbf{h}}^H \mathbf{\Phi}^H \widetilde{\mathbf{h}} \hat{\widetilde{\mathbf{C}}}_1}{\sigma_w^2/2} \tag{4.70}$$

$$= \frac{K}{\sigma_w^2/2} \left| \hat{\widetilde{\mathbf{C}}}_1 \right|^2 \tag{4.71}$$

$$= 2K\text{SNR}. \tag{4.72}$$

Using the right-tail probabilities $Q_{\chi_2^2}(\gamma)$ and $Q_{\chi_2'^2(\eta)}(\gamma)$, P_{FA} and P_D can be expressed as

$$P_{FA} = Q_{\chi_2^2}(\gamma) = 1 - F_{\chi_2^2}(\gamma) \tag{4.73}$$

$$P_D = Q_{\chi_2'^2(\eta)}(\gamma) = 1 - F_{\chi_2'^2(\eta)}(\gamma) \tag{4.74}$$

where $F(\cdot)$ represents the cumulative distribution function (CDF) of the distribution. The threshold in terms of P_{FA} is then given by

$$\gamma = F_{\chi_2^2}^{-1}(1 - P_{FA}) \tag{4.75}$$

and the resulting expression for the ROC is then

$$P_D = 1 - F_{\chi_2'^2(\eta)}\left(F_{\chi_2^2}^{-1}(1 - P_{FA})\right). \tag{4.76}$$

Real Samples

Assume the test statistic from (4.59), where the noise is assumed white, as

$$T(\mathbf{y}) = \frac{1}{K\sigma_w^2}\left(\mathbf{h}^T\mathbf{\Phi}^T\mathbf{y}\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.77}$$

$$= \frac{1}{K\sigma_w^2}\left(\sum_{n=0}^{N-1} a_n y(n)h(n)\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma \tag{4.78}$$

$$= \frac{1}{K\sigma_w^2}\left(\sum_{k=0}^{K-1} y(d_k)h(d_k)\right)^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \gamma, \quad d_k \in D. \tag{4.79}$$

$T(\mathbf{y})$ is the square of the sum of K real Gaussian RVs, which is a single squared Gaussian RV, since $\mathbf{h}^T\mathbf{\Phi}^T\mathbf{y}$ is a linear transformation of $\mathbf{y}$, where

$$\mathbf{y} \sim \begin{cases} \mathcal{N}\left(\mathbf{0}, \sigma_w^2\,\mathbf{\Phi}\,\mathbf{\Phi}^T\right) & \text{under } \mathcal{H}_0 \\ \mathcal{N}\left(\mathbf{\Phi}\,\mathbf{h}\,\mathrm{C}, \sigma_w^2\,\mathbf{\Phi}\,\mathbf{\Phi}^T\right) & \text{under } \mathcal{H}_1\,. \end{cases} \tag{4.80}$$

Therefore, $T(\mathbf{y})$ follows a central chi-squared distribution with 1 degree of freedom under $\mathcal{H}_0$, and a non-central chi-squared distribution with 1 degree of freedom under $\mathcal{H}_1$, expressed as

$$T(\mathbf{y}) \sim \begin{cases} \chi_1^2 & \text{under } \mathcal{H}_0 \\ \chi_1'^2(\eta) & \text{under } \mathcal{H}_1, \end{cases} \tag{4.81}$$

where η is the noncentrality parameter given by

$$\eta = \frac{\hat{\mathrm{C}}_1{}^T\mathbf{h}^T\mathbf{\Phi}^T\mathbf{h}\hat{\mathrm{C}}_1}{\sigma_w^2} \tag{4.82}$$

$$= \frac{K}{\sigma_w^2}\left(\hat{\mathrm{C}}_1\right)^2 \tag{4.83}$$

$$= K\mathrm{SNR}. \tag{4.84}$$

Using the right-tail probabilities $Q_{\chi_1^2}(\gamma)$ and $Q_{\chi_1'^2(\eta)}(\gamma)$, P_{FA} and P_D can be expressed as

$$P_{FA} = Q_{\chi_1^2}(\gamma) = 1 - F_{\chi_1^2}(\gamma) \tag{4.85}$$

$$P_D = Q_{\chi_1'^2(\eta)}(\gamma) = 1 - F_{\chi_1'^2(\eta)}(\gamma) \tag{4.86}$$

where $F(\cdot)$ represents the CDF of the distribution. The threshold in terms of P_{FA} is then given by

$$\gamma = F_{\chi_1^2}^{-1}(1 - P_{FA}) \tag{4.87}$$

and the resulting expression for the ROC is then

$$P_D = 1 - F_{\chi_1'^2(\eta)}\left(F_{\chi_1^2}^{-1}(1 - P_{FA})\right). \tag{4.88}$$

4.4.3 Detection Simulations

The performance metrics presented in Section 4.4.2 are validated below using simulations. For both the theoretical and simulated performance samples are discarded using a random binary window with length $V = 100$ and $K = 50$ elements. This window was selected by generating 100 random binary windows with $V = 100$ and $K = 50$ and determining the maximum sidelobe level (SLL) of the resulting frequency response (calculated over $10 \times V$ points) for all V possible cyclic shifts of the window, and then selecting the window which had the lowest maximum SLL for a cyclic shift.

For each test statistic, the process of simulating the detection performance is essentially the same: appropriate vectors for the signal and the noise are generated, after which samples are discarded using the random window. A test statistic is formed and then compared to the calculated threshold, generating a 1 if there is a detection and a 0 if there is not. This process

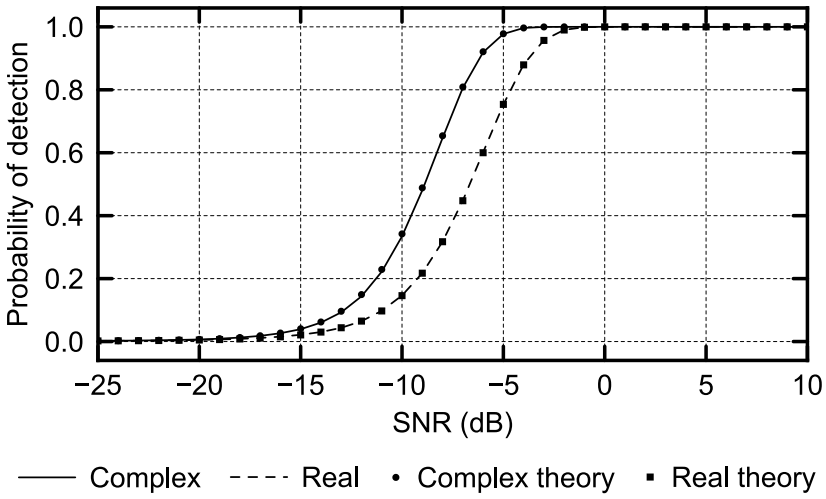

Figure 4.1. Simulated and theoretical probability of detection over SNR curves for the derived test statistics, for $P_{FA} = 0.001$, $V = 100$, and $K = 50$. In each case, the theoretical and simulated performance is practically identical.

is repeated for 1 000 000 windows,[5] and the detections are then summed together to give the probability of detection P_D at one signal-to-noise ratio (SNR) value for the specified value of P_{FA}. Depending on the desired result this entire process is repeated over a range of either SNR or P_{FA} values.

The probability of detection over SNR is given in Figure 4.1, for an SNR range of −25 dB to 10 dB and a P_{FA} value of 0.001. It can be seen that for each simulated test statistic, the resulting P_D curve is so similar to its theoretical counterpart that they are indistinguishable. Additionally, the complex sample case outperforms the real sample case, which is to be expected: for complex-valued data samples, you have twice the amount of data that you do for real-valued samples.

The ROC for each test statistic is given in Figure 4.2, for an SNR value of −10 dB. Figure 4.2(a) shows the ROCs for the scenario where a signal is present (P_D versus the desired P_{FA}). As in Figure 4.1, the simulated ROCs are practically identical to their theoretical counterparts, and the complex sample cases outperform the real sample cases. Figure 4.2(b) shows the ROCs

[5]For the ROC simulations the smallest value of P_{FA} used is 10^{-6}, so 1 000 000 windows should achieve 1 false alarm in the probability calculations. For the simulations where the P_{FA} is set to 0.001, this should achieve an estimated 1 000 false alarms in the probability calculations.

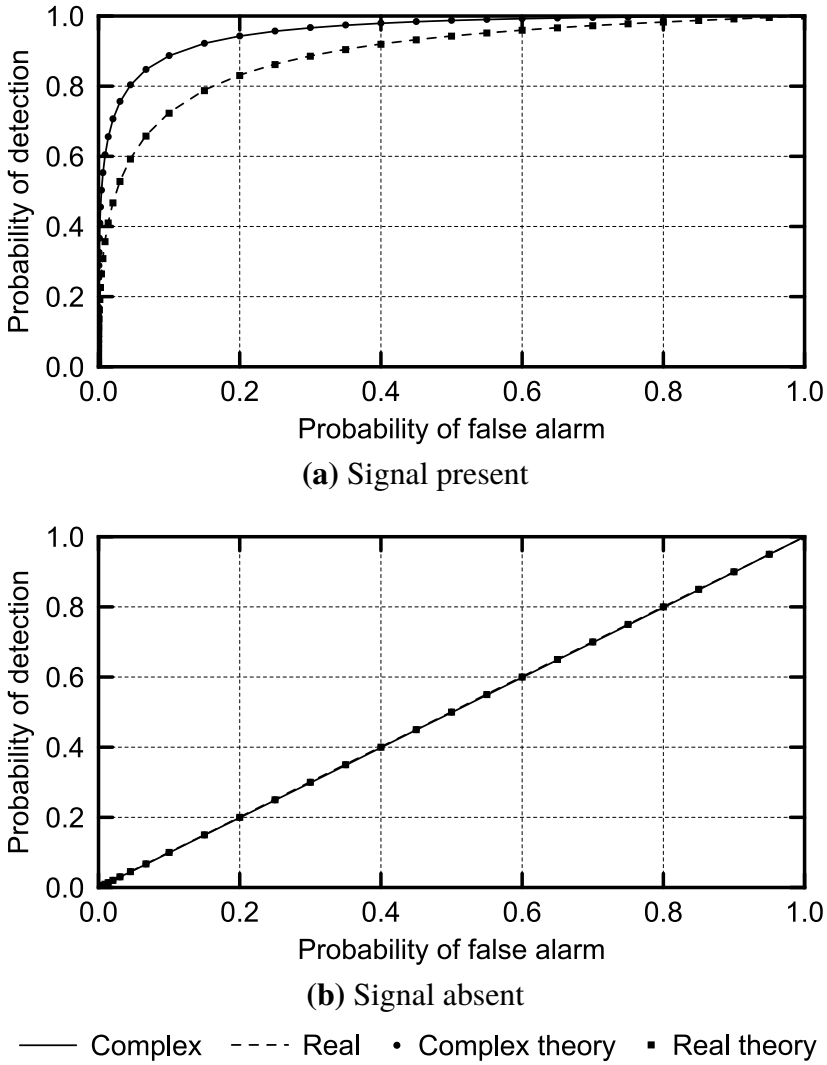

(a) Signal present

(b) Signal absent

Figure 4.2. Simulated and theoretical ROC curves for the derived test statistics, for $V = 100$ and $K = 50$. The absent case helps to verify the simulated probability of false alarm. For all cases SNR $= -10$ dB. In each case, the theoretical and simulated performance is practically identical.

for the scenario where a signal is absent (the simulated P_{FA} versus the desired P_{FA}). Each detection in this case is then in fact a false alarm. All of the ROCs in this case are practically identical, following the desired linear curve which verifies that the simulated P_{FA} is almost indistinguishable from the desired P_{FA}.

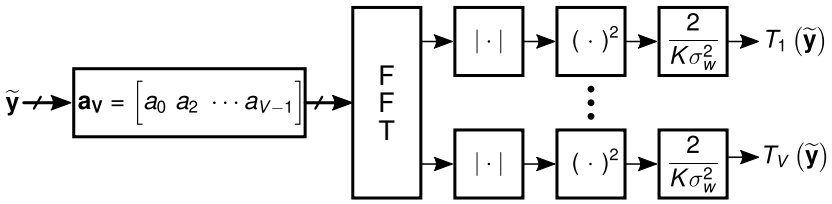

(a) Complex-valued samples

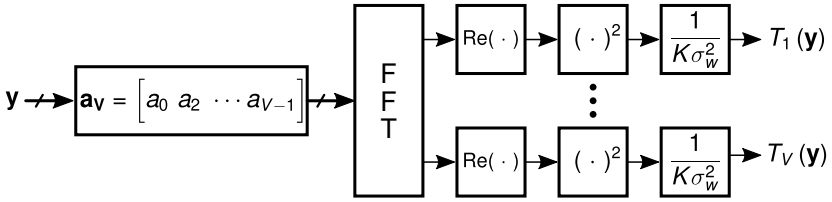

(b) Real-valued samples

Figure 4.3. Practical implementation of the derived test statistics for complex-valued samples and real-valued samples.

4.4.4 Detector Implementation

In Section 4.4.1 it can be seen from (4.43) and (4.62) that for both complex-valued samples and real-valued samples, the derived test statistics are each some variation of the DFT of the received samples at the carrier frequency f_c. However, this requires exact knowledge of the carrier frequency of the signal in (4.10), and since it is already assumed that the amplitude A and the phase ψ are unknown, it would be unreasonable to assume that f_c is somehow known beforehand.

To overcome this, for a sample window of length V a V-point DFT, or more practically a V-point fast Fourier transform (FFT) of the received samples can be taken to produce V filter outputs, where f_c is the frequency at the centre of the FFT bin for each filter output. In the case of complex-valued samples, the squared magnitude of each filter output results in the test statistic in (4.43), and in the case of real-valued samples the square of the real component of each filter output results in the test statistic in (4.62) (each multiplied by the appropriate factors). For both sample types, this approach is shown in Figure 4.3.

This approach has the added benefit of combining the processes of channelisation and detection,

effectively reducing computational complexity compared to having to perform each process separately. This is the final reason that the primary signal model considered in this work is the unknown deterministic signal case. This is also the reason why the sidelobes of the selected binary window function should have an effect on the resulting detection performance: while the test statistic of the frequency bin that the signal is located is only determined by the number of retained samples, for the remaining $V-1$ bins the test statistics will also be affected by the SLL of that particular bin. If there is more than one signal present at a given time, the sidelobes produced by one signal are expected to interfere with those produced by others, and vice versa.

4.5 DETECTION USING DIFFERENCE SETS AND ALMOST DIFFERENCE SETS

Applying DSs and ADSs to the detectors detailed in Section 4.4.4 is a relatively simple procedure. A (V,K,Λ) DS or a (V,K,Λ,t) ADS D can be used to construct the binary window $\mathbf{a}_V$, as in (4.20), with the locations of the K ones in the window function dictated by the elements of the set.

As discussed in Section 4.4.4, implementation of each detector is achieved by taking the product of the window and the samples and applying the appropriate operations shown in Figure 4.3. During these processes a V-point FFT is taken. It is critical that the value of V corresponds to the V of the DS/ADS, because the resulting frequency response only has controlled sidelobes at V sample points, prescribed by (3.19) for a DS and (3.23) for an ADS (see Section 3.3.3). Therefore, if a certain number of frequency bins are required, the DS/ADS should be chosen accordingly.

For each subsequent detection cycle, the sample window will be cyclically shifted by a desired value, as in Figure 3.1. It has already been established in Section 3.3.3 that for any binary

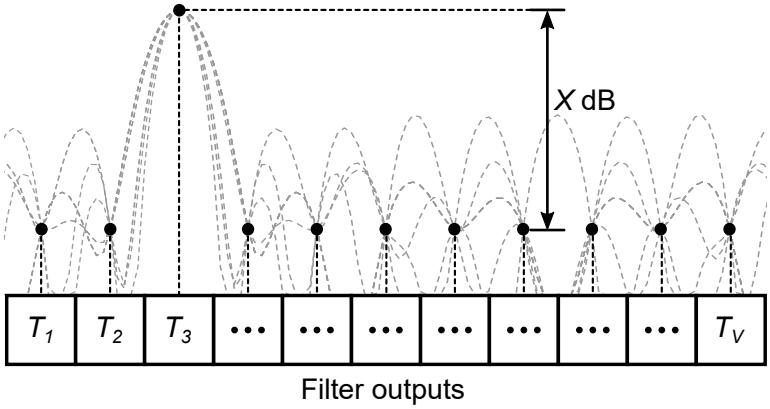

(a) Effect on test statistics

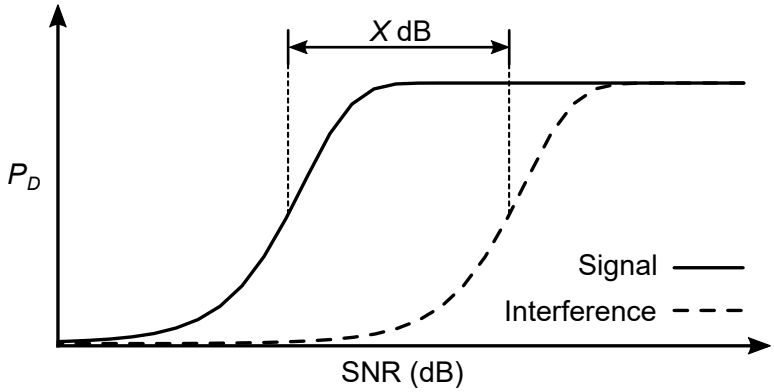

(b) Effect on detection performance

Figure 4.4. The effects of a DS on the derived detector. A signal is located in bin T_3. For the test statistics, the sidelobes of every cyclic shift of the window are plotted. For the detection performance all $V - 1$ interference curves are superimposed onto one another. (Reprinted, with permission, from [10]. ©2018 IEEE.)

window the resulting frequency response will pass through the same sample points irrespective of cyclic shift, and thus the effect on the detection performance should remain the same as the sample window shifts in time.

The effect of the window is depicted in Figure 4.4(a) (for a DS). Assume that there is a signal located in bin T_3 (which is also the test statistic T_3). The remaining $V - 1$ test statistics will be attenuated by the SLL of the corresponding sample point in the frequency response (as they would for any window function). Since the sidelobes do not roll-off as they do for a rectangular window, certain test statistics will actually be amplified compared to the case for a rectangular window, creating interference that could obscure signals in these bins (as is the

case for random binary windows). However, for a window based on a DS or an ADS, the SLLs are more controlled than for the equivalent random binary window.

This means that the impact of this generated interference on the detection performance will also be more controlled. Consider the probability of detection P_D over SNR, shown in Figure 4.1. Each bin in Figure 4.4(a) will have a characteristic for P_D over SNR, the shape of which is identical to the curves in Figure 4.1 (the specific curve is dependent on the sample type); for each bin the resulting curve is shifted in SNR, indicating that for a large enough SNR value you would be able to achieve a detection. The difference in SNR between the curve for the bin that the signal is located in and each remaining bin is given by the SLL value of the bin. For a rectangular window this separation is large, and increases as the window rolls off; for random binary windows and windows based on DSs and ADSs this separation becomes much smaller. However, since the SLLs of a window based on a DS or an ADS are more controlled that those of a random binary window, the resulting interference curves will also be more controlled. Additionally, for a DS the SLL can be calculated directly from the set parameters V, K, and Λ, which is not possible for a random binary window. This concept is illustrated in Figure 4.4(b).

The benefit of consistent interference curves like those shown in Figure 4.4(b) is that one can determine the maximum SNR value required to achieve an acceptable probability of detection in a signal bin, while minimising the probability of detection in an interference bin (much like choosing a threshold value to achieve some P_{FA}). Additionally, using a DS means minimising the detection probability in one interference bin minimises it in all interference bins, which is not possible for the equivalent random window.

4.6 SUMMARY

A mathematical analysis of the process of discarding samples prior to performing signal detection was presented. For a received signal modelled as unknown deterministic in the presence of Gaussian noise, two detectors were derived (one for complex-valued data samples, one for real-valued data samples) for the case where certain samples are discarded using a binary window function.

The application of DSs and ADSs to these derived detectors was also discussed, including the expected effects of the set properties on the detection performance. These expectations are validated by empirical results in Chapter 5.

CHAPTER 5 EMPIRICAL RESULTS

5.1 INTRODUCTORY REMARKS

The performance of a detection system that exploits difference sets (DSs) and almost difference sets (ADSs) for sample selection is considered in this chapter. Empirical results for the complex detector derived in Chapter 4 are presented, including those for both simulated and practical signals. The performance of windows based on DSs and ADSs is compared against that of equivalent random binary windows, to illustrate the benefits that the sets provide. Additionally, the detection performance and computation time of the proposed technique is compared to that of performing detection using a reconstruction algorithm.

The experimental setup, including the signals and sets used and the approaches, is detailed in Section 5.2. The results and subsequent discussion for simulated signals, generated practical signals, and recorded practical signals are provided in Section 5.3, Section 5.4, and Section 5.5, respectively (additional results can be found in Addendum B). A summary is provided in Section 5.6.

The work described in this chapter forms part of a submitted journal paper [10].[1]

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

5.2 EXPERIMENTAL SETUP

Several experiments were conducted in order to evaluate the detection performance achieved when samples are chosen using different binary window functions. These experiments include determining the probability of detection over signal-to-noise ratio (SNR), test statistics at a specific SNR value, and the probability of detection over probability of false alarm to compare the performance achieved using different window functions. Additionally, the detection performance and computation time achieved using the proposed technique is compared to that achieved when performing detection using a reconstruction algorithm. Experiments are evaluated using both simulated and practical signals in order to draw comparisons between them.

The specifics of the signals that are used in the experiments are detailed in Section 5.2.1. The sets and resulting windows that were chosen are described in Section 5.2.2. The exact experiments conducted are summarised in Section 5.2.3 and Section 5.2.4.

5.2.1 Signals

Two major types of signals are evaluated in this work: simulated signals and practical signals. Simulated signals are created and processed entirely in software. Practical signals are captured using a software-defined radio (SDR), and fall into two categories: generated signals, which are created in a laboratory environment and transmitted over-the-air before being received (i.e. custom “real-world" equivalents to the simulated signals), and recorded signals, which are recordings of real-world communications signals.

Simulated Signals

Initial testing was conducted using simulated signals, where both the signal and the noise are generated in software. In this work, a simulated signal is a simple, complex-valued sinusoid of the same form as in (4.24), and the simulated noise is a collection of complex additive white Gaussian noise (AWGN) samples. The simulated signal is multiplied by an appropriate gain value to obtain the desired SNR relative to the calculated noise variance, before being added to the simulated noise to create the final simulated signal.

Generated Practical Signals

In order to provide a "real-world" equivalent to the simulated signals, generated practical signals were created using the setup in Figure 5.1. The transmitter is a Rohde & Schwarz SMX 826.4517.52 signal generator set to produce a simple radio-frequency (RF) carrier with a frequency of 495 MHz. The gain of the signal generator was set to a chosen value in the range of −60 dB to 10 dB (1 dB increments). It is important to note that this signal generator was uncalibrated at the time of use, as calibration is a very costly procedure.

The receiver consists of an Ettus Research Universal Software Radio Peripheral (USRP) N200 SDR [61] fitted with a SBX USRP Daughterboard [62], connected to a computer running Ubuntu 18.04 and GNU Radio. The centre frequency of the USRP was set to 500 MHz, before upconverting in GNU Radio by 5 MHz to shift the centre frequency of the transmitted signal to 500 MHz prior to filtering. This was done to move any leakage from the local oscillator of the USRP out of the band of interest. The bandwidth of the lowpass filter was kept small in order to limit that amount of noise that entered the system. This is due to the fact that the maximum gain that the signal generator can reach is only 10 dB, so in order to obtain larger values of SNR the noise power needed to be kept low. Signals were recorded using this setup

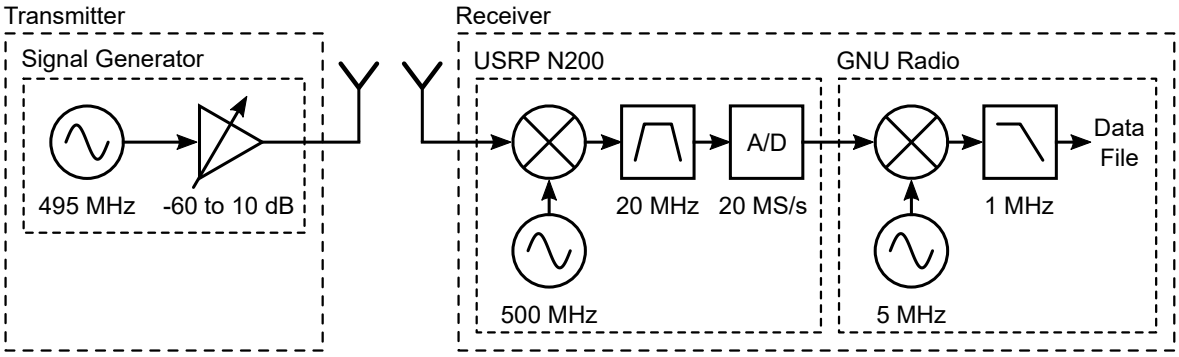

Figure 5.1. Setup for the creation of the generated practical signals.

for the entire gain range of −60 dB to 10 dB. Additionally, reference noise was acquired by switching off the transmitter and terminating the antenna of the receiver. This reference noise was used to determine the noise variance during detection.

Recorded Practical Signals

The generated practical signals are simple RF carriers, so to test the effects of discarding samples on the detection of real-world signals, recorded practical signals were obtained by sampling a portion of a Global System for Mobile Communications (GSM) band using a Nuand bladeRF SDR [63]. GSM signals were chosen since a promising application of this work is detecting cellphones for anti-poaching activities (for use with [2]). During sampling the receiver of the SDR was configured with a centre frequency of 1.748 MHz, a bandwidth of 28 MHz, and a sampling rate of 40 MS/s, which means that the signal was sampled from the uplink portion of the GSM-1800 band. During detection the noise variance was set to 0.01.

5.2.2 Sets and Windows

There are a large number of DSs and ADSs widely available [57, 58], meaning that sets can be selected based on the application at hand (e.g. window length, sidelobe level (SLL), sample thinning). For the experiments in this work (for both simulated signals and practical signals) the two primary sets selected are the DS (400,57,8) and the ADS (400,200,99,100). The sets have the same length but differing numbers of elements: the DS is an example of more drastic thinning (roughly 15% of the samples are retained versus the more common 50% of the ADS), which helps to illustrate the effects of thinning as well as complementary sets. Additionally, since the recorded practical signals in this work are off-the-air GSM signals, these sets were chosen as the 200 kHz channel spacing for GSM [64] means 200 frequency bins are required to achieve a resolution of one channel per bin for the given sampling rate of 40 MS/s. While there are no DSs/ADSs with exactly $V = 200$, there are both a DS and an ADS with exactly $V = 400$.

For completeness, the results for four additional sets are also included: the DSs (31,15,7) and (101,25,6) and the ADSs (30,15,7,22) and (101,26,6,50). These sets help illustrate the effects of even and odd values of window length for DSs/ADSs as well as shorter sets. The results for these sets for the simulated signals and generated practical signals are included in Addendum B; due to their short length they are not applied to the recorded practical signals.

For each (V, K, Λ) DS or (V, K, Λ, t) ADS, four different windows are tested: a window based on the DS/ADS itself, a window based on the complement of the DS/ADS, a filled binary window with $K = V$, and a random binary window where the positions of the K samples are chosen randomly. A filled binary window represents the best case scenario where no samples are discarded (and therefore the best possible detection performance), and random sampling is prevalent in massive thinning of antenna arrays as well as in compressive sensing [6], which is why random binary windows are used for comparison. As in Section 4.4.3, each random window was selected by generating 100 random binary windows with the same length and number of samples (the positions of the sample are randomised for each window) and

determining the maximum SLL of the resulting frequency response (calculated over $10 \times V$ points) for all V possible cyclic shifts of the window, and then selecting the window which had the lowest maximum SLL for a cyclic shift. The frequency responses for all of the windows can be found in Addendum A.

