

STATISTICAL SIGNAL DETECTION WITH
MULTI-SENSOR AND SPARSITY

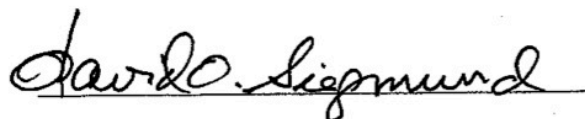
A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL
ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Yao Xie

March 2012

Yao Xie

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.



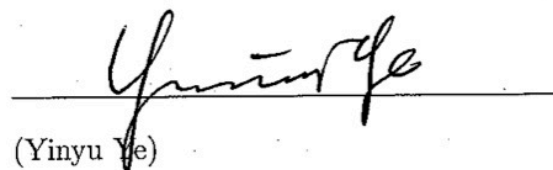
(David O. Siegmund) Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.



(Andrea J. Goldsmith)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.



(Yinyu Ye)

Approved for the Stanford University Committee on Graduate Studies

Abstract

The objective of many signal processing problems is to detect signals buried in a noisy background. Many of these signal detection problems have sparsity structure that can be exploited to reduce noise or reduce complexity. This dissertation will focus on two such problems: multi-sensor sequential change-point detection and multiuser detection, and it will present new methods to exploit sparsity: the mixture sequential detection procedure and reduced-dimension multi-user detection (RD-MUD).

In multi-sensor change-point detection, sensors are deployed to monitor the abrupt emergence of a change-point. The change-point is an event that affects the observations of a subset of sensors simultaneously. Typically the subset of sensors that are affected is unknown, and the level-of-affectedness of each affected sensor is also unknown. The goal is to detect the change-point as soon as possible once it occurs, and rarely make false detections if it does not occur. An empirical observation is that the number of affected sensors is usually small compared to the total number of sensors. This is a form of sparsity. For this problem, traditional methods have not exploited the sparsity structure: they either assume all sensors are affected by the change-point and use observations from all sensors and, hence, include too much noise from the observations of the unaffected sensors, or assume only one sensor is affected, use only

observations from the affected sensor with the highest level-of-affectedness, and ignore observations from other affected sensors. We develop a mixture procedure that exploits this sparsity. In particular, we model this sparsity by assuming that each sensor has a small probability p_0 to be affected by the change-point. The value of p_0 is a guess for p . Based on this model, we form a mixture log generalized likelihood ratio (GLR) statistic and present a mixture detection procedure. The mixture statistic essentially applies a non-linear weighting function, which is parameterized by p_0 , on the log GLR statistic of each sensor before combining them. This nonlinear weighting function automatically emphasizes the statistics from the sensors that are affected by the change-point and suppresses those from the sensors that are not affected. We derive a theoretical approximation for the false alarm rate, which is captured by the average-run-length (ARL), and a theoretical approximation for the expected detection delay. The accuracy of these approximations is verified by numerical studies. We also demonstrate that the mixture procedure is robust against the lack of knowledge of p . Numerical studies compare the new mixture procedure with other proposed procedures.

The multiuser detection (MUD) problem arises in multiuser communication systems, where multiple users communicate simultaneously with a receiver. The receiver receives a signal consisting of a set of known waveforms modulated by the information symbols of the users that is contaminated by noise. The receiver has to determine which users are active and their information symbols. The conventional solutions to the MUD problem all consist of a matched-filter bank (MF-bank) front-end, followed by digital signal processing. The MF-bank front-end uses a set of correlators, where each one correlates the received signal with a signature waveform. Hence the number

of correlators used in the conventional matched-filter bank is equal to the number of users in the system. We present a reduced-dimension multiuser detector (RD-MUD) structure that significantly decreases the number of required correlation branches in the receiver front-end, while still achieving performance similar to that of the conventional matched-filter (MF) bank. RD-MUD exploits the fact that the number of active users is typically small relative to the total number of users in the system and relies on ideas of analog compressed sensing to reduce the number of correlators. The correlating signals used by each correlator are chosen as an appropriate linear combination of the users' spreading waveforms, which in turn are chosen from a large class of spreading codes. We derive the probability-of-error when using two methods for recovery of active users and their transmitted symbols: the reduced-dimension decorrelating (RDD) detector, which combines subspace projection and thresholding to determine active users and sign detection for data recovery; and the reduced-dimension decision-feedback (RDDF) detector, which combines decision-feedback orthogonal matching pursuit for active user detection and sign detection for data recovery. We identify conditions such that error is dominated by active user detection. We then show that the number of correlators needed to achieve a small probability-of-error under these conditions is on the order of the logarithm of the number of users in the system for a given projection method based on random discrete Fourier transform (DFT) matrices. Thus, RD-MUD has significantly fewer correlators than the number of correlators required by MUD using the conventional MF-bank. Our theoretical results take into consideration the effects of correlated signature waveforms as well as near-far issues. The theoretical performance results for both detectors are validated with numerical simulations.

Acknowledgements

I would like to thank my advisor David O. Siegmund for his guidance and friendship throughout my Ph.D. studies. He has not only shared his extensive knowledge with me, but has also demonstrated the meaning of true scholarship. He has challenged me with many interesting problems and inspired me in many ways. He would always ask me to doublecheck my numerical results using various methods such that they are perfectly accurate. I remember every time I asked him whether my paper was ready for submission, he would tell me to hold onto it and push the theoretical depth to a higher level. He said that “If you are truly interested in this area, you have a whole life to work on it. There is no hurry.” He will continue to be my friend and academic role model as I continue along my career path.

I also want to thank my co-advisor Andrea J. Goldsmith for her guidance and support. Her rigor in doing research has truly influenced me and allowed me to develop my critical thinking skills. I have deeply appreciated her great passion for and dedication to research. She gave me countless revisions and would ponder problems so constantly that she would give me hand-written comments on a hotel notepad from her travels. She has taught me how to be a well-rounded researcher, both in producing good quality work and presenting it dynamically and with great clarity through writing and presentations. She has influenced and encouraged me in various

aspects of my career goals. She has taught me to equip myself with knowledge, diligence, and passion before tackling a problem.

I am also very grateful to my committee members (in alphabetical order) - J. Michael Harrison, Benjamin Van Roy and Yinyu Ye. Their comments and questions have helped to improve the quality of my dissertation substantially. Moreover, they have provided me with much encouragement and guidance through classes and personal communications. Professor J. Michael Harrison has taught me queueing theory and stochastic control, Professor Yinyu Ye has taught me various methods in formulating and solving optimization problems, and Professor Benjamin Van Roy has taught me dynamic programming. All of their expertise has been greatly beneficial to my research.

I want to thank the Department of Electrical Engineering. The students, faculty and staff have provided me with a wonderful study and work environment. My thanks also go to the Office of the Vice Provost for Graduate Education (VPGE) at Stanford for their financial support through the General Yao-Wu Wang Stanford Graduate Fellowship, and to the Interconnected Focus Center for their financial support through a research assistantship. I also want to thank Dr. Claude Reichard, Dr. Phil Hubbard, and the Stanford Hume Writing Center and Technical Communication Program for their great help in developing my writing skills for this dissertation.

I also want to thank my mentors during my Master's studies at the University of Florida - Jian Li, and many professors during my Bachelor's studies at the University of Science and Technology of China (USTC). They helped me find my true interests and drew me into the fascinating world of statistical signal processing.

I would like to thank my friends at and outside Stanford for their friendship. I

want to thank Li for his constant love and support and for spending his life with me. Last but not least, I want to thank my parents for their love and care during my pursuit of a Ph.D..

Contents

Abstract	iv
Acknowledgements	vii
1 Introduction	4
1.1 Multi-Sensor Change-Point Detection	6
1.2 Reduced-Dimension Multiuser Detection	11
2 Background	18
2.1 Classic One-Sensor Change-Point Detection	18
2.2 Multi-User Detection	25
2.3 Compressed Sensing	30
3 Multi-Sensor Change-Point Detection	34
3.1 Model and Formulation	35
3.2 Mixture Procedure	37
3.3 Other Detection Procedures	42
3.3.1 Maximizing Over Subset of Sensors	42
3.3.2 Other Procedures	43

3.4	Performance of Mixture Procedure	46
3.4.1	Average Run Length (ARL)	48
3.4.2	Expected Detection Delay	51
3.4.3	Accuracy of Approximations for Mixture Procedure	59
3.5	Numerical Examples	64
3.5.1	Robustness to Choice of p_0	64
3.5.2	Effect of Window Size m_1	66
3.5.3	Comparing Procedures	68
3.6	Parallel Mixture Procedure	73
3.7	Profile-Based Procedure	75
3.7.1	Likelihood Ratio Statistic	76
3.7.2	Theoretical ARL of Profile-Based Procedure	78
3.7.3	Numerical Examples	80
3.8	Conclusions	88
4	Reduced-Dimension Multi-User Detection	89
4.1	System Model	91
4.1.1	Conventional MUD	94
4.2	Reduced-Dimension MUD (RD-MUD)	99
4.2.1	RD-MUD: Front-End	99
4.2.2	RD-MUD: Detectors	102
4.2.3	Choice of \mathbf{A}	113
4.3	Performance of RD-MUD	114
4.3.1	Single Active User	114
4.3.2	Noise Amplification of Subspace Projection	117

4.3.3	Coherence Based Performance Guarantee	119
4.3.4	Bounding Probability-of-Error of RDD and RDDF	124
4.3.5	Comparison with Existing Bounds	126
4.3.6	Lower Bound on Number of Correlators	128
4.4	Numerical Examples	131
4.4.1	Noise-Free Scenario	131
4.4.2	Noisy Scenario	134
4.5	Conclusions	138
5	Summary and Future Work	142
5.1	Summary	142
5.2	Future Work	144
A	Proof for Multi-Sensor Change-Point Detection	147
A.1	Proof of Lemma 4	148
A.2	Multiple Overlapping Sources	151
B	Proof for RD-MUD	154
B.1	Covariance of RD-MUD Output Noise	155
B.2	Derivation of RD-MUD MMSE	156
B.3	Proof of Theorem 1	157
B.4	Proof of Lemma 6	168
	Bibliography	171

List of Tables

2.1	Complexity-per-bit of MUD detectors (N users, K active users). . . .	30
2.2	Required channel-state-information (CSI) by MUD detectors.	30
3.1	Average run length (ARL) of the mixture procedure with $m_1 = 200$. The Monte Carlo results are obtained from 500 trials.	61
3.2	Expected detection delay of a mixture procedure with $m_1 = 200$ such that the theoretical approximation of ARL is equal to 5000. The Monte Carlo results are obtained from 500 trials.	62
3.3	Thresholds for procedures with $m_1 = 200$ such that their ARLs equal 5000. The Monte Carlo results are obtained from 500 trials.	69
3.4	Expected Detection Delays for various procedures with $N = 100$ and $m_1 = 200$, obtained from 500 Monte Carlo trials. The mixture proce- dure has $p_0 = 0.1$. Mei's and the Modified TV procedures use $\tilde{\mu}_n = 1$. The thresholds of these procedures with ARL 5000 are listed in Table 3.3. We set $\mu_n = 1$ for $n \in \mathcal{N}_a$ in the first column, and $\mu_n = 0.7$ for $n \in \mathcal{N}_a$ for the second column. For the third column, we choose post-change means independently and uniformly from the interval $[0.5,$ $1.5]$	70

3.5	Expected detection delays for mixture procedure versus the profile-based procedure for the small example. The mixture procedure uses $p_0 = 0.1$, and the profile-based procedure assumes correct profile-function. They both use $m_1 = 200$. The thresholds are obtained by Approximation 1 for the mixture procedure, and by Monte Carlo simulation using <i>the second method</i> for the profile-based procedure, respectively. The Monte Carlo results are obtained from 500 trials. . .	84
3.6	Expected detection delays for mixture procedure versus the profile-based procedure for the small example. The mixture procedure uses $p_0 = 0.01$, and the profile-based procedure assumes correct profile-function. They both use $m_1 = 200$. The thresholds are obtained by Approximation 1 for the mixture procedure, and by Monte Carlo simulation using <i>the second method</i> for the profile-based procedure, respectively. The Monte Carlo results are obtained from 500 trials. . .	86
4.1	RD-MUD Detection	106

List of Figures

1.1	The five sequences in the picture demonstrate observations from five sensors. The sensor index is n and the time index is t . A change-point happens at time $\kappa = 25$, with location marked by a red star, illustrates an instance where the change-point affects the observations of $M = 3$ out of $N = 5$ sensors. The post-change means of the observed changes are different for these affected sensors.	7
1.2	One example of change-point problem with spatial structure. A change-point that occurs abruptly changes the distributions of a subset of sensors simultaneously. The observations of the affected sensors have larger means. In this example, the different post-change means, or the level-of-affectedness, of the affected sensors are spatially correlated. The sensors that are spatially close tend to have similar post-means.	8
1.3	Multiuser communication network.	12
1.4	Diagram of conventional MUD using MF-bank.	13
3.1	Plot of $g(x; p_0) = \log[1 - p_0 + p_0 e^{(x^+)^2/2}]$ for different p_0	40

3.2	Cumulative distribution function (cdf) of T_{mix} from theoretical approximation versus from numerical simulation. The mixture procedure has $p_0 = 0.1$, $m_1 = 200$, and $b = 19.5$, which corresponds to the a theoretical approximation of ARL that equals 5000. The Monte Carlo results are obtained from 500 Monte Carlo trials.	60
3.3	Comparison of theoretical and simulated expected detection delay of a mixture procedure. For simulation, we use 500 Monte Carlo trials, with $p_0 = p$, $m_1 = 200$, and different b such that the theoretical approximation of ARL is equal to 5000.	62
3.4	Comparison of theoretical and simulated expected detection delay of a mixture procedure. For simulation, we use 500 Monte Carlo trials, with $p_0 = p$, $m_1 = 200$ and $b = 19.5$ such that the theoretical approximation of ARL is equal to 5000.	63
3.5	Comparison of the expected detection delay from Approximation 2, when there is no mismatch, $p_0 = p$, and when $p_0 = 0.1$ so that there is a mismatch between p_0 and p . The thresholds are chosen such that the theoretical approximation of ARL is equal to 5000.	65
3.6	Expected detection delay versus m_1 for the mixture procedure when the theoretical approximation of ARL is equal to 5000. The expected detection delay is obtained from Monte Carlo simulation. The Monte Carlo results are obtained from 500 trials.	67

3.7	Comparison of expected detection delays for various procedures with window size $m_1 = 200$, and (a): $\mu_n = 1$, $n \in \mathcal{N}_a$, (b): $\mu_n = 0.7$, $n \in \mathcal{N}_a$, (c): μ_n , $n \in \mathcal{N}_a$, arbitrarily and uniformly chosen from an interval $[0.5, 1.5]$. The mixture procedure uses $p_0 = 0.1$. Mei's and the modified TV procedures use $\tilde{\mu}_n = 1$. Monte Carlo results are obtained from 500 trials.	72
3.8	Expected detection delay of the parallel mixture procedure with $p_1 = 0.2$ and $p_2 = 0.03$ versus that of a single mixture procedure with $p_0 = 0.1$. The ARLs of the parallel procedure and the single procedure are both 6577. The Monte Carlo results are obtained from 500 trials. . .	74
3.9	Four non-overlapping sources observed by sensors located on a 10 by 10 grid ($N = 100$). The sources have Gaussian profile function given in (3.70) with $\beta = 1/4$. The outer circle for the contour for each source corresponds to 0.1 of the peak value. The dots represent location of the sensors.	83
3.10	Three non-overlapping sources observed by sensors located on a 25 by 25 grid ($N = 625$). The sources have Gaussian profile function given in (3.70) with $\beta = 1$. The outer circle for the contour for each source corresponds to 0.1 of the peak value. The dots represent location of the sensors.	87
4.1	Problem set up of MUD.	94
4.2	Front-end of conventional MUD using MF-bank.	95
4.3	Front-end of RD-MUD.	96
4.4	The diagram of a linear detector.	96

4.5	The diagram of a non-linear detector.	97
4.6	The diagram MUD detector with prewhitening.	109
4.7	A MUD problem with $N = 3$, $M = 2$, $K = 1$ and the received signal is due to the first user. Illustration of the detection by (a) the single-user detection of the MF-bank, and (b) the RD-MUD. The projection of the signature waveform onto the projection subspace results in signature waveform vectors $\{\hat{\mathbf{s}}_n\}$	117
4.8	A MUD problem with $N = 3$, $M = 2$, $K = 1$, and the received signal is due to the first user. Illustration of the detection by the RD-MUD. The mapping of the signature waveform is denoted by $\hat{\mathbf{s}}_n$	118
4.9	Performance of Algorithm-I, in the absence of noise, for $K = 2$, different N , P_e versus M , with marked points equal to: $\log N$, $2 \log N$, $4 \log N$, and $8 \log N$. When $M = 8 \log N$, $P_e = 0$	132
4.10	In the absence of noise, P_e versus M , for $N = 100$, using Algorithm-I for different K	134
4.11	In the absence of noise, P_e versus M , for $N = 100$, using Algorithm-II compared with Algorithm-I when $K = 10$	135
4.12	In the absence of noise, P_e versus M of the RDD detector using the random partial DFT versus using the Gaussian random matrices for $N = 100$ and $K = 6$	136
4.13	Performance of Algorithm-I, P_e versus M for different SNRs, when the signature waveforms are orthogonal, i.e., $\mathbf{G} = \mathbf{I}$. When SNR is greater than 15 dB, the probability-of-error of the MF-bank is less than 10^{-4}	137

4.14	Comparison of P_e versus M , without and with the noise whitening transform, when $N = 100$, $K = 2$, and for two different cases of \mathbf{G} . The conventional decorrelating detector with the noise whitening transform in Fig. 4.14, and the conventional decorrelating detectors with and without the noise whitening transform in Fig. 4.15 have probability-of-error less than 10^{-4}	138
4.15	Comparison of P_e versus M , with and without using pre-whitening, when $N = 100$, $K = 2$, when $\lambda_{\max}(\mathbf{G}^{-1})$ is small. The MF-bank decorrelators with and without whitening have probability-of-error less than 10^{-4} and hence are not shown in the pictures.	139
4.16	Comparison of the conditional probability of error $\mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b} \hat{\mathcal{I}} = \mathcal{I}\}$, for RD-MUD linear detectors, when $\lambda_{\max}(\mathbf{G}^{-1})$ is large. The conditional probability of error for MF-bank decorrelator is less than 10^{-4} and hence is not shown in the pictures.	140
4.17	Comparison of the conditional probability of error $\mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b} \hat{\mathcal{I}} = \mathcal{I}\}$, for RD-MUD linear detectors: RDD, RD-LS, and RD-MMSE. The RD-MMSE detector has the smallest conditional probability of symbol error.	141

Acronyms

ARL Average Run Length

BER Bit Error Rate

BPSK Binary Phase Shift Keying

cdf Cumulative Distribution Function

CUSUM Cumulative Sum Control Chart

DF Decision Feedback

DF-OMP Decision Feedback Orthogonal Matching Pursuit

DFT Discrete Fourier Transform

GLR Generalized Likelihood Ratio

i.i.d. independent and identically distributed

K-L Divergence Kullback-Leibler Divergence

TV Tartakovsky-Veravalli Procedure

MF Matched Filter

MLSE Maximum Likelihood Sequence Detector

MMSE Minimum Mean Square Error

MSE Mean Square Error

MUD Multiuser Detection

OMP Orthogonal Matching Pursuit

pdf Probability Distribution Function

RD-MUD Reduced-Dimension Multiuser Detection

RDD Reduced-Dimension Decorrelating

RDDF Reduced-Dimension Decision-Feedback

RD-MMSE Reduced-Dimension Minimum Mean Square Error

RD-LS Reduced-Dimension Least Squares

SIC Successive Interference Cancellation

SNR Signal-to-Noise Ratio

Standard Notations

x^*	Conjugate of a complex number
x^+	Positive part of x : $x^+ = x$ if $x \geq 0$, otherwise $x^+ = 0$.
x^-	Negative part of x : $x^- = x$ if $x \leq 0$, otherwise $x^- = 0$.
\mathbb{R}	Real numbers
\mathbb{R}^+	Positive real numbers
\mathbb{C}	Complex numbers
\mathbb{R}^N	N -dimensional vector space
$\Re[x]$	Real part of a number
$[\mathbf{x}]_n$	The n th element of a vector \mathbf{x}
$[\mathbf{X}]_{nm}$	The entry of \mathbf{X} at the n th row and the m th column
$\mathbf{X}_{\mathcal{I}}$	Submatrix formed by columns of \mathbf{X} indexed by \mathcal{I}
$\mathbf{x}_{\mathcal{I}}$	Subvector formed by entries of \mathbf{x} indexed by \mathcal{I}
\mathbf{X}^\top	Transpose of a vector or matrix
\mathbf{X}^H	Conjugate transpose of a vector or matrix
\mathbf{X}^{-1}	Inverse of a matrix
$\text{tr}(\mathbf{X})$	Trace of a matrix. Sum of the diagonal entries
\mathbf{I}	Identity matrix
$\lambda_{\max}(\mathbf{X}), \lambda_{\min}(\mathbf{X})$	The maximum and minimum eigenvalues of a matrix
$ \mathbf{X} $	Determinant of a matrix
$\rho(\mathbf{X})$	Spectral norm of a matrix: $\rho(\mathbf{X}) = [\lambda_{\max}(\mathbf{X}\mathbf{X}^H)]^{1/2}$
$(x)^+$	Positive part of x : $(x)^+ = x$ if $x \geq 0$; otherwise $(x)^+ = 0$
$\text{sgn}(x)$	Sign of a number
δ_{nm}	$\delta_{nn} = 1$, and $\delta_{nm} = 0$ for $n \neq m$
$\dot{f}(x), \ddot{f}(x)$	The first-order and second-order differential of a function
$\mathbb{E}\{x\}$	Expectation of a random variable
$\mathbb{E}\{x \mathcal{A}\}$	Conditional expectation of x given event \mathcal{A}
$\mathbb{P}(\mathcal{A})$	Probability of an event \mathcal{A}
$A \subset B$	Set A is a subset of set B
$\{A\} \cap \{B\}$	Intersection of Set $\{A\}$ with set $\{B\}$
$\{A\} \cup \{B\}$	Union of set $\{A\}$ and set $\{B\}$
A^c	Complement of set A

Chapter 1

Introduction

Statistical signal processing is an important topic in modern science and engineering. While many good methods have been established for classic problems, developing efficient signal processing tools for problems in large scale systems that involve large amount of data presents new challenges. For these problems, computational complexity becomes a major issue. For example, in a typical wireless network, there may be thousands of users whose data need to be processed simultaneously. In a sensor network, there may be thousands of sensors deployed to detect an event of interest. In developing signal processing methods for these problems, a major challenge is to process data from all users or sensors without too much complexity either in signal processing algorithms or in digital circuitry. Another issue is non-homogeneity of the data. In a sensor network, the signal may be a local disturbance that only affects a small subset of sensors. All the unaffected sensors observe only noise. If we want to detect efficiently the emergence of such a signal, it is important to use only the observations that contain useful information and to suppress noise from the non-affected sensors. Another requirement for some applications is real-time signal

processing. With this requirement, we cannot afford algorithms with high complexity due to their large processing times.

We can go a long way towards approaching complexity-reduction and real-time processing by exploiting sparsity, a property inherent in many statistical signal detection problems that may take different forms depending on specific settings. This dissertation will focus on two such problems: sequential multi-sensor change-point detection and multi-user detection. The major contributions of the dissertation consist of the following:

- We present a new method for the sequential multi-sensor change-point detection. A mixture procedure, which will be discussed in more detail in Chapter 3, exploits sensor sparsity - the fact that the fraction of sensors affected by the change-point, p , is relatively small, to achieve a quicker detection. We model the sensor sparsity by assuming that the probability for each sensor to be affected by the change-point is p_0 . Based on this model, we form a mixture log generalized likelihood ratio (GLR) statistic, and derive the mixture procedure. The mixture statistic essentially applies a non-linear weighting function, which is parameterized by p_0 , on the log GLR statistic of each sensor before combine them. We develop theoretical approximations to two performance metrics of the mixture procedure and also demonstrate its performance via numerical simulations.
- We present a new reduced-dimension multiuser detection (RD-MUD) method for multiuser detection in communication systems. RD-MUD exploits *user sparsity* to reduce the number of correlators used at the MUD receiver front-end and still achieve a performance comparable to that of the conventional MUD. Here

user sparsity refers to the empirical observation that the number of active users at any particular time is usually small compared to the total number of users. We develop theoretical and numerical results to characterize the performance of the RD-MUD detectors in terms of their probability-of-error. The error is due to an error in detecting the active users or detecting the wrong symbol of an active user.

1.1 Multi-Sensor Change-Point Detection

In multi-sensor change-point detection, sensors are deployed to monitor the abrupt emergence of a signal. The emergence of such a signal will simultaneously change the distribution of the observations from a subset of sensors. Such a signal is called a *change-point*. Mathematically, the multi-sensor change-point problem can be described as follows. Suppose N sensors make sequences of observations: $y_{n,t}$, $n = 1, \dots, N$, and $t = 0, 1, 2, \dots$. At an unknown time κ , a change-point occurs and simultaneously changes the distribution of observations for a subset of sensors, $n \in \mathcal{N}_a$, $t \geq \kappa + 1$. The cardinality of the subset \mathcal{N}_a is M , and $1 \leq M \leq N$. We would like to detect the occurrence of the change-point using observations $y_{n,t}$. The goal is to detect the change-point as soon as possible after it occurs, while keeping the frequency of false alarms as low as possible. In the change-point detection literature, the frequency of false alarms is usually replaced by a related performance metric, the *average-run-length* (ARL), which by definition is the expected time before (incorrectly) announcing a change of distribution when none has occurred.

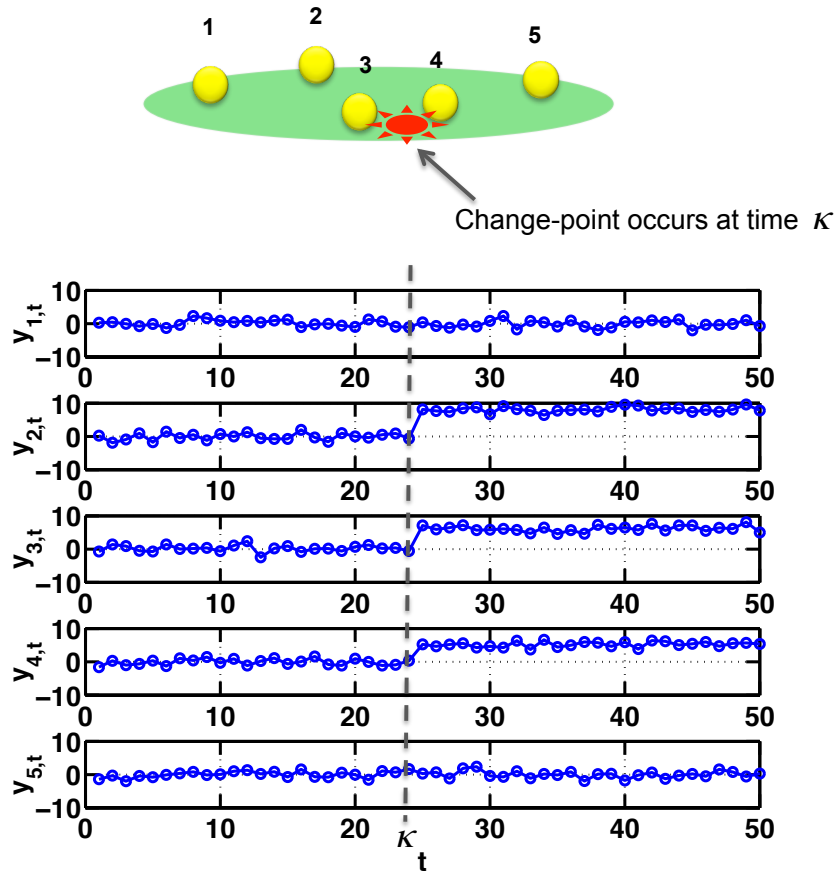


Figure 1.1: The five sequences in the picture demonstrate observations from five sensors. The sensor index is n and the time index is t . A change-point happens at time $\kappa = 25$, with location marked by a red star, illustrates an instance where the change-point affects the observations of $M = 3$ out of $N = 5$ sensors. The post-change means of the observed changes are different for these affected sensors.

Fig. 1.1 illustrates an instance where the change-point affects the observations of $M = 3$ out of $N = 5$ sensors. The observations before the change-point are independent and identically distributed (i.i.d.) with normal distribution that has zero mean and unit variance. A change-point that occurs at $\kappa = 25$ affects sensors with index 2, 3 and 4 and raises expected values of observations from these sensors. In particular, after the change-point, the affected sensors have observations that are i.i.d. with normal distribution that has mean $\mu_n > 0$, $n = 2, 3, 4$; the unaffected sensors have observations that are i.i.d. with the same distribution as before the change-point. The change-point time κ , the subset \mathcal{N}_a , the size M of the subset, and the post-change means μ_n , $n \in \mathcal{N}_a$, of the affected sensors are all unknown.




Figure 1.2: One example of change-point problem with spatial structure. A change-point that occurs abruptly changes the distributions of a subset of sensors simultaneously. The observations of the affected sensors have larger means. In this example, the different post-change means, or the level-of-affectedness, of the affected sensors are spatially correlated. The sensors that are spatially close tend to have similar post-means.

Multi-sensor change-point problems arise in many real world applications. For

example, in network intrusion detection, network anomalies such as TCP/SYN flooding may lead to changes in Internet traffic, and these changes can be modeled as change-points in the network traffic data, as studied in [50]. For TCP/SYN flooding detection, observations consist of data collected at several points in the Internet. The data are the number of TCP/SYN packets received by each destination IP address per unit of time. An attack may dramatically increase the number of TCP/SYN packets for multiple destination IP addresses at the same time, which can be modeled as a change-point. The goal is to provide algorithms to detect a change-point, under computational constraints to process the data on-line. Typically, the number of destination IP addresses we record data simultaneous is very large, which can be up to several thousands and even millions.

Another envisioned application of change-point detection is sensor networks, where multiple sensors are used to make measurements and a change-point may alter the distribution of the measurements from a subset of sensors, as illustrated in Fig. 1.2. An example is the California freeway performance measurement system which comprises of a collection of 25,000 sensors. Each sensor reports the number of vehicles that crossed the sensors within every past five minutes[69]. These measurements are used to generate a real-time traffic map. On average, however, up to 40% of the sensors may fail per day. The goal is to detect the failed sensor quickly in order to generate a map without their measurements.

In multi-sensor change-point detection problems, the level-of-affectedness of each sensor affected by the change-point may be related or not. In some problems, spatially closed sensors tend to have high level-of-affectedness simultaneously. These problems have a spatial structure, which relates the level-of-affectedness of each sensor affected

by the change-point. We call such problems *structured problems*. The spatial structure is determined by known sensor locations and the locations of the sources of the change-point, and decays according to some function. Examples of structured problems include epidemiological monitoring for disease outbreaks [68], fMRI scanning [75], and on-line monitoring of the emergence of a signal with known profile [81]. On the other hand, we call the other problems that do not have a spatial structure *unstructured problems*. An example of the unstructured problem is the network intrusion detection [50] that we discuss earlier. There level-of-affectedness in that example is the mean of the number of TCP/SYN packets for each IP addresses. In that problem, the means tend to be uncorrelated. Non-structured problems have been discussed in [84], [56], [19], [64], and [50] with variations depending on the envisioned applications.

In Chapter 3, we are primarily interested in an unstructured multi-sensor detection problem, where the number of sensors N is large, and the number of sensors affected by the change-point M is relatively small. Moreover, the number and the subset of affected sensors are unknown. The affected sensors and the level-of-affectedness of the affected sensors are unknown. To achieve efficient detection, the detection procedure should use only the observations of the affected sensors, ignore noise from the unaffected sensors, and take into consideration the unknown and non-homogeneous level-of-affectedness of each affected sensor.

There have been a number of suggested methods to find an efficient multi-sensor change-point detection procedure. These methods compute a detection statistic for each sensor, and combine these statistics in various ways, such as Mei's procedure [56] and a procedure suggested by Tartakovsky and Veeravalli [84]. Their works have not exploited the empirical observation that the fraction of sensors affected

by the change-point is typically small. Another drawback of these two procedures is that they assume the distributions of the observations after the change-point is completely prescribed. However, in practice, complete knowledge about the post-change distributions is usually not available. When the true distributions deviate from the assumed ones, the proposed methods both suffer from performance degradation, in particular long detection delays.

We develop a mixture procedure, which models the fact that the fraction of affected sensor by the change-point is p , by assuming that each sensor has a probability p_0 to be affected. The value of p_0 is a guess for p . Based on this model, we form a mixture of the log generalized likelihood ratio (GLR) statistic using observations from each sensor, and derive a mixture detection procedure. The mixture statistic essentially applies a nonlinear weighting function on the log GLR statistic of each sensor before combining them, and automatically emphasizes statistics of the affected sensors and suppresses those of the non-affected sensors. The nonlinear weighting function is parameterized by p_0 . To characterize the performance of the mixture procedure, we derive theoretical approximation for its ARL and expected detection delay. We validate that the approximations have good numerical accuracy. We demonstrate numerically that the mixture procedure is robust against the misspecification of p_0 , and has good performance compared with other existing procedures.

1.2 Reduced-Dimension Multiuser Detection

Multiaccess communication, in which several transmitters share a common channel (Fig. 1.3), is a common protocol in wireless systems today. Examples include cellular networks where mobile telephones transmit to a base station, satellite systems where

ground stations communicate with a satellite, wireless local-area networks (LANs) where many laptops share a common access point, or sensor networks where different sensors communicate with a data collecting central node. A common feature of these communication systems is that the receiver obtains a noisy version of the superposition of the signals sent by the active transmitter, and the signals carry information messages. We call the message sources in the multiaccess channel *users* [96].




Figure 1.3: Multiuser communication network.

In this work, we are interested in detecting the information messages sent by different users, which is called multiuser detection (MUD) (there are other important aspects of the multiaccess communication network, for example, capacity [21][35] and throughput [40]). Multiuser detection (MUD) is a well-studied problem in multiuser communications and signal processing (see, e.g., [26][60][96] for classic reviews with extensive reference lists and [2][74][46] for some recent development in the field). In a typical multiuser system, a number of users communicate simultaneously with a given receiver by modulating information symbols onto their unique signature waveforms, which can be characterized by the chip waveform and the signature sequences (also

called the spreading code) [101]. The received signal consists of a noisy version of the superposition of the transmitted waveforms, and this signal is used to detect the symbols of all users simultaneously. While there has been a large body of work developed for the multiuser detection problem over the last several decades, it is not yet widely implemented in practice, largely due to its complexity. That is the fundamental issue we address in the second part of the dissertation.

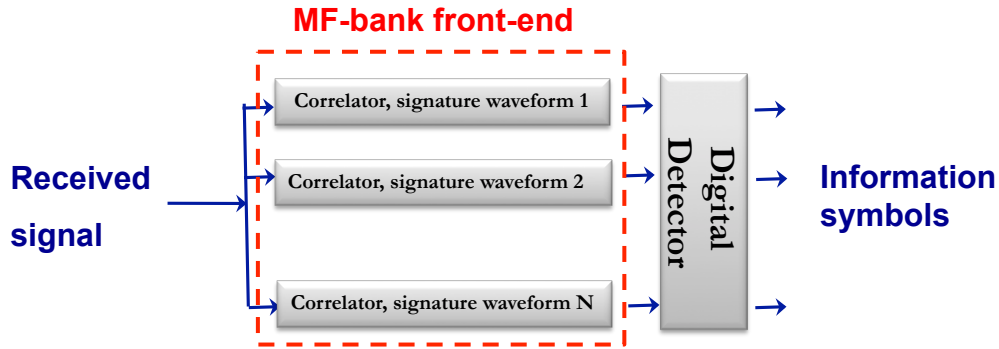


Figure 1.4: Diagram of conventional MUD using MF-bank.

The structure of a MUD detector typically consists of a front-end that processes the analog received signal followed by a digital detector. The front-end is a bank of matched-filters (MF-bank), with each matched filter correlating the received signal with the signature waveform of a different user, as illustrated in Fig. 1.4 (which assumes there are N users). The digital detector can be linear, which linearly transforms the front-end output and detects user data separately in one pass, or it can be nonlinear, which detects user data jointly or iteratively. Both linear and nonlinear MUD have sufficiently high complexity to preclude their wide adoption in real systems. The complexity of MUD arises both in the analog circuitry for decorrelation as well the digital signal processing for data detection of each user. We measure the

first part of the complexity by the number of correlators used at the front-end, and measure the second part of digital signal processing complexity by complexity-per-bit [54], which is the number of real floating point operations required per bit decision. Both the linear and nonlinear detectors require the number of correlators equal to the number of users. The complexity-per-bit of the nonlinear detectors are exponential in the number of users in the system for the maximum likelihood sequence estimator (MLSE), and less for the DF (decision feedback) detector. The complexity-per-bit of the linear detectors are linear in the number of users, which is much lower complexity than that of the nonlinear methods. In a typical communication system, there may be thousands of users. Hence, the complexity of the conventional methods has been a major obstacle for implementing the MF-bank based conventional MUD detectors.

In the second part of the dissertation, Chapter 4, we develop a low complexity front-end for MUD along with corresponding digital detectors. We call this structure a reduced-dimension multiuser detector (RD-MUD). The RD-MUD reduces the number of correlators while still achieving performance similar to that of conventional MUDs that are based on the MF-bank front-end. We reduce complexity by exploiting the fact that at any given time the number of active users, K , is typically much smaller than the total number of users, N . This analog signal sparsity allows us to use techniques from analog compressed sensing, which exploits sparsity in analog signals (see, e.g., [57][30][27][58][28][59][36] and [24] for more details on these techniques). Our RD-MUD has a front-end that correlates the received signal with M correlating signals, where M is much smaller than N . The correlating signals are formed as linear combinations of the signature waveforms via a (possibly complex) coefficient matrix \mathbf{A} , as is done in the analog compressed sensing literature for sparse signal recovery.

Our choice of \mathbf{A} will be shown to be crucial for performance. The output of the RD-MUD front-end can thus be viewed as a projection of the MF-bank output onto a lower dimensional detection subspace. To recover information from this detection subspace, we process the front-end output using algorithms that combine ideas from compressed sensing and MF-bank based conventional MUD. We study two such detectors in detail: the reduced-dimension decorrelating (RDD) detector, a linear detector that combines subspace projection and thresholding to determine active users with a sign detector for data recovery [37][11], and the reduced-dimension decision-feedback (RDDEF) detector, a nonlinear detector that combines decision-feedback orthogonal matching pursuit (DF-OMP) [63][85] for active user detection with sign detection for data recovery in an iterative manner. We present theoretical probability-of-error performance guarantees for these two specific detectors in terms of the coherence of the matrix \mathbf{A} in a non-asymptotic regime with a fixed number of users and active users. Based on these results, we develop a lower bound on the number of correlators M needed to attain a certain probability-of-error performance. For example, if \mathbf{A} is a random partial discrete Fourier transform matrix, the M required by these two specific detectors is on the order of $\log N$. We validate these theoretical performance results via numerical examples.