5.2.3 Approach: Detection Performance

For various combinations of signals and windows, the following experiments were conducted to evaluate and compare the detection performance achieved:

Probability of detection over SNR: The probability of detection was calculated over a defined SNR range, for a specified value of P_{FA}.

Test statistics: The test statistics relative to the calculated detection threshold were calculated for specific values of SNR and P_{FA}, for all of the windows that were processed.

Probability of detection over probability of false alarm: The probability of detection was calculated over a defined P_{FA} range, for a specified value of SNR (the same one used for the test statistics). This is alternatively referred to as the receiver operating characteristic (ROC).

The combinations of inputs and sets as well as the relevant experimental parameters for all of the conducted experiments are summarised in Table 5.1. An entry of "N/A" indicates that no experiment was conducted for that combination of input signal and set, such as calculating the probability of detection over SNR using the recorded practical signal, which is not feasible since the recording contains multiple signals that are not continuous, and have different SNR values.

For all of the experiments, the detector derived in Chapter 4 (for complex-valued data samples) was used to form the test statistics and calculate the thresholds. All data processing was

Table 5.1. The signals, sets, and parameters of the experiments conducted. "N/A" indicates that no experiment was conducted for that combination of inputs.

Inputs	Experimental parameters					
Simulated signals	P_D vs SNR		Test statistics		ROC	
	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}
DS (31,15,7) ADS (30,15,7,22)	–30 to 50	0.001	8	0.001	8	10^{-6} to 1
DS (101,25,6) ADS (101,26,6,50)	–30 to 50	0.001	6	0.001	6	10^{-6} to 1
DS (400,57,8) ADS (400,200,99,100)	–30 to 50	0.001	4 3	0.001	4 3	10^{-6} to 1
Generated practical signals	P_D vs SNR		Test statistics		ROC	
	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}
DS (31,15,7) ADS (30,15,7,22)	–30 to 30	0.001	8	0.001	8	10^{-6} to 1
DS (101,25,6) ADS (101,26,6,50)	–30 to 30	0.001	6	0.001	6	10^{-6} to 1
DS (400,57,8) ADS (400,200,99,100)	–30 to 30	0.001	4 3	0.001	4 3	10^{-6} to 1
Recorded practical signals	P_D vs SNR		Test statistics		ROC	
	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}	SNR (dB)	P_{FA}
DS (31,15,7) ADS (30,15,7,22)	N/A	N/A	N/A	N/A	N/A	N/A
DS (101,25,6) ADS (101,26,6,50)	N/A	N/A	N/A	N/A	N/A	N/A
DS (400,57,8) ADS (400,200,99,100)	N/A	N/A	–*	0.001	N/A	N/A

*The recording consists of multiple signals with different SNR values.

performed in MATLAB. The probabilities were determined by computing the test statistics for a number of windows, summing the detections (the ones) for all of the windows and dividing by the total number of windows used. For the probability of detection over false alarm and test statistic experiments, 100 000 windows were used, since for a P_{FA} of 0.001 this results in 100 false alarms. For the probability of detection over probability of false alarm experiments, 10 000 000 windows were used in the probability calculations, due to the fact that the smallest probability of false alarm value was 10^{-6}, so evaluating 10 000 000 windows is estimated to

result in roughly 10 false alarms. While this might seem low, this P_{FA} value is only used in the ROC calculations and represents an extreme value. Finally, during detection consecutive windows were overlapped sample-by-sample (an overlap of 100%).

For the experiments that require a single SNR value, the value was determined by the DS/ADS that the current window is based on: first the theoretical probability of detection over SNR curve is calculated for the specified P_{FA} using (4.76), then the first SNR value SNR_1 where $P_D > 0.99$ is found, and finally the SNR is determined using (3.21) as

$$\text{SNR} = \text{SNR}_1 + \frac{\rho}{2}. \tag{5.1}$$

It is assumed that the sidelobes will affect the probability of detection, so this way the SNR is set halfway between the peak of the mainlobe and the minimum SLL (at least for a DS). Since this value is used for all of the windows based on a specific DS/ADS (including the complementary windows), this makes the performance of each window relative to the others easier to evaluate.

5.2.4 Approach: Comparison with Detection using Reconstruction

For the simulated and generated practical signals considered in this work, the detection performance achieved using the proposed technique is also compared to that achieved when performing detection using a reconstruction algorithm. Additionally, the computation time of both approaches is measured, in order to compare the benefits and limitations of each.

The reconstruction algorithm chosen was the multiple measurement vector FOCal Undetermined System Solver (M-FOCUSS) algorithm, which is an extension of the FOCUSS class of algorithms designed to process multiple measurement vectors at once [65]. M-FOCUSS was chosen for two reasons; in addition to being designed to process multiple measurement vectors, there exists a robust version of the algorithm known as Regularized M-FOCUSS,

which handles noisy data. This should allow for better performance at low SNR values. This robust algorithm involves the selection of a regularization parameter, which in this work was a value chosen in the range 10^{-10} to 10^{-4} that gave the smallest mean square error (MSE) between the original samples and the reconstructed ones for a given SNR. For details on the implementation of the M-FOCUSS algorithm, the reader is directed to [65].

Detection with reconstruction was tested using multiple different sensing matrices: some are variations of the binary windows used for the proposed technique, and others are quite different. The specific types of sensing matrices tested are as follows:

Random Gaussian (RG) matrix: A $K \times V$ matrix whose entries are normalised Gaussian random variables (RVs) with zero-mean and unit variance.

Random binary (RB) matrix: A $K \times V$ matrix, the rows of which are selected randomly from a $V \times V$ identity matrix. The random sequences used are identical to the random windows used when performing detection without reconstruction.

DS matrix: A $K \times V$ matrix, the rows of which are selected from a $V \times V$ identity matrix according to a DS.

Comparison of the two approaches was achieved using two different, but functionally equivalent algorithms, referred to as detection with reconstruction and detection without reconstruction. Both algorithms take a window of V data samples and apply a sensing matrix; the principle difference is that detection with reconstruction then proceeds with only K values (i.e. it uses a $K \times V$ sensing matrix), while detection without reconstruction proceeds with V values, where $V - K$ of these are 0 (i.e. it uses a $V \times V$ sensing matrix). For detection with reconstruction, M-FOCUSS is then used to reconstruct V data samples in the frequency domain using the K retained time domain samples,[2] while for detection without reconstruction a V-point fast Fourier transform (FFT) is performed. Both of these approaches result in V frequency values, after which the necessary operations are performed to produce V test statistics of the form given

[2]Each M-FOCUSS run actually processes L windows of samples at once, where $L < K$.

in (4.43). After processing a desired number of sample windows, for each window the V test statistics are then compared to the appropriate threshold γ to determine any detections.

The detection performance for both approaches is measured as described in Section 5.2.3 (probability of detection over SNR), with a SNR range of −30 dB to 10 dB (1 dB increments) and a P_{FA} of 0.001, except that only 10 000 windows are processed (this is due to long computation times when running all of the sensing matrices, but should still result in 10 false alarms for the desired P_{FA}). The algorithms were run on a computer with two 6-core, 2.30-GHz Intel Xeon E5-2360 central processing units (CPUs) and 30GB of main memory.

For every SNR value, the computation times for both approaches are measured after producing detections from the 10 000 evaluated windows. For each of these times, the average time achieved over the entire SNR range is then calculated.

5.3 RESULTS: SIMULATED SIGNALS

The results for the simulated signals for the windows based on the DS (400,57,8) and the ADS (400,200,99,100) are given here. The results for the DSs (31,15,7) and (101,25,6) and the ADSs (30,15,7,22) and (101,26,6,22) can be found in Addendum B.1 (except for the comparison with detection using reconstruction, where the results are included in this chapter).

5.3.1 Probability of Detection over SNR

The probability of detection over SNR results for the simulated signals are given in Figures 5.2 and 5.3. Figure 5.2 shows the probabilities of detection over SNR for the DS (400,57,8), and

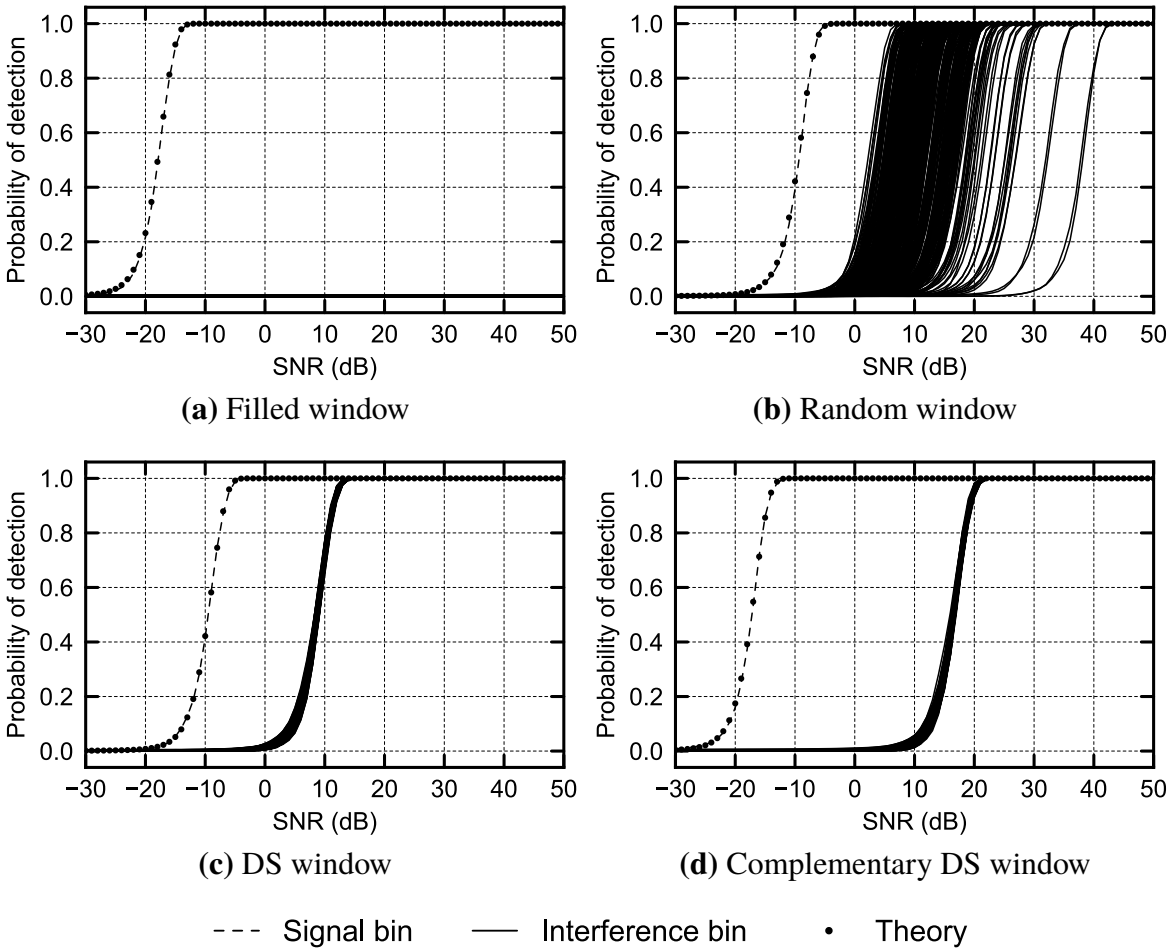

Figure 5.2. Probability of detection over SNR curves for the simulated signals resulting from windows based on the DS (400,57,8). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

Figure 5.3 shows the probabilities of detection over SNR for the ADS (400,200,99,100).

The dashed line in each plot represents the probability of detection over SNR in the frequency bin that the test signal is located in. Each marker represents the theoretical value of P_D at the current SNR for $P_{FA} = 0.001$, given by (4.76); for all of the windows the curves in the signal bin and the theoretical curves are so similar that they are practically indistinguishable. The solid lines represent the probability of detection in the remaining bins as a result of the signal. These will be referred to as interference bins due to the fact that there is no signal present in these bins, making all detections false alarms caused by the sidelobes of the applied window function. These interference false alarms can prevent detection of any weaker signals which

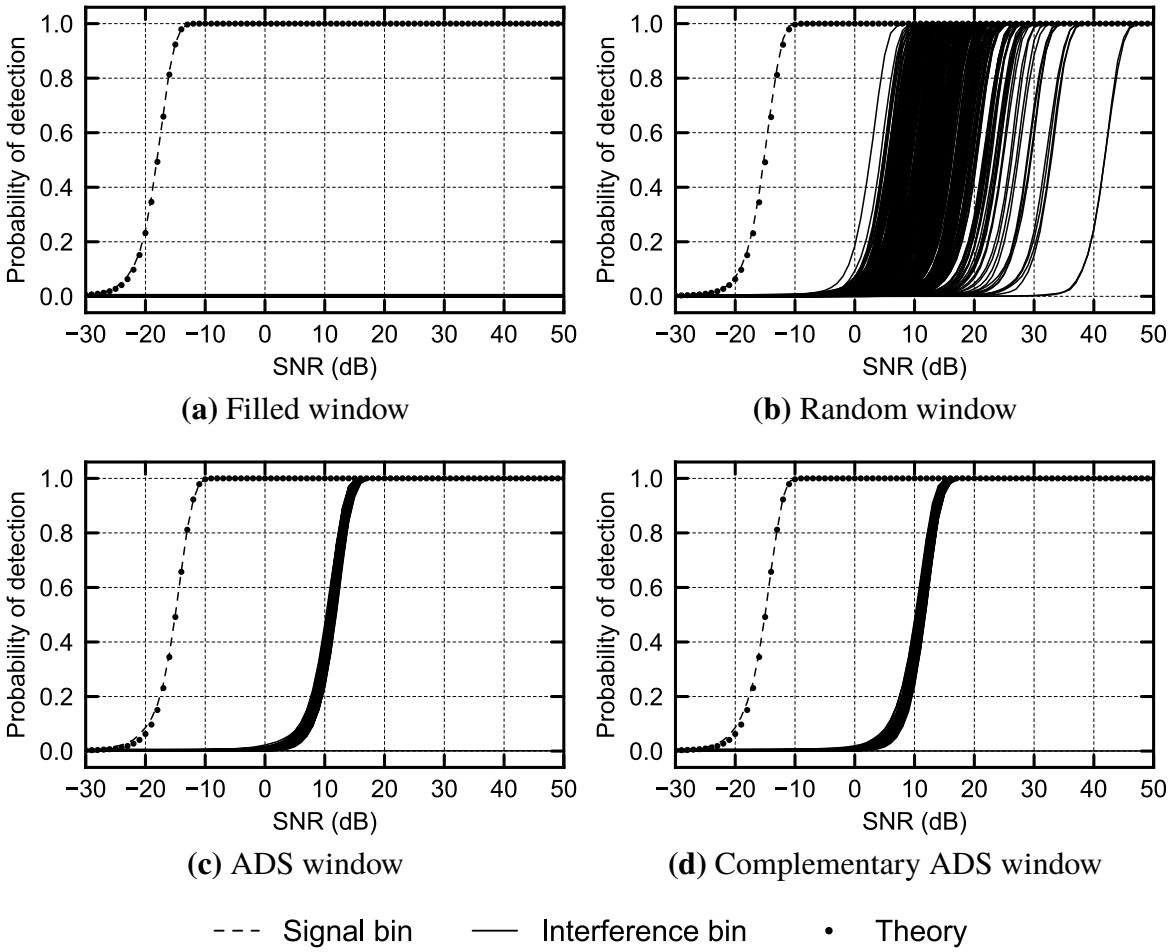

Figure 5.3. Probability of detection over SNR curves for the simulated signals resulting from windows based on the ADS (400,200,99,100). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

may be present in the interference bins because detections in these bins would be attributed to the sidelobes of a stronger signal in another bin. The difference in SNR between the curve for the signal bin and the curve for each interference bin is equal to the normalised SLL value of the sample point that corresponds to that bin in the appropriate frequency response found in Addendum A.

Figures 5.2(a) and 5.3(a) show the best possible detection performance for each set, since all of the samples have been retained (their performance is identical, since $K = V = 400$ in both cases). For these windows the interference bins all have the same probability of detection value, equal to P_{FA}.

For the random and DS/ADS windows of each set, Figures 5.2(b), 5.2(c), 5.3(c), and 5.3(b) show that the performance in the signal bin is identical, which is expected as they all have the same number of retained samples. This means that a random window and a window based on a DS/ADS will perform identically when detecting the same signal in the bin in which the signal is present.

However, the performance in the interference bins varies considerably between the random and DS/ADS windows. Figures 5.2(b) and 5.3(b) show that the random windows produce interference bin curves that are more spread out, and as expected, more random than those from the windows based on DSs and ADSs. While some of the interference bins for the random windows require higher SNR before the window sidelobes trigger detections than the comparable DS/ADS windows (i.e. the effect of the interference is less), other interference bins have sidelobe false alarms at significantly lower SNR than the DS/ADS windows, so weaker signals in these bins could be masked. The DS/ADS windows are thus able to achieve far more predictable and consistent performance across all the detection bins than the random windows.

Something important to note is that because V is even for the ADS (400,200,99,100), the interference bins do not converge to two distinct levels, but are rather somewhat spread out.[3] Additionally a single interference bin appears to be flat, which is probably a result of the null discussed in Section 3.3.3. Fortunately this does not detract from the performance of ADSs with even values of V, seeing as the interference bins are still better controlled than for the equivalent random window, and the interference caused by the null is lower than for the other bins.

Finally, Figures 5.2(d) and 5.3(d) show the effects of complementary sets. While the complement of the ADS (400,200,99,100) is identical to the standard set, and thus the performance is the same, the complement of the DS (400,57,8) produces a larger separation

[3]There are two distinct levels for the ADS (101,26,6,50) where V is odd, as seen in Figure B.2(e) in Addendum B.

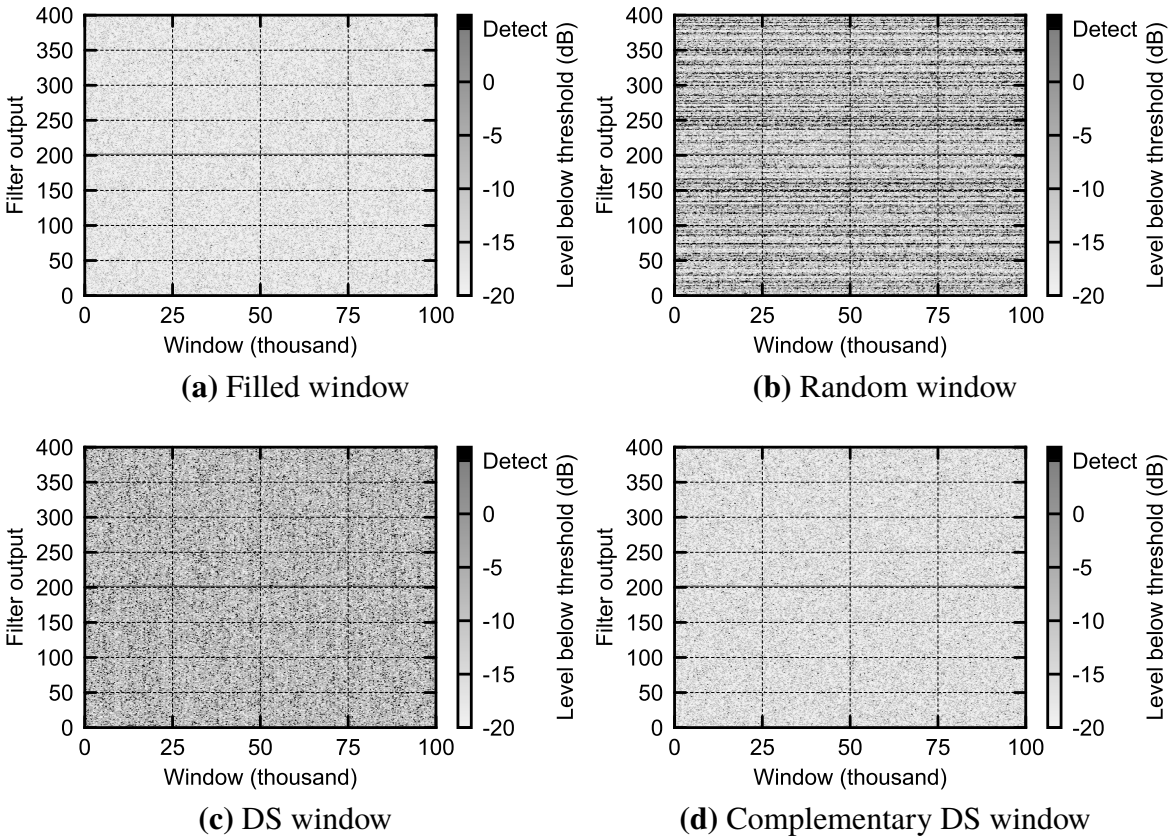

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

Figure 5.4. Test statistics for SNR = 4 dB for the simulated signals resulting from windows based on the DS (400,57,8).

between the curves for the signal bin and the interference bins than the window based on the standard set, as expected since more samples are retained.

5.3.2 Test Statistics

The resulting test statistics for the simulated signals are given in Figures 5.4 and 5.5. Figure 5.4 shows the test statistics for the DS (400,57,8), and Figure 5.5 shows the test statistics for the ADS (400,200,99,100).

For each window the test statistic results are plotted relative to the specific detection threshold

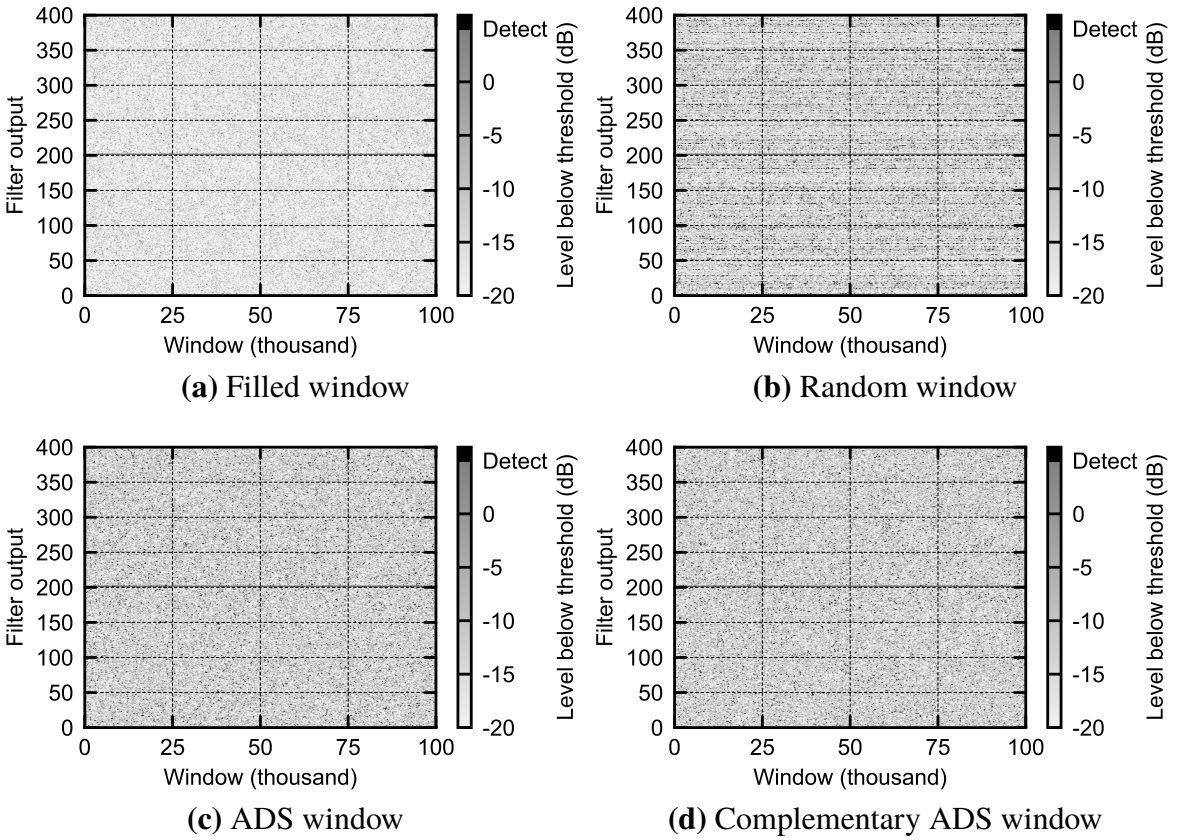

(a) Filled window **(b)** Random window

(c) ADS window **(d)** Complementary ADS window

Figure 5.5. Test statistics for SNR = 3 dB for the simulated signals resulting from windows based on the ADS (400,200,99,100).

γ given by (4.75). For any test statistic that exceeds the threshold, the result is a detection, and is plotted in black. For the windows based on the DS (400,57,8) the SNR is set to 4 dB, and for the windows based on the ADS (400,200,99,100) the SNR is set to 3 dB.

Figures 5.4(a) and 5.5(a) show that for the filled window, the main concentration of detections is in the signal bin. There are some detections in the remaining bins, but these are attributed to actual false alarms, since each bin should have one false alarm for every 1 000 windows for $P_{FA} = 0.001$. As the sidelobes of the windows roll off as expected for a rectangular window, they will only affect a small number of bins to either side of the signal bin.

Figures 5.4(b), 5.4(c), 5.5(b), and 5.5(c) show that for the random and DS/ADS windows the interference from sidelobes has a greater effect. For each window, the signal in the signal bin is still detected, but the sidelobes of the window effectively increase the value of the test

statistics in all of the remaining bins, bringing them closer to the threshold. This has the potential to mask weaker signals in the other bins and prevent them from being detected.

However, the interference from the DS/ADS windows is much more consistent that for the random windows. Figures 5.4(c) and 5.5(c) show that the interference from the sidelobes appears as though the noise floor of the system has been raised. While strong signals can still mask weaker signals with their sidelobes, there do not appear to be any "false targets", which is the case for the random windows. Figures 5.4(b) and 5.5(b) show that the inconsistency of the sidelobes means that certain bins appear to have "false targets" (i.e. attenuated duplicates of the actual signal), which could potentially fool an operator.

For an ADS (with V even or V odd), since the sample points of the frequency response do not converge to one distinct level, the resulting sidelobe interference is not as consistent as for a DS, meaning that inconsistencies between adjacent bins could results in the appearance of "false targets" (this effect can be more easily seen in Figures B.3(e) and B.4(e)). However, it is far less severe than for the equivalent random window.

Finally, Figures 5.4(c) and 5.4(d) show that for a complementary set with more elements than the standard set, the effect of the interference is less severe for the same SNR value.

5.3.3 Probability of Detection over Probability of False Alarm

The probability of detection over probability of false alarm (the ROC) results for the simulated signals are given in Figures 5.6 and 5.7. Figure 5.6 shows the ROCs for the DS (400,57,8), and Figure 5.7 shows the ROCs for the ADS (400,200,99,100).

The plots are similar to the probability of detection over SNR results. The dashed line in each plot represents the ROC in the frequency bin that the test signal is located in. Each

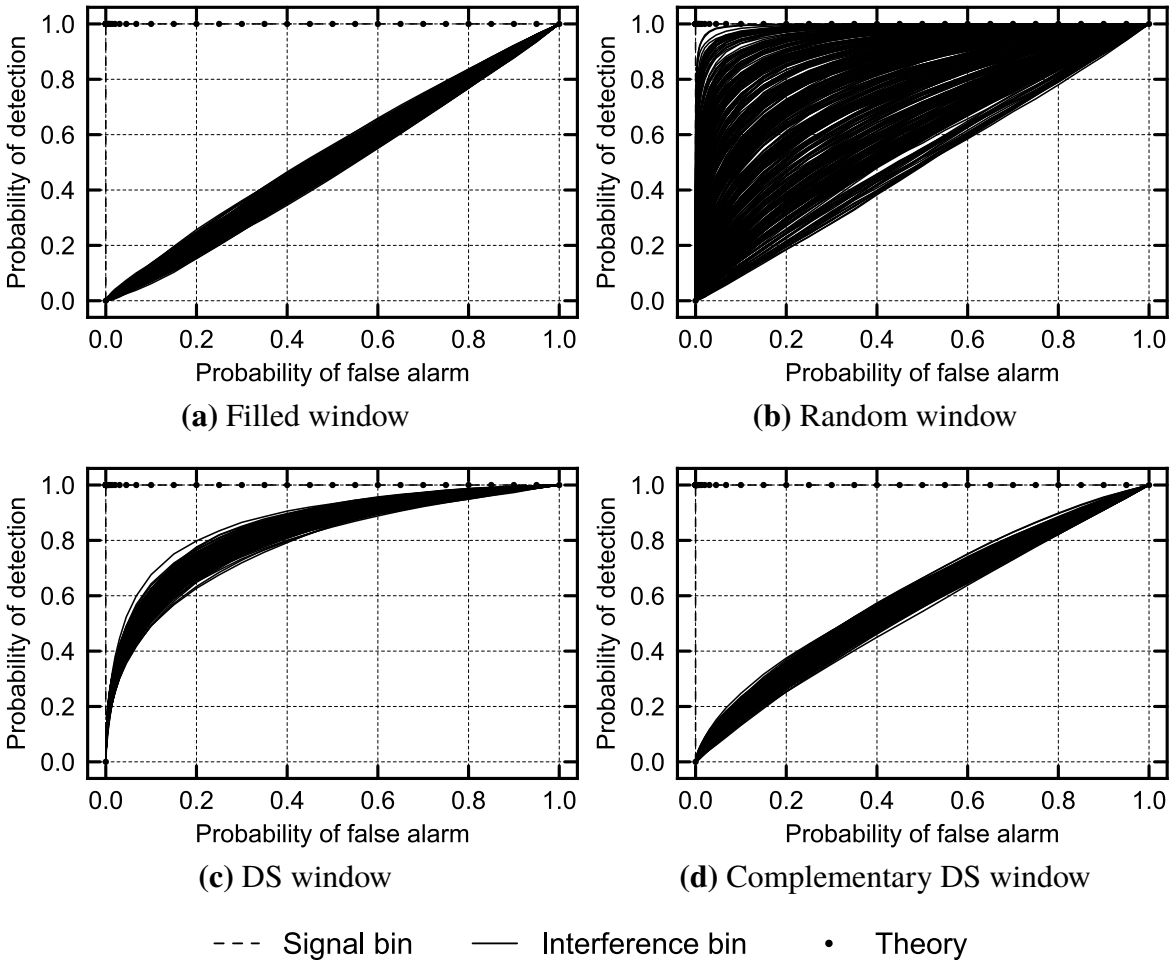

Figure 5.6. ROCs for SNR = 4 dB for the simulated signals resulting from windows based on the DS (400,57,8). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

marker represents the theoretical ROC given by (4.76); for all of the windows the curves in the signal bin and the theoretical curves are so similar that they are practically indistinguishable. The solid lines represent the ROCs in the remaining bins as a result of the signal (i.e. the interference bins). The ROC curves for the signal bin are identical for all windows, because the SNR was chosen high enough that the probability of detection is 1.

Figures 5.6(a) and 5.7(a) show that for the filled windows, the interference curves have a nearly linear characteristic, meaning that $P_D \approx P_{FA}$ for each value of P_{FA}, and any false detections can be attributed to expected false alarms. This means that the value of the threshold can be decreased, in order to allow for detection of weaker signals, and the interference in the

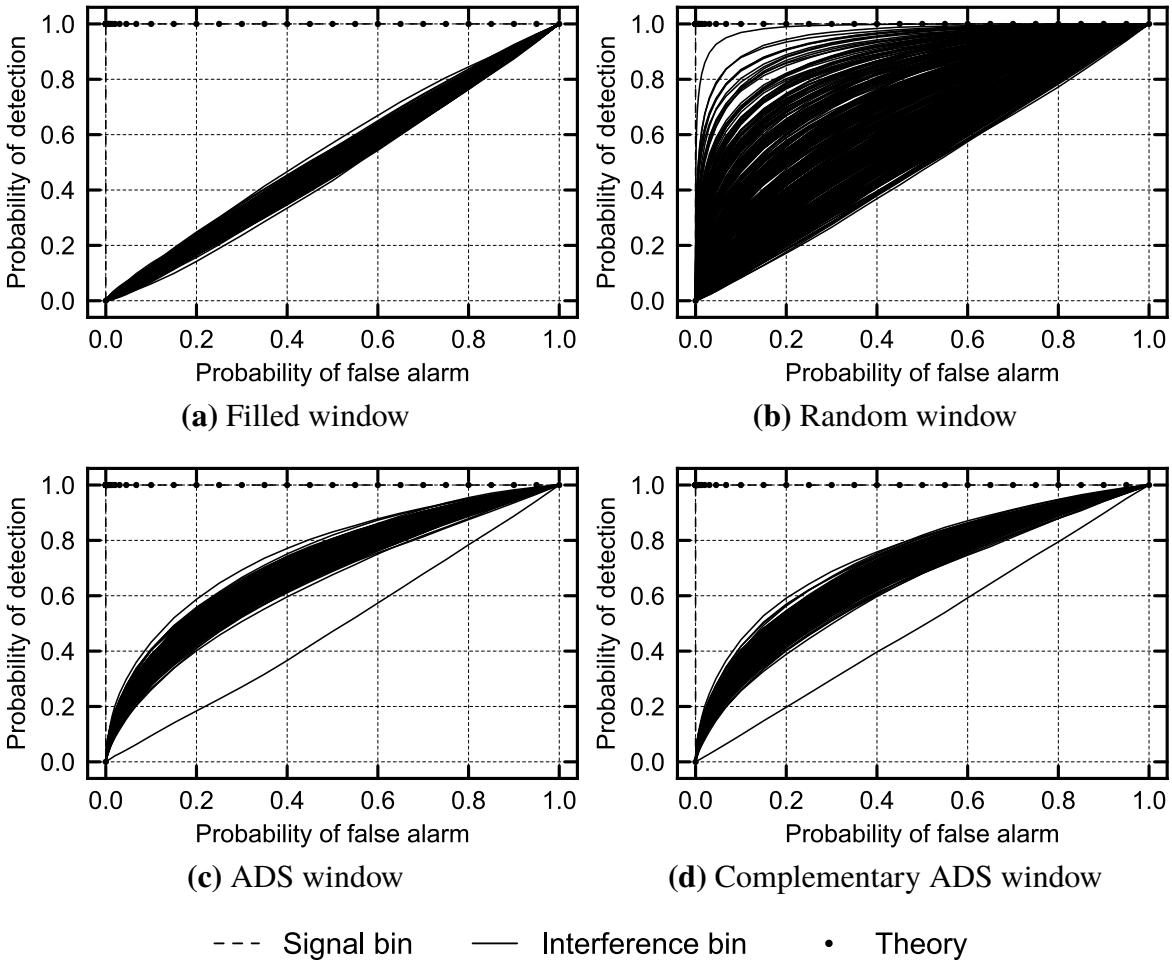

(a) Filled window **(b)** Random window

(c) ADS window **(d)** Complementary ADS window

Figure 5.7. ROCs for SNR = 3 dB for the simulated signals resulting from windows based on the ADS (400,200,99,100). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

remaining bins will increase linearly with the resulting value of P_{FA}.

For the random and DS/ADS windows, Figures 5.6(b), 5.6(c), 5.7(b), and 5.7(c) show that in most of the interference bins these curves deviate sharply from the expected linear trend. For the random windows, while some bins do follow the linear characteristic, most do not, and many reach a saturation point (a P_D value of 1) for very small values of P_{FA}. This means that there is a very limited amount that the probability of false alarm can be increased by in order to improve the detection of weak signals (by consequently lowering the threshold, making it easier to achieve a detection). However, for the DS/ADS windows, while the curves do deviate from the linear trend, it is relatively consistent across bins, and less severely than for many of

the bins for the equivalent random binary window, allowing for a greater increase in P_{FA} to lower the threshold and improve detection performance.

5.3.4 Comparison with Detection using Reconstruction

The results which compare performing detection with reconstruction and detection without reconstruction on the simulated signals are given here. The computation times for the various sensing matrices are summarised in Table 5.2, while the probability of detection over SNR curves are given in Figure 5.8. For detection with reconstruction, results for RG, RB, and DS sensing matrices are included, while for detection without reconstruction, results for RB and DS sensing matrices are included. This is due to the fact that the complex detector derived in Chapter 4 was assumed to be handling binary sensing matrices, not matrices consisting of Gaussian RVs.

Table 5.2. Comparison of the computation times (in seconds) for performing detection with (W) and without (W/O) reconstruction on the simulated signals (Adapted, with permission, from [10]. ©2018 IEEE.).

Parameters $(V, K)^*$	W Recon.			W/O Recon.	
	RG	RB	DS	RB	DS
31, 15	77.719	95.083	91.317	0.0356	0.0362
31, 16	74.623	91.310	76.471	0.0372	0.0362
101, 25	253.48	279.81	298.82	0.1411	0.1394
101, 76	55.290	40.309	39.788	0.1366	0.1371
400, 57	1237.2	1152.2	1212.5	0.5814	0.5808
400, 343	302.55	267.89	259.42	0.5725	0.5681

From Table 5.2, it is clear that the computation times are significantly longer when reconstruction is used, even with M-FOCUSS processing multiple windows at once. Comparing the fastest reconstruction result of 39.788 s to the slowest result for the same set without

reconstruction of 0.1411 s shows that detection with reconstruction is at least 290 times slower than detection without reconstruction (at least for these runs; in [10] a decrease of 120 times was achieved). The computation times are relatively consistent within the two general classes. Detection with reconstruction often takes longer when more samples are discarded, which is anticipated in light of the fact that these additional discarded samples need to be reconstructed. By comparison, the time required by detection without reconstruction is largely independent the number of discarded samples for a specific window length. It is worth noting that these computation times were obtained running the algorithms on server-grade hardware, so it can be assumed that they will significantly increase on a platform with limited size, weight and power (SWAP) resources.

Figure 5.8 shows the resulting detection performance versus SNR for each algorithm, compared to the best-case scenario where no samples are discarded. The curves shown depict the detection performance only for the signal bin in each case. The interference-bin performance is not shown, as for both the filled window and detection with reconstruction the curves would be flat (reconstruction prevents interference), while for detection without reconstruction, the performance can be found throughout Chapter 5 and Addendum B. While detection with reconstruction consistently results in better detection performance than detection without reconstruction, it is still not as good as detection using all of the original samples for most of the sensing matrices. This performance loss can most probably be attributed to reconstruction error. While a performance increase is obtained using M-FOCUSS reconstruction, this increase is almost certainly not worth the large increase in computation time, especially in applications that require real-time (or near real-time) operation; the largest performance difference is around 7 dB in Figure 5.8(e) (where only roughly 14% of samples were retained), for computation over 100 times slower with reconstruction.

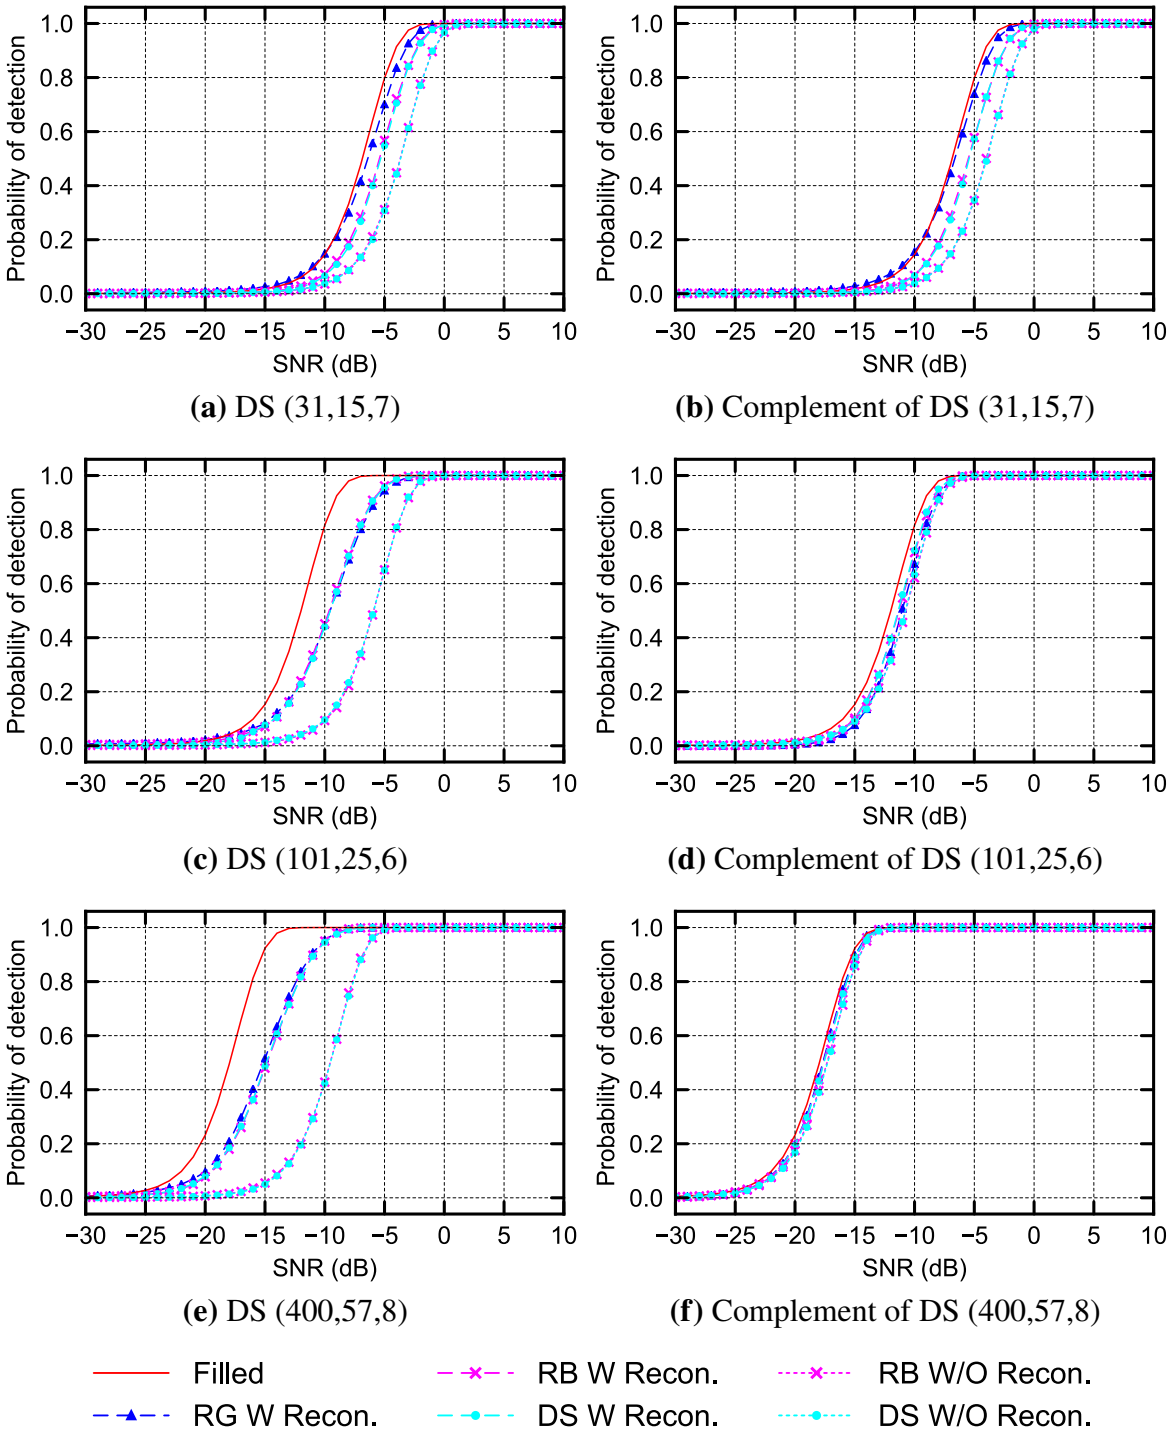

(a) DS (31,15,7)

(b) Complement of DS (31,15,7)

(c) DS (101,25,6)

(d) Complement of DS (101,25,6)

(e) DS (400,57,8)

(f) Complement of DS (400,57,8)

Figure 5.8. Probability of detection over SNR curves for the simulated signals for detection with (W) and without (W/O) reconstruction for sensing matrices based on the chosen DSs. Note that only the curves for the signal bin are compared (i.e. the interference bin curves are not shown). The curve "Filled" is the theoretical filled case where $V = 31$, 101, or 400.

5.4 RESULTS: GENERATED PRACTICAL SIGNALS

The results for the generated practical signals for the windows based on the DS (400,57,8) and the ADS (400,200,99,100) are given here. The results for the DSs (31,15,7) and (101,25,6) and the ADSs (30,15,7,22) and (101,26,6,22) can be found in Addendum B.2 (except for the comparison with detection using reconstruction, where the results are included in this chapter). All of the plots are structured identically to those in Section 5.3 (i.e. the line types and markers represent the same concepts etc.).

5.4.1 Probability of Detection over SNR

The probability of detection over SNR results for the generated practical signals are given in Figures 5.9 and 5.10. Figure 5.9 shows the probabilities of detection over SNR for the DS (400,57,8), and Figure 5.10 shows the probabilities of detection over SNR for the ADS (400,200,99,100).

Comparing the results in Figures 5.2 and 5.3 with those in Figures 5.9 and 5.10 show that the probability of detection results for the simulated signals and the generated practical signals are similar, with certain key differences. First of all, the curves in Figures 5.9 and 5.10 are not as smooth as those in Figures 5.2 and 5.3, and there is a noticeable discrepancy between the "Signal bin" and "Theory" curves in the results for the generated practical signals. This can most probably be attributed to small errors in the reception of the transmitted signals as the SNR is varied, either occurring on the transmitter side due to the fact that the signal generator used is uncalibrated, or on the receiver side due to errors from the SDR. Regardless, these errors only cause slight deviations in the expected curve shape and small SNR offsets.

The second key difference between the probability of detection results for the simulated signals

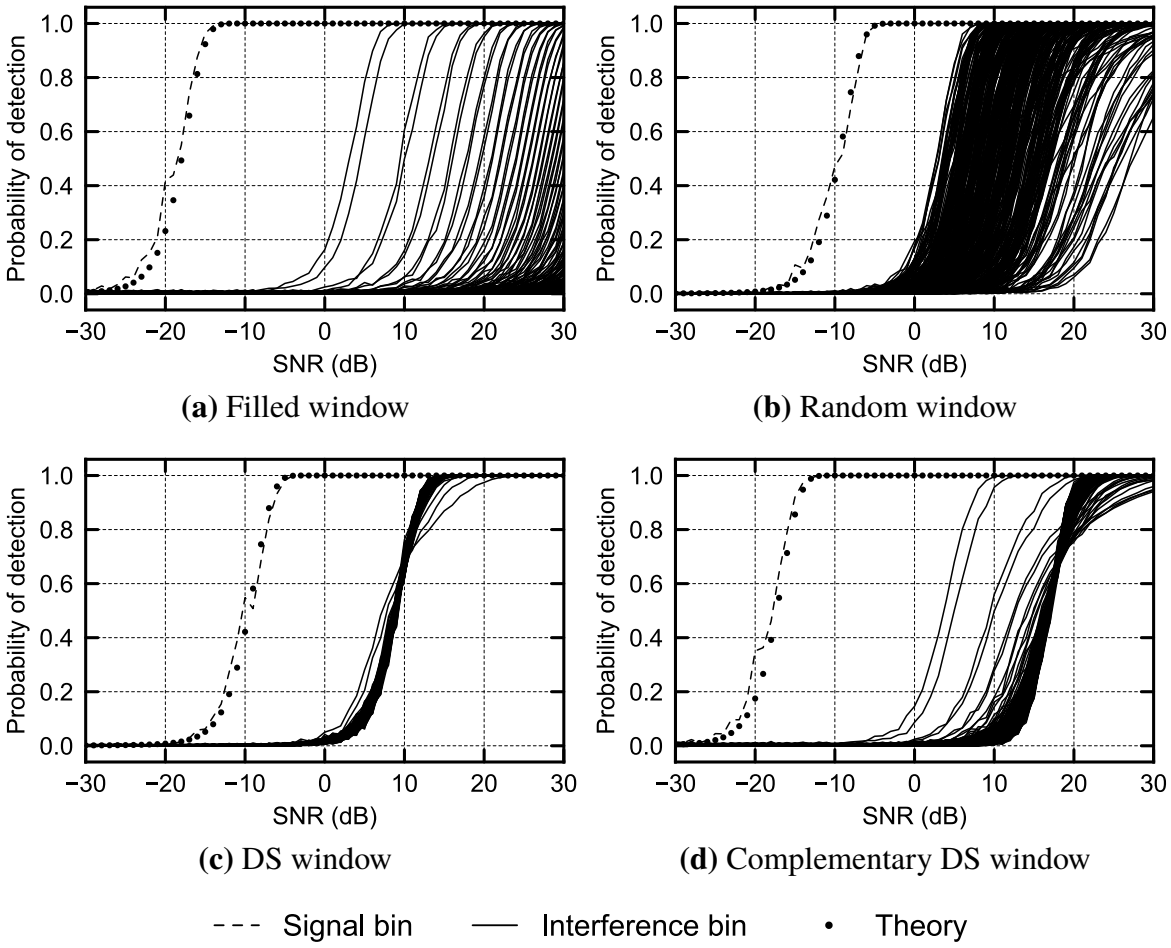

Figure 5.9. Probability of detection over SNR curves for the generated practical signals resulting from windows based on the DS (400,57,8).

and the generated practical signals can be seen by comparing Figures 5.2(a) and 5.3(a) with Figures 5.9(a) and 5.10(a). In Figures 5.9(a) and 5.10(a) interference curves for the filled windows are now visible, as opposed to the flat lines seen in Figures 5.2(a) and 5.3(a). The reason for this is the presence of sidelobes in the bins adjacent to the signal bin owing to the practical width of the signals. This is to be expected for practical signals; while the simulated signals are made up of a single RF carrier at one frequency, the practical signals are "distorted", due to the effects of practical transmission, reception, mixing, filtering, digitisation, and noise.

For the random and DS/ADS windows the results are very similar to those for the simulated signals. The are a few curves that appear distorted, but this is likely due to the window function

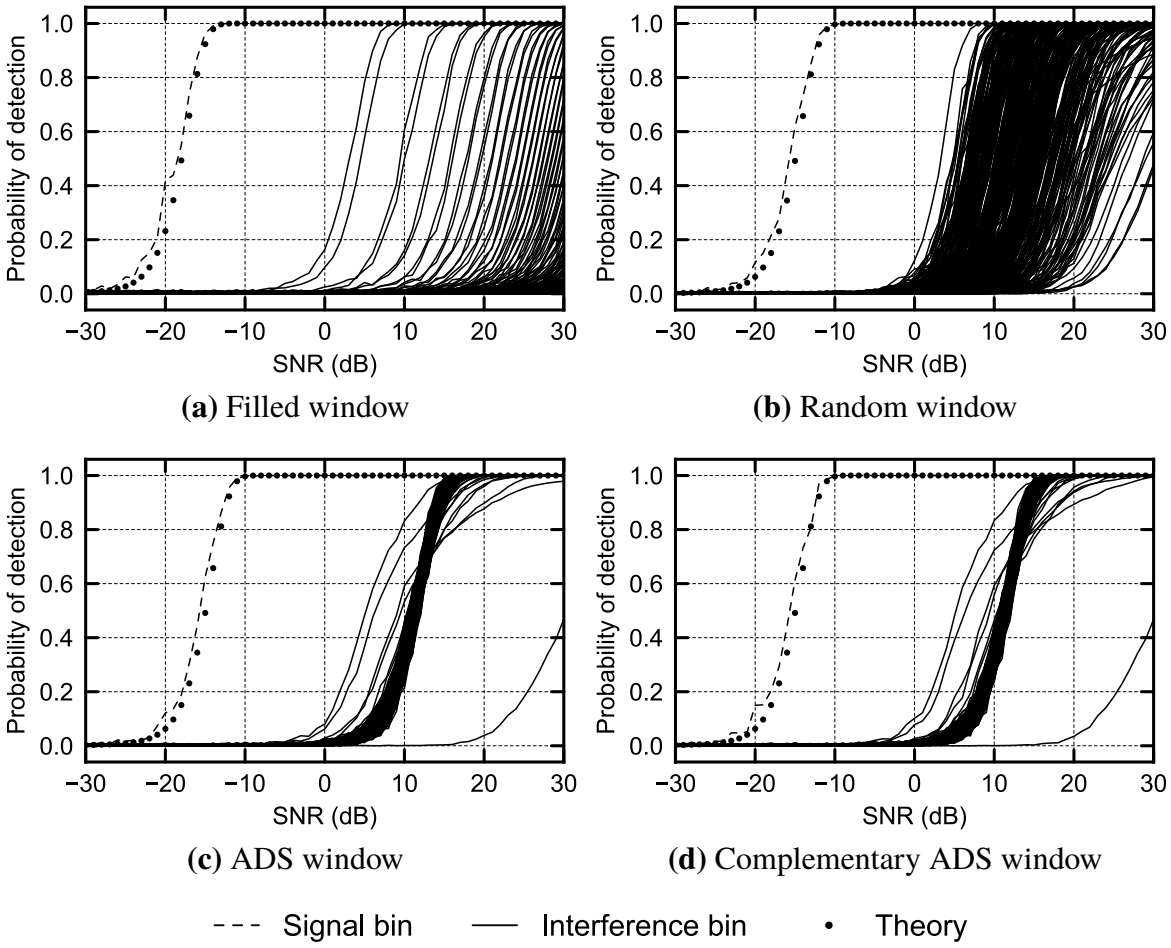

Figure 5.10. Probability of detection over SNR curves for the generated practical signals resulting from windows based on the ADS (400,200,99,100).

interacting with the signal sidelobes present in those bins; while for the simulated signals the signal is only present in one bin, for the generated practical signals some energy now occupies adjacent bins. This distortion appears limited to bins where the signal SLL is higher than the window, and the application of a random or DS/ADS window does not increase the interference. As is the case with the simulated signals, the DS/ADS windows result in more controlled interference than the equivalent random window.