Previous work on MUD based on active user detection falls into two categories: conventional methods and compressed sensing methods. Specifically, to detect active users in the system, conventional MUD approaches use techniques such as the multiple signal classification algorithm (MUSIC) [102], quickest change detection [61], random set theory, sphere detection and Bayesian filter [10][3][4][14]. In particular, [61] focuses on the on-online detection of the entrance of a new user into the system,

and [102][10][14] consider the case when the number of active users is unknown. There has also been work about detecting a subset of active users of interest [71], which belongs to a more general technique called group detection [89]. The group detection technique partitions active users into groups and jointly detects users within a group using the generalized likelihood ratio test. A compressed sensing approach has also been applied for active user detection. However, in contrast to our approach, which processes analog signals, most existing work on exploiting compressed sensing ideas for signal detection based on the original compressed sensing results [18][22] assume discrete signals. In particular, most prior work on MUD exploiting user sparsity applies compressed sensing techniques on discrete signals via matrix multiplication [32][31][47][41][42][106], whereas in our work the compressed sensing techniques are incorporated into the RD-MUD analog front-end. Furthermore, RD-MUD aims at detecting active users as well as their transmitted symbols, whereas prior work [5][6][32][31][47] aims at detecting only the active users and hence the problem is equivalent to support recovery. These prior works establish conditions on the number of correlators M required to achieve a zero probability-of-error of active user detection when the number of users N tends to infinity. While providing important insights into complexity reduction in large systems, they do not answer questions for practical system design with a finite number of users, such as how many correlators should be used to achieve a target probability-of-error. There is another branch of compressed detection work that focuses on detecting the presence of a discrete signal that is sparse in time [43][42][39]. This work is not relevant to our problem since the multiuser signal we consider is sparse in the number of users.

Our RD-MUD consists of two stages: active user detection and data detection of

active users. The first stage is closely related to [9]. However, our problem differs in that the probability-of-error must consider errors in both stages. We derive conditions under which the probability-of-error is dominated by errors in the first stage. Also, the decision-feedback (DF) detector we consider is different in its active user detection because it subtracts out detected data symbols whose values are from a finite alphabet (rather than subtracting out the estimated data symbols whose values are real) from the second stage, which makes detection of the remaining active users easier.

The rest of the dissertation is organized as follows. Chapter 2 reviews some background knowledge about change-point detection and multiuser detection. Chapter 3 presents the mixture procedure for multi-sensor change-point detection. Chapter 4 introduces the reduced-dimension multi-user detection. Finally Chapter 5 concludes the dissertation with discussions about future research.

Chapter 2

Background

In this chapter, we review some fundamentals that will be used for later development. In particular, we will review the classic one-sensor change-point detection problem. Then we will review two topics related to the reduced-dimension multiuser detection (RD-MUD): conventional multi-user detection and compressed sensing.

2.1 Classic One-Sensor Change-Point Detection

In the classical single-sensor change-point detection problem, the sensor observes a sequence of independent random variables y_t , $t = 1, 2, \dots$. If there is no change-point, all the observations, y_t , $t = 1, 2, \dots$ are independent and identically distributed (i.i.d.) with a known probability density function (pdf) f_0 . However, if there is a change-point occurs at κ , initially and up to time κ , the observations $y_1, y_2, \dots, y_\kappa$ are i.i.d. with a known pdf f_0 , and after the change-point, the observations $y_{\kappa+1}, y_{\kappa+2}, \dots$ are i.i.d. with a different pdf f_1 . The parameter κ , the time at which the change-point happens, is unknown. An important question to ask after observing t observations

is whether the change-point occurs, i.e., $\kappa < t$, or whether the change-point never occurs.

There are two different versions of the change-point detection problems, one where the total number of observations T is fixed, and we can detect offline after acquiring all T samples. A second version of the change-point detection has the total number of observations not fixed. The samples y_t , $t = 1, 2, \dots$ are observed sequentially and we have to detect online as we acquire new samples. In this chapter we will focus on the second case, the sequential change-point detection. For sequential change-point detection, the goal is to detect the occurrence of the change-point with a minimal detection delay after κ , subject to having a large number of observations until a change-point is incorrectly detected when there is no change (which is related to the false alarm rate). The two performance metrics: the expected detection delay and the false alarm rate will be defined more precisely below.

Historically the applications of change-point detection have come from quality control. Later the change-point detection problem has been applied to a more general setting. The literature of single-sensor change-point detection is immense. See, for example, [49][48][8][66] for surveys. Solutions proposed for the single-sensor problem are usually either variations of the Page-Lorden CUSUM procedure [62][51] or of the Shiryaev-Roberts procedure [76][72].

Let \mathbb{P}^κ and \mathbb{E}^κ denote the probability and expectation when the change-point occurs at the κ th observation, $\kappa = 0, 1, \dots$. The probability \mathbb{P}^κ is defined on the infinite dimensional product space. For any finite dimensional vector (y_1, \dots, y_t) and

Borel set $\mathcal{B} \subset \mathbb{R}^t$,

$$\mathbb{P}^\kappa\{(y_1, \dots, y_t) \in \mathcal{B}\} = \begin{cases} \int_{\mathcal{B}} f_0(y_1) \cdots f_0(y_\kappa) f_1(y_{\kappa+1}) \cdots f_1(y_t) dy_1 \cdots dy_t, & t > \kappa, \\ \int_{\mathcal{B}} f_0(y_1) \cdots f_0(y_t) dy_1 \cdots dy_t, & 1 \leq t \leq \kappa. \end{cases} \quad (2.1)$$

For a measurable real-valued function defined on this infinite dimensional product space, its expectation is defined as:

$$\mathbb{E}^\kappa[g] = \int g d\mathbb{P}^\kappa, \quad (2.2)$$

provided the integral exists. Let \mathbb{P}^∞ and \mathbb{E}^∞ denote the probability and the expectation when there is no change, i.e., $\kappa = \infty$. The probability \mathbb{P}^∞ is defined on the infinite dimensional product space. For any finite dimensional vector (y_1, \dots, y_t) and Borel set $\mathcal{B} \subset \mathbb{R}^t$,

$$\mathbb{P}^\infty\{(y_1, \dots, y_t) \in \mathcal{B}\} = \int_{\mathcal{B}} f_0(y_1) \cdots f_0(y_t) dy_1 \cdots dy_t, \quad t = 1, 2, \dots. \quad (2.3)$$

The expectation of a measurable real-valued function defined in the infinite dimensional product space is given by

$$\mathbb{E}^\infty[g] = \int g d\mathbb{P}^\infty. \quad (2.4)$$

The detection procedure we are seeking is a stopping rule of the form

$$T = \inf\{t : F_t(y_1, \dots, y_t) \geq b, t = 1, 2, \dots\}, \quad (2.5)$$

for some time-varying function F_t of observations, and a threshold $b > 0$. A desired stopping rule should minimize the expected detection delay and rarely stop when the change-point has not occurred. We usually consider the worst-case expected detection delay after a change-point occurs, which is the largest expected detection delay maximized over all possible change-point times $\kappa = 0, 1, \dots$, denoted as

$$\sup_{k \geq 0} \mathbb{E}^\kappa[T - k | T > k], \quad (2.6)$$

also referred to as the *expected detection delay* in the change-point detection literature. The expected period of false alarm, i.e. detecting a change-point, when there is no change-point, denoted as $\mathbb{E}^\infty[T]$, is referred to as the *average run length* (ARL).

The above change-point detection problem can be formulated as a sequential hypothesis testing. The null hypothesis is that there is no change-point, and hence y_1, y_2, \dots are independent and identically distributed (i.i.d.) with probability density function (pdf) f_0 . An alternative hypothesis is that there is a change-point that occurs at $0 \leq \kappa \leq t$, such that y_1, \dots, y_κ are i.i.d. with pdf f_0 , while $y_{\kappa+1}, \dots, y_t$ are i.i.d. with pdf f_1 . Here $\kappa = 0$ means that the change-point happens from the first observation, and $\kappa = t$ means that the change-point has not occurred at time t . Since the change-point occurs at an unknown time κ , we have to test a composite hypothesis that κ happens somewhere between 0 and t . For such a composite hypothesis test, the log likelihood ratio statistic is given by

$$\max_{0 \leq k \leq t} \sum_{l=k+1}^t \log(f_1(y_l)/f_0(y_l)). \quad (2.7)$$

The log likelihood ratio is a random process in t , with a drift-rate given by $\mathbb{E}[\log(f_1(y_l)/f_0(y_l))]$.

This drift-rate is negative before the change-point occurs $l \leq \kappa$, and is positive after the change-point occurs $l \geq \kappa + 1$. These facts can be shown using Jensen's inequality. Before the change-point occurs $l \leq \kappa$, the drift-rate is given by

$$\begin{aligned} \mathbb{E}[\log(f_1(y_l)/f_0(y_l))] &\leq \log \mathbb{E}[f_1(y_l)/f_0(y_l)] \\ &= \log \int (f_1(u)/f_0(u)) \cdot f_0(u) du = \log 1 = 0, \end{aligned} \tag{2.8}$$

with equality if and only if $f_1 = f_0$. However, we are not interested in the trivial case that $f_1 = f_0$, and hence $\mathbb{E}[\log(f_1(y_l)/f_0(y_l))] < 0$ for $l \leq \kappa$. After the change-point occurs $l \geq \kappa + 1$, the drift rate is given by

$$\begin{aligned} \mathbb{E}[\log(f_1(y_l)/f_0(y_l))] &\geq -\log \mathbb{E}[f_0(y_l)/f_1(y_l)] \\ &= -\log \int (f_0(u)/f_1(u)) \cdot f_1(u) du = -\log 1 = 0, \end{aligned} \tag{2.9}$$

with equality if and only if $f_1 = f_0$. Again since we are not interested in the trivial case that $f_1 = f_0$, hence $\mathbb{E}[\log(f_1(y_l)/f_0(y_l))] > 0$ for $l \geq \kappa + 1$. Hence, intuitively, when there is a change-point, we can detect its occurrence by comparing the log likelihood ratio statistic with a positive threshold b and claiming a detection when it exceeds the threshold. In other words, we can choose the function F_t in (2.5) to be the log likelihood ratio (2.7).

Define $S_t = \sum_{l=1}^t \log(f_1(y_l)/f_0(y_l))$ and $S_0 = 0$. Then the statistic (2.7) can be written as

$$S_t - \min_{0 \leq k \leq t} S_k, \tag{2.10}$$

which is referred to as the CUSUM statistic [8]. The related CUSUM procedure is

defined by a stopping time

$$T = \inf \left\{ t : S_t - \min_{0 \leq k \leq t} S_k \geq b \right\}, \quad (2.11)$$

for a threshold $b > 0$. The CUSUM procedure was first considered by Page in [62].

An important instance of the CUSUM procedure is when the pre-change and post-change distributions are normal. Assume f_0 is normal pdf with zero mean and unit variance, and f_1 is normal pdf with mean μ and unit variance. Then $\log(f_1(y_l)/f_0(y_l)) = y_l\mu - \mu^2/2$, and the CUSUM procedure becomes

$$T = \inf \left\{ t : \max_{0 \leq k \leq t} \sum_{l=k+1}^t (y_l\mu - \mu^2/2) \geq b \right\}. \quad (2.12)$$

When the post-change mean of the normal distribution is unknown, we can use a hypothetical value for μ or use the generalized likelihood ratio (GLR) statistic, where the unknown change-point mean is replaced by its maximum likelihood estimate at time t . If we assume the change-point happens at $\kappa = k$, the maximum likelihood estimate for the post-change mean is given by

$$\hat{\mu}_k = \sum_{l=k+1}^t y_l / (t - k). \quad (2.13)$$

In this case, the detection procedure based on the GLR is given by

$$T = \inf \left\{ t : \max_{0 \leq k \leq t} \left[\sum_{l=k+1}^t y_l \right]^2 / [2(t - k)] \geq b \right\}. \quad (2.14)$$

The properties of this GLR procedure have been analyzed in [78].

The two performance metrics of a change-point detection procedure, i.e., the ARL and the expected detection delay, can be evaluated by several frequently used identities and tools, which are given here for future reference. Wald's identity is particular useful for evaluating detection delays (see, e.g., [80]):

Lemma 1 (Wald's Identity). *Let y_1, y_2, \dots be a sequence of i.i.d. random variables. Define $S_t = \sum_{l=1}^t y_l$. Let T be any integer-valued random variable such that $\{T = t\}$ is an event determined by y_1, \dots, y_t (and is independent of y_{t+1}, \dots) for all $t = 1, 2, \dots$, and assume that $\mathbb{E}[T] < \infty$. Then $\mathbb{E}[S_T] = \mathbb{E}[T]\mathbb{E}[y_1]$.*

Consider the stopping time $T = \inf\{t : S_t \geq b\}$ for an integer t . Then if the process achieves the threshold b at the stopping time T , we must have $S_T = b$, from which we can easily evaluate the expected value of T using Wald's identity in Lemma 1 to obtain $\mathbb{E}[T] = b/\mathbb{E}[y_1]$. However, the random process generally will have an overshoot over the threshold b at the stopping time and hence $S_T \neq b$. To account for this overshoot, we can write Wald's identity as

$$b + \mathbb{E}[S_T - b] = \mathbb{E}[T]\mathbb{E}[y_1], \quad (2.15)$$

and evaluate the expected overshoot $\mathbb{E}[S_T - b]$ via renewal theory (see, e.g., [80]). In particular, the following results from [80] can be used.

Lemma 2 (Siegmund [80]). *With the definition in Lemma 1, if $\mathbb{E}[y_1] > 0$ and $\mathbb{E}[y_1] < \infty$, then as $b \rightarrow \infty$*

$$\mathbb{E}[S_T - b] \rightarrow \frac{\mathbb{E}[S_\tau^2]}{2\mathbb{E}[S_\tau]} = \frac{\mathbb{E}[y_1^2]}{2\mathbb{E}[y_1]} - \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E}[S_n^-], \quad (2.16)$$

where $\tau = \inf\{t : S_t > 0\}$, and $x^- \triangleq |\min\{0, x\}|$.

Lemma [80] can be proven using renewal theory and Theorem 8.43 in [80].

Lemma 3 (Kac). *With the definition in Lemma 1, if $\mathbb{E}[y_1] > 0$ and $\mathbb{E}[y_1] < \infty$,*

$$\mathbb{E} \left[\min_{1 \leq k \leq t} S_k \right] = - \sum_{k=1}^t \frac{1}{k} \mathbb{E}[S_k^-]. \quad (2.17)$$

A proof of Lemma 3 can be found in Section 8.4 and problem 8.14 in [80].

2.2 Multi-User Detection

As we have briefly discussed in Chapter 1, multiuser detection (MUD) has many examples in communications and signal processing. In multiuser systems, the users communicate simultaneously with a given receiver by modulating information symbols onto their unique signature waveforms, which can be characterized by the chip waveform and the signature sequences (also called the spreading code) [101]. The received signal consists of a noisy version of the superposition of the transmitted waveforms. The MUD has to detect the symbols of all users simultaneously.

The channel model associated with the MUD may be synchronous or asynchronous. In the synchronous channel model [54], the transmission rate of all users is the same and their symbol epochs are perfectly aligned. This requires closed-loop timing control or synchronization among all transmitters. In the asynchronous channel model [93][55], user time epochs need not be aligned and the transmitted waveforms arrive at the receiver with different time delays. Allowing users to be asynchronous simplifies system design but complicates the system model. The synchronous channel model

can be viewed as a special case of the asynchronous channel with delays of all users to be the same. In this paper we will focus on the synchronous channel model.

Part of the MUD problem is signature sequence selection, for which there has been a large body of work, both theoretical [97][73][98][99][88][86][44][52][38][87][12] and practical [101][82]. If we require the signature waveforms to be orthogonal, for a system with bandwidth B , the number of orthogonal signature waveforms available (and hence the number of users that the system can support) is approximately $2TB$ [96]. This hard limit on system capacity can be relaxed if we allow nonorthogonal signature waveforms, and instead require the crosscorrelation of the selected signature waveforms to be sufficiently low. In this work, we do not consider optimizing signature waveforms and hence our results will be parameterized by the crosscorrelation properties of the signature waveforms used in our design. When the signature waveforms are nonorthogonal, there is mutual interference among users, which degrades system performance for all users.

An important issue in multiuser systems with nonorthogonal signature waveforms is the *near-far problem* [96]: a strong user with high enough power at the receiver may cause severe performance degradation of a weak user. One of the key challenges in MUD is to design a detector that works well when user signals are received at different power levels. Such detectors are discussed in more detail below.

While there has been a large body of work developed for the multiuser detection problem over the last several decades, it is not yet widely implemented in practice, largely due to its complexity and high-precision A/D requirements. The complexity of MUD arises both in the analog circuitry for decorrelation as well as in the digital signal processing for data detection of each user. We characterize the decorrelation

complexity by the number of correlators used at the receiver front-end, and measure the data detection complexity by the complexity-per-bit [54], which is the number of real floating point operations required per bit decision.

The conventional MUD detection structure consists of a matched-filter (MF) bank front-end followed a linear or nonlinear digital multiuser detector. The MF-bank front-end (shown in Fig. 1.4 and later in Fig. 4.2) is a set of correlators, each correlating the received signal with the signature waveform of a different user. Hence the conventional MUD requires the number of correlators to be equal to the number of users. The MF-bank front-end obtains a set of sufficient statistics for MUD when the receiver noise is Gaussian.

To recover user data from the MF-bank output, various digital detectors have been developed. Verdu in a landmark paper [93] establishes the optimal MUD detector as the maximum likelihood sequence estimator (MLSE), which minimizes the probability of error for symbol detection. The upper bound on the probability of error for symbol detection, or the bit-error-rate (BER), was derived in [93] and the analysis is also given in [92]. Although the MLSE detector can nearly eliminate the degradation in performance due to multiuser interference, it has two main limitations: complexity and the requirement for complete channel state information associated with all users [96]. The complexity-per-bit of the MLSE detector is exponential in the number of users when the signature waveforms are nonorthogonal. To address the complexity issue, other MUD suboptimal detectors have been developed. The MLSE detector is an example of a nonlinear detector that detects symbols of all users jointly. Another example of a nonlinear detector is the decision feedback (DF) detector [91][90], which is based on the idea of interference cancellation and takes various forms. One such

form is the successive interference cancellation (SIC) detector [100][25]. The idea of successive interference cancellation dates back to the information theoretic study of the Gaussian multiple-access channel [20]. The DF detector decodes symbols iteratively and subtracts the detected symbols of strong users first to facilitate detection of weak users. The DF detector also requires complete channel state information but it has less complexity-per-bit than the MLSE detector. The number of correlators required by the DF detector is also equal to the number of users. The DF detector is a good compromise between complexity and performance among all nonlinear and linear techniques (see, e.g., [90]). For this reason we will analyze the DF detector below as an example of a nonlinear detector, but in a reduced dimensional setting.

Linear detection methods, which apply a linear transform to the receiver front-end output and then detect each symbol separately, have lower complexity than nonlinear methods but also worse performance. Linear MUD techniques include the single-user detector, the decorrelating detector and the minimum mean-square-error (MMSE) detector. The single-user detector is the simplest linear detector, which follows the MF-bank front-end with conventional single-user detection in each branch. Thus single-user detection demodulates symbols by comparing the front-end branch output with a threshold [54]. When the signature waveforms are orthogonal and synchronous, the single-user detector coincides with the MLSE detector and it minimizes the probability-of-error [96]. In this case, by correlating with each user's signature waveform, the MF-bank maximizes the output signal-to-noise ratio (SNR) for each individual user. When the signature waveforms are nonorthogonal, users interfere with each other, so the probability-of-error of the single-user detector degrades [26]. A linear detector that eliminates user inference is the decorrelating detector [54][55],

which, for each user, projects the received signal onto the subspace associated with the signature waveform of that user. This projection amplifies noise when the signature waveforms are nonorthogonal. The decorrelating detector also gives the best joint estimate of symbols and amplitudes in the absence of knowledge of the complete channel state information [54], and it maximizes the near-far resistance among all linear detectors [94]. The near-far resistance is a performance measure for the degree of robustness against the near-far problem achieved by a multiuser detector [96]. The MMSE detector [103][104][67] is designed to minimize the mean-square-error (MSE) between symbols and the linearly transformed MF-bank front-end output. The MMSE detector takes into account the background noise and interference, and hence to some extent it mitigates the noise amplification effect of the decorrelating detector in the low and medium SNR regimes [90]. A drawback of the MMSE detector is that it requires complete channel state information and it does not achieve the optimal near-far resistance. When the signal-to-noise power ratio (SNR) goes to infinity, the MMSE detector converges to the decorrelating detector [96]. Because of the many advantages of the decorrelating detector, it has received much attention in the MUD literature and is one of the most common linear detectors in MUD [60]. Hence, in this paper, we will focus on the decorrelating detector as an example of a linear detector in the reduced-dimension setting. The decorrelating detector requires inverting the correlation matrix of the signature waveforms, but this computation can be done offline once the signature waveforms are selected. Its complexity-per-bit (other than inverting the correlation matrix) can be shown to be linear in the number of users [54].

The complexity-per-bit and required channel-state-information (CSI) requirements

for different MUD detectors are summarized in Table 2.1 and Table 2.2, respectively.

Table 2.1: Complexity-per-bit of MUD detectors (N users, K active users).

MUD detectors	Complexity-per-bit
MLSE	Exponential in N when signature waveforms nonorthogonal
DF detector	Proportional to NK
Decorrelating detector	Linear in N
MMSE	Linear in N

2.3 Compressed Sensing

Compressive sensing is a new field that has attracted much interest from theorists and practitioners. Compressive sensing recovers a sparse signal using few measurements (see, e.g., [16][18][15][23]). In the classic compressive sensing setting, the signal is usually represented by a finite dimensional vector $\mathbf{x} \in \mathbb{R}^N$ that has very few non-zero elements. In particular, the vector \mathbf{x} that has only k non-zero entries is called k -sparse. For this sparse vector, obtaining M linear measurements of \mathbf{y} is done by

Table 2.2: Required channel-state-information (CSI) by MUD detectors.

MUD detectors	Required CSI
MLSE	Complete CSI (amplitude and phase)
DF detector	Complete CSI
Decorrelating detector (fading channel)	Phase
Decorrelating detector (non-fading channel)	No CSI
MMSE	Complete CSI

applying a matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x}. \quad (2.18)$$

The matrix \mathbf{A} is designed to reduce the number of measurements M as much as possible while still being able to recover the sparse signal \mathbf{x} from the measurement \mathbf{y} . Mathematically, (2.18) is an underdetermined system, and if the entries of \mathbf{x} are all non-zero, it cannot be uniquely identified from \mathbf{y} . However, when \mathbf{x} has only few non-zero entries, i.e., when it is sparse, various algorithms in compressive sensing can recover \mathbf{x} from \mathbf{y} .

In compressive sensing, one question that arises is how many measurements are needed to guarantee that \mathbf{x} can be reliably recovered. A fact that can be easily verified is that, in the absence of noise, to recover the k -sparse \mathbf{x} from \mathbf{y} , we need at least $2k$ measurements. If we have two different k -sparse vectors \mathbf{x}_1 and \mathbf{x}_2 such that $\mathbf{A}\mathbf{x}_1 = \mathbf{A}\mathbf{x}_2$, we should have $\mathbf{x}_1 = \mathbf{x}_2$. Equivalently, this requires $\mathbf{A}(\mathbf{x}_1 - \mathbf{x}_2) = \mathbf{0}$. Since $\mathbf{x}_1 - \mathbf{x}_2$ has at most $2k$ non-zero entries, to satisfy this requirement, the column rank of \mathbf{A} must be greater than $2k$, or the number of rows of \mathbf{A} should also be at least $2k$. Hence the number of measurements should be greater than $2k$. When there is noise in the measurements: $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, the performance metric for recovery algorithms considered in compressive sensing literature is usually the MSE, which is defined as $\mathbb{E}\{\hat{\mathbf{x}} - \mathbf{x}\}$.

Several fundamental results in compressive sensing (see, e.g., [18][17]) state that when the number of measurements is on the order of $k \log(N/k)$ then, in the absence of noise, it is possible to recover the sparse vector exactly; in the presence of Gaussian noise with zero mean and variance σ^2 , the MSE of the recovery algorithms can be on the order of $k\sigma^2/M \cdot \log(N)$ if the number of measurement M is on the order of

$k \log(N)$.

Analog compressive sensing extends the notion of compressive sensing so that the signals can be not only finite dimensional sparse vectors but also continuous signals that have certain sparse representation (see, e.g., [57][30][27][58][28][59] [36] and [24] for review).

The recovery algorithm is the key to reducing the number of measurements. These algorithms include l_1 relaxation methods, such as Danzig selector [15] and basis pursuit denoising (also known as LASSO) [23][18], and the greedy algorithms, such as thresholding and orthogonal matching pursuit (OMP) [63][85]. The l_1 relaxation method recovers the sparse vector by solving a quadratic optimization problem,

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \gamma \|\mathbf{x}\|_1, \quad (2.19)$$

for some regularization parameter γ . Here $\|\mathbf{x}\|_1$ and $\|\mathbf{x}\|_2$ are the l_1 and l_2 norms of \mathbf{x} , respectively.

Rather than solving optimization problems, the greedy algorithms recover the sparse signal by first estimating the support set \mathcal{I} of \mathbf{x} , and then estimating the values of the non-zero entries using a least-squares (LS) solution:

$$\hat{\mathbf{x}} = \begin{cases} \mathbf{A}_{\mathcal{I}}^\dagger \mathbf{y}, & \text{on the support set } \mathcal{I}; \\ \mathbf{0}, & \text{elsewhere.} \end{cases} \quad (2.20)$$

Here $\mathbf{A}_{\mathcal{I}}$ denotes a submatrix \mathbf{A} formed by columns of \mathbf{A} on the support set \mathcal{I} , and \mathbf{A}^\dagger denotes the Moore-Penrose pseudoinverse of matrix \mathbf{A} .

Greedy algorithms differ in their ways of estimating \mathcal{I} . The standard thresholding

algorithm estimate \mathcal{I} by computing the correlation of the measured signal \mathbf{y} with each column of \mathbf{A} and defining \mathcal{I} as the set of indices of K columns having the highest correlations.

An iterative algorithm which improves on simple thresholding is known as the orthogonal matching pursuit algorithm (OMP). The OMP initializes the estimate of $\hat{\mathbf{I}}^0$ to be an empty set, and then sets the residual vector \mathbf{r}^0 to be \mathbf{y} . Subsequently, at each iteration $k = 1, \dots, K$, OMP finds the index of a single column of \mathbf{A} that is most highly correlated with \mathbf{r}^{k-1} . This index n_k is then added to the support set, so that $\mathcal{I}^k = \mathcal{I}^{k-1} \cup \{n_k\}$. The estimate $\hat{\mathbf{x}}^k$ at the k th iteration is then obtained via the least-squares solution of (2.20) using the support set \mathcal{I}^k , and the residual vector is updated to $\mathbf{r}^k = \mathbf{y} - \mathbf{A}\hat{\mathbf{x}}^k$. A well-known property of the OMP algorithm is that it never chooses the same column index twice [34]. After K iterations, the sparse vector is defined as $\hat{\mathbf{x}}^K$ with sparsity K .

Chapter 3

Multi-Sensor Change-Point Detection

In this chapter, we develop a mixture procedure for the multi-sensor change-point detection problem with an unknown subset of affected sensors and incompletely specified post-change distributions. The mixture procedure exploits the sparsity that the fraction of sensors affected by the change-point is typically small, and models the sparsity by assuming that each sensor has a chance p_0 to be affected. The value of p_0 is a guess for p . Then we form a mixture log generalized likelihood ratio (GLR) statistic for change-point detection. Equivalently, the mixture statistic uses a nonlinear weighting function, which is parameterized by p_0 , to combine the log GLR statistic of each sensor. The nonlinear weighting function automatically emphasizes the statistics of the affected sensors, and suppresses those of the unaffected sensors. We characterize the performance of the mixture procedure by two metrics: the false alarm rate, which we capture by its commonly used surrogate: the average run length (ARL), and the expected detection delay. We also compare the performance of the

mixture procedure with other existing procedures using numerical simulations.

The reminder of the chapter is organized as follows. In Section 3.1 we establish notation for this chapter and present the system model and problem formulation. In Section 3.3 we review several detection procedures and introduce the new mixture procedure. In Section 3.4 we derive approximations to ARL and expected detection delay of the mixture procedure, and demonstrate with numerical examples that these approximations are reasonably accurate. In Section 3.5, we demonstrate that the mixture procedure has good performance compared with other known procedures. Finally Section 3.8 concludes the chapter with some discussion.

3.1 Model and Formulation

Consider a multi-sensor detection system, where N sensors make sequential observations to monitor the occurrence of a change-point. Define a set

$$\mathcal{N} \triangleq \{1, 2, \dots, N\}, \quad (3.1)$$

which contains indices of all sensors. The sensor observations are given by

$$y_{n,t}, \quad n = 1, \dots, N, \quad t = 1, 2, \dots. \quad (3.2)$$

Assume that observations from different sensors are mutually independent. The null hypothesis is that the change-point never occurs. Denote by \mathbb{P}^∞ the probability that the change-point never occurs. Under the null hypothesis, the sensor observations are independent and identically distributed (i.i.d.) with normal distribution that has zero

mean and unit variance, and the pdf of all observations up to time t , $y_{n,l}$, $l = 1, \dots, t$, $n = 1, \dots, N$, equals

$$\prod_{n=1}^N \prod_{l=1}^t \phi(y_{n,l}), \quad (3.3)$$

where $\phi(x) = (2\pi)^{-1/2} \exp\{-x^2/2\}$ is the pdf of the standard normal distribution. Using (3.3), under the null hypothesis, the expectation \mathbb{E}^∞ of a function of all the observations up to time t is defined as

$$\begin{aligned} & \mathbb{E}^\infty[g(y_{1,1}, \dots, y_{N,t})] \\ &= \int \cdots \int g(u_{1,1}, \dots, u_{N,t}) \prod_{n=1}^N \prod_{l=1}^t \phi(y_{n,l}) du_{1,1} \cdots du_{N,t}. \end{aligned} \quad (3.4)$$

The alternative hypothesis is that there exists a change-point $\kappa \geq 0$ and a subset of sensors $\mathcal{N}_a \subset \mathcal{N}$, the distributions of whose observations are affected by the change-point. Denote by \mathbb{P}^κ the probability that the change-point occurs at time κ . Note that this probability in the alternative case depends on \mathcal{N}_a and the values of $\mu_n > 0$, although this dependence is suppressed in the notation. The cardinality of the subset \mathcal{N}_a is M . So in the alternative, the true fraction of affected sensors is given by $p = M/N$. For $n \in \mathcal{N}_a$, the sensor observations after time κ , $y_{n,t}$, $t > \kappa$, are i.i.d. with normal distribution that has mean $\mu_n > 0$ and unit variance. For $n \in \mathcal{N}_a^c$, the sensor observations are not affected by the change-point and $y_{n,t}$, $n \in \mathcal{N}_a^c$, have the same distribution as before the change, i.e. they are i.i.d. with normal distribution that has zero mean and unit variance. Under the alternative hypothesis, the pdf of

all observations up to time t , $y_{n,l}$, $l = 1, \dots, t$, $n = 1, \dots, N$, equals

$$\left[\prod_{n \in \mathcal{N}_a^c} \prod_{l=1}^t \phi(y_{n,l}) \right] \cdot \left[\prod_{n \in \mathcal{N}_a} \prod_{l=1}^{\kappa} \phi(y_{n,l}) \prod_{l=\kappa+1}^t \phi(y_{n,l} - \mu_n) \right]. \quad (3.5)$$

Using (3.5), under the alternative hypothesis, the expectation \mathbb{E}^κ of a function of all the observations up to time t is defined as

$$\begin{aligned} & \mathbb{E}^\kappa[g(y_{1,1}, \dots, y_{N,t})] \\ &= \int \cdots \int g(u_{1,1}, \dots, u_{N,t}) \left[\prod_{n \in \mathcal{N}_a^c} \prod_{l=1}^t \phi(y_{n,l}) \right] \cdot \left[\prod_{n \in \mathcal{N}_a} \prod_{l=1}^{\kappa} \phi(y_{n,l}) \right. \\ & \quad \left. \prod_{l=\kappa+1}^t \phi(y_{n,l} - \mu_n) \right] du_{1,1} \cdots du_{N,t}. \end{aligned} \quad (3.6)$$

Our goal is to find a sequential change-point detection procedure, or equivalently to define a stopping rule T such that for a prescribed large constant $c > 0$, asymptotically as $c \rightarrow \infty$, it satisfies a certain ARL constraint, $\mathbb{E}^\infty[T] \geq c$, and has small expected detection delay $\mathbb{E}^\kappa[T - \kappa | T > \kappa]$.

3.2 Mixture Procedure

To model the fact that the fraction of sensors affected is an unknown quantity p , we assume that each sensor is affected by the change-point with probability p_0 . Here p_0 is a guess for the true fraction p . Such a mixture model has been used and derived for a *fixed sample* change-point detection problem in prior work [105]. Based on this model, we derive a mixture log generalized likelihood ratio (GLR) statistic, and use the statistic to form a detection procedure.

First we derive the mixture log GLR statistic. For sensors that are affected by the change-point, $n \in \mathcal{N}_a$, $t > \kappa$, using (3.3) and (3.5), the log likelihood for observations of sensor n up to time t is given by

$$\begin{aligned} \ell_n(t, \kappa, \mu_n) &= \log \left\{ \frac{\prod_{l=1}^{\kappa} \exp\{-y_{n,l}^2/2\} \prod_{l=\kappa+1}^t \exp\{-(y_{n,l} - \mu_n)^2/2\}}{\prod_{l=1}^t \exp\{-y_{n,l}^2/2\}} \right\} \\ &= \sum_{l=\kappa+1}^t (\mu_n y_{n,l} - \mu_n^2/2). \end{aligned} \quad (3.7)$$

For sensors that are not affected by the change-point, $n \in \mathcal{N}_a^c$, the log likelihood for the observations of sensor n up to time t is one. Since each sensor has a probability p_0 to be affected by the change-point, the mixture log GLR statistics is given by

$$\sum_{n=1}^N \log(1 - p_0 + p_0 \exp[\ell_n(t, \kappa, \mu_n)]). \quad (3.8)$$

When μ_n and the change-point time κ are unknown, we form the log GLR statistic by maximizing (3.8) with respect to these two parameters:

$$\max_{0 \leq \kappa < t} \max_{\mu_n > 0} \sum_{n=1}^N \log(1 - p_0 + p_0 \exp[\ell_n(t, \kappa, \mu_n)]). \quad (3.9)$$

The μ_n that maximizes (3.9) equivalently maximizes (3.7). Hence we take derivative of ℓ_n in (3.7) with respect to μ_n , set it to zero, and solve for the solution. We take the positive part of the solution, since μ_n are assumed to be positive. This results in the maximum likelihood estimate of the post-change mean μ_n as a function of the

current observations up to time t and an assumed change-point occurs at time $\kappa = k$:

$$\hat{\mu}_{n,k} = \left(\sum_{l=k+1}^t y_{n,l} \right)^+ / (t - k), \quad (3.10)$$

where $(x)^+ = x$ if $x \geq 0$, and is equal to zero when $x < 0$. Substitution of (3.10) in (3.7) gives:

$$\ell_n(t, k, \hat{\mu}_n) = \left[\left(\sum_{l=k+1}^t y_{n,l} \right)^+ \right]^2 / [2(t - k)]. \quad (3.11)$$

It will be convenient to define the following

$$S_{n,t} = \sum_{l=1}^t y_{n,l}, \quad (3.12)$$

$$U_{n,k,t} = (t - k)^{-1/2} (S_{n,t} - S_{n,k}),$$

and write (3.11) as

$$\ell_n(t, k, \hat{\mu}_n) = (U_{n,k,t}^+)^2 / 2. \quad (3.13)$$

Using (3.13), we can write the mixture log GLR statistic (3.9) as

$$\max_{0 \leq k < t} \sum_{n=1}^N \log(1 - p_0 + p_0 \exp[(U_{n,k,t}^+)^2 / 2]). \quad (3.14)$$

Define a function with the argument x and parameter p_0 ,

$$g(x; p_0) = \log[1 - p_0 + p_0 e^{(x^+)^2 / 2}]. \quad (3.15)$$




Figure 3.1: Plot of $g(x; p_0) = \log[1 - p_0 + p_0 e^{(x^+)^2/2}]$ for different p_0 .

We note that the mixture statistic (3.14) is equivalent to applying $g(x; p_0)$ as a non-linear weighting function on the log GLR statistic of each sensor (3.11) before combining them. The function (3.15) plays a role in automatically emphasizing the log GLR statistic of the affected sensors and suppressing noise from those of the unaffected sensors. To see this, when the n th sensor is affected by the change-point at κ , for this sensor, its observations $y_{n,t}$, $n \in \mathcal{N}_a$, $t > \kappa$, have positive means μ_n . By definition (3.12), the process $U_{n,k,t}$ also has positive drift for $t > k \geq \kappa$. Thus as t increases, $\exp[(U_{n,k,t}^+)^2/2]$ grows fast and dominates the other term, i.e., $p_0 \exp[(U_{n,k,t}^+)^2/2] \gg 1 - p_0$. In this case, applying $g(x; p_0)$ to the local log GLR of an affected sensor yields $g(U_{n,k,t}; p_0) \approx -\log p_0 + (U_{n,k,t}^+)^2/2$, $n \in \mathcal{N}_a$. On the other hand, for the unaffected sensors $n \in \mathcal{N}_a^c$, $\exp[(U_{n,k,t}^+)^2/2]$ can be bounded with high probability since it is due to noise only. Hence, $g(U_{n,k,t}; p_0)$ can be bounded with high

probability for $n \in \mathcal{N}_a^c$. The function $g(x; p_0)$ for various values of p_0 is illustrated in Fig. 3.1. A special case is when $p_0 = 1$, $g(x; p_0) = (x^+)^2/2$, in which case the function applies no weighting on the log GLR statistics of any sensor. This is consistent with our definition, since $p_0 = 1$ means we believe all sensors observe the change-point and hence we should use GLR statistics of all sensors when combining them.

Usually very small changes are not of interest, so we can modify the definition of the mixture statistic (3.14) as follows. Instead of maximizing over possible change-point time k for $0 \leq k < t$, we can maximize over the most recent m_1 possible change-point time $t - m_1 \leq k < t$, where m_1 is a suitable window size. By applying time-windowing, we only examine possible change-points in the most recent m_1 observations. This also simplifies the computations required to implement the detection procedure. It is possible to require a minimum window size as well, say m_0 , if we are concerned about possible outliers in the data and hence do not allow change-point detection based on too few observations.