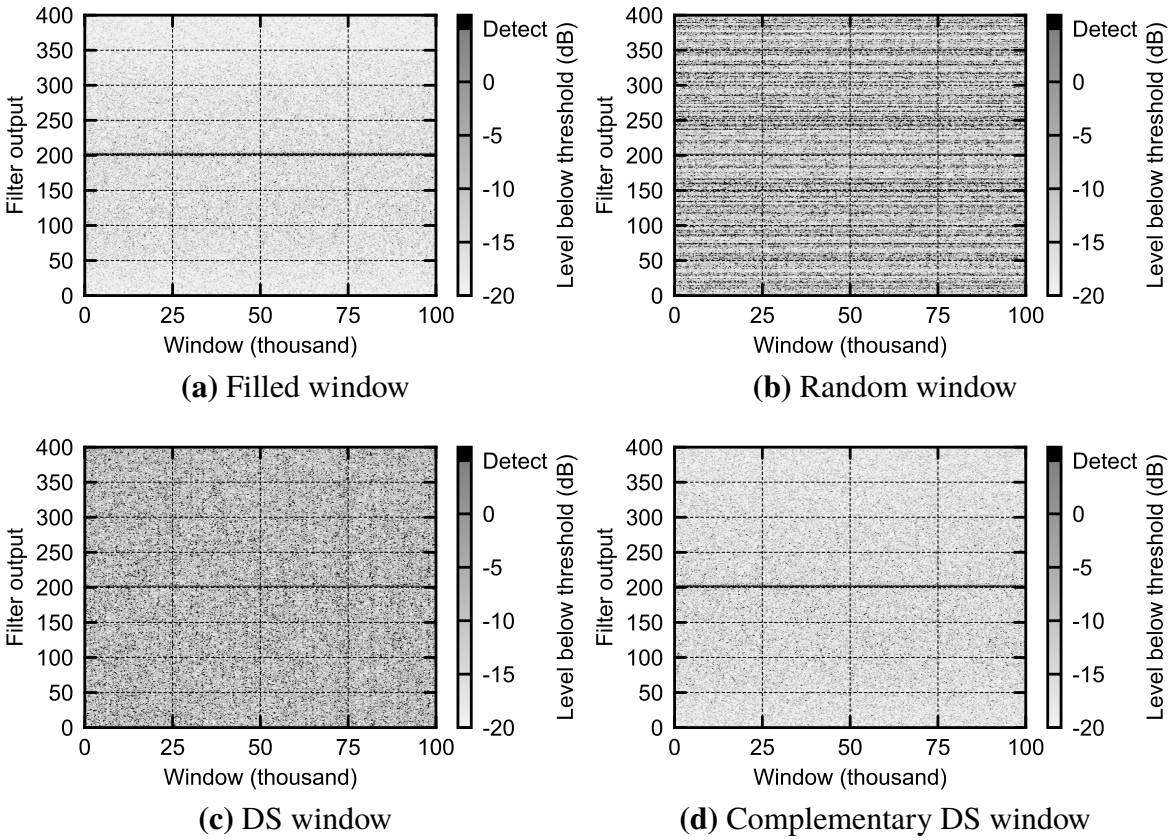

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

Figure 5.11. Test statistics for SNR = 4 dB for the generated practical signals resulting from windows based on the DS (400,57,8).

5.4.2 Test Statistics

The resulting test statistics for the generated practical signals are given in Figures 5.11 and 5.12. Figure 5.11 shows the test statistics for the DS (400,57,8), and Figure 5.12 shows the test statistics for the ADS (400,200,99,100).

Once again, the results for the generated practical signals are very similar to those for the simulated signals. Figures 5.11(a) and 5.12(a) appear to have a similar level of interference to the results in Figures 5.4(a) and 5.5(a), with the interference curves seen in Figures 5.9(a) and 5.10(a) seemingly belonging to the bins adjacent to the signal bin. This helps to confirm that these interference curves are the result of the practical sidelobes of the generated signal. Additionally, in the outer filter outputs the interference appears to decrease, most probably

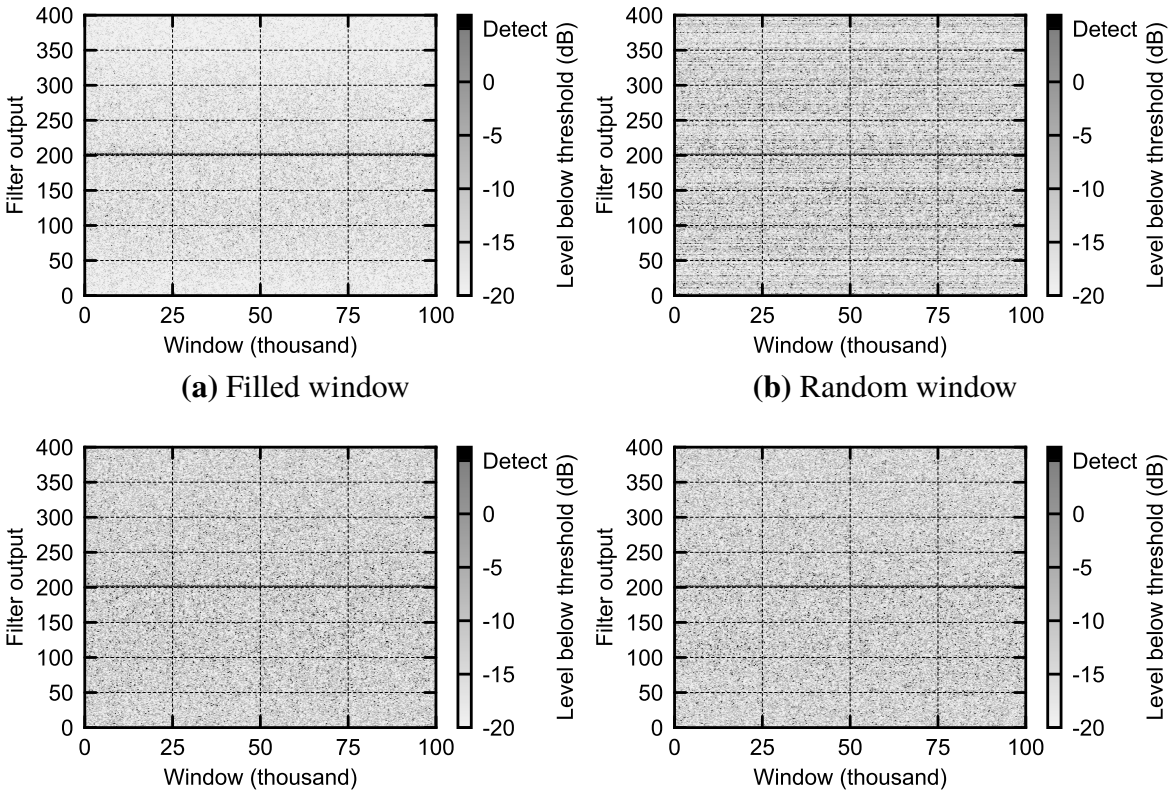

(a) Filled window **(b)** Random window

(c) ADS window **(d)** Complementary ADS window

Figure 5.12. Test statistics for SNR = 3 dB for the generated practical signals resulting from windows based on the ADS (400,200,99,100).

due to the roll-off of the bandpass filter of the SDR.

Figures 5.11(b), 5.11(c), 5.12(b), and 5.12(c) show that the similarities extend to the random and DS/ADS windows, with an increased amount of interference for both window types (the roll-off of the SDR filter is imperceptible), where the irregularity of the random window results in “false targets” that are not present for the DS/ADS windows.

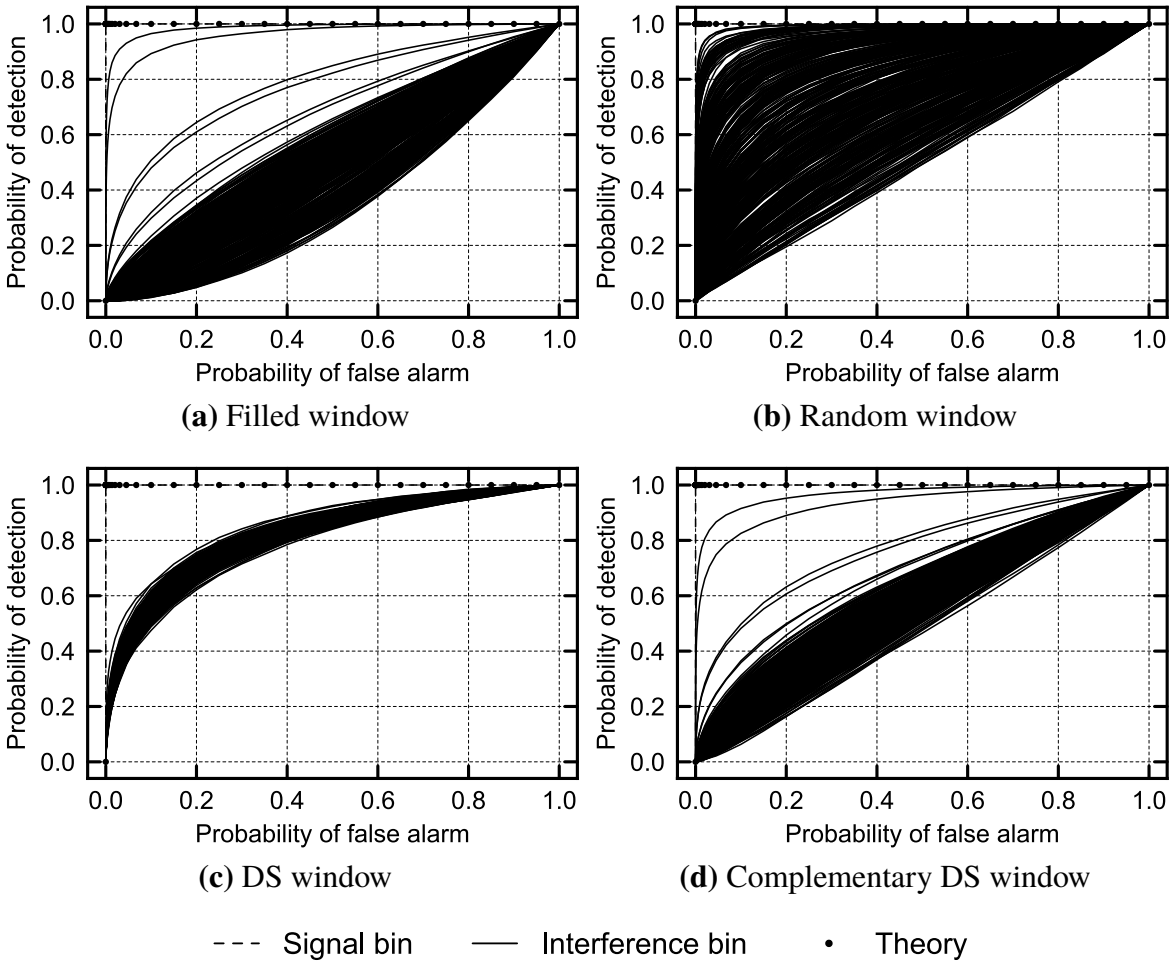

Figure 5.13. ROCs for SNR = 4 dB for the generated practical signals resulting from windows based on the DS (400,57,8).

5.4.3 Probability of Detection over Probability of False Alarm

The probability of detection over probability of false alarm (the ROC) results for the generated practical signals are given in Figures 5.13 and 5.14. Figure 5.13 shows the ROCs for a SNR of 4 dB for the DS (400,57,8), and Figure 5.14 shows the ROCs for a SNR of 3 dB for the ADS (400,200,99,100).

As with the other results, the plots are similar to those for the simulated signals. The practical sidelobes of the signal produces additional interference curves in the plots, especially

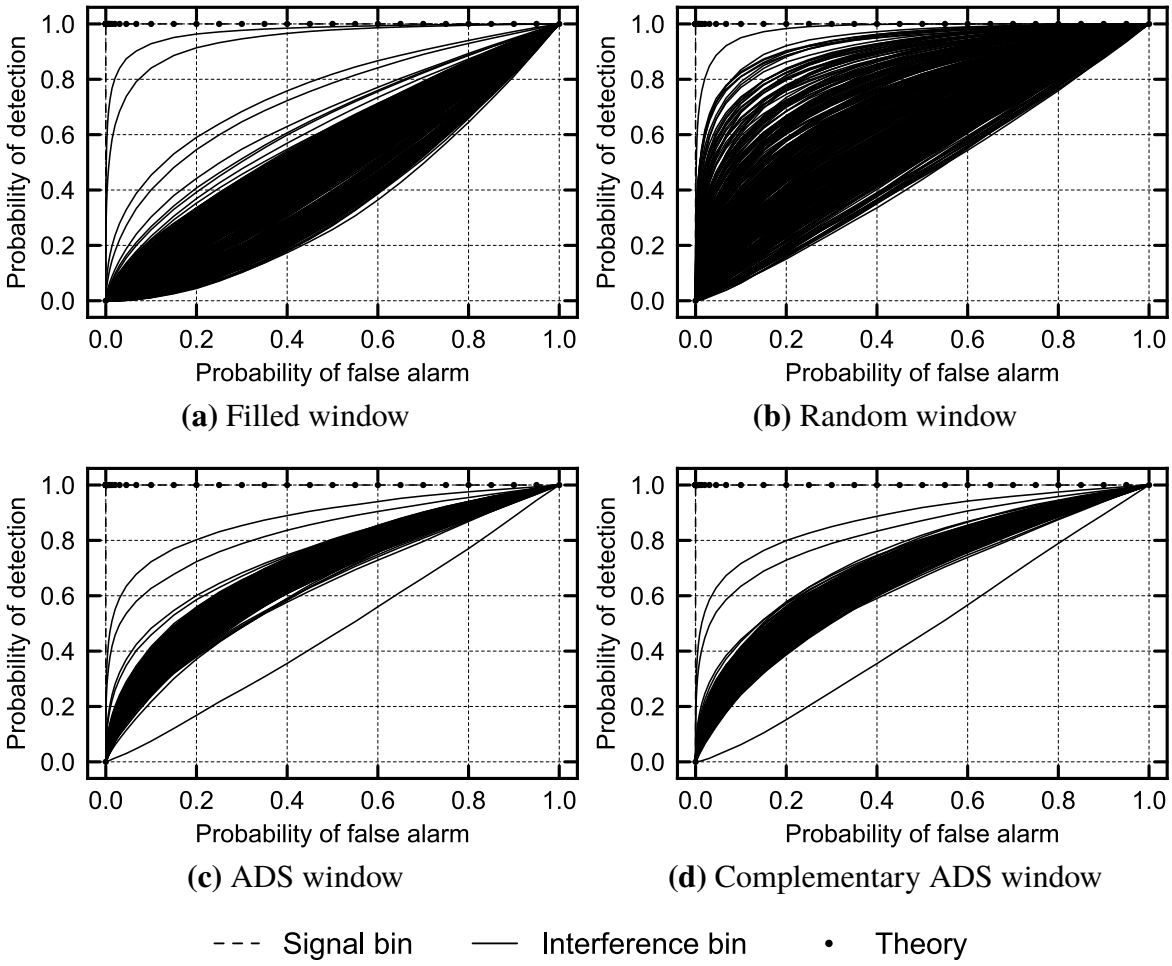

(a) Filled window **(b)** Random window

(c) ADS window **(d)** Complementary ADS window

Figure 5.14. ROCs for SNR = 3 dB for the generated practical signals resulting from windows based on the ADS (400,200,99,100).

in Figures 5.13(a) and 5.14(a), just like those seen in Figures 5.9 and 5.10. For the random and DS/ADS windows, the results are very similar to those for the simulated signals, with the DS/ADS windows again showing more consistent interference than their random counterparts.

5.4.4 Comparison with Detection using Reconstruction

The results which compare performing detection with reconstruction and detection without reconstruction on the generated practical signals are given here. The computation times for the various sensing matrices are summarised in Table 5.3, while the probability of detection over SNR curves are given in Figure 5.15. For detection with reconstruction, results for RG, RB, and DS sensing matrices are included, while for detection without reconstruction, results for RB and DS sensing matrices are included.

Table 5.3. Comparison of the computation times (in seconds) for performing detection with (W) and without (W/O) reconstruction on the generated practical signals (Adapted, with permission, from [10]. ©2018 IEEE.).

Parameters $(V, K)^*$	W Recon.			W/O Recon.	
	RG	RB	DS	RB	DS
31, 15	30.834	35.134	33.103	0.1002	0.1002
31, 16	31.969	32.253	26.491	0.0985	0.0980
101, 25	89.181	107.65	115.75	0.4132	0.4140
101, 76	22.914	9.8684	9.8545	0.4089	0.4118
400, 57	434.40	468.75	507.05	3.8957	3.8933
400, 343	66.615	60.839	59.971	3.9564	5.4544

The computation times in Table 5.3 agree with those in Table 5.2, indicating that detection with reconstruction is still much slower than detection without reconstruction. However, for detection with reconstruction the computation times for the generated practical signals are considerably faster than those for the simulated signals, while for detection without reconstruction the opposite is true. Given these changes, the best comparison is now between 59.971 s with reconstruction to 5.4544 s without reconstruction, which shows that detection with reconstruction is now at least 15 times slower than detection without reconstruction. This is a huge improvement over the results for the simulated signals, although it is still quite a large difference in computation time between the two approaches.

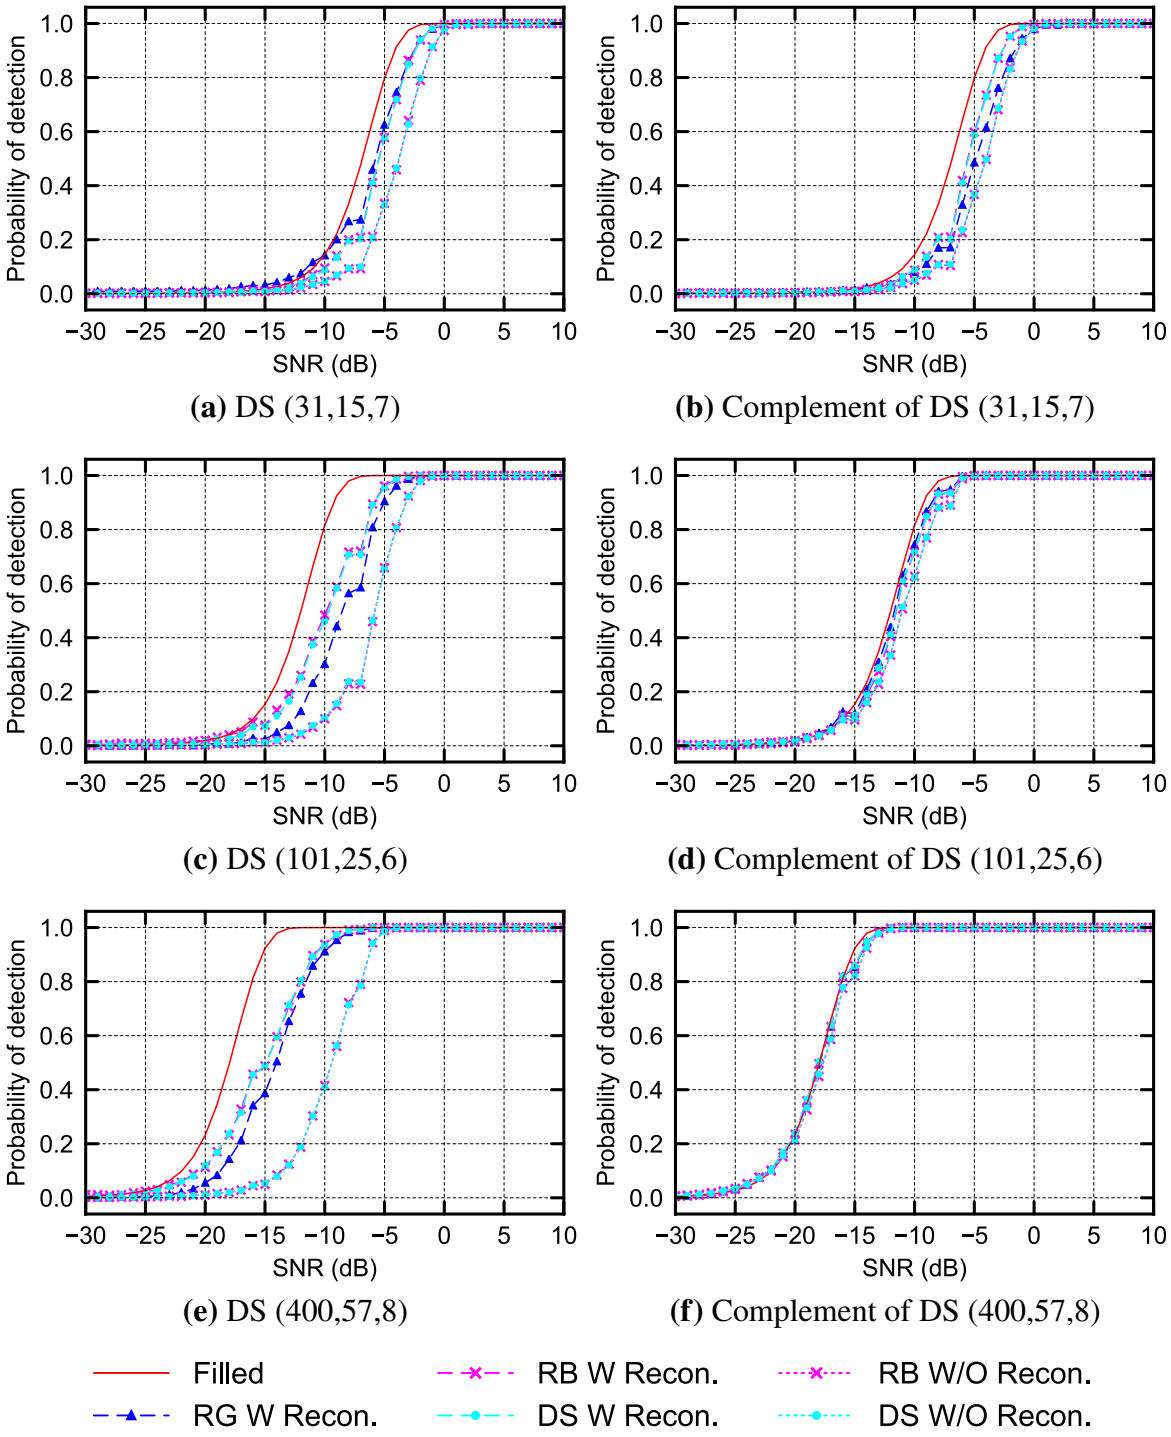

(a) DS (31,15,7)

(b) Complement of DS (31,15,7)

(c) DS (101,25,6)

(d) Complement of DS (101,25,6)

(e) DS (400,57,8)

(f) Complement of DS (400,57,8)

Figure 5.15. Probability of detection over SNR curves for the generated practical signals for detection with (W) and without (W/O) reconstruction for sensing matrices based on the chosen DSs. Note that only the curves for the signal bin are compared (i.e. the interference bin curves are not shown). The curve "Filled" is the theoretical filled case where $V = 31$, 101, or 400.

Figure 5.15 shows that the detection performance is very similar to the results obtained for the simulated signals (with slight deviations due to signal generation inaccuracies). It should be noted that for the generated practical signals, the regularization parameter needed to be made quite small (10^{-20}) in order for the algorithm to produce a solution, which would indicate that M-FOCUSS had difficulty processing the practical data (although this can most probably be attributed to amplitude scaling).

5.5 RESULTS: RECORDED PRACTICAL SIGNALS

The results for the recorded practical signals for the windows based on the DS (400,57,8) and the ADS (400,200,99,100) are given here. As mentioned in Section 5.2.3, only the test statistics are calculated for the recorded practical signals.

5.5.1 Test Statistics

The resulting test statistics for the recorded practical signals are given in Figures 5.16 and 5.17. Figure 5.16 shows the test statistics for the DS (400,57,8), and Figure 5.17 shows the test statistics for the ADS (400,200,99,100). The results are for a small fraction of the recorded signal, to allow the reader to clearly see multiple GSM bursts of various strengths occurring at staggered points in time.

The results for the filled windows are given in Figures 5.16(a) and 5.17(a) (recall the windows, and thus the results, are identical). As expected, this represents the best case scenario with the smallest interference, with both strong and weak GSM bursts being clearly observable. It is even possible to see the roll-off of the bandpass filters of the SDR in the outer bins.

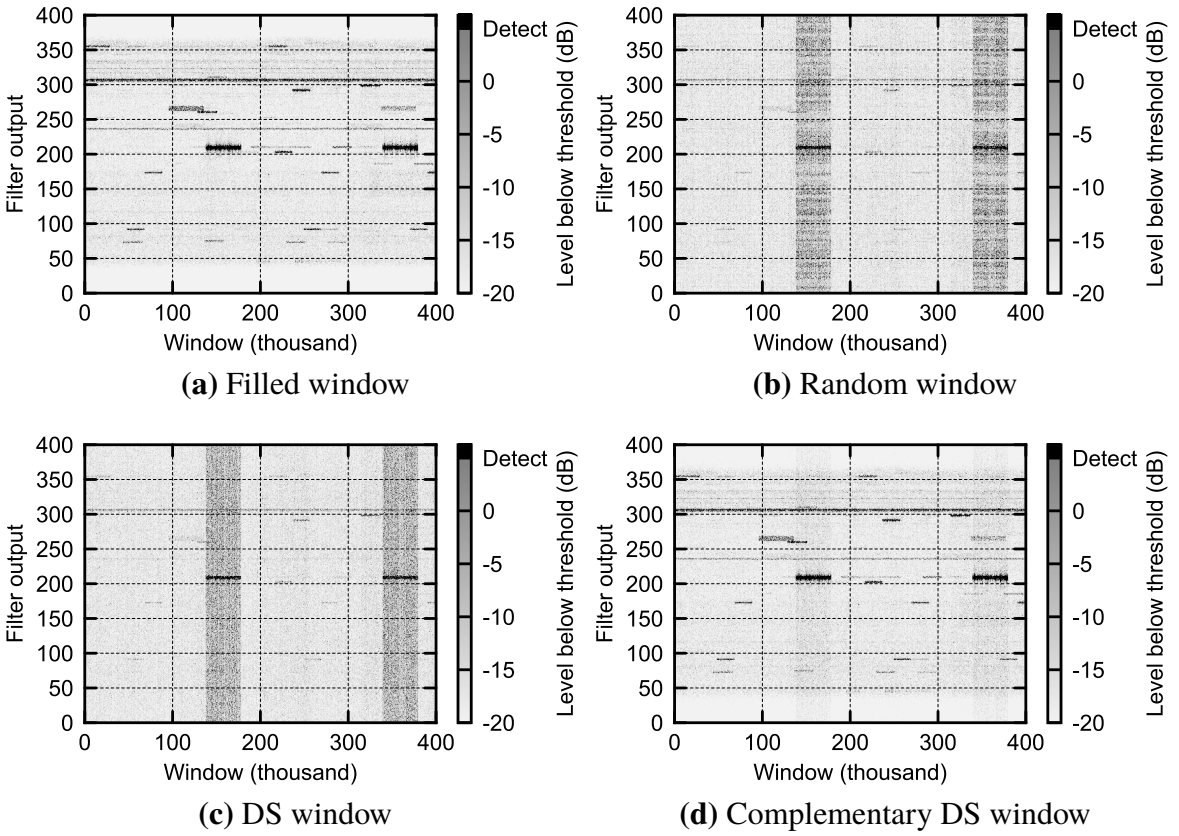

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

Figure 5.16. Test statistics for off-the-air GSM signals, resulting from windows based on the DS (400,57,8). (Reprinted, with permission, from [10]. ©2018 IEEE.)

The results for the random and DS/ADS windows are given in Figures 5.16(b), 5.16(c), 5.17(b), and 5.17(c). Many of the weaker signals are more difficult to detect compared to the case for the filled windows, due to the reduced number of samples, but additionally the masking effect discussed earlier now also comes into play. The way the strong signals mask weaker signals is clearly seen for both types windows when compared to Figures 5.16(a) and 5.17(a), with the result that some of the weaker signals can no longer be detected because they are obscured by the sidelobes of stronger ones. A good example of this is the signal located around filter output 75, between windows 100 000 and 200 000; it is clearly visible in Figures 5.16(a) and 5.17(a), but in Figures 5.16(b), 5.16(c), 5.17(b), and 5.17(c) it is almost imperceptible (it is more visible for the windows based on the ADS, because more samples are retained).