After time-windowing, the mixture procedure compares the log GLR statistic (3.14) with a threshold $b > 0$ and stops whenever the statistic exceeds the threshold. The associated stopping rule T_{mix} for the mixture procedure is given by

$$T_{\text{mix}} = \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \sum_{n=1}^N \log (1 - p_0 + p_0 \exp[(U_{n,k,t}^+)^2/2]) \geq b \right\}. \quad (3.16)$$

3.3 Other Detection Procedures

3.3.1 Maximizing Over Subset of Sensors

In the mixture procedure, we model the fraction of sensors affected by the change-point directly by assuming a parameter p_0 . On the other hand, without introducing the parameter p_0 , we can maximize the log likelihood ratio statistic over all unknown parameters, including the unknown subset of affected sensors. Using (3.5) and (3.3), such a log likelihood ratio is given by:

$$\begin{aligned} & \max_{\mathcal{N}_a \subset \mathcal{N}} \max_{0 \leq k < t} \max_{\mu_n > 0} \\ & \log \left\{ \frac{\left[\prod_{n \in \mathcal{N}_a^c} \prod_{l=1}^t \phi(y_{n,l}) \right] \cdot \left[\prod_{n \in \mathcal{N}_a} \prod_{l=1}^k \phi(y_{n,l}) \prod_{l=k+1}^t \phi(y_{n,l} - \mu_n) \right]}{\prod_{n=1}^N \prod_{l=1}^t \phi(y_{n,l})} \right\} \\ &= \max_{\mathcal{N}_a \subset \mathcal{N}} \max_{0 \leq k < t} \max_{\mu_n > 0} \sum_{n \in \mathcal{N}_a} \sum_{l=k+1}^t (\mu_n y_{n,l} - \mu_n^2/2) \end{aligned} \quad (3.17)$$

$$= \max_{\mathcal{N}_a \subset \mathcal{N}} \max_{0 \leq k < t} \sum_{n \in \mathcal{N}_a} \left[\left(\sum_{l=k+1}^t y_{n,l} \right)^+ \right]^2 / [2(t-k)] \quad (3.18)$$

$$= \max_{0 \leq k < t} \sum_{n=1}^N \left[\left(\sum_{l=k+1}^t y_{n,l} \right)^+ \right]^2 / [2(t-k)], \quad (3.19)$$

where from (3.17) to (3.18) we have used (3.10), and from (3.18) to (3.19) we have used the fact that the terms inside the summation in (3.18) are all nonnegative so maximizing over $\mathcal{N}_a \subset \mathcal{N}$ ends up taking $\mathcal{N}_a = \mathcal{N}$.

Note that when we use a set of assumed values $\{\tilde{\mu}_n\}$ as a surrogate for post-change

means μ_n in forming the detection statistics:

$$\max_{\mathcal{N}_a \subset \mathcal{N}} \max_{0 \leq k < t} \log \left\{ \frac{\left[\prod_{n \in \mathcal{N}_a^c} \prod_{l=1}^t \phi(y_{n,l}) \right] \cdot \left[\prod_{n \in \mathcal{N}_a} \prod_{l=1}^k \phi(y_{n,l}) \prod_{l=k+1}^t \phi(y_{n,l} - \tilde{\mu}_n) \right]}{\prod_{n=1}^N \prod_{l=1}^t \phi(y_{n,l})} \right\}, \quad (3.20)$$

using the similar argument to the above, the procedure based on (3.20) is identical to the modified TV procedure (3.28) that we will discuss in the following section.

3.3.2 Other Procedures

For the model in Section 3.1, there can be other procedures based on two extreme assumptions about the subset of affected sensors: either all sensors are affected or only a few sensor are affected. Several other detection procedures have also been considered in the literature, assuming a set of specific values of μ_n for the post-change means. Assuming a specific set of values of μ_n may also be used to specify a subset of sensors which we believe are affected by the change-point. We will discuss these procedures in the following.

First we derive two procedures using log GLR statistics based on two extreme assumptions about the subset of affected sensors. If we assume all the sensors are affected, we can derive the following procedure, which we refer to as the log GLR procedure:

$$T_{\text{GLR}} = \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \sum_{n=1}^N (U_{n,k,t}^+)^2 / 2 \geq b \right\}, \quad (3.21)$$

where m_1 is the window-length. When $N = 1$, the GLR procedure (3.21) corresponds to the procedure studied in [78] for a single sensor case. The log GLR procedure also corresponds to letting $p_0 = 1$ in the mixture procedure (3.16). The log GLR procedure

is efficient when the change-point affects a large fraction of the sensors.

At the other extreme, if the set \mathcal{N}_a is very small such that only a few sensors are affected by the change-point, a reasonable procedure is to take the maximum over the log GLR statistics (3.13) of each sensor. By using the maximum-likelihood estimates for the post-change means (3.10), we obtain the following max procedure:

$$T_{\max} = \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \max_{1 \leq n \leq N} (U_{n,k,t}^+)^2/2 \geq b \right\}, \quad (3.22)$$

where m_1 is the window length.

The following two procedures have been considered in the literature. They both use a set of assumed values $\{\tilde{\mu}_n\}$ as a surrogate for post-change means μ_n in forming the detection statistics. The set of $\{\tilde{\mu}_n\}$ essentially establish a minimum size of change of interest. The assumed post-change means $\{\tilde{\mu}_n\}$ can be different from $\{\mu_n\}$. This model mismatch may degrade performance of the detection procedures [56] [84].

Mei's procedure [56] forms a one-sensor CUSUM statistic for each sensor as defined in (2.12), and then sums these CUSUM statistics to form the detection statistic:

$$T_{\text{Mei}} = \inf \left\{ t \geq 1 : \sum_{n=1}^N \max_{0 \leq k \leq t} [\tilde{\mu}_n(S_{n,t} - S_{n,k}) - \tilde{\mu}_n^2(t-k)/2] \geq b \right\}. \quad (3.23)$$

For Mei's procedure, there are convenient recursive computational algorithms for the CUSUM statistics of each sensor (3.23). Define

$$W_{n,t} \triangleq \tilde{\mu}_n(S_{n,t} - S_{n,k}) - \tilde{\mu}_n^2(t-k)/2. \quad (3.24)$$

We can write the detection statistics (3.23) for Mei's procedure as $\sum_{n=1}^N \max_{0 \leq k < t} W_{n,t}$.

On the other hand, note that at time t

$$\begin{aligned}
& \max_{0 \leq k \leq t} W_{n,t} \\
&= \max_{0 \leq k \leq t} [\tilde{\mu}_n(S_{n,t} - S_{n,k}) - \tilde{\mu}_n^2(t - k)/2] \\
&= \max \{0, \tilde{\mu}_n(y_{n,t} - \tilde{\mu}_n/2), \dots, \tilde{\mu}_n(y_{n,t} - \tilde{\mu}_n/2) + \dots + \tilde{\mu}_n(y_{n,1} - \tilde{\mu}_n/2)\}.
\end{aligned} \tag{3.25}$$

Hence in the time step $t + 1$, we have

$$\begin{aligned}
& \max_{0 \leq k \leq t+1} W_{n,t+1} \\
&= \max \{0, \tilde{\mu}_n(y_{n,t+1} - \tilde{\mu}_n/2), \dots, \tilde{\mu}_n(y_{n,t+1} - \tilde{\mu}_n/2) + \dots + \tilde{\mu}_n(y_{n,0} - \tilde{\mu}_n/2)\} \\
&= \max \{0, \tilde{\mu}_n(y_{n,t+1} - \tilde{\mu}_n/2) + W_{n,t}\}.
\end{aligned} \tag{3.26}$$

So for Mei's procedure, when we increase the time step from t to $t + 1$, we simply add onto each $W_{n,t}$ the quantity, $\tilde{\mu}_n(y_{n,t+1} - \tilde{\mu}_n/2)$, take its positive part, and sum to form the new detection statistic.

The TV procedure [84] also assumes a set values of the post-change means $\tilde{\mu}_n$. In contrast to Mei's procedure, the TV procedure sums the log likelihood ratio of each sensor first, and then forms the CUSUM statistic (Mei's procedure forms a CUSUM statistic for each sensor). The TV procedure has a stopping rule defined as

$$T_{\text{TV}} \triangleq \inf \left\{ t \geq 1 : \max_{k < t} \sum_{n=1}^N [\tilde{\mu}_n(S_{n,t} - S_{n,k}) - \tilde{\mu}_n^2(t - k)/2] \geq b \right\}. \tag{3.27}$$

However, in the detection statistic of the TV procedure (3.28), when the change-point occurs, the statistics of the sensors that are not affected by the change-point have negative drifts. These negative drifts cancel the positive drifts of the statistics

formed by sensors that are affected by the change-point. To address this limitation, we suggest the following modification to the TV procedure, by taking the positive part of the statistic formed by each sensor before summing them:

$$T_{\text{TV,modified}} \triangleq \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \sum_{n=1}^N [\tilde{\mu}_n(S_{n,t} - S_{n,k}) - \tilde{\mu}_n^2(t-k)/2]^+ \geq b \right\}, \quad (3.28)$$

where m_1 is the window length. Comparing (3.23) and (3.27), we note that the detection statistic of Mei's procedure is greater than that of the modified TV procedure, since interchanging the order of maximizing with that of summing increases the value.

The threshold b of each of the procedures (3.16) - (3.27) defined above are chosen to meet the ARL requirement. In the following section, we will derive an approximation to the ARL of the mixture procedure. This approximation is parameterized by the threshold b . We can numerically solve for the threshold b to meet a given ARL value by inverting the approximation with respect to b . For other procedures, the threshold b is found by numerical Monte Carlo simulations.

3.4 Performance of Mixture Procedure

In this section we study the theoretical performance of the mixture procedure (3.16) using two performance metrics: the average run length (ARL) when there is no change, and the expected detection delay in the extreme case where a change occurs immediately at $\kappa = 0$. This provides an upper bound on the expected detection delay when a change occurs later in the sequence $\kappa > 0$.

We will need the following quantities to characterize the ARL and expected detection delay of the mixture procedure. Let z_1, z_2, \dots be i.i.d. with normal distribution

that has mean $\Delta^2/2$ and variance Δ^2 . Define a random walk

$$\tilde{S}_t \triangleq \sum_{l=1}^t z_l. \quad (3.29)$$

Also define a stopping time

$$\tau(c) = \min\{t \geq 1 : \tilde{S}_t > c\}, \quad (3.30)$$

and a stopped process $\tilde{S}_{\tau(c)}$. The first quantity is

$$\nu(\Delta) = \lim_{c \rightarrow \infty} \mathbb{E}[\exp\{-(\tilde{S}_{\tau(c)} - c)\}]. \quad (3.31)$$

The exact expression for $\nu(x)$ is given by (see page 82 of [80])

$$\nu(x) = 2|x|^{-2} \exp \left\{ -2 \sum_{l=1}^{\infty} l^{-1} \Phi \left(-|x| \sqrt{l}/2 \right) \right\}. \quad (3.32)$$

An approximation for $\nu(x)$ that is convenient for numerical computation is given by [79]:

$$\nu(x) \approx \frac{(2/x)[\Phi(x/2) - 0.5]}{(x/2)\Phi(x/2) + \phi(x/2)}.$$

where $\phi(x)$ and $\Phi(x)$ are the pdf and cdf of the standard normal distribution, with $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$, and $\Phi(x) = \int_{-\infty}^x \phi(z) dz$.

The second quantity is

$$\rho(\Delta) = \lim_{c \rightarrow \infty} \mathbb{E}[\tilde{S}_{\tau(c)} - c] = \frac{\mathbb{E}[\tilde{S}_{\tau(0)}^2]}{2\mathbb{E}[\tilde{S}_{\tau(0)}]}. \quad (3.33)$$

The exact computational expressions and useful approximations for (3.33) are available in [80]. From Lemma 2, one such expression for $\rho(\Delta)$ is given by

$$\rho(\Delta) = \frac{\mathbb{E}[z_1^2]}{2\mathbb{E}[z_1]} - \sum_{l=1}^{\infty} l^{-1} \mathbb{E}[\tilde{S}_l^-] = \frac{\Delta^2}{4} + 1 - \sum_{l=1}^{\infty} l^{-1} \mathbb{E}[\tilde{S}_l^-], \quad (3.34)$$

where the notation x^- denotes the negative part of a number x : $x^- = x$ if $x < 0$ and $x^- = 0$ if $x \geq 0$.

The third quantity is $\mathbb{E} \left[\min_{t \geq 1} \tilde{S}_t \right]$. Using Lemma 3, we can evaluate this quantity as

$$\mathbb{E} \left[\min_{t \geq 1} \tilde{S}_t \right] = \rho(\Delta) - 1 - \Delta^2/4. \quad (3.35)$$

3.4.1 Average Run Length (ARL)

The average run length (ARL) is the average length of intervals between two false-alarms when there is no change-point. It is a common performance metric for false-alarm rate of a sequential detection procedure. To characterize the ARL of the mixture procedure, we define the following:

$$\psi(\theta) = \log \mathbb{E}[\exp\{\theta g(U; p_0)\}], \quad (3.36)$$

where U has a standard normal distribution. Recall the definition (3.15) $g(x; p_0) = \log(1 - p_0 + p_0 e^{(x^+)^2/2})$. The first-order derivative of $g(x; p_0)$ with respect to x is given by:

$$\dot{g}(x; p_0) = p(x^+) e^{(x^+)^2/2} / [1 - p_0 + p_0 e^{(x^+)^2/2}], \quad (3.37)$$

where the dot denotes first-order derivative of the function. Based on (3.37), we define:

$$\gamma(\theta) = \frac{1}{2}\theta^2 \mathbb{E} \{ [\dot{g}(U; p_0)]^2 \exp[\theta g(U; p_0) - \psi(\theta)] \}, \quad (3.38)$$

and

$$f(N, \theta, p_0) = \frac{\gamma(\theta)N^{1/2}}{\theta\{2\pi\ddot{\psi}(\theta)\}^{1/2}} \exp\{-N[\theta\dot{\psi}(\theta) - \psi(\theta)]\}, \quad (3.39)$$

where the double-dot denotes the second-order derivative of the function.

The following approximation, for which we provide a heuristic argument, is the main result of this section:

Approximation 1. Assume $N \rightarrow \infty$ and $b \rightarrow \infty$, with b/N a fixed, positive constant. Define θ by $\dot{\psi}(\theta) = b/N$, with definition (3.15) for $g(x; p_0)$, notation (3.36), (3.38) and (3.39). Then

$$\mathbb{E}^\infty[T_{\text{mix}}] \sim \left[f(N, \theta, p_0) \int_{[2N\gamma(\theta)/m_1]^{1/2}}^{[2N\gamma(\theta)]^{1/2}} y\nu^2(y)dy \right]^{-1}. \quad (3.40)$$

Here the notation $x \sim y$ means $x/y \rightarrow 1$ asymptotically.

Remark: The integrand $y\nu^2(y)$ in (3.40) is integrable at both 0 and ∞ by virtue of the relations $\nu(y) \rightarrow 1$ as $y \rightarrow 0$, and $\nu(y) \sim 2/y^2$ as $y \rightarrow \infty$. Also note that Approximation 1 applies to any general function $g(x; p_0)$ that is second-order differentiable in x .

The following heuristic calculations provide support for Approximation 1. Let

$$Z_{k,t} = \sum_{n=1}^N \log \left(1 - p_0 + p_0 \exp[(U_{n,k,t}^+)^2/2] \right). \quad (3.41)$$

In [105] it was shown that for large m ,

$$\begin{aligned}
& \mathbb{P}^\infty \{T_{\text{mix}} \leq m\} \\
&= \mathbb{P}^\infty \left\{ \max_{t \leq m, m_0 \leq t-k \leq m_1} Z_{k,t} \geq b \right\} \\
&\sim N^2 e^{-N[\theta\dot{\psi}(\theta) - \psi(\theta)]} [2\pi N\ddot{\psi}(\theta)]^{-1/2} |\theta|^{-1} \gamma^2(\theta) \int_{\frac{m_0}{m}}^{\frac{m_1}{m}} \nu^2 ([2N\gamma(\theta)/(mt)]^{1/2}) (1-t) dt/t^2,
\end{aligned} \tag{3.42}$$

where N and b diverge to ∞ at the same rate, and m is also large, but small enough that the right hand side of (3.42) converges to 0. We are primarily interested in the case where $m_0 = 1$, and the maximum window size m_1 is small compared to m . Hence $1 - t \approx 1$ in the integrand. Using change of variables $u = mt$, $z = u/(2N\mu(\theta))$ and then $y = z^{-1/2}$, and the definition of f (3.39), we can rewrite (3.42) as

$$\begin{aligned}
& \mathbb{P}^\infty \{T_{\text{mix}} \leq m\} \\
&\sim mN^2 e^{-N\{\theta\dot{\psi}(\theta) - \psi(\theta)\}} \{2\pi N\ddot{\psi}(\theta)\}^{-1/2} |\theta|^{-1} \gamma^2(\theta) \int_1^{m_1} \nu^2 \left(\sqrt{\frac{2N\gamma(\theta)}{u}} \right) \frac{du}{u^2} \\
&= mN^2 e^{-N\{\theta\dot{\psi}(\theta) - \psi(\theta)\}} \{2\pi N\ddot{\psi}(\theta)\}^{-1/2} |\theta|^{-1} \gamma^2(\theta) \frac{1}{2N\gamma(\theta)} \int_{\frac{1}{2N\gamma(\theta)}}^{\frac{m_1}{2N\gamma(\theta)}} \nu^2 \left(\sqrt{\frac{1}{z}} \right) \frac{1}{z^2} dz \\
&= mN^2 e^{-N\{\theta\dot{\psi}(\theta) - \psi(\theta)\}} \{2\pi N\ddot{\psi}(\theta)\}^{-1/2} |\theta|^{-1} \gamma^2(\theta) \frac{1}{2N\gamma(\theta)} \int_{[2N\gamma(\theta)/m_1]^{1/2}}^{[2N\gamma(\theta)]^{1/2}} \nu^2(y) y^4 \cdot 2y^{-3} dy \\
&= mf(N, \theta, p_0) \int_{[2N\gamma(\theta)/m_1]^{1/2}}^{[2N\gamma(\theta)]^{1/2}} y \nu^2(y) dy.
\end{aligned} \tag{3.43}$$

Let λ_b denote the factor multiplying m in the last line of (3.43). The only restriction on m in (3.43) is that we should have $m\lambda_b \rightarrow 0$ when b and N go to infinity at the

same rate.

To deal with larger m , it can be shown using arguments similar to those in [1] and [81], and is verified numerically in Section 3.4.3, that T_{mix} is asymptotically exponentially distributed. More precisely, for any positive t ,

$$\mathbb{P}^\infty\{\lambda_b T_{\text{mix}} \leq t\} \sim [1 - \exp(-t)]. \quad (3.44)$$

Hence, $\lambda_b \mathbb{E}^\infty[T_{\text{mix}}] \rightarrow 1$, or $\mathbb{E}^\infty[T_{\text{mix}}] \sim 1/\lambda_b$ as N and b go to infinity at the same rate.

3.4.2 Expected Detection Delay

After a change-point occurs, we are interested in the expected number of additional observations required for detection. The maximum expectation detection delay over all possible change-point time $\kappa \geq 0$ happens when the change-point occurs at the first observation $\kappa = 0$. Hence in the change-point detection literature, the expected detection delay when a change-point occurs at the first observation is usually used as the performance metric. In this section, we derive an approximation to the expected detection delay of the mixture procedure.

We continue to use the notation of the proceeding section. In particular $g(x; p_0) = \log[1 - p_0 + p_0 e^{(x^+)^2/2}]$, and U denotes a standard normal random variable. Recall that \mathcal{N}_a denotes the set of sensors at which there is a change, M is the cardinality of this set, and $p = M/N$ is the fraction of sensors that are affected by the change-point. For each $n \in \mathcal{N}_a$, the mean of the observations changes from 0 to $\mu_n > 0$, and for $n \in \mathcal{N}_a$ the distribution of the observations is the same before and after the change-point occurs.

To study the expected detection of the mixture procedure, we need the following definitions. First, define a quantity

$$\Delta = \left(\sum_{n \in \mathcal{N}_a} \mu_n^2 \right)^{1/2}, \quad (3.45)$$

with Δ^2 interpreted as the total energy of the change-point observed by the sensors. Define the Kullback-Leibler (KL) divergence of the post-change distribution for a vector of sensor observations of all sensors with respect to their pre-change distribution as:

$$I(\mu_1, \dots, \mu_N) \triangleq \int \cdots \int f_1(u_1) \cdots f_1(u_N) \left[\sum_{n=1}^N \log(f_1(u_n)/f_0(u_n)) \right] du_1 \cdots du_N, \quad (3.46)$$

and in our case it is given by

$$\begin{aligned} I(\mu_1, \dots, \mu_N) &= \int \cdots \int \phi(u_1 - \mu_1) \cdots \phi(u_N - \mu_N) \left[\sum_{n=1}^N (\mu_n u_n - \mu_n^2/2) \right] du_1 \cdots du_N \\ &= \Delta^2/2. \end{aligned} \quad (3.47)$$

The K-L divergence determines the asymptotic rate of growth of the detection statistic after the change-point κ . From the law of large numbers, we see that a first order approximation of the expected detection delay is $2b/\Delta^2$, provided that the maximum window size m_1 is large compared to this quantity. In the following derivation we assume $m_1 \gg 2b/\Delta^2$.

Our main result of this section is the following approximation, which refines this first order result for the expected detection delay of the mixture procedure.

Approximation 2. As $b \rightarrow \infty$, with other parameters held fixed, the definition (3.15) for $g(x; p_0)$, the definition (3.45) for Δ , (3.33) for $\rho(\Delta)$, and (3.35) for $\mathbb{E} \left[\min_{t \geq 1} \tilde{S}_t \right]$, we have that the expected detection delay when the change-point happens at $\kappa = 0$ is given by

$$\begin{aligned} \mathbb{E}^0[T_{\text{mix}}] = \\ 2\Delta^{-2} \left\{ b - M \log p_0 + \rho(\Delta) - M/2 + \mathbb{E} \left[\min_{t \geq 1} \tilde{S}_t \right] - (N - M)\mathbb{E}[g(U; p_0)] + o(1) \right\}, \end{aligned} \quad (3.48)$$

where U is a normal random variable with zero mean and unit variance.

Remarks:

The first term on the right-hand-side of (3.48) is the first-order expected detection delay. The other terms capture the over-shoot of the detection statistic process over the threshold b , the effect of the non-linear transform $g(x; p_0)$ on the log GLR statistic of each sensor, and the second-order effects of noise. In particular, the positive term $-M \log p_0$ can be interpreted as a penalty because we do not know the subset of affected sensors. The term $\rho(\Delta)$ captures the expected overshoot of the detection statistic process over the threshold b . The term $\mathbb{E} \left[\min_{t \geq 1} \tilde{S}_t \right]$ is due to drift of the detection statistic from the affected sensors. For a residual process defined as the difference of the detection statistic process relative to its drift, the term $-M/2$ is due to the variance of that residual process for the affected sensors. The term $\mathbb{E}[g(U; p_0)]$ is due to noise of the unaffected sensors.

We can provide a heuristic argument to support Approximation 2 based on the idea as follows. We decompose the detection statistic of the mixture procedure (3.16) into two parts: one part due to the affected sensors, and the other part due to the

unaffected sensors. The first part from the affected sensors can further be decomposed into two terms: one term captures the positive drift of the statistic which eventually causes the statistic to exceed the threshold, and the other term that captures the negative drift of the statistic due to our uncertainty about the change-point time κ . Then we take expectation of the term, use Wald's Identity for the positive drift term, and bound the other terms. When applying Wald's Identity, we also take into account that the detection statistic process generally will not hit the threshold exactly and will have an overshoot over the threshold, and this overshoot significantly affects the accuracy of the approximation to the expected detection delay. We capture this overshoot using a technique from non-linear renewal theory [80].

The following heuristic argument provides a support to Approximation 2. For derivation for a similar problem involving a single sequence where $N = 1$, see [65] and [78]. In the following, for convenience, we write $T = T_{\text{mix}}$. Let

$$k_0 = b^{1/2}. \tag{3.49}$$

Recall the notation of $Z_{k,t}$ introduced in (3.41). For $k < T - k_0$, we can write the detection statistic at the stopping time T as follows, up to a term that tends to zero

exponentially fast in probability:

$$\begin{aligned}
Z_{k,t} &= \sum_{n=1}^N g(U_{n,k,T}; p_0) \\
&= \sum_{n \in \mathcal{N}_a} g(U_{n,k,T}; p_0) + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \\
&= \sum_{n \in \mathcal{N}_a} \log \left(p_0 \exp \{ (U_{n,k,T}^+)^2 / 2 \} \left[1 + \frac{1-p_0}{p_0} \exp \{ -(U_{n,k,T}^+)^2 / 2 \} \right] \right) \\
&\quad + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \\
&= \sum_{n \in \mathcal{N}_a} [\log p_0 + (U_{n,k,T}^+)^2 / 2] + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \\
&\quad + \sum_{n \in \mathcal{N}_a} \log \left(1 + \frac{1-p_0}{p_0} \exp \{ -(U_{n,k,T}^+)^2 / 2 \} \right) \\
&= M \log p_0 + \sum_{n \in \mathcal{N}_a} (U_{n,k,T}^+)^2 / 2 + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) + o(1) \\
&= M \log p_0 + \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})^+]^2 / 2(T-k) + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) + o(1).
\end{aligned} \tag{3.50}$$

The residual term $\sum_{n \in \mathcal{N}_a} \log \left(1 + (1-p_0) \exp \{ -(U_{n,k,T}^+)^2 / 2 \} / p_0 \right)$ tends to zero exponentially fast when $b \rightarrow \infty$, because when $b \rightarrow \infty$, $T \rightarrow b/\Delta$, and $n \in \mathcal{N}_a$, $(U_{n,k,T}^+)^2$ grows on the order of $\mu_n^2(T-k) > \mu_n^2 k_0 = \mu_n^2 \sqrt{b}$.

We then use the following simple identity to decompose the second term in (3.50) for the affected sensors into two parts:

$$\begin{aligned}
(S_{n,t}^+)^2 / 2t &= S_{n,t}^2 / 2t - (S_{n,t}^-)^2 / 2t \\
&= \mu_n (S_{n,t} - \mu_n t / 2) + (S_{n,t} - \mu_n t)^2 / 2t - (S_{n,t}^-)^2 / 2t.
\end{aligned} \tag{3.51}$$

From the preceding discussion, we see that $\max_{0 \leq k < T-k_0} Z_{k,T}$ is on the order of b , while

$\max_{T-k_0 \leq k < T} Z_{k,T}$ is on the order of $k_0 = b^{1/2}$. Hence with overwhelming probability the max over all k is attained for $k < T - k_0$, so from (3.51) and (3.50) we have

$$\begin{aligned}
& \max_{0 \leq k < T} Z_{k,t} \\
&= \max_{0 \leq k < T-k_0} \sum_{n=1}^N g(U_{n,k,T}; p_0) + o(1) \\
&= M \log p_0 + \max_{0 \leq k < T-k_0} \left[\sum_{n \in \mathcal{N}_a} \mu_n [(S_{n,T} - S_{n,k}) - (T-k)\mu_n/2] \right. \\
&\quad + \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T-k)\mu_n]^2 / [2(T-k)] \\
&\quad \left. - [(S_{n,T} - S_{n,k})^-]^2 / 2(T-k) + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \right] + o(1) \\
&= M \log p_0 + \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,T} - T\mu_n/2) + \\
&\quad \max_{0 \leq k < T-k_0} \left[- \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - k\mu_n/2) + \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T-k)\mu_n]^2 / [2(T-k)] \right. \\
&\quad \left. - \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})^-]^2 / [2(T-k)] + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \right] + o(1).
\end{aligned} \tag{3.52}$$

The following lemma forms the basis of the rest of the derivation (see Appendix B.1 for details).

Lemma 4. For $k_0 = b^{1/2}$, asymptotically as $b \rightarrow \infty$

$$\begin{aligned}
& \max_{0 \leq k < T-k_0} \left[- \sum_{n \in \mathcal{N}_a} \mu_n \left(S_{n,k} - \frac{k\mu_n}{2} \right) + \sum_{n \in \mathcal{N}_a} \frac{[(S_{n,T} - S_{n,k}) - (T-k)\mu_n]^2}{2(T-k)} \right. \\
& \quad \left. - \sum_{n \in \mathcal{N}_a} \frac{[(S_{n,T} - S_{n,k})^-]^2}{2(T-k)} + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \right] \\
&= \sum_{n \in \mathcal{N}_a} (S_{n,T} - T\mu_n)^2 / 2T + \sum_{n \in \mathcal{N}_a^c} g(U_{n,1,T}; p_0) + \max_{0 \leq k < k_0} \left[- \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - k\mu_n/2) \right] + o(1).
\end{aligned}$$

Using the preceding discussion and (3.52), by taking expectations, letting $b \rightarrow \infty$ and using Lemma 4, we have

$$\begin{aligned}
& \mathbb{E}^0 \left[\max_{0 \leq k < T} \sum_{n=1}^N g(U_{n,k,T}; p_0) \right] \\
&= \mathbb{E}^0 \left[M \log p_0 + \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,T} - T\mu_n/2) + \sum_{n \in \mathcal{N}_a} \frac{(S_{n,T} - T\mu_n)^2}{2T} + \sum_{n \in \mathcal{N}_a^c} g(U_{n,1,T}; p_0) + \right. \\
& \quad \left. \max_{0 \leq k < k_0} \left\{ - \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - k\mu_n/2) \right\} \right] + o(1).
\end{aligned} \tag{3.53}$$

We will compute each term on the right hand side of (3.53) separately. We need the following lemma due to Anscombe and Doeblin (see Theorem 2.40 in [80]), which states that the randomly stopped average of random variables are asymptotically normal distributed under general conditions.

Lemma 5 (Anscombe-Doeblin). *Let y_1, y_2, \dots be i.i.d. with mean μ and variance $\sigma^2 \in (0, \infty)$. Suppose D_c and c are positive integer valued random variables such that*

for some constants $m_c \rightarrow \infty$, D_c/m_c converges to one in probability. Then as $c \rightarrow \infty$

$$\mathbb{P} \left\{ \sqrt{D_c} \left[\frac{1}{D_c} \sum_{l=1}^{D_c} y_l - \mu \right] \leq x \right\} \rightarrow \frac{1}{\sigma} \Phi \left(\frac{x}{\sigma} \right), \quad (3.54)$$

where Φ denotes the standard normal cdf.

We continue bounding the terms in (3.53):

(i) By Wald's Identity [80]:

$$\mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a} \mu_n (S_{n,T} - T\mu_n/2) \right] = \mathbb{E}^0[T] \Delta^2/2. \quad (3.55)$$

(ii) By the Anscombe-Doeblin Lemma, $(S_{n,T} - T\mu_n)/T^{1/2}$ is asymptotically normally distributed with zero mean and unit variance. Hence $\sum_{n \in \mathcal{N}_a} (S_{n,T} - T\mu_n)^2/T$ is asymptotically a sum of independent χ_1^2 random variables, so

$$\mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a} (S_{n,T} - T\mu_n)^2/2T \right] = M/2 + o(1). \quad (3.56)$$

(iii) For the same reason as (ii):

$$\mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a^c} g(U_{n,0,T}; p_0) \right] \rightarrow (N - M) \mathbb{E}^0[g(U; p_0)]. \quad (3.57)$$

(iv) The term $-\sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - \mu_n k/2)$ is a random walk with index k with negative drift $-\Delta^2/2$ and variance Δ^2 . Hence $\mathbb{E}^0 \left[\max_{0 \leq k < k_0} -\sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - k\mu_n/2) \right]$ converges to the expected minimum of this random walk. By symmetry, this random walk has the same distribution as the process $\min_{t \geq 1} \tilde{S}_t$ defined above.

Having finished evaluating the right hand side of (3.53), we now consider the left-hand side. The first order asymptotic behavior of the process $\sum_{n=1}^N g(U_{n,k,T}; p_0)$ is the same as that of $\sum_{n \in \mathcal{N}_a} \mu_n (S_{n,T} - T\mu_n/2)$, which has drift $\Delta^2/2$ and variance Δ^2 . Equivalently, the process $\sum_{n=1}^N g(U_{n,k,T}; p_0)$ has the same distribution as the process \tilde{S}_T in (3.29). By writing

$$\mathbb{E}^0 \left[\max_{0 \leq k < T} \sum_{n=1}^N g(U_{n,k,T}; p_0) \right] = b + \mathbb{E}^0 \left[\max_{0 \leq k < T} \sum_{n=1}^N g(U_{n,k,T}; p_0) - b \right], \quad (3.58)$$

and using nonlinear renewal theory to evaluate the expected overshoot of the process of (3.29) over the boundary ([80], Chapter IX), we obtain

$$\mathbb{E}^0 \left[\max_{0 \leq k < T} \sum_{n=1}^N g(U_{n,k,T}; p_0) - b \right] \rightarrow \rho(\Delta). \quad (3.59)$$

3.4.3 Accuracy of Approximations for Mixture Procedure

We start with examining the accuracy of our approximations to the ARL in (3.40) and the expected detection delay in (3.48) for the mixture procedure. For the following Monte Carlo simulations we use $N = 100$ sensors, $m_1 = 200$ and $\mu_n = 1$ for affected sensors, and the results are obtained from 500 Monte Carlo trials. All sensor observations are contaminated by independent normally distributed noise with zero mean and unit variance.

First we verify our exponential distribution approximation for the stopping time T_{mix} under the null hypothesis. Consider a mixture procedure with $p_0 = 0.1$, $m_1 = 200$, and $b = 19.5$, which corresponds to the theoretical ARL approximation of 5000. The cumulative distribution function (cdf) of T_{mix} from theoretical approximation

is obtained from (3.43), and the cdf from Monte Carlo simulation is obtained from 500 Monte Carlo trials. Fig. 3.2 demonstrates that our the cdf using theoretical approximation based on the exponential distribution assumption fits well with the numerically simulated cdf.

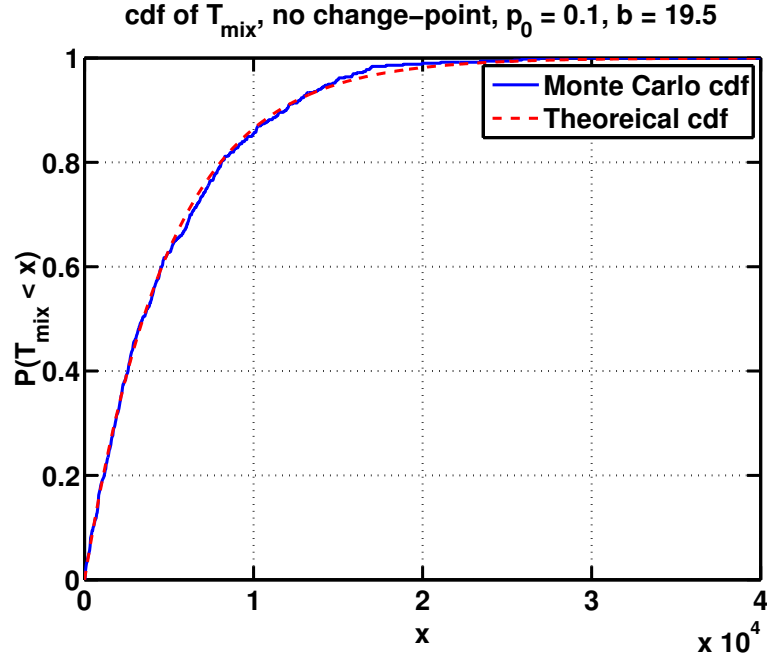


Figure 3.2: Cumulative distribution function (cdf) of T_{mix} from theoretical approximation versus from numerical simulation. The mixture procedure has $p_0 = 0.1$, $m_1 = 200$, and $b = 19.5$, which corresponds to the a theoretical approximation of ARL that equals 5000. The Monte Carlo results are obtained from 500 Monte Carlo trials.

Then we examine the accuracy of our approximation to ARL in Approximation 1. The comparisons between the theoretical approximated and Monte Carlo simulated ARLs for different values of p_0 are given in Table 3.1. Our numerical results in Table 3.1 demonstrate that the Approximation 1 is quite accurate.

Next we examine the accuracy of our approximation to the expected detection

delay in Approximation 2 when the change-point occurs at the first observation $\kappa = 1$. We start assuming the parameter in the mixture procedure equals true parameter value $p_0 = p$. The comparison of the theoretical approximated versus Monte Carlo simulated expected detection delays (Fig. 3.3) demonstrates that Approximation 2 is reasonably accurate in this case. Then we assume that the mixture procedure has no knowledge of the true parameter value and uses $p_0 = 0.1$. Hence, in this case, there is a mismatch between p and the true parameter value: $p_0 \neq p$. The comparison of the theoretical approximated versus Monte Carlo simulated expected detection delays (Fig. 3.4) demonstrates that the Approximation 2 is also reasonably accurate with the mismatch. Some values of the above comparisons are also listed in Table 3.2.

Table 3.1: Average run length (ARL) of the mixture procedure with $m_1 = 200$. The Monte Carlo results are obtained from 500 trials.

p_0	b	Approximation	Monte Carlo
0.3	31.2	5002	5504
0.3	32.3	10002	10221
0.1	19.5	5000	4968
0.1	20.4	10001	10093
0.03	12.7	5001	4830
0.03	13.5	10001	9948

Table 3.2: Expected detection delay of a mixture procedure with $m_1 = 200$ such that the theoretical approximation of ARL is equal to 5000. The Monte Carlo results are obtained from 500 trials.

p	p_0	b	Approximation	Monte Carlo
0.3	0.3	31.2	3.5	3.2 ± 0.7
0.3	0.1	19.5	5.2	5.0 ± 0.9
0.1	0.1	19.5	7.2	6.7 ± 1.7
0.03	0.1	19.5	13.9	14.4 ± 4.4
0.03	0.03	12.7	13.9	14.2 ± 4.4




Figure 3.3: Comparison of theoretical and simulated expected detection delay of a mixture procedure. For simulation, we use 500 Monte Carlo trials, with $p_0 = p$, $m_1 = 200$, and different b such that the theoretical approximation of ARL is equal to 5000.

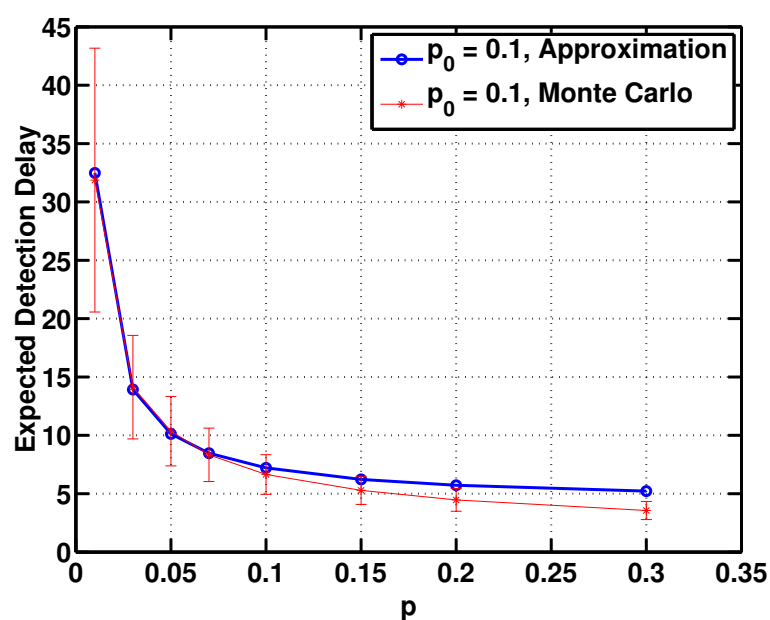


Figure 3.4: Comparison of theoretical and simulated expected detection delay of a mixture procedure. For simulation, we use 500 Monte Carlo trials, with $p_0 = p$, $m_1 = 200$ and $b = 19.5$ such that the theoretical approximation of ARL is equal to 5000.