However, as observed in previous sections, for the windows based on DSs and ADSs the interference is more controlled and consistent, and consequently less severe than for the

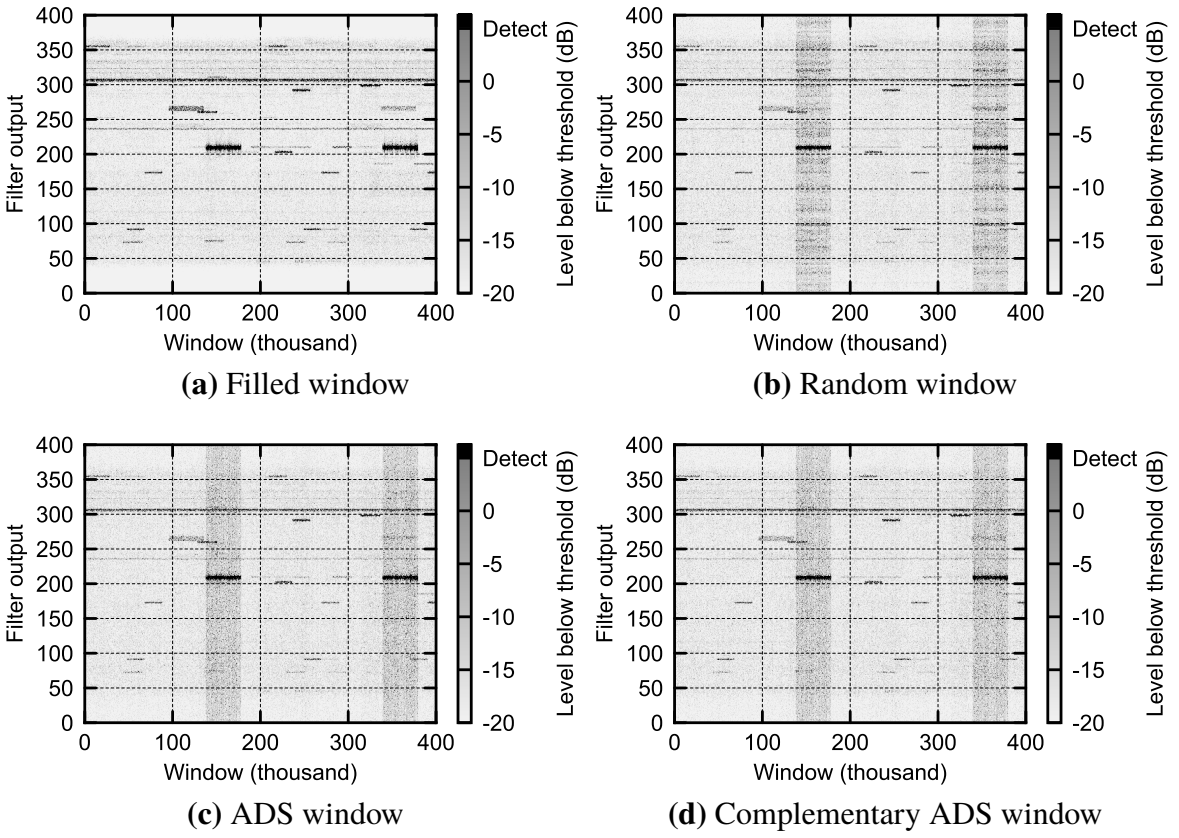

(a) Filled window **(b)** Random window

(c) ADS window **(d)** Complementary ADS window

Figure 5.17. Test statistics for off-the-air GSM signals, resulting from windows based on the ADS (400,200,99,100). (Reprinted, with permission, from [10]. ©2018 IEEE.)

random windows. For the random windows the inconsistency of the interference gives the appearance of multiple false transmissions, which could potentially confuse an operator and mask actual transmissions. For the example signal at filter output 75, between windows 100 000 and 200 000, Figures 5.16(c) and 5.17(c) show that the signal is still distinguishable within the interference, while in Figures 5.16(b) and 5.17(b) the inconsistency of the interference causes the validity of the signal to come into question, as it is surrounded by “false targets”. Another excellent example of this is the transmission located around filter output 260, between windows 300 000 and 400 000, which is two bursts long and aligns perfectly in time with the stronger transmission around window 200. Even for the ADS (400,200,99,100), which has more retained samples and thus lower interference, “false targeting” is an issue for the random window.

Finally, Figure 5.16(d) shows that using the complement of a DS with a large thinning

percentage allows a smaller number of samples to be discarded (14.25% rather than 85.75%), increasing the detection performance and lowering the SLL compared to the standard DS. In this way, it is shown that the use of DSs and ADSs to form windows allows some flexibility in determining the number of samples to be discarded and the resulting sidelobe behaviour.

5.6 SUMMARY

Empirical results comparing the detection performance of random binary windows and windows based on DSs and ADSs, for the complex-sample detector derived in Chapter 4, are presented in this chapter. Results were generated using simulated signals, generated practical signals, and recorded practical signals, for multiple sets (the remaining results can be found in Addendum B).

For each type of input signal, the results point to the same conclusions. For a random binary window and a DS/ADS based window of the same length V and number of elements K, the detection performance is identical in the bin that a signal is located in. However, in the remaining bins the sidelobes of the windows generate interference, which is radically different for each window type. For a random binary window, the resulting interference is completely random and unpredictable, due to the uncontrolled SLL of the frequency response. The randomness of the interference can also result in the appearance of "false targets" (signals that are not actually present). For a window based on a DS or ADS, the interference is far more consistent (more so for a DS than an ADS), and can be predicted (or in the case of an ADS, estimated) using only the parameters of the set. For a DS -based window, the consistency precludes the appearance of false targets.

Additionally, it was shown that there is a significant increase in computation time incurred when a reconstruction algorithm is used, while there is only a small increase in detection performance compared to detection without reconstruction.

CHAPTER 6 CONCLUSIONS AND FUTURE RESEARCH

6.1 GENERAL CONCLUSIONS

Signal detection is a fundamental operation in any signal acquisition receiver. However, many such systems process wideband signals at high frequencies with large bandwidths, which leads to a large number of data samples that require processing. While many systems are designed to handle high data rate processing, there are some scenarios (e.g. [2]) where size, weight and power (SWAP) limitations require that the amount of data samples be reduced.

The objective of this work was to investigate the effects of using difference sets (DSs) and almost difference sets (ADSs) to construct binary windows for discarding samples prior to performing signal detection. The properties of these sets have already seen them successfully applied to the design of massively thinned antenna arrays [7, 9] (as well as other signal processing problems), and the similarity between array thinning and discarded detection suggested that the sets were a promising solution.

The use of DSs and ADSs for detection has been thoroughly investigated. The properties of the sets and their applicability to the detection problem was described mathematically in Chapter 3. A mathematical analysis of detection with discarded samples was conducted in Chapter 4, which also includes an explanation of the effects of the sets on the derived detector.

These effects were then experimentally validated in Chapter 5, using simulated, practically generated, and practically recorded signals. The proposed technique was also compared to an equivalent approach that makes use of a reconstruction algorithm in Chapter 5, to illustrate the trade-offs between the two.

The results of this work, cast in terms of the research questions given in Chapter 1, are summarised below.

Binary windowing effects: For detectors that process frequency domain information, such as the modified version of the periodogram (or spectrogram) derived in Chapter 4, the application of a binary window function prior to detection has been shown to affect the detection performance beyond having fewer samples to work with. The detector used in this work uses a V-point fast Fourier transform (FFT) to produce V test statistics (or bins) for each window of samples. A binary window function will weight these V test statistics according to the sidelobe level (SLL) of the window's frequency response at the corresponding points. In a bin where a signal is located, the detection performance behaves according to theory, and will be identical for any binary window of length V and number of elements K. However, the sidelobes of various windows will differ, and these sidelobes also affect the detection performance. The sidelobes from a signal in one bin can directly affect the detection performance in the remaining bins, with high sidelobes causing interference that can obscure smaller signals. Therefore, construction of a binary window that has a frequency response with controlled sidelobes is shown to be imperative.

Effects of random binary windows: It has been verified that random binary windows have frequency responses with uncontrolled sidelobes (as seen in thinned antenna array design [8]). For a signal located in a particular bin, these sidelobes result in interference in the remaining bins, interference that (as expected) is completely random. As discussed previously, the interference created by large signals can potentially obscure smaller ones, preventing their detection. In addition, for random binary windows the uncontrolled nature of the sidelobes results in interference that is inconsistent between adjacent bins, which can cause the appearance of "false targets" that could be mistaken for actual

signals. These effects increase with increased thinning (i.e. the discarding of more samples).

Effects of DS/ADS binary windows: It has been shown that binary windows based on DSs and ADSs also result in sidelobe interference, however it is nowhere as severe as for random binary windows. The frequency response of a window based on a DS/ADS is far more controlled than that of the equivalent random binary window, as is the resulting interference. For DSs, the interference curves all align at the same position (signal-to-noise ratio (SNR) difference from the signal bin curve), a position that can be determined from the parameters (V, K, Λ) of the set, which is not the case for a random binary window. While for ADSs the interference is not as controlled as for DSs, they still allow for reasonable estimation via the set parameters. The consistency of the sidelobes precludes the appearance of false targets for DSs (and they are far less severe for ADSs). Finally, while sets do not exist for every possible window length V and number of elements K, there are a large number of sets available [57, 58], and the existence of complementary sets allow for a diverse range of windows to be constructed.

In addition to answers to the proposed research questions, this work also produced additional results. These results are summarised below.

Operation on real-world signals: The proposed technique was successfully applied to practical recordings of real-world Global System for Mobile Communications (GSM) signals, proving that the approach can be used in practical applications.

Comparison with reconstruction: The proposed approach of discarding samples and performing detection without a reconstruction algorithm was also compared to an equivalent approach that did make use of reconstruction. It was shown that although lack of a reconstruction algorithm does result in a decrease in detection performance (in addition to the creation of interference), the computational burden of reconstruction is exceedingly high, as theorised.

6.2 FUTURE WORK

This research work was relatively limited in its scope, as the main goal was to prove a concept. Possible expansions that could be made to this work are summarised below.

Practical implementation: The goal of this work was to prove the concept of discarding samples using DSs and ADSs. However, the primary motivation is application to practical systems with limited SWAP requirements, which has not yet been achieved. Therefore, arguably the first expansion to be investigated is whether or not the technique will actually allow for detection to be performed successfully on wideband signals on a platform with limited processing power.

Sub-bin performance: The two detectors derived in Chapter 4 make use of a discrete Fourier transform (DFT) (practically a FFT) to form the test statistics. It was stipulated previously that the length of this DFT/FFT should be V-point, where V is the length of the binary sampling window, because the controlled SLLs associated with DSs and ADSs only occur at V points in the frequency response. This results in V bins (or test statistics) centred at different frequencies.

A potential problem that arises is that of frequency resolution and "sub-bin" performance. If the frequency of the desired signal happens to coincide with the centre frequency of a bin, the expected detection performance will be obtained. However, most often this is not the case, and signals exist at frequencies between adjacent frequency bins. The effects of this, referred to here as "sub-bin" performance, is cause for further investigation.

Additional detectors: In Chapter 4 three different signal models were proposed, but only one was used to subsequently derive two detectors.[1] The effects of DSs and ADSs on the detectors that result from the remaining signal models, as well as additional detectors, have yet to be determined.

[1]This was done for multiple reasons: simplicity, the applicability of the chosen signal model in the given context, and the fact that the resulting test statistics combine the processes of detection and channelisation.

Optimised binary windows: This work has proven the effectiveness of DSs and ADSs as a means of discarding samples prior to signal detection that produces controlled, deterministic results (especially compared to random binary windows). However, there may be other deterministic methods that could be applied to the sample reduction problem to given comparable, or even improved, results. Specifically, the problem of selecting samples within a window to minimise sidelobe interference could be cast as a very interesting optimisation problem.

Weighted binary windows: The binary windows considered in this work do not weight the selected samples, in order to minimise computational complexity. However most window functions, designed to control the resulting frequency response, do. Therefore, the sidelobe performance of windows constructed using DSs/ADSs could be further enhanced by incorporating some form of weighting function.

REFERENCES

[1] C. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.

[2] J. N. Koster, A. Buysse, L. Smith, J. Huyssen, J. Hotchkiss, J. Malangoni, and J. Schneider, "AREND: a sensor aircraft to support wildlife rangers," in *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, San Diago, CA, USA, Jan. 2016, pp. 1–19.

[3] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Upper Saddle River, NJ, USA: Prentice Hall, 1998.

[4] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[5] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[6] S. Qaisar, R. M. Bilal, W. Iqbal, M. Naureen, and S. Lee, "Compressive sensing: from theory to applications, a survey," *Journal of Communications and Networks*, vol. 15, no. 5, pp. 443–456, Oct. 2013.

[7] D. G. Leeper, "Isophoric arrays - massively thinned phased arrays with well-controlled sidelobes," *IEEE Transactions on Antennas and Propagation*, vol. 47, no. 12, pp. 1825–1835, Dec. 1999.

[8] Y. T. Lo, "A mathematical theory of antenna arrays with randomly spaced elements," *IEEE Transactions on Antennas and Propagation*, vol. 12, no. 3, pp. 257–268, May 1964.

[9] G. Oliveri, M. Donelli, and A. Massa, "Linear array thinning exploiting almost difference sets," *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 12, pp. 3800–3812, Dec. 2009.

[10] W. R. Smith and W. P. du Plessis, "Sample reduction during signal detection using difference sets and almost difference sets," *IEEE Transactions on Signal Processing*, submitted for publication.

[11] M. Mishali and Y. Eldar, "Sub-Nyquist sampling," *IEEE Signal Processing Magazine*, vol. 28, no. 6, pp. 98–124, Nov. 2011.

[12] J. Crols and M. Steyaert, "Low-IF topologies for high-performance analog front ends of fully integrated receivers," *IEEE Transactions on Circuits and Systems—Part II: Analog and Digital Signal Processing*, vol. 45, no. 3, pp. 269–282, Mar. 1998.

[13] R. Vaughan, N. Scott, and D. White, "The theory of bandpass sampling," *IEEE Transactions on Signal Processing*, vol. 39, no. 9, pp. 1973–1984, Sep. 1991.

[14] H. J. Landau, "Necessary density conditions for sampling and interpolation of certain entire functions," *Acta Mathematica*, vol. 117, no. 1, pp. 37–52, 1967.

[15] Y.-P. Lin and P. Vaidyanathan, "Periodically nonuniform sampling of bandpass signals," *IEEE Transactions on Circuits and Systems—Part II: Analog and Digital Signal*

Processing, vol. 45, no. 3, pp. 340–351, Mar. 1998.

[16] P. Feng and Y. Bresler, "Spectrum-blind minimum-rate sampling and reconstruction of multiband signals," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'96)*, Atlanta, GA, USA, May 1996, pp. 1688–1691.

[17] C. Herley and P. W. Wong, "Minimum rate sampling and reconstruction of signals with arbitrary frequency support," *IEEE Transactions on Information Theory*, vol. 45, no. 5, pp. 1555–1564, Jul. 1999.

[18] R. Venkataramani and Y. Bresler, "Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals," *IEEE Transactions on Information Theory*, vol. 46, no. 6, pp. 2173–2183, Sep. 2000.

[19] Y. M. Lu and M. N. Do, "A theory for sampling signals from a union of subspaces," *IEEE Transactions on Signal Processing*, vol. 56, no. 6, pp. 2334–2345, May 2008.

[20] M. Mishali and Y. C. Eldar, "Blind multiband signal reconstruction: compressed sensing for analog signals," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.

[21] ——, "Reduce and boost: recovering arbitrary sets of jointly sparse vectors," *IEEE Transactions on Signal Processing*, vol. 56, no. 10, pp. 4692–4702, Oct. 2008.

[22] ——, "From theory to practice: sub-Nyquist sampling of sparse wideband analog signals," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 375–391, Apr. 2010.

[23] M. Mishali, Y. C. Eldar, and A. J. Elron, "Xampling: signal acquisition and processing in union of subspaces," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp.

4719–4734, Jan. 2011.

[24] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, Jun. 2002.

[25] I. Maravić and M. Vetterli, "Sampling and reconstruction of signals with finite rate of innovation in the presence of noise," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2788–2805, Aug. 2005.

[26] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, Mar. 1986.

[27] R. Roy and T. Kailath, "ESPRIT - estimation of signal parameters via rotational invariance techniques," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, Jul. 1989.

[28] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: efficient sampling of sparse bandlimited signals," *IEEE Transactions on Information Theory*, vol. 56, no. 1, pp. 520–544, Jan. 2010.

[29] J. Yoo, S. Becker, M. Monge, M. Loh, E. Candes, and A. Emami-Neyestanak, "Design and implementation of a fully integrated compressed-sensing signal acquisition system," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'12)*, Kyoto, Japan, Aug. 2012, pp. 5325–5328.

[30] Y. C. Eldar and G. Kutyniok, Eds., *Compressed Sensing: Theory and Applications*. Cambridge: Cambridge University Press, 2012.

[31] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk, "Signal processing with compressive measurements," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 445–460, Apr. 2010.

[32] J. Haupt and R. Nowak, "Compressive sampling for signal detection," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'07)*, Honolulu, HI, USA, Apr. 2007, pp. 1509–1512.

[33] B. S. Mysore Rama Rao, S. Chatterjee, and B. Ottersten, "Detection of sparse random signals using compressive measurements," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'12)*, Kyoto, Japan, Mar. 2012, pp. 3257–3260.

[34] E. Lagunas, S. K. Sharma, S. Chatzinotas, and B. Ottersten, "Compressive sensing based energy detector," in *Proceedings of the 24th European Signal Processing Conference (EUSIPCO'16)*, Budapest, Hungary, Aug. 2016, pp. 1678–1682.

[35] E. Lagunas and L. Rugini, "Performance of compressive sensing based energy detection," in *Proceedings of the IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC'17)*, Montreal, QC, Canada, Oct. 2017, pp. 1–5.

[36] J. Cao and Z. Lin, "Bayesian signal detection with compressed measurements," *Information Sciences*, vol. 289, no. 1, pp. 241–253, Dec. 2014.

[37] T. Wimalajeewa and P. K. Varshney, "Sparse signal detection with compressive measurements via partial support set estimation," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 1, pp. 46–60, Mar. 2017.

[38] M. Duarte, M. Davenport, M. Wakin, and R. Baraniuk, "Sparse signal detection from incoherent projections," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'06)*, Toulouse, France, May 2006, pp. 305–308.

[39] H. E. A. Laue, "Demystifying compressive sensing," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 171–176, Jul. 2017.

[40] T. Wimalajeewa, H. Chen, and P. K. Varshney, "Performance analysis of stochastic signal detection with compressive measurements," in *2010 Conference Record of the Forty Fourth Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, USA, Nov. 2010, pp. 813–817.

[41] Z. Tian and G. B. Giannakis, "Compressed sensing for wideband cognitive radios," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'07)*, Honolulu, HI, USA, Apr. 2007, pp. 1357–1360.

[42] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, Jan. 1978.

[43] J. B. Y. Tsui and C. Cheng, *Digital Techniques for Wideband Receivers*, 3rd ed. Edison, NJ, USA: SciTech Publishing, 2015.

[44] B. Steinberg, "Comparison between the peak sidelobe of the random array and algorithmically designed aperiodic arrays," *IEEE Transactions on Antennas and Propagation*, vol. 21, no. 3, pp. 366–370, May 1973.

[45] L. D. Baumert, *Cyclic Difference Sets*. New York, NY, USA: Springer-Verlag, 1971.

[46] M. Hall, *Combinatorial Theory*. Hoboken, NJ, USA: John Wiley & Sons, Inc., 1988.

[47] D. R. Stinson, *Combinatorial Designs: Constructions and Analysis*. New York, NY, USA: Springer-Verlag, 2004.

[48] C. J. Colbourn and J. H. Dinitz, Eds., *Handbook of Combinatorial Designs*, 2nd ed. Boca Raton, FL, USA: Chapman and Hall/CRC, 2007.

[49] C. Ding, M. Golin, and T. Kløve, "Meeting the Welch and Karystinos-Pados bounds on DS-CDMA binary signature sets," *Designs, Codes and Cryptography*, vol. 30, no. 1, pp. 73–84, Aug. 2003.

[50] P. Xia, S. Zhou, and G. Giannakis, "Achieving the Welch bound with difference sets," *IEEE Transactions on Information Theory*, vol. 51, no. 5, pp. 1900–1907, May 2005.

[51] H. Stahl, J. Mietzner, and R. F. H. Fischer, "A sub-Nyquist radar system based on optimized sensing matrices derived via sparse rulers," in *Proceedings of the 3rd International Workshop on Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa'15)*, Pisa, Italy, Jun. 2015, pp. 36–40.

[52] L. Welch, "Lower bounds on the maximum cross correlation of signals," *IEEE Transactions on Information Theory*, vol. 20, no. 3, pp. 397–399, May 1974.

[53] J. A. Davis, "Almost difference sets and reversible divisible difference sets," *Archiv der Mathematik*, vol. 59, no. 6, pp. 595–602, Dec. 1992.

[54] C. Ding, T. Helleseth, and K. Y. Lam, "Several classes of binary sequences with three-level autocorrelation," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2606–2612, Nov. 1999.

[55] C. Ding, T. Helleseth, and H. Martinsen, "New families of binary sequences with optimal three-level autocorrelation," *IEEE Transactions on Information Theory*, vol. 47, no. 1, pp. 428–433, Jan. 2001.

[56] K. T. Arasu, C. Ding, T. Helleseth, P. V. Kumar, and H. M. Martinsen, "Almost difference sets and their sequences with optimal autocorrelation," *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2934–2943, Nov. 2001.

[57] D. Gordon. (2018, Jun.) La Jolla difference set repository. [Online]. Available: https://www.ccrwest.org/diffsets.html

[58] (2018, Jun.) ELEDIA cyclic set repository. [Online]. Available: https://eledia.science.unitn.it/ADSRepository/main_ads_page.html

[59] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 2nd ed. New York, NY, USA: McGraw-Hill, 1984.

[60] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice Hall, 1993.

[61] (2018, Aug.) Ettus Research | USRP N200 software defined radio. National Instruments. [Online]. Available: https://www.ettus.com/product/details/UN200-KIT

[62] (2018, Aug.) Ettus Research | SBX USRP daughterboard. National Instruments. [Online]. Available: https://www.ettus.com/product/details/SBX

[63] (2018, Jun.) Nuand | bladeRF software defined radio. Nuand LLC. [Online]. Available: http://www.nuand.com

[64] European Telecommunications Standards Institute, "Digital cellular telecommunications system (Phase 2+)(GSM); GSM/EDGE Radio transmission and reception," 3GPP TS 45.005 Rev. 14.4.0, Apr. 2018.

[65] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, "Sparse solutions to linear inverse problems with multiple measurement vectors," *IEEE Transactions on Signal Processing*, vol. 53, no. 7, pp. 2477–2488, Jul. 2005.

ADDENDUM A FREQUENCY RESPONSES

The frequency responses for the windows based on the sets specified in Chapter 5 are given here. These frequency responses are directly related to the interference bins found in the detection performance results in Chapter 5 and Addendum B. The work described in this addendum forms part of a submitted journal paper [10].[1]

For each window the cyclic shifts corresponding to the lowest, median, and highest SLL values are plotted against u, which represents normalised frequency in this work. The details of the markers and horizontal lines can be found in Section 3.3.3.

The frequency responses behave similarly to those found in Section 3.3.3. There is no sidelobe roll-off as seen for filled/rectangular windows, and the sample points of the random windows are far less controlled than those of the DS/ADS-based windows. For the DS and complementary DS windows, whether V is even or odd the sample points are constrained to a fixed level dictated by (3.21). For the ADS and complementary ADS windows, if V is odd the sample points are constrained to two distinct levels as dictated by (3.23). However, if V is even the levels are not distinct, and there appears to be a null in the centre of the response at $u = 1.0$, below the minimum limit set by (3.29). The cause of this phenomenon is unknown, but is also outside the scope of this work.

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

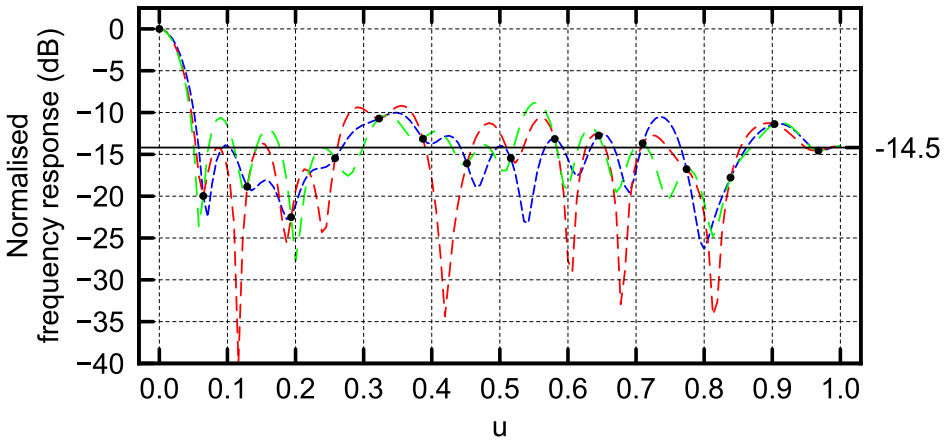

(a) Random window

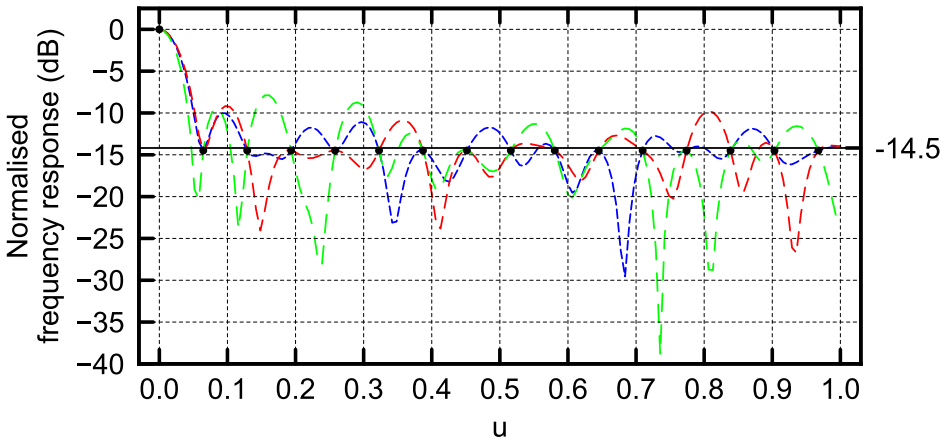

(b) DS window

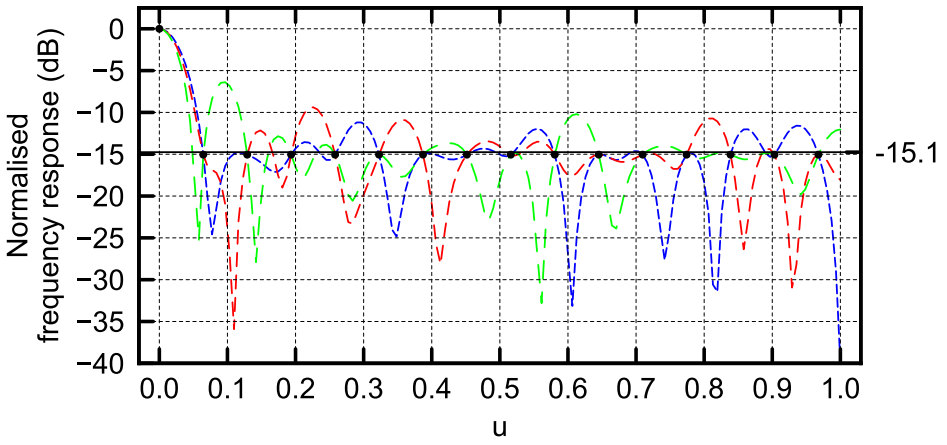

(c) Complementary DS window

Figure A.1. Frequency responses resulting from windows based on the DS (31,15,7). (Reprinted, with permission, from [10]. ©2018 IEEE.)