3.5 Numerical Examples

In this section, we will demonstrate the performance of the mixture procedure and compare it with other detection procedures.

3.5.1 Robustness to Choice of p_0

First we demonstrate that the expected detection delay of the mixture procedure does not deviate much if the assumed p_0 is different from p . Again let $N = 100$ and all $\mu_n = 1$. We compare two mixture procedures. The first mixture procedure is ignorant of the true p . It sets $p_0 = 0.1$ regardless of p and chooses $b = 19.5$ so that the theoretical ARL approximation is equal 5000. The second mixture procedure knows the true fraction of affected sensors, sets $p_0 = p$ and chooses threshold values for each p such that the theoretical approximation of ARL is equal to 5000 for every p values. Fig. 3.5 shows the expected detection delays from Approximation 2 for these two mixture procedures when p varies from 0.01 to 0.3. Note that although the first procedure is ignorant of p and does not perform as well as the second procedure, which knows p , the first procedure does very well in a broad range of values where $0.03 < p < 0.2$. We will return to this point in Section 3.6.




Figure 3.5: Comparison of the expected detection delay from Approximation 2, when there is no mismatch, $p_0 = p$, and when $p_0 = 0.1$ so that there is a mismatch between p_0 and p . The thresholds are chosen such that the theoretical approximation of ARL is equal to 5000.

3.5.2 Effect of Window Size m_1

When approximating the expected detection delay, we assume that the window size m_1 is large enough such that when observations within the window are affected by a change-point, the mixture procedure of (3.16) can reach the threshold. In this example, we study how large m_1 should be for this assumption to hold. Assume $N = 100$, $p_0 = p = 0.03$, and consider two scenarios with $\mu_n = 1$ and $\mu_n = 0.6$, respectively. We vary m_1 from 10 to 100, and find b using Approximation 1 such that the theoretical approximation of ARL is equal to 5000. From the first order analysis, m_1 should be greater than $2b/\Delta^2$ for Approximation 2 to hold, which suggests that m_1 should be greater than 12 for $\mu_n = 1$, and be greater than 32 for $\mu_n = 0.6$. This is verified in Figure 3.6, which shows the expected detection delay from Monte Carlo simulation as m_1 increases. In Figure 3.6, the expected detection delay from simulation converges to that from Approximation 2 when $m_1 > 20$ for $\mu_n = 1$, and $m_1 > 40$ for $\mu_n = 0.6$. Hence, in practice, if we use m_1 sufficiently larger than $4b$, the approximation in (3.48) is guaranteed to hold.




Figure 3.6: Expected detection delay versus m_1 for the mixture procedure when the theoretical approximation of ARL is equal to 5000. The expected detection delay is obtained from Monte Carlo simulation. The Monte Carlo results are obtained from 500 trials.

3.5.3 Comparing Procedures

We compare the expected detection delays for the mixture procedure and other procedures, when their ARLs are all approximately 5000. The ARL for the mixture procedure is obtained from Approximation 1, and the ARLs for other procedures are obtained from 500 Monte Carlo trials. The expected detection delay of all procedures are obtained from 500 Monte Carlo trials. The thresholds for all procedures are listed in Table 3.3. Assume that the number of sensors $N = 100$ and $m_1 = 200$ for all procedures, except for Mei's procedure, which has a recursive implementation and hence needs no windowing. The mixture procedure sets $p_0 = 0.1$ such that it assumes no priori knowledge about the true fraction of affected sensors p . Mei's and the modified TV procedures both use $\tilde{\mu}_n = 1$ such that they both have no priori knowledge about the true post-change means. The expected detection delays are listed in Table 3.4. For the first two columns in the table, we set of post-change means to be the same across all affected sensors. In particular, we set $\mu_n = 1$, $n \in \mathcal{N}_a$ in the first column, and set $\mu_n = 0.7$, $n \in \mathcal{N}_a$ for the second column. For the third column, we choose post-change means arbitrarily and uniformly from the interval $[0.5, 1.5]$. Once these μ_n are chosen, they are fixed for all Monte Carlo trials. The expected detection delays for these three cases are plotted in Fig. 3.7.

Note that the max procedure (3.22) has the smallest detection delay when $p = 0.01$, but it has the greatest delay for p greater than 0.1. The modified TV procedure has a slight advantage in detection delay when p approaches 0.3 and larger values, but much longer delay as p approaches to 0.01. This is expected, since when most sensors are affected, the modified TV procedure collects most energy of the change-point. When only one sensor is affected, the max procedure collects energy most efficiently


in that it excludes noise from the unaffected sensors. The Mei's procedure performs well when p approaches 0.3 and larger values, but it has longer delay than the mixture procedure. The mixture procedure has the smallest detection delay when p is greater than 0.05, and it is only slightly slower than the max procedure when $p = 0.01$. In this respect, the mixture procedure performs best over a wide range of p values.


Table 3.3: Thresholds for procedures with $m_1 = 200$ such that their ARLs equal 5000. The Monte Carlo results are obtained from 500 trials.

Procedure	b	Monte Carlo ARL
Max	12.7	5041
GLR	53.4	4978
Mixture ($p_0 = 0.1$)	19.3	5000
Mei	88.5	4997
Modified TV	41.6	4993


Table 3.4: Expected Detection Delays for various procedures with $N = 100$ and $m_1 = 200$, obtained from 500 Monte Carlo trials. The mixture procedure has $p_0 = 0.1$. Mei's and the Modified TV procedures use $\tilde{\mu}_n = 1$. The thresholds of these procedures with ARL 5000 are listed in Table 3.3. We set $\mu_n = 1$ for $n \in \mathcal{N}_a$ in the first column, and $\mu_n = 0.7$ for $n \in \mathcal{N}_a$ for the second column. For the third column, we choose post-change means independently and uniformly from the interval $[0.5, 1.5]$.

p	method	DD, $\mu_n = 1$	DD, $\mu_n = 0.7$	DD, varying μ_n	
0.01	max	26.2	49.0	$\Delta^2 = 1.87$	14.9
	GLR	53.1	100.2		30.0
	mixture	31.7	61.5		18.5
	Mei	53.4	102.7		35.0
	Modified TV	83.6	226.0		48.2
0.03	max	17.9	32.9	$\Delta^2 = 3.33$	13.7
	GLR	19.1	36.8		17.1
	mixture	14.4	27.2		12.5
	Mei	22.6	40.5		22.2
	Modified TV	26.8	67.8		25.6
0.05	max	15.2	28.0	$\Delta^2 = 6.23$	10.4
	GLR	12.0	22.2		9.9
	mixture	10.1	18.4		7.9
	Mei	15.7	26.4		13.9
	Modified TV	14.9	38.2		12.7
0.1	max	12.4	22.9	$\Delta^2 = 14.22$	8.8
	GLR	6.6	11.8		5.1
	mixture	6.6	11.6		5.0
	Mei	9.5	15.0		8.1
	Modified TV	7.0	16.1		5.3
0.3	max	9.4	16.5	$\Delta^2 = 32.03$	7.8
	GLR	3.0	4.5		2.9
	mixture	3.5	5.6		3.3
	Mei	5.0	7.0		4.9
	Modified TV	3.1	4.2		3.0
0.5	max	8.2	14.7	$\Delta^2 = 51.69$	7.1
	GLR	2.2	3.1		2.2
	mixture	2.7	4.0		2.7
	Mei	3.9	5.0		3.8
	Modified TV	2.3	3.1		2.3
1	max	7.2	12.2	$\Delta^2 = 113.45$	5.7
	GLR	2.0	2.1		2.0
	mixture	2.0	2.6		2.0
	Mei	3.0	3.4		3.0
	Modified TV	2.0	2.1		2.0





(b)



(c)

Figure 3.7: Comparison of expected detection delays for various procedures with window size $m_1 = 200$, and (a): $\mu_n = 1$, $n \in \mathcal{N}_a$, (b): $\mu_n = 0.7$, $n \in \mathcal{N}_a$, (c): μ_n , $n \in \mathcal{N}_a$, arbitrarily and uniformly chosen from an interval $[0.5, 1.5]$. The mixture procedure uses $p_0 = 0.1$. Mei's and the modified TV procedures use $\tilde{\mu}_n = 1$. Monte Carlo results are obtained from 500 trials.

3.6 Parallel Mixture Procedure

We have demonstrated via simulation in Section 3.5.1 that the mixture procedure is not sensitive to the precision of p_0 . However, when N is large, a small difference in p will result in a large difference in the number of sensors affected by a change-point. For instance, when $N = 500$, $p = 0.1$ corresponds on 50 sensors being affected, while $p = 0.01$ corresponds to 5 sensors being affected. In this case, the mixture procedure may be more sensitive to a mis-specification of p_0 . To achieve robustness over a wider range of p , we consider a *parallel mixture procedure* that combines several mixture procedures, each using a different parameter p_0 and monitoring a different range of p values. The thresholds of these individual mixture procedures are chosen such that they have the same ARL. For example, we can use two mixture procedures with a small $p_0 = p_1$ and a large $p_0 = p_2$, respectively, and then choose thresholds b_1 and b_2 such that their ARLs are the same. Denote by $T_{\text{mix}}(p, b)$ the mixture procedure with parameter p and threshold b . Then the parallel procedure claims a detection once any of the mixture procedures reach their thresholds:

$$T_{\text{parallel}} \triangleq \min\{T_{\text{mix}}(p_1, b_1), T_{\text{mix}}(p_2, b_2)\}. \quad (3.60)$$

The ARL of the parallel mixture procedure is smaller than the ARL of each of the individual single mixture procedure. As we will demonstrate in the following, the parallel procedure has a smaller expected detection delay than a single mixture procedure over a wide range of p values.

To compare the performance of the parallel procedure with that of the single mixture procedure, we consider a case with $N = 500$ and $m_1 = 200$. The parallel

procedure uses two mixture procedures with $p_1 = 0.2$ and $p_2 = 0.03$ and threshold values $b_1 = 76.1$ and $b_2 = 26.9$ such that their ARLs are approximately 7000. The resulting ARL of the parallel mixture procedure is 6577. For the single mixture procedure, we use $p_0 = 0.1$ and threshold $b = 51.3$ such that the ARL of the single procedure is also 6577. Fig. 3.8 shows that the expected detection delay of the parallel procedure is smaller than that of the single procedure, especially when p is very small.

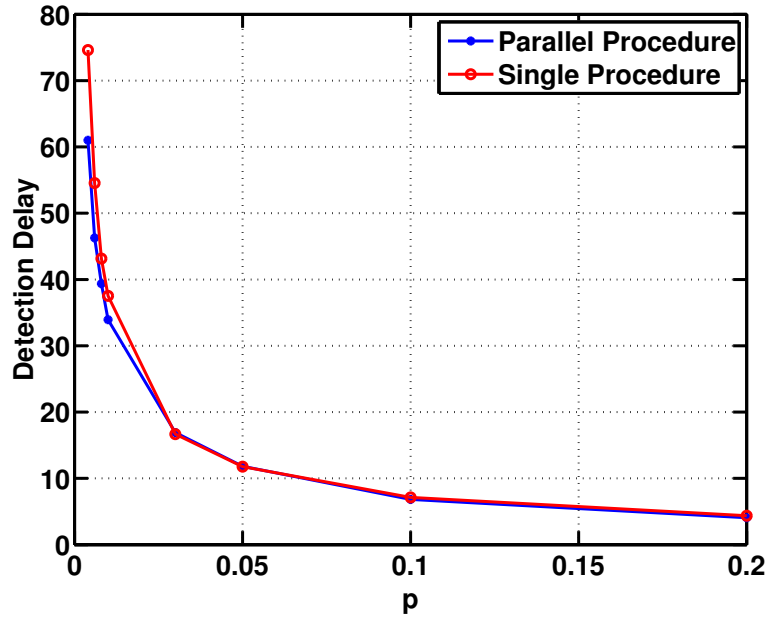


Figure 3.8: Expected detection delay of the parallel mixture procedure with $p_1 = 0.2$ and $p_2 = 0.03$ versus that of a single mixture procedure with $p_0 = 0.1$. The ARLs of the parallel procedure and the single procedure are both 6577. The Monte Carlo results are obtained from 500 trials.

3.7 Profile-Based Procedure

In discussions so far we have assumed that there is no spatial structure relating the post-change means at difference sensors, i.e., the post-change means of each sensor are unrelated to each other. We refer to such problems as unstructured problems. In the following we will consider an alternative scenario, where there is a parameterized profile function that determines the post-change means of each sensor, and hence the post-change means of each sensor are related to each other. We refer to such problems as structured problems. If we have some knowledge about this profile function, we can incorporate this knowledge into the definition of the likelihood ratio statistic and derive a detection procedure, referred to as the profile-based procedure. The profile-based procedure has been proposed and studied in [81]. Here we will review the profile-based procedure and compare it with the mixture procedure.

Consider a multi-sensor system, with the location of the n th sensor given by its coordinates in (for example) two-dimensional the Euclidean space (u_n, v_n) , $n = 1, \dots, N$. Suppose we are only interested in sources located in a region $\mathcal{D} \subset \mathbb{R}^2$. Assume there can be Q sources and these sources occur simultaneously. The post-change mean of the n th sensor takes a specific form:

$$\mu_n = \sum_{m=1}^Q r_m \alpha_{\mathbf{z}_m}(u_n, v_n). \quad (3.61)$$

The two-dimensional vector $\mathbf{z}_m \in \mathcal{D}$ is the (unknown) spatial location of the signal at the m th source, and r_m is the (unknown) amplitude of the m th source. The profile function is defined as

$$\alpha_{\mathbf{z}_m} : \mathbb{R}^2 \rightarrow \mathbb{R}^+, \quad (3.62)$$

and it describes how the amplitude of the m th source decays at the n th sensor. We assume that some knowledge about this profile function is available. For example, $\alpha_{\mathbf{z}}(u, v)$ can be a decreasing function of the Euclidean distance between the source location \mathbf{z} and the sensor location (u, v) . In principle, $\alpha_{\mathbf{z}}(u, v)$ may also depend on finitely many parameters, such as the rate of decay of the function. See [68] or [75] for examples in a fixed sample context.

If the change-point amplitudes r_m are multiplied by a positive constant and the profile $\alpha_{\mathbf{z}_m}(u_n, v_n)$ divided by the same constant, the values of the post-change means μ_n do not change. To remove this ambiguity, it is mathematically convenient to assume that the change-point profile has been standardized to have unit norm. Define a vector $\boldsymbol{\alpha}_{\mathbf{z}} = [\alpha_{\mathbf{z}}(u_1, v_1), \dots, \alpha_{\mathbf{z}}(u_N, v_N)]^\top$. The normalization of the profile function means that we require the norm of the profile function to be one:

$$\boldsymbol{\alpha}_{\mathbf{z}}^\top \boldsymbol{\alpha}_{\mathbf{z}} = 1, \quad (3.63)$$

for any $\mathbf{z} \in \mathcal{D}$.

3.7.1 Likelihood Ratio Statistic

In this section, we will derive the log GLR statistic under the assumption that the profile function (3.61) take a specific form. First, assume that the post-change mean (3.61) is due to only one source:

$$\mu_n = r\alpha_{\mathbf{z}_0}(u_n, v_n). \quad (3.64)$$

Using (3.64), we see that the log likelihood ratio for observations up to time t of all sensors (3.7), with an assumed change-point occurs at time $\kappa = k$, is given by the following:

$$\begin{aligned} l(t, k, r, \mathbf{z}) &= \sum_{n=1}^N \sum_{l=k+1}^t [r\alpha_{\mathbf{z}}(u_n, v_n)y_{n,l} - r^2\alpha_{\mathbf{z}}^2(u_n, v_n)/2] \\ &= r(t-k)^{1/2}\boldsymbol{\alpha}_{\mathbf{z}}^\top \mathbf{U}_{k,t} - (t-k)r^2/2, \end{aligned} \quad (3.65)$$

where $\mathbf{y}_s = [y_s^1, \dots, y_s^N]^\top$ and $\mathbf{U}_{k,t} = [U_{1,k,t}, \dots, U_{N,k,t}]^\top$. In (3.65), we have used the property (3.63) that the profile function has unit norm for any source location $\mathbf{z} \in \mathcal{D}$.

Then we solve for the maximum-likelihood estimate of r from (3.65), by taking the derivative of the log likelihood function (3.65) with respect to r , setting it to zero:

$$(t-k)^{1/2}\boldsymbol{\alpha}_{\mathbf{z}}^\top \mathbf{U}_{k,t} - (t-k)r = 0. \quad (3.66)$$

and solving for an solution, assuming $r > 0$:

$$\hat{r}_k = \left[\frac{\boldsymbol{\alpha}_{\mathbf{z}}^\top \mathbf{U}_{k,t}}{(t-k)^{1/2}} \right]^+. \quad (3.67)$$

Substitution of the maximum likelihood estimate \hat{r}_k into (3.65) leads to

$$l(t, k, \hat{r}_k, \mathbf{z}) = \frac{1}{2}[(\boldsymbol{\alpha}_{\mathbf{z}}^\top \mathbf{U}_{k,t})^+]^2. \quad (3.68)$$

We further maximize the function (3.68) with respect to the change-point time k and source location \mathbf{z} to obtain the log GLR statistic. Using this log GLR statistic, the

profile-based procedure is given by:

$$T_{\text{profile}} = \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \max_{\mathbf{z} \in \mathcal{D}} \frac{1}{2} [(\boldsymbol{\alpha}_{\mathbf{z}}^\top \mathbf{U}_{k,t})^+]^2 \geq b \right\}. \quad (3.69)$$

Note that (3.69) is a matched-filter type of statistic: the statistic $\mathbf{U}_{k,t}$ is matched to a profile function $\boldsymbol{\alpha}_{\mathbf{z}}$.

When the post-change means μ_n are due to multiple sources as described in (3.61), the corresponding log GLR statistic is more complex and is derived in Appendix A.2. We say that the sources are non-overlapping if $\boldsymbol{\alpha}_{\mathbf{z}_m}^\top \boldsymbol{\alpha}_{\mathbf{z}_p} = 0$, for $p \neq m$, i.e., the amplitudes of the sources decay fast enough such that no two sources have profiles overlap with each other. In the later numerical examples in Section 3.7.3, we will focus on the case with non-overlapping sources.

3.7.2 Theoretical ARL of Profile-Based Procedure

In this section, we derive a theoretical approximation to ARL of the profile-based procedure using results presented in [81], when the profile function is given by a Gaussian function:

$$\alpha_{\mathbf{z}}(u, v) = (\sqrt{2\pi\beta})^{-1/2} e^{-[(z_1-u)^2 + (z_2-v)^2]/(4\beta)}, \quad \mathbf{z} = (z_1, z_2) \in \mathbb{R}^2, \beta > 0. \quad (3.70)$$

In (3.70), a sensor is located at (u, v) and a source is located at $\mathbf{z} = (z_1, z_2)$, and the parameter $\beta > 0$ controls of the decay rate of profile function. Define the inner product of two functions in L_2 as

$$\langle f, g \rangle = \int \int_{\mathbb{R}^2} f(u, v) g(u, v) du dv. \quad (3.71)$$

For short-hand we also write $\alpha_{\mathbf{z}}$ for $\alpha_{\mathbf{z}}(u, v)$. It can be verified that $\langle \alpha_{\mathbf{z}}, \alpha_{\mathbf{z}} \rangle = 1$, and that the norm (3.63) of (3.70) is approximately one: $\alpha_{\mathbf{z}}^\top \alpha_{\mathbf{z}} \approx 1$, for $\mathbf{z} \in \mathcal{D} \subset \mathbb{R}^2$. So this function $\alpha_{\mathbf{z}}$ approximately satisfies (3.63), provided β is sufficiently large, the distance between points of the grid is sufficiently small, and hence the sum over all sensor locations is well approximated by integration over the entire Euclidean space [81].

In [81], the theoretical approximation for the ARL of a profile-based procedure with a general profile function is given. We can adapt the result in [81] to find the following result for a profile-based procedure with a Gaussian profile function (3.70):

$$\mathbb{P}^\infty\{T_{\text{profile}} \leq t\} = 2 \cdot t e^{-b/2} (b/4\pi)^{3/2} \cdot \sqrt{2} \int_{\sqrt{b/m_1}}^{\sqrt{b}} x \nu^2(x) dx \cdot \int_D \sqrt{|\langle \dot{\alpha}_{\mathbf{z}}, \dot{\alpha}_{\mathbf{z}}^\top \rangle|} d\mathbf{z}. \quad (3.72)$$

To evaluate the last term in (3.72), which involves an integration corresponding to the surface area of $\alpha_{\mathbf{z}}$ over D , we do the following computation. By (3.70), we have

$$\dot{\alpha}_{\mathbf{z}} = \left[\frac{d\alpha_{\mathbf{z}}}{dz_1}, \frac{d\alpha_{\mathbf{z}}}{dz_2} \right]^\top = (2\pi\beta)^{-1/2} e^{-\frac{1}{4\beta}[(u-z_1)^2 + (v-z_2)^2]} [(u-z_1)/(2\beta), (v-z_2)/(2\beta)]^\top. \quad (3.73)$$

Hence

$$\langle \dot{\alpha}_{\mathbf{z}}, \dot{\alpha}_{\mathbf{z}}^\top \rangle = \begin{bmatrix} \alpha_{11} & \alpha_{12} \\ \alpha_{21} & \alpha_{22} \end{bmatrix}, \quad (3.74)$$

where

$$\begin{aligned}
\alpha_{11} &= 1/(4\beta^2) \int \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(u-z_1)^2} (u-z_1)^2 du = 1/4\beta^2, \\
\alpha_{12} &= 1/(4\beta^2) \int \int \frac{1}{2\pi} e^{-\frac{1}{2}(u-z_1)^2 - \frac{1}{2}(v-z_2)^2} (u-z_1)(v-z_2) dudv = 0 \\
\alpha_{21} &= 1/(4\beta^2) \int \int \frac{1}{2\pi} e^{-\frac{1}{2}(u-z_1)^2 - \frac{1}{2}(v-z_2)^2} (u-z_1)(v-z_2) dudv = 0 \\
\alpha_{22} &= 1/(4\beta^2) \int \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(v-z_2)^2} (v-z_2)^2 dv = 1/4\beta^2.
\end{aligned} \tag{3.75}$$

The determinant of (B.35) is given by $|\langle \dot{\alpha}_{\mathbf{z}}, \dot{\alpha}_{\mathbf{z}}^\top \rangle| = 1/(16\beta^4)$. Hence

$$\int_D \sqrt{|\langle \dot{\alpha}_{\mathbf{z}}, \dot{\alpha}_{\mathbf{z}}^\top \rangle|} d\mathbf{z} = |D|/4\beta^2, \tag{3.76}$$

with $|D|$ denoting the area of D . Since T_{profile} is asymptotically exponentially distributed when b is large, substituting (3.76) into (3.72) and using a similar argument to the argument in Approximation 1, we have as $b \rightarrow \infty$:

$$\mathbb{E}^\infty[T_{\text{profile}}] \sim e^{b/2} (b/4\pi)^{-3/2} \cdot \left[\sqrt{2} \int_{\sqrt{b/m_1}}^{\sqrt{b}} x \nu^2(x) dx \cdot \frac{D}{2\beta^2} \right]^{-1}. \tag{3.77}$$

3.7.3 Numerical Examples

Two Methods to Simulate ARL

We can use two methods to simulate the ARL.

The first method has been used in Section 3.3 to verify the numerical accuracy of our approximation to the ARL of the mixture procedure. The first method directly simulates the detection statistic process. In particular, we generate N sequences of i.i.d. zero-mean normal random variables starting from $t = 1$. We compute the

detection statistic, for example, using (3.16) for the mixture procedure, let the process continue until it hits the threshold, and record the time when it stops. We call such realization a Monte Carlo trial. We repeat such trial for 500 times, compute the mean of the stopping times, and use it as the average run length.

However, the first method, direct simulation of the stopping times to estimate a large average run length is very time consuming, especially when the number of sensors is large. As an alternative, we can also consider the following method to indirectly simulate the average run length, which we refer to as the second method.

In *the second method*, for each Monte Carlo trial, we also generate N sequences of i.i.d. zero-mean normal random variables starting from $t = 1$. But we only continue process for a fixed short time, say $t = 250$. Then we record the value of the process at that time. We repeat this trial for many times (typically 500 times), and compute the frequency that the processes exceed a threshold at the end-point $t = 250$. This frequency can be used to approximate ARL of a detection procedure for the following reasons. We have shown in Section 3.4.3 that we can assume the stopping time of the mixture procedure is asymptotically exponential distributed when the threshold b is large. A similar assumption can be made for other detection procedures when their thresholds are large. Under the exponential assumption, if $\mathbb{E}^\infty[T] = 1/\lambda$, we have $\mathbb{P}^\infty\{T \leq t\} \approx 1 - e^{-t\lambda} \approx t\lambda$, when t is small relative to λ^{-1} . Using this relationship, we can estimate that $\mathbb{E}^\infty[T] \approx t/\mathbb{P}^\infty(T \leq t)$. For example, when $1/\lambda = 5000$, for $t = 250$, the probability is given by $\mathbb{P}^\infty\{T \leq t\} \approx t\lambda = 250/5000 = 0.05$. Hence if we choose the threshold b such that $\mathbb{P}^\infty\{T \leq 250\} \approx 0.05$, the resulted ARL is approximately 5000.

Small Example

Consider a structured problem where the profile function is the Gaussian function (3.70) with parameter $\beta = 1/4$. There are $N = 100$ sensors uniformly distributed on a 10 by 10 grid with locations on integer points (u_n, v_n) with $u_n = -4, -3, \dots, 5$ and $v_n = -4, -3, \dots, 5$. The sensor locations as well as the signal amplitude (which determines the post-change means) contour when there are four sources are illustrated in Fig. 3.9. There can be up to four non-overlapping sources with identical profile function described by (3.70) with $\beta = 1/4$. In comparing the performance of the mixture procedure and the profile-based procedure, we consider four scenarios. For the one-source case, the source amplitude $r_1 = 1$ and the location is $z_1 = (-1, -1)$. For the two-source case, the source amplitudes $r_1 = r_2 = 1/\sqrt{2}$ and locations are $z_1 = (-1, -1)$ and $z_2 = (3, 2)$. For the three-source case, the source amplitudes are $r_p = 1/\sqrt{3}$, $p = 1, 2, 3$ and locations are $z_1 = (-1, -1)$, $z_2 = (3, 2)$, and $z_3 = (-2, 3)$. For the four-source case, the source amplitudes are $r_p = 1/2$, $p = 1, \dots, 4$ and the source locations are $z_1 = (-1, -1)$, $z_2 = (3, 2)$, $z_3 = (-2, 3)$ and $z_4 = (5, -2)$. The amplitudes r_m of the sources in the multiple-signal cases are chosen such that the total energy of the sources adds up to one.

We choose $m_1 = 200$ for both the mixture procedure and the profile-based procedure. The profile-based procedure assumes a profile-function that is identical to the true profile function when forming the detection statistic. When there is one source, approximately 10% sensors are affected by the change-point. Since the sources are non-overlapping, when there are two, three and four sources occurring simultaneously, the fractions of affected sensors are 20%, 30% and 40% for these different cases, respectively. In all the four cases, we set $p_0 = 0.1$ for the mixture procedure such that

it assumes no priori knowledge of the number of sources.




Figure 3.9: Four non-overlapping sources observed by sensors located on a 10 by 10 grid ($N = 100$). The sources have Gaussian profile function given in (3.70) with $\beta = 1/4$. The outer circle for the contour for each source corresponds to 0.1 of the peak value. The dots represent location of the sensors.

Table 3.5: Expected detection delays for mixture procedure versus the profile-based procedure for the small example. The mixture procedure uses $p_0 = 0.1$, and the profile-based procedure assumes correct profile-function. They both use $m_1 = 200$. The thresholds are obtained by Approximation 1 for the mixture procedure, and by Monte Carlo simulation using *the second method* for the profile-based procedure, respectively. The Monte Carlo results are obtained from 500 trials.

	Profile-based, $b = 16.5$	Mixture, $b = 19.3$
One-source	13.5	36.8
Two-source	16.5	38.5
Three-source	18.2	43.8
Four-source	20.3	44.1

The thresholds (listed in Table 3.5) of the profile-based procedure and the mixture procedure are chosen such that their ARLs are about 5000. The threshold of the profile-based procedure is simulated from 500 Monte Carlo trials using *the second method*. The threshold of the mixture procedure is obtained from Approximation 1.

Comparison of the expected detection delay of the profile-based procedure with that of the mixture procedure (listed in Table 3.5) demonstrates that the mixture procedure detects slower than the profile-based procedure. However, in this case, the profile-based procedure assumes a correct profile function. When there is a mismatch between the assumed profile function and the true one, the performance of the profile-based procedure will degrade.

Large Example

Then we consider a structured problem where there are more number of sensors, $N = 625$, distributed over a square grid. The profile function is the Gaussian function (3.70) with parameter $\beta = 1$, so the profile function also decays slower and

has a wider spread than the previous example. The sensors are located on (u_n, v_n) with $u_n = -12, -11, \dots, 12$ and $v_n = -12, -11, \dots, 12$. There can be up to three non-overlapping signal sources occurring simultaneously, and for all scenarios we set the amplitudes of the sources such that the total energy adds up to one. The sensor locations as well as the signal amplitude contour when there are three sources simultaneously are illustrated in Fig. 3.10.

We choose $m_1 = 200$ for both the mixture procedure and the profile-based procedure. The profile-based procedure assumes a profile-function that is identical to the true profile function when forming the detection statistic. When there is one source, approximately $p = 0.016$ sensors are affected. The mixture method assumes $p_0 = 0.01$ for all scenarios regardless of the number of sources.

The thresholds of these procedures are chosen such that their ARLs are approximately 5000. These thresholds are listed in Table 3.6. We also verify the threshold for the profile-based procedure in this example using the theoretical approximation for ARL in (3.77). The threshold obtained by approximation is $b = 29.7$, which is very close to the threshold obtained by Monte Carlo simulation using the second method, which is $b = 28.7$. The threshold of the mixture procedure is obtained from Theorem 1.

The comparison of the expected detection delay of the profile-based procedure with that of the mixture procedure (listed in Table 3.6) demonstrates that for this larger example with wider spread sources, it is even harder for the mixture procedure to obtain a comparable performance to that the profile-based procedure.

Table 3.6: Expected detection delays for mixture procedure versus the profile-based procedure for the small example. The mixture procedure uses $p_0 = 0.01$, and the profile-based procedure assumes correct profile-function. They both use $m_1 = 200$. The thresholds are obtained by Approximation 1 for the mixture procedure, and by Monte Carlo simulation using *the second method* for the profile-based procedure, respectively. The Monte Carlo results are obtained from 500 trials.

	Profile-based, $b = 28.7$	Mixture, $b = 17.3$
One-source	25.6	84.1
Two-source	39.3	62.5
Three-source	50.3	108.6

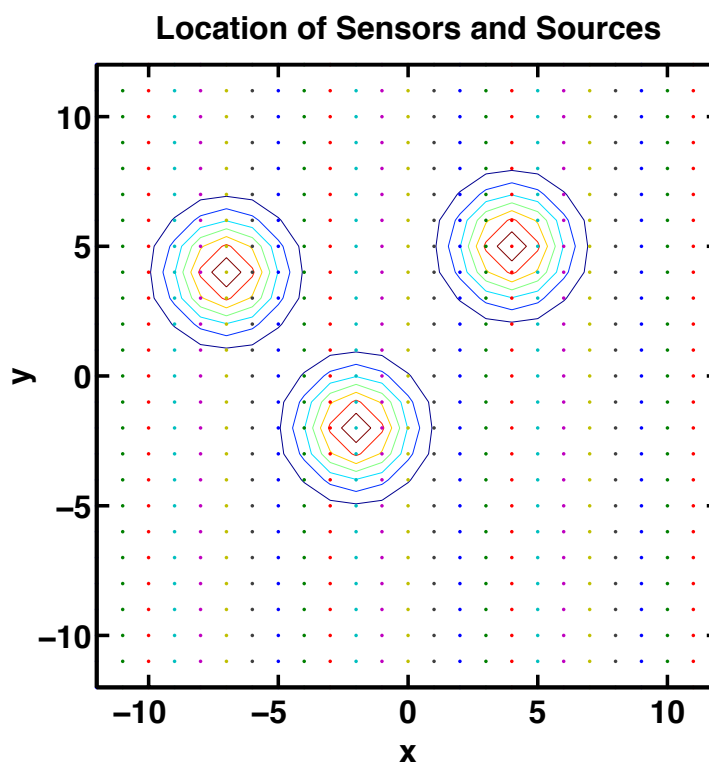


Figure 3.10: Three non-overlapping sources observed by sensors located on a 25 by 25 grid ($N = 625$). The sources have Gaussian profile function given in (3.70) with $\beta = 1$. The outer circle for the contour for each source corresponds to 0.1 of the peak value. The dots represent location of the sensors.

3.8 Conclusions

In this chapter, we have presented a mixture procedure for the multi-sensor sequential change-point detection problem. We assume that the pre-change and post-change sensor observations are normally distributed and the change-point affects the means of the observations. The post-change means as well as the subset of affected sensors are both unknown. The fraction of sensors affected by the change-point is given by a small number p , which is a form of sparsity. The mixture procedure models this sparsity by assuming that each sensor has probability p_0 to be affected by the change-point, where the value of p_0 is a guess for p . Based on this model, the mixture procedure forms a mixture log generalized likelihood ratio (GLR) statistic for change-point detection. The mixture statistic essentially applies a nonlinear weighting function on the log GLR statistic of each sensor before combining them. The mixture statistic is then compared with a threshold, and an alarm is raised whenever the statistic exceeds the threshold. We derive theoretical approximations to two performance metrics of the mixture procedure, the average run length (ARL), which is used to characterize the false alarm rate, and the expected detection delay. These approximations are shown to have good numerical accuracy. We demonstrated the performance of the mixture procedure using numerical simulations, showing that the mixture procedure does not require a precise knowledge of p . We also numerically compared the performance of the mixture procedure with that of the other procedures. We demonstrated that the mixture procedure has a lower expected detection delay than other procedures in the unstructured problems, and higher expected detection delay than a profile-based procedure in the structured problem, when the profile-based procedure assumes a correct profile function.

Chapter 4

Reduced-Dimension Multi-User Detection

We have introduced the notion of multiuser detection (MUD) in Chapter 1. Both linear and nonlinear MUDs have sufficiently high complexity to preclude their wide adoption in deployed systems. They both require the number of correlators at the receiver front-end to be equal to the number of users in the system. The complexity-per-bit of the nonlinear detectors is exponential in the number of users in system for the MLSE detector, and less for the DF detector. The complexity-per-bit of the linear detectors is linear in the number of users, which is much less than that of the nonlinear methods. In a typical communication system, there may be thousands of users. Hence, the complexity of the conventional methods has been a major obstacle for implementing the MF-bank based conventional MUD detectors. Our methods reduce the front-end complexity far below that of the nonlinear and linear techniques in large systems, and their complexity-per-bit is comparable to those of the conventional nonlinear and linear techniques. For example, when the random partial discrete

Fourier transform (DFT) matrix is used to construct the correlating signals, the number of correlators used by our methods is proportional to the logarithm of the number of users and proportional to the square of the number of active users, and the complexity-per-bit of our method is comparable to that of the corresponding conventional MUD detectors.

In this section, we will develop a low complexity front-end for MUD along with a set of appropriate detectors. We call this structure a reduced-dimension multiuser detection (RD-MUD). The RD-MUD reduces the number of correlators while still achieving performance similar to that of conventional detectors based on the MF-bank front-end. We reduce complexity by exploiting an empirical observation that the number of active users, K , is typically much smaller than the total number of users, N , at any given time. We call this observation *user sparsity*. In contrast to prior work on MUD with user sparsity, we use techniques from analog compressed sensing, which exploits sparsity in analog signals. The RD-MUD front-end projects the received signal into a lower dimensional *detection subspace*, and we will develop two detectors to detect active users and recover their data from the detection subspace. The two detectors are the linear reduced-dimensional decorrelating (RDD) detector and the nonlinear reduced-dimensional decision-feedback (RDDF) detector. The performance of RD-MUD depends on the coefficient matrix \mathbf{A} we use to form the correlating signals. We will present theoretical performance guarantees for the performance of RD-MUD in terms of the coherence property of \mathbf{A} . The theoretical results are also validated by numerical results.

The rest of the chapter is organized as follows. Section 4.1 discusses the system model and reviews MF-bank based detectors. Section 4.2 presents the RD-MUD

front-end and detectors. Section 4.3 contains the theoretical performance guarantee of two RD-MUD detectors: RDD and RDDF. Section 4.4 demonstrates performance of these techniques through numerical examples, and finally Section 4.5 concludes the paper.

4.1 System Model

Consider a multiuser system with N users. Each user is assigned a unique signature waveform from a set $\mathcal{S} = \{s_n(\cdot) : [0, T] \rightarrow \mathbb{R}, 1 \leq n \leq N\}$. The signature waveforms are assumed given and known, and possess certain properties discussed in more detail below. Each user modulates its data signal by its signature waveform to transmit its data symbols. The symbols carry information and are chosen from a constellation depending on the specific modulation scheme. The duration of the data symbol T is referred to as the symbol time, which is also equal to the inverse of the data rate for binary modulation.

Define the *inner product* (or *crosscorrelation*) between two real analog signals $x(t)$ and $y(t)$ in L_2 as

$$\langle x(t), y(t) \rangle = \frac{1}{T} \int_0^T x(t)y(t)dt \quad (4.1)$$

over the symbol time T . We also define the L_2 norm of a real analog signal $x(t)$ as

$$\|x(t)\| = \langle x(t), x(t) \rangle^{1/2}. \quad (4.2)$$

Two signals are orthogonal if their crosscorrelation is zero. We assume that the signature waveforms are linearly independent. That is, any linear combination of different signature waveforms cannot be another signature waveform (or its multiple)

in order to avoid cancellation of each other's transmission. The crosscorrelations of the signature waveforms are characterized by the Gram matrix \mathbf{G} , defined as

$$[\mathbf{G}]_{nl} \triangleq \langle s_n(t), s_l(t) \rangle, \quad 1 \leq n \leq N, \quad 1 \leq l \leq N. \quad (4.3)$$

For convenience, we assume that $s_n(t)$ has unit energy: $\|s_n(t)\| = 1$ for all n so that $[\mathbf{G}]_{nn} = 1$. Due to our assumption of linear independence of the signature waveforms, \mathbf{G} is invertible. The signature waveforms typically have low crosscorrelations, so we also assume that the magnitudes of the off-diagonal elements of \mathbf{G} are much smaller than 1.

We consider a synchronous MUD model with Binary Phase Shift Keying (BPSK) modulation [96]. There are K active users out of N possible users transmitting to the receiver. The set \mathcal{I} contains indices of all active users, and its complement set \mathcal{I}^c contains indices of all non-active users. The active users modulate their signature waveforms using BPSK modulation with the symbol of user n denoted by $b_n \in \{1, -1\}$, for $n \in \mathcal{I}$. The n th user transmits its modulated signal at power P_n through a wireless channel with channel amplitude g_n . We assume that the channel amplitude g_n can be estimated and is known to the receiver (CSIR). Define the gain $r_n \triangleq \sqrt{P_n}g_n$ where we assume r_n is also known at the receiver. For simplicity, we assume that g_n is real and, hence, r_n is real and can be negative. The nonactive user can be viewed as transmitting with power $P_n = 0$, or equivalently transmitting zeros: $b_n = 0$, for $n \in \mathcal{I}^c$.

The received signal $y(t)$ is a superposition of the transmitted signals from the

active users, plus white Gaussian noise $w(t)$ with zero-mean and variance σ^2 :

$$y(t) = \sum_{n=1}^N r_n b_n s_n(t) + w(t), \quad t \in [0, T], \quad (4.4)$$

with $b_n \in \{1, -1\}$, $n \in \mathcal{I}$, and $b_n = 0$, $n \in \mathcal{I}^c$. In the presence of delays of the asynchronous channel model, the ideas can be combined with the methods developed in [36] for time-delay recovery from low-rate samples.

The goal of multiuser detection (MUD) is to detect the set of active users, i.e. users with indices in \mathcal{I} , and their transmitted symbols $\{b_n : n \in \mathcal{I}\}$. In practice the number of active users K is typically much smaller than the total number of users N , which is a form of user sparsity. As we will show, this user sparsity enables us to reduce the number of correlators at the front-end and still be able to achieve performance similar to that of a conventional MUD using a bank of MFs. To simplify the detection algorithm, we assume that K is known. To consider joint estimation of K as well as the active users and their symbols greatly complicates analysis for the probability-of-error, since in that case the error will come from three sources: estimating the wrong number of active users, estimating the wrong set of active users, and estimating the wrong symbols. By assuming K is known, we will consider the latter two sources of errors in our analysis. The problem of estimating K can be treated using techniques such as these in [102] and [10].

In the following subsection, we present the structure of conventional MUD detectors using the MF-bank front-end. In Section 4.2 we introduce the new reduced-dimension MUD (RD-MUD) front-end and detectors.




Figure 4.1: Problem set up of MUD.

4.1.1 Conventional MUD

A conventional MUD detector has a front-end consisting of a bank of matched filters followed by a digital detector. In the following we review the MF-bank front-end and various digital detectors.

MF-bank front-end

For single-user systems, the matched filter (MF) passes the signal $s(t)$ through a filter with impulse response $s^*(T - t)$ or, equivalently, multiplies the received signal $s(t)$ with itself and integrates over a symbol time: this maximizes the output SNR of the decision statistic among the class of all linear detectors [96]. The MF-bank is an extension of the MF for multiple users, and it has N MFs in parallel: the n th branch correlates the received signal with the corresponding signature waveform $s_n(t)$,




Figure 4.2: Front-end of conventional MUD using MF-bank.

as illustrated in Fig. 4.2. As stated earlier, the output of the MF-bank is a set of sufficient statistics for MUD when the gains r_n are known [96]. Hence, no information is left in the MF-bank front end of the MUD receiver.

Using the model for the received signal (4.4), the output of the n th correlator of the MF-bank can be written as

$$z_n = \langle y(t), s_n(t) \rangle = r_n b_n + \sum_{l \neq n} [\mathbf{G}]_{nl} r_l b_l + u_n, \quad 1 \leq n \leq N. \quad (4.5)$$

The output noise $u_n = \langle w(t), s_n(t) \rangle$ is a Gaussian random variable with zero mean and covariance $\mathbb{E}\{u_n(t)u_m(t)\} = \sigma^2[\mathbf{G}]_{nm}$ (for derivation see Section 2.9 in [96]). Let $\mathbf{z} = [z_1, \dots, z_N]^\top$, $\mathbf{R} \in \mathbb{R}^{N \times N}$ be a diagonal matrix with $[\mathbf{R}]_{nn} = r_n$, $\mathbf{b} = [b_1, \dots, b_N]^\top$ and $\mathbf{u} = [u_1, \dots, u_N]^\top$. We can express the output of the MF-bank (4.5) in a vector




Figure 4.3: Front-end of RD-MUD.




Figure 4.4: The diagram of a linear detector.

form as:

$$\mathbf{z} = \mathbf{G}\mathbf{R}\mathbf{b} + \mathbf{u}, \quad (4.6)$$

where \mathbf{u} is a Gaussian random vector with zero mean and covariance $\mathbb{E}\{\mathbf{u}\mathbf{u}^H\} = \sigma^2\mathbf{G}$.

MF-bank detection

Conventional MUD detectors based on using the MF-bank output can be classified into two categories: linear and nonlinear detectors. The diagrams of these two detectors are illustrated in Fig. 4.4 and Fig. 4.5. In the literature, typically the basic

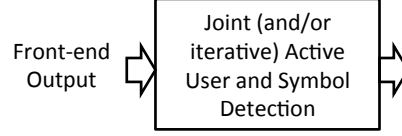


Figure 4.5: The diagram of a non-linear detector.

synchronous MUD model assumes all users are active, i.e. $b_n \in \{1, -1\}$, and hence the goal of the MUD detectors is to detect all user symbols.

The linear detector applies a linear transform to the MF-bank output (illustrated in Fig. 4.4):

$$\mathbf{Tz} = \mathbf{TGRb} + \mathbf{Tu}, \quad (4.7)$$

and detects symbol for each user separately using a sign detector:

$$\hat{b}_n = \text{sgn}(r_n[\mathbf{Tz}]_n), \quad 1 \leq n \leq N. \quad (4.8)$$

Several commonly used linear detectors are the single-user detector, the decorrelating detector and the minimum-mean-square-error (MMSE) detector. The single-user detector [54] is equivalent to having $\mathbf{T} = \mathbf{I}$ in (4.7), and detecting symbols as $\hat{b}_n = \text{sgn}(r_n z_n)$. The decorrelating detector [54] is motivated by the fact that when the signature waveforms are nonorthogonal, even in the absence of noise, the single-user detector (4.8) will not lead to good detection, when there is non-negligible interference between users (captured mathematically in (4.9) by the second term in parenthesis), since generally

$$\hat{b}_n = \text{sgn}(r_n z_n) = \text{sgn} \left(r_n^2 b_n + \sum_{l \neq n, l=1}^N [\mathbf{G}]_{nl} r_n r_l b_l \right) \neq \text{sgn}(r_n^2 b_n) = b_n. \quad (4.9)$$

The decorrelating detector addresses this problem by applying a linear transform $\mathbf{T} = \mathbf{G}^{-1}$ in (4.7). The decorrelating detector can recover symbols perfectly in the absence of noise; however, it also amplifies noise when $\mathbf{G} \neq \mathbf{I}$ and requires \mathbf{G} to be invertible. The minimum mean square error (MMSE) detector is based on the criterion to minimize the mean square error (MSE) between the linear transform of the MF-bank output and symbols. Based on this principle, MMSE uses a linear transform $\mathbf{T} = (\mathbf{G} + \sigma^2 \mathbf{R}^{-2})^{-1}$ in (4.7) [96].

The nonlinear detectors, on the other hand, detect symbols jointly or iteratively, as illustrated in Fig. 4.5. The nonlinear detectors include the maximum likelihood sequence estimator (MLSE) and the successive interference cancellation (SIC) detector [96]. The MLSE achieves the minimum probability-of-error by solving the following optimization problem to detect symbols jointly:

$$\max_{b_n \in \{1, -1\}} 2\mathbf{y}^H \mathbf{R} \mathbf{b} - \mathbf{b}^H \mathbf{R} \mathbf{G} \mathbf{R} \mathbf{b}. \quad (4.10)$$

If the signature waveforms are orthogonal, \mathbf{G} is a diagonal matrix and the optimization problem of (4.10) decouples into N single user problems, and the optimal solution is the sign detector (4.8) (with $\mathbf{T} = \mathbf{I}$). However, when the signature waveforms are nonorthogonal this optimization problem (4.10) is exponentially complex in the number of users [95]. It can be solved by exhaustive search which is computationally very expensive. The SIC detector first finds the active user with the largest gain, detects its symbol, subtracts its effect from the received signal, and iterates the above process using the residual signal. After K iterations, the SIC detector determines all active users. As we will show later, the ideas of SIC and the orthogonal matching pursuit (OMP) in compressed sensing are similar.

4.2 Reduced-Dimension MUD (RD-MUD)

The RD-MUD front-end, illustrated in Fig. 4.3, correlates the received signal $y(t)$ with a set of correlating signals $h_m(t)$, $m = 1, \dots, M$, where M is typically much smaller than N . This is in contrast to the conventional matched filter (MF) bank, which correlates the received signal with the full set of N signature waveforms [96]. The front-end output is processed by either a linear or nonlinear detector to detect active users and their symbols, as shown in Fig. 4.4 and 4.5 for both linear and nonlinear detectors.

4.2.1 RD-MUD: Front-End

The design of the correlating signals $h_m(t)$ is the key for RD-MUD to reduce the number of correlators. To construct these signals, we rely on biorthogonal waveforms. The related notion of the biorthogonal function has been used in analog compressed sensing to construct multichannel filters to sample the analog signal [27]. Define a delta function $\delta_{n,m}$, which is equal to one when $n = m$, and is equal to zero otherwise. The biorthogonal signals with respect to $\{s_n(t)\}$ are defined as a linear combination of all signature waveforms using a weighting coefficient matrix \mathbf{G}^{-1} :

$$\hat{s}_n(t) = \sum_{l=1}^N [\mathbf{G}^{-1}]_{nl} s_l(t), \quad 1 \leq n \leq N, \quad (4.11)$$

and have the property that $\langle s_n(t), \hat{s}_m(t) \rangle = \delta_{nm}$, for all n, m , i.e. the set of signals $\{\hat{s}_m(t)\}$ are biorthogonal to the signature waveforms $\{s_n(t)\}$. This is because

$$\langle s_n(t), \hat{s}_m(t) \rangle = \left\langle s_n(t), \sum_{l=1}^N [\mathbf{G}^{-1}]_{ml} s_l(t) \right\rangle = \sum_{l=1}^N [\mathbf{G}^{-1}]_{ml} \langle s_n(t), s_l(t) \rangle = \sum_{l=1}^N [\mathbf{G}^{-1}]_{ml} [\mathbf{G}]_{ln} = \delta_{nm}. \quad (4.12)$$

Since we assume that the signature waveforms are linearly independent in Section 4.1, \mathbf{G} is invertible. Note that when $\{s_n(t)\}$ are orthogonal, $\mathbf{G} = \mathbf{I}$ and $\hat{s}_n(t) = s_n(t)$.

The RD-MUD front-end uses as its correlating signals $\{h_m(t)\}$ which are linear combinations of the biorthogonal waveforms. The linear combination uses (possibly complex) weighting coefficients a_{mn} that we choose:

$$h_m(t) = \sum_{n=1}^N a_{mn} \hat{s}_n(t), \quad 1 \leq m \leq M. \quad (4.13)$$

The performance of RD-MUD depends on these coefficients $\{a_{mn}\}$. Define a coefficient matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ with $[\mathbf{A}]_{mn} \triangleq a_{mn}$ and denote the n th column of \mathbf{A} as $\mathbf{a}_n \triangleq [a_{1n}, \dots, a_{Mn}]^\top$, $n = 1, \dots, N$. We normalize the columns of \mathbf{A} so that $\|\mathbf{a}_n\| = 1$. The design of the correlating signals is equivalent to the design of the coefficient matrix \mathbf{A} for a given $\{s_n(t)\}$. We will use coherence as a measure of the quality of \mathbf{A} , which is defined as [13][9]:

$$\mu \triangleq \max_{n \neq l} |\mathbf{a}_n^H \mathbf{a}_l|. \quad (4.14)$$

As we will show later, it is desirable that the columns of \mathbf{A} have small correlation such that μ is small. This requirement for small μ also reflects a tradeoff in choosing how many correlators M to use in the RD-MUD front-end. We will show later that with

more correlators, the coherence of \mathbf{A} can be lower and the performance of RD-MUD can be better.

We now derive the output of the RD-MUD front-end. From the choice of the correlating signals (4.13) and the receive signal model (4.4), the output of the m th correlator is given by:

$$y_m = \langle h_m(t), y(t) \rangle \quad (4.15)$$

$$= \left\langle \sum_{n=1}^N a_{mn} \hat{s}_n(t), \sum_{l=1}^N r_l b_l s_l(t) \right\rangle + \left\langle \sum_{n=1}^N a_{mn} \hat{s}_n(t), w(t) \right\rangle \quad (4.16)$$

$$= \sum_{l=1}^N r_l b_l \sum_{n=1}^N a_{mn} \langle \hat{s}_n(t), s_l(t) \rangle + w_m \quad (4.17)$$

$$= \sum_{l=1}^N a_{ml} r_l b_l + w_m, \quad (4.18)$$

where we have defined $w_m \triangleq \left\langle \sum_{n=1}^N a_{mn} \hat{s}_n(t), w(t) \right\rangle$, and used the property that $\langle s_n(t), \hat{s}_m(t) \rangle = \delta_{nm}$. We define the output noise

$$w_m = \sum_{n=1}^N a_{mn} \langle \hat{s}_n(t), w(t) \rangle, \quad 1 \leq m \leq M. \quad (4.19)$$

This is a Gaussian random variable with zero-mean, variance $\sigma_m^2 \triangleq \sigma^2 [\mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H]_{mm}$, and covariance $\rho_{nm} \triangleq \mathbb{E}\{w_n w_m\} = \sigma^2 [\mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H]_{nm}$ (for derivation see Appendix B.1). Denoting $\mathbf{y} = [y_1, \dots, y_M]^\top$ and $\mathbf{w} = [w_1, \dots, w_M]^\top$, we can express the RD-MUD output (4.18) in vector form as

$$\mathbf{y} = \mathbf{A} \mathbf{R} \mathbf{b} + \mathbf{w}, \quad (4.20)$$

where \mathbf{w} is a Gaussian random vector with zero mean and covariance $\sigma^2 \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H$. The vector \mathbf{y} can be viewed as a linear projection of the MF-bank front-end output onto a lower dimensional subspace which we call the *detection subspace*. Since there are at most K active users, \mathbf{b} has at most K non-zero entries. The idea of RD-MUD is that when the original signal vector \mathbf{b} is sparse, with proper choice of the matrix \mathbf{A} , the detection performance for \mathbf{b} based on \mathbf{y} of (4.20) in the detection subspace can be similar to the performance based on \mathbf{z} (4.6), the output of the MF-bank front-end.

4.2.2 RD-MUD: Detectors

We now discuss how to recover \mathbf{b} from the RD-MUD front-end output \mathbf{y} of (4.20) using digital detectors. The model for the output (4.20) of the RD-MUD front-end has a similar form to the observation model in the compressed sensing literature [31][9], except that the noise in the RD-MUD front-end output is colored due to matched filtering at the front-end. Hence, to recover \mathbf{b} , we can adopt the ideas developed in the context of compressed sensing, and combine them with techniques in MF-bank detection.

The linear detector for RD-MUD first estimates active users $\hat{\mathcal{I}}$ using support recovery techniques from compressed sensing [70][31][83]. These support recovery techniques include thresholding [11], which we will describe in more detail later and use for active user detection in the RDD detector and other linear detectors, and orthogonal matching pursuit (OMP) [63][85], which we extend to decision-feedback OMP and use for joint active user and symbol detection in the RDDF detector.

Once the active users are estimated, their symbols $\mathbf{b}_{\hat{\mathcal{I}}}$ can be detected. Based on the estimated index set of active users $\hat{\mathcal{I}}$ (which may be different from \mathcal{I}), we can

write the RD-MUD front-end output model (4.20) as

$$\mathbf{y} = \mathbf{A}_{\hat{\mathcal{I}}} \mathbf{R}_{\hat{\mathcal{I}}} \mathbf{b}_{\hat{\mathcal{I}}} + \mathbf{w}. \quad (4.21)$$

The symbols $\mathbf{b}_{\hat{\mathcal{I}}}$ can be detected from (4.21) by applying a linear transform to the front-end output and detecting symbols separately. The nonlinear detector for RD-MUD detects active users and their symbols jointly (and/or iteratively).

We will focus on recovery based on two algorithms: (1) the reduced-dimension decorrelating (RDD) detector, a linear detector that uses subspace projection along with thresholding [11] to determine active users and sign detection for data recovery; (2) the reduced-dimension decision feedback (RDDBF) detector, a nonlinear detector that combines decision-feedback orthogonal matching pursuit (DF-OMP) for active user detection and sign detection for data recovery. DF-OMP differs from conventional OMP [63][85] in that in each iteration, the binary-valued detected symbols, rather than the real-valued estimates, are subtracted from the received signal to form the residual used by the next iteration. The residual consists of the remaining undetected active users. By subtracting interference from the strongest active user we make it easier to detect the remaining active users. DF-OMP can also be viewed as a successive interference cancellation (SIC) detector [100][25] in the detection subspace: DF-OMP detects each user by computing the inner product of the received signal vector and the signature waveform vector in the projection subspace and subtracts the effect of each user using the data model (4.20), whereas the SIC detector does this using a difference data model (4.6). These two algorithms are summarized in Table 4.1.

Reduced-dimension decorrelating (RDD) detector

The RDD detector works as follows. As per (4.20), the front-end of the RD-MUD projects the received signal $y(t)$ onto the detection subspace as a vector \mathbf{y} . By considering the RD-MUD front-end output when the input signal is $s_n(t)$, we can show that the column \mathbf{a}_n of \mathbf{A} corresponds to the n th signature waveform vector in the detection subspace:

$$\langle h_m(t), s_n(t) \rangle = \sum_{l=1}^N a_{ml} \langle \hat{s}_l(t), s_n(t) \rangle = a_{mn}, \quad 1 \leq m \leq M. \quad (4.22)$$

Considering the detection method of the conventional MUD, a natural strategy for RD-MUD is to match the received signal vector \mathbf{y} and the n th signature waveform vector in the detection subspace by computing their inner product, which is given by $\mathbf{a}_n^H \mathbf{y}$, $n = 1, \dots, N$. To detect active users, we can rank the magnitudes of these inner products and detect the index of the K largest as active users:

$$\hat{\mathcal{I}} = \{n : \text{if } |\Re[\mathbf{a}_n^H \mathbf{y}]| \text{ is among the } K \text{ largest of } |\Re[\mathbf{a}_n^H \mathbf{y}]|, n = 1, \dots, N\}. \quad (4.23)$$

This method (4.23) has also been used for sparsity pattern recovery in the compressed sensing literature (e.g. [31]). To detect their symbols, we use sign detection:

$$\hat{b}_n = \begin{cases} \text{sgn}(r_n \Re[\mathbf{a}_n^H \mathbf{y}]), & n \in \hat{\mathcal{I}}; \\ 0, & n \notin \hat{\mathcal{I}}. \end{cases} \quad (4.24)$$

In detecting active users (4.23) and their symbols (4.24), we take the real parts of the inner products because the imaginary part of $\mathbf{a}_n^H \mathbf{y}$ contains only noise and

interference. To see this, expand

$$\mathbf{a}_n^H \mathbf{y} = r_n b_n + \sum_{l \neq n, l \in \mathcal{I}} r_l b_l (\mathbf{a}_n^H \mathbf{a}_l) + \mathbf{a}_n^H \mathbf{w}, \quad n \in \mathcal{I}, \quad (4.25)$$

$$\mathbf{a}_n^H \mathbf{y} = \sum_{l \in \mathcal{I}} r_l b_l (\mathbf{a}_n^H \mathbf{a}_l) + \mathbf{a}_n^H \mathbf{w}, \quad n \notin \mathcal{I}. \quad (4.26)$$

Recall that symbols b_n and gains r_n are real and only \mathbf{A} can be complex. Hence the term $r_n b_n$ in (4.25), which contains the transmitted symbol, is real, and the rest of the terms in (4.25) and (4.26), which contain noise and interference, are complex. For real \mathbf{A} , (4.23) and (4.24) are equivalent to decisions based on the magnitudes of the inner products.

The RDD detector computes the inner products of the received signal vector and the signature waveform vector $\mathbf{a}_n^H \mathbf{y}$, $n = 1, \dots, N$, which is equivalent to computing $\mathbf{A}^H \mathbf{y}$. This requires MN floating point operations when \mathbf{A} is real (or $2MN$ operations when \mathbf{A} is complex) for detection of $N \log_2 3$ bits (since equivalently we are detecting $b_n \in \{-1, 0, 1\}$). Hence the complexity-per-bit of the RDD detector is proportional to M . For other RD-MUD linear detectors, following the inner product computation, linear processing can be done by multiplying an $N \times N$ matrix that incurs N^2 operations. Hence the complexity-per-bit of other RD-MUD linear detectors is proportional to $M + N$. Since $M \leq N$ in RD-MUD, the complexity for data detection of the RDD detector and other RD-MUD linear detectors is on the same order as that of the conventional linear MUD detector. But the RDD detector and other linear RD-MUD detectors require much lower decorrelation complexity in the analog front-end than the conventional linear detector.

Table 4.1: RD-MUD Detection

RDD detector
<p>Detect active users: find $\hat{\mathcal{I}}$ that contains indices of the K largest $\Re[\mathbf{a}_n^H \mathbf{y}]$.</p> <p>Detect symbols: $\hat{b}_n = \text{sgn}(r_n \Re[\mathbf{a}_n^H \mathbf{y}])$ for $n \in \hat{\mathcal{I}}$, and $\hat{b}_n = 0$ for $n \notin \hat{\mathcal{I}}$.</p>
RDDF detector
<p>Initialize: $\hat{\mathcal{I}}$ is empty, $\mathbf{b}^{(0)} = 0$, $\mathbf{v}^{(0)} = \mathbf{y}$.</p> <p>Iterate K times: $k = 1, \dots, K$.</p> <p>Detect active user: $n_k = \arg \max_n \Re[\mathbf{a}_n^H \mathbf{v}^{(k-1)}]$</p> <p>Detect symbol: $\hat{b}_n^{(k)} = \text{sgn}(r_{n_k} \Re[\mathbf{a}_{n_k}^H \mathbf{v}^{(k-1)}])$, for $n = n_k$, and $\hat{b}_n^{(k)} = \hat{b}_n^{(k-1)}$ for $n \neq n_k$.</p> <p>Update: $\hat{\mathcal{I}}^{(k)} = \hat{\mathcal{I}}^{(k-1)} \cup \{n_k\}$, and $\mathbf{v}^{(k)} = \mathbf{y} - \mathbf{A} \mathbf{R} \mathbf{b}^{(k)}$.</p>

Reduced-dimension decision-feedback (RDDB) detector

The RDDB detector determines active users and their corresponding symbols iteratively. It starts with an empty set as the initial estimate for the set of active user $\hat{\mathcal{I}}^0$, zeros as the estimated symbol vector $\mathbf{b}^{(0)} = \mathbf{0}$, and the front-end output as the residual vector $\mathbf{v}^{(0)} = \mathbf{y}$. Subsequently, in each iteration $k = 1, \dots, K$, the algorithm selects the column \mathbf{a}_n that is most highly correlated with the residual $\mathbf{v}^{(k-1)}$ as the detected active user in the k th iteration, with the active user index:

$$n_k = \arg \max_n |\Re[\mathbf{a}_n^H \mathbf{v}^{(k-1)}]|. \quad (4.27)$$

This index is then added to the active user set $\hat{\mathcal{I}}^{(k)} = \hat{\mathcal{I}}^{(k-1)} \cup \{n_k\}$. The symbol for user n_k is detected with other detected symbols staying the same:

$$b_n^{(k)} = \begin{cases} \text{sgn}(\Re[r_{n_k} \mathbf{a}_{n_k}^H \mathbf{v}^{(k-1)}]), & n = n_k; \\ b_n^{(k-1)}, & n \neq n_k. \end{cases} \quad (4.28)$$

Then the residual vector is updated through

$$\mathbf{v}^{(k)} = \mathbf{y} - \mathbf{A} \mathbf{R} \mathbf{b}^{(k)}. \quad (4.29)$$

The residual vector represents the part of \mathbf{b} that has yet to be detected by the algorithm along with noise. The iteration repeats K times (as we will show, with high probability DF-OMP never detects the same active user twice), and finally the active user set is given by $\hat{\mathcal{I}} = \hat{\mathcal{I}}^{(K)}$ with the symbol vector $\hat{b}_n = b_n^{(K)}$, $n = 1, \dots, N$.

The RDDF detector computes the inner product between the received signal vector and the signature waveform vector $\mathbf{a}_n^H \mathbf{v}^{(k)}$, $n = 1, \dots, N$, for $k = 1, \dots, K$. This requires KMN floating point operations when \mathbf{A} is real ($2KMN$ operations when \mathbf{A} is complex) for detection of $N \log_2 3$ bits. Hence the complexity-per-bit is proportional to KM . Since $M \leq N$, this implies that the complexity for data detection of the RDDF detector is on the same order as that of the conventional DF detector (the complexity-per-bit of the DF detector is proportional to KN). But the RDDF detector requires much lower decorrelation complexity in the analog front-end than the conventional DF detector.

Noise whitening transform

The noise in the RD-MUD output (4.20) is in general colored due to the matched filtering at the front-end. We can whiten the noise at the front-end output by applying a linear transform before detecting active users and symbols, as illustrated in Fig. 4.6. The linear transform to whiten noise in the RD-MUD output is given by $(\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1/2}$, and the whitened output is given by:

$$\mathbf{y}_w \triangleq (\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1/2} \mathbf{y} = (\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1/2} \mathbf{A}\mathbf{R}\mathbf{b} + \mathbf{w}_0, \quad (4.30)$$

where \mathbf{w}_0 is a Gaussian random vector with zero mean and covariance matrix $\sigma^2 \mathbf{I}$. If we define a new measurement matrix

$$\mathbf{A}_w \triangleq (\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1/2} \mathbf{A}, \quad (4.31)$$

then the RDD and RDDF detectors work with the whitened output (4.30) if we replace \mathbf{A} with \mathbf{A}_w and \mathbf{y} with \mathbf{y}_w in (4.23), (4.24), (4.27) and (4.28). While whitening the noise in the RD-MUD front-end output, the noise whitening transform also distorts the signal component. As we will demonstrate via numerical examples in Section 4.4.2, the benefits of noise whitening exceed the impact of the corresponding signal detection only when the signature waveforms $\{s_n(t)\}$ are highly correlated. Since this is typically not the case in multiuser systems due to the interference between users that coexists [96], our analysis will focus on detectors without noise whitening, and the benefits of omitting the noise whitening demonstrated in our numerical results.




Figure 4.6: The diagram MUD detector with prewhitening.

Other RD-MUD linear detectors

Various linear detectors have been developed for the MF-bank front-end output in the conventional MUD setting. In this section we explore some alternate linear detectors other than the decorrelating detector in the reduced-dimension setting. In particular, we will examine the reduced-dimension MMSE (RD-MMSE) and the reduced-dimension least-squares (RD-LS) detectors.

Reduced-dimension MMSE (RD-MMSE) detector:

Similar to the MMSE detector of the conventional MUD, a linear detector based

on the MMSE criterion can be derived for (4.21) based on the active user set estimate $\hat{\mathcal{I}}$ given by (4.23). In other words, we want to find a linear transform \mathbf{M} that minimizes $\mathbb{E}\{\|\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y}\|^2\}$, where the expectation is with respect to the vector of transmitted symbols $\mathbf{b}_{\hat{\mathcal{I}}}$ and the noise vector \mathbf{w} . Following the approach for deriving the conventional MMSE detector [96], we assume that $\mathbf{b}_{\hat{\mathcal{I}}}$ has a distribution that is uncorrelated with the noise \mathbf{w} and $\mathbb{E}\{\mathbf{b}_{\hat{\mathcal{I}}}\mathbf{b}_{\hat{\mathcal{I}}}^H\} = \mathbf{I}$. Adapting the techniques for deriving the conventional MMSE detector [96] and taking into account the noise covariance matrix $\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H$ of the reduced-dimension model (4.21), we obtain the linear transform for the reduced-dimension MMSE (RD-MMSE) detector as:

$$\mathbf{M} = \mathbf{R}_{\hat{\mathcal{I}}}\mathbf{A}_{\hat{\mathcal{I}}}^H(\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}^2\mathbf{A}_{\hat{\mathcal{I}}}^H + \sigma^2\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1}. \quad (4.32)$$

The derivations are given Appendix B.2. Accordingly, the RD-MMSE detector determines symbols as:

$$\hat{b}_n = \begin{cases} \text{sgn}([\mathbf{R}_{\hat{\mathcal{I}}}\mathbf{A}_{\hat{\mathcal{I}}}^H(\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}^2\mathbf{A}_{\hat{\mathcal{I}}}^H + \sigma^2\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1}\mathbf{y}]_n), & n \in \hat{\mathcal{I}}; \\ 0, & n \notin \hat{\mathcal{I}}. \end{cases} \quad (4.33)$$

In summary, the RD-MMSE detector determines active users first through the support recovery method of (4.23) and then uses (4.33) to detect symbols.

Reduced-dimension least squares (RD-LS) detector:

In the reduced-dimension model (4.21), based on the active user set estimate $\hat{\mathcal{I}}$ given by (4.23), the matrix $\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}$ introduces interference when we detect $\hat{\mathbf{b}}_{\hat{\mathcal{I}}}$. From the view of the system of linear equations, (4.21) is an over-determined system with more equations than unknowns since in general we require the number of branches to

be greater than the number of active users $M > K$. Hence we cannot directly invert the matrix $\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}$ to remove interference, mimicking the idea of the decorrelating detector of conventional MUD. Alternately, we can alleviate the effect of interference using the method of least-squares (LS) to find an estimate of $\hat{\mathbf{b}}_{\hat{\mathcal{I}}}$: $\hat{\mathbf{b}}_{\hat{\mathcal{I}}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}\mathbf{x}\|^2$, and detecting symbols as the signs of the estimated vector. The solution is given by $\hat{\mathbf{b}}_{\hat{\mathcal{I}}} = \mathbf{R}_{\hat{\mathcal{I}}}^{-1}(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}$. This corresponds to the maximum likelihood estimate of $\hat{\mathbf{b}}_{\hat{\mathcal{I}}}$ if (a) we ignore the covariance of the noise \mathbf{w} in (4.21) and assume it is white; (b) we know the active users, i.e. $\hat{\mathcal{I}} = \mathcal{I}$, and (c) the gains r_n are known. We call this the reduced-dimension least squares (RD-LS) detector. In summary, the RD-LS detector first detects active users by the support recovery method of (4.23). Since $\text{sgn}([\hat{\mathbf{b}}_{\hat{\mathcal{I}}}]_n) = \text{sgn}([\mathbf{R}_{\hat{\mathcal{I}}}^2 \hat{\mathbf{b}}_{\hat{\mathcal{I}}}]_n)$, the RD-LS detects symbols by:

$$\hat{b}_n = \begin{cases} \text{sgn} \left(r_n \Re \left[(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y} \right]_n \right), & n \in \hat{\mathcal{I}}; \\ 0, & n \notin \hat{\mathcal{I}}. \end{cases} \quad (4.34)$$

We can show that the RD-LS detector (4.34) and the RDD detector (4.24) give quite similar results in low noise and with low coherence of \mathbf{A} . To see this, write $\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}} = \mathbf{I} + \mathbf{E}$, where the symmetric matrix \mathbf{E} has zeros on the diagonal and the off-diagonals are bounded by the coherence μ of \mathbf{A} . As discussed in more detail in Section 4.2.3, for the RDD detector to work well, we choose \mathbf{A} with small μ . When $(K - 1)\mu < 1$, by Gershgorin's Theorem, we have that the spectral norm of the symmetric matrix \mathbf{E} is bounded by $\rho(\mathbf{E}) \leq (K - 1)\mu < 1$. Hence, using Lemma 4 in [29], we can write $(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} = \mathbf{I} + \sum_{n=1}^{\infty} (-\mathbf{E})^n$, and bound the spectral norm of the

series by a small number: $\rho(\sum_{n=1}^{\infty}(-\mathbf{E})^n) \leq (K-1)\mu/[1-(K-1)\mu]$. This means

$$\|(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y} - \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}\| = \left\| \left[\sum_{n=1}^{\infty} (-\mathbf{E})^n \right] \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y} \right\| \leq (K-1)\mu/[1-(K-1)\mu] \|\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}\|. \quad (4.35)$$

When the coherence μ of \mathbf{A} is sufficiently small relative to K and $\|\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}\|$, (4.35) says that the difference between the two vectors $(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}$ and $\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}$ is small. Then when μ is small and with sufficiently small noise, detecting symbols using $(\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{A}_{\hat{\mathcal{I}}})^{-1} \mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}$ in (4.34) is similar to detecting symbols using $\mathbf{A}_{\hat{\mathcal{I}}}^H \mathbf{y}$ in (4.24). As numerically shown in Section 4.4.2, the conditional probability of detecting wrong symbols given the correct support of active users, i.e. $P(\hat{\mathbf{b}} \neq \mathbf{b} | \hat{\mathcal{I}} = \mathcal{I})$, for (4.34) is similar to that for (4.24).

Maximum likelihood detector

The optimal detector that minimizes the probability-of-error for the RD-MUD output is the nonlinear maximum likelihood detector. The maximum likelihood detector finds the active users and symbols by minimizing the likelihood function, or, equivalently, minimizing the quadratic function $\|(\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1/2}(\mathbf{y} - \mathbf{A}\mathbf{R}\mathbf{b})\|^2$. This is equivalent to solving the following integer optimization problem

$$\max_{b_n \in \{-1, 0, 1\}} 2\mathbf{y}^H (\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1} \mathbf{A}\mathbf{R}\mathbf{b} - \mathbf{b}^H \mathbf{R}\mathbf{A}^H (\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)^{-1} \mathbf{A}\mathbf{R}\mathbf{b}, \quad (4.36)$$

where $b_n = 0$ corresponds to the n th user being inactive and this creates an augmented state space with one more state of possible transmitted symbols corresponding to a null symbol “0”. Hence (4.36) is more complex than the conventional maximum likelihood detector for MUD with BPSK modulation (4.10) since we add the “0”

symbol. Similar to the conventional maximum likelihood detector of the conventional MF-bank, the maximization in (4.36) is a combinatorial optimization problem, which can be solved by exhaustive search with complexity-per-bit exponential in the number of users.

4.2.3 Choice of \mathbf{A}

In Section 4.2.1 we have shown that the coefficient matrix \mathbf{A} is our design parameter. In Section 4.2.2 and Section 4.2.2 we have shown that both the RDD and RDDF detectors are based on the inner products between the projected received signal vector and the columns of \mathbf{A} , which correspond to the signature waveform vectors in the detection subspace. Hence, intuitively, for the RDD and RDDF detectors to work well, the inner products between columns of \mathbf{A} , or its coherence defined in (4.14), should be small, since each column of \mathbf{A} represents a signature waveform vector in the detection subspace. Several commonly used random matrices in compressed sensing that have small coherence with high probability are:

- (1) Gaussian random matrices: entries a_{nm} are independent and identically distributed (i.i.d.) with a zero mean and unit variance Gaussian distribution, with columns normalized to have unit norm;
- (2) Randomly sampled rows of a unitary matrix that satisfies $\mathbf{X}\mathbf{X}^H = \mathbf{X}^H\mathbf{X} = \mathbf{I}$. For instance, the random partial discrete Fourier transform (DFT) matrix, which is formed by randomly selecting rows of a DFT matrix \mathbf{F} : $[\mathbf{F}]_{nm} = e^{i\frac{2\pi}{N}nm}$ and normalizing the columns of the sub-matrix, where $i = \sqrt{-1}$.

We will focus on the random partial DFT matrix for the following reason. If we choose the number of correlators equal to the number of users, i.e. $M = N$, there

is no dimension reduction, and the performance of RD-MUD should equal that of the MF-bank. When $M = N$, the random partial DFT matrix becomes the DFT matrix with the property that $\mathbf{A}^H \mathbf{A} = \mathbf{I}$, i.e., $\mathbf{a}_n^H \mathbf{a}_m = \delta_{nm}$. Consequently, in this case, $\{\mathbf{a}_n^H \mathbf{y}\}$, which is a set of statistics that the RDD and RDDF detectors are based on, has the same distribution as the MF-bank output. To see this, write $\mathbf{a}_n^H \mathbf{y} = \mathbf{a}_n^H \left(\sum_{m=1}^N \mathbf{a}_m r_m b_m \right) + \mathbf{a}_n^H \mathbf{w} = r_n b_n + \mathbf{a}_n^H \mathbf{w}$, where $\mathbf{a}_n^H \mathbf{w}$ is a Gaussian random variable with zero mean and covariance $\sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_m = [\mathbf{G}^{-1}]_{nm}$. However, the Gaussian random matrix does not have this property: when $M = N$, $\mathbf{a}_n^H \mathbf{a}_m \neq 0$ for $n \neq m$, and so the performance of RD-MUD using the Gaussian random matrix \mathbf{A} is worse than that using the random partial DFT matrix. This has been validated in our numerical results in Section 4.4.1 where we will show that when M is relatively large, the Gaussian random matrix performs worse than the random partial DFT matrix.

4.3 Performance of RD-MUD

In the following, we study the performance of RD-MUD with the RDD and RDDF detectors. We begin by considering the scenario of a single active user without noise, and then move on to analyze the more general scenario with multiple active users and noise.

4.3.1 Single Active User

The following discussion shows that, when there is only one active user in the absence of noise, the RDD detector can detect the correct active user and symbol by using

only *two* correlators, if every two columns of \mathbf{A} are linearly independent. Later we will also show this is a corollary (Corollary 2) of a more general theorem, Theorem 1 below.