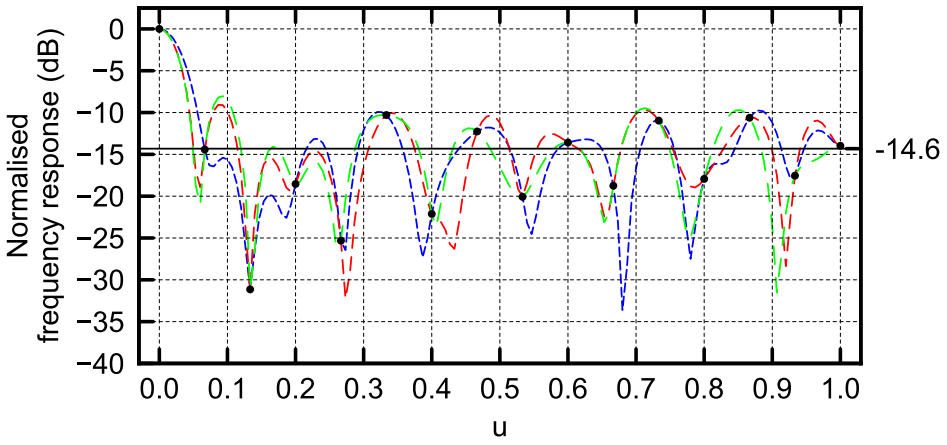

(a) Random window

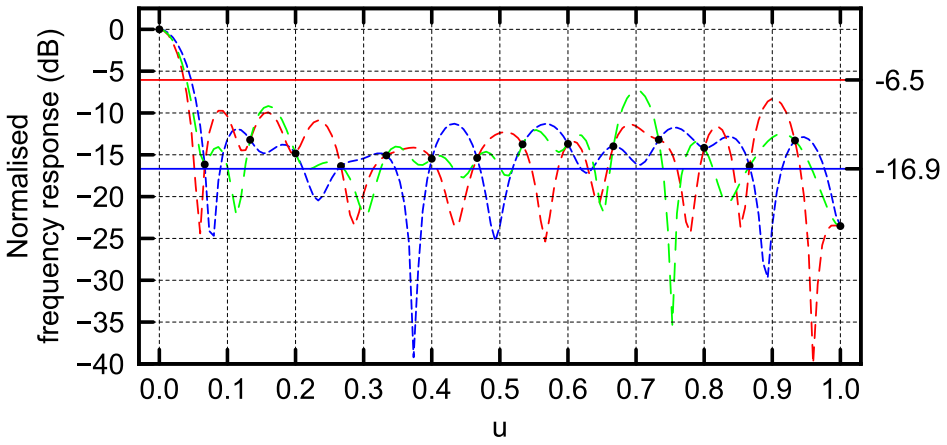

(b) ADS window

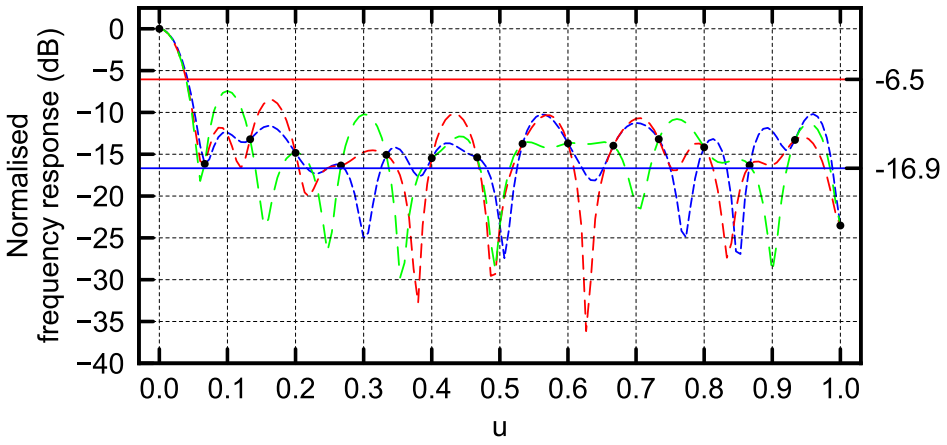

(c) Complementary ADS window

Figure A.2. Frequency responses resulting from windows based on the ADS (30,15,7,22). (Reprinted, with permission, from [10]. ©2018 IEEE.)

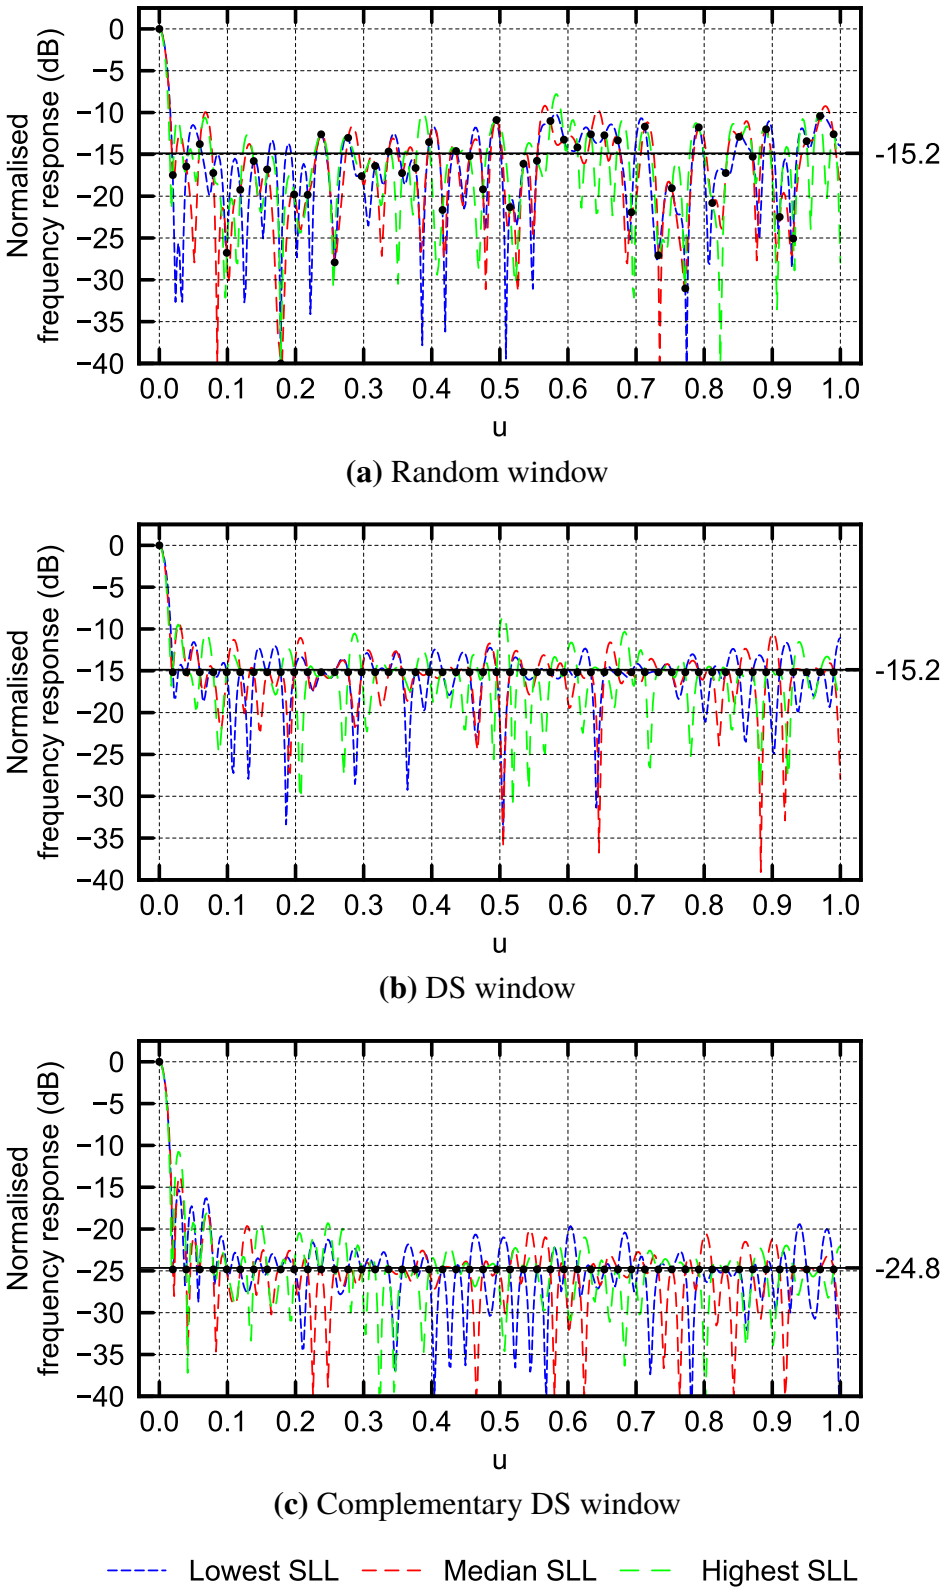

(a) Random window

(b) DS window

(c) Complementary DS window

Figure A.3. Frequency responses resulting from windows based on the DS (101,25,6). (Reprinted, with permission, from [10]. ©2018 IEEE.)

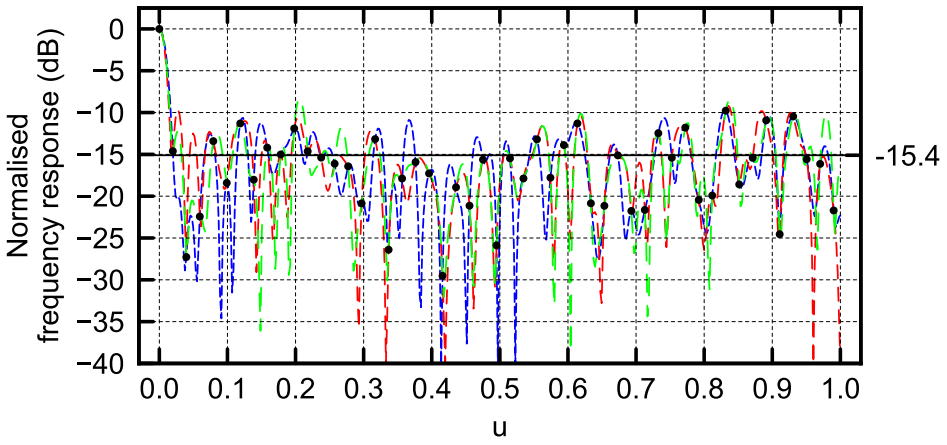

(a) Random window

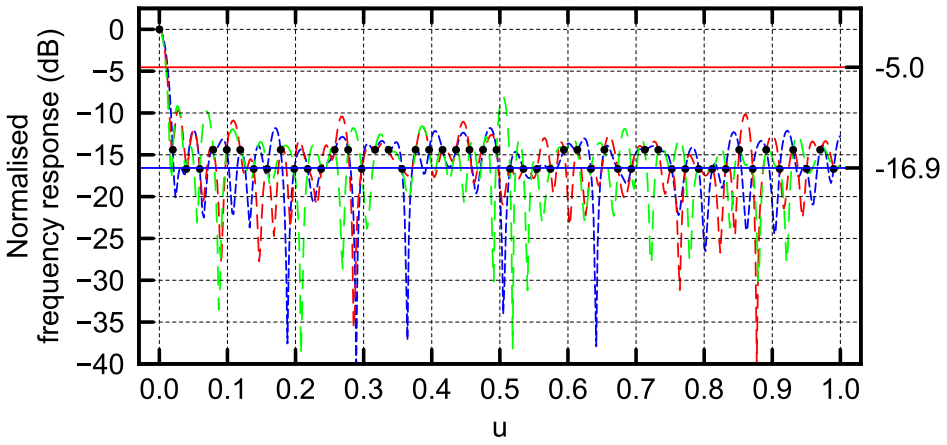

(b) ADS window

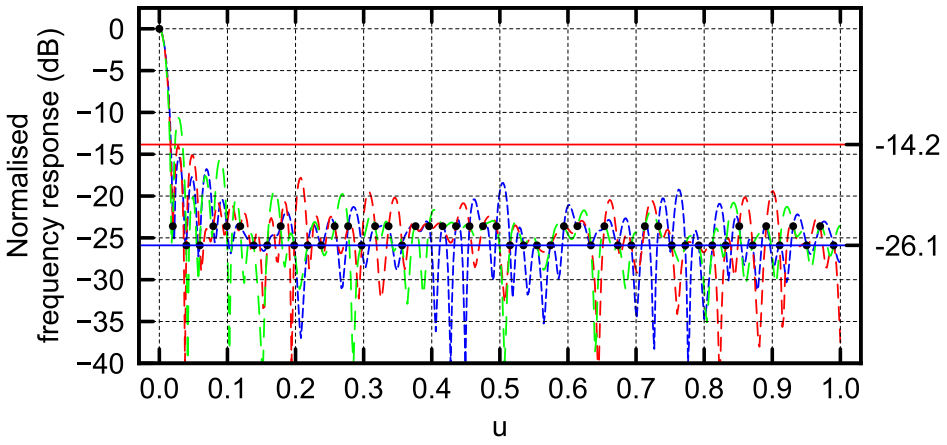

(c) Complementary ADS window

Figure A.4. Frequency responses resulting from windows based on the ADS (101,26,6,50). (Reprinted, with permission, from [10]. ©2018 IEEE.)

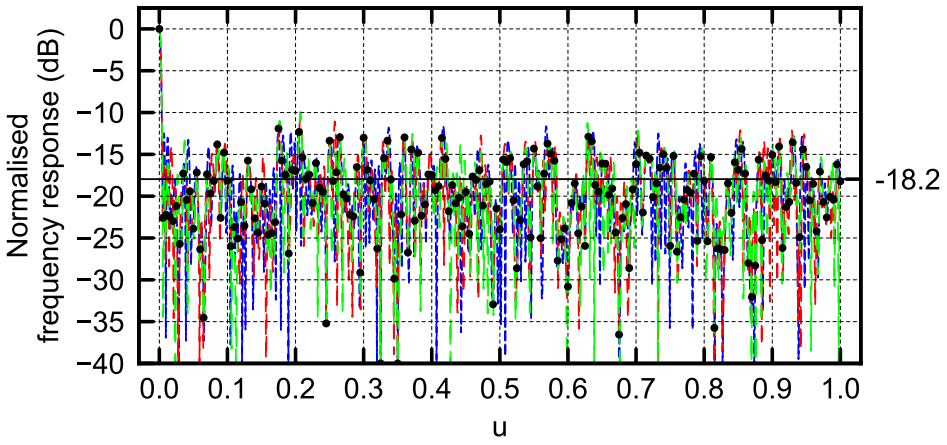

(a) Random window

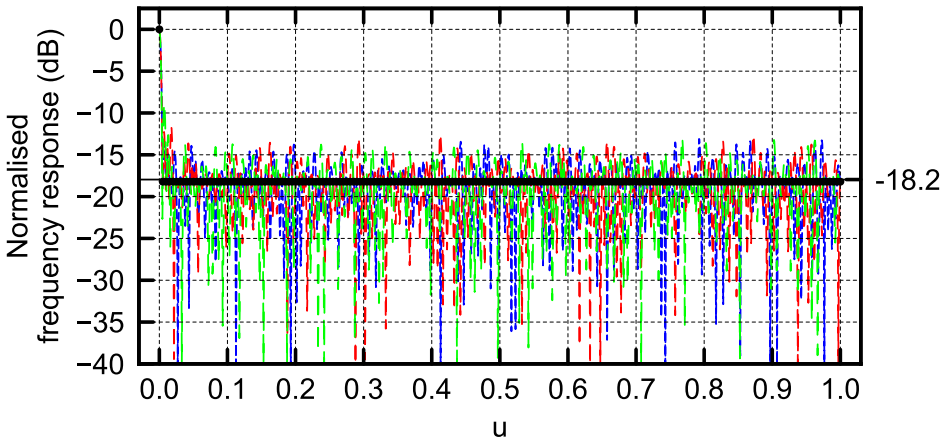

(b) DS window

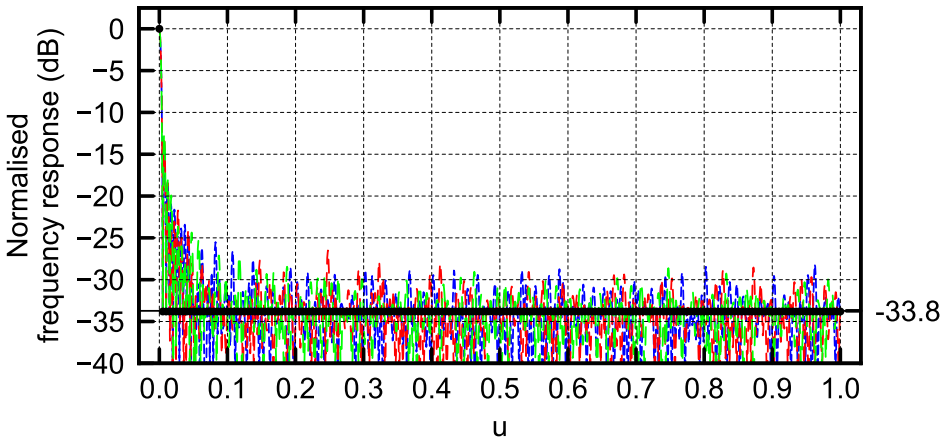

(c) Complementary DS window

Figure A.5. Frequency responses resulting from windows based on the DS (400,57,8).

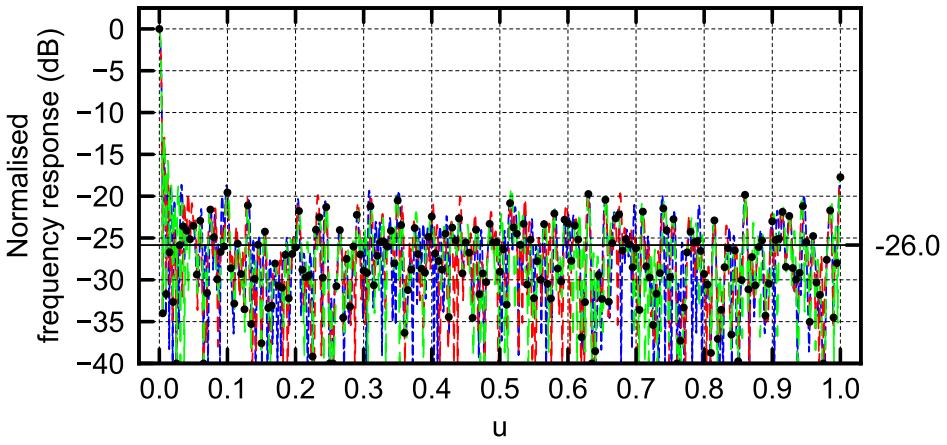

(a) Random window

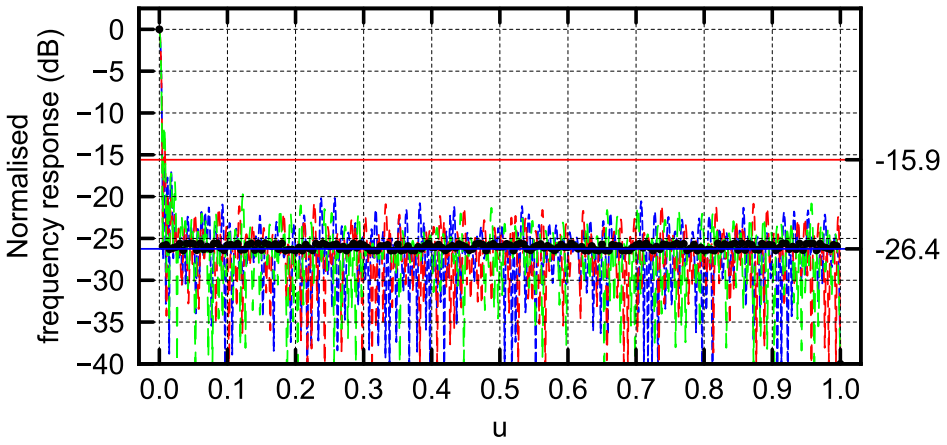

(b) ADS window

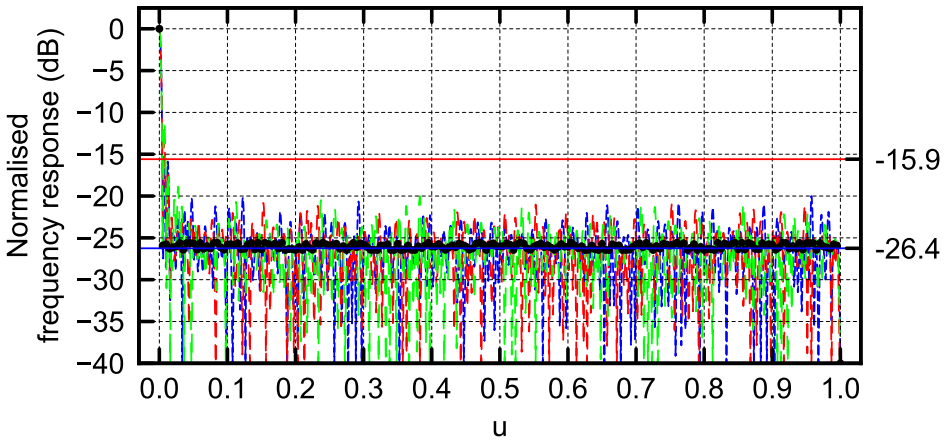

(c) Complementary ADS window

Figure A.6. Frequency responses resulting from windows based on the ADS (400,200,99,100).

ADDENDUM B ADDITIONAL EMPIRICAL RESULTS

The empirical results for binary windows based on several difference sets (DSs) and almost difference sets (ADSs) are given here. Specifically, the results for windows based on the DSs (31,15,7) and (101,25,6) and the ADSs (30,15,7,22) and (101,26,6,50). These results are meant to illustrate the effects of set length and even and odd values of V on detection performance. The results for simulated signals and generated practical signals are given; no results were generated for recorded practical signals, because the sets are so short. The work described in this addendum forms part of a submitted journal paper [10].[1]

The results here support those found in Chapter 5. While both random binary windows and windows based on DSs/ADSs produce interference curves (compared to filled binary windows), for the DS/ADS windows that interference is well controlled and at least moderately predictable compared to the random windows. The non-uniform nature of the interference from the random binary windows gives the appearance of "false targets" in certain frequency bins where the sidelobe level (SLL) is high. Additionally, shorter sets (smaller values of V) produce less interference curves, but their frequency resolution will be lower. The more samples (the larger the value of K) the better the detection performance in the signal bin (a larger P_D can be achieved at a lower signal-to-noise ratio (SNR)), and the larger the separation in SNR between the signal curve and the interference curves (which is one of the primary benefits of complementary DSs/ADSs).

[1]Portions of this chapter are reprinted, with permission, from [10]. ©2018 IEEE.

B.1 SIMULATED RESULTS

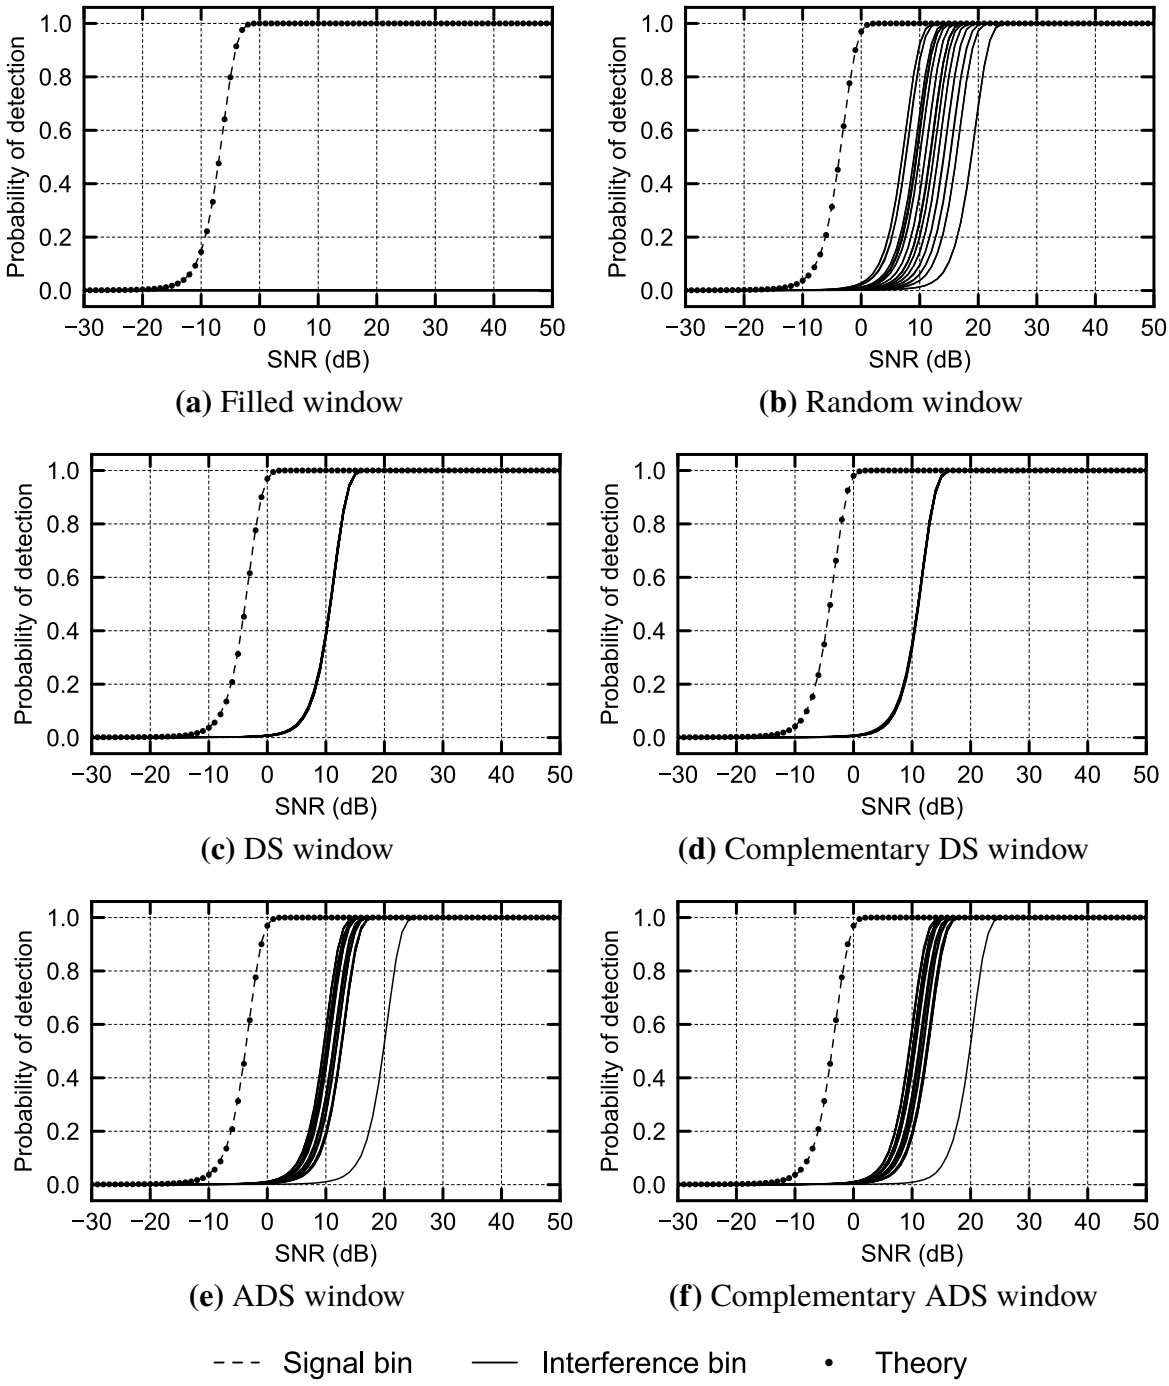

Figure B.1. Probability of detection over SNR curves for the simulated signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable. (Reprinted, with permission, from [10]. ©2018 IEEE.)