Assume there is no noise and only one user with index n_0 is active. In this case $y(t) = r_{n_0} b_{n_0} s_{n_0}(t)$, and by assumption we know only one user is active, i.e. we know $K = 1$. In the conventional MUD, the single-user detector based on the MF-bank detects the active user by finding $\hat{n}_0 = \arg \max_n |\langle y(t), s_n(t) \rangle|$ and the symbol by $\hat{b}_{\hat{n}_0} = \text{sgn}(r_{\hat{n}_0} \langle y(t), s_{\hat{n}_0}(t) \rangle)$. From the Cauchy-Schwarz inequality, for any n ,

$$|\langle y(t), s_n(t) \rangle| = |\langle r_{n_0} b_{n_0} s_{n_0}(t), s_n(t) \rangle| \leq |r_{n_0}| \|s_{n_0}(t)\| \|s_n(t)\| = |r_{n_0}|, \quad (4.37)$$

with equality if and only if $s_n(t) = c s_{n_0}(t)$ for some constant c , and thus $\hat{n}_0 = n_0$. The symbol can also be recovered perfectly, since

$$\hat{b}_{n_0} = \text{sgn}(r_{n_0} \langle y(t), s_{n_0}(t) \rangle) = \text{sgn}(r_{n_0}^2 b_{n_0} \langle s_{n_0}(t), s_{n_0}(t) \rangle) = b_{n_0}. \quad (4.38)$$

In RD-MUD, with two correlators, the RDD detector determines the active user by finding

$$\hat{n}_0 = \arg \max_{n=1, \dots, N} |a_{1n} \langle h_1(t), y(t) \rangle + a_{2n} \langle h_2(t), y(t) \rangle|. \quad (4.39)$$

From the Cauchy-Schwarz inequality,

$$|a_{1n} \langle h_1(t), y(t) \rangle + a_{2n} \langle h_2(t), y(t) \rangle|^2 \leq (a_{1n}^2 + a_{2n}^2) [\langle h_1(t), y(t) \rangle^2 + \langle h_2(t), y(t) \rangle^2], \quad (4.40)$$

with equality if and only if $a_{mn} = c \langle h_m(t), y(t) \rangle = c a_{mn_0} r_{n_0} b_{n_0} = c(n_0) a_{mn_0}$ for both $m = 1, 2$ with some constant $c(n_0)$. If every two columns of \mathbf{A} are linearly

independent, we cannot have two indices n such that $a_{mn} = c(n_0)a_{mn_0}$ for $m = 1, 2$. Also recall that the columns of \mathbf{A} are normalized, $a_{1n}^2 + a_{2n}^2 = \|\mathbf{a}_n\|^2 = 1$. Therefore, the maximum is achieved only for $n = n_0$ and $c(n_0) = 1$, which detects the correct active user. The detected symbol is also correct, since

$$\hat{b}_{n_0} = \text{sgn}(r_{n_0}[a_{1n_0}\langle y(t), h_1(t) \rangle + a_{2n_0}\langle y(t), h_2(t) \rangle]) = \text{sgn}(r_{n_0}^2 b_{n_0}[a_{1n_0}^2 + a_{2n_0}^2]) = b_{n_0}. \quad (4.41)$$

In the presence of noise, detectors in RD-MUD as well as those based on the conventional MF-bank will make detection errors. However, RD-MUD can have a performance similar to the detectors based on the MF-bank, as we now explain using geometric intuition and later prove formally in Section 4.3.3. Consider a scenario with three users having orthogonal signature waveforms, where only the first user is active. Suppose that $b_1 = 1$ and $r_1 = 1$, as illustrated in Fig. 4.7. By correlating with signature waveforms, the MF-bank (Figure 4.7) obtains inner products (4.5) of the received signal $y(t)$ with each of the signature waveforms $\{s_n(t)\}$ and detects based on these inner products. Because of noise, in Fig. 4.7 the received signal $y(t)$ does not coincide with $s_1(t)$. However, when the noise is sufficiently small, as shown in Fig. 4.7, the inner product of the received signal with the first signature waveform is the largest and positive, and hence in this scenario the single-user detector based on the conventional MF-bank detects the correct active user and its symbol. On the other hand, for the same setting, the RDD detector with two correlators (Fig. 4.8) projects the received signal onto the detection subspace via (4.18), and then obtains the decision statistics by computing the inner product between \mathbf{y} , the mapping of the received signal onto the subspace, and \mathbf{a}_n , the projection of each signature waveform onto the subspace. Because of noise, in Fig. 4.8 the projected signal vector \mathbf{y} does

not coincide with the first signature waveform vector \mathbf{a}_1 . However, when the noise is sufficiently small, as shown in Fig. 4.8, the inner product with the first signature vector is still the largest and positive, and hence in this scenario the RDD detector also detects the correct active user and symbol. From the above discussion we see that the RDD detector works well when the columns of \mathbf{A} are nearly orthogonal and the noise is sufficiently small. The former requirement is equivalent to requiring the coherence of \mathbf{A} to be as small as possible. The above discussion applies to the RDDF detector as well since its detection is also based on the inner products in the projection space.




Figure 4.7: A MUD problem with $N = 3$, $M = 2$, $K = 1$ and the received signal is due to the first user. Illustration of the detection by (a) the single-user detection of the MF-bank, and (b) the RD-MUD. The projection of the signature waveform onto the projection subspace results in signature waveform vectors $\{\hat{\mathbf{s}}_n\}$.

4.3.2 Noise Amplification of Subspace Projection

The RDD and RDDF detectors use the set of statistics $\{\mathbf{a}_n^H \mathbf{y}\}$ to detect active users and their symbols, which has noise components $\{\mathbf{a}_n^H \mathbf{w}\}$. We will show that the




Figure 4.8: A MUD problem with $N = 3$, $M = 2$, $K = 1$, and the received signal is due to the first user. Illustration of the detection by the RD-MUD. The mapping of the signature waveform is denoted by \hat{s}_n .

projection onto the detection subspace amplifies noise. To see this, first consider a special case with orthogonal signature waveforms, i.e. $\mathbf{G} = \mathbf{I}$, and hence the noise amplification is not caused by correlated signature waveforms. Assume the random partial DFT matrix is used as the coefficient matrix \mathbf{A} . Using (4.20), the noise variance of the detection statistic for the n th user is given by $\sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n = \sigma^2(N/M)$ by the definition of the random partial DFT matrix in Section 4.2.3. Hence in this special case the noise variance for each user is amplified by a factor $N/M \geq 1$ due to subspace projection.

In general, from (4.20) the detection statistic for the n th user has a noise variance $\sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n$. We will show this noise variance can be greater than σ^2 . First we

bound the output noise variance of the RD-MUD for each user as

$$\sigma^2 \lambda_{\min}(\mathbf{G}^{-1}) \|\mathbf{A}^H \mathbf{a}_n\|^2 \leq \sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n \leq \sigma^2 \lambda_{\max}(\mathbf{G}^{-1}) \|\mathbf{A}^H \mathbf{a}_n\|^2. \quad (4.42)$$

We now show that the upper bound in (4.42) is greater than σ^2 . The factor $\lambda_{\max}(\mathbf{G}^{-1})$ captures the largest possible noise amplification due to correlated signature waveforms. We have $\lambda_{\max}(\mathbf{G}^{-1}) \geq 1$ for the following. The sum of all the eigenvalues $\sum_{k=1}^N \lambda_k(\mathbf{G}) = \text{tr}(\mathbf{G}) = N$ since all the diagonal entries of \mathbf{G} are one. Since all eigenvalues of \mathbf{G} are nonnegative (a property of \mathbf{G} is that it is positive semi-definite), if $\lambda_{\min}(\mathbf{G}) > 1$, the sum of the eigenvalues will exceed N , and hence $\lambda_{\min}(\mathbf{G}) \leq 1$. Since the eigenvalues of \mathbf{G}^{-1} are the inverse of the eigenvalues of \mathbf{G} , $\lambda_{\max}(\mathbf{G}^{-1}) \geq 1$. On the other hand, the factor $\|\mathbf{A}^H \mathbf{a}_n\|^2$ captures the noise amplification due to subspace projection. It is lower-bounded by 1: $\|\mathbf{A}^H \mathbf{a}_n\|^2 = 1 + \sum_{l=1, l \neq n}^N (\mathbf{a}_l^H \mathbf{a}_n)^2 \geq 1 + \min_{l \neq n} |\mathbf{a}_l^H \mathbf{a}_n|^2$. When $M < N$, the inner product $\min_{l \neq n} |\mathbf{a}_l^H \mathbf{a}_n|$ is non-zero, and hence the factor $\|\mathbf{A}^H \mathbf{a}_n\|^2$ will be strictly greater than one. As a result, the upper bound on the noise variance in (4.42) is greater than σ^2 .

In the following section, we will capture this noise amplification more precisely by relating the noise variance of the decision statistic to the performance of the RD-MUD detectors.

4.3.3 Coherence Based Performance Guarantee

In this section, we present conditions under which the RDD and RDDF detectors can successfully recover active users and their symbols. The conditions depend on \mathbf{A} through its coherence and are parameterized by the crosscorrelations of the signature waveform through the properties of the matrix \mathbf{G} . Our performance measure is the

probability-of-error, which is defined as the chance of the event that the set of active users is detected incorrectly, *or* any of their symbols are detected incorrectly:

$$P_e = \mathbb{P}\{\hat{\mathcal{I}} \neq \mathcal{I}\} + \mathbb{P}\{\{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} \neq \mathbf{b}\}\}. \quad (4.43)$$

We will show in the later section that the second term of (4.43) is dominated by the first term when (4.23) and (4.27) are used for active user detection. The noise plays two roles in the P_e of (4.43). First, the noise can be sufficiently large relative to the weakest signal such that a nonactive user is determined as active; second, the noise can be sufficiently large such that the transmitted symbol plus noise is detected in an incorrect decision region and hence decoded in error.

The first error term in (4.43) is related to the probability-of-error for support recovery (see, e.g. [32] in noise, and the first steps of the greedy algorithms to recover support [9]). There are two major differences in our results on this aspect of RD-MUD performance relative to those previous works. First, although noise in the analog signal model (4.4) is white, matched filtering at the RD-MUD front-end introduces colored noise in (4.20). Second, we take into account the second term in (4.43), which has not been considered in previous work. We find the conditions such that the second term of (4.43) is dominated by the first term of (4.43).

Define the largest and smallest channel gains as

$$|r_{\max}| \triangleq \max_{n=1}^N |r_n|, \quad |r_{\min}| \triangleq \min_{n=1}^N |r_n|. \quad (4.44)$$

Our main result is the following theorem:

Theorem 1. *Let $\mathbf{b} \in \mathbb{R}^{N \times 1}$ be an unknown deterministic symbol, $b_n \in \{-1, 1\}$, $n \in \mathcal{I}$,*

and $b_n = 0$, $n \in \mathcal{I}^c$, $n = 1, \dots, N$. Assume that the number of active users K is known. Given the RD-MUD front-end output $\mathbf{y} = \mathbf{A}\mathbf{R}\mathbf{b} + \mathbf{w}$, where $\mathbf{A} \in \mathbb{C}^{M \times N}$ and $\mathbf{G} \in \mathbb{R}^{N \times N}$ are known, and \mathbf{w} is a Gaussian random vector with zero mean and covariance $\sigma^2 \mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H$, if every two columns of \mathbf{A} are linearly independent and the coherence of \mathbf{A} (4.14) satisfies the following condition:

$$|r_{\min}| - (2K - 1)\mu|r_{\max}| \geq 2\sigma\sqrt{2(1 + \alpha)\log N} \cdot \sqrt{\lambda_{\max}(\mathbf{G}^{-1})} \cdot \sqrt{\max_n (\mathbf{a}_n^H \mathbf{A}\mathbf{A}^H \mathbf{a}_n)}, \quad (4.45)$$

for some constant $\alpha > 0$, and $N^{-(1+\alpha)}[\pi(1 + \alpha)\log N]^{-1/2} \leq 1$, then the probability-of-error (4.43) for the RDD detector is upper bounded as:

$$P_e \leq N^{-\alpha}[\pi(1 + \alpha)\log N]^{-1/2}. \quad (4.46)$$

If every two columns of \mathbf{A} are linearly independent and the coherence of \mathbf{A} (4.14) satisfies a weaker condition:

$$|r_{\min}| - (2K - 1)\mu|r_{\min}| \geq 2\sigma\sqrt{2(1 + \alpha)\log N} \cdot \sqrt{\lambda_{\max}(\mathbf{G}^{-1})} \cdot \sqrt{\max_n (\mathbf{a}_n^H \mathbf{A}\mathbf{A}^H \mathbf{a}_n)}, \quad (4.47)$$

for some constant $\alpha > 0$, and $N^{-(1+\alpha)}[\pi(1 + \alpha)\log N]^{-1/2} \leq 1$, then the probability-of-error (4.43) for the RDDF detector is upper bounded by the right hand side of (4.46).

Proof. See Appendix B.3. □

Note in Theorem 1 that the condition of having a small probability-of-error for the RDDF detector is weaker than for the RDD detector. Intuitively, the iterative

approach of decision feedback removes the effect of the largest element in $\mathbf{R}\mathbf{b}$ iteratively, which helps the detection of weaker users. This is an extension of the ideas in SIC for standard MUD, except that now the presence as well as the data of the strongest user is detected, then removed in the detection subspace, which makes it easier to detect the presence and data of the next strongest user, after which the process repeats.

The main idea of the proof is the following. Consider $1 - P_e = P\{\{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}\}$. First we define an event $\mathcal{G} = \{\max_n |\mathbf{a}_n^H \mathbf{w}| < \tau\}$ for a quantity τ proportional to the right hand side in (4.45), and prove that \mathcal{G} occurs with high probability. This bounds the probability that the noise projected onto the detection subspace exceeds τ , i.e. it bounds the tail probability of the projected noise. Then we show that under the condition (4.45), whenever \mathcal{G} occurs, the active users can be correctly detected, which means $\mathcal{G} \subset \{\hat{\mathcal{I}} = \mathcal{I}\}$. On the other hand, we show that under a condition weaker than (4.45), whenever \mathcal{G} occurs, the user data symbols can be correctly detected, which means $\mathcal{G} \subset \{\hat{b}_n = b_n, n \in \mathcal{I}\}$. In other words, condition (4.45) ensures that whenever \mathcal{G} occurs, both the correct set of active users are detected and that their data are correctly decoded. This means that under condition (4.45) for the RDD detector, $\mathcal{G} \subset \{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}$, and thus $P(\mathcal{G}) \leq P\{\{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}\}$, which concludes the proof. A similar but inductive approach is used to prove the performance guarantee for the RDDF detector.

A special case for Theorem 1 is when $\mathbf{A}\mathbf{A}^H = (N/M)\mathbf{I}$, $\max_n (\mathbf{a}_n^H \mathbf{A}\mathbf{A}^H \mathbf{a}_n) = N/M$ and $\mathbf{G} = \mathbf{I}$, $\lambda_{\max}(\mathbf{G}^{-1}) = 1$. This is true when \mathbf{A} is the random partial DFT matrix and the signature waveforms are orthogonal, and hence the noise in (4.20) is white. If we scale σ^2 by M/N , the right hand sides of (4.45) and (4.47) are then identical

to the corresponding quantities in Theorem 4 of [9]. Hence, for the random partial DFT matrix for \mathbf{A} , Theorem 1 has the same conditions as those of Theorem 4 in [9]. However, Theorem 4 in [9] guarantees detecting the correct sparsity pattern of \mathbf{b} (equivalently, the correct active users), whereas Theorem 1 guarantees correct detection of not only the active users but their symbols as well. That is because, as mentioned above, correct detection of these transmitted symbols comes “for free” when the conditions to correctly detect the active users are met.

Remarks:

The term $\max_n(\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n)$ on the right hand side of (4.45) and (4.47) is bounded by

$$1 \leq \max_n(\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n) \leq 1 + (N - 1)\mu^2. \quad (4.48)$$

Equation (4.48) follows because $\max_n(\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n) = \max_n \sum_{l=1}^N (\mathbf{a}_n^H \mathbf{a}_l)^2$, and

$$1 = (\mathbf{a}_n^H \mathbf{a}_n)^2 \leq \max_n \sum_{l=1}^N (\mathbf{a}_n^H \mathbf{a}_l)^2 = 1 + \max_n \sum_{l \neq n} (\mathbf{a}_n^H \mathbf{a}_l)^2 \leq 1 + (N - 1)\mu^2. \quad (4.49)$$

On the other hand, there is a noise phase-transition effect, in the following sense. Conditions (4.45) and (4.47) suggest that for the RDD and RDDF detectors to have P_e as small as (4.46), we need to have

$$|r_{\min}|^2 / \sigma^2 > 8 \log N \lambda_{\max}(\mathbf{G}^{-1}), \quad (4.50)$$

because $\alpha > 0$ and (4.48) holds. If the minimum SNR, i.e. the SNR associated with the minimum gain r_{\min} is not sufficiently high, these algorithms cannot attain small probability-of-error. We illustrate this effect via numerical examples in Section

4.4.2 (a similar effect can be observed in standard MUD detectors using linear or DF detection).

4.3.4 Bounding Probability-of-Error of RDD and RDDF

Theorem 1 provides a condition on how small μ has to be to achieve a small probability-of-error. The condition and the achievable small probability-of-error are related by the constant α . We can eliminate this constant and write Theorem 1 in an equivalent form that gives error bounds for the RDD and RDDF detectors explicitly. Define the minimum signal-to-noise ratio (SNR) in the projection subspace as

$$\text{SNR}_{\min} = \frac{|r_{\min}|^2}{\sigma^2 \lambda_{\max}(\mathbf{G}^{-1})}, \quad (4.51)$$

where the factor $\lambda_{\max}(\mathbf{G}^{-1})$ captures the noise amplification effect in the projection due to nonorthogonal signature waveforms. Also define two factors β_1 and β_2 as

$$\beta_1 \triangleq \frac{[1 - (2K - 1)\mu|r_{\max}|/|r_{\min}|]^2}{\max_n(\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n)}, \quad \beta_2 \triangleq \frac{[1 - (2K - 1)\mu]^2}{\max_n(\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n)}. \quad (4.52)$$

For the RDD detector, we have already implicitly assumed that $1 - (2K - 1)\mu|r_{\max}|/|r_{\min}| \geq 0$, since the right hand side of (4.45) in Theorem 1 is non-negative. For the same reason, for the RDDF detector, we have assumed that $1 - (2K - 1)\mu > 0$. By (4.48) and (4.52), $\beta_1 \leq 1$ and $\beta_2 \leq 1$. We have the following corollary from Theorem 1:

Corollary 1. *Under the setting of Theorem 1, with the definitions (4.51) and (4.52), the probability-of-error for the RDD detector is upper-bounded by*

$$P_{e,\text{RDD}} \leq \frac{2N}{\sqrt{\pi}} \left[\frac{\text{SNR}_{\min}}{2} \cdot \beta_1 \right]^{-1/2} e^{-\frac{1}{4} \frac{\text{SNR}_{\min}}{2} \cdot \beta_1}, \quad (4.53)$$

with $1 - (2K - 1)\mu|r_{\max}|/|r_{\min}| \geq 0$, and the probability-of-error for the RDDF detector is upper bounded by

$$P_{e,\text{RDDF}} \leq \frac{2N}{\sqrt{\pi}} \left[\frac{\text{SNR}_{\min}}{2} \cdot \beta_2 \right]^{-1/2} \cdot e^{-\frac{1}{4} \frac{\text{SNR}_{\min}}{2} \beta_2}, \quad (4.54)$$

with $1 - (2K - 1)\mu > 0$.

Proof. We begin by bounding the probability-of-error of the RDD detector from Theorem 1. Under condition (4.45), by (4.46) the probability-of-error of the RDD detector is bounded by $N^{-\alpha}[\pi \log(N^{1+\alpha})]^{-1/2}$ for some constant $\alpha > 0$. To make the bound tight, we choose α as large as possible such that it still satisfies (4.45). With definitions (4.51) and (4.52), for the RDD detector, we can rewrite (4.45) as

$$[\log N^{(1+\alpha)}]^{1/2} \leq \frac{1}{2} \cdot \left[\frac{\text{SNR}_{\min}}{2} \cdot \beta_1 \right]^{1/2}, \quad (4.55)$$

or equivalently

$$N^\alpha \leq N^{-1} e^{\frac{1}{4} \frac{\text{SNR}_{\min}}{2} \cdot \beta_1}. \quad (4.56)$$

The right hand sides of (4.55) and (4.56) are the largest values for $[\log N^{(1+\alpha)}]^{1/2}$ and N^α we can obtain under (4.45), for given SNR_{\min} , \mathbf{A} , K and N . Combining (4.55) and (4.56) in (4.46), we have (4.53). Similarly, by choosing the largest possible α satisfying (4.47), we derive the bound (4.54) on the probability-of-error of the RDDF detector. \square

Remarks:

For bounds (4.53) and (4.54) to be meaningful, they have to be less than one, and hence SNR_{\min} should be on the order of $\log N$. Also note that the error bounds

(4.53) and (4.54) for the RD-MUD detectors are larger when the signature waveforms are correlated, since these error bounds increase in SNR_{\min} , and SNR_{\min} decreases in $\lambda_{\max}(\mathbf{G}^{-1})$. This implies that the performance of the RD-MUD detectors tends to degrade when signature waveforms are nonorthogonal (the same effect can also be observed for the conventional MUD detectors based on the MF-bank).

By letting the noise variance σ^2 go to zero in (4.53) and (4.54) for the RDD and RDDF detectors, we can derive the following corollary from Theorem 1 (another proof for the RDD detector in this case has been given in Section 4.3.1).

Corollary 2. *Under the setting of Theorem 1, in the absence of noise, the RDD detector can correctly detect the active users and their symbols if $\mu < |r_{\min}|/[|r_{\max}|(2K-1)]$, and the RDDF detector can correctly detect the active users and their symbols if $\mu < 1/(2K-1)$. In particular, if $K = 1$, with $M = 2$ correlators, $P_e = 0$ for the RDDF detector, and if furthermore $|r_{\max}| = |r_{\min}|$, $P_e = 0$ for the RDD detector (which has also been shown in Section 4.3.1).*

Proof. In Theorem 1, if we let σ^2 go to zero, then SNR_{\min} goes to infinity, and the right hand sides of both (4.53) and (4.54) go to zero, i.e. $P_e = 0$, as long as $\beta_1 > 0$ and $\beta_2 > 0$, or equivalently, $1 - (2K-1)|r_{\max}|/|r_{\min}| > 0$ and $1 - (2K-1)\mu > 0$. When $K = 1$, the bound on μ for the RDDF detector becomes $1/(2K-1) = 1$, which is satisfied for any μ as long as $M \geq 2$ (since Theorem 1 also requires linear independence of the columns of \mathbf{A} and this rules out the possibility of $M = 1$). \square

4.3.5 Comparison with Existing Bounds

In this section we compare the bound on the probability-of-error in the literature for the decorrelating detector of the conventional MUD with our bound for the RDD

and RDDF detectors derived from Theorem 1. The decorrelating detector is the counterpart of the RDD detector in the conventional MUD setting. To see this, note that for the RD-MUD front-end, when $M = N$, we can choose the coefficient matrix $\mathbf{A} = \mathbf{I}$ such that the output data model (4.20) is equivalent to the MF-bank decorrelating detector (4.7) (with $\mathbf{T} = \mathbf{G}^{-1}$).

For the decorrelating detector of the conventional MUD, a commonly used performance measure is the probability of error of each user [53][54], which is given by [53] [96]:

$$\mathbb{P}\{\hat{b}_n \neq b_n\} = Q\left(\frac{|r_n|}{\sigma\sqrt{[\mathbf{G}^{-1}]_{nn}}}\right), \quad (4.57)$$

where $Q(x) = \int_x^\infty (1/\sqrt{2\pi})e^{-z^2/2}dz$ is the Gaussian tail probability. To compare (4.57) with the P_e bound defined by (4.43), which consists of both active user detection error and symbol error, we consider the case when all users are active, i.e. $K = N$, and then P_e is only due to symbol error. In this setting, we have, using the union bound and (4.57):

$$P_e = \mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b}\} \leq \sum_{n=1}^N \mathbb{P}\{\hat{b}_n \neq b_n\} \leq NQ\left(\sqrt{\text{SNR}_{\min}}\right) \leq \frac{N}{2\sqrt{\pi}} \left[\frac{\text{SNR}_{\min}}{2}\right]^{-1/2} e^{-\frac{\text{SNR}_{\min}}{2}}, \quad (4.58)$$

where we have also used the fact that $|r_n|/\left[\sigma\sqrt{[\mathbf{G}^{-1}]_{nn}}\right] \geq \sqrt{\text{SNR}_{\min}}$ and $Q(x)$ is decreasing in x , as well as the bound on $Q(x)$ [96] on $Q(x)$ given by

$$Q(x) \leq \frac{1}{x\sqrt{2\pi}}e^{-x^2/2}. \quad (4.59)$$

The bounds on P_e of the RDD and RDDF detectors when $1 \leq K \leq N$ are given in (4.53) and (4.54), respectively. Since $\beta_1 \leq 1$ and $\beta_2 \leq 1$, the error bounds (4.53) for

the RDD detector and (4.54) for the RDDF detector are larger than the bound (4.58) for conventional MUD. This is because the RDD and RDDF detectors have one extra source of error from detecting the wrong set of active users, and also because the noise can be amplified by the projection onto the detection subspace, as discussed in Section 4.3.2. The enlargement of the error bound due to subspace projection is captured by factors β_1 and β_2 for the RDD and RDDF detectors, respectively. These factors reduce the effect of SNR_{\min} in the bounds. Note that β_1 and β_2 increase in μ , and hence we want small μ , which leads to a small error bound for RD-MUD detectors.

A special case is when $K = N$, $\mathbf{A} = \mathbf{I}$. Then $\mu = 0$, $\max_n(\mathbf{a}_n \mathbf{A} \mathbf{A}^H \mathbf{a}_n) = 1$, and by definition (4.52), $\beta_1 = \beta_2 = 1$. At the beginning of this section, we have shown that this corresponds to the decorrelating detector of the conventional MUD. In this case the bounds (4.54) and (4.58) become the same expression:

$$P_e \leq \frac{2N}{\sqrt{\pi}} \left[\frac{\text{SNR}_{\min}}{2} \right]^{-1/2} e^{-\frac{1}{4} \frac{\text{SNR}_{\min}}{2}}. \quad (4.60)$$

Compared with the bound (4.58) for P_e of the conventional decorrelating detector, the bound (4.60) obtained from our result is larger. This can be explained since (4.60) is obtained as a special case of RD-MUD which must also detect active users. As we have shown in the proof for Theorem 1, the error from detecting active users dominates the error from detecting symbols.

4.3.6 Lower Bound on Number of Correlators

Theorem 1 is stated for any matrix \mathbf{A} . If we substitute the expression for coherence of a given \mathbf{A} in terms of its dimensions M and N into Theorem 1, we can obtain

a lower bound on the smallest number of correlators M needed to achieve a certain probability-of-error.

There is a general lower bound on the coherence of any $M \times N$ matrix \mathbf{A} given by [33]:

$$\mu \geq \left[\frac{N-M}{M(N-1)} \right]^{1/2} = M^{-1/2} \left(\frac{N}{N-1} \right)^{1/2} \cdot (1 - M/N)^{1/2} \sim M^{-1/2}, \quad (4.61)$$

when N is large relative to M and N is much larger than 1. In the absence of noise, the upper bound on the coherence in Corollary 2 together with the bound (4.61) imply that, for the RDDF detector to have perfect detection, the number of correlators M should be on the order of $(2K-1)^2$. In the compressed sensing literature, it is known that the bounds obtained using the coherence of the matrix \mathbf{A} may not be as sharp as those obtained using the restricted isometry properties of \mathbf{A} [33]. For example, in compressed sensing, to estimate a sparse vector with κ non-zero entries, the lower bound on the number of measurements required based on the coherence of \mathbf{A} is proportional to κ^2 , while that based on the restricted isometry properties is proportional to κ . This effect is referred to as the “quadratic bottleneck” in the compressed sensing literature [33]. Nevertheless, the coherence properties are easy to evaluate, while evaluating the restricted isometry property of a given matrix \mathbf{A} is in general NP-hard [9]. Also as we demonstrate in the proof of Theorem 1, the coherence is a convenient measure of the user interference level in the detection subspace. For this reason, our result is based on the coherence of matrix \mathbf{A} .

In the compressed sensing literature, the matrix \mathbf{A} is often chosen to be random, in which case its coherence can be bounded in probability. Consider for example a random partial DFT matrix. We have the following result (which can be proven easily

by the complex Hoeffding's inequality [45]):

Lemma 6. *Let $\mathbf{A} \in \mathbb{C}^{M \times N}$ be a random partial DFT matrix. Then the coherence of \mathbf{A} is bounded by*

$$\mu < [4(2 \log N + c)/M]^{1/2}, \quad (4.62)$$

with probability exceeding $1 - 2e^{-c}$, for some constant $c > 0$.

Using Lemma 6, we have the following corollary to Theorem 1:

Corollary 3. *Consider the setting of Theorem 1, where \mathbf{A} is a random partial DFT matrix. Suppose the number of correlators satisfies the following lower bound for the RDD detector*

$$M \geq 4 \left[\frac{(2K - 1)|r_{\max}|}{|r_{\min}| - 2\tau} \right]^2 (2 \log N + c), \quad (4.63)$$

or satisfies the following smaller lower bound for the RDDF detector

$$M \geq 4 \left[\frac{(2K - 1)|r_{\min}|}{|r_{\min}| - 2\tau} \right]^2 (2 \log N + c), \quad (4.64)$$

for some constants $c > 0$ and $\alpha > 0$, and $|r_{\min}| > 2\tau$, for τ defined in (B.8). Then the probability-of-error P_e of the RDD detector or the RDDF detector is bounded by

$$1 - (1 - N^{-\alpha}[\pi(1 + \alpha) \log N]^{-1/2})(1 - 2e^{-c}), \quad (4.65)$$

for some constant $\alpha > 0$.

This corollary says that to attain a small probability-of-error, the number of correlators needed by the RDD and RDDF detectors is on the order of $\log N$, which is much smaller than that required by the conventional MUD using a MF-bank, which is on the order of N .

4.4 Numerical Examples

As an illustration of the performance of RD-MUD, we present some numerical examples mainly of both the RDD and RDDF detectors. The results are obtained from 10^5 Monte Carlo trials. For each trial, we generate a Gaussian random noise vector as well as a random partial DFT matrix for \mathbf{A} , and form the signal vector according to (4.20). To simplify, we assume that the gains for all the users are the same: $|r_{\min}| = |r_{\max}| \triangleq r = 1$. First we consider noise-free scenarios with an increasing number of users N for a fixed number of active users K , and then with increasing K for a fixed N . Next we consider two noisy scenarios with orthogonal waveforms $\mathbf{G} = \mathbf{I}$ and nonorthogonal waveforms $\mathbf{G} \neq \mathbf{I}$.

4.4.1 Noise-Free Scenario

In the absence of noise, from (4.7) of the MF-bank, the conventional decorrelating detector has output $\mathbf{z} = \mathbf{R}\mathbf{b}$. The conventional decorrelating detector determines the active users by choosing the K largest of $|z_n|$, which is equivalent to choosing the K largest of $\{|r_n b_n|\}$ in the absence of noise. Recall that the inactive users have $r_n = 0$, which means the conventional decorrelating detector can correctly detect the active users. It then detects symbols by $\hat{b}_n = \text{sgn}(r_n^2 b_n) = b_n$. For the above reasons, the conventional decorrelating detector has $P_e = 0$ in the absence of noise.

P_e vs. M , as N increases

Fig. 4.9 shows the P_e of the RDD detector as a function of M , for fixed $K = 2$, and different values of N . The data points marked on the curves correspond to $M = \log N$, $2 \log N$, and $4 \log N$. When $M = 8 \log N$, $P_e = 0$ for the RDD detector for all values

of N . This example clearly demonstrates the $\log N$ scaling factor in Corollary 3.

P_e vs. M , as N increases

Fig. 4.9 shows the P_e of the RDD detector as a function of M , for fixed $K = 2$, and different values of N . The data points marked on the curves correspond to $M = \log N$, $2 \log N$, and $4 \log N$. When $M = 8 \log N$, $P_e = 0$ for the RDD detector for all values of N . This example clearly demonstrates the $\log N$ scaling factor for the required number of correlators in Corollary 3.




Figure 4.9: Performance of Algorithm-I, in the absence of noise, for $K = 2$, different N , P_e versus M , with marked points equal to: $\log N$, $2 \log N$, $4 \log N$, and $8 \log N$. When $M = 8 \log N$, $P_e = 0$.

P_e vs M , as K increases

Fig. 4.10 demonstrates the P_e of the RDD detector as a function of M , for a fixed $N = 100$, and different values of K . The points marked on the curves correspond to $M = (K \log N)/2$, $K \log N$, $1.5K \log N$, $2K \log N$ and $\min\{N, 3K \log N\}$. Clearly, the number of correlators needed to obtain $P_e < 10^{-4}$ increases as K increases. When $K = 10$, the RDD detector needs about 80 correlators to obtain $P_e < 10^{-4}$. The number of correlators needed to achieve a small probability-of-error can be improved by using the RDDF detector. As shown in Fig. 4.11, when $K = 10$, the RDDF detector uses only 60 correlators to obtain $P_e < 10^{-4}$.

Fig. 4.10 also demonstrates how Corollary 2 can be used to estimate the number of correlators needed to achieve a small probability-of-error. Corollary 2 says that we need $\mu < 1/(2K - 1)$ in the absence of noise to have perfect detection. When $K = 2$, this requires $\mu < 1/3$. We then obtain an estimate for coherence of the random partial DFT matrix with $N = 100$ and various M , by averaging over 10^5 trials, and find that when M is about 30 the coherence is less than $1/3$. This is consistent with Fig. 4.10, which shows that when $M = 28$, P_e is on the order of 10^{-4} .

Random partial DFT vs. Gaussian random matrices

We compare the performance of the RDD detector using the random partial DFT matrix versus using the Gaussian random matrix for \mathbf{A} (defined in Section 4.2.3). In Fig. 4.12, the probability-of-error of the Gaussian random matrix converges to a value much higher than zero, whereas that of the random partial DFT matrix converges to zero (the value achieved by the conventional decorrelating detector), when M increases to N .




Figure 4.10: In the absence of noise, P_e versus M , for $N = 100$, using Algorithm-I for different K .

4.4.2 Noisy Scenario

Next, we consider noisy scenarios. For comparison, we also consider the conventional decorrelating detector, which corresponds to the RDD detector with $M = N$ as we explained in Section 4.3.5.

P_e vs. M , as SNR increases, $\mathbf{G} = \mathbf{I}$

We study P_e versus M for the RDD detector as SNR increases when the signature waveforms are orthogonal and $\mathbf{G} = \mathbf{I}$ and hence the noise in (4.20) is white. In this case $\text{SNR}_{\min} = |r_{\min}|^2/\sigma^2 = r^2/\sigma^2$, which is denoted as SNR in Fig. 4.13. Assume $N = 100$ and $K = 2$. In Fig. 4.13, when SNR increases, the curves converge to the noise-free curve for $K = 2$ shown in Fig. 4.10, and to the noise-free curve for

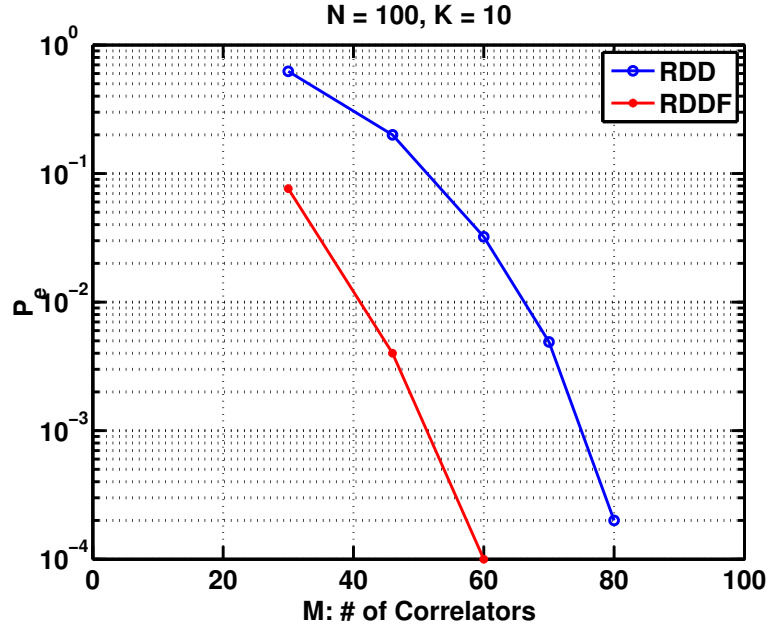


Figure 4.11: In the absence of noise, P_e versus M , for $N = 100$, using Algorithm-II compared with Algorithm-I when $K = 10$.

$N = 100$ shown in Fig. 4.9. Note that there is a noise phase-transition effect in Fig. 4.10, which is discussed in the Remarks of Section 4.3.3. The analysis in (4.50) implies that for $N = 100$ and $\mathbf{G} = \mathbf{I}$, we need SNR to be at least 15.7dB to obtain a small P_e , which is consistent with Fig. 4.10.

P_e vs. M , performance of noise prewhitening transform, $\mathbf{G} \neq \mathbf{I}$

Next we consider a scenario when the signature waveforms are nonorthogonal $\mathbf{G} \neq \mathbf{I}$. We generate an arbitrary symmetric \mathbf{G} with ones on the diagonal and fix it in the Monte Carlo trials. In the first case we consider highly correlated signature waveforms with $\lambda_{\max}(\mathbf{G}^{-1}) = 493.9595$. In the second case we consider nearly orthogonal signature waveforms with $\lambda_{\max}(\mathbf{G}^{-1}) = 4.0771$. Then we compare the P_e of the RDD




Figure 4.12: In the absence of noise, P_e versus M of the RDD detector using the random partial DFT versus using the Gaussian random matrices for $N = 100$ and $K = 6$.

detector in these two cases without and with the noise whitening transform in Section 4.2.2. Assume $N = 100$, $K = 2$, $r = 1$ and $\sigma = 0.1$. In Fig. 4.14, when the signature waveforms are highly correlated, the noise whitening transform significantly reduces P_e for large M . In this case, the conventional decorrelating detector without the noise whitening transform has a non-negligible probability-of-error, and that with the noise whitening transform has a probability-of-error less than 10^{-4} . In Fig. 4.15, when the signature waveforms are nearly orthogonal, the noise whitening transform does not reduce P_e much. In this case, the conventional decorrelating detector without and with the noise whitening transform both have probability-of-error less than 10^{-4} . We also verified that using the noise whitening transform cannot achieve the probability-of-error that is obtained with orthogonal signature waveforms $\mathbf{G} = \mathbf{I}$. This is because




Figure 4.13: Performance of Algorithm-I, P_e versus M for different SNRs, when the signature waveforms are orthogonal, i.e., $\mathbf{G} = \mathbf{I}$. When SNR is greater than 15 dB, the probability-of-error of the MF-bank is less than 10^{-4} .

the noise whitening transform distorts the signal component.

P_e vs. M , RD-MUD linear detectors

In this example, we compare the performance of RD-MUD linear detectors when $\mathbf{G} \neq \mathbf{I}$. In Theorem 1 we have proven that the error is dominated by that from active user detection. So we compare the performance of these RD-MUD linear detectors using their conditional probability of symbol error given the correct detection of active users $\mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b} | \hat{\mathcal{I}} = \mathcal{I}\}$. Assume $N = 100$, $K = 2$, $r = 1$, $\sigma = 0.1$, and let \mathbf{G} take the same form as those used in the previous example. In Fig. 4.16 and Fig. 4.17, the performance of the RDD detector is similar to that of the RD-LS detector (the explanation is given in Section 4.2.2). In comparison, the RD-MMSE detector has

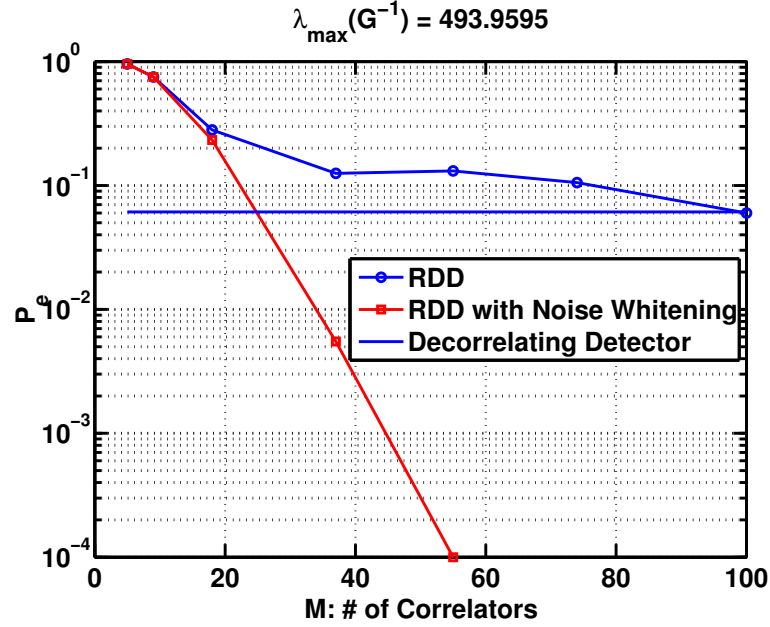


Figure 4.14: Comparison of P_e versus M , without and with the noise whitening transform, when $N = 100$, $K = 2$, and for two different cases of \mathbf{G} . The conventional decorrelating detector with the noise whitening transform in Fig. 4.14, and the conventional decorrelating detectors with and without the noise whitening transform in Fig. 4.15 have probability-of-error less than 10^{-4} .

smaller conditional probability of error, especially in Fig. 4.16 with highly correlated signatures. This improvement is because the linear transform (4.33) of the RD-MMSE detector alleviates the effect of correlated signature waveforms by including an inversion of \mathbf{G} in the linear transform. The conditional probability-of-error of the conventional decorrelating detector is less than 10^{-4} .

4.5 Conclusions

We have developed a reduced dimension multiuser detection (RD-MUD) structure, which decreases the number of correlators at the front-end of a MUD receiver by




Figure 4.15: Comparison of P_e versus M , with and without using pre-whitening, when $N = 100$, $K = 2$, when $\lambda_{\max}(\mathbf{G}^{-1})$ is small. The MF-bank decorrelators with and without whitening have probability-of-error less than 10^{-4} and hence are not shown in the pictures.

exploiting the fact that the number of active users is typically much smaller than the total number of users in the system. Motivated by the idea of analog compressed sensing, the RD-MUD front-end projects the received signal onto a lower dimensional detection subspace by correlating the received signal with a set of correlating signals. The correlating signals are constructed as linear combinations of the signature waveforms using a coefficient matrix \mathbf{A} , which determines the performance of RD-MUD and is our key design parameter. Based on the front-end output, RD-MUD detectors recover active users and their symbols in the detection subspace. We have studied in detail two such detectors. The reduced-dimension decorrelating (RDD) detector, which is a linear detector that combines subspace projection along with thresholding for active user detection and sign detection for data recovery. The reduced-dimension




Figure 4.16: Comparison of the conditional probability of error $\mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b} | \hat{\mathcal{I}} = \mathcal{I}\}$, for RD-MUD linear detectors, when $\lambda_{\max}(\mathbf{G}^{-1})$ is large. The conditional probability of error for MF-bank decorrelator is less than 10^{-4} and hence is not shown in the pictures.

decision feedback (RDDF) detector is a nonlinear detector that combines decision-feedback orthogonal matching pursuit (DF-OMP) for active user detection with sign detection for data recovery. We have shown that to achieve a desired probability-of-error, the number of correlators used by the RD-MUD can be much smaller than that used by the conventional MUD, and the complexity-per-bit of the RD-MUD detectors is not higher than their counterpart in the conventional MUD setting. In particular, when the random partial DFT matrix is used for the coefficient matrix \mathbf{A} and the RDD and RDDF detectors are used for detection, the RD-MUD front-end requires the number of correlators proportional to the log of the number of users, whereas the conventional MF-bank front-end requires the number of correlators equal to the number of users in the system. We have obtained theoretical performance guarantees




Figure 4.17: Comparison of the conditional probability of error $\mathbb{P}\{\hat{\mathbf{b}} \neq \mathbf{b} | \hat{\mathcal{I}} = \mathcal{I}\}$, for RD-MUD linear detectors: RDD, RD-LS, and RD-MMSE. The RD-MMSE detector has the smallest conditional probability of symbol error.

for the RDD and RDDF detectors in terms of the coherence of \mathbf{A} , which are validated via numerical examples.

Chapter 5

Summary and Future Work

5.1 Summary

We have shown that sparsity can be exploited in statistical detection problems with two specific cases: the multi-sensor change-point detection problem and the multi-user detection problem.

In multi-sensor change-point detection, we assume that there are multiple sensors that make sequences of observations in parallel. An abrupt emergence of a change-point alters the distribution of the observations for a subset of sensors simultaneously. Our goal is to detect the change-point as soon as possible after it occurs, and to minimize the rate of making false alarms. The sparsity in this problem takes the form that the fraction of sensors affected by the change-point, p , is relatively small, which can be viewed as the *sparsity of affected sensors*. We model this sparsity by assuming that each sensor is affected with a small probability p_0 , where p_0 is a guess for p . We then derived a mixture log generalized likelihood ratio (GLR) statistic for change-point detection. Our model leads to a nonlinear weighting function, which weights

the log GLR statistic of each sensor before combining them. The nonlinear weighting function automatically emphasizes the statistic formed by sensors that are affected by the change-point, and suppresses those formed by sensors that are not affected. We derived approximations to two performance metrics of the mixture procedure, the false alarm rate - the average run length (ARL), as well as the expected detection delay. We also demonstrated that the performance of the mixture procedure compared with other existing procedures using numerical examples. In summary, the mixture procedure uses a parameter p_0 to model and exploit the sparsity in multi-sensor change-point detection problems.

In multi-user detection (MUD), multiple users transmit information messages using their signature waveforms with a common receiver. The receiver receives a superposition of the transmitted signals and has to decode messages simultaneously. In this problem, the sparsity takes the form of *user sparsity*: the number of active users is small relative to the total number of users in the system. We exploited the user sparsity to reduce complexity of the detector at the receiver, in particular, to reduce the number of correlators at the front-end of the receiver. We proposed a lower complexity detection method, reduced dimension multi-user detection (RD-MUD), that uses a much fewer number of correlators than the conventional method and still achieves a similar performance. In particular, we showed that by proper choice of the design parameter of RD-MUD, which is the coefficient matrix for the correlating signals, we can achieve a similar performance to that of the classic detection method based on the matched-filter (MF) bank. We showed that when the random partial discrete Fourier transform (DFT) matrix is used as the coefficient matrix, the number of correlators needed by RD-MUD to obtain a small probability-of-error is on the

order of logarithm of the number of users.

5.2 Future Work

There are many extensions to the work we presented. For the multi-sensor change-point detection problem, we can consider a closely related procedures. We have shown in Section 3.2 that the mixture procedure is equivalent to applying a non-linear weighting function on the log GLR statistic of each sensor. As an alternative, we can consider the following non-linear weighting function:

$$\tilde{g}(U_{n,\kappa,t}; p_0) \triangleq [\log p_0 + (U_{n,\kappa,t}^+)^2/2]^+. \quad (5.1)$$

The choice of the function (5.1) is motivated by the following argument. We can write the mixture statistic (3.14) as

$$g(U_{n,\kappa,t}; p_0) = \log(1 - p_0 + p_0 \exp[(U_{n,\kappa,t}^+)^2/2]) = \log(1 - p_0 + \exp[\log p_0 + (U_{n,\kappa,t}^+)^2/2]). \quad (5.2)$$

When $\exp[\log p_0 + (U_{n,\kappa,t}^+)^2/2] \gg 1 - p_0$, which is true when $\log p_0 + (U_{n,\kappa,t}^+)^2/2 \geq 0$ and is relatively large, we can approximate (5.2) as $\log p_0 + (U_{n,\kappa,t}^+)^2/2$; otherwise the value of (5.2) is very close to zero. Hence, we can view $\tilde{g}(U_{n,\kappa,t}; p_0)$ in (5.1) as an approximation for $g(U_{n,\kappa,t}; p_0)$. The non-linear weighting function $\tilde{g}(x; p_0)$ is numerically more stable than $g(x; p_0)$, since it avoids calculating an exponential function with large argument as required by $g(x; p_0)$, although it can have a slightly longer detection delay. Also, $\tilde{g}(x; p_0)$ is of interest, because all the quantities starting with $\psi(\theta)$ for the ARL approximation, Approximation 1, can be evaluated analytically in

closed form without numerical integrations.

For reduced-dimensional multiuser detection, our results are based on binary modulation and can be extended to higher order modulation with symbols taking more possible values. With higher order modulation, however, the conditions to guarantee correct symbol detection may be stronger than the conditions to guarantee correct active user detection.

Another aspect of RD-MUD that we could extend further is the digital detectors for RD-MUD front-end. We have derived two digital detectors, the reduced-dimension decorrelator (RDD) and the reduced-dimension decision-feedback (RDDF). These two detectors both borrow ideas from the greedy algorithms for sparsity pattern recovery. As an alternative, we can also consider l_1 relaxation methods (as we reviewed in Section 2.3, and also in [9]) to estimate active users and their symbols. From the RD-MUD output model (4.20), the l_1 relaxation method solves the following convex optimization problem to estimate \mathbf{b}

$$\text{minimize}_{\mathbf{b}} \|\mathbf{y} - \mathbf{A}\mathbf{R}\mathbf{b}\|_2^2 + \lambda \|\mathbf{b}\|_1, \quad (5.3)$$

where the $\|\mathbf{x}\|_1 = \sum_{n=1}^N |\mathbf{x}_n|$ is the l_1 norm of an N -dimensional vector \mathbf{x} , and $\lambda > 0$ is a regularization parameter. The above optimization problem can be solved efficiently by various numerical optimization solvers (see e.g., [9]). Then we can take the support of the solution to (5.3) as the active users, and the signs of the corresponding entries as symbols for active users. The condition on \mathbf{A} such that (5.3) can obtain small probability-of-error is subject to further investigation.

We have obtained some performance guarantees for RD-MUD, such as the conditions on the coherence of the coefficient matrix for RD-MUD detectors to obtain

a small probability-of-error. However, we have not addressed the issue of finding an optimal coefficient matrix to achieve the minimum probability-of-error. For example, we could formulate an optimization problem, to minimize the upper-bounds on the probability-of-error of the RDD and RDDF detectors, which we have derived in (4.54), with respect to matrix \mathbf{A} . However, in general, finding the optimal deterministic coefficient matrix is an open problem in the compressive sensing literature (see [7] and some discussion therein).

Lastly, the reduced-dimension multi-user detection leads to the notion of approximate sufficient statistics. It is well-known that the MF-bank front-end obtains a set of sufficient statistics for multi-user detection (MUD). On the other hand, the front-end of RD-MUD uses fewer correlators and maps the original sufficient statistic into a lower dimension subspace. With proper design, RD-MUD can obtain a performance approximate to that based on the sufficient statistics. In this sense, the RD-MUD front-end obtains a set of approximate sufficient statistics for MUD. It is an interesting theoretical direction to generalize this notion of approximate sufficient statistics to other problems.

Appendix A

Proof for Multi-Sensor Change-Point Detection

A.1 Proof of Lemma 4

Proof. We use the following identity:

$$\max_{0 \leq k < T-k_0} h_k = \max \left\{ \max_{0 \leq k < k_0} h_k, \max_{k_0 \leq k < T-k_0} h_k \right\}, \quad (\text{A.1})$$

which holds for any sequence h_k . In particular,

$$\begin{aligned} h_k = & - \sum_{n \in \mathcal{N}_a} \mu_n (S_{n,k} - k\mu_n/2) + \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T-k)\mu_n]^2 / [2(T-k)] \\ & - \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})^-]^2 / 2(T-k) + \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0). \end{aligned} \quad (\text{A.2})$$

First we evaluate the term $\max_{0 \leq k < k_0} h_k$ in (A.1). Since when $b \rightarrow \infty$, $T \sim b/\Delta$. For $0 \leq k < k_0$, $k/T < k_0/T = \sqrt{b}/T = \Delta/\sqrt{b} \rightarrow 0$. Also, $S_{n,k}/S_{n,T} \rightarrow k\Delta/(T\Delta) \rightarrow 0$ by the law of large number and the previous argument. Hence we have

$$\begin{aligned} & \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T-k)\mu_n]^2 / [2(T-k)] \\ &= \sum_{n \in \mathcal{N}_a} \frac{[S_{n,T}(1 - S_{n,k}/S_{n,T}) - T(1 - k/T)\mu_n]^2}{2T(1 - k/T)} \\ &\rightarrow \sum_{n \in \mathcal{N}_a} (S_{n,T} - T\mu_n)^2 / (2T). \end{aligned} \quad (\text{A.3})$$

Again, for $0 \leq k < k_0$, $k/T < k_0/T = \sqrt{b}/T = \Delta/\sqrt{b} \rightarrow 0$. Also, $S_{n,k}/S_{n,T} \rightarrow k\Delta/(T\Delta) \rightarrow 0$ by the law of large number. So by rewriting the term in the following way, we have

$$\frac{[(S_{n,T} - S_{n,k})^-]^2}{2(T-k)} = \frac{[(S_{n,T}(1 - S_{n,k}/S_{n,T}))^-]^2}{2T(1 - k/T)} \rightarrow \frac{[(S_{n,T})^-]^2}{2T}. \quad (\text{A.4})$$

Then we write the term (A.4) as

$$\frac{[(S_{n,T})^-]^2}{2T} = \left[\left(\frac{S_{n,T} - T\mu_n}{\sqrt{2T}} + \frac{\sqrt{T}\mu_n}{\sqrt{2}} \right)^- \right]^2. \quad (\text{A.5})$$

By the Anscombe-Doeblin Lemma [80], $(S_{n,T} - T\mu_n)/T^{1/2}$ is asymptotically normally distributed with zero mean and unit variance. As $b \rightarrow \infty$, $T \sim b/\Delta$. Since $\mu_n > 0$ for $n \in \mathcal{N}_a$, when $b \rightarrow \infty$, $T \rightarrow \infty$, and hence $\left(\frac{S_{n,T} - T\mu_n}{\sqrt{2T}} + \frac{\sqrt{T}\mu_n}{\sqrt{2}} \right)^- \rightarrow 0$. As a result, as $b \rightarrow \infty$, the left-hand side of (A.5) tend to 0. Consequently, $[(S_{n,T} - S_{n,k})^-]^2/[2(T - k)] \rightarrow 0$.

Since $g(x; p_0) = \log[1 - p_0 + p_0 \exp((x^+)^2/2)]$ as a function of p_0 is monotonically increasing in p , we have

$$g(x; p_0) \leq (x^+)^2/2. \quad (\text{A.6})$$

Hence the last term in (A.2) is upper-bounded by:

$$g(U_{n,k,T}; p_0) = \log(1 - p_0 + p_0 \exp[(U_{n,k,T}^+)^2/2]) \leq [(S_{n,T} - S_{n,k})^+]^2/[2(T - k)]. \quad (\text{A.7})$$

Using a similar argument to (A.3), when $b \rightarrow \infty$, for $k \leq k_0$, we have

$$[(S_{n,T} - S_{n,k})^+]^2/[2(T - k)] \rightarrow [(S_{n,T})^+]^2/(2T). \quad (\text{A.8})$$

By the bounded convergence theorem,

$$\sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \rightarrow \sum_{n \in \mathcal{N}_a^c} \log \{1 - p_0 + p_0 \exp[(S_{n,T}^+)^2/(2T)]\} = \sum_{n \in \mathcal{N}_a^c} g(U_{n,0,T}; p_0). \quad (\text{A.9})$$

Hence when $b \rightarrow \infty$,

$$\max_{0 \leq k < k_0} h_k \rightarrow \max_{0 \leq k < k_0} \left\{ - \sum_{n \in \mathcal{N}_a} \mu_n(S_{n,k} - k\mu_n/2) \right\} + \sum_{n \in \mathcal{N}_a} (S_{n,T} - T)^2 / (2T) + \sum_{n \in \mathcal{N}_a^c} g(U_{n,0,T}; p_0). \quad (\text{A.10})$$

Next we evaluate $\max_{k_0 \leq k < T-k_0} h_k$ in (A.1). Note that when $b \rightarrow \infty$,

$$\begin{aligned} - \sum_{n \in \mathcal{N}_a} \mu_n(S_{n,k} - k\mu_n/2) &\rightarrow -\Delta k \leq -\Delta(T - k_0) \\ &\sim -\Delta(b/\Delta - \sqrt{b}) = -b(1 - \Delta b^{-1/2}) \rightarrow -\infty, \end{aligned} \quad (\text{A.11})$$

and the other three terms are bounded with high probability. By Markov inequality and (A.3), we have that for a constant $c > 0$,

$$\begin{aligned} &\mathbb{P}^0 \left\{ \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T - k)\mu_n]^2 / [2(T - k)] \geq c \right\} \\ &\leq \mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k}) - (T - k)\mu_n]^2 / [2(T - k)] \right] / c = M/c, \end{aligned} \quad (\text{A.12})$$

and

$$\begin{aligned} &\mathbb{P}^0 \left\{ \sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})^-]^2 / [2(T - k)] \geq c \right\} \\ &\leq \mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})^-]^2 / [2(T - k)] \right] / c \\ &\leq \mathbb{E}^0 \left[\sum_{n \in \mathcal{N}_a} [(S_{n,T} - S_{n,k})]^2 / [2(T - k)] \right] / c \leq M/c. \end{aligned} \quad (\text{A.13})$$

By monotonicity (A.6) and Markov inequality,

$$\mathbb{P}^0 \left\{ \sum_{n \in \mathcal{N}_a^c} g(U_{n,k,T}; p_0) \geq c \right\} \leq \mathbb{P}^0 \left\{ \sum_{n \in \mathcal{N}_a^c} (S_{n,T} - S_{n,k})^2 / [2(T - k)] \geq c \right\} \leq (N - M) / c. \quad (\text{A.14})$$

Hence $\max_{k_0 \leq k < T - k_0} h_k \rightarrow -\infty$ as $b \rightarrow \infty$. Substitute this result and (A.10) into (A.1), we have

$$\max_{0 \leq k < T - k_0} h_k = \max \left\{ \max_{0 \leq k < k_0} h_k, \max_{k_0 \leq k < T - k_0} h_k \right\} \rightarrow \max_{0 \leq k < k_0} h_k, \quad (\text{A.15})$$

which concludes the proof for Lemma 4. \square

A.2 Multiple Overlapping Sources

When there are Q sources, the profile is given by

$$\mu_n = \sum_{m=1}^Q r_m \alpha_{\mathbf{z}_m}(u_n, v_n), \quad (\text{A.16})$$

with r_m specifying the amplitude and \mathbf{z}_m specifying the location of the source. The loglikelihood function is given by (3.65) with μ_n replaced by the new expression in

(A.16):

$$\begin{aligned}
l(t, k, \{r_m\}, \mathbf{z}) &= \sum_{n=1}^N \sum_{l=k+1}^t \left[y_{n,l} \sum_{m=1}^Q r_m \alpha_{\mathbf{z}_m}(u_n, v_n) \right] - \frac{1}{2} \left(\sum_{m=1}^Q r_m \alpha_{\mathbf{z}_m}(u_n, v_n) \right)^2 \\
&= \left[\sum_{m=1}^Q r_m (t-k)^{1/2} \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t} \right] - \frac{t-k}{2} \left[\sum_{m=1}^Q r_m^2 \right] \\
&\quad - \frac{t-k}{2} \left[\sum_{m \neq p} r_m r_p \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \boldsymbol{\alpha}_{\mathbf{z}_p} \right].
\end{aligned} \tag{A.17}$$

Setting the derivative of the log-likelihood function (A.17) with respect to each r_m to zero gives us a set of linear equations:

$$r_m + \frac{1}{2} \sum_{p \neq m} r_p \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \boldsymbol{\alpha}_{\mathbf{z}_p} = (t-k)^{-1/2} \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t}, \quad m = 1, \dots, Q. \tag{A.18}$$

We can write (A.18) as:

$$\mathbf{A} \mathbf{r} = \mathbf{b}_{k,t}, \tag{A.19}$$

with M equations and M unknowns. The vector $\mathbf{r} = [r_1, \dots, r_Q]^\top$. The matrix \mathbf{A} has ones on the diagonal, and the entry on the m th row and the n th column is given by $\frac{1}{2} \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \boldsymbol{\alpha}_{\mathbf{z}_n}$. The m th element of the vector $\mathbf{b}_{k,t}$ is given by $(t-k)^{-1/2} \boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t}$. If \mathbf{A} is invertible (otherwise the sources cannot be uniquely identified), from (A.19) we can solve an estimate for the vector of source amplitudes $\hat{\mathbf{r}} = \mathbf{A}^{-1} \mathbf{b}_{k,t}$, for a given set of source locations. By plugging the estimate $\hat{\mathbf{r}}$ into the likelihood function (3.65), and maximizing the resulted statistic with respect to k and all possible source locations $(\mathbf{z}_1, \dots, \mathbf{z}_Q)$ in the set \mathcal{D} , we obtain the GLR statistics when we have multiple overlapping sources.

If we assume the sources are non-overlapping, then $\boldsymbol{\alpha}_{\mathbf{z}_m}^\top \boldsymbol{\alpha}_{\mathbf{z}_p} = 0$, for $p \neq m$. Then we have

$$\hat{r}_m = \left[\frac{\boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t}}{(t-k)^{1/2}} \right]^+, \quad (\text{A.20})$$

for a candidate source location \mathbf{z}_m . If we plugging this amplitude estimate into (A.17), we have the following log GLR statistic:

$$\sum_{m=1}^Q \frac{1}{2} \left(\left[\frac{\boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t}}{(t-k)^{1/2}} \right]^+ \right)^2, \quad (\text{A.21})$$

and the corresponding profile-based procedure is given by

$$T_{\text{profile}} = \inf \left\{ t \geq 1 : \max_{t-m_1 \leq k < t} \max_{\mathbf{z}_1, \dots, \mathbf{z}_Q \in \mathcal{D}} \sum_{m=1}^Q \frac{1}{2} \left(\left[\frac{\boldsymbol{\alpha}_{\mathbf{z}_m}^\top \mathbf{U}_{k,t}}{(t-k)^{1/2}} \right]^+ \right)^2 \geq b \right\}. \quad (\text{A.22})$$

Appendix B

Proof for RD-MUD

B.1 Covariance of RD-MUD Output Noise

Consider the covariance of the output noise at the n th and the m th branches in the RD-MUD front-end:

$$\begin{aligned}\rho_{nm} &= \mathbb{E}\{w_n w_m\} = \mathbb{E} \left\{ \sum_{k=1}^N \sum_{l=1}^N a_{nk} a_{ml} \langle \hat{s}_k(t), w(t) \rangle \langle \hat{s}_l(t), w(t) \rangle \right\} \\ &= \sum_{k=1}^N \sum_{l=1}^N a_{nk} a_{ml} \mathbb{E} \{ \langle \hat{s}_k(t), w(t) \rangle \langle \hat{s}_l(t), w(t) \rangle \}\end{aligned}\tag{B.1}$$

We want to show that $\rho_{nm} = \sigma^2 [\mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H]_{nm}$. We have

$$\begin{aligned}& \mathbb{E} \{ \langle \hat{s}_k(t), w(t) \rangle \langle \hat{s}_l(t), w(t) \rangle \} \\ &= T^{-2} \int_0^T \int_0^T \hat{s}_k(t) \hat{s}_l(u) \mathbb{E} \{ w(t) w(u) \} dt du \\ &= T^{-2} \int_0^T \int_0^T \hat{s}_k(t) \hat{s}_l(u) \sigma^2 \delta(t - u) dt du \\ &= \sigma^2 T^{-1} \int_0^T \hat{s}_k(t) \hat{s}_l(t) dt \\ &= \sigma^2 \left\langle \sum_n [\mathbf{G}^{-1}]_{nk} s_n(t), \sum_m [\mathbf{G}^{-1}]_{ml} s_m(t) \right\rangle \\ &= \sigma^2 \sum_n \sum_m [\mathbf{G}^{-1}]_{nk} [\mathbf{G}^{-1}]_{ml} \langle s_n(t), s_m(t) \rangle \\ &= \sigma^2 \sum_n \sum_m [\mathbf{G}^{-1}]_{nk} [\mathbf{G}^{-1}]_{ml} [\mathbf{G}]_{nm} = \sigma^2 [\mathbf{G}^{-1}]_{lk}\end{aligned}\tag{B.2}$$

Substituting this back into (B.1), we have

$$\rho_{nm} = \mathbb{E}\{w_n w_m\} = \sigma^2 \sum_{k=1}^N \sum_{l=1}^N a_{nk} a_{ml} [\mathbf{G}^{-1}]_{lk} = \sigma^2 [\mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H]_{nm}.\tag{B.3}$$

Setting $n = m$, we have

$$\sigma_m = \mathbb{E}\{w_m^2\} = \sigma^2[\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H]_{mm}. \quad (\text{B.4})$$

This completes the proof.

B.2 Derivation of RD-MUD MMSE

Proof. Given the active user index set $\hat{\mathcal{I}}$ obtained from (4.23), we define $\mathbf{W} = \mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}^2\mathbf{A}_{\hat{\mathcal{I}}}^H + \sigma^2\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H$, and $\bar{\mathbf{M}} = \mathbf{R}_{\hat{\mathcal{I}}}\mathbf{A}_{\hat{\mathcal{I}}}^H\mathbf{W}^{-1}$. We want to show that $\bar{\mathbf{M}} = \arg \min_{\mathbf{M}} \mathbb{E}\{\|\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y}\|^2\}$. Using the same method for deriving the conventional MMSE detector of the MF-bank [96], we assume that $\mathbf{b}_{\hat{\mathcal{I}}}$ has a distribution that is uncorrelated with the noise \mathbf{w} and that $\mathbb{E}\{\mathbf{b}_{\hat{\mathcal{I}}}\mathbf{b}_{\hat{\mathcal{I}}}^H\} = \mathbf{I}$. Based on $\hat{\mathcal{I}}$, we refer to the model (4.21). Since $\|\mathbf{x}\|^2 = \text{tr}(\mathbf{x}\mathbf{x}^H)$, we can write the MSE as $\mathbb{E}\{\|\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y}\|^2\} = \text{tr}(\mathbb{E}\{(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})^H\})$. Now we expand

$$\begin{aligned} & \mathbb{E}\{(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})^H\} \\ &= \mathbb{E}\{\mathbf{b}_{\hat{\mathcal{I}}}\mathbf{b}_{\hat{\mathcal{I}}}^H\} - \mathbb{E}\{\mathbf{b}_{\hat{\mathcal{I}}}\mathbf{y}^H\}\mathbf{M}^H - \mathbf{M}\mathbb{E}\{\mathbf{y}\mathbf{b}_{\hat{\mathcal{I}}}^H\} + \mathbf{M}\mathbb{E}\{\mathbf{y}\mathbf{y}^H\}\mathbf{M}^H \\ &= \mathbf{I} + \mathbf{M}(\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}^2\mathbf{A}_{\hat{\mathcal{I}}}^H + \sigma^2\mathbf{A}\mathbf{G}^{-1}\mathbf{A}^H)\mathbf{M}^H - \mathbf{R}_{\hat{\mathcal{I}}}\mathbf{A}_{\hat{\mathcal{I}}}^H\mathbf{M}^H - \mathbf{M}\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}}. \end{aligned} \quad (\text{B.5})$$

It can be verified that $\mathbf{M}\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}} = \mathbf{M}\mathbf{W}\bar{\mathbf{M}}^H$. Hence from (B.5), we have

$$\begin{aligned}
& \mathbb{E}\{(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})^H\} \\
&= \mathbf{I} + \mathbf{M}\mathbf{W}\mathbf{M}^H - \bar{\mathbf{M}}\mathbf{W}\mathbf{M}^H - \mathbf{M}\mathbf{W}\bar{\mathbf{M}}^H \\
&= \mathbf{I} - \bar{\mathbf{M}}\mathbf{W}\bar{\mathbf{M}}^H + (\mathbf{M} - \bar{\mathbf{M}})\mathbf{W}(\mathbf{M} - \bar{\mathbf{M}})^H \\
&= \mathbf{I} - \mathbf{R}_{\hat{\mathcal{I}}}\mathbf{A}_{\hat{\mathcal{I}}}^H\mathbf{W}^{-1}\mathbf{A}_{\hat{\mathcal{I}}}\mathbf{R}_{\hat{\mathcal{I}}} + (\mathbf{M} - \bar{\mathbf{M}})\mathbf{W}(\mathbf{M} - \bar{\mathbf{M}})^H.
\end{aligned} \tag{B.6}$$

Since \mathbf{W} is a positive semidefinite matrix, the trace of the second term in (B.6) is always nonnegative. So we conclude that the matrix \mathbf{M} that minimizes the MSE $\text{tr}(\mathbb{E}\{(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})(\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y})^H\})$ is $\bar{\mathbf{M}}$, or equivalently, $\mathbb{E}\{\|\mathbf{b}_{\hat{\mathcal{I}}} - \mathbf{M}\mathbf{y}\|^2\}$, as required. \square

B.3 Proof of Theorem 1

The proof of Theorem 1 for both the RDD and RDDF detectors are closely related. To exploit this similarity, we first prove several lemmas that are useful for both results. First, we will demonstrate that the random event

$$\mathcal{G} \triangleq \left\{ \max_{1 \leq n \leq N} |\mathbf{a}_n^H \mathbf{w}| < \tau \right\} \tag{B.7}$$

occurs with high probability, where

$$\tau \triangleq \sigma \sqrt{2(1 + \alpha) \log N} \cdot \sqrt{\lambda_{\max}(\mathbf{G}^{-1})} \cdot \sqrt{\max_n (\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n)}, \tag{B.8}$$

and $\alpha > 0$. Then we show that when \mathcal{G} occurs, both algorithms can detect the active users and their symbols. The proofs follow the arguments in [9] with modifications to account for the fact that \mathbf{w} is colored noise, and the error can also be caused by

incorrect symbol detection. However, as we will show, the error probability of active user detection dominates the latter case.

Sidak's lemma [77] states the following:

Lemma 7 (Sidak's lemma). *Let $[X_1, \dots, X_n]^\top$ be a vector of random variables having the n -dimensional normal distribution with zero means, arbitrary variances $\sigma_1^2, \dots, \sigma_n^2$, and an arbitrary correlation matrix $[\mathbf{P}]_{mk} = \rho_{mk}$. Then, for any positive numbers c_1, \dots, c_n ,*

$$\mathbb{P}\{|X_1| \leq c_1, |X_2| \leq c_2, \dots, |X_n| \leq c_n\} \geq \mathbb{P}\{|X_1| \leq c_1\} \cdot \mathbb{P}\{|X_2| \leq c_2, \dots, |X_n| \leq c_n\}. \quad (\text{B.9})$$

Lemma 8. *Suppose that \mathbf{w} is a Gaussian random vector with zero mean and covariance $\sigma^2 \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H$. If $N^{-(1+\alpha)} [\pi(1+\alpha) \log N]^{-1/2} \leq 1$ for some $\alpha > 0$, then the event \mathcal{G} of (B.7) occurs with probability at least one minus (4.46).*

Proof. The random variables $\{\mathbf{a}_n^H \mathbf{w}\}_{n=1}^N$ are jointly Gaussian, with means equal to zero, variances σ_n^2 equal to $\sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n$, and covariances ρ_{nm} between the n th and m th random variables equal to $\sigma^2 \mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_m$. Define

$$\hat{\tau} \triangleq \sigma [2(1+\alpha) \log N]^{1/2} \cdot \left[\max_n (\mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n) \right]^{1/2}, \quad (\text{B.10})$$

and an event

$$\hat{\mathcal{G}} \triangleq \left\{ \max_{1 \leq n \leq N} |\mathbf{a}_n^H \mathbf{w}| < \hat{\tau} \right\}. \quad (\text{B.11})$$

Using Sidak's lemma, we have

$$\mathbb{P}\{\hat{\mathcal{G}}\} = \mathbb{P}\{|\mathbf{a}_1^H \mathbf{w}| < \hat{\tau}, \dots, |\mathbf{a}_N^H \mathbf{w}| < \hat{\tau}\} \geq \prod_{n=1}^N \mathbb{P}\{|\mathbf{a}_n^H \mathbf{w}| < \hat{\tau}\}. \quad (\text{B.12})$$

Since $\mathbf{a}_n^H \mathbf{w}$ is a Gaussian random variable with zero mean and variance σ_n^2 , the tail probability of the colored noise can be written as

$$\mathbb{P}\{|\mathbf{a}_n^H \mathbf{w}| < \hat{\tau}\} = 1 - 2Q\left(\frac{\hat{\tau}}{\sigma_n}\right). \quad (\text{B.13})$$

By the bound (4.59) on $Q(x)$, (B.13) can be bounded as

$$\mathbb{P}\{|\mathbf{a}_n^H \mathbf{w}| < \hat{\tau}\} \geq 1 - \eta_n, \quad (\text{B.14})$$

where $\eta_n \triangleq \sqrt{\frac{2}{\pi}} \cdot \frac{\sigma_n}{\hat{\tau}} e^{-\hat{\tau}^2/(2\sigma_n^2)}$. Define

$$\begin{aligned} \sigma_{\max} &\triangleq \max_n \sigma_n = \sigma \left[\max_n (\mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n) \right]^{1/2}, \\ \eta_{\max} &\triangleq \sqrt{\frac{2}{\pi}} \frac{\sigma_{\max}}{\hat{\tau}} e^{-\hat{\tau}^2/(2\sigma_{\max}^2)}. \end{aligned} \quad (\text{B.15})$$

Since $\sigma_{\max}/\hat{\tau} = [2(1+\alpha)\log N]^{-1/2}$ by the definition of $\hat{\tau}$, we have $\eta_{\max} = \sqrt{\frac{2}{\pi}}[2(1+\alpha)\log N]^{-1/2}e^{-(1+\alpha)\log N}$. It is easy to show that η_n increases as σ_n increases. Hence $\eta_n \leq \eta_{\max}$. When $\eta_{\max} \leq 1$, we can use the inequality $(1-x)^N \geq 1-Nx$ when $x \geq 0$ and substitute the value of η_{\max} to write (B.12) as

$$\mathbb{P}\{\hat{\mathcal{G}}\} \geq \prod_{n=1}^N (1-\eta_n) \geq (1-\eta_{\max})^N \geq 1-N\eta_{\max} = 1-N^{-\alpha}[\pi(1+\alpha)\log N]^{-1/2}, \quad (\text{B.16})$$

which holds for any $\eta_{\max} \leq 1$ and $N \geq 1$.

Next we show that $\hat{\tau} \leq \tau$. Note that

$$\mathbf{a}_n^H \mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H \mathbf{a}_n \leq \|\mathbf{A}^H \mathbf{a}_n\|^2 \lambda_{\max}(\mathbf{G}^{-1}) \leq [\max_n (\mathbf{a}_n^H \mathbf{A} \mathbf{A}^H \mathbf{a}_n)] \lambda_{\max}(\mathbf{G}^{-1}). \quad (\text{B.17})$$

From inequality (B.17) and definitions (B.8) for τ and (B.10) for $\hat{\tau}$, we obtain $\hat{\tau} \leq \tau$.

Hence

$$\mathbb{P}\{\mathcal{G}\} = \mathbb{P}\{\max_n |\mathbf{a}_n^H \mathbf{w}| < \tau\} \geq \mathbb{P}\{\max_n |\mathbf{a}_n^H \mathbf{w}| < \hat{\tau}\} = \mathbb{P}\{\hat{\mathcal{G}}\}. \quad (\text{B.18})$$

Combining (B.16) and (B.18), we conclude that $\mathbb{P}\{\mathcal{G}\}$ is greater than one minus the expression (4.46), as required. \square

The next lemma shows that, under appropriate conditions, ranking the inner products between \mathbf{a}_n and \mathbf{y} is an effective method of detecting the set of active users. The proof of this lemma is adapted from Lemma 3 in [9] to account for the fact that the signal vector \mathbf{y} here can be complex as \mathbf{A} can be complex. Since only the real part contains all the useful information, to prove this lemma, we basically follow the proof for Lemma 3 in [9] while using the following inequality whenever needed: $|\Re[\mathbf{a}_n^H \mathbf{a}_m]| \leq |\mathbf{a}_n^H \mathbf{a}_m| \leq \mu$ for $n \neq m$, and $|\Re[\mathbf{a}_n^H \mathbf{w}]| \leq |\mathbf{a}_n^H \mathbf{w}|$.

Lemma 9. *Let \mathbf{b} be a vector with support \mathcal{I} which consists of K active users, and let $\mathbf{y} = \mathbf{A} \mathbf{R} \mathbf{b} + \mathbf{w}$ for a Gaussian noise vector \mathbf{w} with zero mean and covariance $\mathbf{A} \mathbf{G}^{-1} \mathbf{A}^H$. Define $|r_{\max}|$ and $|r_{\min}|$ as in (4.44), and suppose that*

$$|r_{\min}| - (2K - 1)\mu|r_{\max}| \geq 2\tau. \quad (\text{B.19})$$

Then, if the event \mathcal{G} of (B.7) occurs, we have

$$\min_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| > \max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|. \quad (\text{B.20})$$

If, rather than (B.19), a weaker condition holds:

$$|r_{\max}| - (2K - 1)\mu|r_{\max}| \geq 2\tau. \quad (\text{B.21})$$

Then, if the event \mathcal{G} of (B.7) occurs, we have

$$\max_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| > \max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|. \quad (\text{B.22})$$

Proof. We begin by deriving a lower-bound for $\min_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|$ under the event \mathcal{G} :

$$\begin{aligned} \min_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| &= \min_{n \in \mathcal{I}} \left| b_n r_n + \sum_{m \neq n} b_m r_m \Re[\mathbf{a}_n^H \mathbf{a}_m] + \Re[\mathbf{a}_n^H \mathbf{w}] \right| \\ &\geq \min_{n \in \mathcal{I}} \left(|b_n||r_n| - \left| \sum_{m \neq n} b_m r_m \Re[\mathbf{a}_n^H \mathbf{a}_m] \right| - |\Re[\mathbf{a}_n^H \mathbf{w}]| \right) \\ &\geq |r_{\min}| - \max_{n \in \mathcal{I}} \sum_{m \neq n} |b_m||r_m||\mathbf{a}_n^H \mathbf{a}_m| - \max_{n \in \mathcal{I}} |\mathbf{a}_n^H \mathbf{w}| \\ &> |r_{\min}| - (K - 1)\mu|r_{\max}| - \tau, \end{aligned} \quad (\text{B.23})$$

where we have used the triangle inequality, the fact that $|b_n| = 1$, $|r_{\min}| \leq |r_n| \leq |r_{\max}|$, $|\Re[\mathbf{a}_n^H \mathbf{a}_m]| \leq |\mathbf{a}_n^H \mathbf{a}_m| \leq \mu$ for $n \neq m$, and $|\Re[\mathbf{a}_n^H \mathbf{w}]| \leq |\mathbf{a}_n^H \mathbf{w}|$. On the other hand, we can similarly expand and upper-bound $\max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|$, under the event

\mathcal{G} , as

$$\begin{aligned}
\max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| &= \max_{n \notin \mathcal{I}} \left| \sum_{m \in \mathcal{I}} b_m r_m \Re[\mathbf{a}_n^H \mathbf{a}_m] + \Re[\mathbf{a}_n^H \mathbf{w}] \right| \\
&\leq \max_{n \notin \mathcal{I}} \sum_{m \in \mathcal{I}} |b_m| |r_m| |\Re[\mathbf{a}_n^H \mathbf{a}_m]| + \max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{w}]| \\
&< K\mu |r_{\max}| + \tau.
\end{aligned} \tag{B.24}$$

Combining (B.23) and (B.24), we have that under the event \mathcal{G} ,

$$\min_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| > |r_{\min}| - (2K - 1)\mu |r_{\max}| - 2\tau + \max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|. \tag{B.25}$$

So when \mathcal{G} occurs, under the condition (B.19), we obtain (B.20).