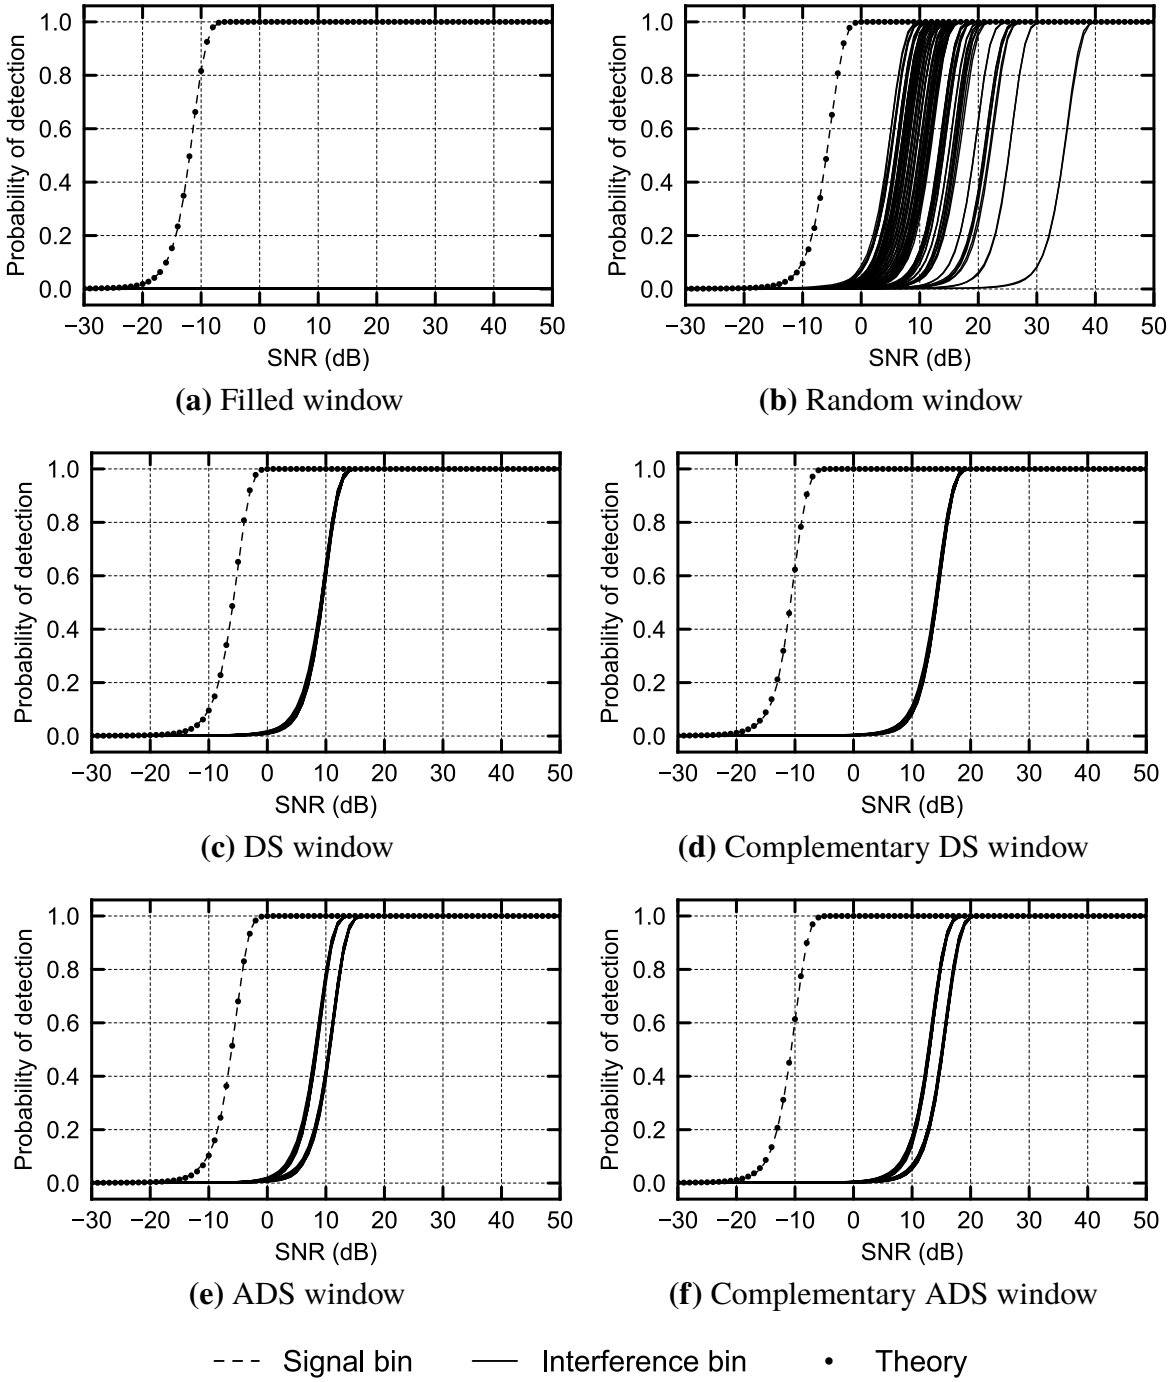

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

(e) ADS window

(f) Complementary ADS window

Figure B.2. Probability of detection over SNR curves for the simulated signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable. (Reprinted, with permission, from [10]. ©2018 IEEE.)

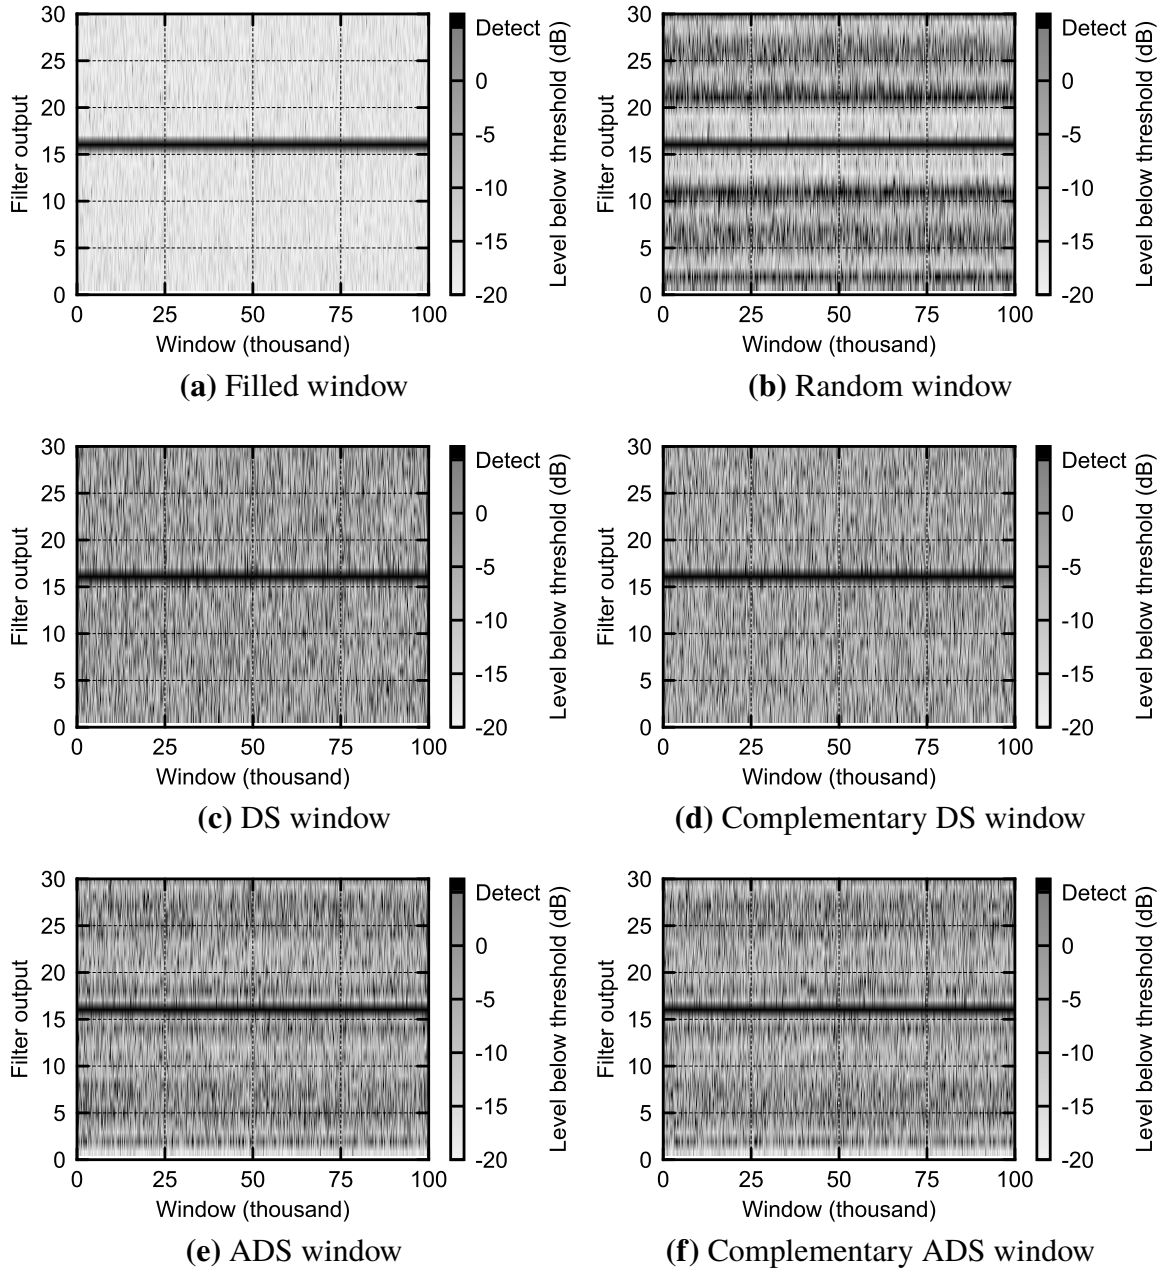

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

(e) ADS window

(f) Complementary ADS window

Figure B.3. Test statistics for SNR = 8 dB for the simulated signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS).

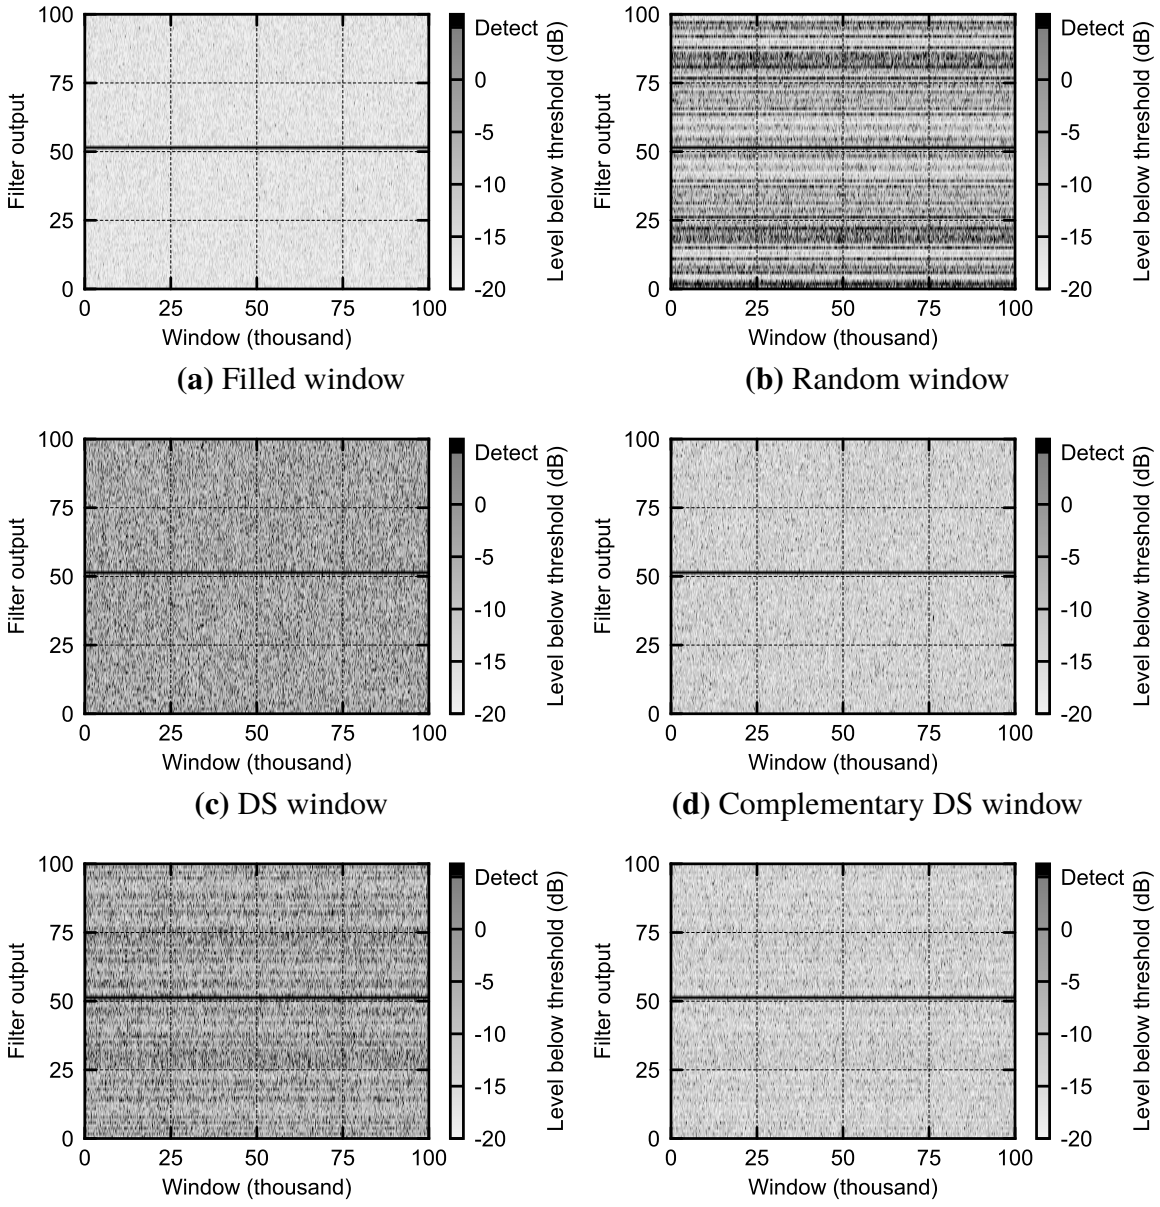

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

(e) ADS window

(f) Complementary ADS window

Figure B.4. Test statistics for SNR = 6 dB for the simulated signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS).

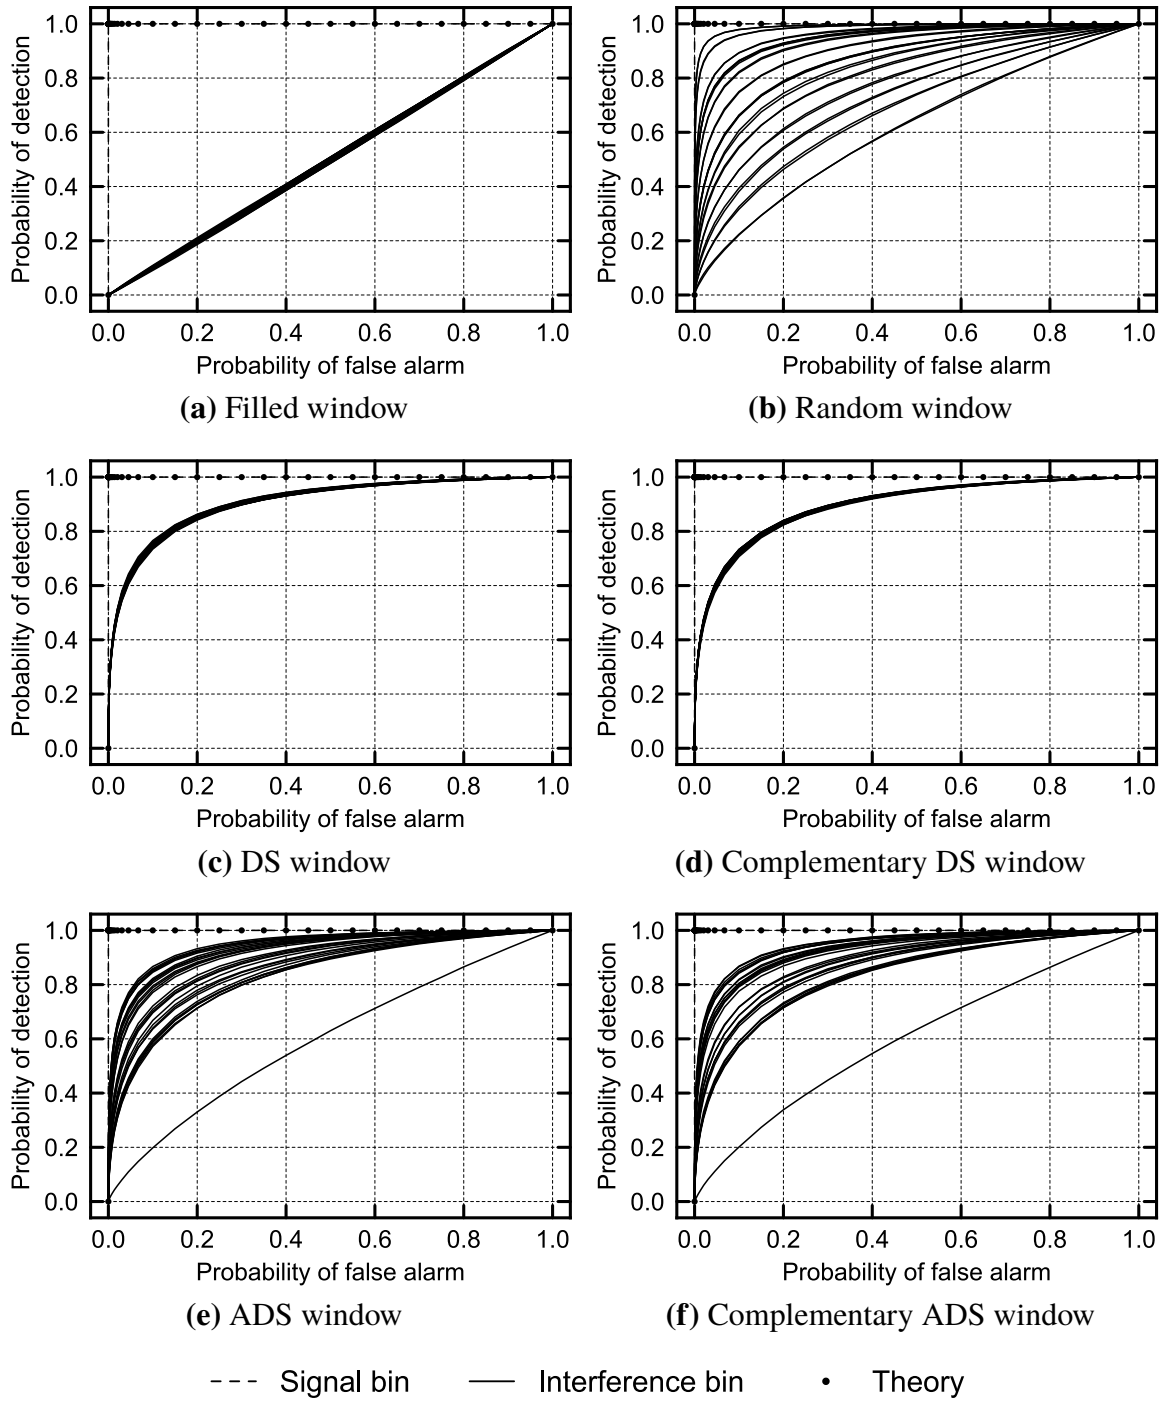

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

(e) ADS window **(f)** Complementary ADS window

Figure B.5. receiver operating characteristics (ROCs) for SNR = 8 dB for the simulated signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

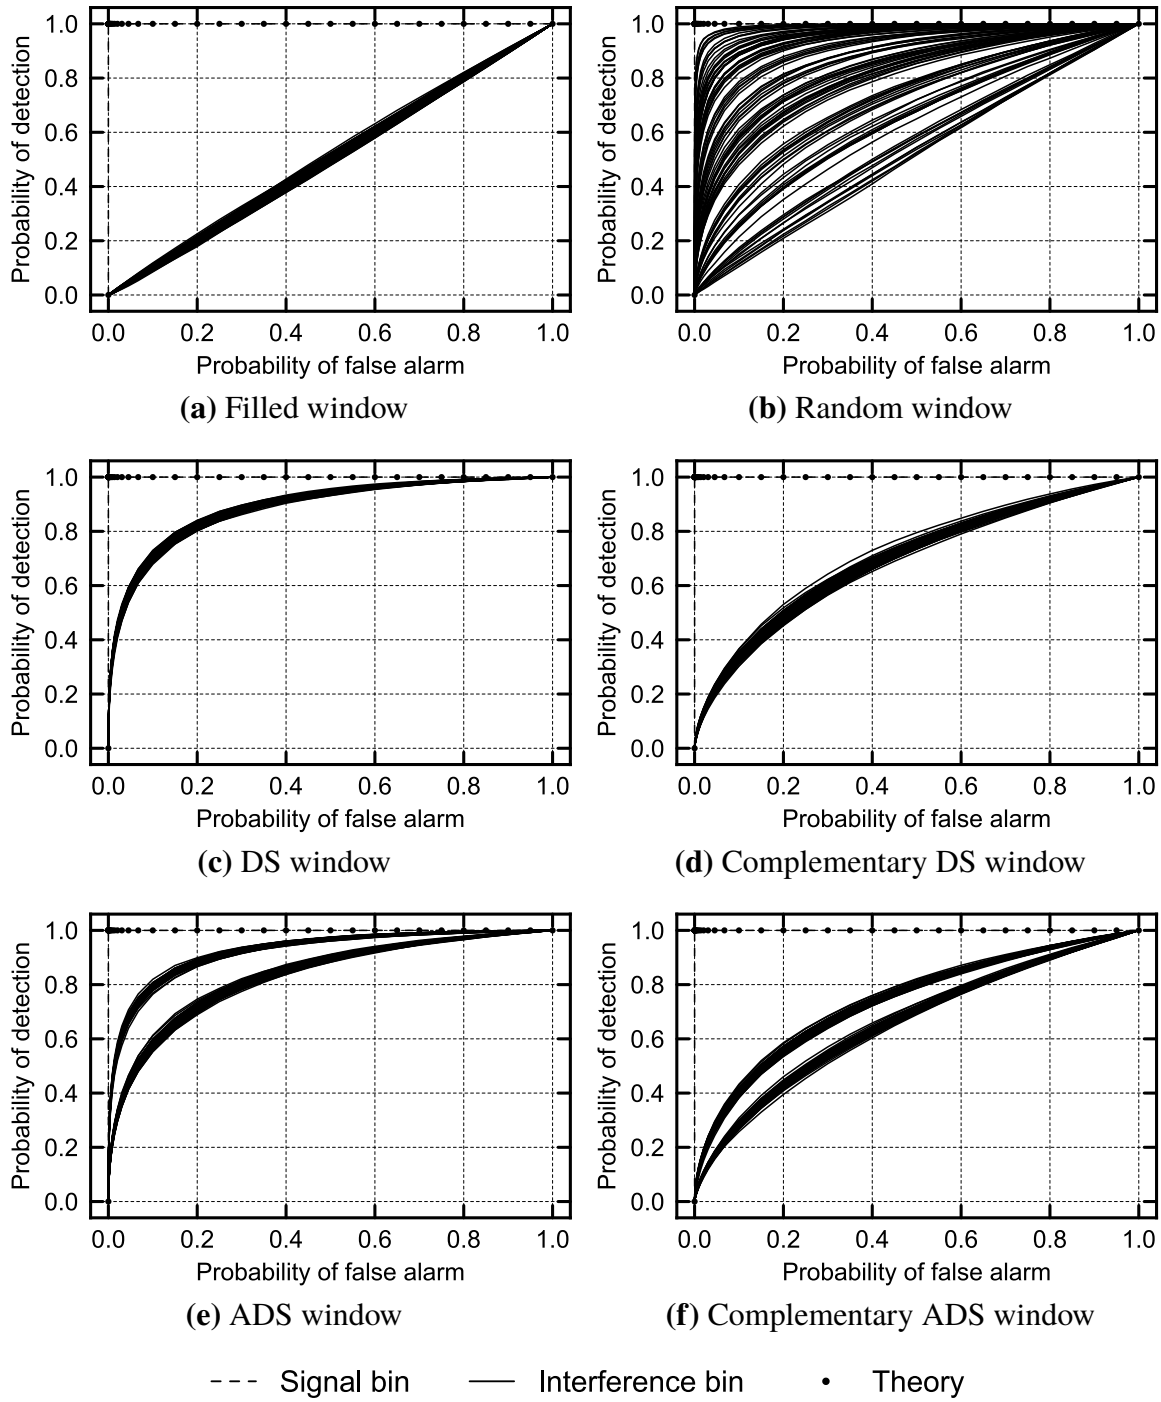

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

(e) ADS window **(f)** Complementary ADS window

Figure B.6. ROCs for SNR = 6 dB for the simulated signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