Similarly, when \mathcal{G} occurs, we expand and lower-bound $\max_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|$. Assume that n_0 is the index achieving the largest absolute gain: $|r_{n_0}| = |r_{\max}|$. Then under event \mathcal{G} :

$$\begin{aligned}
\max_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| &\geq |\Re[\mathbf{a}_{n_0}^H \mathbf{y}]| = \left| b_{n_0} r_{n_0} + \sum_{m \neq n_0} b_m r_m \Re[\mathbf{a}_{n_0}^H \mathbf{a}_m] + \Re[\mathbf{a}_{n_0}^H \mathbf{w}] \right| \\
&\geq |r_{\max}| - \sum_{m \neq n_0} |b_m| |r_m| |\Re[\mathbf{a}_{n_0}^H \mathbf{a}_m]| - |\Re[\mathbf{a}_{n_0}^H \mathbf{w}]| \\
&> |r_{\max}| - (K - 1)\mu |r_{\max}| - \tau.
\end{aligned} \tag{B.26}$$

Combining (B.26) and (B.24), we have that under the event \mathcal{G} ,

$$\max_{n \in \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]| > |r_{\max}| - (2K - 1)\mu |r_{\max}| - 2\tau + \max_{n \notin \mathcal{I}} |\Re[\mathbf{a}_n^H \mathbf{y}]|. \tag{B.27}$$

So when \mathcal{G} occurs, under the condition (B.21), we obtain (B.22), as required. \square

The following lemma demonstrates that the sign detector can effectively detect transmitted symbols for the RDD and RDDF detectors.

Lemma 10. *Let \mathbf{b} be a vector with $b_n \in \{1, -1\}$, for $n \in \mathcal{I}$ and $b_n = 0$ otherwise, and let $\mathbf{y} = \mathbf{A}\mathbf{R}\mathbf{b} + \mathbf{w}$ for a Gaussian noise vector \mathbf{w} with zero mean and covariance $\sigma^2 \mathbf{A}\mathbf{G}^{-1} \mathbf{A}^H$. Suppose that*

$$|r_{\min}| - (K - 1)\mu|r_{\max}| \geq \tau. \quad (\text{B.28})$$

Then, if the event \mathcal{G} occurs, we have

$$\text{sgn}(r_n \Re[\mathbf{a}_n^H \mathbf{y}]) = b_n, \quad n \in \mathcal{I}. \quad (\text{B.29})$$

If, instead of (B.28), a weaker condition

$$|r_{\max}| + |r_{\min}| - 2(K - 1)\mu|r_{\max}| \geq 2\tau \quad (\text{B.30})$$

holds, then under the event \mathcal{G} , we have

$$\text{sgn}(r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{y}]) = b_{n_1}, \quad (\text{B.31})$$

for

$$n_1 = \arg \max_n |\Re[\mathbf{a}_n^H \mathbf{y}]|. \quad (\text{B.32})$$

Proof. To detect correctly, for $b_n = 1$, $\Re[r_n \mathbf{a}_n^H \mathbf{y}]$ has to be positive, and for $b_n = -1$, $\Re[r_n \mathbf{a}_n^H \mathbf{y}]$ has to be negative. First assume $b_n = 1$. We expand $\Re[r_n \mathbf{a}_n^H \mathbf{y}]$, find the lower-bound and the condition such that the lower bound is positive. Substituting in

the expression for \mathbf{y} , using the inequality that $x + y + z \geq x - |y| - |z|$, under the event \mathcal{G} , we obtain

$$\begin{aligned} \Re[r_n \mathbf{a}_n^H \mathbf{y}] &= |r_n|^2 + \sum_{m \neq n} b_m r_n r_m \Re[\mathbf{a}_n^H \mathbf{a}_m] + r_n \Re[\mathbf{a}_n^H \mathbf{w}] \\ &\geq |r_n| |r_{\min}| - \sum_{m \neq n} |r_n| |r_m| |\Re[\mathbf{a}_n^H \mathbf{a}_m]| - |r_n| |\Re[\mathbf{a}_n^H \mathbf{w}]| \\ &> |r_n| [|r_{\min}| - (K-1)\mu|r_{\max}| - \tau]. \end{aligned} \quad (\text{B.33})$$

From (B.33), $\Re[r_n \mathbf{a}_n^H \mathbf{y}] > 0$ for $n \in \mathcal{I}$ if (B.28) holds and $b_n = 1$. Similarly, we can show for $b_n = -1$, under event \mathcal{G} , if (B.28) holds, $\Re[r_n \mathbf{a}_n^H \mathbf{y}] < 0$. Hence if (B.28) holds we obtain (B.29).

Recall that n_0 is the index of the largest gain: $|r_{n_0}| = |r_{\max}|$. Due to (B.32), we have

$$|\Re[\mathbf{a}_{n_1}^H \mathbf{y}]| \geq |\Re[\mathbf{a}_{n_0}^H \mathbf{y}]|. \quad (\text{B.34})$$

We will show that under the event \mathcal{G} , once (B.30) holds, then $\text{sgn}(r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{y}]) \neq b_{n_1}$ leads to a contradiction to (B.34). First assume $b_{n_1} = 1$. If $\hat{b}_{n_1} = \text{sgn}(r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{y}]) \neq b_{n_1}$, then

$$\hat{b}_{n_1} = \text{sgn} \left(r_{n_1}^2 + \sum_{m \neq n_1} b_m r_{n_1} r_m \Re[\mathbf{a}_{n_1}^H \mathbf{a}_m] + r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{w}] \right) = -1. \quad (\text{B.35})$$

So the expression inside the sgn operator of (B.35) must be negative. Since $r_{n_1}^2 > 0$, we must have

$$\sum_{m \neq n_1} b_m r_{n_1} r_m \Re[\mathbf{a}_{n_1}^H \mathbf{a}_m] + r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{w}] < 0. \quad (\text{B.36})$$

Multiplying the left-hand-side of (B.34) by $|r_{n_1}|$, and using the equality $|x| \cdot |y| = |xy|$,

we obtain

$$\begin{aligned}
|r_{n_1}| |\Re[\mathbf{a}_{n_1}^H \mathbf{y}]| &= |r_{n_1}| \left| r_{n_1} + \sum_{m \neq n_1} b_m r_m \Re[\mathbf{a}_{n_1}^H \mathbf{a}_m] + \Re[\mathbf{a}_{n_1}^H \mathbf{w}] \right| \\
&= \left| r_{n_1}^2 + \sum_{m \neq n_1} b_m r_{n_1} r_m \Re[\mathbf{a}_{n_1}^H \mathbf{a}_m] + r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{w}] \right|.
\end{aligned} \tag{B.37}$$

Due to (B.35), the last line of (B.37) inside the $|\cdot|$ operator is negative. Using the fact that $r_{n_1}^2 > 0$ and (B.36), and the identity $|x + y| = -(x + y) = |y| - x$ when $x + y < 0$ and $y < 0$, under the event \mathcal{G} , we obtain that

$$\begin{aligned}
|r_{n_1}| |\Re[\mathbf{a}_{n_1}^H \mathbf{y}]| &= \left| \sum_{m \neq n_1} b_m r_{n_1} r_m \Re[\mathbf{a}_{n_1}^H \mathbf{a}_m] + r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{w}] \right| - r_{n_1}^2 \\
&< |r_{n_1}| (K - 1) \mu |r_{\max}| + |r_{n_1}| \tau - |r_{n_1}| |r_{\min}| \\
&= |r_{n_1}| [(K - 1) \mu |r_{\max}| + \tau - |r_{\min}|].
\end{aligned} \tag{B.38}$$

On the other hand, multiply the right-hand-side of (B.34) by $|r_{n_1}|$. Similarly, using the equality $|x| \cdot |y| = |xy|$ and triangle inequality, under the event \mathcal{G} , we obtain

$$\begin{aligned}
|r_{n_1}| |\Re[\mathbf{a}_{n_0}^H \mathbf{y}]| &= \left| r_{n_1} r_{n_0} b_{n_0} + \sum_{m \neq n_0} b_m r_{n_1} r_m \Re[\mathbf{a}_{n_0}^H \mathbf{a}_m] + r_{n_1} \Re[\mathbf{a}_{n_0}^H \mathbf{w}] \right| \\
&> |r_{n_1}| [|r_{\max}| - (K - 1) \mu |r_{\max}| - \tau].
\end{aligned} \tag{B.39}$$

Combining (B.38) and (B.39), we have that once (B.30) holds, if $b_{n_1} = 1$, then $\text{sgn}(r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{y}]) = -1$ leads to $|\Re[\mathbf{a}_{n_1}^H \mathbf{y}]| < |\Re[\mathbf{a}_{n_0}^H \mathbf{y}]|$, which contradicts (B.34), and hence $\text{sgn}(r_{n_1} \Re[\mathbf{a}_{n_1}^H \mathbf{y}]) = 1$. A similar argument can be made for $b_{n_1} = -1$, which completes the proof.

□

We are now ready to prove Theorem 1. The proof for the RDD detector is obtained by combining Lemmas 8, 9 and 10. Lemma 8 ensures that the event \mathcal{G} occurs with probability at least as high as one minus (4.46). Whenever \mathcal{G} occurs, Lemma 9 guarantees by using (4.23), that the RDD detector can correctly detect active users under the condition (4.45), i.e. $\mathcal{G} \subset \{\hat{\mathcal{I}} = \mathcal{I}\}$. Finally, whenever \mathcal{G} occurs, Lemma 10 guarantees that, based on the correct support of active users, their transmitted symbols can be detected correctly under the condition (B.28), i.e. $\mathcal{G} \subset \{\hat{b}_n = b_n, n \in \mathcal{I}\}$. Clearly condition (B.28) is weaker than (4.45), since (4.45) can be written as $|r_{\min}| - (K-1)\mu|r_{\max}| \geq \tau + (\tau + K\mu|r_{\max}|) > \tau$, and hence if (4.45) holds then (B.28) also holds. In summary, under condition (4.45), $\mathcal{G} \subset \{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}$, and $1 - P_e = \mathbb{P}\{\{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}\} \geq \mathbb{P}\{\mathcal{G}\}$, which is greater than one minus (4.46), which concludes the proof for the RDD detector.

We now prove the performance guarantee for the RDDF detector adopting the technique used in proving Theorem 4 in [9]. First we show that whenever \mathcal{G} occurs, the RDDF detector correctly detects an active user in the first iteration, which follows from Lemmas 8 and 9. Note that (4.47) implies (B.21), and therefore, by Lemma 9, we have that by choosing the largest $|\Re[\mathbf{a}_n^H \mathbf{y}]|$, the RDDF detector can detect a correct user in the set \mathcal{I} . Second, we show that whenever \mathcal{G} occurs, the RDDF detector correctly detects the transmitted symbol of this active user. Note that (4.47) also implies (B.30), since (4.47) can be written as $|r_{\min}| \geq 2\tau/[1 - (2K-1)\mu]$, which implies $|r_{\max}| \geq 2\tau/[1 - (2K-1)\mu]$, and hence $|r_{\max}| + |r_{\min}| - 2(K-1)\mu|r_{\max}| \geq 2\tau[1 - 2(K-1)\mu]/[1 - (2K-1)\mu] + |r_{\min}| > 2\tau$, since $[1 - 2(K-1)\mu]/[1 - (2K-1)\mu] \geq 1$. Therefore, by Lemma 10, using a sign detector, we can detect the symbol correctly. Consequently, the first step of the RDDF detector correctly detect the active user

and its symbol, i.e. $\mathcal{G} \subset \{\hat{\mathcal{I}}^{(1)} \subset \mathcal{I}, b_{n_1}^{(1)} = b_{n_1}\}$.

The proof now continues by induction. Suppose we are currently in the k th iteration of the RDDF detector, $1 \leq k \leq K$, and assume that correct users and their symbols have been detected in all the $k - 1$ previous steps. The k th step is to detect the user with the largest $|\Re[\mathbf{a}_n^H \mathbf{v}^{(k)}]|$. Using the same notations as those in Section 4.2.2 and by definition of $\mathbf{v}^{(k)}$, we have

$$\mathbf{v}^{(k)} = \mathbf{A}\mathbf{R}(\mathbf{b} - \mathbf{b}^{(k-1)}) + \mathbf{w} = \mathbf{A}\mathbf{R}\mathbf{x}^{(k-1)} + \mathbf{w}, \quad (\text{B.40})$$

where $\mathbf{x}^{(k-1)} \triangleq \mathbf{b} - \mathbf{b}^{(k-1)}$. This vector has support $\mathcal{I}/\mathcal{I}^{(k-1)}$ and has at most $K - k + 1$ non-zero elements, since $\mathbf{b}^{(k-1)}$ contains correct symbols at the correct locations for $(k - 1)$ active users, i.e. $b_n^{(k-1)} = b_n$, for $n \in \mathcal{I}^{(k-1)}$. This $\mathbf{v}^{(k)}$ is a noisy measurement of the vector $\mathbf{A}\mathbf{R}\mathbf{x}^{(k-1)}$. The data model in (B.40) for the k th iteration is identical to the data model in the first iteration with \mathbf{b} replaced by $\mathbf{x}^{(k-1)}$ (with a smaller sparsity $K - k + 1$ rather than K), \mathcal{I} replaced by $\mathcal{I}/\mathcal{I}^{(k-1)}$, and \mathbf{y} replaced by $\mathbf{v}^{(k)}$. Since

$$|r_{\max}^{(k)}| \triangleq \max_{n \in \mathcal{I}/\mathcal{I}^{(k-1)}} |r_n| \geq |r_{\min}|, \quad (\text{B.41})$$

we have that under condition (4.47) this model (B.40) also satisfies the requirement (B.21). Consequently, by Lemma 9, we have that under the event \mathcal{G} ,

$$\max_{n \in \mathcal{I}/\mathcal{I}^{(k-1)}} |\Re[\mathbf{a}_n^H \mathbf{v}^{(k-1)}]| > \max_{n \in (\mathcal{I}/\mathcal{I}^{(k-1)})^c} |\Re[\mathbf{a}_n^H \mathbf{v}^{(k-1)}]|. \quad (\text{B.42})$$

Therefore, in the k th iteration, the RDDF detector can detect an active user correctly, i.e. $\mathcal{G} \subset \{\hat{\mathcal{I}}^{(k)} \subset \mathcal{I}\}$, and no index of active users that has been detected before

will be chosen again. On the other hand, since (4.47) can be written as $|r_{\min}| \geq 2\tau/[1 - (2K - 1)\mu]$, from (B.41) this implies $|r_{\max}^{(k)}| \geq 2\tau/[1 - (2K - 1)\mu]$, and hence $|r_{\max}^{(k)}| - (2K - 1)\mu|r_{\max}^{(k)}| \geq 2\tau$, and consequently $|r_{\max}^{(k)}| - (2K - 2)\mu|r_{\max}^{(k)}| + |r_{\min}| \geq 2\tau$. Consequently, condition (B.30) is true for (B.40). Then by Lemma 10, we have that under the event \mathcal{G} ,

$$\text{sgn}(\Re[r_{n_k} \mathbf{a}_{n_k}^H \mathbf{v}^{(k-1)}]) = b_{n_k}, \quad (\text{B.43})$$

i.e. $\mathcal{G} \subset \{b_{n_k}^{(k)} = b_{n_k}\}$. By induction, since no active users will be detected twice, it follows that the first K steps of the RDDF detector can detect all active users and their symbols, i.e.

$$\mathcal{G} \subset \cup_{k=1}^K \{\hat{\mathcal{I}}^{(k)} \subset \mathcal{I}, b_{n_k}^{(k)} = b_{n_k}\} = \{\hat{\mathcal{I}}^{(K)} = \mathcal{I}, b_n^{(K)} = b_n, n \in \mathcal{I}^{(K)}\}. \quad (\text{B.44})$$

Note that condition (B.21) is weaker than (4.47), since (4.47) can be written as $|r_{\min}|[1 - (2K - 1)\mu] \geq 2\tau$, which implies $|r_{\max}|[1 - (2K - 1)\mu] \geq 2\tau$. This further implies $|r_{\max}|[1 - 2(K - 1)\mu] + |r_{\min}| \geq 2\tau$, since $1 - 2(K - 1)\mu \geq 1 - (2K - 1)\mu$ and $|r_{\min}| \geq 0$. Consequently, under condition (4.47), from (B.44), $\mathcal{G} \subset \{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}$, and $1 - P_e = \mathbb{P}\{\{\hat{\mathcal{I}} = \mathcal{I}\} \cap \{\hat{\mathbf{b}} = \mathbf{b}\}\} \geq \mathbb{P}\{\mathcal{G}\}$ which is greater than one minus (4.46), which concludes the proof for the RDDF detector. This completes the proof of Theorem 1.

B.4 Proof of Lemma 6

Proof of Lemma 6 requires the following theorem in probability, which bounds the tail probability of a sum of bounded independent random variables:

Theorem 2 (Complex Hoeffding's Inequality). [45] Let x_1, x_2, \dots, x_N be complex-valued independent random variables satisfying $|x_n| \leq c_n$, $n = 1, \dots, N$, for some constant $c_n \geq 0$ almost surely. Define $S_N = \sum_{n=1}^N x_n$, then for any $t > 0$,

$$\mathbb{P}\{|S_N - E(S_N)| \geq t\} \leq 4 \exp\left(-\frac{t^2}{4 \sum_{n=1}^N c_n^2}\right). \quad (\text{B.45})$$

The l th column of the partial DFT matrix $\mathbf{A} \in \mathcal{C}^{M \times N}$ is given by: $[\mathbf{a}_l]_m = e^{i\frac{2\pi k_m l}{N}}/\sqrt{M}$, with k_m , $m = 1, \dots, M$, are mutually different and randomly selected over $\{1, \dots, N\}$. Now the inner product of the l th and p th columns of \mathbf{A} is given by

$$\mathbf{a}_l^H \mathbf{a}_p = \sum_{m=1}^M e^{i\frac{2\pi}{N} k_m (p-l)} / M, \quad (\text{B.46})$$

which is sum of independent bounded random variables $e^{i\frac{2\pi}{N} k_m (p-l)} / M$ whose absolute values are less than $1/M$. The mean of the sum is zero, $\mathbb{E}\{\mathbf{a}_l^H \mathbf{a}_p\} = 0$, since k_m is uniformly distributed over $\{1, \dots, N\}$.

Hence we can apply the complex Hoeffding's inequality for some $\nu > 0$:

$$\mathbb{P}\{|\mathbf{a}_l^H \mathbf{a}_p| \geq \nu\} \leq 4 \exp\left(-\frac{\nu^2}{4M/M^2}\right) = 4 \exp\left(-\frac{\nu^2 M}{4}\right). \quad (\text{B.47})$$

Now we consider all such inner products and use the union bound:

$$\begin{aligned} \mathbb{P}\left\{\max_{l < p} |\mathbf{a}_l^H \mathbf{a}_p| \geq \nu\right\} &\leq \frac{N(N-1)}{2} 4 \exp\left(-\frac{\nu^2 M}{4}\right) \\ &< 2N^2 \exp\left(-\frac{\nu^2 M}{4}\right). \end{aligned} \quad (\text{B.48})$$

Now let $e^{-c} = N^2 \exp\left(-\frac{\nu^2 M}{4}\right)$ for some constant $c > 0$, we have $\nu = \sqrt{\frac{4}{M}(2 \log N + c)}$.

Hence with probability exceeding $1 - 2e^{-c}$, we have

$$\mu \triangleq \max_{l < p} |\mathbf{a}'_l \mathbf{a}_p| \leq \nu = \sqrt{\frac{4}{M}(2 \log N + c)}, \quad (\text{B.49})$$

which proves Lemma B.4.

Bibliography

- [1] D. Aldous. *Probability approximations via the Poisson clumping heuristic*. Springer, 1 edition, Nov. 1988.
- [2] J. G. Andrews. Interference cancellation for cellular systems: A contemporary overview. *IEEE Wireless Comm. Mag.: on advanced receiver for CDMA*, 2005.
- [3] D. Angelosante, E. Biglieri, and M. Lops. Multiuser detection in dynamic environment - Part II: joint user identification and parameter estimation. *IEEE Trans. Info. Theory*, 55(5):2365 – 2374, May 2009.
- [4] D. Angelosante, E. Biglieri, and M. Lops. Low-complexity receivers for multiuser detection with an unknown number of active users. *Signal Processing*, 90(5):1486 – 1495, May 2010.
- [5] L. Applebaum, W. Bajwa, M. F. Duarte, and R. Calderbank. Multiuser detection in asynchronous on-off random access channels using lasso. In *Proc. 48th Annu. Allerton Conf. Comm., Control, and Computing*, Monticello, IL, Sept. 2010.

- [6] L. Applebaum, W.U. Bajwa, M.F. Duarte, and R. Calderbank. Asynchronous code-division random access using convex optimization. *Submitted for journal publication and arXived*, Jan. 2011.
- [7] W. U. Bajwa. *New information processing theory and methods for exploiting sparsity in wireless systems*. PhD thesis, University of Wisconsin-Madison, Madison, WI., June 2009.
- [8] M. Basseville and Igor V. Nikiforov. *Detection of abrupt changes: Theory and Applications*. Prentice Hall Information Theory and System Sciences Series, Englewood Cliffs, N.J., April 1993.
- [9] Z. Ben-Haim, Y. C. Eldar, and M. Elad. Coherence-based performance guarantees for estimating a sparse vector under random noise. *IEEE Trans. Signal Process.*, 58(10):5030 – 5043, Oct. 2010.
- [10] E. Biglieri and M. Lops. Multiuser detection in dynamic environment - Part I: user identification and data detection. *IEEE Trans. Info. Theory*, 53(9):3158 – 3170, Sept. 2007.
- [11] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Appl. Comput. Harmon. Anal.*, 27:265 – 274, 2009.
- [12] S. Buzzi, H. V. Poor, and A. Zappone. Transmitter waveform and widely linear receiver design: Noncooperative games for wireless multiple-access networks. *IEEE Trans. Info. Theory*, 56(10):4874 – 4892, Oct. 2010.

- [13] T. Cai, L. Wang, and G. Xu. Stable recovery of sparse signals and an oracle inequality. Tech. Rep. [Online]. Available: <http://www-stat.wharton.upenn.edu/tcai/paper/Stable-Recovery-MIP.pdf>, U. Penn., 2009.
- [14] A. T. Campo, A. G. Fabregas, and E. Biglieri. Large-system analysis of multiuser detection with an unknown number of users: A high-SNR approach. *IEEE Trans. Info. Theory*, 57(6):3416 – 3428, June 2011.
- [15] E. Candes and T. Tao. Rejoinder: The Danzig selector: statistical estimation when p is much smaller than n . *Ann. Statist.*, 35(6):2392 – 2404, Dec. 2007.
- [16] E. J. Candes. Compressive sampling. *Proc. of International Congress of Mathematicians, Madrid, Spain*, 2006.
- [17] E. J. Candes and Y. Plan. Near-ideal model selection by l_1 minimization. *Ann. Statist.*, 37(5A):2145 – 2177, 2009.
- [18] E. J. Candes and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Info. Theory*, 52(12):5406 – 5424, Dec. 2006.
- [19] M. Chen, S. Gonzalez, A. Vasilakos, H. Cao, and V. C. M. Leung. Body area networks: a survey. *Mobile Networks and Applications*, DOI:10.1007/s11036-010-0260-8:1 – 23, Aug. 2010.
- [20] T. Cover. Some advances in broadcast channels. In *Advances in Communication Systems*. Academic, New York, NY, USA, 1975.
- [21] T. M. Cover and J. A Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications and Signals Processing, 2 edition, 2006.

- [22] D. L. Donoho. Compressed sensing. *IEEE Trans. Info. Theory*, 52(4):1289 – 1306, April 2006.
- [23] D. L. Donoho, M. Elad, and V. N. Temlyakov. Stable recovery of sparse over-complete representations in the presence of noise. *IEEE Trans. Inf. Theory*, 52(1):6 – 18, 2006.
- [24] M. F. Duarte and Y. C. Eldar. Structured compressed sensing: from theory to applications. *submitted to IEEE Trans. Signal Process. and arXived*, 2011.
- [25] A. Duel-Hallen. Decorrelating decision feedback multiuser detector for synchronous code-division multiple-access channel. *IEEE Trans. Comm.*, 41(2):285 – 290, Feb. 1993.
- [26] A. Duel-Hallen, J. Holtzman, and Z. Zvonar. Multiuser detection for CDMA systems. *IEEE Personal Communications*, pages 46–58, April 1995.
- [27] Y. C. Eldar. Compressed sensing of analog signals in shift-invariant spaces. *IEEE Trans. Signal Process.*, 57(8):2986–2997, August 2009.
- [28] Y. C. Eldar. Uncertainty relations for shift-invariant analog signals. *IEEE Trans. Info. Theory*, 55(12):5742 – 5757, Dec. 2009.
- [29] Y. C. Eldar, P. Kuppinger, and H. Bolcskei. Block-sparse signals: Uncertainty relations and efficient recovery. *IEEE Trans. on Signal Process.*, 58(6):3042 – 3054, June 2010.
- [30] Y. C. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE Trans. Info. Theory*, 55(11):5302 – 5316, Nov. 2009.

- [31] A. K. Fletcher, S. Rangan, and V. K. Goyal. Necessary and sufficient conditions on sparsity pattern recovery. *IEEE Trans. Info. Theory*, 55(12):5758 – 5772, Jan. 2009.
- [32] A. K. Fletcher, S. Rangan, and V. K. Goyal. On-off random access channels: A compressed sensing framework. *submitted to IEEE Trans. Info. Theory and arXived.*, March 2010.
- [33] M. Fornasier and H. Rauhut. Compressive sensing. In O. Scherzer, editor, *Handbook of Mathematical Methods in Imaging*, chapter 2. Springer, 2011.
- [34] J. J. Fuchs. Recovery of exact sparse representations in the presence of bounded noise. *IEEE Trans. Inf. Theory*, 51(10):3601 – 3608, 2005.
- [35] A. E. Gamal and Y.-H. Kim. *Network information theory*. Cambridge University Press, Nov. 2011.
- [36] K. Gedalyahu and Y. C. Eldar. Time-delay estimation from low-rate samples: A union of subspaces approach. *IEEE Trans. on Signal Process.*, 58(6):3017 – 3031, June 2010.
- [37] R. Gribonval, B. Mailhe, H. Rauhut, K. Schnass, and P. Vandergheynst. Average case analysis of multichannel thresholding. In *Proc. IEEE ICASSP07*, Honolulu, 2007.
- [38] T. Guess. CDMA with power control and sequence design: the capacity region with and without multidimensional signaling. *IEEE Trans. Info. Theory*, 50(11):2604 – 2619, Nov. 2004.

- [39] D. Guo and C.-C. Wang. Multiuser detection of sparsely spread CDMA. *IEEE Journal On Selected Areas in Comm.*, 26(3):421 – 431, April 2008.
- [40] P. Gupta and P. R. Kumar. Capacity of wireless networks. *IEEE Trans. Info. Theory*, 46(2):388 – 404, March 2000.
- [41] J. Haupt and R. Nowak. Signal reconstruction from noisy random projections. *IEEE Trans. Info. Theory*, 52(9):4036 – 4080, Sept. 2006.
- [42] J. Haupt and R. Nowak. Compressive sampling for signal detection. *IEEE International Conference on Acoustics Speech and Signal Processing*, 3:1509–1512, 2007.
- [43] J. Haupt and R. Nowak. A generalized restricted isometry property. Technical report, Univ. of Wisconsin, Department of ECE, 2007.
- [44] R. W. Heath, T. Strohmer, and A. J. Paulraj. On quasi-orthogonal signatures for CDMA systems. *IEEE Trans. Info. Theory*, 52(3):1217 – 1226, Mar. 2006.
- [45] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. of Amer. Stat. Asso.*, 58(301):13 – 30, March 1963.
- [46] M. L. Honig, editor. *Advances in multiuser detection*. Wiley Series in Telecommunications and Signal Processing. Wiley-IEEE Press, Aug. 2009.
- [47] Y. Jin, Y.-H. Kim, and B. D. Rao. Support recovery of sparse signals. *submitted to IEEE Trans. Info. Theory and arXived*, March 2010.
- [48] T. L. Lai. Sequential changepoint detection in quality control and dynamic systems (with discussion). *J. Roy. Statist. Soc. Ser. B*, 57:613 – 658, 1995.

- [49] T. L. Lai. Sequential analysis: Some classical problems and new challenges (with discussion). *Statistica Sinica*, 11(2):303 – 408, Apr. 2001.
- [50] C. Lévy-Leduc and F. Roueff. Detection and localization of change-points in high-dimensional network traffic data. *The Annals of Applied Statistics*, 3(2):637–662, 2009.
- [51] G. Lorden. Procedures for reacting to a change in distribution. *Annals of Mathematical Statistics*, 42:1897–1908, Dec. 1971.
- [52] J. Luo, S. Ulukus, and A. Ephremides. Optimal sequence and sum capacity of symbol asynchronous CDMA systems. *IEEE Trans. Info. Theory*, 51(8):2760 – 2769, Aug. 2005.
- [53] R. Lupas and S. Verdu. Asymptotic efficiency of linear multiuser detectors. In *Proc. of 25th Conf. on Decision and Control (CDC)*, pages 2094 – 2100, Dec. 1986.
- [54] R. Lupas and S. Verdu. Linear multiuser detectors for synchronous code-division multiple-access channel. *IEEE Trans. Info. Theory*, 35(1):123 – 136, Jan 1989.
- [55] R. Lupas and S. Verdu. Near-far resistance of multiuser detectors in asynchronous channels. *IEEE Trans. Comm.*, 38(4):496 – 508, April 1990.
- [56] Y. Mei. Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*, 97(2):419 – 433, 2010.
- [57] M. Mishali and Y. C. Eldar. Reduce and boost: Recovering arbitrary sets of jointly sparse wideband analog signals. *IEEE Sel. Topics Signal Process.*, 56(10):4692 – 4702, Oct. 2008.

- [58] M. Mishali and Y. C. Eldar. Blind multiband signal reconstruction: Compressed sensing for analog signals. *IEEE Trans. Signal Process.*, 57(3):993 – 1009, Mar. 2009.
- [59] M. Mishali and Y. C. Eldar. From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals. *IEEE Journal of Selected Topics in Signal Process.*, 4(2):375 – 391, April 2010.
- [60] S. Moshavi. Multi-user detection for DS-CDMA communications. *IEEE Comm. Mag.*, 34(10):124 – 136, Oct. 1996.
- [61] T. Oskiper and H. V. Poor. Online activity detection in a multiuser environment using a matrix CUSUM algorithm. *IEEE Trans. Info. Theory*, 48(2):477 – 493, Feb. 2002.
- [62] E. Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, June 1954.
- [63] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications with wavelet decomposition. In *Proc. 27th Asilomar Conf. Signals, Systems*, pages 40 – 44, Nov. 1993.
- [64] A. Petrov, B. L. Rozovskii, and A. G. Tartakovsky. Efficient nonlinear filtering methods for detection of dim targets by passive systems. *Submitted to Multitarget-Multisensor Tracking: Applications and Advances*, IV, 2003.
- [65] M. Pollak and D. Siegmund. Approximations to the expected sample size of certain sequential tests. *Ann. Statist.*, 3(6):1267 – 1282, 1975.

- [66] H. V. Poor and O. Hadjiladis. *Quickest Detection*. Cambridge University Press, 2008.
- [67] H. V. Poor and S. Verdu. Probability of error in MMSE multiuser detector. *IEEE Trans. Info. Theory*, 43(3):858 – 871, May 1997.
- [68] D. Rabinowitz. *Detecting clusters in disease incidence*, pages 255 – 275. Change-point Problems. IMS, Hayward, CA, 1994.
- [69] R. Rajagopal, X. Nguyen, S.C. Ergen, and P. Varaiya. Distributed online simultaneous fault detection for multiple sensors. In *Int. Conf. Info. Processing in Sensor Networks (IPSN)*, pages 133 – 144, May 2008.
- [70] G. Reeves and M. Gastpar. A note on optimal support recovery in compressed sensing. In *Forty-Third Asilomar Conf. on Sig., Sys, and Computers*, pages 1576 – 1580, Nov. 2009.
- [71] G. Ricci and M. K. Varanasi. Blind multiuser detection of L out of K active users of an N user synchronous CDMA system. In *IEEE 49th Vehicular Tech. Conf.*, pages 1445 – 1449, May 1999.
- [72] S. W. Roberts. A comparison of some control chart procedures. *Technometrics*, 8(3):411 – 430, 1966.
- [73] M. Rupf and J. L. Massey. Optimal sequence multisets for synchronous code-division multiple-access channels. *IEEE Trans. Info. Theory*, 40(4):1261 – 1266, July 1994.
- [74] C. Schlegel and A. Grant. *Coordinated multiuser communications*. Springer, May 2006.

- [75] K. Shafie, B. Sigal, D. Siegmund, and K. Worsley. Rotation space random fields with an application to fmri data. *Ann. Statist.*, 31:1732 – 1771, 2003.
- [76] W. A. Shiryaev. On optimal methods in quickest detection problems. *Theory Prob. Appl.*, 8:22 – 46, Jan. 1963.
- [77] Z. Sidak. Rectangular confidence regions for the means of multivariate normal distributions. *J. of Amer. Stat. Asso.*, 62(318):626 – 633, Jun. 1967.
- [78] D. Siegmund and E. S. Venkatraman. Using the generalized likelihood ratio statistic for sequential detection of a change-point. *Ann. Statist.*, 23(1):255 – 271, 1995.
- [79] D. Siegmund and B. Yakir. *The statistics of gene mapping*. Springer, 2007.
- [80] D. O. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer Series in Statistics. Springer, Aug. 1985.
- [81] D. O. Siegmund and B. Yakir. Detecting the emergence of a signal in a noisy image. *Statistics and Its Inference*, 1:3–12, 2008.
- [82] M. Simon, J. Omura, R. Scholtz, B. Levitt, J. K. Omura, R. A. Scholtz, and M. K. Simon. *Spread spectrum communications handbook*. McGraw-Hill Professional, Sept. 2001.
- [83] G. Tang and A. Nehorai. Performance analysis for sparse support recovery. *IEEE Trans. Info. Theory*, 56(3):1383 – 1399, Mar. 2010.
- [84] A. G. Tartakovsky and V. V. Veeravalli. Asymptotically optimal quickest change detection in distributed sensor. *Sequential Analysis*, 27(4):441–475, 2008.

- [85] J. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inf. Theory*, 50(10):2231 – 2242, Oct. 2004.
- [86] J. A. Tropp, I. S. Dhillon, and R. W. Heath. Finite-step algorithms for constructing optimal CDMA signature sequences. *IEEE Trans. Info. Theory*, 50(11):2916 – 2921, Nov. 2004.
- [87] A. M. Tulino and S. Verdu. *Random matrix theory and wireless communications*, chapter 3.1. Now Publishers Inc, June 2004.
- [88] S. Ulukus and R. D. Yates. Iterative construction of optimum signature sequence sets in synchronous CDMA systems. *IEEE Trans. Info. Theory*, 47(5):1989 – 1998, July 2001.
- [89] M. K. Varanasi. Group detection for synchronous Gaussian code-division multiple-access channels. *IEEE Trans. Info. Theory*, 41(4):1083 – 1096, July 1995.
- [90] M. K. Varanasi. Decision feedback multiuser detection: A systematic approach. *IEEE Trans. Info. Theory*, 45(1):219 – 240, Jan. 1999.
- [91] M. K. Varanasi and B. Aazhang. Multistage detection in asynchronous code division multiple access communications. *IEEE Trans. Comm.*, 38(4):509 – 519, April 1990.
- [92] M. K. Varanasi and B. Aazhang. Near-optimal detection in synchronous code-division multiple-access systems. *IEEE Trans. Comm.*, 39(5):725 – 736, May 1991.

- [93] S. Verdu. Minimum probability of error for asynchronous Gaussian multiple-access channels. *IEEE Trans. Info. Theory*, 32(1):85 – 96, Jan. 1986.
- [94] S. Verdu. Near-far resistant receivers for DS/SSMA communications. Technical report, U.S. Army Research Proposal, Contract DAAL03-87-K-0062, Princeton University, 1986.
- [95] S. Verdu. Computational complexity of optimum multiuser detection. *Algorithmica*, 4:303 – 312, 1989.
- [96] S. Verdu. *Multiuser Detection*. Cambridge University Press, 1998.
- [97] S. Verdu and S. Shamai (Shitz). Spectral efficiency of CDMA with random spreading. *IEEE Trans. Info. Theory*, 45(2):622 – 640, March 1999.
- [98] P. Viswanath and V. Anantharam. Optimal sequences and sum capacity of synchronous CDMA systems. *IEEE Trans. Info. Theory*, 45(6):1984 – 1991, Sept. 1999.
- [99] P. Viswanath and V. Anantharam. Optimal sequences for CDMA under colored noise: a Schur-saddle function property. *IEEE Trans. Info. Theory*, 48(6):1295 – 1318, June 2002.
- [100] A. J. Viterbi. Very low rate convolutional codes for maximum theoretical performance of spread-spectrum multiple-access channels. *IEEE Journal On Selected Areas in Comm.*, 8(4):641 – 649, May 1990.
- [101] A. J. Viterbi. *CDMA: principles of spread spectrum communication*. Addison-Wesley, Apr. 1995.

- [102] W.-C. Wu and K.-C. Chen. Identification of active users in synchronous CDMA multiuser detection. *IEEE Journal On Selected Areas in Comm.*, 16(9):1723 – 1735, Dec. 1998.
- [103] Z. Xie, R. T. Short, and C. K. Rushforth. A family of suboptimal detectors for coherent multiuser communications. *IEEE Journal On Selected Areas in Comm.*, 8(4):683 – 690, May 1990.
- [104] J. Yang and S. Roy. On joint transmitter and receiver optimization for multiple-input-multiple-output (MIMO) transmission systems. *IEEE Trans. Comm.*, 42(12):3221 – 3231, Dec. 1994.
- [105] N. Zhang, B. Yakir, and D. O. Siegmund. Detecting simultaneous variant intervals in aligned sequences. *Submitted to Annals of Applied Statistics*, 2010.
- [106] H. Zhu and G. B. Giannakis. Exploiting sparse user activity in multiuser detection. *IEEE Trans. on Comm.*, 59(2):454 – 465, Feb. 2011.