B.2 PRACTICAL RESULTS: GENERATED SIGNALS

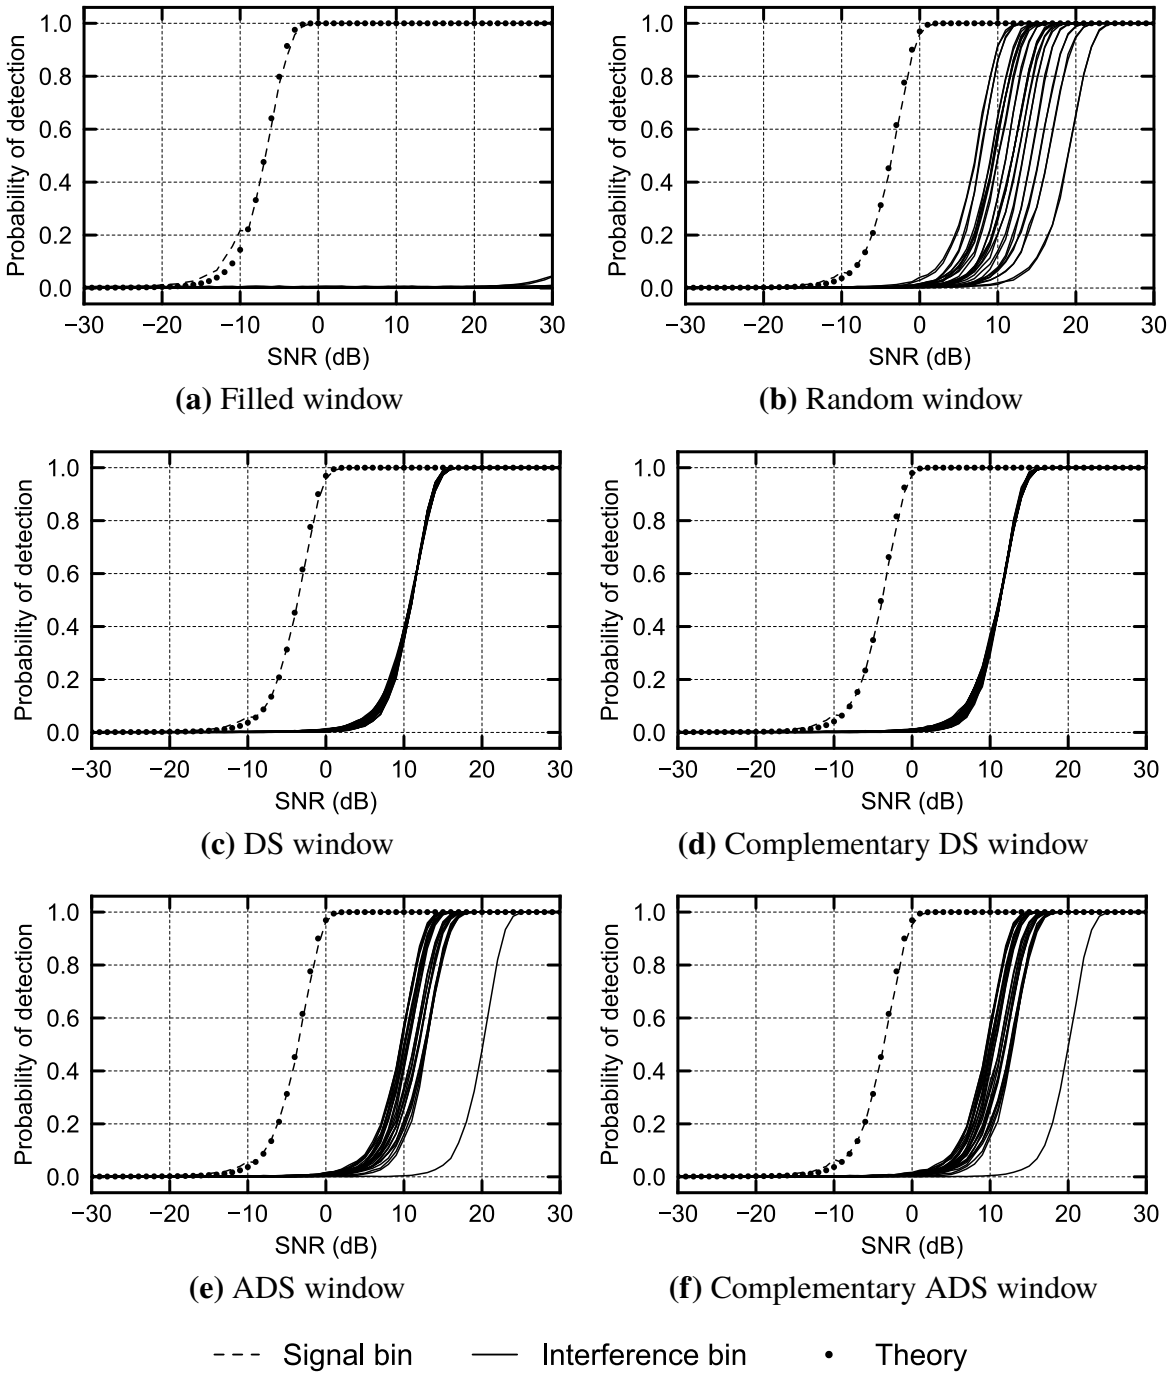

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

(e) ADS window

(f) Complementary ADS window

Figure B.7. Probability of detection over SNR curves for the generated practical signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

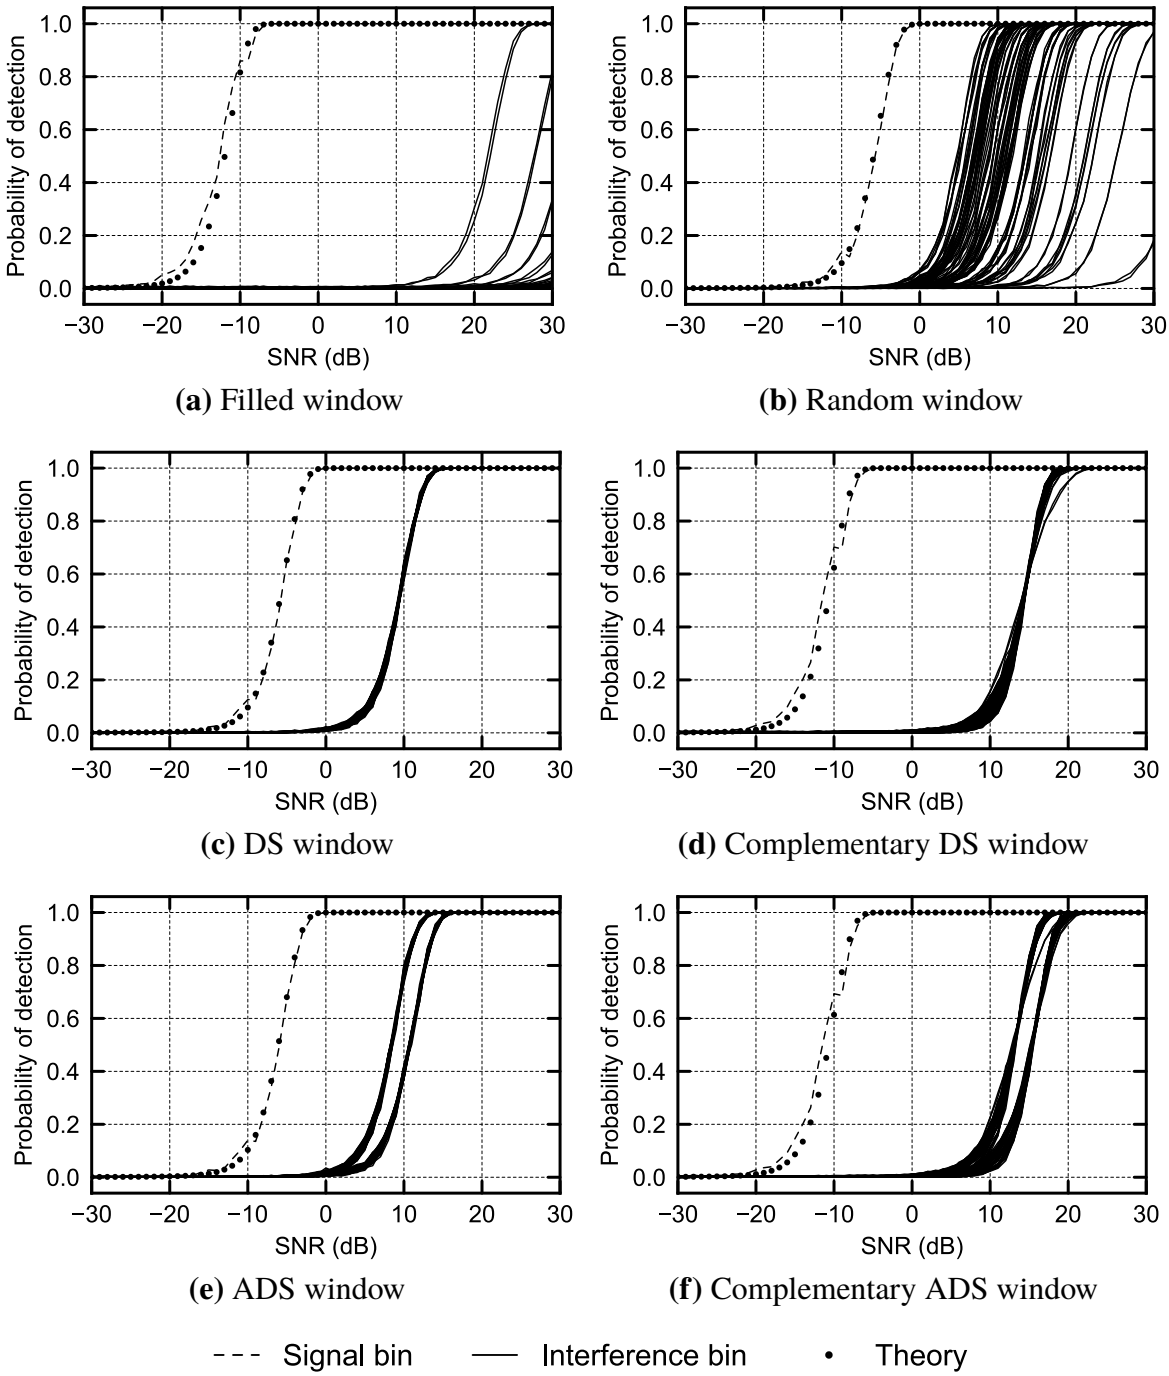

Figure B.8. Probability of detection over SNR curves for the generated practical signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

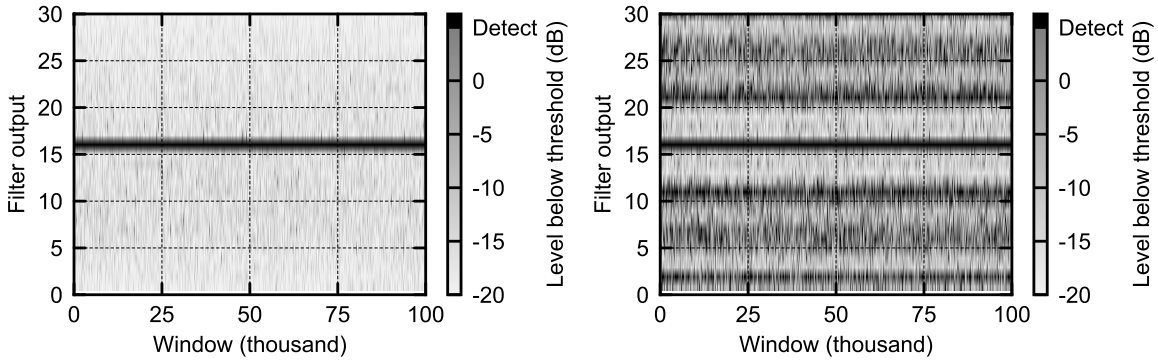

(a) Filled window

(b) Random window

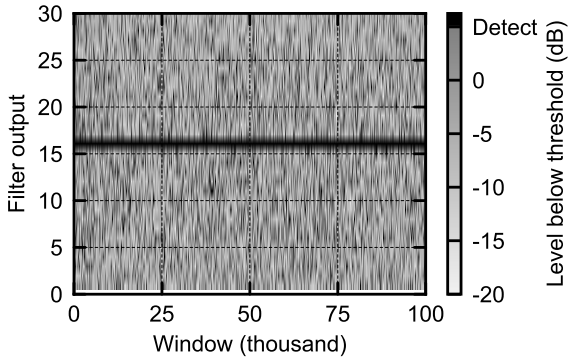

(c) DS window

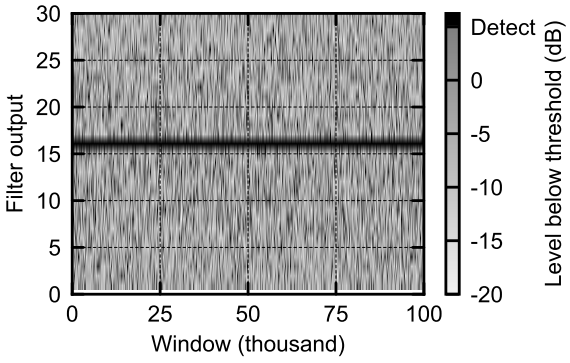

(d) Complementary DS window

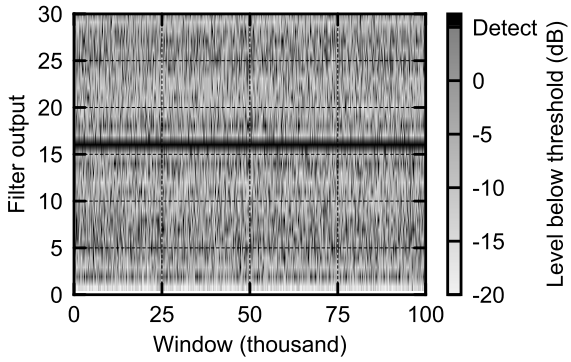

(e) ADS window

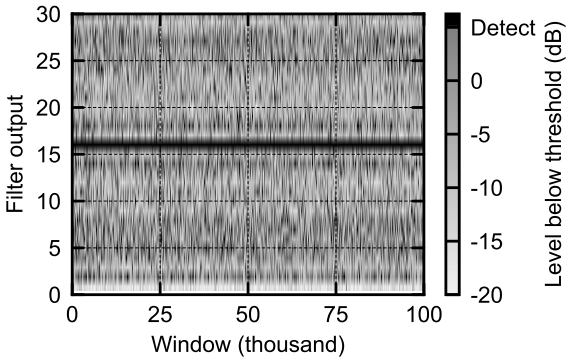

(f) Complementary ADS window

Figure B.9. Test statistics for SNR = 8 dB for the generated practical signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS).

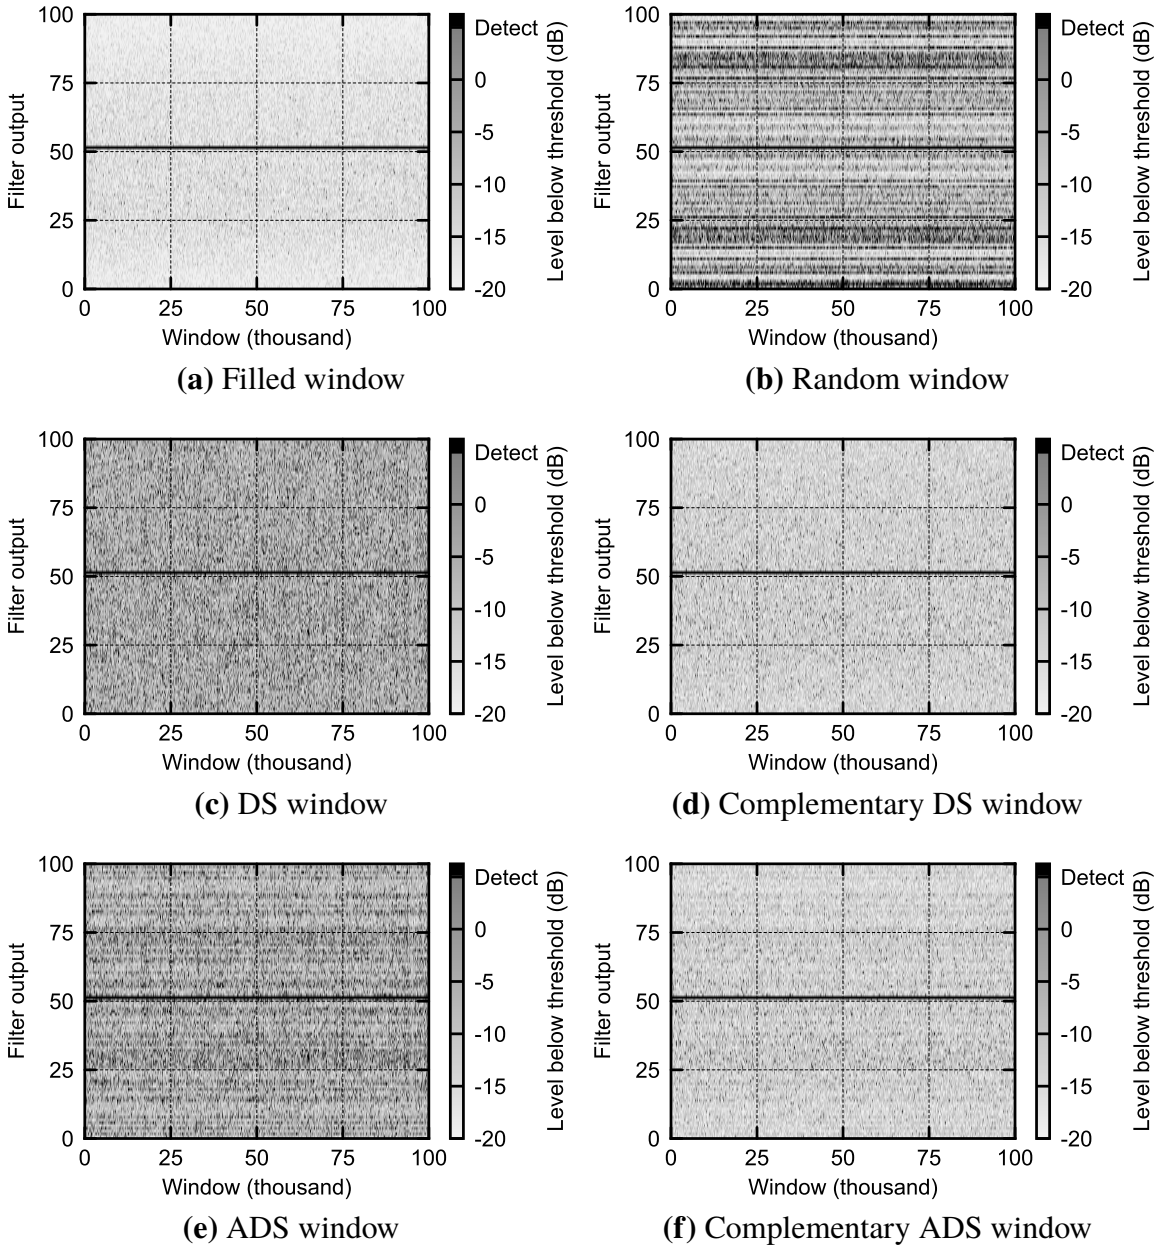

(a) Filled window

(b) Random window

(c) DS window

(d) Complementary DS window

(e) ADS window

(f) Complementary ADS window

Figure B.10. Test statistics for SNR = 6 dB for the generated practical signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS).

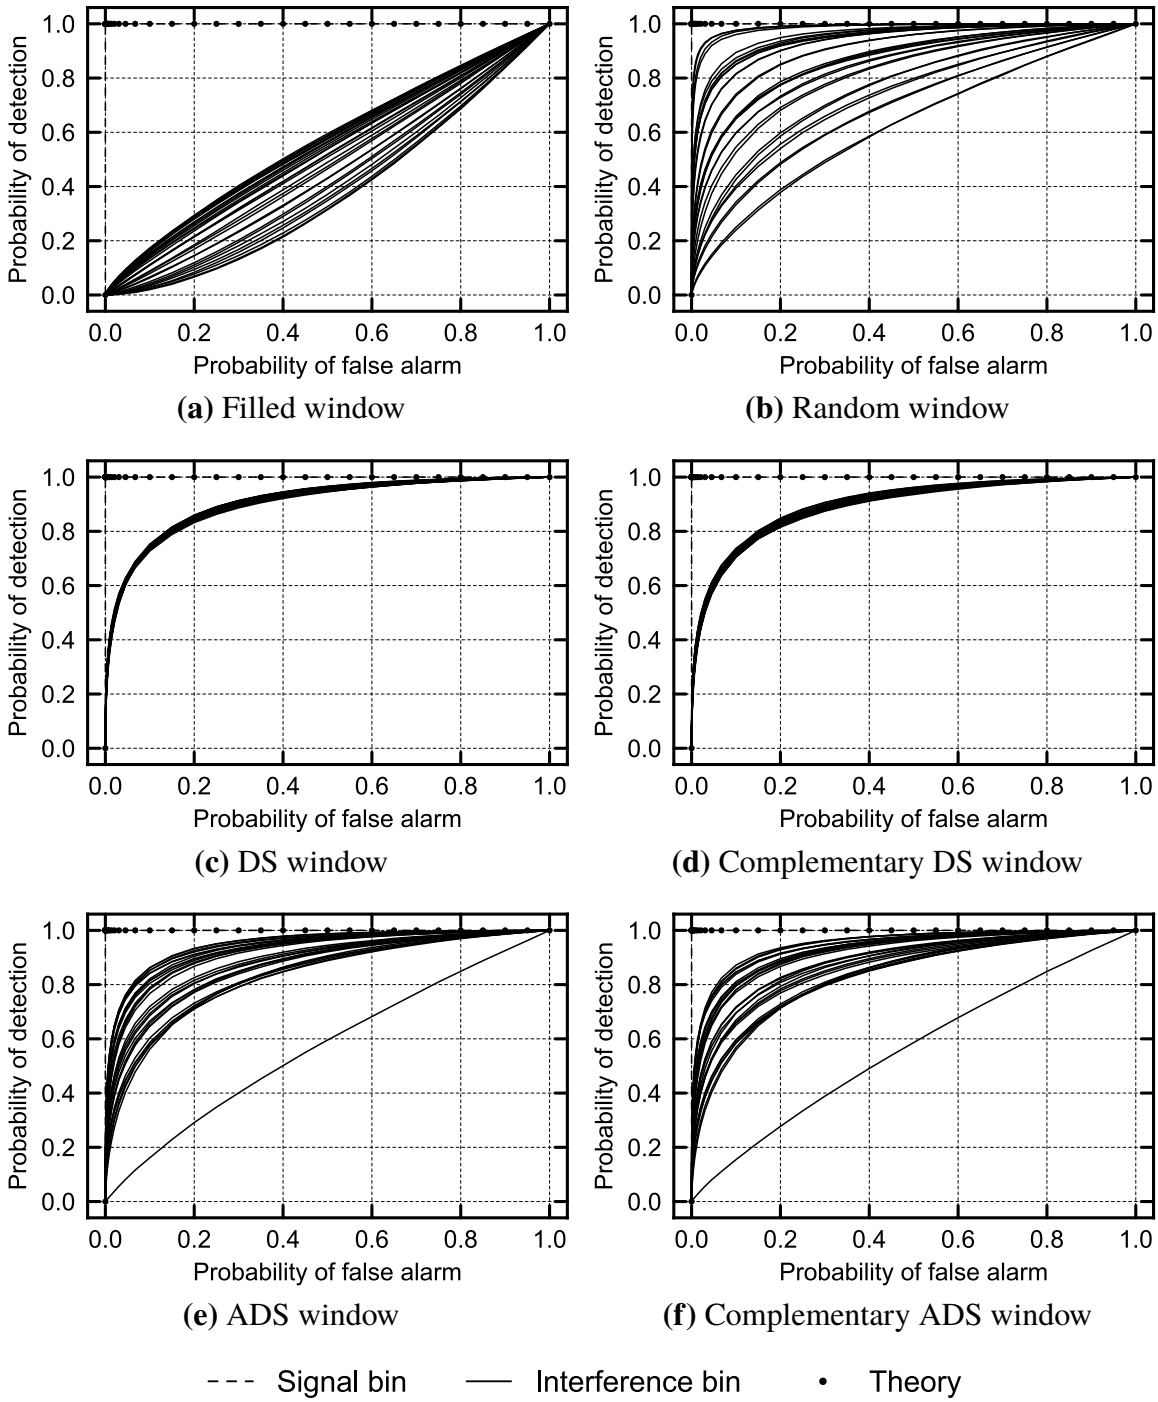

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

(e) ADS window **(f)** Complementary ADS window

Figure B.11. ROCs for SNR = 8 dB for the generated practical signal resulting from windows based on the DS (31,15,7) and the ADS (30,15,7,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.

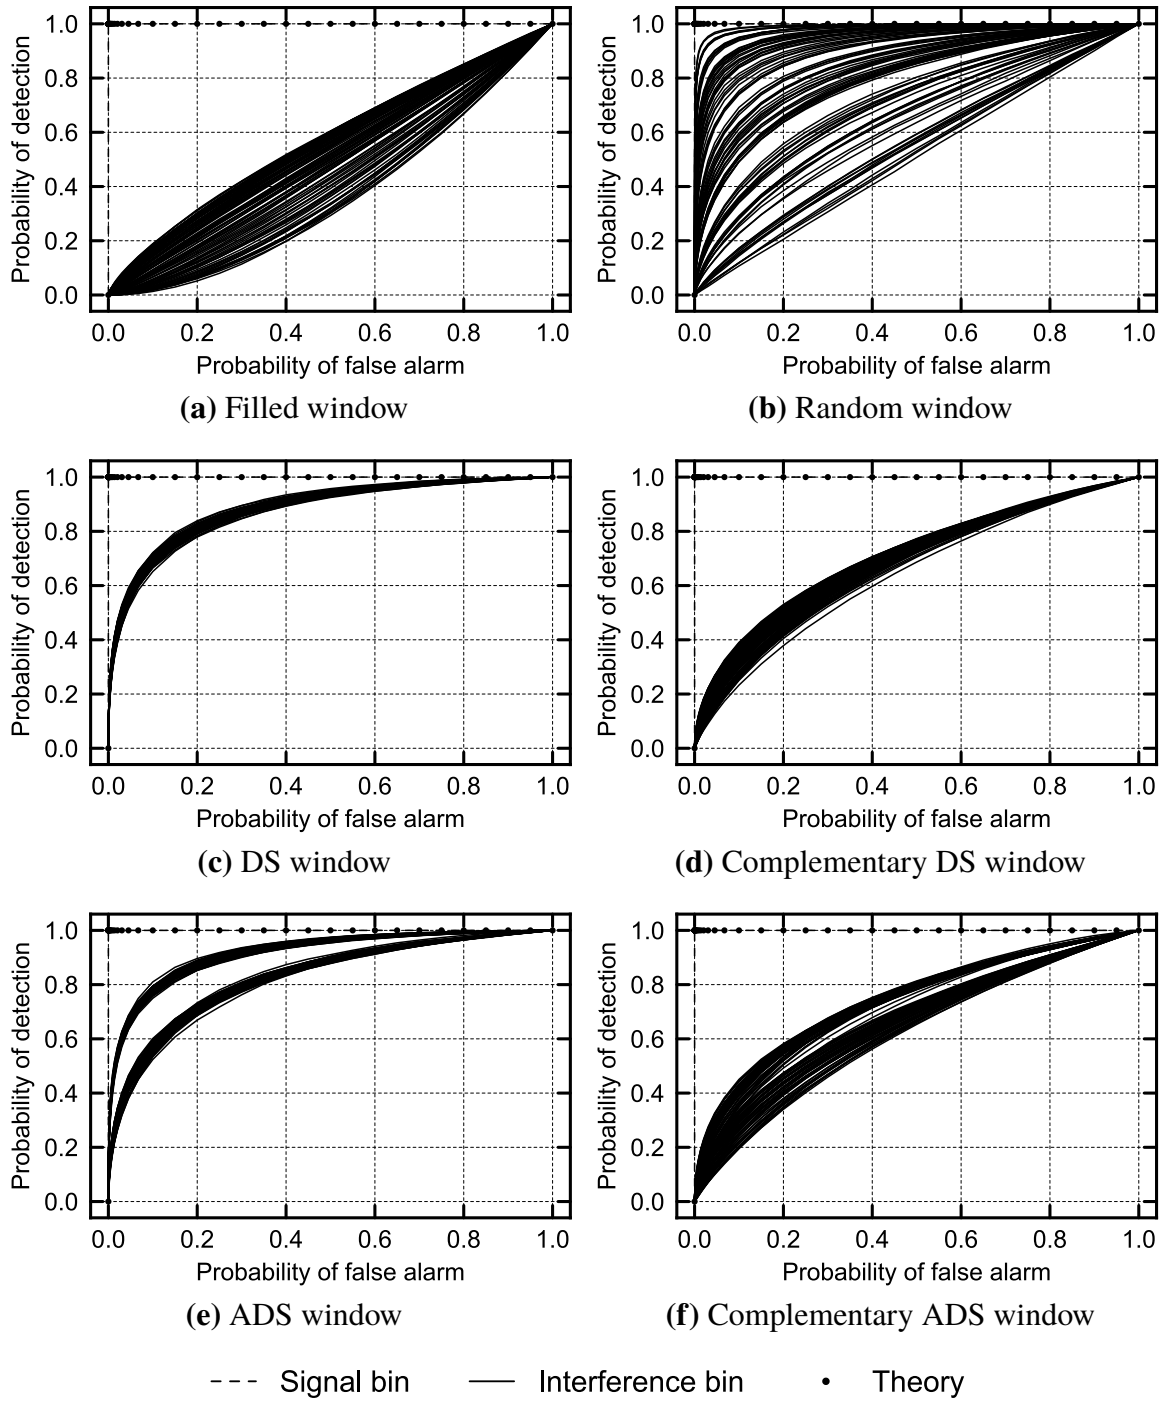

(a) Filled window **(b)** Random window

(c) DS window **(d)** Complementary DS window

(e) ADS window **(f)** Complementary ADS window

Figure B.12. ROCs for SNR = 6 dB for the generated practical signal resulting from windows based on the DS (101,25,6) and the ADS (101,26,6,22) (The filled and random windows are based on the DS). Note that the "Signal bin" and "Theory" curves are so similar that they are indistinguishable.